



Overland  
Storage

# SnapScale®

## *Administrator's Guide*

For Appliances Running  
RAINcloudOS® Version 4.1



December 2014  
10400455-003



©2008-14 Overland Storage, Inc. All rights reserved.

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

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

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

BizNAS™ and QuadPak™ are trademarks of Tandberg Data, Inc.

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

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

#### PROPRIETARY NOTICE

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

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

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

FW 4.1.110

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](http://www.overlandstorage.com)

## Audience and Purpose

This guide is intended for system and network administrators charged with installing and maintaining a SnapScale cluster running RAINcloudOS version 4.1 on their network. It provides information on the installation, configuration, security, and maintenance of the SnapScale cluster and nodes. It is assumed that the administrator is familiar with the basic concepts and tasks of multi-platform network administration.

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

- The RAINcloudOS 4.1 Web Management Interface
- SnapServer Manager (SSM)
- CA Antivirus software

RAINcloudOS version 4.1 comes preinstalled on all new SnapScale X-Series nodes. It can also be upgraded from a previously installed version of RAINcloudOS version 3.0 or later.

## Product Documentation & Software Updates

SnapScale product documentation and additional literature are available online, along with the latest release of the RAINcloudOS version 4.1.

Point your browser to:

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

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

## Overland Technical Support

For help configuring and using your SnapScale cluster, email our technical support staff at:

[techsupport@overlandstorage.com](mailto:techsupport@overlandstorage.com).

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

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

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

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

## Japanese Voluntary Control Council for Interference (VCCI)

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

VCCI- A

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

## Conventions

This document exercises several alerts and typographical conventions.

Convention	Description & Usage
 <b>WARNING</b>	A <i>Warning</i> contains information concerning personal safety. Failure to follow directions in the Warning could result in bodily harm or death.
<b>WARNUNG</b>	Eine <i>Warnung</i> enthält Informationen zur persönlichen Sicherheit. Das Nichtbeachten der Anweisungen in der Warnung kann zu Verletzungen oder zum Tod führen.
<b>AVERTISSEMENT</b>	Un <i>avertissement</i> contient des informations relatives à la sécurité personnelle. Ignorer les instructions dans l'avertissement peut entraîner des lésions corporelles ou la mort.
 <b>CAUTION</b>	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.
 <b>IMPORTANT</b>	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.
<b>Item_name</b>	Words in this special boldface font indicate the names of buttons or page names found in the Web Management Interface.
Ctrl-Alt-r	This type of format details the keys you press simultaneously. In this example, hold down the <b>Ctrl</b> and <b>Alt</b> keys and press the <b>r</b> key.
<b>NOTE</b>	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, <b>Setup &gt; User &gt; Password</b> indicates that you should click the <b>Setup</b> tab, then the <b>User</b> secondary tab, and finally the <b>Password</b> button to accomplish a task.
<i>Courier Italic</i>	A variable for which you must substitute a value.
<b>Courier Bold</b>	Commands you enter in a command-line interface (CLI).

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

# Contents

## Preface

Conventions .....	4
-------------------	---

## Chapter 1: Overview

SnapScale Conventions .....	11
SnapScale Node Requirements .....	12
RAINcloudOS Specifications .....	13
What's New in RAINcloudOS 4.1 .....	14
RAINcloudOS 4.0 Features .....	15
Using SnapServer Manager with SnapScale .....	16
SnapServer Manager Installation .....	16
Client and Storage Networks .....	16
Node/Switch Cabling Example .....	17
Node Port Configurations .....	17
X2 Node Configurations .....	17
X4 Node Configurations .....	19

## Chapter 2: Initial Setup and Configuration

Connecting for the First Time .....	21
Connect Using the Node Name .....	21
Connect Using SSM .....	22
Setup a New SnapScale Cluster (via Wizard) .....	23
Step 1 – Select SnapScale Nodes .....	24
Step 2 – Client Network Configuration Overview .....	25
Step 3 – Choose Client Network Static TCP/IP Settings .....	26
Step 4 – Configure Node Static IP Addresses .....	27
Step 5 – Basic SnapScale Properties .....	29
Step 6 – Set Date and Time .....	30
Step 7 – Summary Verification & Cluster Creation .....	31
Join an Existing SnapScale Cluster (via Wizard) .....	36
Web Management Interface .....	37
Alert Messages .....	40
Site Map .....	41
Contact, Hardware & Software Information .....	41

## Chapter 3: SnapScale Settings

SnapScale Properties .....	43
Date/Time .....	44
Secure Shell .....	45
Disable SSH .....	45

