



Overland
Storage

SnapServer®

Administrator's Guide

For GuardianOS™ Version 6.5 on
SnapServers and Expansion Arrays



June 2011
10400327-002



©2011 Overland Storage, Inc. All rights reserved.

Overland®, Overland Data®, Overland Storage®, LibraryPro®, LoaderXpress®, Multi-SitePAC®, NEO®, NEO Series®, PowerLoader®, Protection OS®, REO®, REO 4000®, REO Series®, Snap Care®, SnapDisk®, SnapServer®, StorAssure®, and XchangeNOW® are registered trademarks of Overland Storage, Inc.

GuardianOS™, SnapWrite™, Snap Enterprise Data Replicator™, SnapSAN™, and SnapServer Manager™ are trademarks of Overland Storage, Inc.

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

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

PROPRIETARY NOTICE

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

Overland Storage provides this manual as is, without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Overland Storage may make improvements or changes in the product(s) or programs described in this manual at any time. These changes will be incorporated in new editions of this publication.

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

FW 6.5

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

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

Audience and Purpose

This guide is intended for system and network administrators charged with installing and maintaining SnapServers running GuardianOS 6.5 on their network. We assume the administrator is familiar with the basic concepts and tasks of multi-platform network administration.

This guide provides information on the installation, configuration, security, and maintenance of SnapServers running GuardianOS 6.5. It also provides information on installing and using the following utilities and software components:

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

GuardianOS 6.5 can be added as an upgrade to older SnapServer N2000, 650, 620, 550, 520, 510, 410, 210, or 110 systems.

Some of the information presented in this manual (particularly the Storage Configuration and Expansion sections) applies only to SnapServers with four (4) or more drives. Users of the SnapServer 110 and SnapServer 210 are encouraged to consult the *User's Guide for SnapServer 110 and 210* as their primary reference and to refer to this Administrator's Guide for advanced guidelines.

Product Documentation

Overland Storage GuardianOS and SnapServer product documentation and additional literature are available online at:

<http://support.overlandstorage.com/support/snapserver-nas.htm>

Overland Technical Support

For help configuring and using your product, search for help at:

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

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

Our Overland Storage Technical Support staff is also available to assist you by phone at:

1.877.654.3429 (Toll-free and active only in the U.S. and Canada)

1.858.571.5555 x5 (Worldwide access)

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

http://support.overlandstorage.com/support/overland_care.html

In addition, European customers can contact our United Kingdom office at:




+44 (0) 118-9898050

9:00 A.M. to 5:00 P.M. (GMT) Monday through Friday

Conventions

This document exercises several typographical conventions and alerts.

Alerts

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

Typographical Conventions

Convention	Description & Usage
Boldface	Words in boldface indicate items to select such as menu items or command buttons.
Ctrl-Alt-r	This type of format details the keys you press simultaneously. In this example, hold down the Ctrl and Alt keys and press the r key.
NOTE	A Note indicates neutral or positive information that emphasizes or supplements important points of the main text. A note supplies information that may apply only in special cases, for example, memory limitations or details that apply to specific program versions.
Menu Flow Indicator (>)	Words with a greater than sign between them indicate the flow of actions to accomplish a task. For example, Setup > Passwords > User indicates that you should press the Setup button, then the Passwords button, and finally the User button to accomplish a task.

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

Software Updates

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

<http://support.overlandstorage.com/support/snapserver-nas.htm>

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

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

Finding More Information

Product documentation related to GuardianOS SnapServers and expansion arrays are listed below. The current versions of all these documents are always available from the Overland Storage Knowledge Base (<http://support.overlandstorage.com/kb>).

Source	Location	Content
Quick Start Guide	Product Packaging and Web	Provides complete instructions for installing the server into a rack and connecting the server to the network. Also contains links to warranty registration and information.
Quick Start Guide Translations	Product Packaging and Web	Quick Start Guide translated into French, Italian, German, Spanish, and Russian.
EULA	Product Packaging and Web	End User License Agreement for GuardianOS.
Administrator Guide	Web	Provides an overview of the configuration, maintenance, and troubleshooting of SnapServers, the administration of the CA Antivirus software, the installation of third-party backup agents, and detailed instructions on using the Web Management Interface.
SnapServer Online Help	Web Management Interface and Overland Website	Basic troubleshooting information embedded in the software with direct links to additional information from the Administrator Guide.
Setup Guide	Web	Lists hardware specifications and initial configuration for SnapServers and SnapDisk expansion arrays.
User Guide for SnapServer 110 and 210	CD packaged with the product	Provides an overview of the configuration and maintenance of the SnapServer 110 and 210.

Electrostatic Discharge Information

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

- Transport products in static-safe containers such as conductive tubes, bags, or boxes.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free stations.
- Cover the library with approved static-dissipating material.
- Use a wrist strap connected to the work surface and properly-grounded tools and equipment.

- Keep the work area free of non-conductive materials such as foam packing materials.
- Make sure you are always properly grounded when touching a static-sensitive component or assembly.
- Avoid touching pins, leads, or circuitry.

Contents

Chapter 1 – Overview

GuardianOS Specifications	1-1
What's New in GuardianOS	1-3
Using SnapServer Manager	1-5
Connecting to the Server for the First Time	1-7
Using the Initial Setup Wizard for Remote Access	1-8
Scheduling Data Protection Tasks	1-10
Additional Setup Options	1-10
SnapExtensions	1-11

Chapter 2 – Network Access

Viewing Current Network Settings	2-2
TCP/IP Options	2-3
Configuring TCP/IP Settings	2-4
Default Network Protocol Settings	2-8
Windows Networking Configuration	2-8
Support for Windows Network Authentication	2-9
Connecting from a Windows Client	2-10
Connecting a Mac OS X Client Using SMB	2-11
Configuring Windows Networking	2-11
NFS Access	2-13
Apple Networking Configuration	2-15
FTP/FTPS Access	2-17
HTTP/HTTPS Access	2-18
Web View	2-21
Print Server	2-21
Configuring the Printer	2-21
Configuring the Client	2-22
Monitoring Print Jobs Remotely	2-23
Deleting a Printer	2-23

Chapter 3 – Users and Groups

Default User and Group Settings	3-1
UID and GID Assignments	3-2
Local Users and Groups	3-2
Guidelines for Local Authentication	3-3
Local User Management Procedures	3-4
Local Group Management Procedures	3-6
NIS Domain	3-7

Chapter 4 – Storage and Expansion

Storage Configuration Overview	4-1
--------------------------------------	-----

Snapshot Space Considerations	4-2
Creating a RAID Configuration	4-2
Adding Capacity to a SnapServer with Expansion Arrays	4-2
RAIDs	4-3
Creating a New RAID	4-7
Managing Disk Drives in a RAID	4-8
Using the RAID Storage Guides	4-9
RAID Groups	4-10
Volumes	4-11
Procedures	4-14
Creating a Volume	4-15
Editing/Expanding an Existing Volume	4-16
Quotas	4-17
Setting User Quotas	4-18
Data Migration	4-19
Preserving Permissions	4-22
Maintenance Tools	4-23
Expansion Arrays	4-24
SnapDisk E2000	4-24
Snap Expansion S50	4-25
Managing Expansion Array Storage	4-25
Integrating Orphan Expansion Units	4-27
Disks and Units	4-27
Replacing Disk Drives on a RAID	4-27
Adding Disk Drives to a RAID	4-28
Adding New Disk Drives to Increase Capacity	4-29
Hot Swapping Disk Drives	4-30

Chapter 5 – iSCSI Disks

Configuring iSCSI Initiators	5-2
iSCSI Configuration for Microsoft Windows using MS Initiator	5-3
Configuring the QLogic iSCSI Initiators for Microsoft Windows	5-7
iSCSI Configuration for Linux and UNIX	5-8
iSCSI Configuration for Novell NetWare	5-8
iSCSI Configuration for VMware	5-9
iSCSI Configuration for Mac	5-10
iSCSI Configuration on the SnapServer	5-10
Creating iSCSI Disks	5-12
Creating an iSCSI Disk	5-14
Editing an iSCSI Disk	5-15
Configuring VSS/VDS for iSCSI Disks	5-15
Configuring iSNS	5-18

Chapter 6 – Share and File Access

Configuring Share and Folder Security Overview	6-1
Components and Options	6-3
SnapTrees and Security Models	6-4
ID Mapping	6-8
Shares	6-9
Creating a Share	6-11
Editing or Deleting a Share	6-13
Configuring Share Access	6-13
Share Access Behaviors	6-14

Setting User-based Share Access Permissions	6-14
NFS Access	6-15
Creating Home Directories	6-17
Windows ACLs	6-18
Setting File and Directory Access Permissions and Inheritance (Windows)	6-19
Security Guides	6-19
Chapter 7 – Snapshots	
Snapshot Management and Usage	7-1
Estimating Snapshot Pool Requirements	7-2
Adjusting Snapshot Pool Size	7-3
Accessing Snapshots	7-3
Coordinating Snapshot and Backup Operations	7-4
Snapshot Procedures	7-5
Creating Snapshots	7-6
Editing or Rolling Back a Snapshot	7-7
Editing Snapshot Schedules	7-7
Managing the Snapshot Pool	7-8
Chapter 8 – Disaster Recovery	
Backing Up Server and Volume Settings	8-1
SnapDRImage Usage Scenario	8-4
Backing Up the NetVault Database Directory	8-4
Recovering the NetVault Database	8-5
Disaster Recovery Procedural Overview	8-6
Cloning a Server	8-7
Chapter 9 – CA Antivirus Software	
Antivirus Dependencies	9-1
Launching the CA Antivirus GUI	9-2
The Local Scanner View	9-3
Scan Job Configuration and Scheduling	9-3
Defining Scan Jobs	9-4
Running a Manual Scan Job	9-4
Scheduling a Scan Job	9-5
Signature Updates	9-5
Updating SnapServers that have Internet Access	9-6
Updating a SnapServer that does not have Internet Access	9-7
Distributing Updates from One Server to Another	9-7
Verifying Download Events	9-9
Alert Options	9-9
The Move Directory	9-9
Log View	9-10
Chapter 10 – Unicode	
Unicode and Protocol Interaction	10-1
How Snapshots Interact with Unicode	10-3
Backing Up Unicode Servers	10-3
Unicode and Expansion Arrays	10-4
Chapter 11 – Web Management Interface	
Server Status and Site Map	11-2

Server Name	11-3
Date/Time	11-3
Email Notification	11-4
SNMP	11-5
SSH Secure Shell	11-7
UPS	11-7
Registering Your Server	11-9
OS Update	11-9
Snap Finder	11-11
Host File Editor	11-11
Shutdown and Restart	11-12
Powering SnapServers On and Off	11-12
System Status	11-12
Active Users	11-13
Open Files	11-13
Using the Event Log	11-13
Tape Information	11-14
RAID Settings	11-14

Appendix A – Backup and Replication Solutions

Integrated Backup Solutions for the SnapServer	A-2
BakBone NetVault	A-2
Snap Enterprise Data Replicator (Snap EDR)	A-3
Off-the-Shelf Backup Solutions for the SnapServer	A-4
Preparing to Install a Backup Solution	A-4
Preinstallation Tasks	A-5
Installing the CA ARCserve Agent	A-5
Installing the Symantec Backup Exec RALUS Agent	A-7
Installing the Symantec NetBackup 6.5 Client	A-10
Installing the EMC NetWorker Client	A-11
Backup of iSCSI Disks	A-15

Appendix B – Command Line Interface

SnapCLI Syntax	B-1
SnapCLI Commands	B-3
Scripts in SnapCLI	B-8

Appendix C – Troubleshooting SnapServers

The Meaning of LED Indicators	C-1
System Reset Options	C-9
Maintenance Mode	C-10
Networking Issues	C-11
Miscellaneous Issues	C-13
Phone Home Support	C-15

Appendix D – GuardianOS Ports

Glossary

Index

SnapServers are designed as flexible, low-maintenance network file servers optimized for performance and efficiency. SnapServers run GuardianOS, an operating system built to maximize file I/O throughput across multi-network protocols. To this end, all unnecessary system control and processing functions that are associated with a general-purpose server have been removed. This guide applies to SnapServers and expansions running GuardianOS version 6.5 including supported SnapServers that you can upgrade to GuardianOS 6.5 (models N2000, 650, 620, 550, 520, 510, 410, 210, or 110).

Note: GuardianOS 6.5 does not support the legacy Windows NT Domain mode. Please configure the server for either Active Directory or workgroup mode before upgrading to GuardianOS 6.5.

Topics in Overview:

- GuardianOS Specifications
- What’s New in GuardianOS
- Using SnapServer Manager
- Connecting to the Server for the First Time
- Using the Initial Setup Wizard for Remote Access
- Scheduling Data Protection Tasks
- Additional Setup Options
- SnapExtensions

GuardianOS Specifications

These specifications apply to all SnapServers and expansion arrays running version 6.5 of GuardianOS.

Feature	Specification
Network Transport Protocols	TCP/IP UDP/IP
Network Block Protocols	iSCSI
Network File Protocols	Microsoft Networking (CIFS/SMB) UNIX Network File System (NFS) 2.0/3.0/4.0 Apple Filing Protocol (AFP) v2.0/v3.1 Hypertext Transfer Protocol (HTTP/HTTPS) File Transport Protocol (FTP/FTPS)

Feature	Specification
Network Client Types	Microsoft Windows 2000/XP/2003 R2/2008 R2/Vista/7 Mac OS X 10.x Sun Solaris 10 HP-UX 11 AIX 5.3/6 Red Hat Linux 9.0 Red Hat Enterprise Linux (RHEL) 4.x/5.x Red Hat Fedora Core 5.x/6.x SuSE Linux Enterprise Server (SLES) 10.x/11.x
Server Emulation	Windows 2000/2003/2008 AppleShare 6.0 Network File System (NFS) 2/3/4 Windows Print Server IPP Print Server
Network Security	CA Antivirus software Microsoft Active Directory Service (ADS) (member server) UNIX Network Information Service (NIS) File and Folder Access Control List (ACL) Security for Users and Groups Secure Sockets Layer (SSL v2/3) 128-bit Encryption Target Challenge Handshake Authentication Protocol (CHAP) for iSCSI SMTP Authentication and support for email encryption (STARTTLS and TLS/SSL encryption protocols)
Data Protection	Snapshots for immediate or scheduled point-in-time images of the file system Support for local backup with Symantec NetBackup/Backup Exec Remote Media Server for Linux (plus BakBone NetVault on SnapServers upgraded from older versions) Support for network backup with Symantec NetBackup/Backup Exec, CA ARCserve, EMC NetWorker, or BakBone NetVault 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
System Management	Browser-based remote administration tool called the Web Management Interface SnapCLI for volume system deployment SnapServer Manager utility (platform independent) SNMP (MIB II and Host Resource MIB) User disk quotas for Windows, UNIX/Linux, Mac, FTP/FTPS Group disk quotas for UNIX/Linux Environmental monitoring Email event notification and SNMP trap notification Data migration

Feature	Specification
RAID Options	<ul style="list-style-type: none"> • RAID 0 (drive striping): Large virtual drive with data striped across all drives of the array to provide maximum performance with no loss in usable capacity. Does not provide data protection. • RAID 1 (drive mirroring): One or more drives duplicate one drive for maximum data protection. Available only on systems with two (2) or more drives. • RAID 5 (drive striping with parity): For each array, the size of one drive is reserved for parity. Provides good performance and space utilization with one-drive fault tolerance. Available only on systems with four (4) or more drives. • RAID 6 (drive striping with two parity drives): Like a RAID 5 except that two drives are used for parity rather than one. Provides moderate performance and reasonable space utilization with two-drive fault tolerance. Available only on systems with four (4) or more drives. • RAID 10 (striped mirroring): A combination of RAID 0 and RAID 1. Provides high performance and fault tolerance. Available only on systems with four (4) or more drives. • Global or local hot spare support • Instant Capacity Expansion (ICE): Logically groups RAIDs for dynamic online scalability.
DHCP Support	Supports Dynamic Host Configuration Protocol (DHCP) for automatic assignment of IP addresses

What's New in GuardianOS

The following tables list the new and changed features since GuardianOS v5.2.

What's New in GuardianOS v6.5

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

GuardianOS 6.5 has the following new features and functionality:

Feature	New Functionality
Windows Server 2008 R2 Enhanced Support	GuardianOS 6.5 now supports the ability to join a Windows Server 2008 R2 Domain.
Unicode Support	GuardianOS now runs exclusively in Unicode mode. All systems upgrading from GuardianOS 5.2 or 6.0 that previously did not have Unicode enabled will be converted to Unicode.
iSCSI Write Cache Modification	The iSCSI Disk Properties page of the Web Management Interface now features an "Enable Write Cache" checkbox to allow modification of the write cache setting of an existing iSCSI disk.
NFS Configuration	The NFS page of the Web Management Interface now features a "Domain Name" field for NFSv4. This is the domain used by NFSv4 for ID mapping. The default is "localdomain."

Feature	New Functionality
LDAP Signing	GuardianOS now supports LDAP signing to guarantee authenticity of LDAP communication with Active Directory domain controllers. LDAP signing with GuardianOS is controlled by both the Web Management Interface and CLI. There are three possible values for LDAP signing: plain (no signing - default), sign, or seal.
Zero Configuration IP Assignment	SnapServers now use Zero Configuration IP Assignment (also known as Link-Local addresses). This replaces the 10.10.10.10 static IP address used in previous releases when no DHCP service is available. The Zero Configuration address range is 169.254.1.0 to 169.254.254.255. SnapServer Manager (SSM) will discover SnapServers that are configured with a Zero Configuration address. SnapServer Manager may be required to access the administration page of a SnapServer that is using a Zero Configuration IP address.
SnapServer Command Line Interface (SnapCLI) Changes	The following SnapCLI commands have been updated or changed: <ul style="list-style-type: none"> • fscheck – Check or repair user or root file systems; previously named volume fscheck. • checkroot - Added support to perform root file system checks, previously named volume fscheck-root. This command requires a system reboot. • raid-speed-limit – Added support to permanently (or temporarily) change the maximum RAID re-sync speed. • windows – The set subcommand has been expanded to support enable and disable of LDAP signing.
GuardianOS 6.5 Discontinued Features	<p>Windows NT Domain Support – GuardianOS 6.5 does not support SnapServers as part of Windows NT Domains.</p> <p>BakBone NetVault: Backup for GuardianOS – The BakBone Software product is no longer included with GuardianOS</p> <p>AppleTalk Protocol – The AppleTalk protocol is no longer supported in GuardianOS 6.5. AFP over TCP/IP remains supported.</p> <p>DHCP Server –The capability for GuardianOS to provide DHCP services for the network to which it is attached has been removed.</p> <p>SnapServer Command Line Interface (SnapCLI) – The following SnapCLI commands have been removed:</p> <ul style="list-style-type: none"> • apple – The subcommands supporting AppleTalk configuration have been removed. TCP/IP is the only public option. • dhcp – The support for SnapServers to host DHCP has been removed. The dhcp command has been removed. • registration – Removed the registration command. • windows – Support for Windows NT Domains has been removed. The subcommands supporting NT Domain configuration have been removed. • unicode – Removed the Unicode command as SnapServers with GuardianOS 6.5 have Unicode enabled.

What's New in GuardianOS v6.0

GuardianOS 6.0 has the following new features and functionality:

Feature	New Functionality
Support for Multiple Ethernet Ports	With the installation of an Ethernet card, the SnapServer N2000 supports up to 6 Ethernet ports.
Write Cache Option	Write cache can now be disabled on a volume, allowing data to be written directly to the disk. Note: For this feature to be available, all drives on the volume must support disabling of write cache.
Email Authentication and Encryption Capability	SMTP Authentication and Secure Connection have been added to the SnapServer email capabilities.
Wake-on-LAN	SnapServer N2000 supports Wake-on-LAN on Ethernet1 and Ethernet2.

What's New in GuardianOS v5.2

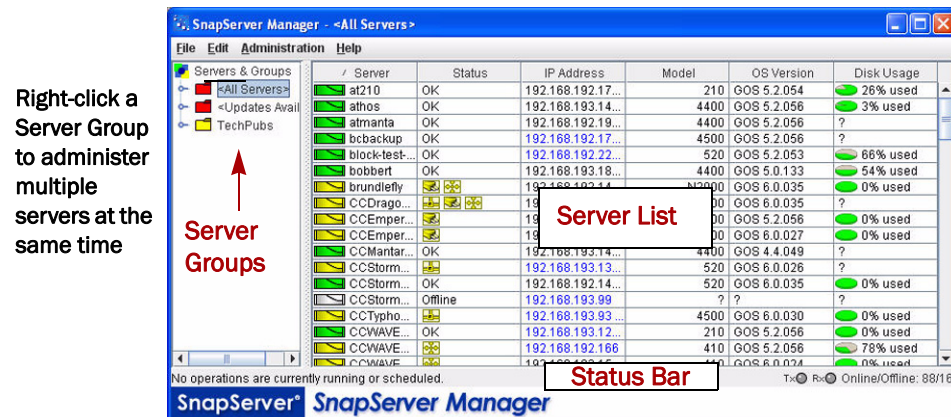
GuardianOS 5.2 has the following new features and functionality:

Feature	New Functionality
VSS/VDS Support for iSCSI	VSS (Volume Shadow Copy Service) and VDS (Virtual Disk Service) Hardware Providers have been added for SnapServer iSCSI targets. VSS provides a mechanism by which application-consistent snapshots of iSCSI targets may be taken without performing full application (or system) shutdown, for backup or other purposes. The VDS feature allows a Windows administrator to natively manage iSCSI storage, using any VDS compliant management console application.
Password Policies	The administrator can now set password policies for local users to establish requirements, expiration dates, and automatic logout.
User Interface Enhancements	The User Interface now comes in three color schemes: green slate, azure sea, and golden desert.
Windows 2008 Domain Support	Windows domains hosted by Windows 2008 servers are now supported.
Support for 128-bit SMB Encryption	GuardianOS now supports 128-bit encrypted communication with SMB clients and servers.
File Security Viewing	When logged in as an administrator, files and folders in Web View now display a security icon (key) that, when clicked, shows security information about the file/folder.
Root Filesystem Check	The Web Management Interface now provides the ability to check the root filesystem for errors and repair if found.

Using SnapServer Manager

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

copy, and configure settings for groups of GuardianOS SnapServers in a single operation.



Installing SnapServer Manager

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

Launching SnapServer Manager

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

Operating System	Procedure
Microsoft Windows XP/2000/2003/Vista/2008/7	Click Start . Point to Programs (or All Programs)> SnapServer Manager, then select SnapServer Manager.
Mac OS v10.3 or higher	Open the SnapServer Manager folder and double-click the SnapServer Manager icon.
UNIX/Linux	For default options: “cd” to home directory, then run the SnapServer Manager command: ./Snap_Server_Manager If you selected not to create links: “cd” to home directory, then cd to the SnapServer Manager directory, and run the SnapServer Manager command: ./Snap_Server_Manager

Multiserver Administration

Multiserver administration is available only for GuardianOS SnapServers.

- **Simultaneous application of settings to server groups** – You can organize GuardianOS servers into functional groups and apply settings to all servers in the group simultaneously.
- **Comparing settings across servers** – SSM can compare settings across any number of GuardianOS servers and identify when settings differ among servers. For example, comparing protocol access configuration for a group of servers may reveal that settings are consistent for Windows, NFS, and AFP but that differences exist among servers in HTTP/HTTPS and FTP/FTPS settings.

- **Copying settings from one server to one or more different servers** – SSM can copy selected settings (TCP/IP, SNMP, SMB, etc.) from any GuardianOS server to one or more different GuardianOS servers.
- **Scheduling operations to run during offpeak hours** – Operations can be scheduled to run on multiple GuardianOS servers during offpeak hours.
- **Automatic email notification of completed operations** – You can configure SSM to send an operations report (CSV format) upon completion of any operation.
- **Automatic notification of available GuardianOS updates** – SSM is by default configured to check daily for applicable updates to the servers it has discovered and display an alert, notifying the administrator of the available updates.

SSM Feature Licensing

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

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

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

Connecting to the Server for the First Time

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

To Connect Using the Server Name

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

1. Find the **server name**.

The default server name is “SNAPnnnnnnn,” where *nnnnnnn* is the server number. For example, the name of a SnapServer N2000 with a server number of 610019 is SNAP610019. The server number is a unique, numeric-only string that appears on a label affixed to the top of the server in the left front corner.

2. In a Web browser, enter the **server URL**.

For example, enter “<http://SNAPnnnnnnn>” (where *SNAPnnnnnnn* is the server name).

3. Press **Enter** to open the Web View screen.

4. Log into the **Web Management Interface**.
In the login dialog box, enter **admin** as the user name and **admin** as the password, then click OK.
5. Complete the **Initial Setup Wizard**.

To Connect to a SnapServer Using SSM

1. Launch SSM.
SSM discovers all SnapServers on its local network segment and displays their server names, IP addresses, and other status information in the main console. If you do not have a DHCP server, there might be a delay before the server appears on the network.
Note: To distinguish multiple SnapServers, you may need to find their default server names as explained in the previous procedure.
2. If using a DHCP server, proceed to **Step 3**; otherwise, assign an **IP address** to the new server.
 - a. In SSM, right-click the **server name**.
 - b. Select **Set IP Address**.
 - c. Enter an IP address and a subnet mask, then click **OK**.
3. In SSM, right-click the server name and select **Launch Web Administration**.
4. Log into the **Web Management Interface**.
In the login dialog box, enter **admin** as the user name and **admin** as the password, then click OK.
5. Complete the **Initial Setup Wizard**.

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

Using the Initial Setup Wizard for Remote Access

The first time you connect to a SnapServer via the Web Management Interface, the Initial Setup Wizard runs. The Initial Setup Wizard consists of several screens that allow you to change the server name, set the date and time, set the administrator password, configure TCP/IP settings for the primary Ethernet port (by default Ethernet1), and optionally configure storage space if none is configured.

Server Name

The default server name is SNAPnnnnnnn, where nnnnnnn is the server number. 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.

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 EDR operations. Edit the settings according to local conditions.

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

Changing 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 secure password immediately in the fields provided.

Note: A password must consist of 1 to 15 alphanumeric characters and is case sensitive.

Gathering TCP/IP Addressing Information

SnapServers are preset to acquire an IP address from a DHCP server. If you wish to assign a static IP instead, assemble the following information:

- The IP address for the SnapServer (required)
- The subnet mask (required)
- The default gateway IP address
- The DNS IP address
- WINS server(s) IP address(es)

Detecting Disk Drives Installed

All detected disk drives are displayed. If necessary, you can insert more hot-pluggable drives until you have the configuration you desire.

Choosing RAID Type and Snapshot Space

Select either a RAID type (such as RAID 6) or to manually configure the storage later. If a RAID type is selected, the setup program continues with the creation of the RAID, a volume that uses 80% of the RAID space, and a snapshot space that uses the remaining 20%.

For additional information on the spaces used, refer to [Chapter 7, "Snapshots."](#)

GuardianOS runs from a protected partition which consumes about 12 GB of space from each disk on systems shipped with GuardianOS 6.5 or later preinstalled.

Example:

To calculate the capacity of a SnapServer 410 with 2 TB total capacity (4x500 GB) in its default state, consider both the hardware and software configuration:

- The four 500 GB disk drives each provide 480 GB of formatted capacity.
- The four disks when joined in a RAID 5 configuration net 1.44 TB of capacity for the RAID (three drives plus one hot spare).
- The snapshot space is 20% of the space available on the RAID, reducing the space on the RAID for the data volume by 288 GB.
- This results in a data volume capacity of 1.152 TB for this SnapServer 410/2 TB.

Scheduling Data Protection Tasks

Scheduling backups, snapshots, and antivirus scans, and creating a disaster recovery image preserves your server configuration and protects your data from loss or corruption. Snapshots can be taken to provide a point-in-time image of files and changes to files to help in quickly recovering from accidental deletion or modification, or to facilitate performing an offline tape backup of an active data partition.

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

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

Create a disaster recovery image (DRImage) on the **Maintenance > Disaster Recovery** page. This DRImage should be created after the server configuration is complete, and can be used to recover the server or a replacement server to the configured state. See [“Disaster Recovery” on page 8-1](#) for detailed information on creating and using disaster recovery images.

GuardianOS contains built-in support for SnapEDR (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 A-1](#).

Additional Setup Options

Once the basic options are configured, other options can then be set up.

Migrating Data from Legacy Servers to the SnapServer

The Data Migration utility can be used to copy or move data from any computer supporting CIFS/SMB or NFS (v2 and v3) directly to a SnapServer. Access the utility by selecting **Maintenance > Data Migration**. For more information, see [“Data Migration” on page 4-19](#).

Configuring the SnapServer as a Print Server

Your SnapServer can be configured to emulate either a Windows or an IPP print server to manage USB-connected printers. To configure your SnapServer as a print server, select **Server > Printing** in the Web Management Interface. For more information, see [“Print Server” on page 2-21](#).

Configuring the SnapServer as a Simple Web Server

When the SnapServer is configured with a web root, the browser opens to a user-definable directory and optionally automatically loads a default HTML page when a user connects with a web browser to the root of the server (for example, `http://[servername]` or `http://[ipaddress]`). To configure a web root on the SnapServer, select **Network > Web** in the Web Management Interface. For more information, see [“Using WebRoot to Configure the SnapServer as a Simple Web Server” on page 2-19](#).

Configuring an APC-Brand UPS

Overland Storage recommends that you use a UPS with SnapServers and expansion arrays to protect your data from unforeseen power outages. SnapServers are compatible with USB- and network-based, APC-brand uninterruptible power supplies that allow you to take advantage of the automatic shutdown capability (some serial-only APC UPS's are also supported by using the IOGear GUC232A USB to Serial Adapter Cable). For instructions on configuring your APC-brand UPS device, navigate to the **Server > UPS** screen and click the **Help** icon see UPS.

Wake-on-LAN Support

Note: Available on the SnapServer N2000 only.

Wake-on-LAN, the Ethernet computer networking standard that allows a powered-off computer to be powered on by a network signal, is automatically enabled (and cannot be disabled) for Ethernet1 and Ethernet2. Wake-on-LAN is activated when another computer on the same LAN sends a “magic packet” to the SnapServer using the SnapServer Manager or other program designed to send magic packets. Wake-on-LAN only works for SnapServer N2000 systems and is not effective for any expansion that may be attached to the system.

SnapExtensions

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

To access SnapExtensions, click the SnapExtensions icon () from any page in the Web Management Interface.

Note: You may have a different set of SnapExtensions available to you than are listed in the following table if you have installed other SnapServer software, independent of the current operating system release.

Feature	Description
CA Antivirus	Preinstalled antivirus software that is fully functional out-of-the-box. For information on configuring the software, see “CA Antivirus Software” on page 9-1 . Note: A separate license is required on some platforms.
BakBone NetVault 8.2 (or later)	If previously installed and licensed on a GuardianOS system, GuardianOS 6.5 will support it. For information on installing and configuring NetVault, see “BakBone NetVault” on page A-2 , and the documentation included with the NetVault CD.
Snap EDR Management Console and Agent	Utility included with your SnapServer that synchronizes, transfers, backs up, and restores files between Windows, UNIX, and GuardianOS systems. Comes with a 45 day trial license, but requires a license for each SnapServer thereafter. For more information, see “Snap Enterprise Data Replicator (Snap EDR)” on page A-3 .

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

Topics in Network Access:

- [Viewing Current Network Settings](#)
- [TCP/IP Options](#)
- [Configuring TCP/IP Settings](#)
- [Default Network Protocol Settings](#)
- [Windows Networking Configuration](#)
- [NFS Access](#)
- [Apple Networking Configuration](#)
- [FTP/FTPS Access](#)
- [HTTP/HTTPS Access](#)
- [Print Server](#)

IMPORTANT: The default settings enable access to the SnapServer via all protocols supported by the SnapServer. As a security measure, disable all 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.

Viewing Current Network Settings

The **Network > Information** screen displays the server's current network settings. One column appears for each Ethernet port. Field definitions are given in the following table:

Ethernet Interface Information	
Port Name	The name of the Ethernet interface (for example, Ethernet1).
Enabled	Yes or no.
TCP/IP Settings Obtained From	DHCP or Static.
IP Address	The unique 32-bit value that identifies the server on a network subnet. This address consists of a network address, optional subnet address, and host address. It displays as four addresses ranging from 1 to 255, separated by periods (.).
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.
Primary WINS Server	The Windows Internet Naming Service server, which locates network resources in a TCP/IP-based Windows network by automatically configuring and maintaining the name and IP address mapping tables.
Secondary WINS Servers	Secondary Windows Internet Naming Service servers.
Ethernet Address	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port.
Speed Status	10 Mbps, 100 Mbps, or 1000 Mbps.
Duplex Status	Half-duplex: two-way data flow, only one way at a time. Full-duplex: two-way data flow simultaneously.
Bonding Status	Standalone, Load balance, Failover, Switch Trunking, or Link Aggregation.
Gateway Information	
Default Gateway	The network address of the gateway is the hardware or software that bridges the gap between two otherwise unroutable networks. It allows data to be transferred among computers that are on different subnets.
DNS Information	
Domain Name	The ASCII name that identifies the Internet domain for a group of computers within a network.
Primary DNS	The IP address of the primary Domain Name System server that maintains the list of all host names.
Secondary DNS #1	Secondary Domain Name System server #1.
Secondary DNS #2	Secondary Domain Name System server #2.

TCP/IP Options

GuardianOS SnapServers ship with one or more Gigabit Ethernet (GbE) ports. The following table describes TCP/IP options; default settings appear in italics.

Option	Setting	Description
TCP/IP Addressing	<i>DHCP</i>	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 necessary.
Network Bonding		Network bonding treats two or more ports as a single channel for failover or load balancing purposes. Note: Only applicable to servers with more than one Ethernet port.
	<i>Standalone</i>	The default state <i>Standalone</i> is the absence of network bonding and treats each port as a separate interface. Note: If you have more than two ports, you can have a mixture of <i>standalone</i> and <i>bonded</i> ports. For example, on a 4-port system, one port is <i>standalone</i> and three ports are <i>bonded</i> into a load balanced configuration.
	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 Ethernet1) as the primary network interface and one or more Ethernet ports are held in reserve as backup interfaces. 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, one of the backup ports assumes its network identity. The ports should be connected to different switches (though this is not required). Note: Failover mode provides switch fault tolerance, as long as the ports are connected to different switches.
	Switch Trunking	This mode groups multiple physical Ethernet links to create one logical interface. Provides high fault tolerance and fast performance between switches, routers, and servers.
	Link Aggregation (802.3ad)	Like Switch Trunking, this mode groups multiple physical Ethernet interfaces to create one logical interface, and provides high fault tolerance and fast performance between switches, routers, and servers. Uses Link Aggregation Control Protocol (LACP) to autonegotiate trunk settings.
Enable Ethernet	<i>Checked</i>	By default, all Ethernet ports are enabled, whether they are used or not.
	Unchecked	Ports other than the Primary Interface (by default Ethernet1) can be disabled by selecting the port and unchecking the Enable Ethernet checkbox. However, a bonded Ethernet port cannot be disabled, nor can a disabled Ethernet port be placed in bonded mode. Note: The primary Ethernet port must always be enabled. GuardianOS will not allow you to disable it.

Option	Setting	Description
Speed/ duplex	Auto	The default setting of Auto enables automatic negotiation of the speed and duplex settings based on the physical port connection to a switch. The speed setting establishes the rate of transmission and reception of data. The duplex setting allows the Ethernet port to transmit and receive network packets simultaneously. Note: Auto is the only allowable setting for a Gigabit port.
	Fixed	The SnapServer may also be set to fixed speed/duplex setting: 10Mbps/half; 10Mbps/full; 100Mbps/half; 100Mbps/full. Note: To prevent connectivity problems when changing to a fixed setting, see “Changing from Auto to a Fixed Link Setting” on page 2-6.
Primary Interface		By default, the primary Ethernet port is Ethernet1 and it cannot be disabled. However, the Primary Interface can be changed to a different Ethernet port by selecting the Ethernet port you want to make Primary and putting a check in 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.

Configuring TCP/IP Settings

TCP/IP settings are configured on the **Network > TCP/IP** screen of the Web Management Interface. This screen displays information about the server's Ethernet ports, including:

Column	Description
Port/Bond	A list of the Ethernet Ports or Bonds. Click a port or bond to display or modify configuration details.
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 <i>not available</i> if unknown. Whether the IP address was obtained by <i>DHCP</i> or is <i>Static</i>.
Bond Type	<ul style="list-style-type: none"> <i>Standalone</i>—The default <i>Standalone</i> setting treats each port as a separate interface, effectively disabling network bonding. Network bonding treats two or more ports as a single channel for failover or load balancing purposes. <i>Load Balance (ALB)</i>—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.

Column	Description
	<ul style="list-style-type: none"> • <i>Failover</i>—This mode uses one Ethernet port (by default Ethernet1) as the primary network interface and a second 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"> • <i>Switch Trunking</i>—This mode groups multiple physical Ethernet links to create one logical interface. Provides high fault tolerance and fast performance between switches, routers, and servers.
	<ul style="list-style-type: none"> • <i>Link Aggregation (802.3ad)</i>—Like Switch Trunking, this mode groups multiple physical Ethernet interfaces to create one logical interface, and provides high fault tolerance and fast performance between switches, routers, and servers. Uses Link Aggregation Control Protocol (LACP) to autonegotiate trunk settings.
Modified	<p>Indicates whether configuration for one or more interfaces has been changed and needs to be applied to take effect:</p> <ul style="list-style-type: none"> • <i>Yes</i>—One or more parameters for the interface have been modified. • <i>No</i>—No parameters for the interface have been modified.

Issues in TCP/IP Configuration

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

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

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

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

If you connect only one port, use the default primary port (Ethernet1). If you use Ethernet2 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).

Connect the SnapServer to the Network via a Switch

While it is possible to connect a SnapServer to the network via a hub, this configuration unduly restricts the performance of the server for the following reasons:

- Hubs do not support full-duplex. You can employ full duplex only when the SnapServer is connected to a switch.
- Hubs do not support Gigabit speeds. Attempting to force a Gigabit setting when the SnapServer is cabled to a hub will produce unintended consequences.

100 Mbps/half duplex is the best performance possible when connected to a hub.

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

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

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

Configure the Switch for Load Balancing

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

Changing from Auto to a Fixed Link Setting

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

- Configure the fixed setting in the Web Management Interface first; and
- Configure the switch to the same fixed setting.

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.

Procedures

To Edit TCP/IP Settings

Edit settings as described in the following table, and then click **OK** to update network TCP/IP settings immediately.

1. Select the appropriate port from the **Ethernet Ports** column. Edit the following information on the screen that opens.