Connect to the CLI using SSH .....	46
UPS Protection .....	46
Edit UPS Properties .....	47
Configure UPS Protection .....	48
Add Network UPS Device .....	48
Change Network UPS Device .....	49
Delete Network UPS Device .....	49

## Chapter 4: Network Settings

Network Information .....	51
Client Network Information .....	51
Storage Network Information .....	53
TCP/IP Networking .....	54
Bonding Options .....	56
Guidelines in TCP/IP Configuration .....	56
Edit Storage Network Properties .....	57
Utility IP Address .....	59
Windows/SMB Networking .....	61
Support for Windows/SMB Networking .....	62
Support for Windows Network Authentication .....	62
Connect from a Windows Client .....	63
Connect a Mac OS X Client Using SMB .....	64
Configure Windows/SMB Networking .....	64
NFS Access .....	67
Assigning Share Access to NFS Users .....	68
Enable NFS Access to the Cluster .....	68
Configure NFSv4 Access .....	69
LDAP/NIS Domains .....	71
LDAP vs. NIS Overview .....	71
Configuring LDAP .....	71
Configuring NIS .....	73
FTP/FTPS Access .....	74
Supported FTP Clients .....	74
Configure FTP/FTPS Access .....	74
Connect via FTP/FTPS .....	75
SNMP Configuration .....	75
Default Traps .....	76
Supported Network Manager Applications and MIBs .....	76
Configure SNMP .....	77
Web Access .....	78
Configuring HTTP/HTTPS .....	78
Configure the SnapScale Cluster as a Simple Web Server Using Web Root .....	79
iSNS Configuration .....	82
Configure the iSNS Settings .....	82
Update iSNS Registration .....	83

## Chapter 5: Storage Options

Peer Sets .....	85
Peer Sets and Recovery .....	86
Peer Set Utilization .....	87
Peer Set Basics .....	87

Peer Sets Page .....	88
Spare Disks Page .....	90
Spare Distributor .....	91
Data Balancer .....	93
Volumes .....	96
Creating Volumes .....	96
Edit Volume Properties .....	99
Deleting Volumes .....	100
Quotas .....	101
Default Quotas .....	101
Quotas for Volume Page .....	103
Add Quota Wizard .....	105
Editing or Removing Quotas .....	108
Snapshots .....	109
Creating Snapshots .....	110
Adjusting Snapshot Space .....	113
Accessing Snapshots .....	114
Scheduling Snapshots .....	114
Edit Snapshot Properties .....	115
iSCSI Disks .....	115
Configuring iSCSI Initiators .....	115
iSCSI Configuration on the SnapScale .....	116
Create iSCSI Disks .....	119
Edit iSCSI Disk Properties .....	121
Delete an iSCSI Disk .....	121
Configuring VSS/VDS for iSCSI Disks .....	122
Nodes .....	125
Edit Node Properties .....	126
Node Drives .....	127
Adding Nodes .....	127
Removing Nodes .....	133
Node Identification .....	134
Disks .....	135
Replacing Drives .....	136
Adding Drives .....	136

## Chapter 6: Security Options

Security Considerations .....	142
Guidelines for Local Authentication .....	142
User and Group ID Assignments .....	143
Security Guides .....	143
Security Guide for Windows Active Directory .....	144
Security Guide for Entire Volume Access .....	145
Security Guide for Folder Access on Volume .....	146
Shares .....	147
Create Shares .....	148
Edit Share Properties .....	150
Delete Shares .....	152
Configuring Share Access .....	153
Local Users .....	161
Create a User .....	162

Edit User Properties .....	163
Local User Password Policies .....	165
Assign User to Group .....	166
Delete Local User .....	167
Local Groups .....	168
Create New Group .....	169
Edit Group Properties .....	170
Specify Users in Group .....	171
Delete Group .....	172
Security Models .....	173
Managing Volume Security Models .....	173
ID Mapping .....	175
Add Mapping .....	175
Change Mapping .....	178
Auto Mapping .....	181
Remove Mappings .....	182
Remove Missing ID Mappings .....	185
Filesystem Updates .....	187
Home Directories .....	188
Configure Home Directories .....	189

## Chapter 7: System Monitoring

System Status .....	192
Activity .....	193
Active Users .....	194
Open Files .....	195
Network Monitor .....	196
Event Log .....	201
Filter the Log .....	201
Protocol Manager .....	202
SnapScale Settings .....	203
Tape .....	205