Option	Settings
Enable Ethernet	<p>If a port other than the Primary Interface (by default Ethernet1) is selected, a checkbox will allow you to enable (check) or disable (uncheck) the port if it is unused. However, a bonded Ethernet port cannot be disabled, nor can a disabled Ethernet port be placed in bonded mode.</p> <p>Note: The Primary Interface port must always be enabled. GuardianOS will not allow you to disable it.</p>
TCP/IP Addressing	<p>This defaults to the DHCP-assigned setting. To assign a static IP address, select <i>Use The Settings Below</i> and enter the following information:</p> <ul style="list-style-type: none"> • IP address (required) • Subnet mask (required) <p>Optionally, you may also specify the default gateway, domain name server, domain name, and WINS server as appropriate.</p>
Speed/Duplex	<p>Accept the default setting <i>Auto</i> to have the SnapServer automatically negotiate speed and duplex settings with the switch. This setting is recommended as the optimal solution for port settings, and it <i>must</i> be selected to take advantage of a full-Gigabit connection.</p> <p>If you wish, you can also select a specific speed/duplex setting. If you manually configure speed/duplex settings, you must also configure the hub or switch with the same settings.</p>
Primary Interface	<p>By default, the Primary Ethernet port is Ethernet1 and it cannot be disabled. To change the Primary Interface to an Ethernet port other than Ethernet1, select the Ethernet port you want to make Primary and put a check in the Primary Interface box.</p>

2. Click **OK** to close the Configuration page, then click **OK** again to apply the changes.

To Create a Bond

1. Select a bonding mode by clicking the **Create Bond** button. On the page that opens, select:

Option	Settings
Bond Type	<p>Choose one of the following settings:</p> <ul style="list-style-type: none"> • <i>Load Balance (ALB)</i> – enables all selected ports to share the network load. • <i>Failover</i> – enables other selected ports to automatically take over the connection if the primary port fails. Only one port is active at any given time. • <i>Switch Trunking</i> and <i>Link Aggregation (802.3ad)</i> – these two options group multiple Ethernet ports into one logical Ethernet port for high speed and fault tolerance. <p>Ports not joined to a bond are configured as Standalone and have separate interfaces (one IP address per port).</p>
Select Ports to include in Bond	<p>Select the port(s) you want to include in the Bond from the Ports not in Bond column and use the Add button to move them to the Ports in Bond column.</p>

- Click **OK** to close the Configuration page, then click **OK** again to apply the changes.

Default Network Protocol Settings

SnapServers are preconfigured to allow multi-platform access in heterogeneous Windows, UNIX/Linux, and Mac environments. The following table summarizes the SnapServer's default network protocol access configuration.

Protocol	Default	Comments
Windows (CIFS/SMB)	Enabled	Allows access to Windows clients via the workgroup <i>Workgroup</i> .
UNIX (NFS)	Enabled	Allows universal access to all computers running NFS without client address restrictions.
Apple (AFP)	Enabled	Allows access for Mac clients over a TCP/IP AFP network using the default zone.
FTP/FTPS	Enabled for FTP, FTPS, and Anonymous User	<ul style="list-style-type: none"> Allows users to access files via FTP or FTPS. Allows access using the anonymous user account, which is mapped to the SnapServer's local guest user account.
HTTP/HTTPS (Internet/Intranet)	Enabled	Allows users to access files via HTTP or HTTPS using a Web browser.
Secure Shell (SSH)	Enabled	Required when installing a supported backup agent, using the Command Line Interface, or troubleshooting under the direction of a technical support representative. Using SSH for any other purpose is not supported and may void your warranty.

Note: As a security measure, disable any network protocols not required in your network environment.

Windows Networking Configuration

Windows SMB and security settings are configured on the **Network > Windows** screen of the Web Management Interface.

Topics include:

- Support for Windows Networking (SMB)
- Support for Windows Network Authentication
- Connecting from a Windows Client
- Connecting a Mac OS X Client Using SMB

Support for Windows Networking (SMB)

The default settings make the SnapServer available to SMB clients in the workgroup named *Workgroup*. Language support is set to North America/Europe (code page 850); opportunistic locking is enabled, as is participation in master browser elections. See the online help for details in configuring these options.

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

Windows Networking File and Folder Name Support

In Windows networking, most file and directory names are transmitted as a two-byte (16-bit) UCS-2 character set. However, this is not true in every case. Some are still sent via a single byte character set. The Language Support option selected for Windows networking clients is used only to enable the server to accept file and folder names in a single byte character set.

Support for Microsoft Name Resolution Servers

The SnapServer supports both of the Microsoft name resolution services: Windows Internet Naming Service (WINS) and Domain Name System (DNS). However, when you use a domain name server with a Windows Active Directory (ADS) server, make sure the forward and reverse name lookup is 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** screen gives you access to a separate and distinct Hidden share option that hides a share from SMB, AFP, HTTP, HTTPS, and FTP clients. (However, shares are not hidden from NFS clients, which cannot connect to shares that aren't visible. To hide shares from NFS clients, consider disabling NFS access on hidden shares.) For new shares, select **Create Share** and click the **Advanced Share Properties** button to access the Hidden share option. For existing shares, select the share, click **Properties**, and click **Advanced Share Properties** to access the Hidden share option.

Support for Windows Network Authentication

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

Windows Networking Options

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

Option	Description
Workgroup	In a workgroup environment, users and groups are stored and managed separately on each server in the workgroup.

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

Kerberos Authentication

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

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

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

Interoperability with Active Directory Authentication

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

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

Guest Account Access to the SnapServer

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

Connecting from a Windows Client

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

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

Connecting a Mac OS X Client Using SMB

Mac OS X clients can connect using SMB as well as AFP. Specify the server in the Connect to Server window as one of the following:

- `smb://servername`
- `smb://ipaddress`

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

Configuring Windows Networking

Windows SMB and security settings are configured from this page. Before performing the configuration procedures provided here, be sure you are familiar with the information provided in [“Support for Windows Networking \(SMB\)”](#) on page 2-8 and [“Support for Windows Network Authentication”](#) on page 2-9.

Procedures

To Edit Windows (SMB) Access Settings

Edit settings as described in the following table, and then click **OK** to update Windows network settings immediately.

Note: These settings apply to joining domains. For more information about joining domains, see [“To Join an Active Directory Domain”](#) on page 2-12.

Option	Settings
Enable Windows SMB	Check the Enable Windows Networking (SMB) checkbox to enable SMB. Clear the checkbox to disable SMB.
Member Of	Select <i>Workgroup</i> or <i>Active Directory Domain</i> .
Workgroup or Domain Name	The default settings make the SnapServer available in the workgroup named <i>Workgroup</i> . Enter the workgroup or domain name to which the server belongs. If you join a Windows domain through Advanced Security (Security > Security Guides > appropriate option), the domain name you entered displays here and can be changed only via the Advanced Security screen.
Administrator Name and Password	If joining a domain, enter the user name and password of a user with domain join privileges (typically an administrative user).
Organizational Unit	To create a machine account at a different location than the default, enter a name in the Organizational Unit field. By default, this field is blank, signaling the domain controller to use a default defined within the controller. Note: Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([organizationalunit]/[sub-unit1]/[sub-unit1a])
LDAP Signing	Set ADS domain LDAP signing to Plain (no signing), Sign, or Seal, as appropriate for your domain. Default setting is Plain.
Enable Guest Account	Check the Enable Guest Account checkbox to allow unknown users or users explicitly logging in as “guest” to access the SnapServer using the guest account. Clear the option to disable this feature.
Enable Opportunistic Locking	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the checkbox to disable opportunistic locking.

Option	Settings
Enable this Server as the Master Browser	Enabled by default. The SnapServer can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the checkbox if you plan to install this server in a Windows environment and you want this server to be able to serve as the Master Browser for a workgroup. Clear the checkbox to disable this feature.
Allow Root Authentication	Check the Allow Root Authentication checkbox to allow a root login on the selected server.
Disable NetBIOS over TCP/IP	Some administrators may wish to disable NetBIOS over TCP/IP. Select the checkbox to disable NetBIOS; clear the checkbox to leave NetBIOS enabled. Note: If you disable NetBIOS and you are joining a domain, you must enter the domain name as a fully qualified domain name (for example, actdirdomainname.companyname.com). A short form such as ActDirDomName will not work.
Enable Trusted Domains	SnapServers recognize trust relationships established between the domain to which the SnapServer is joined and other domains in a Windows environment by default. Select the checkbox to toggle this feature. Note: SnapServers remember trusted domains. That is, if this feature is disabled and then activated at a later time, the previously downloaded user and group lists, as well as any security permissions assigned to them, will be retained.

To Join an Active Directory Domain

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

Option	Description
Member Of	Select <i>Active Directory Domain</i> .
Workgroup/Domain Name	Enter the name of the domain. Note: Windows 2000 domain controllers must run SP2 or later.
Organizational Unit	To create a machine account at a different location than the default, enter a name in the Organizational Unit field. By default, this field is blank, signaling the domain controller to use a default defined within the controller. Note: Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([organizationalunit]/[sub-unit1]/[sub-unit1a])
Disable NetBIOS over TCP/IP	Some administrators may wish to disable NetBIOS over TCP/IP. Select the checkbox to disable NetBIOS; clear the checkbox to leave NetBIOS enabled. Note: If you disable NetBIOS, you must enter the domain name as a fully qualified domain name (for example, actdirdomainname.companyname.com). A short form such as ActDirDomName will not work.
Administrator Admin Password	Enter a user name and password with sufficient administrative privileges to allow a remote computer to join the domain.

Option	Description
Enable Trusted Domains	<p>SnapServers recognize trust relationships established between the domain to which the SnapServer is joined and other domains in a Windows environment by default. Select the checkbox to toggle this feature.</p> <p>Note: SnapServers remember trusted domains. That is, if this feature is disabled and then activated at a later time, the previously downloaded user and group lists, as well as any security permissions assigned to them, will be retained.</p>
Administrator Admin Password	Enter a user name and password with sufficient administrative privileges to allow a remote computer to join the domain.

NFS Access

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

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

NFS client access to shares can be specified by navigating to the **Security > Shares** screen and clicking the **NFS Access** link next to the share. Because you are in Unicode mode, you must configure the SnapServer's protocol for the code page being used. See "NFS" on page 10-2 for more information.

Support for NFS

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

Assigning Share Access to NFS Users

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

Supported Protocols

SnapServers support these versions of the NFS protocol:

Protocol	Version	Source
NFS	2.0, 3.0, 4.0 ¹	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

1. NFSv4 ACLs are not supported.

Procedures

To Enable NFS Access to the Server

Check the **Enable NFS** checkbox to enable NFS, then check the version(s) you want to enable (NFSv2, NFSv3, and NFSv4). Leave the checkbox blank to disable access to this server via the NFS protocol.

Choose the desired Unicode **Client code page** from the drop-down menu, and then click **OK** to update NFS network settings.

To Configure NFSv4 Access to the Server

1. Check the **Enable NFS** and **Enable NFSv4** checkboxes.

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

2. Select the security you want to apply:

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

3. If you select Unix Kerberos security, complete the following:

Note: The service will not start unless the TCP/IP domain name ([page 2-3](#)) is set up exactly the same as the ketchup.

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 Note: This file can have any name the administrator wishes to give it. If you do not have a keytab file for the SnapServer: -- create a host and NFS principle for the SnapServer on the KDC -- generate a keytab file -- save it to a location the client administering the SnapServer can access.

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

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

To Mount Shares from NFS Clients

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

1. To mount an NFS client, enter one of the following commands:

- a. To mount via NFSv2 or NFSv3:

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

where *server_name* is the name or IP address of the server, *share_name* is the name of the share you want to mount, and *local_mount* is the name of the mount target directory.

- b. To mount via NFSv4 with standard NFS host-based security, modify the above command as follows:

```
mount -t nfs4 server_name:/share_name /local_mount
```

- c. To mount via NFSv4 with Kerberos security, enter one of the following commands, depending on the level of security desired:

- Authentication only (client is authenticated via Kerberos):

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

- Authentication and integrity (client is authenticated and data contains a cryptographic checksum that guarantees the data has not changed):

```
mount -t nfs4 -o sec=krb5i server_name:/share_name /  
local_mount
```

- Authentication, integrity, and privacy (client is authenticated, data integrity is guaranteed, and data itself is encrypted):

```
mount -t nfs4 -o sec=krb5p server_name:/share_name /  
local_mount
```

Note: Syntax can vary depending upon the operating system.

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

Apple Networking Configuration

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

Note: Mac OS X users can also connect to the SnapServer using Windows networking (SMB). See [“Connecting a Mac OS X Client Using SMB” on page 2-11](#).

AFP Configuration Guidelines

Consider the following when configuring access for your AFP clients.

Terminology

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 Chooser. 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 file system.
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 ." Mac users 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** 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 Chooser 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** screen of the Web Management Interface is enabled by default. This option allows Mac applications 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.

Procedures

To Edit AFP Access

Edit settings as described in the following table, and then click **OK** to update network AFP settings immediately.

Options	Usage
Enable Apple Filesharing (AFP)	Check the Enable Apple Filesharing (AFP) checkbox to enable AFP; leave the checkbox blank to disable AFP access.
Add Server Name to Apple Shared Folder Names	Select this option to identify shares to AFP clients using both the server name and share name. Clear the checkbox to display only the share name.

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's servername. For example: <i>windowsdomain\username</i> or <i>snap12345\username</i>

To Connect from a Mac (OS X)

Using the Connect to Server window, enter one of the following:

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

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

FTP/FTPS Access

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

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

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

Supported FTP Clients

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

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

Procedures

To Configure FTP/FTPS Access

Edit settings as described below, and then click **OK** to update network FTP settings immediately.

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

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

To Connect via FTP/FTPS

To connect to the server through standard FTP, enter the server's name or IP address in the FTP Location or Address box of a web browser or FTP client application.

- To connect via a command line, enter:
ftp *server_name*
- To connect via a Web browser, enter:
ftp://*server_name*
(where *server_name* is the name or IP address of the server)

To connect to the server through FTPS:

- Configure your FTPS client application to use explicit FTPS (such as, FTPES or "Auth TLS").
- Enter the SnapServer's server name or IP address.

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.

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.

HTTP/HTTPS Access

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

Note: To access the CA Antivirus configuration interface (on the **Snap Extensions** screen), HTTP must be enabled.

GuardianOS supports the following browsers: Microsoft Internet Explorer (7.0 or later), Apple Safari (2.0 or later), Mozilla FireFox (2.0 or later), and Google Chrome (1.0 or later).

Configuring HTTP/HTTPS

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

To Require Web Authentication

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

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

To Enable HTTP Access to the Server

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

Option	Description
Enable (non-secure) HTTP Access	Check the Enable HTTP Access checkbox to enable non-secure HTTP access. Leave the checkbox blank to disable access to the server via HTTP. Notes 1. This option applies to both Web View and Web Root modes. 2. To access the CA Antivirus configuration interface, HTTP must be enabled.

To Connect via HTTPS or HTTP

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

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

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

2. Press Enter.

The Web View screen opens.

Using WebRoot to Configure the SnapServer as a Simple Web Server

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

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

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

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


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

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

Configuring Web Root

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

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

2. Map a drive to the share you have designated as the web root share and upload your HTML file(s) to the root of the directory, making sure the file name(s) is listed in the Index Files box.

Accessing the Web Management Interface when Web Root is Enabled

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

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


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

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

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

Web View

Web View opens when the user accesses a SnapServer using a Web browser, unless the administrator has enabled the Web Root feature (see “[Using WebRoot to Configure the SnapServer as a Simple Web Server](#)” on page 2-19). This screen displays a list of all shares to which the user has access. Users can navigate the share structure to locate and view or download files, but they cannot modify or upload files.

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

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

To Switch to a Different User

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

To Change a User Password

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

Print Server

The SnapServer can be configured to emulate a Windows print server for locally-attached USB printers. 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 SnapServer as a print server is a two-part process:

- Configuring the printer on the SnapServer.
- Configuring the client to print via the SnapServer.

Configuring the Printer

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

1. Connect a printer to one of the USB ports on the SnapServer.
2. Power ON the printer.
3. In the SnapServer's Web Management Interface, navigate to **Server > Printing**. A list of currently defined USB printers is displayed. To add the new printer, click **Add Local Printer**.
4. The SnapServer will detect the new printer and it should appear as an option in the **Local Printer Device** drop-down list. Select that printer.

5. Give the printer a name, and complete Description and Location information as desired. Click **OK**. The printer will appear in the list on the main printing page.

Configuring the Client

Next, add the printer to a Windows, Mac, or Linux client, enabling you to print via the Snap Server. 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 set up.

Adding the 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.

Adding the 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.

Monitoring Print Jobs Remotely

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

To Pause the Printer

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. Click the **Pause Printer** button to pause all print jobs.

Note: When the printer is paused, the button will become a **Resume Printer** button, which you can click to resume printing.

To Cancel or Check the Status of Print Jobs

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

Deleting a Printer

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

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

Authentication validates a user's identity by requiring the user to provide a registered login name (User ID) and corresponding password. SnapServers ship with predefined local users and groups that allow administrative and guest user access to the server via all protocols. Administrators may choose to join the SnapServer to a traditional Windows Active Directory domain, and Windows clients can then authenticate to the server using their domain credentials. To accommodate NFS clients, the SnapServer can also join an NIS domain, and the SnapServer can look up user and group IDs maintained by the domain. For authentication control beyond the guest account, Mac and FTP client login credentials can be created locally on the sever.

Topics in Users and Groups:

- [Default User and Group Settings](#)
- [UID and GID Assignments](#)
- [Local Users and Groups](#)
- [NIS Domain](#)

Default User and Group Settings

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, "guest" is disabled in SMB.

A *local user* or *group* is one defined locally on a SnapServer using the Web Management Interface. The default users and groups listed below cannot be modified or deleted.

Default Local Users and Groups	
admin	The admin user account is used to log into the Web Management Interface. The default password for the admin account is also <i>admin</i> .
guest	The guest user account requires no password.
AllLocalUsers	The AllLocalUsers group account includes all local users created on the SnapServer.
AllUsers	The AllUsers group account includes all local, Windows domain, and NIS users.
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.
Domain	
Windows	The SnapServer can join a Windows Active Directory domain.
NIS	The SnapServer can join an NIS domain and function as an NIS client.

UID and GID Assignments

The SnapServer uses the POSIX standard to assign a user ID (UID) and group ID (GID), in which each user and group must have an ID. This requirement applies to all users and groups on the SnapServer, including local, Windows, and NIS users and groups.

If you join the SnapServer to a Windows or NIS domain, IDs are automatically assigned on a 'first come, first served' basis. 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 in use by NIS or the Windows domain, or that is less than 101, you will get an error message.
- When the server automatically generates UIDs or GIDs for imported Windows domain users or groups, UIDs or GIDs that are already in use by local and NIS users will be skipped.
- When NIS domain users and groups are imported, the SnapServer will discard any UIDs that are less than 101 or are in conflict with UIDs already in use by local or Windows domain users and groups.

The NIS “nobody” user IDs (UID 65534 and 65535) are reserved. They are not mappable to other IDs, nor is another ID mappable to “nobody.”

GuardianOS offers ID Mapping, which allows mapping of Windows users to local or NIS users to provide unified permission assignments to users of different protocols. For more information on ID Mapping, please see [“ID Mapping” on page 6-8](#).

Local Users and Groups

Local users or groups are created using the **Security > Local Users** and **Security > Local Groups** screens in the Web Management Interface. Local users and groups are used for

administrative (and non-domain guest) access to the server. Network clients can initially access the server using the guest account, but if you require a higher degree of control over individual access to the file system for these clients, you must create local accounts (or, in the case of Windows, use Windows Active Directory security).

This information is divided into three sections:

- Guidelines for Local Authentication
- Local User Management Procedures
- Local Group Management Procedures

Guidelines for Local Authentication

These password authentication guidelines are for both users and groups.

Duplicating Client Login Credentials for Local Users and Groups

To simplify user access for Windows Workgroup or Mac clients, duplicate their local client logon credentials on the SnapServer. That is, create local accounts on the SnapServer that match those used to log on to client workstations. This strategy allows users to bypass the login procedure when accessing the SnapServer.

CAUTION: This strategy applies only to local users. Do not use duplicate domain user credentials if joined to an Active Directory domain.

Default Local Users and Groups

Default users and groups *admin*, *guest*, and *admingrp* appear on the list of users or groups on the User or Group Management screens, but they cannot be deleted or modified (although the admin password can be changed). As you would expect, the default local users and groups do appear on the Share Access and Quotas screens.

Changing Local UIDs or GIDs

The SnapServer automatically assigns and manages UIDs and GIDs. Because you may need to assign a specific ID to a local user or group in order to match your existing UID/GID assignments, the SnapServer makes these fields editable.

Password Policies

To provide additional authentication security, set password character requirements, password expiration dates, and lockout rules for local users.

Local users can also be individually exempted from password expiration and character requirement policies. The built-in *admin* user is exempt from all password policies.

Note: Local users with expired passwords can change their passwords at <http://<snapservername>/changepassword>.

Local Account Management Tools

The SnapServer offers the following tools for creating, modifying, and editing local user and group accounts.

Function	Navigation Path
Local User Management	Navigate to the Security > Local Users screen, from which you can create, view, edit, and delete local users. You can also set user password policy, including password character requirements, maximum number of allowed logon failures, and password expiration settings.
Local Group Management	Navigate to the Security > Local Groups screen, from which you can create, view, edit, and delete local groups.

Local User Management Procedures

Use these procedures when creating or configuring a local user.

To Create a Local User

1. On the **Security > Local Users** screen, click **Create**.
2. On the screen that opens, enter the following information:

Option	Description
Name	Use up to 50 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.
Password Verify	Type the chosen password again for verification.
User ID (UID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see "UID and GID Assignments" on page 3-2 .
Disable User Login	Select this checkbox to disable the user login. The user's information will remain in the system, but login rights will be denied. The user login can be enabled by deselecting the checkbox. This checkbox can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
Exempt from Password Expiration and Character Requirements	Select this checkbox to exempt this user from password expiration and character requirement policies. Note: This checkbox is only visible if Password Policy is enabled.
Grant Admin Rights	Select this checkbox to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **OK**.

To Edit Local User Account Information

1. On the **Security > Local Users** screen, select the user you want to edit and click **Properties**.

2. On the screen that opens, enter the following information:

Option	Description
Name	Cannot be modified.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case sensitive. Use up to 15 alphanumeric characters. Leave this field blank to keep the existing password.
Password Verify	Type the chosen password again for verification. Leave this field blank to keep the existing password.
User ID (UID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see “UID and GID Assignments” on page 3-2 . Note: Changing a user's UID may alter file system 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	Select this checkbox to disable the user login. The user's information will remain in the system, but login rights will be denied. The user login can be re-enabled by deselecting the checkbox. This checkbox can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
Exempt from Password Expiration and Character Requirements	Select this checkbox to exempt this user from password expiration and character requirement policies. Note: This checkbox is only visible if Password Policy is enabled.
Grant Admin Rights	Select this checkbox to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click OK.

To Set Password Policy for Local Users

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 3-4](#) for procedures to set password policy for local users. Also, the built-in *admin* user is automatically exempt from all password policies.

1. On the **Security > Local Users** screen, click the **Password Policy** button.
2. On the screen that opens, enter the following information:

Option	Description
Enable Password Policy	Select this checkbox to enable password policy for local users, then select from the following restrictions.
Character Requirements	Select the alpha/numeric/special character requirements for the password from the drop-down list.
Minimum Number of Characters	Check the checkbox to enable the policy, then enter the minimum number of characters required for the password.

Option	Description
Login Attempts Permitted	Check the checkbox to enable the policy, then enter the number of times a user can fail to login before the system locks the user out. Note: To unlock a user, clear the Disable user login checkbox for the user in the Local Users page.
Minutes to Automatically Re-enable a Disabled Login	If you have defined a limit to the number of times a user can fail to log in, you can also check this checkbox and enter a time period after which the system will allow the user to log in again. Note: This will save the administrator from having to manually re-enable the user.
Days until Password Expires	Check the checkbox to enable the policy, then enter the number of days before the password must be changed. Note: Local users with expired passwords can change their passwords at <a href="http://<snapservername>/changepassword">http://<snapservername>/changepassword .

3. Click **OK**.

To Add or Remove Users from Groups

1. On the **Security > Local Users** screen, select a user and click **Groups**.
2. Displayed are the group settings for the selected user. To add the user to a group, select the group and click **Add**. To delete the user from a group, select the group and click **Remove**.
3. Click **OK** to save your changes and return to the Local Users screen.

To Delete Local Users

1. On the **Security > Local Users** screen, select the user to be deleted and click **Delete**.
2. The confirmation screen will display. Click **Yes** to delete the selected user. Click **No** to cancel the deletion.

Local Group Management Procedures

Use these procedures when creating or configuring a local group.

To Create a new Local Group

1. On the **Security > Local Groups** screen, click **Create**.
2. On the screen that opens, enter the following information:

Option	Description
Group Name	Use up to 24 alphanumeric characters and the underscore.
GID	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see "UID and GID Assignments" on page 3-2 .

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

To Edit Local Group Properties

1. On the **Security > Local Groups** screen, select the group you want to edit and click **Properties**.
2. On the screen that opens, you can change the GID. For information on available UID ranges, see [“UID and GID Assignments” on page 3-2](#).

Note: Changing a group's GID may alter file system 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**.

To Add Users to or Remove Users from a Group

1. After creating a new group, or when editing an existing group, add and remove users by selecting the desired group and clicking **Users**.
2. Add users by selecting the user and clicking **Add**.
3. Delete users by selecting the user and clicking **Remove**.
4. Click **OK** when finished.

To Delete a Group

1. On the **Security > Local Groups** screen, select the group to be deleted and click **Delete**.
2. The confirmation screen will display. Click **Yes** to delete the selected group, or click **No** to cancel the deletion.

NIS Domain

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

Guidelines for Configuring NIS

Unless UID/GID assignments are properly handled, NIS users and groups may fail to display properly. For guidelines on integrating compatible SnapServer UIDs, see [“UID and GID Assignments” on page 3-2](#).

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

Procedure

To Join an NIS Domain

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

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

Note: Much of the configuration discussion presented here is not applicable to SnapServers with fewer than four (4) drives. For SnapServer 110 and 210, see the *User's Guide for SnapServer 110 and 210* for storage configuration guidelines.

Topics in Storage and Expansion:

- [Storage Configuration Overview](#)
- [RAIDs](#)
- [Volumes](#)
- [Quotas](#)
- [Data Migration](#)
- [Maintenance Tools](#)
- [Expansion Arrays](#)
- [Disks and Units](#)

Storage Configuration Overview

The SnapServer's flexible storage architecture allows for a wide variety of implementations to suit many different storage needs. In some cases, administrators can configure storage to maximize capacity by modifying the configuration of the SnapServer, or attaching one or more expansion arrays.

RAID 5 and RAID 6 typically give the best balance of performance, capacity, and data protection in most cases. However, in some cases when backup is ongoing or very frequent, a RAID 0 configuration may be most appropriate to maximize capacity and performance. In cases where multiple expansion arrays are attached to the SnapServer, a combination of a RAID 1 or RAID 10 and hot spares may be the optimal configuration for the SnapServer.

Storage Configuration Examples

The following examples and procedures detail three common alternative storage configurations:

- [Snapshot Space Considerations](#)
- [Creating a RAID Configuration](#)
- [Adding Capacity to a SnapServer with Expansion Arrays](#)

Snapshot Space Considerations

A snapshot is a consistent, stable, point-in-time image of a volume (filesystem). The snapshot provides a way for the administrator to restore a volume to a previous state without resorting to tape backup. Because snapshots depend on successive snapshots for part of their content, enough disk space must be reserved within a RAID for the storage of multiple snapshot data sets. In the default storage configuration of many SnapServers, twenty percent of the RAID capacity is allocated to the snapshot pool.

For more details about snapshots, refer to [Chapter 7, "Snapshots."](#)

Creating a RAID Configuration

The basic building block of the server's storage is the disk array, or RAID. There can be one or more RAID's per server, each containing one or more disks. Disks must be members of a RAID in order to make their storage available, but an individual disk can be configured into single-drive RAID 0 to replicate a JBOD configuration.

Initial Setup prompts you to create a RAID using disks on the head unit with a RAID level of your choice. A volume is then automatically created consuming 80% of the RAID, and the remaining 20% is allocated for the snapshot space.

Note: Alternatively, you can skip this portion of Initial Setup and manually configure your storage using the Web Management Interface. Manual configuration is also required to configure storage on expansion units, and to execute more advanced storage configurations.

Example

In the following example, two-drive mirrors (RAID 1s) are created on a SnapServer 650. Capacity is reduced in the head unit to provide redundancy and high data availability. Capacity is regained by adding expansion arrays in the following procedure. This configuration can sustain up to two drive failures, making it more fault tolerant than a RAID 5. This procedure does not retain any of the data already present on the server.

1. On the **Storage > RAID Sets** page, click **Create RAID Set** to create a new RAID.
2. Create a two-disk RAID 1 with Drive 1 and 2 for the GuardianOS, antivirus, and backup software to reside on.
3. Create a volume and share for the RAID 1.
If user data will be stored on this volume, configuring snapshots, antivirus, and a Disaster Recovery Image is strongly recommended.
4. Create a second two-disk RAID 1 with Drives 3 and 4.
Configure the volume without snapshots for use as an iSCSI target, or to store data, optionally configuring snapshots and a Disaster Recovery Image.

Adding Capacity to a SnapServer with Expansion Arrays

When adding an expansion array, a RAID and volume must be created on the expansion array before using the storage space. The existing volume on the host SnapServer can be extended to provide more capacity for network users in a seamless fashion, or a new volume can be created to host data. In the following procedure, an expansion array will be used to increase the capacity of a SnapServer with a four-drive RAID 5 on the SnapServer and an unconfigured expansion array. This procedure will retain any data that is already present on the server.

1. Shut down the SnapServer and attach the expansion array, as described in the expansion array's *Quick Start Guide*.
2. Power on the expansion array and restart the SnapServer.
3. Log into the Web Management Interface for the SnapServer.
4. When the server has completed booting, navigate to **Storage > Volumes** and click the volume you want to extend.
5. Check the boxes for four or more disks on the expansion array, click **Next**, then click **Yes** on the confirmation screen.

RAIDs

RAIDs are created, viewed, edited, and deleted from the **Storage > RAID Sets** screen of the Web Management Interface. Before configuring RAIDs, consider the following information on the SnapServer's RAID implementation.

Note: Much of the configuration discussion presented here is not applicable to SnapServers with fewer than four (4) drives. For SnapServer 110 and 210, see the *User's Guide for SnapServer 110 and 210* for storage configuration guidelines.

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 hot 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 hot 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 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.

Local and Global Hot Spares

A *hot spare* is a disk drive that can automatically replace a damaged drive in a RAID 1, 5, 6, or 10. Designating a disk drive as a hot 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 hot spare to rebuild itself without administrator intervention. SnapServers offer two kinds of hot spares: local and global.

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

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 RAID configurations, a hot-inserted non-GuardianOS or unassigned drive will be automatically configured as a global hot spare. To enable automatic incorporation of unassigned drives, go to the **Storage > RAID Sets** screen and click the **RAID Settings** button.

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 the I/O activity falls to a very low disk activity. Once the activity rises above the *idle threshold*, the background scan stops and waits for the activity to fall to the idle threshold again before resuming. As a result, there should be minimal to no impact on performance. Once the disk scan has completed a pass on a given RAID set, it waits a certain 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** screen and click the **RAID Settings** button.

Notes:

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

RAID Management Tools

SnapServers use the following tools for configuring and monitoring RAIDs:

Function	Navigation Path
Ongoing Maintenance	Navigate to the Storage > RAID Sets screen, from which you can create, assess, edit, and delete RAIDs. You can also disable or enable the Background Disk Scan and the automatic assignment of GuardianOS-partitioned unused disks to a degraded RAID.
Email Notification	The server can notify you when a RAID is degraded, failed, or has experienced another error or maintenance condition. This allows you to take action to ensure workflows are not disrupted (Server > Email Notification).

You can view the status of your RAID sets on the **Storage > RAID Sets** and **Monitor > System Status** screens.

RAID Set Procedures

To Assess RAID Status

You can view the status of your RAID sets on the **Storage > RAID Sets** screen. Click the RAID set to see additional status details about member drives of that specific RAID.

Label	Description
RAID Set	The name of each RAID

Label	Description
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 that the RAID was configured with.
Group	The name of the RAID Group to which the RAID belongs
Size	The total capacity of the RAID
Unallocated	The total storage space not allocated to a volume

To Create a New RAID

Click **Create RAID Set**. The process involves first selecting a RAID type, and then selecting the disk drives to include in the RAID.

To Remove, Repair, or Add Disk Drives or Hot Spares to a RAID

Click a RAID name. You can remove, repair, or add disk drives to an existing RAID 1, 5, 6, or 10.

To Delete a RAID

Deleting a RAID deletes all data stored on its disk drives and removes all volumes and shares using the RAID. If you have already begun using your SnapServer, you must back up all data before you delete a RAID (or a volume).

Click a RAID set name. On the next screen, click **Delete RAID**, and then follow the onscreen instructions.

To Manage Global Spares

Click **Global Spares**. The Global Spares screen displays all disks available for use, or that are in use, as global spares.

To enable a disk as a global spare, check the checkbox next to the desired disk and click **OK**. More than one disk can be checked at a time.

To disable or delete a disk assigned as a global spare, clear the checkbox next to the disk and click **OK**.

To Disable the Background Disk Scan

The background disk scan is enabled by default. To disable it, click the **RAID Settings** button. On the RAID Settings screen, click to clear the checkbox next to **Enable continuous disk scan during idle I/O system time**, then click **OK**.

To enable the disk scan, click to add a check in the box and click **OK**.

To Enable the Automatic Assignment of Unused Disks

Click the **RAID Settings** button. On the RAID Settings screen, click the checkbox next to **Enable automatic incorporation of an unused disk into a degraded RAID set**, then click **OK**.

To disable the automatic incorporation, clear the checkbox and click **OK**.

Creating a New RAID

The SnapServer offers two methods of creating a new storage configuration: (1) the RAID Storage Guides let you quickly create a RAID, a volume to the RAID, and a single share to the volume; or (2) a sequence of Storage screens requires a little more time, but allows you full control over volume, share, and share access definitions. The wizard is the fastest way to create a RAID, but the Storage screens offer more choices and greater flexibility.

- To create a RAID using the wizard, navigate to the **Storage > Storage Guides** screen and follow the onscreen instructions.
- To create a new RAID, navigate to the **Storage > RAID Sets** screen, and click **Create RAID**.

Step 1: Select the RAID Type

Choose the type of RAID you want. For information about the different RAIDs available, please see [“Factors in Choosing a RAID Type” on page 4-3](#).

Note: Each RAID type requires a minimum number of available drives. A RAID type will be available only if its minimum drive requirement is met.

Step 2: Choose Devices

1. Select the drives you want to include in the RAID.

CAUTION: In some SnapServer storage system configurations, you may be able to select disk drives of different sizes. Because all drives within a RAID must be the same capacity, using mixed-capacity drives in the same RAID will result in wasted capacity. For example, if a RAID is configured with the drives listed in the following table, over half of the capacity of the larger drives will go unused.

Drive	Raw Capacity	Actual Used Capacity	Usage
Drive1	160 GB	160 GB	100%
Drive 2	160 GB	160 GB	100%
Drive 4	250 GB	160 GB	64% of 250 GB
Drive 6	250 GB	160 GB	64% of 250 GB

2. If you are creating any RAID except a RAID 0, you can configure one of the selected drives as a hot spare by selecting the *I want a local spare* or *I want a global spare* option button. If you do not want a hot spare, select *I do not want a hot spare*.

Note: For more information on local vs. global hot spares, see “Local and Global Hot Spares” on page 4-4.

3. Click **Next** to confirm your configuration.

Step 3: Confirm RAID Configuration

In most cases, this screen simply calculates the capacity of the RAID you are about to create. Review the settings and do one of the following:

- To change the configuration, click the **Prev** button.
- Click **Next** to create the RAID

If you selected mixed-capacity drives for the RAID, the interface calculates the amount of disk space that will go unused. Do one of the following:

- To change the configuration, click the **Prev** button.
- To ignore the warning and create the RAID, click **Next**.
- To have the system attempt a more efficient configuration, select the *View System Recommendation* checkbox and click **Next**. The screen refreshes, and the system may display an alternative configuration in addition to your original configuration. If so, select the one you want using the option buttons at the bottom of the screen, and click **Next**.

Step 4: Done

Do one of the following:

- To create a volume for the new RAID, click **Create Volume Now**.
- To return to the RAID screen, click **Create Volume Later**.

Managing Disk Drives in a RAID

To manage disk drives or delete a RAID, navigate to the **Storage > RAID Sets** screen, and click the name of a RAID. You can remove, repair, or add disk drives to any existing RAID set except a RAID 0.

To Remove a Disk Drive (Member or Hot Spare)

Click the **Remove** link in the Actions column. (This option is not available for a RAID 0.)

CAUTION: Removing a disk drive causes the RAID to run in degraded mode and puts your data at risk.

To Repair a RAID

Click **Repair** and follow the onscreen instructions to select another drive to incorporate into the RAID as a replacement for a failed member drive. This button appears only when a drive has failed or been removed, and the RAID is in degraded mode. The RAID Status now reads *Resyncing* while the new drive is incorporated into the RAID. It reads *OK* once the incorporation is complete.

To Delete a RAID

Click the RAID you want to delete, click **Delete RAID**, and then follow the onscreen instructions.

To Add a Disk Drive (Member or Hot Spare) to a RAID

Click **Add Disk**, and then follow the onscreen instructions. If you are adding a drive to a RAID 1, you have the option to designate the drive as a global hot spare, a local hot spare, or a member. If you are adding a drive to a RAID 5, 6, or 10, only the hot spare options are available.

Using the RAID Storage Guides

RAID types available for creation on your SnapServer appear as active links on this page. To start the RAID Storage Guides, click the RAID type you want to create. A list of available disk drives appears.

Note: Each RAID type requires a minimum number of available drives. A RAID type will be available only if its minimum drive requirement is met.

Step 1: Create RAID Set

After selecting the RAID type, select the drives for the RAID from the list provided. If you are creating any RAID except a RAID 0, you can configure one of the selected drives as a hot spare by selecting **I want to add a local spare** or **I want to add a global spare**. If you do not want to add a hot spare, select **I do not want a hot spare**. Click **Next**.

Note: For more information on the two types of hot spares, see [“Local and Global Hot Spares” on page 4-4](#).

Step 2: Confirm RAID Set Configuration

Review the information on this screen to verify the RAID configuration, and then click **Next** to open the Create Volume screen.

Step 3: Create a Volume

The volume capacity defaults to the entire storage space of the RAID. You can enter a lesser number if you plan to create additional volumes on the RAID. The space reserved to store snapshot data defaults to 20%. Adjust this figure as necessary, click **Next** and then click **Yes** to create the volume and open the Create Share screen.

Note: For assistance in estimating your snapshot storage requirements, see [“Estimating Snapshot Pool Requirements” on page 7-2](#).

Step 4: Create a Share to the Volume

Make a note of the share name and click **Next** to create a share that allows access to all users via all protocols. (You can modify the share access definition by navigating to the **Security > Shares** screen and clicking the name of the share.)

Step 5: Done

Review the new storage configuration and click **OK** to return to the main Storage screen.

RAID Groups

Two RAID groups 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 groups and allows users to continue working as usual with no interruption.

Grouped RAID groups must be the same type; you can group two RAID 1s or two RAID 5s (for example, you cannot group a RAID 1 and a RAID 5).

Note: Only RAID groups of the same PE (physical extent) size can be grouped. If you are growing the volume on one RAID to use free capacity on another RAID, you will only be allowed to select from those RAID groups that can be grouped.

Also consider the following:

Adding an Expansion Array

In a common scenario, a four-drive SnapServer configured as a RAID 5 is nearing full utilization. The administrator decides to add an expansion array. The administrator creates a RAID 5 on the expansion array, groups it with the existing RAID on the SnapServer, and then expands the size of the original volumes using the new storage from the expansion array.

Grouping RAID groups with other Grouped RAID groups

Just as RAID groups can be grouped, individual groups of RAID groups can be brought together to form an even larger group. For example: A 1 TB SnapServer is running out of capacity. Two 1 TB 12-drive expansion arrays are attached to the SnapServer to provide increased capacity. You can configure a RAID 5 on each of the expansion arrays, then group them together. The resulting RAID group can then be grouped with the RAID on the SnapServer, allowing network users to take advantage of the full capacity of the head and expansion arrays with no loss of capacity.

Deleting Grouped RAID groups

Deleting the RAID Group will delete all member RAID groups, all their volumes and shares, and all their data. If one RAID becomes inaccessible for any reason, the entire RAID group 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 host SnapServer and an expansion array and one of the RAID groups 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 RAID groups becomes corrupted and remains unrecoverable, the data in the other RAID will also be lost.