## Chapter 8: Maintenance

Shutdown and Restart .....	207
Manually Powering SnapScale Nodes On and Off .....	207
Data Import .....	208
Setting Up a Data Import Job .....	209
Stopping an Import Job .....	211
Recreating an Import Job .....	211
Preserving Permissions .....	212
OS Update .....	213
Update the RAINcloudOS .....	213
Online/Rolling Updates Notes .....	216
Update Notification Option .....	218
Last OS Update .....	221
Support .....	221
Registering Your Cluster .....	222
Maintenance Tools .....	225
Email Notification .....	226
Host File Editor .....	228

Delete SnapScale Cluster .....	229
--------------------------------	-----

## Chapter 9: Misc. Options

Home Pages .....	232
Home Page .....	232
Administration Page .....	234
SnapExtensions .....	236
BitTorrent Sync .....	236
Snap EDR .....	239
Snap Finder .....	239
Edit Snap Finder Properties .....	241
Change Password .....	242
Changing Your Password .....	242
Management Interface Settings .....	243

## Appendix A: Backup Solutions

Backup and Replication Solutions .....	245
Snap Enterprise Data Replicator .....	245
Snap EDR Usage .....	246
Configuring Snap EDR .....	246
Scheduling Jobs in Snap EDR .....	247
Backup via SMB or NFS .....	247
Off-the-Shelf Backup Solutions .....	247
Utility IP Address .....	247

## Appendix B: Security and Access

Security Model Rules .....	248
Security Model Management .....	249
Special Share Options .....	249
Hiding Shares .....	250
Share Level Permissions .....	250
Where to Place Shares .....	250
File and Share Access .....	250
NFS Share Access .....	251
Snapshot Access .....	251
Snapshot Shares and On Demand File Recovery .....	251
Creating a Snapshot Share .....	251
File-level Security .....	252
Security Personalities and Security Models .....	252
Windows ACLs .....	252

## Appendix C: Troubleshooting

LED Indicators .....	254
SnapScale X2 Node LEDs .....	254
X2 LED States .....	255
SnapScale X4 Node LEDs .....	256
Network Reset .....	257
Performing System Resets Without Network Access .....	258
Networking Issues .....	258
Miscellaneous Issues .....	260

Phone Home Support .....	261
--------------------------	-----

## Appendix D: RAINcloudOS Ports

## Appendix E: Command Line Interface

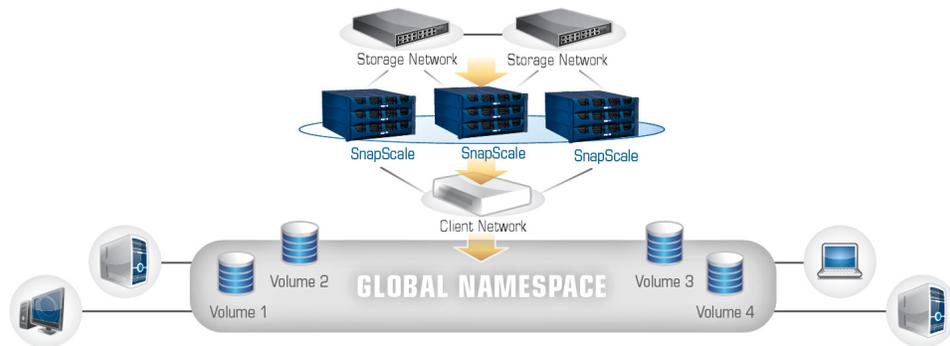
SnapCLI Syntax .....	267
SnapCLI Procedures .....	268
SnapCLI Commands .....	269
Scripts in SnapCLI .....	273
Running a SnapCLI Script .....	273
Sample Script .....	273

## Master Glossary & Acronym List

## Index

SnapScale is a flexible, scalable, low-maintenance network-attached storage cluster composed of a redundant array of independent nodes running RAINcloudOS (ROS). This guide applies to SnapScale nodes running RAINcloudOS version 4.1.

Offering user-selectable levels of data redundancy, SnapScale uses File-level Striping to write data across multiple nodes and drives simultaneously for instant protection and high availability. With a SnapScale cluster, volumes can be configured, created, provisioned, and grown on demand. Special features such as Data Balancer redistributes files to optimize performance and Spare Distributor evenly distributes spare drives across nodes. Files can be accessed either through NFS or CIFS/SMB protocols. SnapScale Flexible Volumes automatically adjust capacity so they only occupy as much space as their data requires.



**Topics in Overview:**

- [SnapScale Conventions](#)
- [SnapScale Node Requirements](#)
- [RAINcloudOS Specifications](#)
- [What's New in RAINcloudOS 4.1](#)
- [RAINcloudOS 4.0 Features](#)
- [Using SnapServer Manager with SnapScale](#)
- [Client and Storage Networks](#)
- [Node Port Configurations](#)

## SnapScale Conventions

The SnapScale cluster supports three or more nodes hosting redundant sets of data for data protection. An Administrator can configure, add, or remove nodes on demand to change storage requirements. The overall storage system is able to easily grow from three nodes to meet your needs.

Peer sets are created using two or three drives (based on redundancy choices) located on different nodes. Each peer set member has the same data and metadata as its peers.

There are three different states for SnapScale nodes:

- **Uninitialized node** – an independent node that has not yet been joined to a SnapScale cluster.
- **SnapScale node** – a healthy node that is a member of a fully-configured SnapScale cluster. Both 2U nodes with up to 12 drives and 4U nodes with 36 drives are available.
- **Management node** – a SnapScale node with special duties involved in managing the cluster. The Management node is selected automatically by the RAINcloudOS when the cluster boots. Should that management node fail, another currently available node is automatically chosen to become the new Management node. This Management node also hosts peer sets with metadata and data just like all other SnapScale nodes.

Other key concepts include:

- **Management IP** – the IP address through which the administrator accesses the Web Management Interface of the current Management node.
- **Peer set** – a set of two or three disks (each on a separate node) that have mirrored data for redundancy.
- **Cluster Name** – the name visible to network clients and used to connect to the cluster (similar to a server name), and resolvable to node IP addresses via round robin DNS.
- **Cluster Management Name** – the hostname resolvable to the Management IP for Web Management Interface access or Snap EDR configuration.
- **Data Replication Count** – an administrator-specified, cluster-wide count of the number of mirrored copies of data within the cluster. The Data Replication Count can be either “2” or “3” and determines the number of drives in a peer set.

A SnapScale cluster consists of two separate networks:

- **Client Network** – used exclusively for client access. Clients can connect to any node to access data anywhere on the cluster.
- **Storage Network** – an isolated network used exclusively by the cluster for inter-node communications. This includes:
  - Heartbeat (node health/presence) sensing.
  - Synchronization of peer set members.
  - Data transfer between nodes to facilitate clients reading from and writing to files.

## SnapScale Node Requirements

The following table details the basic requirements for cluster nodes:

Requirement	Detailed Description
<b>Minimum number of nodes</b>	A SnapScale cluster must have a minimum of three (3) nodes to operate normally.
<b>No expansion units</b>	A SnapScale node cannot have any expansion units attached to it.
<b>Minimum number of disks per node</b>	Each node must have a minimum of four disks. Additional disks can be added as needed.

Requirement	Detailed Description
<b>Maximum size of file on cluster</b>	While the system reports total free space across the entire cluster, the maximum file size at any given time is dictated by free space on the least-utilized peer set. This is reported in the Web Management Interface.
<b>Common Storage network</b>	To form or join a SnapScale cluster, each Uninitialized node must be connected to the same Storage network as the other nodes.
<b>Storage network links</b>	To form or join a SnapScale cluster each Uninitialized node must have connectivity (active link) on both Storage network ports.
<b>Storage network usage</b>	Only a single cluster can use a given Storage network.
<b>Client network separate from Storage network</b>	The Client and Storage networks must be on different (independent) networks, and the Storage network must be isolated from all other networks.
<b>Nodes must be running same RAINcloudOS version</b>	To form a SnapScale cluster, all nodes must be running the same version of RAINcloudOS.  To join an already configured SnapScale cluster, an Uninitialized node must have the same version of RAINcloudOS as the other SnapScale nodes: <ul style="list-style-type: none"> <li>• If the Uninitialized node has an <b>older</b> version of the RAINcloudOS, the Uninitialized node must be upgraded to the later version.</li> <li>• If the Uninitialized node has a <b>newer</b> version of the RAINcloudOS, then all SnapScale nodes must be upgraded to the later version. (The node can be reinstalled with a version matching the cluster if the hardware supports it.)</li> </ul>
<b>Adding nodes</b>	When adding nodes to an existing cluster, the number of nodes added at one time should be at least the same number as the Data Replication Count. This ensures the new nodes and cluster are efficiently utilizing increased storage space.
<b>Disk requirements</b>	All disks in the cluster must be the same type of disk (such as SAS) and same rotational speed.