Snapshot Pools are Combined

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

Two RAID groups at a Time

To group more than two RAID groups, create a RAID group with two RAID groups, then group the RAID group with each RAID one at a time.

RAID Group Procedures

To Assess RAID Group Status

You can view your RAID group status from either the **Storage > RAID Sets** or **Monitor > System Status** screen.

Label	Description
RAID Set	The name of each RAID is listed.
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 is offline. • <i>Degraded</i> – A drive has failed or been removed.
Group	The name of the RAID Group to which the RAID belongs
Size	The total capacity of the RAID is listed.
Unallocated	The total storage space not allocated to a volume or the RAID's snapshot pool is listed.

To Create a RAID Group

The **Group RAID** button becomes available when there are RAIDs of the same type on the server able to be grouped. Click **Group RAID** and then follow the onscreen instructions.

Note: Only RAIDs of the same PE size can be grouped. The Web Management Interface will notify you if you attempt to group two RAIDs with different PE sizes.

To Add a RAID to an Existing RAID Group

Click a RAID group name, and then click **Add RAID**. Existing RAIDs of the same type appear on the screen that opens. Select a RAID, click **Next**, and then follow the onscreen instructions.

To Delete a RAID Group

Deleting a RAID group deletes all data stored on its disk drives. If you have already begun using your SnapServer, you must back up all data before you delete a RAID (or a volume).

Click **Delete Group**, and then follow the onscreen instructions.

Volumes

Volumes are created, viewed, edited, and deleted from the **Storage > Volumes** screen of the Web Management Interface. Consider the following facts and guidelines when planning your volume configuration.

More details are available in the “[Procedures](#)” section of Volumes.

Volumes and the Snapshot Pool

The default capacity settings for the filesystem and future snapshot use is 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 Pool Requirements](#)” on page 7-2.

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 5-15.

Deleting Volumes

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

The NetVault for GuardianOS Database Directory (NVDB), containing files that keep track of the data you back up; the 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, NetVault will need to be reconfigured, Antivirus will need to be re-enabled, and Snap EDR will need to be reinstalled when a new volume with sufficient space exists.

To retain NVDB information, you must back up the NVDB directory (see [page 8-4](#)) before you delete the volume, create your new storage configuration, and then restore the directory.

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

To reactivate Snap EDR functionality after creating a new volume, download the Snap EDR package from the SnapServer website and install it on the server using the OS Update feature. Then click the **Snap EDR** link in the Site Map (under Misc.) and click the **Start** button.

Note: If you delete a volume, you will also delete any iSCSI disks that reside on that volume.

Expanding Volume Capacity

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

- **Adding Unallocated Capacity** – If there is unallocated capacity remaining on the RAID, you can add this capacity to the volume by editing the Volume size field or clicking the **Grow to Max. Size** button, and then clicking **OK**.
- **Creating a New RAID** – If all capacity on the RAID is allocated, and either: (1) a sufficient number of drives to create a new RAID exists, or (2) a RAID of the same type with excess capacity exists, the **Expand Volume** button appears. Click this button to create an additional RAID, group the RAID with the existing RAID, and expand the volume into the space on the new RAID.

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

A volume can be expanded up to 16 TBs, either as a standalone volume or as a volume group.

Security Models, SnapTrees, and Volumes

Note: For detailed information on SnapTrees and file access, see [“Chapter 6, “Share and File Access.”](#)

Volumes are created with the Windows security model (which can be changed in the **Security > SnapTrees** page or when creating a share to point to the volume root). Directories created in the root of a volume (SnapTree directories) in the Web Management Interface are automatically assigned either a Windows- or a UNIX-style security model, based on the security model of the parent volume (this can also be subsequently changed in the SnapTrees page or when creating a share pointing to them). The security model determines the file-level security scheme that will apply to files and folders within the volume or SnapTree directory.

Configuring Write Cache

Note: Not related to write cache on iSCSI disks. For information about configuring write cache on iSCSI disks, see [“Write Cache Options with iSCSI Disks”](#) on page 5-11.

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

Notes:

- 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.
- Not all disk drives support disabling write cache. If any of the volume's drives are IDE drives, you will not have the option to disable write cache for that volume.

Checking Filesystems

Filesystems on individual volumes can be checked for errors and repaired, if necessary. The root volume can also be checked, and any errors found will automatically be repaired. Since GuardianOS automatically checks the root volume for errors if any of a number of triggers occurs (for example, a power outage, failure of the volume to mount, etc.), it is recommended that the root filesystem check feature only be used when directed by a Technical Support representative.

See [To Check the Filesystem on a Volume](#) for the procedure to check a volume's filesystem. See [“To Check the Root Filesystem”](#) on page 4-23 for the procedure to check the root filesystem.

Volume Management Tools

The SnapServer offers several tools for monitoring and controlling how storage space on a volume is used.

Maintenance Tools

Function	Navigation Path
Ongoing Maintenance	Navigate to the Storage > Volumes screen, from which you can create, view, edit, and delete volumes.
Email Notification	The server can notify you when a volume is full. This allows you to increase volume size or take other actions to ensure workflows are not disrupted (Server > Email Notification).
Volume Usage	You can view the current utilization totals for each volume, from the Storage > Volumes screen.
Quotas	Use quotas (Storage > Quotas) to limit the amount of storage space on a volume that specific users or groups can consume. See Quotas for more information.

Procedures

To Create a New Volume

Click **Create Volume**. The process involves first defining volume parameters, then confirming your settings.

To Assess Volume Status

You can view volume status from the **Storage > Volumes** screen.

Label	Description
Volume	The name of each volume
RAID Set	The RAID on which the volume was created
Status	Current condition of the volume: <ul style="list-style-type: none"> • <i>Active</i> — The volume is online and accessible. • <i>Inactive</i> — The volume is offline. • <i>Rollback</i> — A snapshot rollback operation is in progress.
Used %	Indicates the current utilization percentage for the volume
Size	Current total capacity of the volume
Free	Amount of free space remaining in the volume
Quotas	Indicates whether quotas are enabled or disabled

You can use the **Server > Email Notification** screen to set up an alarm to notify you when a volume is approaching full capacity.

To Add Capacity to or Rename a Volume

Click a volume name. You can modify the name or increase (but not decrease) its size either by allocating more space to the volume from the existing RAID or by expanding the volume into a new RAID Group.

To Delete a Volume

First select the volume name, then click **Delete Volume** on the screen that displays. Follow the onscreen instructions to delete one or more volumes.

CAUTION: Deleting a volume deletes all data residing on the volume, including any data on any iSCSI Disks it supports. To preserve the data, back it up before deleting the volume. Deleting the volume also deletes any Quotas, Security Settings, Shares, and ACLs associated with the volume being deleted.

To Disable/Enable Write Cache on a New Volume

If the disk drives in the RAID or RAID group hosting the volume support disabling write cache, an **Enable Write Cache** checkbox will be displayed when you are creating a volume (**Storage > Volumes > Create Volume**). The box is checked by default. To disable write cache, uncheck the box. To re-enable write cache, check the box.

Note: If you do not have a configured, online UPS and you re-enable write cache, you will receive a warning that the volume's data may be at risk during a power outage.

To Disable/Enable Write Cache on an Existing Volume

1. From the **Storage > Volumes** page, click the volume you want to modify.
2. If all drives in the RAID or RAID group hosting the volume support disabling write cache, an **Enable Write Cache** checkbox will be displayed. Uncheck the box and click **OK**.

Note: Disabling write cache on an existing volume requires the volume to be unmounted and remounted. Any active users, connections, or iSCSI disks will be disconnected.

Creating a Volume

To create a volume, navigate to the **Storage > Volumes** screen. Creating a volume is a simple process of defining the name, RAID location, and size of the volume. When finished, you can continue to the Create Share screen or exit to the Volumes screen.

Note: When you create a new volume, it is important that you create at least one share to the root of the volume exclusively for backup purposes.

Step 1: Select Volume Parameters

To begin the process, click **Create Volume**. On the screen that opens, define the volume's parameters as described in the following table, and then click **OK**:

Option	Description
RAID Set	If necessary, select the RAID set on which to create the volume. The screen refreshes to display the options and information appropriate to the RAID's remaining capacity.
Volume Name	Accept the default volume name or enter a new one. To rename the volume, use up to 20 alphanumeric characters, including the hyphen, but starting with an alphanumeric character.
Volume Size	The capacity of the volume defaults to the total remaining capacity on the RAID. If you plan to have only one volume on the RAID and do not plan to use snapshots, accept the default value. If you plan to have more than one volume, or if you need to reserve space for the snapshot pool, adjust the value accordingly.

Option	Description
Enable Write Cache	<p>If the disk drives in the RAID or RAID group hosting the volume support disabling write cache, an Enable Write Cache checkbox will be displayed. The box is checked by default. To disable write cache, uncheck the box. To re-enable write cache, check the box.</p> <p>Note: If you do not have a configured, online UPS and you re-enable write cache, you will receive a warning that the volume's data may be at risk during a power outage.</p>
Snapshot Size	<p>The Snapshot Size drop-down menu is a shortcut that allows you to add a percentage of the capacity being allocated to the new volume to the snapshot pool. This feature defaults to 20%, the recommended amount of space to reserve for snapshots.</p> <p>Note: This figure <i>does not</i> indicate how much space is already allocated for snapshot storage. To see the current allocation for snapshots, navigate to the Snapshots screen, click the Snapshot Space button, then click the RAID Set.</p> <p>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 Estimating Snapshot Pool Requirements.</p>

Step 2: Confirm Volume Configuration

Review your settings and click **Create Volume** to create the volume.

Step 3: Done

On the completion screen, do one of the following:

- To create a share for the new volume, click **Create Share**.
- To return to the Volumes screen, click **Close**.

Editing/Expanding an Existing Volume

You can modify the volume name or increase (but not decrease) its size in one of two ways: (1) if unused capacity remains on the RAID, you can enlarge the volume using the Size field or by clicking the **Grow to Max. Size** button; or (2) if all capacity on the existing RAID is already allocated *and* the system has spare capacity either in an existing RAID or unassigned drives, you can enlarge the volume by clicking the **Expand Volume** button. This button launches a wizard that allows administrators to create a RAID Group (see RAID Groups) and automatically expand the volume into the new capacity.

Note: The **Expand Volume** button only appears if there is a sufficient number of unassigned drives of the same capacity to create a new RAID of the same type as the one containing the volume, or if there's another existing RAID of the same type with no volumes.

To Rename a Volume

Enter the new name using up to 20 alphanumeric characters, including hyphens (but not spaces), starting with an alphanumeric character, and then click **OK**.

To Add Capacity by Using Unallocated Space from the Existing RAID

In the Size field, increase the volume capacity up to the total amount of unallocated capacity or click the **Grow to Max. Size** button, and then click **OK**.

To Add Capacity by Creating a New RAID (Expand Volume)

1. On the Volumes - Edit screen, click **Expand Volume**.
2. On the Expand Volume screen, select either an existing RAID, or select an appropriate number of disk drives to create a new RAID.

CAUTION: SnapServers have a hard limit on the size of a volume. This limit is shown in the Max Size field. Comparing the current size of the volume to the maximum size, determine the ceiling on the capacity you can add to the volume. Any capacity that exceeds this ceiling cannot be used by the expanded volume.

3. Click **OK** to review the information on the confirmation screen, and then do one of the following:
 - Click your browser's **Back** button to make changes.
 - Click **Expand Volume** to complete the process.

Quotas

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

Default Quota Assignments

For users and groups, there are no pre-assigned default quotas on the SnapServer. When quotas are enabled on the SnapServer, you can assign a default quota for all users, or allow all users to have unlimited space on the volume.

Unless you assign individual user or group quotas, all users and groups will receive the default quota.

How the SnapServer Calculates Usage

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

Procedure

To Enable or Disable Quotas and Assign a Default Quota for a Volume

1. Navigate to the **Storage > Quotas** screen. This screen displays all volumes on the SnapServer. The enabled column indicates if a volume has been enabled for quotas by displaying **Yes** for enabled, or **No** for not enabled. To change the settings for a volume, click **Yes** or **No**.
2. The Enable Quotas screen opens. To enable quotas, check the checkbox. To disable quotas, clear the checkbox.

3. To set a default quota, enter a number (MB) in the **Limit user to** field, or select to allow the user to use unlimited space.
4. Click **OK**.

If you are enabling quotas that could not be unmounted, you are then prompted to reboot the server. Click **Restart**.

Note: Disabling quotas for a volume deactivates any existing quota assignments you have made but does not delete them. If you re-enable quotas, previous quota settings will be re-established.

Setting User Quotas

You can set individual quotas on a per-volume basis for any local user, Windows domain user, or NIS user or group known to the SnapServer.

Note: Specific individual user quotas always override the default quota.

To Set Default User Quota Limits

When quotas are enabled on a volume, the default user quota is set and applies to all users of the volume. To change the default quota that applies to all users of the volume:

1. Navigate to the **Storage > Quotas > Yes/No** link in the Enabled column. Edit the user default limit by selecting one of the following:
 - **Allow user to consume the entire disk (no limit)** - this allows all users to have an unlimited amount of space to use on the SnapServer.
 - **Limit user to** - this sets a limit to the amount of space users can have on the volume. Enter a whole number. If you do not assign a default quota to users, all users will be able to use unlimited space on the volume.
2. Click **OK** when finished.

To set quotas for individual users, for example a user who needs more space than the default quota allows, you must set a new quota that only applies to that user. See [To Add, Modify, or Remove Individual User or NIS Group Quotas](#) for more information.

To Add, Modify, or Remove Individual User or NIS Group Quotas

To set quotas for individual users or NIS groups, for example a user or group who needs more space than the default quota allows, you must set a new quota that only applies to that individual user or group of users.

1. Click the link for the volume on which you want to assign quotas to view a list of all users known to the SnapServer. If the number of users exceeds a single screen's display limit, the table spans multiple screens.
2. To assign specific quotas to a single user or NIS group, click the user or group.
3. Set the limit for the user or group, select no limit, or select the default quota, then click **OK**. This limit will only apply to the user or group specified.
4. You should see the new limit reflected in the users table. If you do not see the new limit, click **Refresh**.
5. Click **Close**.

Data Migration

Use the Data Migration feature to migrate data from a legacy SnapServer or other computer that supports CIFS or NFS (v2 or v3) to a new SnapServer. The Data Migration (DM) feature can be used to copy or move files and folders from a server on the network (source) to a SnapServer (target).

To access the Data Migration utility, navigate to **Maintenance > Data Migration**.

If an error is encountered during migration (for example, a file or folder is locked and cannot be migrated), the DM utility records the error in a log, and continues the operation. When the migration is completed, the administrator can view the log of migration errors. Once the errors have been corrected, the user returns to the DM main screen, and recreates the migration. With the exception of the password, all fields will still be populated with the specifications of the last job.

The following migration options can be specified:

- Copy or move data
- Include subfolders
- Overwrite existing files
- Preserve the original permissions settings

Note: If you elect to preserve original permissions settings, be sure to review [“Preserving Permissions” on page 4-22](#).

- Verify migrated data

Note: If you elect to verify migrated data, all data will be read twice, once for migration and once for comparison to the copied data. This could be a lengthy process.

Procedures

Setting Up a Data Migration Job

Before setting up a migration job, be sure to specify a user identity for the operation that will have full access to all files on the source, regardless of permissions set:

- For Windows migration, specify an administrator or member of the Windows server/domain administrators group.
- For NFSv2/3 migration, consider using the user root, and configuring the NFS export on the source to `no_root_squash` for the IP Address of the SnapServer for the duration of the migration.

To create a data migration job, perform the following procedure:

Note: Only one migration job can run at a time.

1. On the Data Migration page, complete the required information for source (legacy server) and target (SnapServer).

Option	Description
Network Protocol	<p>Protocol that the SnapServer uses to connect to the source server. Select:</p> <ul style="list-style-type: none"> • Windows (SMB for Windows servers or GuardianOS servers with source data on a Windows SnapTree: default) <p>Note: If you are migrating via SMB, SMB must also be enabled on the target SnapServer (go to Network > Windows to enable SMB in GuardianOS).</p> <ul style="list-style-type: none"> • NFS (NFSv2/3 for UNIX/Linux-based servers or GuardianOS servers with source data on a UNIX SnapTree)
User Identification	<ul style="list-style-type: none"> • If Windows was selected as the protocol, provide the Authentication Name and Password (Windows user name and password to log in to the server over SMB). • If NFS was selected as the protocol, provide the User Name (SnapServer local user name or NIS user, representing the UID used to perform the operation over NFS).
Host	Enter the name or IP address of the source computer you are migrating data from.
Source Share/Export and Path	<p>Specify the share (Windows) or export (NFS) on the source server containing the data you want to migrate and the path to the file or folder you want to migrate. If you are migrating the entire share, you can leave the Path field blank.</p> <p>Note: Wildcards are not supported when specifying the source share or path to migrate.</p>
Target Volume and Path	Specify the volume and path on the target SnapServer where you want the data migrated.
Migration Type	<p>Options are to Copy data (source data is maintained) or Move data (source data is removed after copy, or after verification if Verify migrated data is enabled). The default is Copy.</p> <p>Note: If you select to Move rather than Copy data, it is strongly recommended that you also select to Verify migrated data.</p>

Option	Description
Migration Options	<p>Check to select or uncheck to deselect the following migration options:</p> <ul style="list-style-type: none"> • Include all sub-folders. If the folder you select for migration contains sub-folders, selecting this option will migrate all files and folders underneath this folder (checked by default). If disabled, <i>only</i> the files directly in this folder will be migrated. • Overwrite existing target files & folders. If files/folders on the target share identical names with files/folders on the source, checking this option overwrites those files/folders during migration (checked by default.) • Preserve file/folder permissions. Selecting this option will retain the source permissions when the files/folders are migrated to the SnapServer target (unchecked by default). Note: Before selecting this option, be sure to review Preserving Permissions. • Verify migrated data. Selecting this option will cause all source data to be read twice, once to write to the target SnapServer and once to perform a binary comparison with the data written to the SnapServer (unchecked by default). If enabled, and if the Migration Type is <i>Move</i>, files on the source will only be removed after verification. Otherwise, files will be removed immediately after copying them to the SnapServer. If you select to move files rather than copy them, it is strongly recommended that you enable the Verify migrated data option. If a file mismatch occurs during verification, the target file is moved to a <i>data_migration_verify_failures</i> directory on the root of the same volume. Check the failed file to determine the problem, then run the migration again with Overwrite existing target files & folders deselected (so you don't re-copy files that have already been copied and verified). Note: Depending upon how much data is being migrated, verifying migrated data can be a lengthy process.
Email Notification	<p>Clicking the email notification link will take you to the Email Notification screen (for more information, see Email Notification). Fill in notification information and check the box next to Administrative Operation Event in order to receive an email when the migration operation is complete.</p>

2. Once you have completed the migration information, click the **Start Migration** button to begin the migration. You can see the progress of the migration, an estimated time until completion, and the Migration log on the Data Migration page as it is compiling.
3. When the migration is complete, click the **View Log** button to see details of all errors. Click the **Data migration error log link** to download the entire log.

Stopping a Migration Job

To stop the migration at any time, click the **Stop Migration** button on the Data Migration page. If a file was in the process of being copied, the partially-copied file on the target will be removed.

Recreating a Migration Job

The data migration log records all errors that occurred during migration. You can migrate files and folders that were not migrated during the original job because of an error condition (for example, the file was locked).

1. Review the Data migration errors log and correct all error conditions.
2. Reopen the Data Migration page. All fields (except the password) for the last migration will still be entered on the page.
3. Click **Start Migration** to run the migration again. By default, all files will be re-migrated. If you want only to migrate those files that failed to migrate the first time, you can disable the **Overwrite existing target files** option. However, make sure that any problematic files during the first migration are deleted from the target SnapServer so they will be re-migrated.

Note: If a migration failed, it is strongly recommended that you enable the **Verify migrated data** option for the re-migration.

Preserving Permissions

The types of permissions retained will differ, depending on which of the following migration scenarios is applied:

Migrating from a Windows Security Model to a Windows SnapTree

If you are migrating from a Windows server (or other type of server that follows the Windows security model) to a Windows SnapTree on a SnapServer, permissions will be retained exactly as they exist on the source. However, as is the case when moving files with permissions between Windows servers, permissions for users that are unknown on the target server will be retained but not enforced. This includes permissions for:

- Local users on the source machine.
- Domain users for domains unknown to the SnapServer (for example, trusted domains, if the SnapServer is not configured to support trusted domains).
- Certain built-in Windows users and groups.

Migrating from a UNIX Security Model to a UNIX SnapTree

If you are migrating from a UNIX server to a UNIX SnapTree, UNIX permissions for UIDs/GIDs are copied exactly from source to target; thus, identities of the users and groups will be best retained if the SnapServer belongs to the same NIS domain as the UNIX server.

Migrating Between Conflicting Security Models

When migrating from a Unix source to a Windows SnapTree, Unix permissions will be retained and the security personality on the resulting files and directories will be Unix.

However, when migrating from a Windows source to a Unix SnapTree, permissions cannot be retained (since Unix SnapTrees are required to be Unix personality throughout). Files and directories will inherit the Unix personality and will have a set of default Unix permissions.

Migrating from a GuardianOS Server

When migrating from one GuardianOS server to another, it is recommended that you maintain the same security model on the target server that you have on the source.

- If your source server uses a Windows SnapTree and has permissions assigned to Windows domain users, use a Windows connection for migration. Windows permissions will be retained exactly as they are on the source, with the same enforcement limitations for unknown users as for migration from Windows servers (see Migrating from a Windows Security Model to a Windows SnapTree).

Note: If migrating from a pre-5.0 GuardianOS server, Windows permissions will be retained verbatim, but may have different meaning due to the differences between the pre-5.0 POSIX ACL security model and the Windows security model introduced in 5.0.

- If your source server uses a UNIX SnapTree and has permissions assigned to local or NIS users, use an NFS connection for migration.

Note: Local users that have UNIX permissions on the source will not be created on the target with the same UIDs.

Migrating from a SnapOS Server

When migrating from a SnapOS Server to a GuardianOS server, permissions will not be correctly retained.

Maintenance Tools

These tools provide general-purpose server maintenance for both volume and root filesystems.

To Check the Root Filesystem

CAUTION: Checking the root filesystem requires a reboot of the server.

1. In Maintenance > Tools, click **Check Root Filesystem**.
2. On the page that opens, click the **Check Root Filesystem** button.
3. Click **Yes** when informed that a reboot is required.
4. After the server reboots, to view a log of the results, click the **View Log** button.

To Check the Filesystem on a 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** to be checked.

3. Choose the **type** of check:
 - **Do not repair errors** – Checks for errors, but does not repair them. It is recommended that you do this periodically, especially following a power outage or any other unconventional incident.
 - **Repair errors** – Repairs standard filesystem errors. It is recommended that you run this level if you suspect filesystem damage may have occurred (for example, if a previous **Do not repair errors** operation reported filesystem errors).
 - **Repair errors (aggressive)** – Attempts to repair severe filesystem corruption.

CAUTION: It is only recommended that you run this level if you have been advised to do so by SnapServer Technical Support, or if **Repair errors** has failed to solve the problem and you are willing to risk loss of data.

4. Click **Check Filesystem**.

Checking a filesystem may require a reboot of the server in some circumstances. If prompted that a reboot is required, click **Yes**.
5. To view a log of the results, click the **View Log** button after the filesystem check completes.

Expansion Arrays

Note: This section only applies to SnapServer models that can attach an expansion array. See the *Configuration and Hardware Options Guide* for expansion array options.

Note: If GuardianOS detects an expansion unit that is not integrated with the SnapServer, a warning displays across the top of the Disks/Units screen with a link to information about the orphan expansion unit.

To increase the capacity of a SnapServer, Overland Storage offers the SnapDisk E2000 and the Snap Expansion S50 expansion arrays. Details on installing a SnapDisk E2000 or a Snap Expansion S50 are provided in the quick start guide that comes packaged with the array. The guide is also available for download from:

<http://www.snapserver.com/support>.

Topics in Expansion Arrays

- SnapDisk E2000
- Snap Expansion S50
- Managing Expansion Array Storage
- Integrating Orphan Expansion Units

SnapDisk E2000

The SnapDisk E2000 is a 2U expansion array with up to twelve SATA II or SAS disk drives, or a combination of SAS and SATA disk drives up to a maximum of 12. It ships as a set of unassigned disks with no RAID configuration. Up to five SnapDisk E2000s can be connected to a SnapServer N2000.

Note: Specific configurations are recommended when SAS and SATA drives (or drives with different rotational speeds) are combined in the same expansion array. Be sure to review [“Adding New Disk Drives to Increase Capacity”](#) on page 4-29 before configuring a mixed-drive array.

A SnapDisk E2000 expansion array is accessed and managed through the SnapServer to which it is connected. The expansion array has no physical connection to the network. After the SnapDisk E2000 is installed and powered on (see the *E2000 Quick Start Guide* for details), the array's disk drives appear as unassigned drives, allowing the administrator to configure RAID configurations as necessary.

Snap Expansion S50

The Snap Expansion S50 storage subsystem is a 2U expansion array with up to twelve SAS or SATA II disk drives, or a combination of SAS and SATA disk drives up to a maximum of 12. It ships as a set of unassigned disks with no RAID configuration. Up to seven Snap Expansion S50s can be connected to a SnapServer 520, 550, 620, or 650.

Note: Specific configurations are recommended when SAS and SATA drives (or drives of different rotational speeds) are combined in the same expansion array. Be sure to review [“Adding New Disk Drives to Increase Capacity” on page 4-29](#) before configuring a mixed-drive array.

A Snap Expansion S50 expansion array is accessed and managed through the SnapServer to which it is connected. The expansion array has no physical connection to the network. After the S50 is installed and powered on (see the quick start guide for details), the array's disk drives appear as unassigned drives, allowing the administrator to configure RAID configurations as necessary.

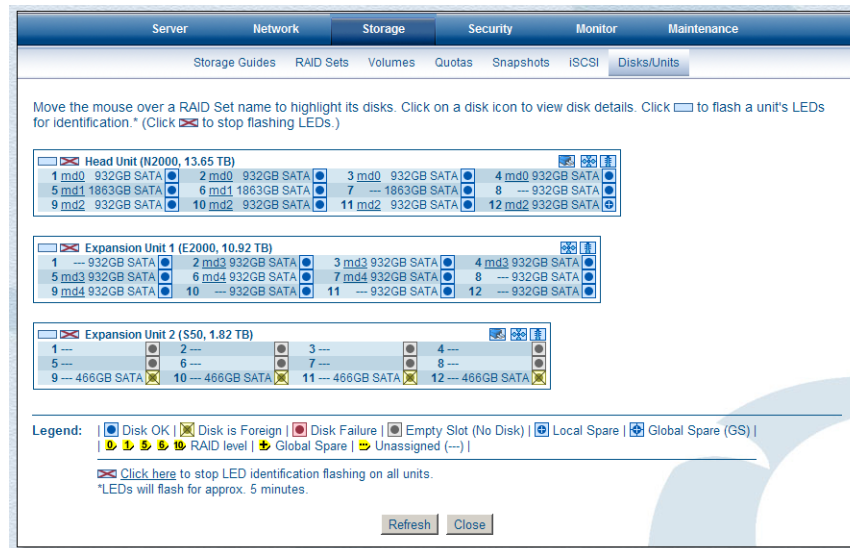
Preparing the SnapServer

Some SnapServers ship with an HBA installed for connectivity to one or more expansion arrays. If your server already has an expansion HBA, no further preparation (other than preparing rack space) is necessary. To connect an expansion array to a SnapServer that does not have an expansion HBA, you will need to purchase and install the HBA, available from an authorized SnapServer reseller.

Managing Expansion Array Storage

Disk drives on expansion arrays are not preconfigured, but are shipped as unassigned disk drives, allowing administrators to configure the array as appropriate.

The **Storage > Disks/Units** screen displays the head unit and any expansion arrays attached to the head unit. For more information about the Disk/Units screen, please see “[Disks and Units](#)” on page 4-27.



The disk drives of an expansion array are completely integrated into the host SnapServer’s logic. The default RAID configurations can be deleted and the internal and external disk drives recombined as necessary. For example, to create one large RAID, you could delete the existing RAIDs on both the host server and the expansion array, then combine all drives into one high-capacity storage system.

This configuration 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 inaccessible. This configuration also increases the potential for multiple drive failures in a single RAID. See “[RAID Groups](#)” on page 4-10 for information on how to avoid this.

CAUTIONS:

- Host server disk drives and expansion array disk drives are logically interchangeable, but they are not physically interchangeable. That is, you cannot physically take a disk drive from an expansion array and place it in a host SnapServer. SnapServer disk drives contain GuardianOS-specific data that is lacking on expansion array disk drives.
- Do not mix drives of different capacity in a RAID 1, 5, 6, or 10. The redundancy schemes in these RAID types limit capacity usage in all member drives to the capacity of the smallest member disk drive. For example, if a RAID consists of one 160 GB disk drive and three 250 GB disk drives, the RAID can use only 160 GB on each disk drive. In this case, the total RAID capacity is approximately 640 GB (4 x 160) rather than the expected 910 GB (160 + [3 x 250]).
- Do not mix drives of different rotational speeds in the same column. See “[Adding New Disk Drives to Increase Capacity](#)” on page 4-29 for illustrations of supported and unsupported drive configurations.

Integrating Orphan Expansion Units

Expansion units that have been discovered by GuardianOS (for example, are physically connected to the SnapServer) but have not been integrated with the SnapServer are listed in the Orphan Expansion Units table:

Property	Description
Expansion Unit	A description of the unit
Status	The status of the unit (for example, orphan)
Serial Number	The expansion unit's serial number
Origin	The serial number of the server with which the expansion unit was last incorporated

If you want to use the expansion unit with the SnapServer, click the checkbox next to the orphan expansion unit you want to integrate, and click **OK**.

CAUTION: Before integrating an orphan expansion unit, be sure that it is compatible with the SnapServer (for example, data on the expansion unit is compatible with the SnapServer configuration, Unicode settings are the same, etc.).

Disks and Units

The Disks/Units screen is a graphic representation of RAID configuration and disk status on your server. The legend explains the meaning of each icon.

- Move the mouse over a RAID set name to highlight all disks within the RAID set.
- Click a RAID set name to view or edit the RAID set.
- Click a disk icon to view disk details.
- Click a unit's LED icon to flash the unit's LEDs for identification.

Note: The LEDs will continue to flash for five minutes. To stop a unit's flashing LED, click that unit's LED icon with a red 'X'. To stop flashing LEDs for all units, click the link at the bottom of the Disks/Units page.

Expansion arrays, if attached to your server, will also be displayed here.

Note: If GuardianOS detects an expansion unit that is not integrated with the SnapServer, a warning displays across the top of the Disks/Units screen with a link to information about the orphan expansion unit. Also, the orphan expansion unit will be highlighted on the screen.

Topics in Disks and Units

- [Replacing Disk Drives on a RAID](#)
- [Adding Disk Drives to a RAID](#)
- [Adding New Disk Drives to Increase Capacity](#)
- [Hot Swapping Disk Drives](#)

Replacing Disk Drives on a RAID

This section describes how to safely remove and replace drives to a degraded RAID. After a fresh drive is inserted into the drive bay, you must use the Web Management Interface to add it to a RAID.

How RAIDs React to Disk Drive Removal

- **RAID 0 (nonredundant)** – Removing a disk drive from a RAID 0 causes the RAID 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)** – Removing a disk drive from a two-drive RAID 1 or a RAID 5, 6, or 10 places the RAID 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, which can tolerate a two-drive failure before it loses redundancy).

Note: If you configure a RAID 1, 5, 6, or 10 with a hot spare, the array automatically starts rebuilding with the hot spare when one of the disk drives fails or is removed.

To Replace a Disk Drive

The following procedure assumes 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.

1 Physically remove the failed disk drive, and insert a new one in its place.

See detailed instructions that come with the disk drive for replacing it.

2 After the drive is removed, navigate to the **Storage > Disks/Units** screen.

The status of the new drive now reads *Unassigned*.

Note: If you have enabled the *automatic incorporation of an unused disk* feature, the drive you insert (a raw drive, a drive with a non-GuardianOS partition, or an unassigned GuardianOS-partitioned drive) will be automatically incorporated into the RAID. Skip Step 2.

3 Add the disk drive logically to a degraded RAID.

Navigate to the **Storage > RAID Sets** screen and click the name of the degraded RAID. On the RAID's Edit screen, click **Repair** to view a list of available drives. Select a drive from the list, and click **Next**. On the confirmation screen, click **Next**. You return to the RAID's Edit screen. The status of the RAID now reads *Resyncing*, and the status of the newly added drive shows as *Hot Spare*.

Adding Disk Drives to a RAID

This section describes how to safely add drives to an existing RAID 1, 5, 6, or 10. On SnapServers, after a fresh drive is inserted into a drive bay, you must use the Web Management Interface to add it to a RAID.

How RAIDs React to Disk Drive Additions

- **RAID 0 (nonredundant)** – You cannot add a drive to a RAID 0. To reconfigure a RAID 0, you must delete the RAID and then recreate it.
- **RAID 1 (redundant)** – You can add a new drive to a RAID 1 as either a hot spare or as a new member. Adding a disk drive to a RAID 1 does not add storage capacity. The new member simply creates an additional copy of the original drive.
- **RAID 5, RAID 6, or RAID 10 (redundant)** – You can add a hot spare to a RAID 5; RAID 6, or RAID 10. However, you cannot add a new drive as a new member.

To Add a Drive to an Existing RAID 1, 5, 6, or 10 Using the Web Management Interface

1. Navigate to the **Storage > RAID Sets** screen and click the name of the RAID (except RAID 0) to which you want to add a drive.
2. On the screen that opens, click **Add Disk**. If you are adding to a RAID 1, select either **Hot Spare** or **Member** at the top of the screen.
3. Select one or more drives to add to the configuration, and then click **Next**.
4. On the confirmation screen, click **Add Disk**.


To Reintegrate Orphaned Disk Drives

An *orphan* disk drive occurs in the following circumstances: (1) a working drive from a RAID is accidentally removed from the server; or (2) the RAID 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 on the RAID in the **Storage > RAID Sets** screen, then click the **Repair** link next to the drive in question.


Adding New Disk Drives to Increase Capacity

For those servers and expansion arrays that ship with fewer than the maximum number of disk drives, additional drives can be added to the server or expansion array to increase capacity. Drives of different rotational speed (for example, SAS and SATA drives) can be combined in the same server. However, they cannot be combined in the same column, and it is recommended that columns of same-type drives be grouped together. If you are combining drives with different rotational speeds, use the figures below to plan where to place the disk drives.

Recommended Disk Drive Configurations



Unsupported Disk Drive Combinations



Do NOT include drives with different RPM rates in the SAME column.

Hot Swapping Disk Drives

The term *hot swap* refers to the ability to remove and add components to a system without the need to turn off the server or interrupt client access to files.

When to Hot Swap Disk Drives

When available storage space is not at a premium, most administrators prefer to configure a RAID with a hot spare that automatically takes the place of a failed drive. This solution assures that client access to filesystems is not interrupted. In environments where configuring a hot spare is not possible, you may need to hot swap a drive.

Hot Swapping Disk Drives

You can hot swap disk drives on SnapServer RAID 1, 5, 6, or 10 by following the two basic steps outlined next:

1 Remove the failed drive from its bay, and insert the new drive.

The procedures for the physical removal and replacement of a disk drive for SnapServers are explained in the following sections.

Note: If you have enabled the *automatic incorporation of an unused disk* feature, the drive you insert (a raw drive, a drive with a non-GuardianOS partition, or an unassigned GuardianOS-partitioned drive) will be automatically incorporated into the RAID. Skip Step 2.

2 Configure the new drive as part of the RAID.

When you remove a drive from a SnapServer, the affected RAID transitions to degraded mode. It remains in degraded mode until the newly inserted drive is configured as a member of the RAID via the Web Management Interface. For details on this procedure, see [“Adding Disk Drives to a RAID” on page 4-28](#).

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-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 initiator(s) to the SnapServer.

Topics in iSCSI Disks:

- [Configuring iSCSI Initiators](#)
- [iSCSI Configuration on the SnapServer](#)
- [Creating iSCSI Disks](#)
- [Configuring VSS/VDS for iSCSI Disks](#)
- [Configuring iSNS](#)

iSCSI Disk Limitations

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

For Additional Information

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

- **SnapServer Online Help** – Available from the **Storage > iSCSI screen**, the online help provides details on creating and managing iSCSI disks on SnapServers.
- **RFC3720: Internet Small Computer System Interface (iSCSI)** – Detailed specification for the iSCSI protocol, available from <http://www.ietf.org>.
- **RFC4171: Internet Storage Name Service (iSNS)** – Detailed specification for the iSNS protocol, available from <http://www.ietf.org>.

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

Configuring iSCSI Initiators

Overland Storage has qualified a number of software initiators, PCI cards, and drivers to interoperate with SnapServers. See the [iSCSI support](#) page on our website for the latest information on supported versions of these software and hardware initiators.

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

- iSCSI Configuration for Microsoft Windows using MS Initiator
- Configuring the QLogic iSCSI Initiators for Microsoft Windows
- iSCSI Configuration for Linux and UNIX
- iSCSI Configuration for Novell NetWare
- iSCSI Configuration for VMware
- iSCSI Configuration for Mac

iSCSI Configuration for Microsoft Windows using MS Initiator

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

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

- On pre-Vista operating systems, Microsoft does not support dynamic disks for use with the Microsoft iSCSI initiator. Overland Storage recommends:
 - Using a QLogic QLA4010 or QLA4050 HBA which supports dynamic disks.
 - Using only “basic” disks with the Microsoft initiator to avoid unexpected behavior and possible data loss when connecting to iSCSI targets in a SnapServer.
- To extend the size of a basic disk on pre-Vista operating systems, use the `diskpart.exe` utility as described in “[Using the Microsoft Diskpart Utility to Grow iSCSI Basic Disks](#)” on page 5-5 or refer to Microsoft KB article [325590](#). The Microsoft knowledge base can be found at <http://support.microsoft.com>. On Vista, Windows 2008, and Windows 7 systems, use the disk management tool to resize the disks.

Configuring Microsoft Services Installed on iSCSI Disks to Start Automatically

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

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

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

Configuring the Server to Persistently Connect

1. Create an iSCSI disk on the SnapServer (see “[Creating iSCSI Disks](#)” on page 5-12).
2. From the Target tab of the Initiator's Property dialog box, select the Target, click the **Logon** button, check the **Automatically restore this connection when the system reboots** box to make this a persistent target, then click **OK** to log in to the SnapServer target.
3. Use the Disk Administrator to configure all volumes on top of the disks.
4. From the Bound Volumes/Devices tab on the Property dialog box, click **Bind All** to allow the iSCSI service to configure the list of persistent volumes. If you are running Windows XP, Windows 2003 Server, Vista, Windows 7, or Windows

2008 Server, your iSCSI disks will now start automatically on reboot. If you are running Windows 2000 Server, you must continue to the following procedure and edit the registry to make services dependent on the iSCSI Initiator service.

Editing the Windows Registry for MS Exchange Server or MS SQL Server (Windows 2000 only)

1. Install Exchange Server 2003 and configure it to use the iSCSI disk as the location to store database files.
2. On a Windows workstation running Windows 2000, enter the following on the command line:
`regedt32`
3. Navigate to the Key:
 - For Exchange Server:
`HKey_Local_Machine > System > Current Control Set > Services > lanmanserver`
 - For SQL Server:
`HKey_Local_Machine > System > Current Control Set > Services > MSSQLServer`
4. If the value `DependOnService` already exists, double-click it. If it does not, create it:
 - a. Select **Add Value** from the Edit menu to open the Add Value dialog box.
 - b. In the **Name** field, enter **DependOnService** and click **OK**.
5. In the Data box that opens, enter **msiscsi**, click **OK**, and then close the registry.
6. Reboot the Windows server.