## RAINcloudOS Specifications

These specifications apply to SnapScale nodes running RAINcloudOS version 4.1:

Feature	Specification
<b>Network Transport Protocols</b>	<ul style="list-style-type: none"> <li>• TCP/IP (Transmission Control Protocol/Internet Protocol)</li> <li>• UDP/IP (User Datagram Protocol/Internet Protocol)</li> </ul>
<b>Network File Protocols</b>	<ul style="list-style-type: none"> <li>• Microsoft Networking (CIFS/SMB1/SMB2)</li> <li>• Unix Network Filesystem (NFS) 2.0/3.0/4.0</li> <li>• Hypertext Transfer Protocol (HTTP/HTTPS)</li> </ul>

Feature	Specification
<b>Network Security</b>	<ul style="list-style-type: none"> <li>• Microsoft Active Directory Service (ADS) (member server)</li> <li>• Unix Network Information Service (NIS) user/group UID/GID translation</li> <li>• LDAP user/group UID/GID translation</li> <li>• File and Folder Access Control List (ACL) Security for Users and Groups</li> <li>• Secure Sockets Layer (SSL v2/3) 128-bit Encryption</li> <li>• Target Challenge Handshake Authentication Protocol (CHAP) for iSCSI</li> <li>• SMTP Authentication and support for email encryption (STARTTLS and TLS/SSL encryption protocols)</li> </ul>
<b>Network Client Types</b>	<ul style="list-style-type: none"> <li>• Microsoft Windows 2003/2003 R2/2008 SP2/2008 R2 /XP SP3/ Vista SP2/7/8/2012</li> <li>• Mac OS X 10.5/10.6/10.7/10.8/10.9</li> <li>• Sun Solaris 10 and 11</li> <li>• HP-UX 11</li> <li>• AIX 5.3/6</li> <li>• Red Hat Enterprise Linux (RHEL) 4.x/5.x/6.x</li> <li>• Novell SuSE Linux Enterprise Server (SLES) 10.x/11.x</li> </ul>
<b>Data Protection</b>	<ul style="list-style-type: none"> <li>• Snapshots for immediate or scheduled point-in-time images of the cluster filesystem</li> <li>• Support for local backup with Symantec NetBackup/Backup Exec Remote Media Server for Linux</li> <li>• Support for network backup with Symantec NetBackup/Backup Exec, CA ARCserve, or EMC NetWorker</li> <li>• 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</li> </ul>
<b>DHCP Support</b>	<ul style="list-style-type: none"> <li>• Supports Dynamic Host Configuration Protocol (DHCP) for automatic assignment of IP addresses</li> </ul>
<b>System Management</b>	<ul style="list-style-type: none"> <li>• Browser-based administration tool called the Web Management Interface</li> <li>• SnapServer Manager utility (platform independent)</li> <li>• Read-only CLI support</li> <li>• SNMP (MIB II and Host Resource MIB)</li> <li>• User disk quotas for Windows, Unix/Linux, FTP/FTPS</li> <li>• Group disk quotas for Unix/Linux</li> <li>• Environmental monitoring</li> <li>• Email event notification</li> <li>• Data importation (migration)</li> </ul>

## What's New in RAINcloudOS 4.1

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

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

Feature	New Functionality
<b>NFSv4 Support</b>	In addition to NFSv3 default support, the option now exists to enable NFSv4.
<b>Online/Rolling Updates</b>	Added a new OS update process that allows the cluster to still be online and available during the entire update process. However, performance will be reduced and some administrative features will not be available.
<b>Snapshots Performance Optimization</b>	File write performance is significantly improved when one or more snapshots exist for a volume.
<b>LDAP Integration</b>	Lightweight Directory Access Protocol (LDAP) can be used to look up user/group names and UIDs/GIDs for quota assignment, ID mapping, and home directories.
<b>BitTorrent Sync</b>	BitTorrent Sync can be used to replicate data between the SnapScale cluster and other servers or workstations.
<b>Windows Security Model</b>	Volumes can be configured with a security model that permits only Windows personality file system permissions (ACLs).
<b>Performance Monitor</b>	Completely redesigned with a new interface and functionality including the ability to keep more than just one hour of data enabling the user to review network performance information going back several weeks or months.

## RAINcloudOS 4.0 Features

**NOTE:** For details and descriptions of all the new features of the operating system, see the [Product Release Notes on the Overland SnapScale website](#).