Configuring Shares to iSCSI Disks

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

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

Mounting an iSCSI Disk Without a Drive Letter

To complete this procedure, you must create and format an iSCSI target on the SnapServer and connect to this iSCSI disk using the Microsoft initiator. You must also have an existing NTFS volume on a local disk within the Windows server, initiating the connection.

1. Right-click My Computer and select **Manage**.
2. The new formatted volume will appear in the Disk Management window.
3. Right-click the **New Volume** and select **Change Drive Letter and Paths**.

4. Click **Remove** in the Change Drive Letter and Paths for (New Volume) dialog, and click **Yes** to confirm drive letter removal.
5. Right-click the **New Volume** again and select **Change Drive Letter and Paths**.
6. Select **Add** in the Change Drive Letter and Paths for (New Volume) dialog.
7. In the Add Drive Letter or Path dialog, select **Mount in the following empty NTFS folder**.
8. Create a folder or enter the path to the one that will be shared from the Windows server and select **OK**.
9. Select **OK** in the Add Drive Letter or Path dialog. This returns you to the Disk Management window.

You will see the icon of a disk in place of the folder icon in the File Management window.
10. Create a share to the iSCSI disk in the standard method, then reboot the Windows machine and verify that the share is persistent.

Configuring Dynamic Disks to Persistently Reconnect

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

Using the Microsoft Diskpart Utility to Grow iSCSI Basic Disks

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

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

Preparing to Expand a Microsoft Basic iSCSI Disk

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

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

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

3. Log off the iSCSI volume that is to be expanded.
 - a. Open the Microsoft initiator tool.

- b. Under Connected Targets, highlight the specific iSCSI disk(s) you want to expand.
 - c. Click **LogOff**. This will log you off the specific target.
4. Verify that you have additional space available on the SnapServer to expand an existing volume
 - a. Open the browser-based Web Management Interface for the SnapServer from a client on the network.
 - b. Navigate to **Storage > iSCSI**.
 - c. Select the iSCSI disk you intend to expand.
If you have not disconnected from the iSCSI disk at the host, you will be unable to proceed to the configuration page.
 - d. From the configuration screen, ensure that you have additional space on the volume to expand the selected iSCSI disk.
 - e. Make changes to the iSCSI disk size as desired.
 - f. Click **OK**. The disk should now reflect the larger size.

Expanding the Basic Disk on the Microsoft Host

1. Open the Microsoft initiator tool.
2. Under Available Targets, highlight the specific iSCSI disk(s) you expanded in the previous procedure.
3. Click **LogOn**. This will connect the initiator to the selected iSCSI target.
4. Close the Microsoft initiator tool.
5. Open the Disk Management tool by right-clicking My Computer and selecting **Manage**. In the Computer Management GUI, select **Disk Management**.

The disk will automatically reattach, and the additional expanded space in the iSCSI disk will appear as unallocated space on the same disk.

Expanding an iSCSI Volume using the Microsoft Diskpart Utility

1. In the Start menu, select **Run** and enter **cmd** in the Run dialog to open a command-line window.
2. Enter the command:
diskpart
3. To show all the available disks on the host, enter:
list disk
4. Identify the specific disk you are expanding.
5. To show all the available volumes on the host, enter:
list volume
6. Identify the specific volume you are expanding.
7. Enter the necessary data:
 - a. **select disk *n***
where *n* is the disk number that Diskpart indicated from the list command.

b. select Volume *n*

where *n* is the volume number that Diskpart indicated from the list command.

c. extend size=*n*

where *n* is the number of megabytes you want to expand the disk.

For example, if you are adding 10 GBs to an existing disk of 100 GBs, use the following command:

d. extend size=10240 (the number is in megabytes, 1024MBs = 1GB)

The Disk Management GUI will show the newly expanded disk size.

8. Exit the Computer Management tool.
9. Restart the necessary application services.

Configuring the QLogic iSCSI Initiators for Microsoft Windows

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

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

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

iSCSI Configuration for Linux and UNIX

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

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

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

Using CHAP Authentication to Enable Multiple Linux Systems to Share iSCSI Disks Securely on a SnapServer

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

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

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

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

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

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

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

iSCSI Configuration for Novell NetWare

Consider the following information before implementing iSCSI on NetWare servers:

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

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

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

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

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

iSCSI Configuration for VMware

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

On connecting to the SnapServer targets, the VMware ESX 3.5 Server initiator will find all iSCSI disks and automatically log into them. If iSCSI disks are shared across multiple servers, you can use CHAP authentication to restrict the number of iSCSI disks the VMware initiator can access. See "Creating iSCSI Disks" on page 12 for more information. The VMware vSphere 4.0 Server initiator provides the option for Static Discovery, allowing you to enter the IP addresses of only those targets you want the VMware initiator to access.

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

Using the VI Client to Configure iSCSI Services

Follow the instructions in the *VMware Server Configuration Guide*, available from <http://www.vmware.com>

to configure your iSCSI service. Use the VI Client to:

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

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

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

iSCSI Configuration for Mac

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

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

iSCSI Configuration on the SnapServer

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

Basic Components of an iSCSI Network

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

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

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

Isolate iSCSI Disks from Other Resources for Backup Purposes

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

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

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

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

Note: While some third-party, agent-based backup packages could *technically* back up an iSCSI disk on the SnapServer, the result would be inconsistent or corrupted backup data if any clients are connected during the operation. Only the client can maintain the file system 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 Containing an iSCSI Disk** – Running a GuardianOS snapshot on a volume 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 5-15](#) for more information).

iSCSI Multi-Initiator Support

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

Note: GuardianOS v5.1 and later support Windows 2003 and Windows 2008 Server failover clustering.

The warning message *Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable Multi-Initiator Support if your environment or application supports it* occurs when the checkbox for Support Multi-Initiator is selected. 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, see [“Configuring Write Cache” on page 4-13](#).

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.

Notes:

- 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 the *Volume is Full* Message

When an iSCSI disk is created, the volume allocates the specified capacity to the disk. If all volume 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
```

Creating 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 5-10](#) before you begin creating iSCSI Disks.

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

Click **VSS/VDS Access** to add VSS/VDS clients to the SnapServer. See [“Configuring VSS/VDS for iSCSI Disks” on page 5-15](#) for more information.

To Create a New iSCSI Disk

Click **Create iSCSI Disk**. The process involves first defining iSCSI parameters, setting up security, and then confirming your settings.

To use CHAP authentication

1. Click to put a check in the **Enable CHAP Logon** box.
2. Enter a user name and target secret (password). 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.

GuardianOS Support for CHAP Security (Target Only)

CHAP is a network login protocol that uses a challenge-response mechanism to control iSCSI initiator access to an iSCSI target. GuardianOS supports target authentication, in which the initiator must provide the same CHAP user name and password (or “target secret”) that was configured on the target SnapServer iSCSI disk. Other forms of CHAP authentication are not currently supported.

To View iSCSI Disk Status Information

You can view iSCSI disk status information from the **Storage > iSCSI** screen.

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

To Configure iSNS

Go to the **Network > iSNS** screen, from which you can configure **iSNS**.

To Edit an iSCSI Disk

Click an iSCSI disk name. You can increase (but not decrease) its size and enable or disable CHAP logon.

Note: You cannot edit an iSCSI disk if an initiator is connected. The hostname and IQN name of all connected initiators will be displayed.

To Delete an iSCSI Disk

The system will not allow the deletion of an iSCSI disk when clients are connected (the hostname and IQN name of all connected initiators will be displayed). After disconnecting all client initiators, click the iSCSI Disk **name link** on the **Storage >**

iSCSI page to view the iSCSI Disk Properties page. From the properties page, click **Delete iSCSI Disk** (which is followed by a confirmation screen) to delete the iSCSI disk.

Creating an iSCSI Disk

To create an iSCSI disk, define the name and capacity of the disk.

Step 1: Select iSCSI Parameters

Edit the following options as defined below and click **Next**:

Option	Description
iSCSI Disk Name	Accept the default name or enter a new one. To rename the iSCSI Disk, use up to 20 alphanumeric characters.
Volume	If your configuration includes multiple volumes, select a volume to host the iSCSI Disk. The screen refreshes, displaying the capacity of the selected volume and restoring all fields to default values.
Size	The size of the iSCSI disk defaults to the total remaining capacity on the selected volume. 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, and do not set the size of the iSCSI disk to consume the entire volume. The required Snap volume space for VSS snapshots is 10% of the size of the iSCSI disk per snapshot.
Support Multi-Initiators	Check this box if you want your iSCSI Disk to allow multiple initiator connections. Data corruption is possible if this option is checked. See "iSCSI Multi-Initiator Support" on page 5-11 for more information.
Enable Write Cache	Selected by default, the Write Cache option significantly enhances performance; but 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. Note: Write cache can be disabled on an iSCSI Disk-by-iSCSI Disk basis. Disabling write cache for an iSCSI Disk does <i>not</i> disable 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
Enable CHAP Logon	Select Enable CHAP Logon to enable CHAP authentication for access to the iSCSI Disk, and then enter a user name and target secret. Currently, GuardianOS supports target CHAP authentication only.

Step 2: Confirm iSCSI Configuration

Review your settings and click **Create iSCSI Disk**. On the confirmation screen click **Create iSCSI Disk** to create the iSCSI disk and be returned to the **Storage > iSCSI** screen.

Editing an iSCSI Disk

Click an iSCSI disk name. On the screen that opens, you can increase (but not decrease) its size, and enable or disable CHAP logon. Edit the following fields as appropriate and click **OK**.

Field	Description
Size	<p>Enter the desired capacity (size) of the iSCSI disk.</p> <p>Note: The available space remaining on the volume is listed next to this field. If no available space remains, increase the size of the volume if possible and try again.</p> <p>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, and do not set the size of the iSCSI disk to consume the entire volume. The required Snap volume space for VSS snapshots is 10% of the size of the iSCSI disk per snapshot.</p>
Support Multi-Initiators	<p>Check this box if you want your iSCSI Disk to allow multiple initiator connections. Data corruption is possible if this option is checked. See "iSCSI Multi-Initiator Support" on page 5-11 for more information.</p>
Enable CHAP Logon	<p>Select Enable CHAP Logon to enable CHAP authentication for access to the iSCSI Disk, and then enter a user name and target secret. Currently, GuardianOS supports target CHAP authentication only.</p>

CAUTION: The consistency of the internal file system 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.

Configuring VSS/VDS for iSCSI Disks

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

Note: VSS/VDS operations are supported on iSCSI disks created using GuardianOS 5.2 and later.

- 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 the iSCSI disk resides on. They are not related to GuardianOS snapshots as described in [Snapshots](#).

Note: VSS iSCSI snapshot rollback 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

For more information on using VSS and VDS, see the Online Help.

Procedures

Backing up an iSCSI Disk using VSS Snapshots

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

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

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

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

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

1 Add the VSS client to the SnapServer.

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

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

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

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

- a Depending on the Windows client, locate *SnapServerToolInstall32.exe* or *SnapServerToolInstall64.exe* on the Overland website:

<http://support.overlandstorage.com/support/snapserver-nas.htm>

- b Double-click the executable 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 for the iSCSI disk on the SnapServer using the Web Management Interface.

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

2 Add VDS clients to the SnapServer.

- a From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
- b Click **Add**.
- c Add the hostname of the VDS client you wish to grant access and click **Add** (the hostname is not case sensitive). The client hostname should appear in the VSS/VDS Clients 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 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://support.overlandstorage.com/support/snapserver-nas.htm>

- 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.

3. Create and configure the iSCSI disk using *Storage Manager for SANs* (or other VDS compliant application).

Note: RAID terminology differs somewhat between GuardianOS and *Storage Manager for SANs*. Following are the equivalents:

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

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

Deleting VSS/VDS Client Access

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

Configuring iSNS

Microsoft iSNS Server can be used for the discovery of targets on an iSCSI network. The iSNS software package installs a *readme* file that contains extensive release notes on bug fixes and current iSNS limitations. Be sure to read these notes before attempting to use the service.

1. Install the iSNS service on a Windows server.

Follow the instructions provided in the iSNS *readme* file. Note the IP address of the server or workstation on which the iSNS service is installed.

2. Configure iSNS on the SnapServer

On the **Network > iSNS** screen, check to select the **Enable iSNS** box, enter the IP address of the iSNS workstation, and then click **OK**. The iSNS port default value of 3205 can be changed on this screen as well (if changing the port is supported).

3. Configure iSNS in the iSCSI initiator.

Run the initiator software and configure the iSNS service from the iSNS Servers tab.

For example, from a Windows client:

- When using the Microsoft initiator, run the Microsoft initiator software, select the iSNS Servers tab, and click **Add**. Enter the name or address of the iSNS server, and then click **OK**.
- When using the QLogic4010/4050 initiator, right-click the QLogic adapter and select **Properties**. Select the Discovery Configuration tab, and check **Perform Discovery**. Check **Use iSNS Server**, enter the server name or IP Address, and click **OK**.

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

SnapServer has implemented features to accommodate the disparate methods used by the SMB and NFS protocols for sharing data. At the share level, administrators can assign read-write or read-only share access to individual Windows (and local) users and groups. Administrators can also edit the NFS *exports* file to control how shares are exported to NFS client machines.

The SMB and NFS protocols also handle file-level permissions differently. Administrators can choose to apply Windows- or UNIX-style file-level permissions to entire volumes or to directories at the root of a volume (such as SnapTree directories). These security-based directory structures are referred to as “SnapTrees.”

Files and directories in a Windows SnapTree can have either a Windows or UNIX security personality, depending on the network protocol used to create the file or change permissions on it. Files in a UNIX Snap Tree always have the UNIX security personality and can only be set by NFS clients.

Topics in Share and File Access:

- [Configuring Share and Folder Security Overview](#)
- [Components and Options](#)
- [SnapTrees and Security Models](#)
- [ID Mapping](#)
- [Shares](#)
- [Configuring Share Access](#)
- [Creating Home Directories](#)
- [Windows ACLs](#)
- [Security Guides](#)


Configuring Share and Folder Security Overview

SnapServers support file access in Windows, UNIX, and Apple networks, as well as access via FTP and HTTP. Although GuardianOS runs on an optimized Linux kernel and has many Linux characteristics, the cross-platform features make it very different than a pure Linux distribution. Systems running GuardianOS are storage appliances dedicated to file services. Administrators should not expect the same behavior as a pure Linux system when administering a SnapServer.

By default, volumes are created with the Windows/Mixed security model (Windows-style ACLs for files created by SMB clients and UNIX-style permissions for files created by other protocols and processes), and allow all users to create, delete, and

configure permissions on their own files and to access files and directories created by other users.

New shares are created by default with full read-write access to all users, subject to the file system permissions on the share target directory. The first step to securing a SnapServer is to specify access at the individual share level. Administrators can assign Read/Write or Read-Only share access to individual Windows (and local) users and groups.

Security permissions that have been applied to files and folders can be viewed from the *Web View* page of the Web Management Interface. For users with admin rights, a key icon  appears next to each file and folder in the share. Clicking this icon displays a popup box with security information about the file or folder.

Hidden Shares

There are three ways a share can be hidden in GuardianOS:

- Name the share with a dollar-sign (\$) at the end. This is the traditional Windows method of hiding shares; however, it does not truly hide the share since Windows clients themselves filter the shares from share lists. Other protocols can still see dollar-sign shares.
- Hide the share from all protocols (except NFS) by navigating to **Security > Shares > Create Share > Advanced Share Properties** and selecting the **Hide this Share** checkbox, or by selecting a share, clicking to expand **Advanced Share Properties**, and selecting the **Hide this Share** checkbox. When a share is hidden this way, the share is invisible to clients, and must be explicitly specified to gain access.

Note: Hidden shares are not hidden from NFS, which cannot access invisible shares. To hide shares from NFS, consider disabling NFS access to the hidden shares.

- Disable individual protocol access to certain shares by navigating to **Security > Shares > Create Share > Advanced Share Properties** and enabling/disabling specific protocols, or by selecting a share, clicking to expand **Advanced Share Properties**, and enabling/disabling specific protocols.

File and Directory Permissions

GuardianOS supports two “personalities” of file system security on files and directories:

- UNIX: Traditional UNIX permissions (rwx) for owner, group owner, and other.
- Windows ACLs: Windows NTFS-style file system 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).

The security personality of a file or directory is dependent on the security model of the SnapTree or Volume in which the file or directory exists (see [“SnapTrees and Security Models” on page 6-4](#)).

Share Level Permissions

Share-level permissions on GuardianOS are applied cumulatively. For example, if the user “j_doe” has Read-Only share access and belongs to the group “sales”, which has

Read/Write share access, the result is that the user “j_doe” will have Read/Write share access.

Note: Share-level permissions only apply to non-NFS protocols. NFS access is configured independently by navigating to the **Security > Shares** page, selecting from the table the NFS Access level for the share, and modifying the client access as desired.

Where to Place Shares

For security and backup purposes, it is recommended that administrators restrict access to shares at the root of a volume to administrators only. After initialization, all SnapServers have a default share named *SHARE1* that points to the root of the default volume *vol0*. The share to the root of the volume should only be used by administrators as a “door” into the rest of the directory structure so that, in the event that permissions on a child directory are inadvertently altered to disallow administrative access, access from the root share is not affected. This also allows one root share to be targeted when performing backups of the server. If it is necessary to have the root of the volume accessible, using the Hidden option helps ensure only those that need access to that share can access it.

SnapTrees

SnapTrees are directories that can be configured for the Windows/Mixed or UNIX security model. SnapTrees make a specific directory structure follow the rules of the specified security model, which indicates which file permission personality will be present on files by default, and whether that personality can be changed by users when changing permissions. All top level volume directories, as well as all directories inside the first level of a volume, are considered SnapTrees. For more information, see [“SnapTrees and Security Models” on page 6-4](#).

NFS Share Access

When controlling share access for NFS clients, administrators can limit client access to the shares independently of share level permissions that apply to other protocols. Access is controlled on a per-share basis. To set the NFS access, navigate to **Storage > Shares**. In the Shares table, click in the **NFS Access** column of the share you want to modify. Changes made on this screen 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.

Components and Options

Shares are created and share access is granted using the Web Management Interface. File-level permissions are configured from a Windows or UNIX/Linux workstation.

The following table summarizes the components, options, and tools available for setting up share and file security on SnapServers.

Component	Options
Security Models (SnapTrees)	Volumes and directories created in the root of a volume have one of two security models: Windows/Mixed or UNIX. The security model determines the rules regarding which security personality will be present on files and directories created by the various protocols and clients, and whether the personality of files and directories can be changed by changing permissions. These directories are referred to as SnapTrees, and their security models can be configured from the Security > SnapTrees screen.
Shares	<p>Shares are created on the Security > Shares screen. When creating a share, you must set the following options:</p> <ul style="list-style-type: none"> • Name Select a name for the new share. • Volume Select a volume from the drop-down list. • Path Browse to the directory you want to use as the root of the share or type in the path to the share. If the path does not exist, when you click Browse or OK, you will be asked if you want to create it. • Security Model If you create a share pointing to a volume or a SnapTree directory, a security model may be selected. • Share Access User access to the share can be restricted or full read/write access. <p>By clicking to expand Advanced Share Properties, you can set the following options:</p> <ul style="list-style-type: none"> • Hidden Option The Hidden option allows you to hide a share from clients connecting from SMB, HTTP/HTTPS, AFP, and FTP (but not NFS) protocols. • Protocol Access Client access to the share can be restricted to specific protocols. As a security precaution, disable any protocols not needed by users of the share. • Snapshot Share The snapshot share allows access (using identical security) to snapshots of the data that the new share references.
Share Access	Share-level access allows users/groups/clients to connect to a share and is configured from the Security > Share Access screen. Users and groups known to the system can be given Full Access or Read Only (R) access to the share.
Share NFS Access	The Web Management Interface provides a window into the <i>exports</i> file for defining how a share is exported to NFS clients.
File Permissions	File-level permissions define what actions users and groups can perform on files and directories, and are set from a Windows client for a Windows SnapTree; and from a UNIX/Linux client for a UNIX SnapTree.

SnapTrees and Security Models

Volumes and directories created on the root of a volume are assigned one of two security models: Windows/Mixed or UNIX. The security model determines the rules regarding which security personality will be present on files and directories created by the various protocols and clients, and whether the personality of files and directories can be changed by changing permissions. These directories are referred to as SnapTrees.

- **Creating a SnapTree Directory** – SnapTree directories are created either from the **Security > SnapTrees** screen in the Web Management Interface or from a client from any of the network protocols. SnapTrees created either by clients or in the Web Management Interface will default to the security model of the parent volume.

Note: The security model of a SnapTree directory may differ from the personality of the directory (a Windows/Mixed SnapTree may have the UNIX personality, and vice-versa).

- **Toggling Security Models** – The security model applied to a volume or SnapTree directory can be changed from the **Security > SnapTrees** screen, or when creating a share pointing to a volume root or SnapTree directory. When changing security models, the corresponding personality (such as, Windows for Windows/Mixed and UNIX for UNIX) is applied to the SnapTree directory itself with a default permission, and can optionally be propagated with a default permission to all files and directories inside the SnapTree.

SnapTree Functionality

The following table describes the behavior of SnapTrees and Security Models.

Function	Description
SnapTree Directory Ownership	<p>Default ownership differs according to the method used to create the SnapTree directory:</p> <ul style="list-style-type: none"> • 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 Personality of Files and Directories	<p>Files and directories created by clients inside SnapTrees will acquire security personality and permissions according to the rules of the SnapTree security model.</p> <p>Windows/Mixed SnapTree</p> <ul style="list-style-type: none"> • 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). <p>UNIX SnapTree</p> <ul style="list-style-type: none"> • 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 SnapTree. All files and directories always have the UNIX personality.
SnapTree File System Permissions	<p>Security model and permissions differ according to the method used to create the SnapTree directory:</p> <ul style="list-style-type: none"> • 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 <i>Full Access</i> 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 will be 777 (rwxrwxrwx). If created in a Windows/Mixed volume, permissions will allow all users to create, delete, and change permissions on files created inside the SnapTree, and will grant full control to administrators.

Function	Description
 toggling Security Models	<p>Changes to a SnapTree's security model can optionally be propagated to the corresponding personality with a default permission to all files and directories underneath the SnapTree.</p> <p>When changing the security model on a SnapTree:</p> <ul style="list-style-type: none"> • If changing from Windows to UNIX, all files and directories will be changed to be owned by <i>admin</i> and <i>admingrp</i>, 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.
 Mixing SnapTrees	You can create SnapTrees of different security models on the same volume.

Procedures

To Change a Volume's Security Model

The security model assigned to a volume determines the default security model applied to new SnapTrees directories created in the first level of the volume.

- You can change the security model for a volume such that existing SnapTree subdirectories and normal files and directories retain their current security models and personalities, but new SnapTrees created on the volume default to the new mode.
- Or you can change the security model for the volume and recursively update all existing SnapTrees to the new model at the same time, as well as update all other subfiles and subdirectories to the security personality that corresponds to the selected security model (along with a default permission).

In all cases, changing the security model of a volume changes the security personality of the volume root itself to match the corresponding security model, with a default permission.

1. Click the volume's security model.
2. Select the new security model from the drop-down list.
3. On the page that appears, do one of the following:
 - To change the model only for the volume, leave the **Apply security model to all SnapTrees on this volume** checkbox cleared and click **OK**.
 - To change the security model for the volume and for all SnapTrees on the volume as well, select the **Apply security model to all files and folders on this volume** checkbox and then click **OK**.
4. On the confirmation page, click **Apply Security Model**.
5. Click **Close** to return to the SnapTree page.

To Create a New SnapTree Directory

1. Select a volume on which to create the new SnapTree.
2. Click **Create SnapTree**.

3. Enter a name for the root-level directory, select a security model, and then click **OK**.
4. The confirmation screen appears. Click **Close** to return to the SnapTree Page.

To Change a SnapTree Directory's Security Model

Changing a SnapTree's security model changes its personality to the corresponding security personality, and can optionally change the personality on the files and folders inside the SnapTree to the corresponding security personality, with a default permission.

1. Click the security model for the volume you want to change.
2. Select the new security model from the drop-down menu and click **OK**.
3. On the confirmation page, click **Apply Security Model**.
4. Click **Close** to return to the SnapTree page.

ID Mapping

ID mapping allows users and groups that exist on Windows domains to share user IDs with local or NIS users and groups. This results in the same permissions and quota consumption applying to both the Windows domain user and the local or NIS user.

Example:

John Smith is a local user on a SnapServer, as well as having a user ID on a Windows domain. John's quota for the SnapServer has been set to 200 MB. The administrator of the SnapServer maps the Windows domain user identification for John Smith to the local identification for John Smith, giving both IDs access to John's 200 MB.

Procedures

Configure ID Mapping

1. Map a new ID:
 - a. Select a Windows user or group from the list and click **Map**.

If the desired user does not appear in the list, use the search field to locate the user. Select a domain to search and enter search criteria (for example, name or domain/name), then click the **binoculars** icon to see a subset of all users and groups in that domain. Or click the **world** icon to display all users and groups in that domain.

Note: Search filters without wildcards will search for all entries containing the string you enter in the search field rather than looking for exact matches. For example, if you enter 'abc' as your search criterion, all users and groups containing 'abc' in the name will be identified.
 - b. Select the local or NIS user you want to map the Windows user to, and click **Map User**.
 - c. The domain user now appears in the list as a mapping to another local or NIS user and UID.
2. Remove a mapping:
 - a. Select the user you wish to unmap and click **Remove Mapping**.

- b. The confirmation screen appears. Click **Remove Mapping** to remove the mapping (click **No** to cancel the action).
3. Use the Auto-map feature to generate a list of ID mappings that have the same name as your Local or NIS users and groups:
 - a. Click **Auto Map**.
 - b. Domain, local, and NIS user lists are compared. The matches are automatically queued.
4. Click **View Auto Mappings** to continue.

A page is displayed summarizing your changes.
5. Click **OK** to confirm.

A page is displayed providing an option to apply these mapping changes to existing files and folders on the file systems.
6. Click **Update Filesystem** to start the propagation process.

Clicking **Do Not Update Filesystem** will save your mapping changes but will not apply them to the file systems.

Note: The propagation process may take a long time, depending upon the number of files and folders you have on your server.

Remove all Mappings

1. The **Remove All Mappings** button allows you to remove all ID mappings on the SnapServer. Click this only if you want to remove all ID mappings.
2. A confirmation page appears. Click **Remove All Mappings**.
3. A page is displayed providing an option to remove mappings from all existing files and folders on the file systems. Click **Update Filesystem** to start the propagation process. Clicking **Do Not Update Filesystem** will remove your mappings but will not propagate to the file systems.

Note: The propagation process may take a long time, depending upon the number of files and folders you have on your server.

Note that all ownership and permissions for a given mapped user are converted to the Windows Domain user—they will not necessarily revert to their previous state prior to the original ID mapping.

Shares

Shares are created, viewed, edited, and deleted from the **Security > Shares** screen of the Web Management Interface. The shares table lists all of the shares on the SnapServer, and describes the share properties. Guidelines for creating shares are provided below. Be sure to review them before configuring shares on the SnapServer.

Property	Description
Share	Name of each share
Volume	The volume the share points to
Path	The directory path on the volume

Property	Description
Access	The user-level access defined for that share: <ul style="list-style-type: none"> • Full—if AllUsers has full access • Restricted—If AllUsers does not have full access, or any other user has Read Only access
NFS Access	The NFS access defined for that share: <ul style="list-style-type: none"> • Default—if all hosts have read-write access • User—If not all hosts have read-write access
Protocols	The network protocols enabled for the share (SMB, NFS, AFP, HTTP/HTTPS, FTP/FTPS) <p>Note: As a security measure, disable any protocols not required for your network environment.</p>
Attributes	Attributes for the share: <ul style="list-style-type: none"> • S - has snapshot share • H - hidden share • W - webroot share

The default share (SHARE1) maps to the root of the volume and grants access to all users and groups over all protocols.

Guidelines

Consider the following guidelines when creating or deleting shares.

Maintain at Least One Share at the Root of Each Volume

A share to the root of a volume is recommended for backup purposes. Security for any share at the root of the volume should be given special consideration. Any user or group that has access to the root of a volume will have access to EVERY file and subdirectory on that volume unless there is a specific ACL in place precluding that access. In general, access to a share at the root of a volume should only be granted to a system administrator or backup operator.

Hidden Shares

A *hidden* share is hidden from clients connecting from the SMB, HTTP, AFP, and FTP (but not NFS) protocols. For example, assume SHARE1 is set as hidden. Windows users will not see the share when viewing the server through Network Neighborhood, or when performing a `net view \\servername` on the SnapServer.

For more information, see [“Configuring Share and Folder Security Overview” on page 6-1](#).

Snapshot Shares

A *snapshot share* provides access to all current snapshots of a volume. Just as a share provides access to a portion of a live volume, a snapshot share provides access to the same portion of the file system on any archived snapshots of the volume. You create a snapshot share by selecting the *Create Snapshot Share* checkbox in the course of creating or editing a share.

Security Models, SnapTrees, and Shares

In the course of creating a share that points to a volume or to a directory on the root of the volume (aka SnapTree directory), you must assign a security model to the volume or SnapTree directory. Thereafter, security models for these entities are managed on the **Security > SnapTrees** screens.

NIS Users

When a SnapServer is connected to a UNIX domain, NIS users do not appear in the list of users under **Security > Shares > Access**. NIS user properties cannot be modified from the SnapServer. However, it is possible to assign quotas to NIS users and groups from the **Storage > Quotas** page in the Web Management Interface.

Procedures

To Create a New Share

As a security measure, disable any protocols not relevant to your network environment.

Click **Create Share**. The process involves selecting the share name, volume, directory path for the share, security model, and network access protocols.

To Edit Share Properties

Select a share name. On the page that opens, you can modify its name, description, network access protocols, and snapshot share settings.

An 'S' in parentheses (S) indicates an active snapshot share, an 'H' in parentheses (H) indicates the share is hidden.

To Set Up User-based Share Security

Click the link in the **Access** column next to the share you want to configure. The Share Access screen displays. You can set access levels for the share, as well as grant or deny access to specific users and groups.

Note: To add a new user to a share, you must first create the user, then add that user to the share. Please see [“Local Users and Groups” on page 3-2](#) for information on creating new users.

To Set Up NFS Share Security

Click the link in the **NFS Access** column next to the share you want to configure. The NFS Share Access screen displays. You can configure NFS access to the share using standard Linux “exports” file syntax.

To Delete a Share

Select a share name. On the page that opens, click **Delete Share** and follow the onscreen instructions to delete the share. You can only delete one share at a time. Deleting a share does not delete any directories or files.

Creating a Share

Creating a share involves selecting the volume, security model, and directory path for the share and then defining share attributes and network access protocols.

Step 1: Name the Share and Select a Volume

Name the share in the field provided. Accept the default share name or enter a new one. To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including spaces). Choose the volume you need from the drop-down menu.

Step 2: Select a Path

Select from the following options:

- **To create a share to the entire volume** – The current Path field defaults to the root path of the volume. Simply leave it blank if this is the desired configuration.
- **To create a share to a folder on the volume** – Click **Browse** to display folders on the selected volume. Browse to the folder you wish to point the share to, click on the folder, then click **OK**. Or type a path into the field. If the path you enter does not exist, when you click **Browse** or **OK**, you will be asked if you wish to create it.

Note: If you wish to create a new folder inside any other folder, type the folder name into *New Folder Name* and click **Create Folder**.

Step 3: Enter a Description

If desired, enter a description of the share. This is an opportunity to clarify the purpose of the share.

Step 4: Set Security Model

To choose a security model, select either Windows or UNIX from the drop-down menu. 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 [“SnapTrees and Security Models” on page 6-4](#).

Step 5: Set User-based Share Access

Choose either **Create share with full read and write access for all users**, or **Create share with Admin-only access and proceed to Share Access page** to configure the share access. For more information, see [“Configuring Share Access” on page 6-13](#).

Note: If selecting **Create share with Admin-only access** and if the share has NFS enabled, be sure to configure the NFS Access settings afterward.

Step 6: Configure Advanced Settings

To further configure the share, click **Advanced Share Properties**. Enter the following, and then click **Create Share**:

Option	Description
Hide this Share	Select this option if you want the share to be hidden from network browsing.

Option	Description
Protocols	Select the access protocols for the share: Windows (SMB), Linux/Unix (NFS), Apple (AFP), Web View (HTTP/HTTPS), and FTP/FTPS.
Snapshot Share	To create a snapshot share, select the <i>Create Snapshot Share</i> checkbox. 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, select the <i>Hide Snapshot Share</i> checkbox. If desired, enter a unique name for the snapshot share. Use up to 27 alphanumeric characters (including hyphens and spaces).

Editing or Deleting a Share

Change parameters as necessary, and then click **OK**.

Option	Description
Share Name	Accept the default share name or enter a new one. If you change the default, observe the following guidelines: <ul style="list-style-type: none"> Make sure the share name is unique to this server To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including hyphens and spaces).
Description	If desired, enter a description of the share. This is an opportunity to clarify the purpose of the share.
Hide this share	Select this option if you want the share to be hidden from network browsing.
Protocols	Select the access protocols for the share: Windows (SMB), Web View (HTTP), NFS, and FTP.
Snapshot Share	The option that displays depends on whether a snapshot share currently exists. <p>To create a snapshot share, select the <i>Create Snapshot Share</i> checkbox.</p> <ul style="list-style-type: none"> To hide the snapshot share from the SMB, HTTP, AFP, and FTP protocols, select the <i>Hide Snapshot Share</i> checkbox. If desired, enter a unique name for the snapshot share. Use up to 27 alphanumeric characters (including hyphens and spaces). <p>To remove a snapshot share, do the following:</p> <ul style="list-style-type: none"> Select the <i>Remove Snapshot Share</i> checkbox.

Deleting a Share

Click **Delete Share**. You can only delete one share at a time. Deleting a share does not delete any directories or files.

Configuring Share Access

GuardianOS supports share-level as well as file- and directory-level permissions (see [“Windows ACLs” on page 6-18](#)) for all local and Windows domain users and groups.

Topics in Configuring Share Access:

- Share Access Behaviors
- Setting User-based Share Access Permissions
- NFS Access

Share Access Behaviors

Administrators tasked with devising security policies for the SnapServer will find the following share access behaviors of interest:

- **Share access defaults to full control** – The default permission granted to users and groups when they are granted access to the share is full control. You may restrict selected users and groups to read-only access.
- **User-based share access permissions are cumulative** – An SMB, AFP, HTTP, or FTP user's effective permissions for a resource are the sum of the permissions that you assign to the individual user account and to all of the groups to which the user belongs in the Share Access page. For example, if a user has read-only permission to the share, but is also a member of a group that has been given full-access permission to the share, the user gets full access to the share.
- **NFS access permissions are not cumulative** – An NFS user's access level is based on the permission in the NFS access list that most specifically applies. For example, if a user connects to a share over NFS from IP address 192.168.0.1, and the NFS access for the share gives read-write access to * (All NFS clients) and read-only access to 192.168.0.1, the user will get read-only access.
- **Interaction between share-level and file-level access permissions** – When both share-level and file-level permissions apply to a user action, the more restrictive of the two applies. 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 except where a more restrictive file-level permission applies. The user has read-only access to File A.

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

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

Setting User-based Share Access Permissions

Share permissions for Windows, Apple, FTP, and HTTP users are configured from **Security > Shares** by clicking the link in the **Access** column next to the share you want to configure. Share permissions for NFS are configured and enforced independently. See [“NFS Share Access” on page 6-3](#) for more information.

User-based share access permissions apply to users connecting over SMB, AFP, HTTP, and FTP. Users and groups with assigned share access permissions appear in the list to the left (*Users/groups with access to*) and those without assigned access permissions appear in the list to the right (*Users/groups without access to*).

The default permission granted to users and groups when they are granted access to the share is full access. You may restrict selected users and groups to read-only access.

Share-Level Access Permissions

Full access	Users can read, write, modify, create, or delete files and folders within the share.
Read-only	Users can navigate the share directory structure and view files.

1 To add share access permissions for a user or group:

- Select a user or group from *Users/groups without access to*.
- Select either **Full Access** or **Read Only** from the drop-down list.
- Click **Add**.

Note: To search for a user or group, type the name in the Search box and click the **binoculars** icon. Search filters without wildcards are treated as substring searches and will find all entries containing the string you enter in the search field rather than looking for exact matches. For example, if you enter 'abc' as your search criterion, all users and groups containing 'abc' in the name will be identified. To clear a search and show the complete list, click the **world** icon.

2 To remove share access permissions for a user or group:

- Select a user or group from *users/groups with access to*.
- Click **Remove**.

3 To change access permissions for a user or group:

- Select a user or group from *users/groups with access to*.
- Select either **Full Access** or **Read Only** from the drop-down list and click the **Change Access** button.

4 Click **OK** to save share permissions.

NFS Access

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.

The NFS Access text box is a window into the client access entries in GuardianOS's *exports* file. This file serves as the access control list for file systems that may be exported to NFS clients. You can use the Add Host controls as described below to assist in making entries to the file, or you can directly edit the text box. After all entries are made, click **OK** to return to the **Security > Shares** screen.

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

The SnapServer Exports File Default Options

The default entry provides read-write access to all NFS clients (including NFSv4, if Kerberos security is not enabled). The options are explained in the table following the entry:

*** (rw, insecure, async, root_squash, no_all_squash)**

Value	Description
Asterisk	All NFS clients
ro	The directory is shared read only.
rw	The client machine will have read and write access to the directory.
insecure	Turns off the options that require requests to originate on an Internet port less than IPPORT_RESERVED (1024).
root_squash	Forces users connected as root to interact as the "nobody" user (UID 65534). This is the GuardianOS default.
no_root_squash	no_root_squash means that if root is logged in on your second machine, it will have root privileges over the exported filesystem. By default, any file request made by user root on the client machine is treated as if it is made by user nobody on the server. (Exactly which UID the request is mapped to depends on the UID of user "nobody" on the server, not the client.) If no_root_squash is selected, then root on the client machine will have the same level of access to the files on the system as root on the server. This can have serious security implications, although it may be necessary if you want to perform any administrative work on the client machine that involves the exported directories. You should not specify this option without a good reason.
async	Tells a client machine that a file write is complete - that is, has been written to stable storage - when NFS has finished handing the write over to the file system.
no_all_squash	Allows non-root users to access the nfs export with their own privileges.

Configuring Export Strings for NFSv4 with Kerberos Security

Share access for NFSv4 clients can be enforced either by the traditional NFS host method (described in The SnapServer Exports File Default Options) or via Kerberos.

If Kerberos is enabled, access is applied uniformly to all Kerberos-authenticated NFSv4 clients connected using the matching Kerberos option. Host-based access as described in The SnapServer Exports File Default Options still applies to NFSv2 and v3 clients when Kerberos is enabled, but it does not apply to NFSv4 clients.

When UNIX Kerberos security is enabled for NFSv4, the following entries are automatically added to the NFS Access settings for each NFS-enabled share:

```
gss/krb5 (rw, insecure, async, root_squash, no_all_squash)
gss/krb5i (rw, insecure, async, root_squash, no_all_squash)
gss/krb5p (rw, insecure, async, root_squash, no_all_squash)
```

These give read-write access to Kerberos-authenticated NFSv4 users connecting via:

- Standard Kerberos (**gss/krb5**)
- Kerberos with data integrity checksumming (**gss/krb5i**)
- Kerberos with protection/encryption (**gss/krb5p**).

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

Using the Add Host Controls

1. Select one or both 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**.

Creating Home Directories

The Home Directories feature creates a private directory for every local or Windows domain user that accesses the system. When enabling Home Directories (from the **Security > Home Directories** page), the administrator creates or selects a directory to serve as the home directory root. When a user logs in to the server for the first time after the administrator has enabled Home Directories, a new directory named after the user is automatically created inside the home directory root, and is configured to be accessible only to the specific user and the administrator.

Depending on the protocol, home directories are accessed by users either via a user-specific share, or via a common share pointing to the home directory root.

Home directories are supported for SMB, NFS, AFP, HTTP/HTTPS, and FTP/FTPS. They are accessed by clients in the following manner:

- For SMB, AFP, and HTTP/HTTPS, users are presented with a virtual share named after the user name. The virtual share is visible and accessible only to the user. Users are not limited only to their virtual shares; all other shares on the server continue to be accessible in the usual fashion.
- For NFS, the home directory is exported. When a user mounts the home directory root, all home directories will be visible inside the root, but the user's home directory will be accessible only by the user and the administrator.

Note: If desired, UNIX clients can be configured to use a Snap Home Directory as the local user's system home directory. Configure the client to mount the home directory root for all users, and then configure each user account on the client to use the user-specific directory on the SnapServer as the user's home directory.

- For FTP/FTPS, local users will automatically be placed in their private home directory when they log in. Access to the home directory is facilitated through a share pointing to a parent directory of the home directory, so users can still change to the top-level directory to access other shares.

If ID Mapping is enabled, domain users and local users mapped to the same user will be directed to the domain user's home directory. In some cases, data in the local user's home directory will be copied to the domain user's home directory:

- If a local user home directory accumulates files before the local and domain users are mapped, and if the domain user's home directory is empty, the local user's files will be copied to the domain user's home directory the first time the local user connects after the users are mapped.

- If both the local and domain user home directories accumulate files before the local and domain users are mapped, the files in the local user's home directory will not be copied to the domain user's home directory.

To configure home directories

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

Field	Description
Enable Home Directories	Click to put a check in this box to enable Home Directories for local users. Remove the check to disable.
Volume	Select the volume where the Home Directories will be located. Tip 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	Put a check by 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.

Windows ACLs

GuardianOS fully supports Windows NTFS-style file system ACLs, including configuration, enforcement, and inheritance models. Inside Windows/Mixed SnapTrees, files created and managed by Windows clients have the Windows security personality and behave just as they would on a Windows server. Clients can use the standard Windows 2000, 2003, XP, Vista, or Windows 7 interface to set directory and file permissions for local and Windows domain users and groups on the SnapServer.

Permissions are enforced for the specified users in the same manner for all client protocols, including non-SMB clients that normally have the UNIX security personality. However, if a non-SMB client changes permissions or ownership on a Windows personality file or directory (or deletes and recreates it), the personality will change to UNIX with the UNIX permissions specified by the client.

Note: Group membership of NFS clients is established by configuring the local client's user account or the NIS domain. Group membership of SnapServer local users or users ID-mapped to domain users is not observed by NFS clients. Therefore, ACL permissions applied to groups may not apply as expected to NFS clients.

Default File and Folder Permissions

When a file or directory is created by an SMB client, the owner of the file will be the user who created the file (except for files created by local or domain administrators, in which case the owner will be the "Administrators" group, mapped to the local `admingrp`), and the ACL will be inherited per the inheritance ACEs on the parent directory's ACL. The owner of a file or directory always implicitly has the ability to change permissions, regardless of the permissions established in the ACL. In addition, members of the SnapServer's local admin group, as well as members of Domain Admins (if the server is configured to belong to a domain) always implicitly have *take ownership* and *change ownership* permissions.

Setting File and Directory Access Permissions and Inheritance (Windows)

Access permissions for files and directories with the Windows security personality are set using standard Windows 2000, 2003, XP, Vista, 2008, or 7 security tools.

GuardianOS supports:

- All standard generic and advanced access permissions that can be assigned by Windows clients.
- All levels of inheritance that can be assigned to an ACE in a directory ACL from a Windows client.
- Automatic inheritance from parent directories, as well as the ability to disable automatic inheritance from parents.
- Special assignment and inheritance of the CREATOR OWNER, CREATOR GROUP, Users, Authenticated Users, and Administrators built-in users and groups.

To Set File and Directory Permissions and Inheritance (Windows)

1. Using a Windows 2000, 2003, XP, Vista, 2008, or 7 client, map a drive to the SnapServer, logging in as a user with change permissions for the target file or directory.
2. Right-click the file or directory, choose **Properties**, and then select the **Security** tab.
3. Use the Windows security tools to add or delete users and groups, to modify their permissions, and to set inheritance rules.

To View File and Directory Permissions and Inheritance (Web View)

1. Connect to the SnapServer *Web View*:
 - a In your browser, enter `http://[servername]`
 - b Log in as a user with admin rights on the SnapServer using the Switch User link.

Note: If Web Root is enabled, log in to the administrative interface via

`http://[servername]/config`

then point your browser directly to a share to browse via

`http://[servername]/[sharename]`

2. Browse *Web View* and click on the key icon to view security configuration on files and directories.

Security Guides

Security guides are designed to assist you in setting up security for your SnapServer. To navigate through the guides:

Use	To
Next	Proceed to the next page.
Back	Return to the previous screen.
Finish/OK	Return to the Security Guides menu.

The following guides are available:

Use Windows Active Directory Security

This security guide provides steps for configuring your SnapServer to use Windows Active Directory Security for Microsoft Networking. Once configured, the SnapServer will accept Microsoft networking (and Apple networking if desired) users and groups that are part of the domain. These users and groups can be granted (or denied) access rights for SnapServer network shares.

The SnapServer will need the name of your Active Directory domain, the name and password of an administrative user within your Active Directory domain, and the name of the organizational unit within the Active Directory tree in which the SnapServer will appear.

For more information about Windows Active Directory, please see [“Support for Windows Network Authentication”](#) on page 2-9.

Share-level Access to an Entire Volume

This security guide provides steps for allowing users share-level access to a whole volume on the SnapServer. You will need to know which user to grant access to and which volume they are to access.

Share-level Access to a Folder on a Volume

This security guide provides steps for allowing users share-level access to a folder on a volume on the SnapServer. You will need to know which user to grant access to and which folders they are to access.

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

Note: The Snapshot feature described here does not apply to snapshots for iSCSI disks. Supported Windows servers can create native snapshots of iSCSI disks using VSS. For more information, see “[Configuring VSS/VDS for iSCSI Disks](#)” on page 5-15.

Topics in Snapshots:

- [Snapshot Management and Usage](#)
- [Estimating Snapshot Pool Requirements](#)
- [Adjusting Snapshot Pool Size](#)
- [Accessing Snapshots](#)
- [Coordinating Snapshot and Backup Operations](#)
- [Snapshot Procedures](#)

[Isolate iSCSI Disks from Other Resources for Backup Purposes](#)

Snapshot Management and Usage

This section describes snapshot components and dependencies.

The Snapshot Pool

Snapshots are stored in a RAID in a *snapshot pool*, or space reserved within the RAID for this purpose. Each RAID on the system contains only one snapshot pool. This pool contains all snapshot data for all volumes on the RAID. For more information, see “[Estimating Snapshot Pool Requirements](#)” on page 7-2.

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. 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 screen and clicking the name of a snapshot. For instructions on accessing snapshot shares, see [“Accessing Snapshots” on page 7-3](#).

Rolling a Volume Back to a Previous State

If you need to restore an entire file system 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 file system to a previous state simply by selecting the snapshot and clicking the **Rollback** button. During the rollback operation, data on the volume will be inaccessible to clients.

CAUTIONS:

- 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 invalidate the NetVault for GuardianOS NVDB directory for the volume, and may also disable the antivirus software. If you are using these features, take the necessary precautions as described in [“Volumes” on page 4-11](#).
-

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 [“Coordinating Snapshot and Backup Operations” on page 7-4](#) for more information.

Snapshots and Backup Optimization

When you back up a live volume directly, files that reference other files in the system may become out-of sync in relation to each other. The more data you have to back up, the more time is required for the backup operation, and the more likely these events are to occur. By backing up the snapshot rather than the volume itself, you greatly reduce the risk of archiving inconsistent data. For instructions, see [“Coordinating Snapshot and Backup Operations” on page 7-4](#).

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 5-15](#).

Estimating Snapshot Pool Requirements

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

The default configuration allocates 80 percent of RAID capacity to the volume and 20 percent to the snapshot pool. You can adjust the size of the 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 pool by default, consider decreasing snapshot pool capacity and reallocating the capacity to the file system. To adjust the size of the snapshot pool, navigate to the **Storage > Snapshots** screen, click the **Snapshot Space** button, then click the Raid Set for the snapshot pool you want to adjust.

The number of snapshots that a RAID 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(s) since the snapshot was created

The following table describes minimum and maximum allocation cases.

Allocate about 10% of RAID if	Allocate about 25% of RAID 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

Adjusting Snapshot Pool Size

The current size of the snapshot pool for each RAID (or RAID group) can be viewed by navigating to the **Storage > Snapshots** screen and clicking the **Snapshot Space** button, then clicking the Raid Set. On the screen that opens, you can adjust the size of the pool as necessary. In addition, there are two other processes that may affect the size of the snapshot pool:

- **Creating a Volume** – In the course of creating a new volume, a drop-down menu allows you to add a percentage of the capacity being allocated to the new volume to the snapshot pool. 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 Pool Requirements.
- **Creating a RAID Group** – When two RAID sets are grouped, their snapshot pools are added together. For example, if RAID A with a snapshot pool of 50 GB is grouped with RAID B with a snapshot pool of 25 GB, the resulting RAID group will have a snapshot pool of 75 GB. Depending on the purpose you had in mind when grouping the RAID sets, the result of combining the two snapshot pools may or may not be desirable, and you will need to readjust the size as described previously.

Accessing Snapshots

Snapshots are accessed via a snapshot share. Just as a share provides access to a portion of a live volume (or file system), a snapshot share provides access to the same portion of the file system on all current snapshots of the volume. The snapshot share's path into snapshots mimics the original share's path into the live volume.

Creating a Snapshot Share

You create a snapshot share by selecting the **Create Snapshot Share** option in the course of creating a live-volume share, 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 Internet Explorer or other file browser, two shares display:

```
SALES
SALES_SNAP
```

The first share provides access to the live volume, and the second share provides access to any archived snapshots. Other than read-write settings (snapshots are read-only), a snapshot share inherits access privileges from its associated live-volume share.

Note: The same folders appear on the Web View screen when you connect to the SnapServer using a Web browser; however, the snapshot share folder does not provide access to the snapshot; it will always appear to be empty. You can prevent the snapshot share from displaying on this Web View screen by selecting the **Hide Snapshot Share** option when creating or editing a share.

Accessing Snapshots Within the Snapshot Share

A snapshot share contains a series of directories. Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created. For example, assume the snapshot share named *Sales_SNAP* contains the following four directories:

```
latest
2011-02-25.120000
2011-03-01.000100
2011-03-07.020200
```

The *latest* directory always points to the most recent snapshot (in this case, 2011-03-07.020200, or March 7th, 2011, at 2:02 a.m.). A user may view an individual file as it existed at a previous point in time or even roll back to a previous version of the file by creating a file copy to the current live volume.

Note: The latest subdirectory is very useful for setting up backup jobs as the name of the directory is always the same and always points to the latest available snapshot.

Coordinating Snapshot and Backup Operations

Like backups, snapshots can be scheduled to recur at a designated time and interval. 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.

1 Create a snapshot for each volume 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:

- Put a check in the **Create Recovery File** checkbox to ensure that the ACL, extended attributes, and quota information are captured and appended to the snapshot. This step is needed because many backup packages do not back up native ACLs and quotas. Placing this information in a recovery file allows all backup packages

to include this information. If the volume needs to be restored from tape, or the entire system needs to be recreated from scratch on a different server, this information may be required to restore all rights and quota information.

- Offset the snapshot and backup schedules such that the backup does not occur until you are sure the snapshot has been created. (The snapshot itself does not require much time, but creating the recovery file may take up to 30 minutes, depending on the number of files in the volume.) For example, assuming you schedule nightly backups for a heavily used volume at 3:00 a.m., you might schedule the snapshot of the volume to run every day at 2:30 a.m., allowing half an hour for the snapshot to run to completion.

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** screen, 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**.

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 at the root of the volume. For example, assume the snapshot share named `SHARE1_SNAP` contains the following four directories

```
latest
2011-02-25.120000
2011-03-01.000100
2011-03-07.020200
```

Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created. However, the `latest` directory always points to the latest snapshot (in this case, 2011-03-07.020200, or March 7th, 2011, at 2:02 a.m.). In this case, configuring the backup software to copy from:

```
\SHARE1_SNAP\latest
```

ensures that the most recently created snapshot is always archived.

Snapshot Procedures

Use the **Storage > Snapshots** page as the starting point for creating, scheduling, rolling back, and deleting snapshots. Be sure to read [“Snapshot Management and Usage” on page 7-1](#) before you begin using snapshots.

To Create a New Snapshot

Click **Create Snapshot**. The process involves first defining snapshot parameters, and then scheduling when and how often to run the snapshot.

Do not take more snapshots than your system can store, or more than 250 snapshots. Under normal circumstances, between nine and ten snapshots are sufficient to safely back up any system.

To Edit, Delete, or Roll Back a Snapshot

Click the snapshot's name. You can edit the snapshot's name, schedule, and duration, and you can roll back the snapshot to a volume, or delete the snapshot.

To Edit a Snapshot Schedule

Click the **Snapshot Schedules** button, and then click the snapshot name. You can modify all snapshot parameters.

To Adjust Space Reserved for Snapshots

Click the **Snapshot Space** button, then click the Raid Set for the snapshot space you want to adjust. You can adjust the amount of space allotted for snapshots on each RAID or RAID group.

Creating Snapshots

Creating a snapshot involves first defining the snapshot and then scheduling the snapshot. For regular data backup purposes, create a recurring snapshot that runs at an administrator-configured time and interval. 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 necessary.

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

To Create a New Snapshot

1 Create the snapshot definition.

Complete the following to define the snapshot:

- Name the snapshot (15 character maximum).
- Identify the source volume.

2 Run the snapshot now or schedule it to run later.

Click either the **Create Snapshot Now** option button to run the snapshot immediately or the **Create Snapshot Later** option button to schedule the Snapshot for a later time. If you wish to schedule the Snapshot at a later time, complete the following:

- Schedule a date and time to run the snapshot.
- To repeat a snapshot periodically, select **Recurring** and specify the repeat interval in hours, days, weeks, or months.
- See [“Scheduling Snapshots” on page 7-2](#) for more information on scheduling snapshots.

3 Specify the duration of the snapshot.

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

4 Specify whether to create a recovery file.

If you plan to create a backup from the snapshot and want to save file system security configuration in the backup, put a check in the **Create Recovery File** checkbox. (See the [“Coordinating Snapshot and Backup Operations” on page 7-4](#) for information on coordinating snapshots and backup operations.)

5 Create the snapshot.

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

Editing or Rolling Back a Snapshot

Consider the following before performing a rollback operation:

- Performing a rollback on a volume may invalidate the NetVault for GuardianOS nvdb directory for the volume. For this reason, it is important to synchronize NVDB directory backup operations with your snapshot schedule
- A rollback operation may also disable the antivirus software. If you are using these features, take the necessary precautions, as described in “Volumes” on page 4-11.

A rollback can disable Snap EDR and result in its removal. If this occurs, download the Snap EDR package from the [SnapServer website](#), reinstall it using the OS Update feature, then re-enable and configure it from the Snap Extensions page.

To Roll Back a Snapshot

Click **Rollback**. On the confirmation screen, click **Rollback**. During the rollback operation, access to the volume is blocked.

To Edit a Snapshot

Modify the properties as desired, then click **OK**.

To Delete a Snapshot

Click **Delete Snapshot**. On the confirmation screen, click **Yes**.

Editing Snapshot Schedules

The Scheduled Snapshots table lists details on the snapshots set to run at a future time. Use this screen to view, edit, or delete snapshot schedules.

See “Scheduling Snapshots” on page 7-2 for general information on scheduling snapshots.

To Reschedule a Snapshot

1. Click a schedule name to modify any of the following parameters:
 - Start time date and time
 - Whether to create multiple snapshots and the recurrence interval.
 - Snapshot duration
 - Whether a recovery file should be created
2. Click **OK**.

The Scheduled Snapshots table reflects your changes.

To Delete a Scheduled Snapshot

Click **Delete Snapshot Schedule**, then click **Delete Snapshot Schedule** on the confirmation screen.

Managing the Snapshot Pool

Snapshots require a certain amount of dedicated RAID capacity. You can increase or decrease the reserved snapshot space as necessary.

1 Review current snapshot space size and usage.

The Snapshot Space table shows the current allocation for each RAID, as well as how much of that space is already in use.

2 Estimate snapshot space requirements.

The number of snapshots that the SnapServer can support is a function of the space reserved for the snapshots, the duration of snapshots you create, and the amount and type of write activity to the volumes since the snapshot was created.

Allocate about 10% of RAID if

- 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

Allocate about 25% of RAID if

- 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

3 At Storage > Snapshots, click the RAID Set you want to modify, enter the new allocation and click OK.

Enter the snapshot space (using no commas), and then select the appropriate unit of measurement from the drop-down list.

Disaster recovery entails creating the files you need to recover a SnapServer's configuration information, such as network and RAID configurations, as well as volume-specific information, such as ACLs and quota settings.

It also includes what to do if all access to the data on a SnapServer is cut off due to a hardware or software failure. Focus is placed on the procedures for:

- Reinstalling the SnapServer operating system
- Restoring the server to its original configuration with data intact

These files are then used to restore any SnapServer to its original state. The disaster recovery feature can also be used to clone one server to another by restoring the disaster recovery image from one server to another server.

Topics in Disaster Recovery:

- [Backing Up Server and Volume Settings](#)
- [Backing Up the NetVault Database Directory](#)
- [Recovering the NetVault Database](#)
- [Disaster Recovery Procedural Overview](#)
- [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 **Maintenance > Disaster Recovery** screen 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.

The SnapDRImage File and the Volume Files

Details on the SnapServer disaster recovery files and the information they contain are as follows:

- **SnapDRImage** — The SnapServer disaster recovery image saves server-specific settings such as network, RAID, volume and share configuration, local user and group lists, and snapshot schedules, and Snap EDR Management Console settings (if applicable). There is one SnapDRImage file per server, residing on the root directory of the first volume at the following path: `\\server_name\volume_name`

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, residing at the following path: `\\server_name\volume_name\.os_private`

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 **Maintenance > Disaster Recovery** page. This DRImage should be created after server configuration is complete, and can be used to recover the server or a replacement server to the configured state.

Before you create the disaster recovery files, make sure you have completed the following activities:

- You have completely configured the SnapServer. If you subsequently make any major changes to the configuration of your server, you must repeat the procedures described in this section to have an up-to-date SnapDRImage.

Note: You may want to record, in an off-server location, the following information about the configuration of your server: (1) the server name; (2) the number of RAID's; (3) the number of volumes; and (4) the size of each volume. If the disaster recovery fails, having this information may be useful in recreating the original configuration of the server.

- You have devised and implemented a data backup strategy. It is recommended that you make a backup of your system regularly, from the root of the share for each volume, and store it in an off-server location. This ensures that the most current data is backed up and available for use with a disaster recovery.

Use the following procedure to create and secure the disaster recovery files:

1 Create the disaster recovery files.

Navigate to the **Maintenance > Disaster Recovery** screen. Select the **Create Recovery Image** option button and click **OK** to create the SnapDRImage file and the volume files in a single operation.

2 Copy the files to a safe place off the server.

Once the recovery image has been made, click the **Download Recovery Image** button to download the SnapDRImage file to a safe location on another server or backup medium. (See The SnapDRImage File and the Volume Files for file names and paths.) This strategy ensures that if the file system 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.

3 Take no action regarding the volume-specific files.

These files will be copied to tape as part of your regular volume backup procedures.

Restoring Original Server and Volume Configurations

To restore the original configurations to the server, two separate operations are required and they must be run sequentially. After you start any of the recovery processes, you will see the Disaster Recovery Status screen.

CAUTION: Do not try to navigate back from this screen during the recovery process. Activity is restricted to this screen so that the recovery operation will not be interrupted.

1 Restore server settings.

Select the **Recover System Settings** option button and click **OK** to open the Server Recovery screen. Use the **Browse** button to navigate to the SnapDRImage file. Click **Recover** to start the operation. If the recovery file contains Snap EDR application settings, you are asked if you want to include those settings. Check the settings you want to recover, and click **Recover**. Once the server configuration recovery operation is complete, you can start the volume configuration recovery operation.

2 Restore volume ACL and quota configurations.

Select the **Recover Volume Security Settings** option button and click **OK** to open the Server Recovery screen. Select the volumes you want to restore. (Volumes that do not have a recovery file attached will appear as unavailable.) The creation date of the recovery file on a volume indicates when the recovery image was generated. Click **Recover** to start the operation, and then follow the onscreen instructions.

Note: Restoring the server and volume configurations will only occur if the system can accomplish it without compromising data. If the current configuration is different than previous SnapDRImage files, the restore will fail. A successful restoration requires either that the current RAID and volume configuration exactly match the saved one or that there is no current configuration (such as, you are using raw drives) so that the saved configuration can be recreated with the available raw drives. View the log file if the recovery operation fails.

Rejoining the Server to a Windows Domain

If you are restoring server settings to either the same physical server or to a replacement server, the server will automatically rejoin the Windows domain it was a member of before the SnapDRImage was applied as long as the servername is the same as the current servername. If you have changed the servername, you will have to manually join the server to the desired Windows domain. Navigate to **Network > Windows** to rejoin the server to a domain.

SnapDRImage Usage Scenario

The SnapDRImage contains the server configuration settings for a specific server. These settings may be useful in situations other than disaster recovery.

Reset Server Configuration After Swapping Out Components

In the extreme situation that one of the major components of your SnapServer fails (example: the mother board), and new one has been provided to you, it is possible to reset the sever configuration using the SnapDRImage. Follow the recovery procedures to utilize this option.

Cloning Servers

You can use SnapDRImages to copy server configurations from one server to another. For example: if you have a SnapServer N2000 configured for peak performance in your network environment, you can create a SnapDRImage of this server, then apply the SnapDRImage to a new SnapServer N2000. The server settings and configuration would be identical to the first.

Note: Cloning a disk configuration will only succeed if the destination server has sufficient storage resources to duplicate the original configuration.

Backing Up the NetVault Database Directory

This section details the use of the NetVault Database plug-in and offers various tips for its use.

Note: NetVault is only available for servers upgraded from GuardianOS 5.2.

Backup Recommendations

It is important to note that the NetVault Database can be backed up at any time as long as no other NetVault jobs controlled by this server are running. With this in mind, the following points are recommended when backing up the NetVault Database:

- **Perform Regular Backups** – The data contained in the NetVault Database is integral to NetVault operations, but it also frequently changes as NetVault functions; therefore, it is recommended that frequent, regular backups of the NetVault Database be performed (for example, daily, once all other backups have completed).
- **Target Specific Media for a NetVault Database Backup** – In the event that the NetVault Database needs to be recovered, the specific piece of media targeted can be easily located to perform the recovery.

To Back up the NVDB Directory

1. From the NetVault Server (either locally or remotely), open the NetVault Backup window by clicking the **Backup** button on the command toolbar. The NetVault Backup window displays the list of available clients in the Selections tab.
2. Right-click the NetVault Server (acting as a client to itself) and select **Open** from the pop-up menu.
3. The available plug-ins will be displayed. Right-click the NetVault Database Plug-in and select **Open** from the pop-up menu that appears.

4. A single selectable item will be revealed: the NetVault Database. Select the checkbox to the left of this item.

Note: There are no Backup Options available for use with this plug-in.

5. The remaining tab selections (Schedule, Target Advanced Options) contain additional options that can be set as desired.
6. Enter a suitable name for the job in the Job Title box and start the backup job by clicking the **Submit** button on the command toolbar.

Note: Only clients successfully added via the NetVault Client Management window will display.

Recovering the NetVault Database

This section summarizes the procedure necessary for recovering the NetVault Database (NVDB) from tape. For instructional details, see the NetVault for GuardianOS documentation that shipped with your SnapServer.

Pre-Restore Requirements

Before restoring the database, perform the following steps on the SnapServer acting as the NetVault Server:

1. Completely reinstall and configure the same version of GuardianOS that the server was running. The OS installation will also reinstall the NetVault Server software.
2. If necessary, navigate to the **SnapExtensions** screen and re-enable the NetVault for GuardianOS software.
3. Remove all media from the device(s) used by the NetVault Server, except the media that contains the backup saveset needed for the recovery of the NVDB.
4. Add all devices previously added to the NetVault Server through the use of the Device Management window.
5. From the Device Management window, the media containing the backup saveset will be recognized as FOREIGN in its designated drive or library slot. Scan the media before proceeding with the restore operation.

Restore Recommendations

The following recommendations are offered for the process of recovering the NVDB:

- **Perform a Full Recovery of the NetVault Database** – Although NetVault offers provisions for recovering individual elements of the NVDB, it is recommended that a full recovery be performed. If recovering individual components, it is strongly recommended that this be performed under the guidance of BakBone Technical Support.
- **Do Not Monitor Job Progress During a Recovery** – It is strongly recommended that all NetVault windows be closed, and the NetVault GUI be closed during the recovery of the NVDB, as this may interfere with the process.

Restore Procedure

1. Access the Restore window from the NetVault GUI by clicking the **Restore** button in the command toolbar.

2. Double-click the **NetVault Server** that the desired backup was performed from to open it.
3. Plug-ins (and APMs) used to conduct successful backups on the selected client will be displayed. Double-click the **NVDB Plug-in** to open it.
4. All of the backup savesets are created using the NVDB Plug-in display. Locate the desired saveset, right-click it and select **Open** from the pop-up menu.
5. All of the various components that make up the NVDB will display. Items with checkboxes at their left are single items that can be selected for inclusion, while items without checkboxes can be double-clicked to browse their individual contents.
6. For a full database restore, select each item in the tree. Additionally, open up root items to display their contents by double-clicking them, and then select all of their contents (for example, Events, Notification and Reports Database items).
7. Select the **Restore Options** tab and make sure that the Blank Reports Database Table option is selected.
8. Other tab selections (for example, Schedule and Advanced Options) contain additional options that can be set as desired.
9. Enter a suitable name for the job in the Job Title box and start the restore job by clicking the **Submit** button.
10. The job will now run and the backed-up version of the NVDB will be restored over the one created with the recent installation of NetVault.
11. Once the NVDB has restored successfully, it is necessary to restart NetVault Services via the NetVault Configurator. During the restore procedure these services are automatically stopped.

Disaster Recovery Procedural Overview

The procedure described in this section for responding to a catastrophic event is general in nature and may result in the loss of data. Should such an event actually occur, the exact procedure to follow will vary according to environmental conditions. Overland Storage strongly recommends that you contact a technical service representative before proceeding.

This section describes a worst-case scenario:

- The operating system has failed, (for example, due to a malicious attack to the root file system), and you cannot access the server.
- The data has been corrupted and must be restored from tape.
- Technical support has deemed your server unsalvageable and provided you with a new, unconfigured server.

Restoring Previous Server Settings to a New Server

After Technical Support has supplied you with a new server, you can restore the settings from the previous server to the new server. Any third-party license keys you have not purchased through Overland Storage will be lost. If you have installed data

replication or management utilities such as Snap EDR, you will need to re-install and/or relicense them for use with the new server.

Note: If you are restoring Snap EDR Management Console settings, you must recreate the RAID and volume configuration that matches the DRI settings, then install and enable the Snap EDR Management Console. As an alternative, you can first restore just the system settings, install Snap EDR, and then restore just the Snap EDR settings.

You will also need to reschedule snapshots as well reconfigure CA Antivirus.

Note: If you are restoring the DRImage to the same server, all your license keys should be intact. You will still need to reschedule your snapshots and CA Antivirus.

1. When you connect to the new server, navigate to **Maintenance > Disaster Recovery**, select **Recover System Settings** and click **OK**.
2. Click the **Browse** button and navigate to the SnapDRImage you made of the previous server, then click **Recover**.
3. The server reboots and the settings are restored. To view the log, click the date link on the Disaster Recovery screen after the server has rebooted.
4. After restoring your server settings, rejoin the server to the Windows domain if necessary.
5. Now you can replace your data from tape backup. If the backup doesn't retain permission and ownership settings, you can restore these by selecting **Recover Volume Security Settings** on the **Maintenance > Disaster Recovery** screen.

Note: If you are restoring from any backup other than NetVault, you will need to recover the volume settings.

Cloning a Server

The Disaster Recovery process can be used to clone a server in order to apply the same configuration to one or more servers. To clone a server:

1. Create a disaster recovery image on the source server (refer to Creating the SnapDRImage and Volume Files).
2. Copy the disaster recovery files from the source server to a client.
3. Perform a disaster recovery restore procedure to each of the clone target servers using the disaster recovery files from the source server (refer to Restoring Previous Server Settings to a New Server).

The CA Antivirus software is preinstalled on all GuardianOS SnapServers. By default, the software is enabled on most SnapServers, but no scan jobs or signature updates have been scheduled. (The server will, however, check for signature updates whenever the server boots.) These and other antivirus configuration and management tasks are performed using the CA Antivirus GUI, accessed from the **SnapExtensions > CA Antivirus** screen of the Web Management Interface. This section outlines the major steps in configuring the antivirus software. See the GUI online help for detailed descriptions of all options.

Topics in CA Antivirus Software:

- [Antivirus Dependencies](#)
- [Launching the CA Antivirus GUI](#)
- [The Local Scanner View](#)
- [Scan Job Configuration and Scheduling](#)
- [Signature Updates](#)
- [Alert Options](#)
- [The Move Directory](#)
- [Log View](#)

Note: Antivirus functions or options not relevant to the SnapServer have been disabled in the configuration GUI.

Antivirus Dependencies

The SnapServer implementation of CA Antivirus software includes the following features:

HTTP Access and Antivirus Configuration

To access the CA Antivirus configuration interface, HTTP must be enabled on the **Network > Web** screen.

Resetting the Server Date and Time

If the current server date and time are changed to an earlier date and time (**Server > Date/Time**), the change does not automatically propagate to any scheduled antivirus

operations. To synchronize scheduled antivirus operations with the new date and time settings, you must reschedule each operation.

Note: New jobs may be affected by the time change. Be sure to check that new jobs have been executed if a date or time change has been made to the server.

Storage Configuration and the Antivirus Software

The antivirus software resides on the largest volume (that existed at the time the software was installed). If you delete this volume, the CA Antivirus software will also be deleted. The SnapServer automatically reinstalls the antivirus software on the largest remaining volume on the system.

Note: The antivirus reinstallation process does not preserve custom antivirus configuration settings. Make a note of any such settings before deleting a volume.

Launching the CA Antivirus GUI

The CA Antivirus software on SnapServers upgraded to GuardianOS 6.5 is enabled by default. Some situations, such as deleting a volume or performing an upgrade procedure, may require you to re-enable the software. To learn how the antivirus software interacts with other GuardianOS software components, see [“Antivirus Dependencies” on page 9-1](#).

Note: Antivirus functions or options not relevant to the SnapServer have been disabled in the configuration GUI;

Launching the CA Antivirus Browser Interface

The first time you connect to the CA Antivirus GUI, it may take from 30 seconds to several minutes for the application to load, depending on the speed of your connection.

1. If you need to enable the antivirus software, go to **SnapExtensions > CA Antivirus**, click the checkbox next to **Enable**, and click **OK**.
2. Click the **Configure Antivirus** link. The splash screen opens first, followed momentarily by the GUI login dialog box.
3. Enter the same administrative user name and password (case sensitive) you have established for the GuardianOS Web Management Interface, and then click **Login**. The antivirus GUI box opens.

The Local Scanner View

Use the Local Scanner view ([Figure 9-1](#)) to scan a SnapServer for infected drives, folders, files, or disks on demand.




Figure 9-1: Local Scanner View

Component	Description
Root Directory	Displays the directory structure of the SnapServer. As in Windows Explorer, click folder icons to navigate the structure and display subfolders and files in the right-hand pane.
Move Folder	May contain infected files. The administrator can instruct the software to automatically move infected files to this directory. For more information, see “Scan Job Configuration and Scheduling” on page 9-3 .
Scheduled Scan Jobs	Scan Jobs you schedule appear in this folder. For more information, see “Scheduling a Scan Job” on page 9-5 .

Scan Job Configuration and Scheduling

You can run scan jobs on demand or you can configure scan jobs to run periodically. This section outlines the process of configuring and running manual and scheduled scans. For detailed descriptions of all scanning options, see the CA Antivirus online help.

Note: You may not want to include Snapshot shares (see [“Snapshot Management and Usage” on page 7-1](#)) as part of your virus scan. Because access to an archived version of the file system provided by a snapshot share is read-only, you cannot treat or move any infected file; you would have to delete the entire snapshot to effect a cure. A more useful approach is to always scan your file system for viruses before running a snapshot. Adjust your antivirus scan schedule to synchronize with your snapshot schedule so that any infected files are cured (repaired) or removed before the snapshot is scheduled to run.

Topics in Scan Job Configuration and Scheduling:

- Defining Scan Jobs
- Running a Manual Scan Job
- Scheduling a Scan Job

Defining Scan Jobs

This section provides an overview of the major choices available in configuring scan jobs. Access these options by selecting **Local Scanner Options** from the Scanner Menu.

Choosing an Infection Treatment (Scan Tab)

You can instruct the software to perform one of the following file actions when an infected file is found:

File Actions	Description
Report Only	(Default) Reports when an infection is found.
Delete File	Deletes an infected file.
Rename File	Renames an infected file with an AVB extension. Infected files with the same name are given incremental extensions (for example, FILE.O.AVB, FILE.1.AVB, and so on). After a file is renamed with an AVB-type of extension, it is not scanned subsequently.
Move File	Moves an infected file from its current directory to the Move directory for quarantine.
Cure File	Attempts to cure an infected file automatically. Choosing this setting enables the File Options button. Click this button to display the Cure Action Options and specify how the Cure File option performs.

Note: The *System Cure* option is not available on SnapServers.

Setting the Type of Files to Scan (Selections tab)

Use the Selections tab options to choose the types of objects to scan, the types of file extensions to include or exclude from a scan, and the types of compressed files to scan.

- **File Extensions** – You can choose to scan files regardless of extension, or select specific types of extensions to include or exclude.
- **Compressed Files** – To scan compressed files, select the *Scan Compressed Files* checkbox, and then click **Choose Type** to specify the compressed file extension types.

Filtering File Information for Logs (Manual Scans Only)

You can specify the types of events that are written to a log. Check the *Infected files* option to put information in the log about files that are found to be infected. Check the *Clean files* option to put information in the log about files that are scanned and are not infected. Check the *Skipped files* option to put information in the log about files that have been excluded from the scan.

Running a Manual Scan Job

Before running a local scan job, confirm that the scanner options are correctly configured as described in the previous section “[Defining Scan Jobs](#)” on page 9-4.

- 1 In Local Scanner View, select the folders you want to scan.**

The left-hand pane displays the directory structure of the SnapServer. A red check mark on a folder or file indicates that it is selected for scanning. (By default, all directories and files are selected for scanning.) Click folders or files to toggle file/folder selection on or off.

- 2 Run the scan.**

Select **Scanner > Start Scanning**. The interface is unavailable for further configuration while the scan is in progress. The scan results display in the lower pane of the Local Scanner View, and the action taken with each file is listed in the Status column.

Scheduling a Scan Job

A scan job is configured and scheduled in the Schedule New Scan Job dialog box. To open this dialog box, select the **Scanner > Schedule Scan Job > Create** command.

- 1 Set scan options in the Scan and Selection tabs.**

These options are summarized in “[Defining Scan Jobs](#)” on page 9-4.

- 2 Schedule the scan.**

The Schedule tab allows you to set a start date and a repeat interval for the scan.

- 3 Select the directories to scan.**

The Directories tab lists all paths that currently exist on the server. You can remove or add new paths as desired. You can also use the Exclude Directories tab to achieve the same result.

- 4 Click OK.**

You can view scheduled scan jobs by clicking the **Scheduled Scan Jobs** folder in the Local Scanner View. To edit a job, right-click it and select **Options**.

Signature Updates

Signature updates contain the latest versions of the signature files that recognize the latest infections. They also contain the latest engine versions, which do the work of looking for infections. Signature updates are made available on a regular basis by Computer Associates.

These updates are cumulative, so they contain everything from all previous file updates, plus the newest information on the latest infections. If you have missed a recent update, you only need to collect the latest signature file to have the most up-to-date protection.

SnapServers are preconfigured to download signature updates from the CA FTP site at <ftp://ftpav.ca.com/pub/inoculan/scaneng>. By default, no signature updates are scheduled. The antivirus software will, however, check for signature updates

whenever the server is powered on. To update SnapServers that do not have Internet access, the following methods are available:

Method	Description
FTP	Use FTP to download the update files from the Computer Associates FTP site. You can also use FTP to distribute signature updates from one SnapServer (or any FTP server) to another. Note: When using FTP, the user name and password are passed as clear text.
UNC	Use UNC to distribute signature updates from one SnapServer to another (or from any arbitrary SMB or Windows server). Note that for UNC to work, you must have the Enable Guest Account option enabled (Network > Windows) on the SnapServer on which the signature updates reside. Note: Alternatively, you can distribute updates to SnapServers from any Windows/SMB server. If using this method, make sure the guest account on the chosen server exists, is enabled, and has a blank password.
Local Path	As part of the procedure to provide signature updates to the SnapServer with no Internet access, you can connect to a local path relative to the root (for example, /shares/SHARE1/virusdefs). Note that the path to the share is case sensitive.

Topics in Signature Updates:

- Updating SnapServers that have Internet Access
- Updating a SnapServer that does not have Internet Access
- Distributing Updates from One Server to Another
- Verifying Download Events

Updating SnapServers that have Internet Access

If your SnapServers have direct access to the Internet, you only need to schedule the downloads to set up automatic signature updates. If access to the Internet is routed through a proxy server, you may also need to specify the name of the proxy server. Both procedures are explained below:

To Schedule Signature Update Downloads

1. Choose **Scanner > Signature Update Options**.
2. On the Schedule tab, click **Enable Scheduled Download**. Select the initial download date and time, then select how often to repeat the download.
3. Click **OK**.

To Specify a Proxy Server

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
2. Select *FTP* in the list box, then click **Edit**.
3. In the Proxy Name field, enter the IP address of the proxy server, then click **OK**.

Updating a SnapServer that does not have Internet Access

If you have SnapServers that do not have Internet access, use the following procedures to download the signature files to a machine with Internet access and then copy them to the SnapServer.

Note: When retrieving signature updates, the antivirus software attempts to connect to all the sites in the site list in the order they are listed. To avoid delays or superfluous error messages, delete the default FTP option from the list on SnapServers that have no Internet access.

1. Using a workstation with Internet access, go to <ftp://ftpav.ca.com/pub/inoculan/scaneng> and download the following files.
 - All *.tar files containing the word *Linux*, for example, *fi_Linux_i386.tar* and *ii_Linux_i386.tar*
 - All *.txt files containing the string *Sig*, for example, *Siglist.txt* and *Siglist2.txt*
2. Using a method appropriate to your environment, copy the update files to the SnapServer.
3. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
4. Click the **Add** button, then select *Local Path* from the Method drop-down menu.
5. In the Path field, enter the path to the directory on the server on which the update file resides. If you are using a SnapServer, the path would be similar to the following:

/shares/SHARE1/sigfiles

where *SHARE1/sigfiles* is the share path to the directory containing the signature update files.
6. Click **OK**. The path appears in the list box.
7. Click **Download Now**.

Distributing Updates from One Server to Another

When retrieving signature updates, the antivirus software attempts to connect to all the sites in the site list in the order they are listed. To avoid delays or superfluous error messages, delete the default FTP option from the list on SnapServers that have no Internet access.

To Distribute Files via UNC

If you have more than one SnapServer with no Internet access, you can perform the previous procedure on just one of them (or any Windows/SMB server), and then

configure your other SnapServers to get the update from that server automatically via SMB by specifying the UNC of the server containing the signature files.

Notes: The following conditions must be met in order to distribute updates using UNC:

- The correct Signature files must have been downloaded to the root of the share being used for updates.
- The server containing the Signature updates must have the Guest account enabled (**Network > Windows**) in GuardianOS. For other SMB/CIFS servers, the Guest account must have no password, and there may be additional requirements (for example, Windows servers must allow anonymous connections).
- The share and Signature files must be accessible to the Guest account.
- The server name used in the UNC must be resolvable by the server running CA Antivirus.

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
2. Click the **Add** button, and select **UNC** in the Method list box.
3. Enter the path to the SnapServer (or Windows/SMB server) to which the update files have been downloaded (see previous procedure) using the following format:

```
\\server_name\share_name
```

where *server_name* is the name of the server, and *share_name* is the name of the share providing access to the files. (On a SnapServer, the update files must reside on the root of the share.)
4. Click **OK**. The path you entered appears in the Download Sources list box.
5. Click **Download Now**.

To Distribute Files via FTP

If you have more than one SnapServer with no Internet access, you can perform the FTP download procedure on just one of them (or any FTP server), and then configure your other SnapServers to get the signature updates from that server automatically via FTP.

1. Navigate to **Scanner > Signature Update Options**, and click the **Incoming** tab.
2. Click the **Add** button, and select **FTP** in the Method list box.
3. Enter the following information regarding the server on which the update file resides as follows:
 - In the Host Name field, enter the IP address.
 - In the User Name and Password fields, enter the admin user name and password.
 - In the Remote Path field, enter the path to the directory in which the file resides. If you are using a SnapServer, the path would be similar to the following:

```
/shares/SHARE1/sigfiles
```

where *SHARE1/sigfiles* is the share path to the directory containing the signature update files.

4. Click **OK**. The path you entered appears in the Download Sources list box.

5. Click **Download Now**.

Verifying Download Events

Use the following procedure to verify download and distribution events.

1. Select **View > Log Viewer**.
2. In the left-hand pane, select **Distribution Events**. Distribution events are listed in the upper right-hand pane in chronological order.
3. Select a distribution event. The details of the distribution event display in the lower pane.

Alert Options

Alert options allow you to tailor the notification information that is provided to the Alert Manager, cut down on message traffic, and minimize the dissemination of notifications that are not critical. To set alert options, select **Alert Options** from the Scanner menu. The Alert Options dialog box contains the following tabs:

Tab	Description
Report	Use the Alert Report options to specify where to send notification information, and the Report Criteria options to manage how frequently messages from the General Event Log are reported. Note: The Local Alert Manager option is not supported on SnapServers.
Alert Filter	Use the Alert Filter options to manage notification severity levels, and to determine what types of messages should be passed to the Alert Manager. Note: In the Custom Notification Module, the <i>Realtime Server</i> and <i>Admin server</i> settings have no effect on SnapServers.

The Move Directory

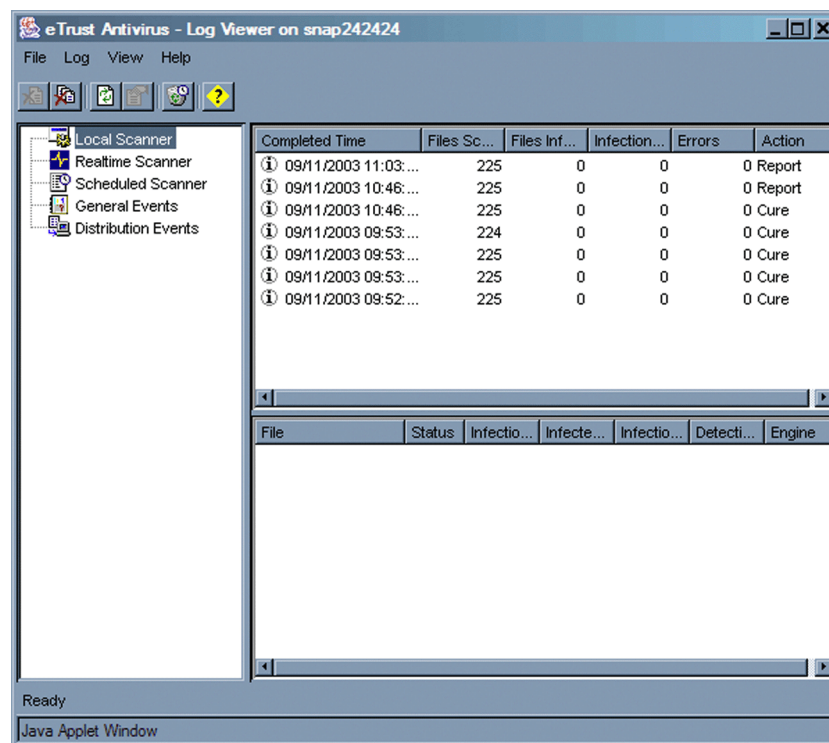
You can configure scans to move infected files to the move folder (**Scanner > Local Scanner** options). To view infected files, click the **Move** directory on the left-hand pane of the Local Scanner View. To manage a moved file, right-click the file and select from the following options:

Option	Description
Restore	This option removes the file from the Move Folder and restores it to its original location with its original name and type.
Restore as	This option displays a dialog box that allows you to change the directory location and file name. You can rename a file and isolate it safely in a different location. You may want to use this option, for example, if you do not have another source for the data and you need to look at the file. Or you may have a file that you want to analyze. Note: To restore a file to a different directory, you must prepend the path to the directory with the string “/shares.” For example, to restore a file to the SHARE1/sales directory, enter the path as follows: /shares/SHARE1/sales

Option	Description
Restore and Cure	This option allows you to restore the selected item back to the original folder it was in, and cure it. This option is useful if you update the signature files after items have been put in the Move folder. If a cure is provided that you did not have available, you can get the latest signature update and use this option to restore and cure an infected item.
Delete	This option deletes the infected file; no warning or confirmation message is displayed.

Log View

The Log View provides easy access to detailed information on scan, distribution, and other events. To access this view select **Log View** from the View menu.



Option	Description
Local Scanner	Displays summary information about scan jobs that have run.
RealTime Scanner	Not supported.
Scheduled Scanner	Displays summary information on scheduled scans that have run.
General Events	Displays the Event log for a given day. Click a date to view all events that occurred that day.
Distribution Events	Displays distribution events by date. Click a date to view detailed information on the distribution event in the lower pane.

Unicode defines a universal means of representing characters in all languages. In the case of SnapServers, this allows better interoperation of varying languages using different alphabets and character sets in file and user names. More information is available at <http://www.unicode.org>.

CAUTION: In GuardianOS 6.5 and beyond, Unicode is always enabled and cannot be disabled. Upgrading to GuardianOS 6.5 automatically enables Unicode if not previously enabled, and may alter the functionality of some third party applications and SnapExtensions that do not fully support Unicode. Review the sections below.

Topics in Unicode:

Unicode and Protocol Interaction

Extended characters in filenames are encoded on the SnapServer file system using UTF8, a method of representing all Unicode characters. However, network protocols and clients vary in their support of Unicode and UTF8, which has ramifications in the way they interact with one another when sharing files with extended characters in filenames.

The following sections describe how different protocols interact with extended characters.

SMB

Most Windows and Mac OS X clients, as well as the SMB protocol, support the majority of Unicode characters. Therefore, in general, all characters written by Windows and Mac OS X clients will be properly retained and visible to other Windows and Mac OS X clients and Unicode-compliant protocols.

However, if there are characters on the file system that are invalid UTF8 or are otherwise not mappable to the Unicode encoding method (UCS2) used by the SMB protocol, an escape sequence will display in the file name of the file being read. Escape sequences begin with “{!^”. The following two characters are the hexadecimal value of the characters in the filename; for example, you might see “{!^AB” in a file name. Windows and Mac OS X clients can edit such files, and the names will be retained in their original form when written back to the file system.

AFP

Mac OS X and higher use the same method to represent Unicode characters as the SnapServer: UTF8. Information written to the server from Mac OS X or higher will be encoded with UTF8 and should be viewed correctly from the Mac OS UI. However, similarly to SMB clients, characters in filenames that are incompatible with UTF8

will be returned with an escape sequence. Escape sequences begin with `{!^`. The following two characters are the hexadecimal value of the characters in the filename; for example, you might see `{!^AB` in a file name. Mac OS X clients can edit such files, and the names will be retained in their original form when written back to the file system.

NFS

The NFS protocol is not Unicode-compliant or -aware. Additionally, there is no means for the SnapServer to determine what method is being used by the client to represent extended characters. Currently, the code pages most commonly used in Linux environments are: UTF8, 8859-1, 8859-15, and EUC-JP. The SnapServer then must make an assumption to enable it to translate to and from UTF8 on the file system. Therefore, when in Unicode mode, you must configure the SnapServer's NFS protocol for the code page being used by NFS clients. Code page options include ISO-8859-1, ISO-8859-15, EUC-JP, and UTF8.

Any extended characters on the file system that cannot be translated to the configured NFS code page will be returned to the NFS client with an escape sequence. Escape sequences begin with `{!^`. The following two characters are the hexadecimal value of the characters in the filename; for example, you might see `{!^AB` in a file name.

FTP

FTP only supports ASCII characters by specification. Some clients bend the specification to allow extended characters, but there is no standard means of representing them. Therefore, no translation is performed on extended characters for FTP clients — all filenames are written to and read from the file system as a “bag-of-bytes”. This has two ramifications: extended characters written to the file system by other protocols will be visible to FTP clients as garbled characters; and FTP clients are able to write invalid UTF8 characters to the file system. For the latter case, when other protocols encounter invalid UTF8 characters on the file system (which normally can only be written by FTP), the characters will be returned in an escape sequence. Escape sequences begin with `{!^`. The following two characters are the hexadecimal value of the characters in the filename; for example, you might see `{!^AB` in a file name.

HTTP

HTTP integrates easily with Unicode and the SnapServer. If invalid UTF8 characters are encountered on the file system, the characters will be returned with an escape sequence. Escape sequences begin with `{!^`. The following two characters are the hexadecimal value of the characters in the file name; for example, you might see `{!^AB` in a file name.

CA Antivirus

CA Antivirus is not Unicode-aware. While the CA Antivirus UI displays garbled characters for extended characters when Unicode has been enabled, it can still scan files, find viruses, clean viruses, move, and rename virus-infected files.

How Snapshots Interact with Unicode

Snapshots taken before your SnapServer was upgraded to GuardianOS 6.5 (or later) with Unicode enabled are not compatible with the upgraded SnapServer. It is not recommended that a pre-Unicode snapshot be used to restore a post-Unicode server.

Note: It is recommended, if you have snapshots on your server from pre-Unicode conversion, you delete all snapshots once the server has been upgraded to Guardian 6.5 or later.

Backing Up Unicode Servers

Backing up a Unicode-enabled SnapServer requires you to use specific methods depending on the type of client you have in use. It is recommended that like languages be used across the backup process. For example: Russian files on a localized Russian server should be backed up with a Unicode-compliant/Russian localized backup application. Mixing languages between applications can result in data corruption.

The following table gives an overview of how Unicode interacts with backup applications:

Application	Officially Supports Unicode	UI Displays Correct Filenames	Backups and Restores Unicode Data
Bakbone NetVault over client ¹	No	No	Yes
Symantec Backup Exec 10.d, 11.d, 12.d over SMB only	No	Yes	Yes
Symantec NetBackup 6.5	No	No	Yes
CA ARCserve 11.5, 12.0 over SMB only ⁴	No	No	Yes
EMC Networker v7.3, 7.4 over SMB only	No	No	
Snap EDR over Sync only	No	No	Yes

1. The UI displays garbage but the data is intact.

Backing Up Using Unicode-Enabled Windows Clients

When backing up using a Unicode-enabled Windows client, connect and backup using SMB. It is recommended that you use Symantec Backup Exec to backup via Unicode-enabled Windows clients, but any Unicode-compliant Backup application should also work.

Backing Up Using Unicode-Enabled UNIX Clients

Most Unicode-enabled UNIX clients run one of four language codes: UTF8, 8859-1 (US), 8859-15 (Europe), or EUC-JP (Japan). In each of these situations, it is important to backup via the UNIX client with a language compliant backup application. Mixing languages (example: having a Japanese UNIX server and a Chinese backup application) will lead to data corruption. If you do not have language compliant backup applications, do not back up using UNIX.

Backing Up Using Unicode-Enabled Mac OS Clients

Mac text encoding UTF8 is supported by Mac OS 10.1.4 AFP 3 and later. For Unicode to function properly, your version of Mac OS must fully support AFP 3.

It is important to back up via the Mac OS client with a language compliant backup application. Mixing languages (example: having Russian files on a server, then using a German backup application) will lead to data corruption.

Unicode and Expansion Arrays

Unicode is enabled on all servers and arrays running GuardianOS 6.5. If a non-Unicode expansion array is attached to a SnapServer running GuardianOS 6.5 or later, the expansion array is automatically converted to Unicode when it is incorporated with the SnapServer. When an expansion array is converted to Unicode, it stays converted to Unicode.

Web Management Interface

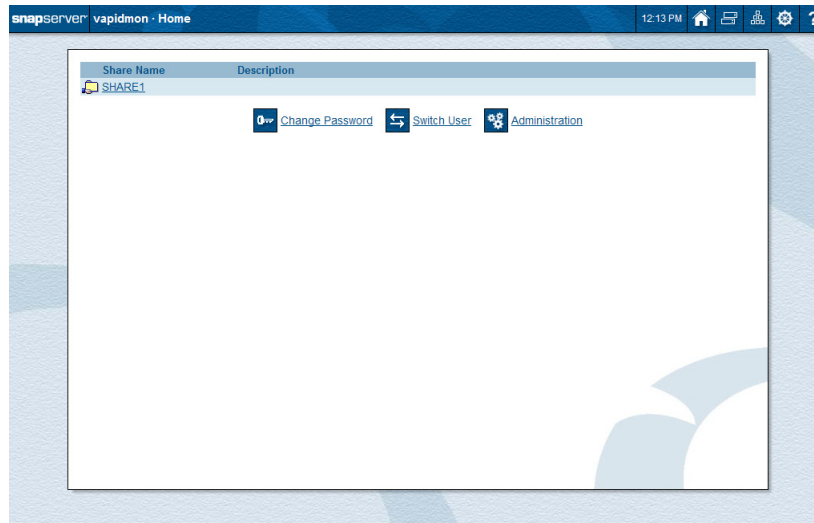
The Web Management Interface has a wide variety of features to monitor and maintain a SnapServer. Information is included here on the following topics:

Topics in Web Management Interface








- [Server Status and Site Map](#)
- [Server Name](#)
- [Date/Time](#)
- [Email Notification](#)
- [SNMP](#)
- [SSH Secure Shell](#)
- [UPS](#)
- [Registering Your Server](#)
- [OS Update](#)
- [Snap Finder](#)
- [Host File Editor](#)
- [Shutdown and Restart](#)
- [Powering SnapServers On and Off](#)
- [System Status](#)
- [Active Users](#)
- [Open Files](#)
- [Using the Event Log](#)
- [Tape Information](#)
- [RAID Settings](#)


Server Status and Site Map

When you first access the Web Management Interface, the home page is displayed. The shared drives are listed at the top, three options are shown below that, and navigation buttons are shown on the right side of the title bar (see the next table).



The home page displays the following:

Option	Description
Change Password 	Click to access the password change page. Passwords are case sensitive. Use up to 15 alphanumeric characters.
Switch User 	Click to log out and open the login dialog box to log in as a different user.
Administration 	Click to assess the status of (or configure) the server.
Navigation Buttons:    	<p>The following Navigation buttons are present in the upper right on every Web Management Interface screen:</p> <p>Home – Click this to return to the home page and site map. If you click this icon while viewing the Home page, you will return to the Web View page.</p> <p>Snap Finder – Click this to view a list of all SnapServers found on your network and to specify a list of remote servers that will be used to discover SnapServers on other subnets. You can access these servers by clicking on the listed IP address.</p> <p>SnapExtensions – Click this to view the SnapExtensions screen, where you can acquire licenses for and configure third party applications.</p> <p>Site Map – Click this to view a Site Map of Web Management Interface, where you can navigate directly to all the major utility pages. The current page is highlighted.</p>

Option	Description
	Help – Click this to access the help files for the UI page you are viewing.
UI Appearance	Click the UI Settings button to choose a solid-colored background for the Web Management Interface by checking the option box. Uncheck the box to show the textured-graphic background.

Server Name

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

Option	Description
Server Name	The default server name is SNAPnnnnnn, where nnnnnn is your server number. For example, the default name for a SnapServer with the serial number 242424 would be SNAP242424. 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. Note: The server number can be found on the Monitor > System Status screen.
Server Comment	Optionally, add a comment (for example, server location) specific to the server. This comment will be displayed in Windows Network Neighborhood.

Date/Time

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

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

To Configure the Date/Time Settings

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

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

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

Click the **Advanced** button for the following additional options:

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

Click **OK** when finished.

Email Notification

To configure the server to send email alerts in response to system events, navigate to the **System > Email Notification** screen. To set up email alerts, you will need: (1) the SMTP server's IP address; and (2) the email address of each recipient to receive an alert.

To Configure Email Notification

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

Option	Description
Enable Email Notification	To enable email notification, check the Enable Email Notification check box.
SMTP Server	Enter a valid SMTP server IP address or host name.
SMTP Port	Enter a port number for the SMTP server or accept the default.
Use Authenticated SMTP	Check this box to require authentication when an email is sent to the SMTP server by the SnapServer. Provide an authentication User Name and Password in the fields that appear when the feature is enabled.
Use Secure Connection	Check this box to encrypt emails from the SnapServer. STARTTLS and TLS/SSL encryption protocols are supported.
Email Address of Sender:	Choose: <ul style="list-style-type: none"> • The default address (<i>servername@domain</i>) where the <i>domain</i> is the DNS domain name. If there is no DNS domain name, then the server's IP address for Eth0 will be used (<i>servername@ipaddress</i>) • Specify a specific sender.
Recipient(s)	Enter one or more email addresses to receive the notifications. One address is required. Three additional email addresses can be added.

Option	Description
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 Migration or Unicode operation has finished or experienced an error. • License event – One of the trial licenses included on the SnapServer is about to expire. A notification email will be sent 14 days before the license expires. One day before the license expires another email will be sent. It is recommended that, if you are not acquiring a license key for the SnapExtension that is expiring, you turn off the SnapExtension.
Send a Test Alert	To verify your settings, check Send a test email , then click OK .

SNMP

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

Default Traps

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

- **coldStart** – Whenever SNMP is enabled and the server boots.
- **linkDown** – An Ethernet interface has gone off-line.
- **linkUp** – An Ethernet interface has come online.
- **authenticationFailure** – An attempt to query the SNMP agent using an incorrect public or private community string was made, and resulted in a failure.

- **enterpriseSpecific** – SnapServer-generated traps that correspond to the error-level, warning-level, and fatal-error-level traps of GuardianOS. These traps contain a descriptive message that helps to diagnose a problem using the following OID's:
 - 1.3.6.1.4.1.6411.2000.1000.1:loglevel 0 syslog messages (“emergency”)
 - 1.3.6.1.4.1.6411.2000.1001.1:loglevel 1 syslog messages (“alert”)
 - 1.3.6.1.4.1.6411.2000.1002.1:loglevel 2 syslog messages (“critical”)
 - 1.3.6.1.4.1.6411.2000.1003.1:loglevel 3 syslog messages (“error”)

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

Supported Network Manager Applications and MIBs

SnapServers respond to requests for information in MIB-II (RFC 1213) and the HostResources 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.

To Configure SNMP

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

Option	Description
Enable SNMP	To enable SNMP, check the Enable SNMP checkbox. Leave the check box blank to disable SNMP.
Public Community	To enable SNMP managers to read data from this server, enter the name of one or more public communities, or accept the default <i>public</i> .
Private Community	To enable SNMP managers to remotely configure this server, enter the name of one or more private communities, or accept the default <i>private</i> . Create unique public and private names. As a precaution against unauthorized access, Overland Storage recommends that you create your own public and private community names.
Server Location	Enter information that helps a user identify the physical location of the server. For example, you might include a street address for a small business, a room location such as <i>Floor 37, Room 308</i> , or a position in a rack, such as <i>rack slot 12</i> .
Contact Person	Enter information that helps a user report problems with the server. For example, you might include the name and title of the system administrator, a telephone number, pager number, or email address.
Enable SNMP Traps	Check the Enable SNMP Traps check box to enable traps. Clear the check box to disable SNMP traps.
IP Address 1-4	Enter the IP address of at least one SNMP manager in the first field as a trap destination. You can enter up to three additional IP addresses.
Send a Test Trap	To verify your settings, check the Send a test trap . . . check box, then click OK .

SSH Secure Shell

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

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

To Disable SSH

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

To connect to the CLI using SSH

1. Make sure your remote machine has an SSH client application installed.

Note: Free or low-cost SSH applications are available over the Internet.

2. Connect to the server using its IP address, and 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.

3. You will automatically be placed in the CLI shell.

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

UPS

Note: If you are not using a UPS and your server supports disabling write cache, consider disabling write cache to help protect your data in case of a power outage. For more information, see [Configuring Write Cache](#).

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

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

To Configure One (Primary) UPS Device

Complete the following fields, and then click **OK**.

Option	Description
Enable UPS Support	Check the Enable UPS Support check box to enable; leave the check box blank to disable UPS support.
Automatically restart server	Check this box to automatically restart the server when power has been restored or the UPS comes back online. Leave the check box blank to manually start the server after a power failure.
Use a single USB-connected UPS device	Select this option button to use a USB-connected APC UPS device or serial UPS with USB to serial adapter cable. Note: If using a serial UPS with a USB-to-serial adapter cable, reboot the SnapServer after connecting the cable to the server to properly initialize the connection to the UPS.
APC Status	Under the selected UPS connection type, an APC status field will display the following possible values: Unknown, No Connection, Low Battery, On Battery, and Online.
Use the following network-connected UPS device(s)	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.

Click **OK** to finish.

If you are configuring a secondary UPS device, continue to the next section.

To Configure a Secondary UPS Device

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

Option	Description
Secondary UPS device (optional)	Check the Secondary UPS device check box to enable; leave the check box blank to disable secondary UPS support.
IP Address	Enter the IP address of the network UPS device.

Option	Description
APC User Name (for authentication)	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.
APC Status	An APC status field will display the following possible values: Unknown, No Connection, Low Battery, On Battery, and Online.
Low Battery Alert	Select one of the following: <ul style="list-style-type: none"> • Either UPS Device: Select this option to allow shutdown upon receipt of a message from either of the two specified UPS servers. • Both UPS Devices: Select this option to allow shutdown only upon receipt of one message from each of the two specified UPS servers.

Registering Your Server

Registering your server activates your warranty and allows you to create and track service requests. Registration also provides access to GuardianOS upgrades, third-party software, and exclusive promotional offers.

Note: Warranty information is available at <http://www.snapserver.com/support>.

To Register Your Server

Note: To use this feature, access to the external Internet is required.

Go to **Server > Registration** and click the “Click here” link to launch the online registration page. The same page is also used to update your registration information.

Once you have registered, you will receive a confirmation email.

OS Update

Use this screen to install updates to GuardianOS and other installed software, and to configure GuardianOS to automatically check for updates to GuardianOS and Snap EDR.

Information about the last GuardianOS update is listed at the bottom of the page, and may include the status of the update, product and version, and the completion time.

CAUTION: Do not interrupt the update process. You may severely damage the server if you interrupt a software update operation.

To Update the GuardianOS Software

1. Click the **Check for Updates** button. If an update is available, follow the onscreen instructions to download it.

Note: If the server does not have access to the Internet, download the latest GuardianOS image or other software package from the [Overland Storage web site](#).

2. Click **Browse** on the OS Update screen, locate the downloaded file, and select it.
3. Click **OK**. The SnapServer uploads the software package and then prompts you to reboot the server to perform the upgrade. Or click **Cancel** to stop the update from beginning.

Update Notification

You can configure GuardianOS to display an alert when GuardianOS or Snap EDR updates are available for the server. When enabled, Update Notification checks weekly for GuardianOS or EDR updates that are applicable to the server. If updates are available, a banner alert will display just below the menu bar on all Web Management Interface pages.

Note: You can choose to hide the banner by clicking the *Remind me later* or *Hide this message* link on the banner. If *Remind me later*, the server will display the banner after the next check for updates; if *Hide this message*, the server will hide the banner for the update in question until a later version is released.

Configuring Update Notification

1. Click the **Update Notification** button.
2. Click to put a check in the **Enable Automatic Update Notification** check box.
3. If your environment requires using a proxy server for external web-based communication, check the **Use a proxy server for HTTP communication** check box and complete the Proxy Host and Proxy Port fields.
4. Click **OK**.

Checking for Updates

Click the **Check for Updates** button to force the server to immediately search for applicable updates. If an update is available, it will be displayed with information about it and a link to download the software.

Snap Finder

Snap Finder displays a list of all the active SnapServers on your network, as well as the status of each. Click the server name (if you have IP address resolution) or IP address of a server to access it through the web based interface.

Note: You can sort the columns (ascending or descending order) by clicking the column head.

Identification	Description
Server Name	Current name of the server. The default server name is SNAPnnnnnn, where nnnnnn is your server number (for example, SNAP242424).
Status	The status of the server (for example, OK, fan failure, or power failure).
IP Address	The IP address of the server.
OS Version	The version of GuardianOS currently loaded on the SnapServer.
Model	The SnapServer model (for example, N2000).
Number	The Server Number derived from the MAC address of the primary Ethernet port, used as part of the default server name.
Avail Cap	The available capacity on the server.
Total Cap	The total capacity on the server.

To enable remote discovery of SnapServers on a different subnet or to display a warning icon for servers with an enabled Ethernet port that has no link, click the **Properties** button to open the Snap Finder Properties page.

Snap Finder Properties

From this screen you can select to display a warning icon for servers with an enabled Ethernet port that has no link and enable remote discovery of SnapServers on a different subnet. Complete the following fields and then click **OK** to return to the Snap Finder screen:

Option	Description
Display warning if any of a server's ethernet ports has no link	Check to display a warning icon in the Status column for any servers that have an enabled Ethernet port with no link. By default, this box is unchecked.
Enable Remote Server Discovery	Check to enable remote discovery of SnapServers on a different subnet.
Add Server	Click the Add button and enter the server's host name or IP Address to add it to the list of remote discovery servers.
Delete Server	Select a server in the Remote Discovery Server column, click the Delete button, and click Delete when asked to confirm the deletion.

Host File Editor

Use this screen to identify backup or media servers in the SnapServer's hosts file. This screen allows you to supply a hostname-to-ip address mapping that persists across system reboots.

Click **Add Host File Entry**, complete the fields as described on the following table, and then click **Add Host File Entry**.

Option	Description
IP Address	The IP address of the backup server.
Host Name	Enter at least one of the following: <ul style="list-style-type: none"> In the Host Name field, enter the fully qualified address for the backup server, using the format: <i>myserver.mydomain.com</i>. In the Alias field, enter an abbreviated address for the backup server, using the format: <i>myserver</i> (optional). <p>Note: Your backup software may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.</p>

Shutdown and Restart

Use the Maintenance > Shutdown/Restart screen to reboot or shut down the server. Click one of the following button icons:

- **Shutdown** – Shuts down and powers off the server.
- **Restart** – Reboots the server.

Powering SnapServers On and Off

Use the power button on the front of the server to power on and power off the server.

To turn on the server, press the power button on the front of the server.

The server takes a few minutes to initialize. A green power LED indicates that the system is on.

To turn off the server, press and release the power button to begin the shutdown process. Do not depress this button for more than four seconds.

Notes: SnapServers have a persistent power state. When a physical loss of power occurs, the SnapServer returns to the same operation it had when the power went out. Therefore, if the system is powered down prior to a power loss, it will remain powered down when the power is restored.

System Status

Use this screen to assess the status of the SnapServer and any attached expansion arrays.

Host SnapServer Status

Field	Description
Server Name	Current name of the server. The default server name is SNAPnnnnnn, where nnnnnn is your server number (for example, SNAP112358).
Server Model	Server model number
OS Version	The version of GuardianOS currently loaded on the SnapServer.
Server Number	Number derived from the MAC address of the primary Ethernet port, used as part of the default server name.

Field	Description
CPU	Details on the server's central processing unit.
Memory	Amount of system RAM.
Ethernet Status	Details on the server's ethernet connections.
Uptime	Length of time since last reboot.
Ambient Temperature	The temperature of the space around the SnapServer.
CPU Temperature	Current CPU temperature.
Power Supply	The status of power supply module(s)
Fan Status	The status of fan module(s).

Expansion Array Status

Field	Description
Expansion Unit	EXTN1, EXTN2, etc.
Expansion Model	Snap Expansion S50, etc.
Serial Number	The serial number of the expansion array
Ambient Temperature	The temperature of the space around the expansion array.
Power Supply	The status of the power supply
Fan Status	The status of fan modules.

Active Users

Use this screen to view read-only details on the active users logged on to the server. Information available on this screen 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 screen to view read-only details on the open files in use on this server.

Using the Event Log

Use the **Monitor > Event Log** screen to view a log of operations performed on the server. Entries are color coded according to severity as described in the following table:

Color	Entry Type
Yellow	Warning
Red	Errors
(no color)	Informational or Unclassified

To Filter the Log

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

Option	Description
Severity	Select the type of entries you want to view.
Display Last	Enter the number of days' entries (24-hour periods) you want to view
Most Recent First	Select to start the list with the most recent entry, deselect to start the list with the oldest entry.

To Erase All Log Entries

Click **Clear Log** to erase all log entries.

Tape Information

Use this screen to view read-only details on the SCSI and USB tape devices attached to the server. Information presented on this screen includes:

Field	Description
Device Model	The manufacturer's model for the device.
Device Type	Type of tape device: either Sequential-Access (tape drive) or Medium-Changer (for example, robotic arm for a tape library).
Device Name	Name of the device node to which the device is bound.
Connection	Identifies the connection type: SCSI or USB.
Bus	Bus number indicating which physical interface (for example, SCSI card) the device is connected to.
ID	ID number (SCSI only)
LUN	LUN identifier (SCSI only)

RAID Settings

Use this screen to enable or disable the background disk scan and to incorporate unassigned disks.

Note: RAID Settings functions do not apply to SnapServers with fewer than four (4) drives.

Automatic Incorporation of Hot-Swapped Drives

If a RAID 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.

Note: Hot-swapped drives cannot be automatically incorporated with a RAID 0.

Background Disk Scan

The background disk scan checks the integrity of the RAID's 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 the I/O activity falls to a very low disk activity. Once the activity rises above the *idle threshold*, the background scan stops and waits for the activity to fall to the idle threshold again before resuming. As a result, there should be minimal to no impact on performance. Once the disk scan has

completed a pass on a given RAID set, it waits a certain period of time before starting again.

Notes:

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

Backup and Replication Solutions

This appendix provides a brief description of the supported backup solutions and, where applicable, gives instructions on how to install the solutions on the SnapServer. GuardianOS supports several backup methods, including third-party off-the-shelf backup applications and applications that have been customized and integrated with GuardianOS on the SnapServer.

Note: Unicode limits some backup applications' ability to function with the SnapServer. Refer to "Unicode" on page 10-1 for more information.

GuardianOS	Backup and Replication Solutions					
	BakBone NetVault	Snap EDR	CA BrightStor	EMC NetWorker	Symantec Backup Exec	Symantec NetBackup
Snap to Backup Server via installed agent	X ¹		X	X	X	X
Snap to Backup Server via network protocol			X		X	
SnapServers to SCSI-attached tape drive ² (disk-to-tape backup)	X				X	X
SnapServers to SAS-attached tape drive (disk-to-disk backup, N2000 only)					X	X
SnapServers to USB-attached tape drive (disk-to-tape backup)	X				X	X
Security Meta Data Backup	X	X				

1. The NetVault agent was preinstalled on SnapServers running GuardianOS 3.0 through 5.2.
 2. SCSI not applicable to the SnapServers 110 and 210. They support USB only.

Topics in Backup and Replication Solutions

- [Integrated Backup Solutions for the SnapServer](#)
- [Off-the-Shelf Backup Solutions for the SnapServer](#)
- [Backup of iSCSI Disks](#)

Integrated Backup Solutions for the SnapServer

The following backup solutions are preinstalled and/or customized for the SnapServer:

- BakBone NetVault (on GuardianOS 5.2 or earlier)
- Snap Enterprise Data Replicator (Snap EDR)

Note: Depending on the SnapServer platform, the above applications may require a license to activate.

BakBone NetVault

Note: NetVault is only available on SnapServers upgraded from GuardianOS 5.2.

BakBone NetVault is a scalable, enterprise-wide backup solution for GuardianOS, Windows, Linux, and UNIX operating systems with the following functionality:

- **Near-line storage** – The SnapServer manages backup jobs, locally storing the backup images on disk using virtual tape library technology. Eight virtual drives with up to 100 GB capacity (total) are supported. (Additional capacity up to 1.1 TB can be added with additional licenses.)
- **Direct-attached storage** – Data from up to five clients is backed up to a standalone, SCSI-attached tape device attached to a SnapServer. Supports up to four tape drives.

Some earlier SnapServers models shipped with BakBone's NetVault server software preinstalled with a Workgroup Server license. This license supports backup and recovery of data to the SnapServer from up to 5 heterogeneous clients. (Additional clients can be added with optional licenses.)

Note: Some SnapServers require an additional license to support NetVault.

For additional information on installing and configuring NetVault, see the documentation included with the NetVault CD.

To Enable NetVault for GuardianOS

To enable NetVault, click the **Bakbone NetVault** link on the SnapExtensions page, check the **Enable** box, and click **OK**.

Adding Clients to the NetVault Management GUI

The **Add Clients** button allows you to specify the name or IP address of the workstation on which the NetVault management GUI has been installed.

1. Click **Add Clients** on the BakBone NetVault page.
2. Enter the management workstation's DNS name or IP address.
3. Enter the NetVault client password created during installation for the client. The password field cannot be left blank.
4. Click **Add**.

You can add multiple workstations by completing the fields and repeatedly selecting the **Add** button. This should only be done when management of your NetVault implementation **MUST** be managed from more than one workstation. Each client added in this way consumes a NetVault node license. Other NetVault Clients can be added using the NetVault Management GUI.

Snap Enterprise Data Replicator (Snap EDR)

Snap EDR provides server-to-server synchronization by moving, copying, or replicating the contents of a share from one SnapServer to another share on one or more different SnapServers. It comes preinstalled on some servers with a 45-day free trial, or it can be downloaded from the [SnapServer website](#).

Snap EDR consists of a Management Console and a collection of Agents. The Management Console is installed on a central system. It coordinates and logs the following data transfer activities carried out by the distributed Agents:

- Replicates files between any two systems including Windows, Linux, and Mac Agents.
- Transfers files from one source host to one or more target hosts
- Transfers files from multiple hosts to a single target host, and stores the files on a local disk or locally attached storage device.
- Backs up data from remote hosts to a central host with locally-attached storage.
- Restores data from a central storage location to the remote hosts from which the data was originally retrieved.

Configuring Snap EDR for GuardianOS

To configure the SnapServer as a Snap EDR Management Console or an Agent, do the following:

1. Click the **Snap EDR** link in the Site Map (under **Extras**).
2. Select either the **Configure as the Management Console** or **Configure as the Agent** button.

Note: If you are configuring the server as an Agent, you must provide the server name or IP Address of the Management Console.

3. Once the server is configured, a screen appears with the following options:

Option	Description
Click here to configure jobs	Opens the Management Console where jobs can be scheduled.
Stop Service	Stops all services.
Restart Service	Restarts all services. Caution: Use only if you have encountered a problem, and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service
Disable Service on System Boot	By default, when a user reboots the SnapServer, services automatically restart. Select Disable Service on System Boot if you do not want the Snap EDR service to start up automatically. Note: When the disable service option is selected, the Enable Service on System Boot button appears.
Uninstall Service	Uninstalls all components of Snap EDR.

Scheduling Jobs in Snap EDR

To schedule jobs, click the **Snap EDR** link in the Site Map (under **Extras**). For complete information on using Snap EDR, see the *Snap EDR Administrator's Guide*, available on the [SnapServer website](#).

Off-the-Shelf Backup Solutions for the SnapServer

Note: These backup packages do not support the backup of Windows ACLs or legacy POSIX ACLs. If you use one of these packages, Overland Storage strongly recommends you create a SnapServer disaster recovery image (see “[Creating the SnapDRImage and Volume Files](#)” on page 8-2) before you perform a backup.

In addition to the integrated backup solutions, GuardianOS supports a number of off-the-shelf backup packages that the user can install on the SnapServer, including:

- CA ARCserve 11.5, 12.0
- EMC NetWorker 7.3, 7.4
- Symantec Backup Exec 11d, 12, 12.5
- Symantec NetBackup 6.5

Agent Installation Procedures:

- [Preparing to Install a Backup Solution](#)
- [Preinstallation Tasks](#)
- [Installing the CA ARCserve Agent](#)
- [Installing the Symantec Backup Exec RALUS Agent](#)
- [Installing the Symantec NetBackup 6.5 Client](#)
- [Installing the EMC NetWorker Client](#)

Preparing to Install a Backup Solution

Before performing one of the backup solution installation procedures described here, make sure you have the following information and tools:

- **Backup and media server IP addresses** – Most backup agents need to know the IP addresses of the backup and media servers you plan to use with the SnapServer. Use the **Maintenance > Host File Editor** screen in the SnapServer's Web Management Interface to supply a host-name-to-ip-address mapping that persists across system reboots.
- **SnapServer is seen by Backup software as a UNIX/Linux client** – When you configure a backup server to see the agent or client running on the SnapServer, assume the server is a UNIX or Linux client.
- **The agent/client files required by your backup software** – Typically, these files are either provided on your backup software's User CD or are available for download from the manufacturer's website. You will need to copy these files (usually delivered in a compressed format, for example, as **.rpm*, **.tgz*, or **tar* files) to the SnapServer.
- **A secure shell (SSH) client** – To remotely install any backup solution on the SnapServer, you must have an SSH client installed on a remote workstation. The SnapServer SSH implementation requires SSH v2. If you do not already have an SSH client application installed, you can download one from the Internet.

Note: The commands you must enter via SSH to install your backup agent are case sensitive; pay careful attention to the capitalization of commands, and enter them exactly as shown.

- **Location of the SnapServer backup and restore path** – Backup servers often request the path for backup and restore operations on the SnapServer. When you configure a backup server to see the agent or client running on the SnapServer, use the following path:

`/shares/sharename`

where *sharename* is the name of the share to be backed up. If you have accepted the default SnapServer configuration, the correct path is as follows:

`/shares/SHARE1`

- **Backup user account is configured to be exempt from password policies (if applicable)** – If the backup application uses a specific local SnapServer user account to perform backups, ensure that the user is exempt from password expiration policies, if enabled (see the Online Help for procedures to set password policy for local users).

Preinstallation Tasks

Perform the following tasks prior to installing any solution:

1 Identify backup and media servers to the SnapServer.

In the Web Management Interface, navigate to the **Maintenance > Host File Editor** screen and click **Add**. In the screen that opens, enter the IP address of the backup or media server; or enter one or both of the following as required by your backup software:

- **Host name (long form)** Enter the fully qualified address for the backup server using the *myserver.mydomain.com* format.
- **Host name (short form)** Enter an abbreviated address for the backup server using the *myserver* format.

Click **OK**. The entry appears on the Host Editor screen. Repeat this procedure for each backup and media server you plan to use.

2 Make sure SSH is Enabled on the SnapServer.

Navigate to the **Server > SSH** screen, make sure the **Enable SSH** box is checked, and then click **OK**. SSH is immediately available.

CAUTION: To maintain security, we recommend disabling SSH when it is not in use.

3 On a client computer connected to the SnapServer, create a directory called *agent*.

You must create a directory to which you will copy the agent files. Create this directory on a client computer connected to the SnapServer. For purposes of illustration, the procedures described here assume that this directory is called *agent*.

4 Copy the agent/backup files to the *agent* directory.

Using a method appropriate to your environment, copy the agent/client files to the directory you just created for this purpose.

Installing the CA ARCserve Agent

This section explains how to install the CA ARCserve Agent versions 11.5 and 12.0.

Notes:

- This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called *agent* (to which to copy your agent/client files) on the default share (SHARE1), such that the path to the directory is */shares/SHARE1/agent*.
- Installing the ARCserve backup agent on a SnapServer requires three agent (**.rpm*) files. These agent files are available from your ARCserve CD, but some ARCserve CDs may not contain all the required files. To obtain the files you need, contact Computer Associates. If you have questions about the agent configuration, refer to your CA ARCserve documentation.

Prepare the SnapServer

1. Connect to the SnapServer via SSH.

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

2. At the prompt, log in as admin, using the password you created for this account during the initial setup of the server.
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

4. To change to superuser, enter the following command, and press Enter:

```
su -
```

5. At the prompt, enter the admin user password, and press Enter.

6. To change to the agent directory, type the following command and press Enter:

```
cd /shares/SHARE1/agent
```

7. To unpack the tar file to get the agent files, type the following command and press Enter:

```
tar -zxvf Linux.tar.Z
```

Note: If you later delete the volume this directory is on, you will need to reinstall the agent.

8. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table.