With the release of RAINcloudOS 4.0, the following features and functionality are now available:

Feature	New Functionality
<b>iSCSI Support</b>	SnapScale can now create and host iSCSI disk targets on the cluster file system. These iSCSI disks can register with an iSNS server, and can also be managed by Windows VSS/VDS.
<b>SMB2, FTP/FTPS, and SNMP Support Added</b>	SMB2, FTP/FTPS, and SNMP are all now supported in RAINcloudOS.
<b>Improved Network Monitoring</b>	The <b>Network Monitor</b> page provides additional information including high-water marks, network activity for the whole cluster, and clearer labels.
<b>Added User/Group Quotas per Volume</b>	Storage consumption and file count quotas can now be configured for users and NIS groups per volume.
<b>SMB 2.0/2.1</b>	The server supports Server Message Block (SMB) version 2.0/2.1 for improved compatibility and performance for later Windows OS versions.

Feature	New Functionality
Data Balancer & Spare Distributor Improved	<b>Data Balancer</b> (formerly <b>Capacity Balancer</b> ) redistributes files to optimize performance. <b>Spare Distributor</b> (formerly the <b>Spare Disk Balancer</b> ) evenly distributes spare drives across nodes. Both have been improved for faster results.

## Using SnapServer Manager with SnapScale

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

The screenshot shows the SnapServer Manager application window titled "SnapServer Manager - <All Servers>". The interface includes a menu bar (File, Edit, Administration, Help), a sidebar for "Servers & Groups" with a tree view showing "All Servers" and two Snap servers, and a main table of server details. A "Server Group" label is overlaid on the sidebar, a "Server List" label is overlaid on the table, and a "Status Bar" label is overlaid on the bottom status area.

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

At the bottom of the window, the status bar displays: "No operations are currently running or scheduled." and "Online/Offline: 115/18".

### SnapServer Manager Installation

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

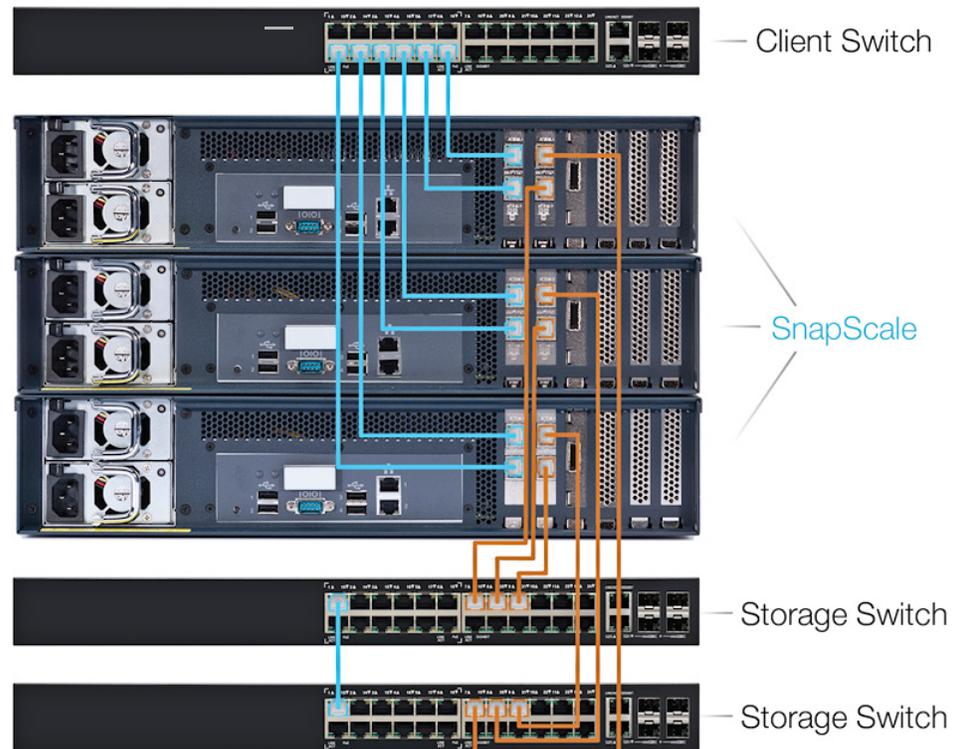
Refer to the *SnapServer Manager User Guide* for details on discovering and configuring SnapScale clusters.

## Client and Storage Networks

SnapScale requires two separate networks to function correctly: A public network (Client) and a private network (Storage). To support failover, two Storage network switches must be connected together (using a 1GbE or 10GbE cable between the switches). Each of the two Storage network ports on the node need to be connected to a different Storage switch.