```
cd /hd
```

```
ls (lists all volumes)
```

```
df -h (shows volume usage)
```

9. Change directory to the volume with the most available space.

```
cd [volumename]
```

where **[volumename]** is the volume with the most available space

10. Create a directory **arcserve** on that volume.

11. Create the following symbolic links from the new directories in **arcserve** to the **/opt** directory:

```
ln -s /hd[volumename]/arcserve /opt/CA
```

Install CA ARCserve Agent

1. To install the agent files, enter the following command at the prompt, and press Enter:

```
rpm --nodeps -Uvh babagtux.rpm *lic*.rpm
```

2. Once the license is installed, run the Install script by entering the following command at the prompt and pressing Enter:

```
./install
```

Answer the prompts using the defaults.

Note: You are installing the Linux Client Agent.

3. To change to the agent directory, enter the following command, and press Enter:

```
cd /opt/CA/BABuagent/
```

4. To run the setup program, enter the following command, and press Enter:

```
./uagentsetup
```

The ARCserve agent is now installed.

5. Enter the following command to run the script that will edit the agent.cfg file:

```
fix-arcserv
```

6. Close the SSH client and return to the Web Management Interface. To start the newly installed backup agent, navigate to the **Maintenance > Shutdown/Restart** screen, and click **Restart**.
7. Delete the agent files you copied to the SnapServer because they are no longer needed.
8. To verify the success of the installation, use your backup management software to configure and run a test backup.

Uninstall ARCserve Agent

1. If you still have the tar or install directory that you copied to the SnapServer when you installed the ARCserve Agent, the uninstall script will be in that directory. If you do not have the directory or tar, copy the files again from the ARCserve CD or get them from Computer Associates.
2. Make sure you have the script **uninstall**. Type the following and follow the prompts:

```
./uninstall
```

Note: Choose Option 1 to uninstall.

3. Uninstall the license **rpm** by typing the following:

```
rpm -e ca-lic
```

4. Verify that ARCserve Agent has been uninstalled by typing the following and verifying that you do not see the agents:

```
rpm -qa | grep BAB
```

Installing the Symantec Backup Exec RALUS Agent

To install the Backup Exec RALUS agent, do the following:

Prepare the SnapServer

1. Connect to the server over SSH.

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

2. Log in as admin (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

4. Change to root by entering the following command:

```
su -
```

Give the root password (same as admin password).

5. Select a volume on which to put a directory called *ralus*.

Note: If you later delete the volume the *ralus* directory is on, you will need to reinstall the agent.

```
cd /hd
```

```
ls [lists all the volumes]
```

```
df -h [shows volume usage]
```

6. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory to the volume with the most available space.

```
cd [volumename]
```

where *[volumename]* is the name of the volume with the most available space.

7. Create a directory *ralus* on that volume:

```
mkdir ralus
```

8. In the *ralus* directory, create 3 directories called *VRTS*, *VRTSralus*, and *VRTSvxms*.

```
cd ralus
```

```
mkdir VRTS VRTSralus VRTSvxms
```

```
ls [to verify that the directories are there]
```

9. If CA Antivirus has been installed, you will have an */opt* directory. If it has not been installed, create an */opt* directory:

```
mkdir /opt
```

10. Create the following symbolic links from the new directories in *ralus* to the */opt* directory:

```
ln -s /hd/[volumename]/ralus/VRTS /opt
```

```
ln -s /hd/[volumename]/ralus/VRTSralus /opt
```

```
ln -s /hd/[volumename]/ralus/VRTSvxms /opt
```

where *[volumename]* is the name of the volume with the most available space.

11. Use the host file editor (**Maintenance > Host File Editor** screen) to add all the Backup Exec servers to **/etc/hosts** on the SnapServer, and verify that the agent server can ping the main Backup Exec server.

Note: Do not edit the **/etc/hosts** file directly with a text editor.

Install Backup Exec RALUS Agent

1. From a network client, create a *ralusinstall* directory on SHARE1 of the SnapServer, then copy the RALUS agent tar file or contents of the RALUS agent CD to the directory.
2. If you copied the files from the CD, proceed to Step 3.. If you downloaded the files from the Symantec website, in SSH, extract the files:

```
cd /shares/SHARE1/ralusinstall  
tar -zxvf [filename].tar.gz
```

where *[filename]* is the name of the Backup Exec tar file.

3. Install the agent:

```
cd /shares/SHARE1/ralusinstall
```

(or other directory containing the CD contents)

```
./installralus
```

Follow the installation instructions, accepting the default options.

Note: During the installation process, you may see an error message about the failure to add root to the *beoper* group. This error will be resolved in the following step.

4. Add the user *root* to the group *beoper* manually (or any other local SnapServer user you wish to use to perform backups):

```
cli group member add group-name=beoper user-name=root
```

Note: If using a local SnapServer user account other than *root* or *admin*, and if password policies are enabled, configure the user to be exempt from password expiration. See “Users and Groups” on page 3-1.

5. Start the Backup Exec RALUS agent by rebooting the SnapServer either through the Web Management Interface (**Maintenance > Restart**), or by typing:

```
/etc/rc.d/init.d/VRTSralus.init start
```

6. Verify that using Backup Exec, you can create a job using the UNIX agent:

- a. Create a GuardianOS Root login account on the Backup Exec server.
- b. Connect as *root* (the password will be the same as the admin account password).
- c. Create a job and choose the Unix agent representing the SnapServer.
- d. Verify that you can connect to the agent, configure a job, and run the job.

Uninstall the Backup Exec RALUS AGENT

1. To uninstall the RALUS Agent, you will need the tar or install directory that you copied to the SnapServer when you installed the Agent (follow Steps 1 through 3 of “Install Backup Exec RALUS Agent”). Make sure you see the script

```
uninstallralus
```

2. Type:

```
./uninstallralus
```

Follow the prompts.

3. Verify that the Symantec RALUS agent has been uninstalled by typing the following command:

```
rpm -qa | grep VRTS
```

Installing the Symantec NetBackup 6.5 Client

Note: This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called *agent* (to which to copy your agent/client files) on the default share (SHARE1), such that the path to the directory is */shares/SHARE1/agent*.

To install the Symantec NetBackup 6.5 Client, do the following:

Prepare the SnapServer

1. Connect to the server over SSH.

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

2. Log in as admin (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

4. Change to root by entering the following command:

```
su -
```

Give the root password (same as admin password).

5. Select a volume on which to put a directory called *openv*.

Note: If you later delete the volume the *openv* directory is on, you will need to reinstall the agent.

```
cd /hd
```

```
ls [lists all the volumes]
```

```
df -h [shows volume usage]
```

6. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory to the volume with the most available space.

```
cd [volumename]
```

where *[volumename]* is the name of the volume with the most available space.

```
ls [lists what is on that volume]
```

7. Create a directory called *openv* on that volume:

```
mkdir openv
```

8. Create a “symbolic” link to the *openv* directory in the */usr/* directory:

```
ln -s hd/[volumename]/openv /usr/
```

where *[volumename]* is the name of the volume with the most available space.

- Use the host file editor (**Maintenance > Host File Editor** screen) to add the NetBackup servers to **/etc/hosts** on the SnapServer. Verify that you can ping the NetBackup server.

Install NetBackup v6.5 Client

- Using a network client, copy the directory called **NBClients** from the Client CD to a directory on a share (for example, **SHARE1** or **Agent**) on the SnapServer.

- In SSH, install the files:

```
cd /shares/SHARE1/NBClients/catalog/anb
./client.inst
```

Follow the instructions, choosing RedHat Linux (choose 2.6 kernel version, if available) as the type.

- Once the NetBackup Client is installed, reboot the server using the Web Management Interface (**Maintenance > Restart**) to start the client service.
- Verify that you can configure the UNIX client:
 - Create a policy and add the SnapServer as a client.
 - Look at the client list to verify that the SnapServer client is listed.

Uninstall the NetBackup v6.5 Client

- Log in to the client system as the root user.
- Navigate to the volume where you installed the NetBackup directory.

```
cd /hd/vol_mnt[X] /
rm -rf /usr/opensv/
rmdir opensv/
```

- Remove the NetBackup entries in the client's **/etc/services** file.

Locate the lines, marked by the following strings and delete them:

```
# NetBackup services#.....# End NetBackup services #
```

- Remove the NetBackup services by deleting the files for **bpcd**, **vnetd**, **vopied**, and **bpjava-msvc** in the **/etc/xinetd.d/** directory.

```
rm -rf /etc/xinetd.d/bpcd
rm -rf /etc/xinetd.d/vnetd
rm -rf /etc/xinetd.d/vopied
rm -rf /etc/xinetd.d/bpjava-msc
```

- Restart the SnapServer services by either rebooting or typing:

```
/etc/rc.d/init.d/xinetd reload
```

Installing the EMC NetWorker Client

Note: This procedure assumes that you are using the default SnapServer configuration; and you have created a directory called *agent* (to which to copy your agent/client files) on the default share (**SHARE1**), such that the path to the directory is **/shares/SHARE1/agent**.

This section describes how to install the EMC NetWorker UNIX/Linux client, as well as special procedures EMC NetWorker users must follow in order to perform backup and restore operations on the SnapServer.

Prepare the SnapServer

1. Connect to the server over SSH.

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

2. Log in as admin (using the password for the admin account).
3. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

4. Change to root by entering the following command:

```
su -
```

Give the root password (same as admin password).

5. Select a volume on which to put a directory called *networker*.

Note: If you later delete the volume the *networker* directory is on, you will need to reinstall the agent.

```
cd /hd
```

```
ls [lists all the volumes]
```

```
df -h [shows volume usage]
```

6. Determine which volume has the most available space by looking at the **Avail** column in the volume usage table. Change directory to the volume with the most available space.

```
cd [volumename]
```

where *[volumename]* is the name of the volume with the most available space.

7. Create a directory *networker* on that volume:

```
mkdir networker
```

8. In the *networker* directory, create the following directories called *opt*, *usr*, and *opt/usr*.

```
cd networker
```

```
mkdir opt usr opt/usr
```

```
ls [to verify that the directories are there]
```

9. If CA Antivirus has been installed, you will have an */opt* directory. If it has not been installed, create an */opt* directory:

```
mkdir /opt
```

10. Create links from the *networker* working volume to the root filesystem:

```
ln -s /hd/vol_mnt[X]/networker/nsr/
```

```
ln -s /hd/vol_mnt[X]/networker/opt/nsr /opt/
```

```
ln -s /hd/vol_mnt[X]/networker/usr /usr/
```

where *vol_mnt[X]* is the NetWorker installation target volume.

11. Modify the SnapServer environment by editing `/etc/profile` as follows:

```
cp /etc/profile /etc/profile.nwbk
echo PATH=$PATH:/hd/vol_mnt[X]/networker/usr/bin:/hd/vol_mnt[X]/
networker/usr/sbin:/hd/vol_mnt[X]/networker/usr/lib >> /etc/profile
```

where `vol_mnt[X]` is the NetWorker installation target volume.

- Note:** Be sure to enter `>>` in the command rather than `>` or you will overwrite the file rather than append to the `/etc/profile` script. If you need to redo Step 11, copy the backup to the original using the command `cp /etc/profile.nwbk /etc/profile` and then edit the file again.

12. To implement the changes, enter the following command:

```
source /etc/profile
```

Install the EMC NetWorker Client

1. Connect to the SnapServer via SSH, and log in as admin, using your admin user password.

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

2. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

3. To change to superuser, enter the following command, and press Enter:

```
su -
```

4. At the prompt, enter the admin user password, and press Enter.

5. Use the `cd` command to change to the directory in the share, for example:

```
cd /shares/SHARE1/agent
```

6. To unpackage the client files, enter the following commands:

```
tar xvfz nw_linux86.tar.gz
```

7. To install the NetWorker Agent rpm, enter the following command:

```
rpm -Uvh --nodeps --relocate=/usr=/hd/vol_mnt[X]/NetWorker/usr/
lgtoclnr-x.x-x.1686.rpm
```

where `vol_mnt[X]` is the NetWorker installation target volume and `x.x-x` is the version number.

8. To start the EMC NetWorker daemon, enter the following command at the console:

```
/etc/rc.d/init.d/networker start
```

The NetWorker client is now installed.

9. Close the SSH client, return to the Web Management Interface. To start the newly installed backup agent, navigate to the **Maintenance > Shutdown/Restart** screen, and click **Restart**.
10. Delete the client files you copied to the SnapServer because they are no longer needed.

11. To verify the success of the installation, use your backup management software to configure and run a test backup.

Backup and Restore Operations with the EMC NetWorker Client

This section describes special procedures EMC NetWorker users must use in order to perform backup and restore operations on the SnapServer.

Add the SnapServer as a Root User

For backup operations, NetWorker requires that the SnapServer be configured as a root user. To add the SnapServer root user as one of the administrators, use the following procedure:

1. Open the NetWorker Administrator application.
2. Click the **Configuration** tab.
3. Click the **User Groups** menu item.
4. Click on the **Administrators** group.
5. In the Configuration box, add one of the following:

```
user=root@hostname
```

where *hostname* is the host name of the SnapServer for each SnapServer.

Or, enter:

```
user=root
```

to add root for all SnapServers.

6. Click **OK**.

Recover and Retrieve Operations

The EMC NetWorker administrative interface does not support data recovery operations from a remote client for a Linux-based operating system such as GuardianOS. To recover data, you must execute one of the following CLI commands from a SSH client.

- **Recover** – The **recover** command restores data from a normal backup job.
- **Nsrretrieve** – The **retrieve** command restores data from an archive.

Use either the **recover** or the **retrieve** command exactly as described below. For more details on these commands, see the *EMC NetWorker Command Reference*.

1. Connect to the SnapServer via SSH, and log in using the admin user name and password.

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

2. You are placed into the CLI shell. However, a standard Linux shell must be used to install the agent. To launch a shell, type the following command and press Enter:

```
osshell
```

3. To change to superuser, enter the following command, and press Enter:

```
su -
```

4. At the prompt, enter the admin user password, and press Enter.

5. **To recover data from a normal backup operation**, enter one of the following commands, and press Enter:
 - To recover data to its original location:


```
recover -s backupservername -c snapservername -f -i "/shares/SHARE1/data/" -a
```

 where `/shares/SHARE1/data` is the path of the data you are restoring.
 - To recover data to a different location:


```
recover -s backupservername -c snapservername -f -i -a R -d "/shares/SHARE1/relocated_data/" "/shares/SHARE1/Data/"
```

 where `/shares/SHARE1/relocated_data/` is the path to the new target location for the restore operation; and where `/shares/SHARE1/Data/` is the path of the data you are restoring.

6. **To retrieve data from an archival backup operation**, enter one of the following commands, and press Enter:
 - To retrieve data to its original location:


```
nsrretrieve -f -i -s backupservername -A annotation "/shares/SHARE1/data/"
```

 where `/shares/SHARE1/data/` is the path of the data you are restoring.
 - To retrieve data to different location:


```
nsrretrieve -f -iR -d "/shares/SHARE1/new_dir" -s backupservername -A "annotation" "/shares/SHARE1/Data/"
```

 where `/shares/SHARE1/new_dir` is the path to the new target location for the restore operation; where `annotation` is the name of the EMC NetWorker backup; and `/shares/SHARE1/Data/` is the path of the data you are restoring.

Backup of iSCSI Disks

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 5-1 for instructions.

Using Backup Exec to Take 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.

Command Line Interface

GuardianOS includes a command line interface (SnapCLI) accessible through SSH. Using the CLI, users can access information about most of the SnapServer configuration parameters and perform configuration and maintenance functions without using the GuardianOS Web Management Interface or SSM.

Note: Some administrative tasks must still be performed using the Web Management Interface. The CLI is intended as a convenient way to perform some functions; it is not intended as an alternative to using the Web Management Interface.

Topics in Command Line Interface

- [SnapCLI Syntax](#)
- [SnapCLI Commands](#)
- [Scripts in SnapCLI](#)

SnapCLI Syntax

SnapCLI command syntax uses three parameters: COMMANDS, ARGUMENTS, and OPTIONS. To generate commands in SnapCLI, use the following syntax:

COMMAND [**ARGUMENT**] [**OPTIONS**]

where **COMMAND** is the name of one of the SnapCLI commands, **ARGUMENT** is an action available for that command, and **OPTIONS** are additional parameters for the command.

Once logged into the CLI, there are several ways of displaying information about available parameters.

Type	To
?	see an overview of the CLI, with a list of available commands and a description of command syntax.
<i>{command}</i> help	see a description of that particular command's function and a list of options available for the command.
tab	finish the command you have started to type (such as, tab-complete).
<i>{command}</i> tab	list any arguments and/or options available for that command.

For example, to see a list of available commands once you have logged into SnapCLI, type “?” at the prompt.

To see a description of a specific command, type the command name (for example, `date`) + “help” or “?”:

Command	Arguments and Options	Descriptions
date	timezones	- list available time zones
	get	- get server date/time
	set [OPTIONS]	- set server date/time
	- [day=1-31]	- day of month
	- [month=1-12]	- month of year
	- [year=1900-current]	- year
	- [hour=0-23]	- hour
	- [minute=0-59]	- minutes
	- [second=0-59]	- seconds
	- [timezone=1- 40]	- timezone (use the command date timezones to get a list of timezones)

In this instance, to set the date to February 27, 2011, enter:

```
date set day=27 month=2 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 entered the command `web`. Two arguments would be available, one with options:

Command	Arguments and Options	Descriptions
web	get	- get WEB properties
	set [OPTIONS]	- set WEB properties
	- require-webview-auth =(yes no)	- require HTTP/HTTPS clients to authenticate in order to access the server
	- non-secure-http =(yes no)	-enable/disable non-secure HTTP access

Thus, the following command string:

```
web set require-webview-auth=yes non-secure-http=no
```

sets HTTP/HTTPS properties on the SnapServer to require clients to authenticate in order to access the server and to disable non-secure HTTP access.

Procedures

Logging into SnapCLI

1. Make sure your client has an SSH v2 client application installed. For more information about SSH, see “[SSH Secure Shell](#)” on page 11-7.

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 *admin*gp).

You will automatically be placed in the CLI shell.

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

Exiting SnapCLI

To exit SnapCLI, type **exit**. The SSH session will close.

SnapCLI Commands

The following table presents a list of the available SnapCLI commands and a brief description of the function of each.

Command	Description
activeusers	Display active users
apple get	Display apple network settings
apple set	Update apple network settings
date get	Get the current date/timezone information
date set	Set the current date/timezone information
date timezones	List the available timezones (used in conjunction with the date set command)
diskunits	Get status information of all the disk units on the server
domain get	Get the domains known to the SnapServer and their properties
domain list	List the domains known to the SnapServer
dri create	Create a Disaster Recovery Image (dri)
dri recover system	Restore a Disaster Recovery Image (dri)
dri recover volume	Restore a Disaster Recovery Volume Image (dri)
email get	Get email notification settings
email set	Set email notification settings
event clear	Clear all events in the System Event Log
event get	Display the System Event Log
factorydefaults	Reset the SnapServer's settings back to the factory defaults, will reboot
fscheck	Check or repair the user or root file system
ftp get	Get the current ftp settings, including anonymous user access
ftp set	Set the current ftp settings, including anonymous user access
globalspares list	List global hot spares
globalspares remove	Remove a disk from the global spares list
globalspares add	Add a disk to the global spares list
group create	Create a local group
group delete	Delete a local group
group get	Get available groups with their associated information

Command	Description
group list	List available groups
group set	Change the properties of a local group
group member add	Add a group member to a local group
group member delete	Delete a group member from a local group
group members get	Get a list of the members of a local group
group members list	List the members of a local group
homedirs get	Get Home Directory configuration information
homedirs set	Set Home Directory configuration information
hostfile add	Add a host file entry
hostfile delete	Delete a host file entry
hostfile get	Get information for a specific host file entry
hostfile set	Set information for a specific host file entry
hostfile list	List all host file entries
idmap auto map	View/Save auto-generated ID mappings
idmap count	Count number of ID mappings
idmap group get	Get ID mapping for a (windows domain) group
idmap group remove	Remove ID mapping for a (windows domain) group
idmap group set	Set ID mapping for a (windows domain) group to a local or NIS group
idmap list	List all ID mappings
idmap remove all	Remove all ID mappings
idmap update files	Update file system for ID mapping changes
idmap update status	View status of ID mapping update file system operation
idmap user get	Get ID mapping for a (windows domain) user
idmap user remove	Remove ID mapping for a (windows domain) user
idmap user set	Set ID mapping for a (windows domain) user to a local or NIS user
iscsi create	Create an iscsi disk
iscsi delete	Delete an iscsi disk
iscsi get	Get iscsi disk properties
iscsi set	Set iscsi disk properties
isns get	Get configuration settings for iSNS server
isns set	Set configuration settings for iSNS server
jumboframe get	Get jumbo frame settings for all interfaces
jumboframe list	List jumbo frame settings for all interfaces
jumboframe set	Set jumbo frame settings for all interfaces
name get	Get the name of the SnapServer
name set	Set the name of the SnapServer
netinfo	Get information about the Ethernet interface
nfs get	Get SnapServer NFS Properties

Command	Description
nfs set	Set SnapServer NFS Properties
nis get	Get current NIS settings
nis set	Set current NIS settings
ntp get	Get NTP client settings
ntp set	Set NTP client settings
ntp_server get	Get NTP Server settings
ntp_server set	Set NTP Server settings
openfiles	List the Open Files
osupdate get	Display status of last OS update
osupdate load	Perform an OS update
passwordpolicy get	Display Password Policy settings and status
passwordpolicy set	Update Password Policy settings
phonehome	Send configuration details to SnapServer Technical Support
proxy get	Display the HTTP proxy properties
proxy set	Set the HTTP proxy properties
quota list	List user or group quotas for a volume
quota get	Get quota settings for a volume
quota set	Set quota settings for a volume
quota group get	Get volume quota limit & usage for a specific group
quota group set	Set volume quota limit & usage for a specific group
quota user get	Get volume quota limit & usage for a specific user
quota user set	Set volume quota limit & usage for a specific user
raid list	List available raids
raid create	Create a raid set
raid delete	Delete a raid set
raid get	Get raid set properties
raid add disk	Add a disk to a raid set
raid remove disk	Remove a disk from a raid set
raid repair	Repair a degraded raid set
raidsettings get	Get auto-incorporation and back-round disk settings
raidsettings set	Set the auto-incorporation and background disk properties
raid-speed-limit get	Get the current setting for the RAID sync speed limit.
raid-speed-limit set	Change the maximum RAID sync or resync speed. Use with caution.
reboot	Reboot the SnapServer
securitymodel get	Get the security model on a SnapServer Volume
securitymodel set	Set the security model on a SnapServer Volume
share create	Create a share
share delete	Delete a share
share get	View a share

Command	Description
share rename	Rename a share
share set	Modify a share
share list	List available shares
share access get	Get access list for the share
share access set	Set access list for the share
share access delete	Delete access permission of the specified user/group for the share
share nfsaccess get	Get NFS access permission of the host for the specified share
share nfsaccess set	Set NFS access permission of the host for the specified share
share nfsaccess delete	Delete NFS access permission of the host for the specified share
shutdown	Shutdown the SnapServer
slidingwindow get	Get sliding window settings for a specific interface
slidingwindow set	Set sliding window settings for a specific interface
slidingwindow list	List sliding window settings for all interfaces
snapex	Perform a control operation on the snap extension
snapshot create later	Create a new snapshot schedule
snapshot get	Get snapshot properties
snapshot set	Set properties for the specified snapshot
snapshot list	Get list of snapshots
snapshot create now	Create a new one time snapshot to be run immediately
snapshot delete	Delete specified snapshot
snapshot sched delete	Delete specified snapshot schedule
snapshot sched get	Get specified snapshot schedule
snapshot sched set	Set specified snapshot schedule
snapshot sched list	List current snapshot schedules
snapshot pool get	Get snapshot pool properties
snapshot pool set	Set snapshot pool properties
snapshot pool list	List current snapshot pools
snapshot rollback	Start a rollback for the specified snapshot
snmp get	Get SNMP parameters
snmp set	Set SNMP parameters
ssh get	Get current SSH settings
ssh set	Enable and Disable SSH. Note: Turning off SSH while running the command line will 'kick' the user off the system and they won't be able to log back into the command line until SSH is re-enabled via the SnapServer Web Administration
syslog all	Create a tar file of syswrapper and all third party logs
syslog edr	Create a tar file of Snap EDR logs
syslog netvault	Create a tar file of NetVault logs

Command	Description
syslog s2s	Create a tar file of S2Sv2 logs
syslog syswrapper	Create a tar file of syswrapper only
systemstatus	Get system status information for the server
tape list	List the SCSI tape devices
tape settings get	Display current SCSI tape device settings
tape settings set	Update SCSI tape device settings
tcpip get	Get TCP/IP parameters
tcpip set	Set TCP/IP parameters. Note: Changing the parameters of the Ethernet interface over which the user is currently running the SSH/command line session may result in the user being disconnected.
tcpip create bond	Create a bond and set TCP/IP properties.
tcpip delete bond	Remove a TCP/IP bond.
updatenotification get	Get update notification properties
updatenotification set	Set update notification properties
updatenotification check	Check to see if updates are available
ups get	Get UPS settings and status
ups set	Set UPS settings
user create	Create a local user
user delete	Delete a local user
user get	Get available users with their associated information
user list	List available users
user set	Change the properties of a local user
user lock	Lock the specified user.
user unlock	Unlock the specified user.
version	Display current version information, including the Server Number. Note: This is the same information displayed in the Web Administration "About" box
volume list	List of the volumes defined on the SnapServer
volume get	Get a specific volume's properties
volume create	Create a new logical volume
volume edit	Edit an existing logical volume
volume delete	Delete a logical volume
volume write-cache	Enable or disable write cache on a volume.
vxxaccess list	List hostnames with VSS/VDS access
vxxaccess add	Add hostname of VSS/VDS client requiring access to this server
vxxaccess delete	Delete access for a VSS/VDS client hostname
web get	Get current HTTP Web access settings
web set	Enable or Disable HTTP access to Web Administration interface
clear	Clear the screen

Command	Description
exit	Quit the command line, log off, and exit ssh/bash session. Note: If user has started another shell, the command 'exit' will return them to the SnapServer command line shell.
history	Print the history of commands typed into the SnapServer command line
less	With a file name, this command allows the user to view any file on the system. It should only be used for 'text' files.
Quit	Quit the command line, log off, and exit the ssh/bash session

Scripts in SnapCLI

Administrative tasks can be automated with shell scripts that call SnapCLI commands.

Running a SnapCLI Script

- 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.
- Log in to the SnapCLI (see Logging into SnapCLI for instructions).
- Type `osshell` to get a bash prompt.
- 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.

- 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
fi
```

```
return 1
else
echo "Group '$1' already exists."
fi
```

```
return 0
}
```

```
# usage: 'adduser2group <user_name> <group_name>'
adduser2group()
{
```

Add the user to the group

```
# if both the user and the group exist add the user as a member of this group
if $CLI user get user-name="$1" > /dev/null 2>&1; then
if $CLI group get group-name="$2" > /dev/null 2>&1; then
echo "Adding user '$1' to group '$2' ..."
$CLI group member add user-name="$1" group-name="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Adding user '$1' to group '$2' failed."
    fi
fi
fi
```

```
return 1
fi
```

```
fi
```

```
return 0
}
```

```
# usage: 'mkshare <share_name> <share_volume>'
mkshare()
{
```

Create a share

```
# if the share does not exist create it
if ! $CLI share get share-name="$1" > /dev/null 2>&1; then
echo "Creating share '$1' ..."
$CLI share create share-name="$1" share-volume="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Creating share '$1' failed."
    fi
return 1
else
echo "Share '$1' already exists."
fi

return 0
}
```

```
# usage: 'rmuser <user_name>'
rmuser()
{
```

Delete the user

```
# if the user exists then delete it
if $CLI user get user-name="$1" > /dev/null 2>&1; then
echo "Deleting user '$1' ..."
    $CLI user delete user-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of user '$1' failed."
    fi
return 1
else
    echo "User '$1' does not exist."
fi

return 0
}
```

```
# usage: 'rmgroup <group_name>'
rmgroup()
{
```

Delete the group

```
# if the group exists then delete it
if $CLI group get group-name="$1" > /dev/null 2>&1; then
echo "Deleting group '$1' ..."
    $CLI group delete group-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of group '$1' failed."
    fi
return 1
else
    echo "Group '$1P' does not exist."
fi

return 0
}
```

```
# usage: 'rmshare <share_name>'
rmshare()
{
```

Delete the share

```
# if the share exists delete it
if $CLI share get share-name="$1" > /dev/null 2>&1; then
echo "Deleting share '$1' ..."
```

```
        $CLI share delete share-name="$1" > /dev/null 2>&1
        if [ $? -ne 0 ]; then
            echo "Deletion of share '$1' failed."
        fi
    return 1
else
    echo "Share '$1' does not exist."
fi
return 0
}
```

Create a user, group, and share; then add the user to the group

```
#####
#   Main   #
#####

# create a user, a group and a share and add the user to the group
mkuser "$USER" "$PASSWORD"
mkgroup "$GROUP"
adduser2group "$USER" "$GROUP"
mkshare "$SHARE" "$VOLUME"

#remove the group, the user and the share
rmgroup "$GROUP"
rmuser "$USER"
rmshare "$SHARE"
```


Troubleshooting SnapServers

Basic techniques for identifying and resolving common hardware and networking issues are described here.

Topics in Troubleshooting SnapServers

- [The Meaning of LED Indicators](#)
- [System Reset Options](#)
- [Maintenance Mode](#)
- [Networking Issues](#)
- [Miscellaneous Issues](#)
- [Phone Home Support](#)

Additional Resources

Resource	Description
Knowledge Base	Search for solutions to specific issues by clicking the Knowledge Base link on the SnapServer support page: http://www.snapserver.com/kb
Hardware Components	Purchase additional hardware components from authorized SnapServer resellers. To locate a reseller in your area, select the How to Buy tab on the SnapServer home page: http://www.snapserver.com
Field Service Documents	Find a list of the hardware components available for your SnapServer or expansion array by navigating to the server or expansion array model: http://www.snapserver.com Procedures to install or replace components are available from the SnapServer support page: http://www.overlandstorage.com/support/crscd

The Meaning of LED Indicators

LED indicators provide information on the status of basic connectivity, disk drives, fan modules, and power supply modules.

- [SnapServer N2000 and SnapDisk E2000 Status and Drive Light Behavior](#)
- [SnapServer 110/210 Status and Drive Light Behavior](#)
- [SnapServer 410 Status and Drive Light Behavior](#)
- [SnapServer 500/600 Series Status and Drive Light Behavior](#)

- [Snap Expansion S50 Enclosure, Disk Drive, APC Module, and Controller Behavior](#)

SnapServer N2000 and SnapDisk E2000 Status and Drive Light Behavior

The SnapServer N2000 has one System light, two Network lights (Ethernet1, left; Ethernet2, right), and two disk lights per disk drive, as shown in the following illustration.

The SnapDisk E2000 has one System light and two Drive lights per drive.




Figure C-1: SnapServer N2000 LEDs


The LEDs operate as described in the following tables:

System LED	Description
Solid green	The unit is powered on but GuardianOS is not running.
Blinking green (N2000 only)	GuardianOS is booted and operating normally.
Network LEDs (N2000 only)	Description
Solid green	The server is active and connected to the network.
Off	The port is disconnected or the Ethernet cable is not connected or linked to an active switch.

Disk LEDs	Status	Description
Top LED	Off (SATA drive) Solid Blue (SAS drive)	Disk drive installed properly but is not active
	Blinking Blue	Disk drive installed properly and is active
Bottom LED	Solid Red	Disk drive error
All Drive LEDs (E2000 only)	Blinking simultaneously	UID identification from Disks/Units page of Web Management Interface.

Power Supply Module Indicator Lights


The LED on a SnapServer N2000 and E2000 power module is identified in the following illustration.



Power	Description
Solid green	The module is operating properly.
Solid amber	The module has failed, is not connected, or the server has been turned off.
Off	

SnapServer 110/210 Status and Drive Light Behavior

The server has two status lights, one network light, and one disk light, as shown in the following illustration:



Power and System LEDs

These status lights are located to the right of the power button. Looking at the server from the front, the lights appear in the following order, from left to right: Power LED, Status LED, Network LED, and Disk LED.

The LEDs operate as described in the following tables:

Power LED	Description
Solid green	The server is powered on.
Off	The server is powered off.


Status LED	Description
Blinking green	The server is operating normally.
Blinking amber	A thermal or other system problem was detected.
Blinking amber and green	The server is in Maintenance Mode.

Network LED	Description
Solid green	The server is active and connected to the network.
Off	The port is disconnected or the Ethernet cable is not connected or linked to an active switch.

Disk LED	Description
Blinking green	Disk drive is active.
Solid amber	Disk drive error.
Off	No disk drive activity.

SnapServer 410 Status and Drive Light Behavior

The server has two status lights, two network lights, and two lights for each of the four disk drives, as shown in the following illustration:



Overland Storage recommends that you become familiar with the operation of these lights.

Power, System, and LAN LEDs

These status lights are located to the right of the power button. Looking at the server from the front, the lights appear in the following order, from left to right: power LED, system LED, LAN 1 (Ethernet1) LED, and LAN 2 (Ethernet2) LED. The Disk Drive LEDs run along the bottom of the bezel, two LEDs for each disk drive.

The LEDs operate as described in the following tables:

Power LED	Description
Solid green	The server is powered on.
Off	The server is powered off.

System LED	Description
Blinking green	The server is operating normally.
Blinking amber	A thermal or other system problem was detected.
Blinking amber and green	The server is in Maintenance Mode.

LAN 1 and LAN 2 LEDs	Description
Solid green	The server is active and connected to the network.
Off	The port is disconnected or the Ethernet cable is not connected or linked to an active switch.


Disk Drive LEDs

Disk drive LEDs are located along the bottom of the bezel, two LEDs for each drive. For all disk drive LEDs, the left light indicates drive status; the right light indicates drive activity. They operate as follows:

Status LED (left)	Activity LED (right)	Description
Off	Off	Drive is not present.
Solid green	Blinking green	Disk drive installed properly and is active
Solid amber	Off	Disk drive installed, but not working correctly

SnapServer 500/600 Series Status and Drive Light Behavior

The server has two status lights, two network lights, two lights for each of the four disk drives, and an identification light, as shown in the following illustration:



Overland Storage recommends that you become familiar with the operation of these lights.

Power, System, and LAN LEDs

These status lights are located to the right of the power button. Looking at the server from the front, the lights appear in the following order, from left to right: power LED, system LED, LAN 1 (Ethernet1) LED, and LAN 2 (Ethernet2) LED. The Disk Drive LEDs are below the status lights, and the identification (Unit ID) LED is to the right of the LCD display.

The LEDs operate as described in the following tables:

Power LED	Description
Solid green	The server is powered on.
Off	The server is powered off.

System LED	Description
Double-blink green	The server is booting up.
Triple-blink green	The server is shutting down.
Solid or blinking amber at boot time	A problem was detected. The server will not boot.
Blinking amber during normal operation	A thermal or other system problem was detected.
Blinking amber and green	The server is in Maintenance Mode.

LAN 1 and LAN 2 LEDs	Description
Solid green	The server is active and connected to the network.
Off	The port is disconnected or the Ethernet cable is not connected or linked to an active switch.

Unit ID (UID) Front and Back LEDs	Description
Blue	Unit ID is on and identifies the unit (front and back).
Off	Unit ID has not been turned on.


Disk Drive LEDs

Disk drive LEDs on these SnapServers are located beneath the status lights on the bezel. For all disk drive LEDs, the left light indicates drive status; the right light indicates drive activity. They operate as follows:

Status LED (left)	Activity LED (right)	Description
Solid green	Off (SATA drive) Solid green (SAS drive)	Disk drive installed properly but is not active
Solid green	Blinking green	Disk drive installed properly and is active
Solid amber	Off	Disk drive installed, but not working correctly
Off	Off	No disk drive installed

Power Supply Module Indicator Lights

The LED on a 500/600 Series power module is identified in the following illustration.




Power	Description
Solid green	The module is operating properly.
Blinking green	The module has failed or is not connected.
Solid amber	The module has failed, is not connected, or the server has been turned off.
Off	

Snap Expansion S50 Enclosure, Disk Drive, APC Module, and Controller Behavior

This section describes the LED indicators on the Snap Expansion S50 enclosure, disk drives, and APC modules.

Enclosure LEDs

On the right front of the Snap Expansion S50 (as shown in the following illustration) are four LEDs that indicate the status of the enclosure.




These LEDs operate as described in the following table:

LED	Condition	Description
1	On	Enclosure power on.
2	On	Fault on enclosure. When a failure occurs on a controller or APC unit, the enclosure LEDs indicate an enclosure fault.

LED	Condition	Description
3	Solid green	Host Link. A solid green LED indicates communication with the SnapServer.
4	Rapidly flashing green	Unit ID. Identifies the expansion unit when you click the expansion unit ID link in the Disks and Units screen.

Disk Drive LEDs

Each disk drive has three LEDs that indicate the status of the disk drive.




These LEDs operate as described in the following table:

LED	Condition	Description
1	N/A	Not used.
2	Solid green	Drive present and OK.
	Solid amber	Drive failed.
	Off	Drive not present
3	Solid green	Drive present and idle.
	Green random flash	I/O activity on disk drive.
	Off	Drive not present.

APC Unit LEDs

Each APC unit has two LEDs that indicate status.




These LEDs operate as described in the following table:

LED	Condition	Description
Power	Off	Enclosure not powered on.
	Green (solid)	APC unit functioning normally.
Fault	Off	APC unit functioning normally.
	Yellow (flash)	APC unit failure predicted.
	Yellow (solid)	APC unit failed.

When an APC unit fails, the enclosure LED on the front of the unit also indicates the failure.

Controller LEDs

The controller has two LEDs that indicate status.



These LEDs operate as described in the following table:

LED	Condition	Description
Master	On	Controller is current Master for enclosure.
Fault	On 5-10 seconds only	Enclosure is powering on.
	On continuously	Fault condition exists on controller

System Reset Options

Often the first thing to try in resolving anomalous behavior on a SnapServer is to reset the server to factory defaults. This section provides information about the following ways to reinstall or reset the system defaults.

Resetting the SnapServer to Factory Defaults

GuardianOS allows you to reset different components of the system. Default settings can be found in the default configuration sections of this guide.

CAUTION: Each reset option requires a reboot 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** screen 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 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 protocol settings to factory defaults, eliminates all shares to all volumes, and returns settings for server name, date and time, users, groups, 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 and SnapTrees are all set to the Windows/Mixed security model. All files and directories are set to the Windows personality with a Windows ACL that gives full access to Administrators, read access to Everyone, file/directory create access to Everyone (for directories), and full access to the owner (owners are retained in the reset operation).
- You cannot initiate a reset to defaults if a Snaptree conversion is in progress.
- Rebooting or shutting down the server in the middle of an ACL reset will halt the operation, and it will not recommence on reboot.

Performing System Resets Without Network Access

Should access to the server be lost, the Reset or LCD panel buttons can be used to reset server settings and re-establish connectivity.

To Perform a Limited Reset Using the Reset Button

On the SnapServer N2000, the Reset button is located below the Power button on the server flange. On all other SnapServers, the Reset button is accessed via a small hole next to the Power button on the front of the server. Verify that the server is fully booted (as indicated by the System LED blinking once per second), and push the **Reset** button. The system will reboot after about a minute. As a part of the reset and reboot process, the SnapServer does the following:

- Clears user-defined settings such as DHCP configuration
- Resets the server name to its default setting (**SNAP<server_number>**)
- Resets network speed and bonding settings to their defaults
- Resets the Administrator password to the default (**admin**)
- Resets the web server to allow http

Maintenance Mode

The SnapServer may enter Maintenance Mode (System LED blinking amber and green) when GuardianOS has been compromised and is in need of repair or reinstallation. The two functions available in Maintenance Mode should only be used under the direction of Overland Technical Support:

- **Repair** – Reapplies the GuardianOSImage, but preserves system settings.
- **Fresh install** – Reinstalls GuardianOS, overwriting any previous configurations and destroying all disk partitions.

CAUTION: Because of significant changes introduced in GuardianOS 6.5, a fresh install of GuardianOS 6.5 should not be performed on a SnapServer running an older version of GuardianOS. Failure to follow this guideline can result in total failure of the SnapServer to start, even into Maintenance Mode. The Fresh install option should only be performed with the same version of GuardianOS currently installed on the SnapServer, and only under the direction of Overland Technical Support.

Note: This file is available for download by entitled users from the SnapServer support site: <http://docs.overlandstorage.com/snapserver>.

Networking Issues

These are some of the networking issues you may encounter when using your SnapServer.

The Server Cannot Be Accessed over the Network

Inaccessibility may be caused by a number of reasons. To resolve this issue, use one of the following methods:

- Verify that you have the correct IP address of the server, and try to connect again.
- Verify that the LED for the primary Ethernet port is lit. (This light indicates network connectivity.) If the light is not lit, perform the following in order:
 - The most likely cause is the physical connection. Check for a loose or damaged cable, or poor connections in the port connector.
 - This problem may also be caused by a mismatch between the settings on the switch or hub and the settings on the SnapServer Ethernet port. These settings must match. To resolve the problem, make sure the port settings on the hub or switch match the settings for the primary port as configured on the **Network > TCP/IP** screen of the Administrator Tool. Use the autonegotiate setting on both the switch and the server port.

You Have No Access to the SnapServer via HTTP

When trying to access the SnapServer via HTTP, the Web browser times out. The server can be accessed using the ping command or Windows Explorer.

- HTTP and HTTPS are both enabled by default on SnapServers. Try typing HTTPS in the Web address rather than HTTP. If you are able to access the server via HTTPS, you can re-enable HTTP on the **Network > Web** screen.
- If you cannot access the server via HTTPS, try resetting the server as described on [“Resetting the SnapServer to Factory Defaults” on page C-9](#).

An Access Denied Message Appears after Configuring Microsoft Domain Security

Customers who have configured local users and local groups with the same name as their domain users and groups can have security conflicts if they integrate with Microsoft Domain Security. The SnapServer will authenticate the users as local SnapServer users before authenticating through the Domain. However, the Domain users/groups may be the ones that had been granted access to the shares.

Be careful not to add local users or groups that are duplicates of those that are found on the Windows domain controller.

The SnapServer Does Not Operate Properly on a Network Running Gigabit-Full-Duplex

For Gigabit Ethernet to operate properly, both the switch and the SnapServer's primary Ethernet port must be set to *Auto* (autonegotiate). Any other setting will result in unexpected behavior and reduced performance.

The Network Does Not Have a DHCP Server and the SnapServer IP Address Is Unknown

Install SnapServer Manager (available from the [SnapServer support page](#) on the Overland Storage website) onto a client workstation on the same subnet as the

SnapServer. You can then use the utility to discover all SnapServers on that network segment, and to assign static IP addresses as necessary.

Apple Users Cannot Log into the SnapServer as Windows Users

To allow Apple users to access a SnapServer, replicate their user names and passwords locally on the SnapServer.

Problems Occur with Domain Controller Authentication

You are receiving the following errors in your error log:

```
SMB: Domain Controller unavailable
```

```
SMB: Username not connected to Domain Controller
```

This means that either your Domain Controller is down, or the SnapServer is unable to reach it. Because it cannot communicate with the Domain Controller, it is not able to authenticate the user. Check to make sure the Domain Controller is online, is consistently reachable via the network, and that users can authenticate to the Domain Controller.

You Start Your SnapServer but Cannot See It on the Network

Ensure that the Ethernet cable is connected securely to both the network port and the server's primary Ethernet port. Also, check to see that the Link light on the front of the SnapServer is lit (solid green). If the Link light is off, this is normally caused by a mismatch between the switch/hub and the Ethernet port on the SnapServer. To resolve this problem, verify that all settings (if using multiple Ethernet ports) on the switch/hub match the setting on the server. When the server is shipped from the factory, both ports are set to autonegotiate. Therefore, the switch/hub *must* be set to autonegotiate to initially connect to the server.

SnapServers are configured to acquire an IP address from a DHCP server. If no DHCP server is found on the network, the SnapServer defaults to an IP address in the range of 169.254.xxx.xxx and is labeled ZeroConf in SSM. While you may not be able to see the server on your network, you can discover the SnapServer using either the default server name or the SSM utility (available at our external download site: <http://www.overlandstorage.com/SSM>). Use the server name method if you are installing one SnapServer on the network. Use SSM if you are installing two or more SnapServers, or if your network does not have IP-to-name resolution services.

You Try to Mount to a Share on Your SnapServer from Your Linux Workstation and You Receive an RPC Timeout Message

Check the firewall configuration to your Linux workstation. Be sure you have not blocked the ability to receive TCP or User Datagram Protocol (UDP) communications. If problems persist, contact Overland Storage Technical Support.

You Receive an Access Denied Message When Attempting to Mount a Share on Your SnapServer from a Linux Workstation

If you are logged in as *root* on your workstation and NFS is enabled on your SnapServer, this message can be misleading, causing you to look for security issues, when in fact it could be a command syntax issue. For example, the common Linux mount command:

```
mount 192.168.32.124:SHARE1 /mnt
```

is missing a forward slash (/) in the command, which will return an Access Denied message. The correct syntax should be the following:

```
mount 192.168.32.124:/SHARE1 /mnt
```

Note: The share name is case sensitive.

You Cannot Log in as Root to the SnapServer

GuardianOS allows you to log in as root over SMB. If this operation has failed or you have trouble logging in, be sure that you have enabled root login in the **Network > Windows** page. Also note that the root account password is tied to the admin account password. If you cannot log in as root, change the password for the admin account on the **Network > Windows** screen. Use the admin password to log in as root.

You Are Unable to See Your Domain Users When Trying to Set Up Windows Security Permissions on File Folders

The SnapServer (GuardianOS) has joined the Active Directory domain properly, and you can see the domain users when you set Share permissions from the browser-based Web Management Interface.

Make sure the Windows client (PC) you are trying to set permissions from is assigned a valid DNS server. You can check your Windows client using the **ipconfig** command from a command prompt.

Miscellaneous Issues

These are some miscellaneous issues you may encounter when using your SnapServer.

Back Up Applications

You Backed Up Your Snapshot Share, Are Now Attempting to Restore It, and the Operation Fails

A snapshot share is read-only. You can restore the data to a read-write accessible share.

The NetVault Client Cannot Connect to the NetVault Server on the SnapServer

Occasionally, after enabling NetVault for GuardianOS for the first time, the NetVault for GuardianOS Server may not start properly. If this happens, the NetVault client application may not be able to connect to the NetVault for GuardianOS server running on the SnapServer. To resolve this issue, simply disable and then re-enable the NetVault for GuardianOS Server via the **SnapExtensions > BakBone NetVault** screen.

BakBone NetVault Restore Limitations for UNIX SnapTrees

File and directory permissions will be restored when using BakBone NetVault. However, when Windows file and directory permissions are restored to a UNIX SnapTree on a SnapServer, the Windows-style extended permissions are removed to preserve proper UNIX Snaptree permissions.

Other Issues

A Problem Occurred While Booting. The System is Offline and the Status LED is Blinking Amber and Green

The SnapServer has booted into Maintenance (Recovery) Mode. This may be due to a boot failure in the previous boot attempt. Try booting again. If the server still returns to Maintenance Mode, call Technical Support.

Power to the SnapServer Is Unexpectedly Cut Off Due to a Power Outage

Overland Storage recommends that you use an uninterruptible power supply (UPS) with the SnapServer. If you did not have a UPS attached to the server at the time of the power outage, do the following:

1. On SnapServers with no on/off switch, remove the power cables. On Snap Expansion S50, turn off the power switches on the back of the unit.
2. Once the power is restored and stabilized, turn the power supplies back on and reboot the server.

Once the SnapServer boots, it begins resynchronizing the RAIDs if necessary. You can use the server during the resynchronization, but performance will be a little slower than normal. Do not remove drives, however, while the server is resynchronizing the RAID.

The Server Is Not Responding to File Requests or Configuration Commands

Call your SnapServer technical support representative.

You Have Problems Seeing the Tape Library Tape Device, Not the Robotic Arm

When you have problems seeing the actual tape device rather than the robotic arm, it is most likely due to the Tape Loader being configured for Sequential Access. Change the Tape Loader to Random or Mixed Mode.

The Admin Password to the Web Management Interface Is Not Available

You can perform a limited reset to defaults, which includes the admin password (as described in [“Performing System Resets Without Network Access” on page C-10](#)); then use the Web Management Interface to set a new password.

The SnapServer 510, 520, 550, 620, or 650 LCD is Flashing

A flashing LCD indicates a server panic. In some cases, rebooting the server may solve the problem. However, if this condition occurs more than once, try resetting the system as described in [“Performing System Resets Without Network Access” on page C-10](#).

You Can Not Delete Files or Folders From an iSCSI Disk

If an iSCSI disk is mounted to a folder, not a letter drive, in Windows you will not be able to delete files and folders inside that mount point. The Windows Recycle Bin does not understand mount points, so to avoid this problem either mount iSCSI disks to letter drives on your Windows OS, or hold down the shift key while deleting folders or files.

Phone Home Support

Once your SnapServer has been registered, Phone Home Support becomes available for use. Phone Home Support emails system logs and files that contain information useful for troubleshooting purposes to Overland Storage technical support. You can use the **Monitor > Support** screen this screen to open a new case with technical support; or, in the course of working to resolve an issue, a technical support representative may ask you to fill out and submit this page. If a case is already in progress, you will need to enter the case number provided by the technical support representative.

Notes: Phone Home Support interacts with two fields on the **Server > Email Notification** screen: (1) To use Phone Home Support, you must enter a valid SMTP server IP address on the Email Notification screen; and (2) the first email address listed in the Recipient(s) field populates the Admin Email Address field on the Support screen.

Complete the following fields as appropriate, then click **OK**:

Text Field	Description
Subject	(Required) Enter a concise description that identifies the issue.
Case	(Required) Select <i>New Case</i> if you are emailing technical support for the first time. Select <i>Existing Case</i> if you have previously contacted technical support concerning the issue.
Case Number	If you selected <i>Existing Case</i> above, enter the case number provided by technical support.
Reply-to Address	(Required) This field defaults to the first email address entered as a recipient on the Server > Email Notification screen. If necessary, enter at least one email address that will serve as the contact email address for this issue. To receive a copy of the email and system information attachment, select the <i>Cc Admin</i> checkbox.
Comments	(Required) Enter additional information that will assist in the resolution of the problem.

GuardianOS Ports

The following table outlines the ports used in GuardianOS.

Port #	Layer	GOS Feature	Name	Comment
1	DDP		rtpm	Routing Table Management Protocol
1	TCP & UDP		tcpmux	TCP port service multiplexer
2	DDP		nbp	Name Binding Protocol
21	TCP & UDP	Network > FTP	ftp	File Transfer Protocol (FTP) port; sometimes used by File Service Protocol (FSP)
22	TCP & UDP	Server > SSH	ssh	Secure Shell (SSH) service
25	TCP & UDP	Server > Email Notification	smtp	Simple Mail Transfer Protocol (SMTP)
67	TCP & UDP	Network > TCP/IP	bootps	Bootstrap Protocol (BOOTP) services; also used by Dynamic Host Configuration Protocol (DHCP) services
68	TCP & UDP	Network > TCP/IP	bootpc	Bootstrap (BOOTP) client; also used by Dynamic Host Control Protocol (DHCP) clients
80	TCP & UDP	Web Management Interface	http	HyperText Transfer Protocol (HTTP) for World Wide Web (WWW) services
81	TCP	Web Management Interface	HTTP	Hypertext Transport Protocol
88	TCP & UDP	Network > NFS	Kerberos	Kerberos Security (NFSv4)
111	TCP & UDP	<ul style="list-style-type: none"> • Networking > NFS • Assist • SnapServer Manager 	sunrpc	Remote Procedure Call (RPC) Protocol for remote command execution, used by Network Filesystem (NFS) and SnapServer Manager

Port #	Layer	GOS Feature	Name	Comment
123	TCP & UDP	Server > Date/Time > Advanced	ntp	Network Time Protocol (NTP)
137	TCP & UDP	Network > Windows	netbios-ns	NETBIOS Name Services used in Red Hat Enterprise Linux by Samba
138	TCP & UDP	Network > Windows	netbios-dgm	NETBIOS Datagram Services used in Red Hat Enterprise Linux by Samba
139	TCP & UDP	Network > Windows	netbios-ssn	NETBIOS Session Services used in Red Hat Enterprise Linux by Samba
161	TCP & UDP	Network > SNMP	snmp	Simple Network Management Protocol (SNMP)
162	TCP & UDP	Network > SNMP	snmptrap	Traps for SNMP
389	TCP & UDP	Network > Windows	ldap	Lightweight Directory Access Protocol (LDAP)
443	TCP & UDP	<ul style="list-style-type: none"> • Web Management Interface • SnapServer Manager • SnapExtension > Snap EDR 	https	Secure Hypertext Transfer Protocol (HTTP).
445	TCP & UDP	Network > Windows	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 File System (NFS)

Port #	Layer	GOS Feature	Name	Comment
2050	UDP	Network > NFS	mountd	
2599	UDP	<ul style="list-style-type: none"> • Snap Finder • SnapServer Manager 	Sysbroker	Multicast Discovery
3052	TCP	Server > UPS		Port for monitoring UPS status
3205	TCP	Network > iSCSI	iSNS	
3260	TCP	Network > iSCSI	iSCSI	
8001	TCP	SnapExtension > SnapEDR	SnapEDR	External Communications
8002	TCP	SnapExtension > SnapEDR	SnapEDR	External Communications
8003	TCP	SnapExtension > 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
10000	TCP	SnapExtension > BakBone NetVault	NetVault	
10001	TCP	Snap Extension	Snap Extension	Shutdown Port
12000	TCP & UDP	Network > Apple	afp2overtcp	Second NIC
12168	TCP	CA Antivirus	inoweb	Admin Interface
16384	UDP		Sysbroker	Random Port
16388	UDP		Sysbroker	Random Port
20031	TCP	SnapExtension > BakBone NetVault	NetVault	Listening 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	SnapExtension > SnapEDR	SnapEDR	External Communications Port
49229	TCP	SnapExtension > SnapEDR	SnapEDR	External Communications Port
1024 - 65535	TCP & UDP	<ul style="list-style-type: none"> • Network > NFS • Network > FTP 	NFS FTP (passive)	Dynamically allocated in runtime for user connections

Glossary

Term	Definition
access permissions	A rule associated with a share, a file, or a directory to regulate which users can have access to the share and in what manner.
ACL (Access Control List)	The list that controls access to directories and files. Each ACL includes a set of access control entries, which contain the metadata that the system uses to determine access parameters for specified users and groups.
ADS (Active Directory Service)	The preferred authentication method for Windows XP, Windows 2000, Windows 2000 Advanced Server, and Windows 3000 network users. This authentication allows Active Directory users to connect to shares on the SnapServer. The SnapServer supports the Microsoft Windows 2000 family of servers that run in native ADS mode.
agent	A program that performs some information-gathering or processing task in the background. SnapServers support Data Protection Agents and can be configured as SNMP agents.
algorithm	A sequence of steps designed to solve a problem or execute a process.
AllLocalUsers group	The default group for all local users on SnapServers. Local users are set up by the SnapServer administrator. Network users or Windows domain users are not part of the AllLocalUsers group.
AllUsers group	A collection of all users. The SnapServer automatically maintains the AllUsers group.
array	A series of objects, all of which are the same size and type. In a server context, an array refers to the grouping of hard drives into a RAID set.
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.
bonding	A technology that treats two ports as a single channel, with the network using one IP address for the server. SnapServers support load balancing and failover bonding modes.
CA Antivirus	The antivirus software bundled with the SnapServer.

Term	Definition
chaining	A native SnapServer technology in which all snapshots of a volume depend on successive snapshots for part of their content.
channel	A communications path between two computers or devices.
CHAP (Challenge Handshake Authentication Protocol)	CHAP verifies the identity of the peer using a three-way handshake.
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 (Common Internet File System)	Also known as SMB (Server Message Block) . The default Windows protocol for communication between computers. A specification for an Internet file access protocol that complements HTTP and FTP and reduces access time.
daemon	A process that runs in the background.
default gateway	The router used when there is otherwise no known route to a given subnet.
degraded	A RAID state caused by the failure or removal of a disk drive in which data is consistent, but there is no redundancy.
DHCP (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. The SnapServer can be configured to perform as a DHCP server and assign IP addresses with a single subnet.
directory	A virtual folder used to organize files. Also called a folder.
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.
disk	A rigid platter, usually constructed of aluminum or mylar, with a magnetic surface that allows the recording of data, that is stored inside the drive.
DNS server (Domain Name System server)	The 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).

Term	Definition
Ethernet address	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet interface.
Ethernet port	The port that houses the network card to provide Ethernet access to the computer.
event	Any significant occurrence in the system that may require notifying a system administrator or adding an entry to a log.
failover	A strategy that enables one Ethernet port to assume the role of another port if the first port fails. If a port fails on a SnapServer, the second port assumes its network identity (if the two Ethernet cards have been configured for failover). When the port comes back online, the original identities are restored. Failover is possible only in a multi-Ethernet configuration.
FTP (File Transfer Protocol)	A standard Internet protocol that provides a way to exchange files between computers on the Internet. By default, a SnapServer is set up to be an FTP server.
full-duplex	A type of transmission that allows communicating systems to both transmit and receive data simultaneously.
gateway	The hardware or software that bridges the gap between two network subnets. It allows data to be transferred among computers that are on different subnets.
GID (group IDs)	On a SnapServer, the unique ID assigned to each group for security purposes.
GuardianOSImage.gsu	An image file used to upgrade GuardianOS.
half-duplex	A type of transmission that transfers data in one way at a time.
hidden share	A share that restricts the display of the share via the Windows (SMB), Web View (HTTP/HTTPS), FTP, and AFP protocols.
host name	The unique name by which a computer is known on a network. It is used to identify the computer in electronic information interchange.
hot spare (local or global)	A disk drive that can automatically replace a damaged drive in a RAID 1 or 5. If one disk drive in a RAID fails or is not operating properly, the RAID automatically uses the hot spare to rebuild itself without administrator intervention. A <i>local</i> hot spare is associated with and available only to a single RAID. A <i>global</i> hot spare is associated with a single RAID, but may be used for any RAID in the system.
hot swapping	The ability to remove and add disk drives to a system without the need to power down or interrupt client access to file systems.
HTTP (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 (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.
I/O (Input/Output)	The operation of transferring data to or from a device, typically through an interface protocol like CIFS, NFS, or HTTP. The SnapServer presents a file system to the user and handles block I/O internally to a RAID array.

Term	Definition
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.
IP (Internet Protocol) address	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.
iSCSI (Internet SCSI)	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.
Jukebox	A robotic tape backup device that stores numerous tape drives and uses a mechanical arm to bring the drive to a station for reading and writing.
JVM (Java Virtual Machine)	Software that converts Java bytecode into machine language and executes it. A JVM allows an application such as SnapServer Manager written in Java to run on any operating system.
Kerberos	<p>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.</p> <p>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.</p>
LCD (Liquid Crystal Display)	An electronic device that uses liquid crystal to display messages on some SnapServers.
LED (Light-Emitting Diode)	An electronic device that lights up when electricity is passed through it.
Linux	A UNIX-like operating system that was designed to provide personal computer users a free or very low-cost operating system comparable to traditional and usually more expensive UNIX systems. GuardianOS is based on the Linux operating system.
load balancing	A process available only in multi-Ethernet configurations. The Ethernet port transmission load is distributed among two or more network ports (assuming the cards are configured for load balancing). An intelligent software adaptive agent repeatedly analyzes the traffic flow from the server and distributes the packets based on destination addresses.
local group/local user	A group/user defined locally on a SnapServer using the Web Management Interface. The local user is defined by the server administrator. Windows domain, ADS, and NIS users are not considered local.

Term	Definition
MAC (Media Access Control)	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. SnapServers with dual-Ethernet ports can respond to a request with either port and have two unique MAC addresses.
maintenance mode	A series of HTML screens in the GuardianOS Web Management Interface that allows you to perform repair, upgrade, or reinstall GuardianOS in a disaster recovery situation.
MIB (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, 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 file system that is available.
multihomed	A SnapServer that is connected to two or more networks or has two or more network addresses.
NAS (Network Attached Storage)	Hard disk storage that is set up with its own network address as opposed to being attached to the department computer that is serving applications to a network's workstation users. By removing storage access and its management from the department server, both application programming and files can be served faster because they are not competing for the same processor resources. The NAS device is attached to a local area network (typically an Ethernet network) and assigned an IP address.
NetVault for GuardianOS	A comprehensive backup solution that is preinstalled on SnapServers running GuardianOS 2.6 through 5.2 to support backup and restore operations to a local tape drive.
NFS (Network File System)	A client/server application that allows a computer user to view and optionally store and update files on a remote computer as though they were on the user's own computer. The user's system needs to have an NFS client and the other computer needs the NFS server. The SnapServer is configured as an NFS server by default.
NIS (Network Information Service)	A network naming and administration system for smaller networks that was developed by Sun Microsystems. NIS+ is a later version that provides additional security and other facilities. The SnapServer accepts NIS users and groups.
node	Any device, including servers, workstations, or tape devices, that are connected to a network; also the point where devices are connected.
NVDB (NetVault Database) directory	A NetVault for GuardianOS database directory stored on the SnapServer that holds records for the media and backups performed.
orphan	A disk drive that has become disconnected from its RAID either by accidental removal of the drive or the intermittent failure of the drive.

Term	Definition
parity	Error correction data. RAID 5 stores equal portions of each file on each disk and distributes parity information for each file across all disks in the group. This distributed parity allows the system to recover from a single disk drive failure.
Permissions	A security category, such as no access, read-only, or read-write, that determines what operations a user or group can perform on folders or files.
PoP (Proof of Purchase)	The number used to obtain a license key for an upgrade to third-party applications.
POSIX (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. Pre-GuardianOS 5.0 SnapServers use Extended POSIX ACLs.
protocol	A standardized set of rules that specifies the format, timing, sequencing, and/or error checking for data transmissions.
public access share	A share that allows all users read/write access to the file system.
quota	A limit on the amount of storage space on a volume that a specific user or NIS group can consume.
RAID (Redundant Array of Independent Disks)	A collection of disk drives that act together as a single storage system. Different RAID types provide different levels of data protection.
RAID 0 (Striped)	Distributes data evenly among all disks in the array. This technique, called data striping, results in fast access speeds because it uses multiple physical devices to store the data. However, RAID 0 offers no redundancy and does not accept hot spares. If a single disk drive fails, every file in the RAID is rendered unavailable.
RAID 1 (Mirrored)	Stores data on one disk drive and copies it to another drive in the RAID. A RAID 1 must contain at least two disk drives: one for the data space and one for redundancy. Although the data space in a RAID 1 can never be larger than a single drive, some administrators prefer to add a third drive (either as a hot spare or a member) for additional redundancy. RAID 1 is the most secure method for storing mission-critical data because there is no catastrophic data loss when a disk fails. However, RAID 1 is the most expensive and least efficient storage method.
RAID 5 (Striping with Parity)	Distributes data evenly among all disks in the array, and maintains parity information (error correction data) that allows the system to recover from a single disk drive failure. RAID 5 provides the best combination of performance, usability, capacity, and data protection.
RAID 6 (Striping with Dual Parity)	Similar to RAID 5 except that two drives maintain parity information for greater redundancy. System can recover from two drive failures. Provides high reliability and data protection but write performance speed is impacted by the dual parity drives.
RAID 10 (Striped Mirroring)	RAID 10 is two or more RAID 1's striped together to provide greater redundancy and higher performance than a simple RAID 1.

Term	Definition
recurring snapshot	A snapshot that runs at an administrator-specified time and interval.
restrict anonymous	<p>A Windows feature in which anonymous users cannot list domain user names and enumerate share names. Microsoft has provided a mechanism in the Registry called restrict anonymous for administrators to restrict the ability for anonymous logon users (also known as NULL session connections) to list account names and enumerate share names.</p> <p>The implementation of the restrict anonymous mechanism may prevent the SnapServer from obtaining the list of account names it needs to authenticate Windows domain users.</p>
resynchronization	A RAID state that describes the process of integrating a new drive into the RAID.
rollback	A snapshot feature that allows the administrator to restore a volume to a previous state as archived in a snapshot without resorting to tape.
SCSI (Small Computer System Interface)	A parallel interface standard used to attach peripheral devices, such as robotic libraries, to computers.
serial number	The ten-character alphanumeric number assigned by the manufacturer at the factory.
server number	A numeric derived from the MAC address of your SnapServer's primary Ethernet port that is used to uniquely identify a SnapServer.
share	A virtual folder that maps to the root of a volume or a directory on the volume. Permissions are assigned to a share that determine access for specific users and groups.
share access	Permissions granted or denied to users and groups that control user and group access to the files.
SMB (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.
Snap EDR	A SnapExtension that copies the contents of a share from one SnapServer to another share on one or more SnapServers. Snap EDR is designed to work with SnapServers and other SnapServer Storage Solutions.
SnapServer Manager (SSM)	A Java-based utility for discovering and monitoring SnapServers.
SnapDRImage	The SnapServer disaster recovery image that saves server-specific settings such as server name, network, RAID, volume and share configuration, local user and group lists, and snapshot schedules.
SnapExtension	A Java application that extends a SnapServer's functionality. SnapExtensions are produced both by SnapServer and third-party vendors.
snapshot	A consistent, stable, point-in-time image of a volume (file system) used for backup purposes.

Term	Definition
snapshot pool	Disk space reserved within a RAID for the storage of snapshot data. In the default storage configuration of many SnapServers, twenty percent of the RAID capacity is allocated to the snapshot pool.
snapshot share	A virtual folder that allows access to all current snapshots at the same directory level as the original share on which it is based.
SnapTree Directory	A directory residing in the root of a volume that is assigned a Windows- or UNIX-style security model. The security model determines the file-level security scheme that will apply to files, folders, and subdirectories within the SnapTree directory.
SNMP (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 (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 (Secure Sockets Layer)	A technology that provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.
standalone	A network bonding mode which treats each port as a separate interface. This configuration should be used only in multihomed environments in which network storage resources must reside on two separate subnets.
static IP address	An IP address defined by the system administrator rather than by an automated system, such as DHCP. The SnapServer allows administrators to use DHCP-assigned or statically assigned IP addresses.
striping	A RAID storage technique that distributes data evenly among all disks in the array.
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.
TCP/IP (Transmission Control Protocol/Internet Protocol)	A commonly used networking protocol that supports the interconnection of different network operating systems.
trap	A signal from the SnapServer 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.
UI (User Interface)	The User Interface is the graphical and textual presentation called Web Management Interface that displays the GuardianOS interface in your web browser.
UID (User IDs)	A unique ID assigned to each user on a SnapServer for security purposes.
unassigned	The state of a disk drive that is seated in a bay but has not been incorporated into a RAID.

Term	Definition
UNC (Universal Naming Convention)	In a network, a way to identify a shared file in a computer without having to specify (or know) the storage device it is on. In the Windows OS, the UNC name format is as follows: $\backslash \textit{server_name} \backslash \textit{share_name} \backslash \textit{path} \backslash \textit{file_name}$
UPS (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 (Uniform Resource Locator)	A Web address.
Virtual Disk Service (VDS)	Microsoft VDS is a service that extends existing storage capabilities of Windows Server operating systems.
volume	A logical partition of a RAID's storage space that contains a file system. In the default storage configuration of many SnapServers, eighty percent of the RAID capacity is allocated to the default volume.
Volume Shadow Copy Service (VSS)	Microsoft VSS provides a mechanism for creating consistent point-in-time copies of data known as shadow copies.
Web Management Interface	A Web-based utility used for configuration and ongoing maintenance, such as monitoring server conditions, configuring email alerts for key events, or for SNMP management.
Web View	The Web-browser screen that opens when users access a SnapServer using their Web browsers, and displays a list of all shares.
Windows domain authentication	Windows-based networks use a domain controller to store user credentials. The domain controller can validate all authentication requests on behalf of other systems in the domain. The domain controller can also generate encrypted challenges to test the validity of user credentials. Other systems use encrypted challenges to respond to CIFS/SMB clients that request access to a share.
WINS (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) **PR-iv**

A

Access

- problems with **C-11**
- users and groups **3-1**
- Windows ACLs **6-18**

Access Denied Message C-11

ACLs

- defined **1-1**
- backing up **8-2**
- resetting to defaults **C-10**
- setting file-level permissions (Windows) **6-19**

Active Directory

- defined **1-1**
- and name resolution servers **2-9**
- joining AD domain **2-12**
- SnapServer interoperability with **2-10**

Admin password

- default **3-2**
- resetting forgotten **C-9, C-14**

AFP, see Mac OS

alert definitions PR-iv

Antivirus

- and volume deletion **4-12**
- dependencies on other software components **9-1**
- distributing updates **9-7**
- enabling **1-11**
- excluding snapshots from **9-3**
- HTTP requirement **9-1**
- launching configuration GUI **9-2**
- scan job configuration **9-3**
- using logs **9-10**

APC-Brand UPS, see UPS

Authentication

- default settings **3-1**

HTTPS/HTTP 2-18

Kerberos 2-10

local 3-2

NIS domain 3-7

UID and GID assignments 3-2

Autonegotiation 2-7

B

Backing up Unicode servers 10-3

Backup

- coordinating with Snapshots **7-4**
- identifying backup/media servers to the SnapServer **A-5**
- inability to back up iSCSI Disks **5-10, 7-2**
- iSCSI Disks **5-10**
- of NetVault directory **8-4**
- of server and volume settings **8-1**
- supported third-party **A-1**

backup.acl 8-2

backup.qta.groups 8-2

backup.qta.users 8-2

C

CA Antivirus, see Antivirus

CA ARCserve, installing agent A-5

CA Unicenter TNg 11-6

Chooser, see Mac OS

Client access, configuring

- Apple **2-15**
- FTP **2-17**
- HTTPS/HTTP **2-18**
- NFS **2-13**

Cloning a server 8-7

Command Line Interface B-1

Running Scripts **B-8**

Syntax **B-1**

Connecting

- to SnapServers 1-7
- conventions, typographical PR-iv
- customer support PR-iii

D**Defaults**

- admin password 3-1
- authentication 3-2
- protocol access 2-8
- resetting to factory C-9
- TCP/IP 2-3

Disaster Recovery

- backing up server and volume settings 8-1
- creating recovery files 8-2
- procedures 8-6

Disk drives

- adding 4-28
- automatic incorporation 4-4, 11-14
- hot swapping 4-30
- reintegrating orphaned 4-29
- replacing 4-27

Documents

- related to SnapServers PR-v

Domain, joining

- NIS 3-8

download website link PR-v**E****electrostatic discharge information PR-v****E-mail notification of server events 11-4****EMC NetWorker**

- installing agent A-11
- special backup and restore operations A-14

Ethernet, see *Gigabit Ethernet***Expand Volume button 4-12****Exports file, NFS 6-1****F****Factory defaults, resetting to C-9****Failover, see *Network bonding*****Features, changes from earlier releases 1-3****Field Service Documents C-1****Files, setting permissions for 6-18****FTP**

- configuring access 2-17
- connecting via 2-18
- defaults 2-8

G**Gigabit Ethernet**

- autonegotiation required 2-4
- switch requirement 2-6

Global hot spares 4-4**Groups**

- creating local 3-6
- file-level access for 6-18
- joining NIS domain 3-8
- quotas for NIS 4-17
- setting share-level access for 6-14

GuardianOS

- specifications 1-1

GuardianOS Ports D-1**H****Half-duplex 2-7****Hardware Components, purchasing new C-1****Hidden Shares 6-10****Host File Editor A-5****Hot spares 4-4****Hot swapping**

- automatic incorporation 4-4, 11-14
- disk drives 4-30

HP Open View 11-6**HTTPS/HTTP**

- configuring 2-18

I**Initial Setup Wizard 1-8****Internal temperature, e-mail notification of 11-5****IP address**

- setting 2-4
- using SSM to discover 1-8

IPP port number, Linux 2-22**IPP port number, Windows 2-22****iSCSI Disks 5-1**

- configuring iSNS 5-18

creating 5-12
iSNS 5-18

K

Kerberos 2-10
Knowledge Base C-1

L

LEDs, understanding C-1
Load balancing
 configuring server for 2-4
 restrictions on 2-3
Local hot spares 4-4
Login
 to antivirus GUI 9-2

M

Mac OS 10-1
 Chooser 2-16
 configuring client access 2-15
 error messages for C-12
 launching SnapServer Manager on 1-6
 Unicode 10-3
menu flow indicator PR-iv
Multihomed configurations 2-5

N

NetVault
 backing up NVDB directory 8-4
 restoring NVDB directory 8-5
Network bonding
 cabling requirements for 2-5
 defaults & options 2-3
Networking
 problems with access C-11
 reset to factory defaults C-9
NFS access
 and read-only share access 2-13
 and share-level permissions 6-15
 configuring 2-13
 to hidden shares 6-10
NFS exports file 6-1

O

Orphaned disk drives 4-29
os_private address 8-2
Overland technical support PR-iii

P

Password
 default for admin account 3-2
Paths
 connecting via web browser 2-19
 for backing up snapshots 7-5
 for distributing antivirus updates 9-7, 9-8
 for restoring a “cured” file 9-9
 to SnapDRImage 8-2
 to volume disaster recovery files 8-2
Permissions
 share- and file-level interaction 6-14
 file-level
 default behavior 6-18
 share-level
 setting 6-14
Phone home support C-15
Print Server
 Adding the Network Printer to the Client 2-22
 Adding the Print Server to a Client Machine 2-22
 Canceling Print Jobs 2-23
 Configuring the Printer on the SnapServer 2-21
 Deleting a Printer from the SnapServer 2-23
 IPP port number, Linux 2-22
 IPP port number, Windows 2-22
 list of supported printers 2-21
 Monitoring Print Jobs 2-23
 Pausing the Printer 2-23
product documentation PR-iii

Q

Quotas
 assigning and managing 4-17
 backing up configuration 8-2

R

RAID
 types defined 1-6
 assessing status of 4-5

- choosing 4-3
- creating 4-7
- creating & monitoring 4-5
- effect of deleting on antivirus software 9-2
- grouping 4-10
- scrubber 4-4, 11-14
- using wizard 4-9
- Reboot**
 - from Web Management Utility 11-12
 - setting up alert for 11-5
- rejoining servers to a Windows Domain** 8-3
- Reset factory defaults** C-9
- Reset Options** C-9
- Resynchronization, setting alert for completion of** 11-5

S

Security

- local authentication 3-2
- resetting default ACLs for volumes C-10
- Windows ACLs 6-18
- Security Model**
 - resetting for volume 6-4
- Server and volume settings, backing up** 8-1
- Server name, discovering** 1-7
- Server registration**
 - via Web Management Utility 11-9
- Setup wizard, see Initial Setup Wizard**
- Shared-hub configurations** 2-6
- Shares**
 - backing up configuration 8-2
 - creating 6-11
 - hidden 6-10
 - snapshot shares 6-10
- Simple Network Management Protocol. See SNMP**
- Single-subnet configuration** 2-5
- SnapCLI** B-1
 - Running Scripts B-8
 - Syntax B-1
- SnapDRImage** 8-2
- SnapServer Manager** 1-5
- SnapServers**
 - backup and restore path A-5
 - configuring e-mail notification of server events 11-4
 - connecting to 1-7

- setting e-mail alerts for 11-5
- Snapshot shares** 7-4
- Snapshots**
 - autobackup of volume settings 8-2
 - coordinating with backup jobs 7-4
 - creating 7-6
 - estimating storage requirements for 7-2
 - excluding from antivirus scans 9-3
 - excluding iSCSI Disks from 5-11
 - ways to adjust pool size 7-3
- Snapshots and Unicode** 10-3
- SnapTree** 6-4
- SNMP** 11-6
 - software update PR-v
- Specifications, GuardianOS** 1-1
- Speed/duplex options** 2-4
- SSH** A-5
- Standalone** 2-5

T

TCP/IP

- configuration guidelines 2-5
- configuring 2-4
- options 2-3
- technical support** PR-iii
- Tivoli NetView** 11-6
- Troubleshooting** C-1
- typographical conventions** PR-iv

U

- Unicode** 10-1, 10-3
 - AFP protocols 10-1
 - Mac OS 10-3
 - SMB protocols 10-1
 - UNIX 10-3
 - Windows Client 10-3
- UNIX**
 - Unicode 10-3
- UPS** 1-11
 - configuring 11-7
 - enabling support for 11-7
- Users**
 - creating local 3-4
 - file-level access for 6-18
 - quotas for 4-17

setting share-level access for **6-14**

V

Volumes

- and antivirus software **4-12**
- and NetVault database directory **4-12**
- as distinct from Macintosh volume **2-16**
- assessing status of **4-14, 5-13**
- backing up configuration **8-1**
- capacity reached alert **11-5**
- creating **4-15**
- effect of deleting on antivirus software **9-2**
- expanding capacity of **4-12**
- management tools **4-13**
- tracking usage of **11-13**
- using quotas to control usage **4-17**

W

Wake-on-LAN Support **1-11**

Web Server **2-19**

Web View **2-21**

WebRoot **2-19**

Windows **10-3**

- connecting from a client **2-10**
- enabling guest account access **2-11**
- file and folder name support **2-9**
- guest account access **2-10**
- issues with PDC **C-12**
- name resolution server support **2-9**
- security, joining
 - active directory domain **2-12**
- see also *Active Directory*
- see also *Authentication*

Windows Client

- Unicode **10-3**