## Node/Switch Cabling Example

The following example shows three dual 10GbE card X2 nodes and how to connect them to the network switches. The cables used to connect to the Client side of the network (blue) originate from the Client 10GbE card in slot 1. Two cables are used to connect both ports of each node to the Client switch.



The cables used to connect to the Storage side of the network (orange) originate from the Storage 10GbE card in slot 2. For each node, one cable is used to connect a one Storage port of each node to one of the two Storage switches used for failover.

For connections between 10GbE cards and 10GbE switches, use either direct-attached copper cables or fibre cables with SFP+ modules pre-installed in the card and switch ports.

**NOTE:** If using fibre cables, you must use Overland-approved SFP+ modules. With the cluster powered OFF, insert the modules into the card and switch ports. Connect the fibre cable between the two SFP+ modules and restore power to the cluster.

## Node Port Configurations

Both the X2 and X4 nodes come in three different configurations: 1GbE ports (both Client and Storage ports), a single 10GbE card (with 1GbE Client ports), and dual 10GbE cards.

**NOTE:** If desired, optional 10GbE cards can be added later to upgrade the node.

### X2 Node Configurations

**Basic 1GbE.** At the rear of the X2 node, the 1GbE ports connected directly to the motherboard are configured to access the Client and Storage networks.



Configuration	Node GbE Ports	Network Switch
Basic 1GbE X2	Ports 1 & 2	Client (public)
	Slot 2 1GbE Card (ports 3 & 4)	Storage (private)

**Single 10GbE.** The single-card 10GbE configuration uses the two 1GbE ports for the Client connection and the two 10GbE ports on the card for the Storage connections.



Configuration	Node GbE Ports	Network Switch
Single 10GbE X2	Ports 1 & 2	Client (public)
	Slot 2 10GbE Card (ports 3 & 4)	Storage (private)

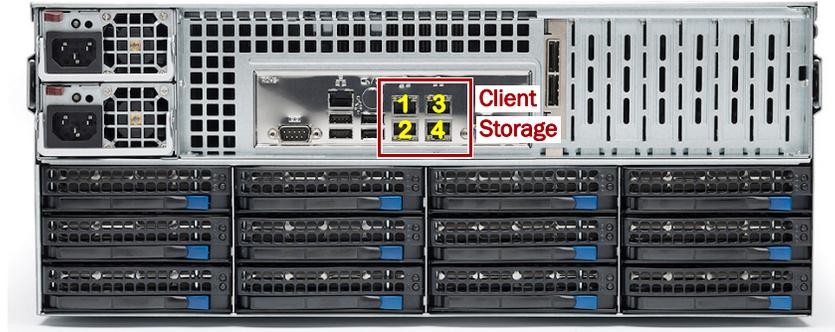
**Dual 10GbE.** The dual-card configuration uses the left 10GbE card ports for the Client connections and the right 10GbE card ports for the Storage connections. The 1GbE ports are not used.



Configuration	Node GbE Ports	Network Switch
Dual 10GbE X2	Slot 1 10GbE Card (ports 1 & 2)	Client (public)
	Slot 2 10GbE Card (ports 3 & 4)	Storage (private)

## X4 Node Configurations

**Basic 1GbE.** At the rear of the 1GbE configuration, there are two sets of 1GbE ports connected directly to the motherboard for connecting to the switches. The top two are for the Client network; bottom two for Storage.



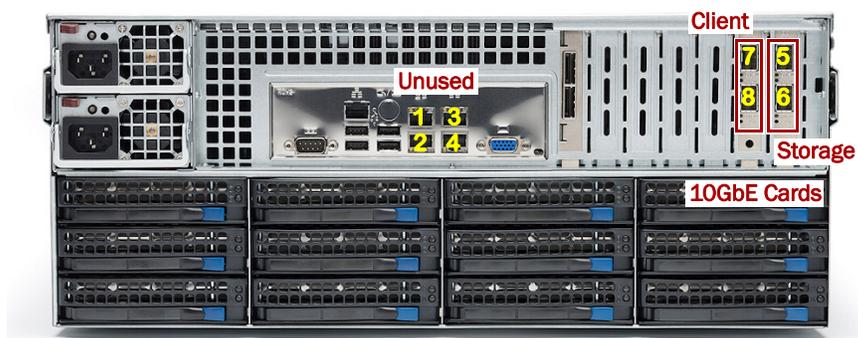
Configuration	Node GbE Ports	Network Switch
Basic 1GbE X4	Ports 1 & 2	Client (public)
	Ports 3 & 4	Storage (private)

**Single 10GbE.** The single-card 10GbE configuration uses the two 1GbE ports for the Client connection and the two 10GbE ports on the card for the Storage connections.



Configuration	Node GbE Ports	Network Switch
Single 10GbE X4	Ports 1, 2, 3, & 4	Client (public)
	Slot 7 10GbE Card (ports 5 & 6)	Storage (private)

**Dual 10GbE.** The dual-card configuration uses the left 10GbE card ports for the Client connections and the right 10GbE card ports for the Storage connections. The 1GbE ports are not used.



Configuration	Node GbE Ports	Network Switch
Dual 10GbE X4	Slot 6 10GbE Card (ports 7 & 8)	Client (public)
	Slot 7 10GbE Card (ports 5 & 6)	Storage (private)

# Initial Setup and Configuration

This section covers the initial setup and configuration of an individual SnapScale node running RAINcloudOS 4.1. It also addresses how to use that node to set up a SnapScale cluster of three or more nodes, or to add the node to an existing SnapScale cluster.

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

## Topics in Setup and Configuration:

- [Connecting for the First Time](#)
- [Setup a New SnapScale Cluster \(via Wizard\)](#)
- [Join an Existing SnapScale Cluster \(via Wizard\)](#)
- [Web Management Interface](#)

## Connecting for the First Time

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

### Connect Using the Node Name

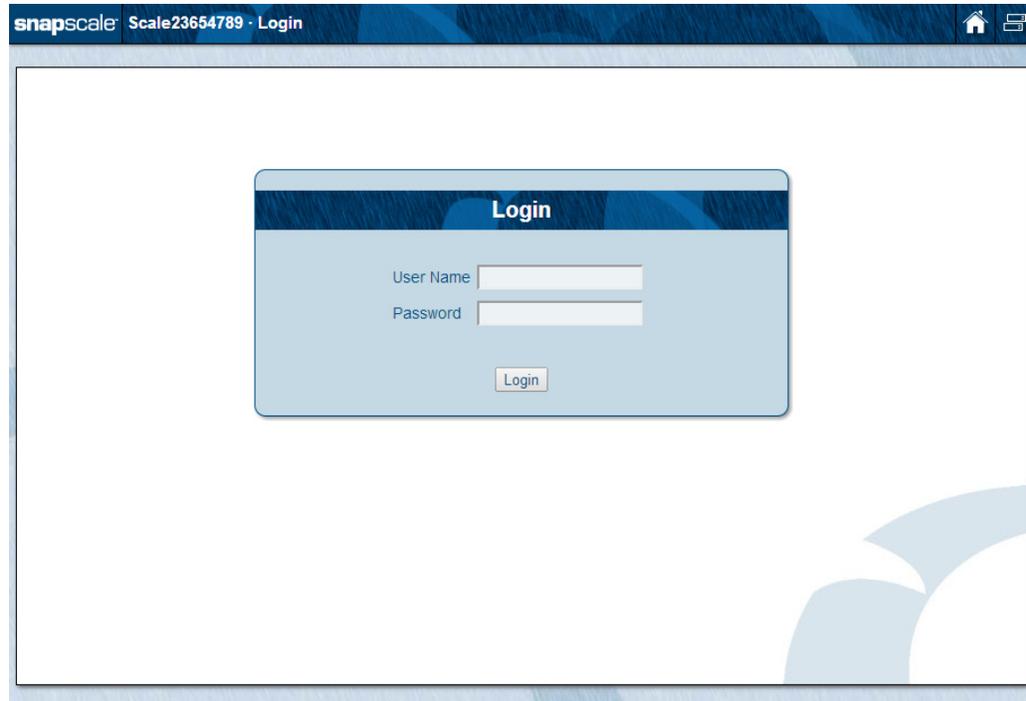
This procedure requires that name resolution services (via DNS or an equivalent service) be operational.

**NOTE:** Any node that is selected to be part of a cluster can be used to create the cluster.

1. Find the **node name** of an Uninitialized node that is to be used to create a new SnapScale cluster.

A SnapScale node name is of the format “Nodennnnnnnnn,” where *nnnnnnnnn* is the node number. The number is a unique, numeric-only string that appears on a label affixed to the bottom of the unit.

2. In a web browser, enter the **URL** to connect to the node.  
For example, enter “http://Nodennnnnnnnn” (using the node name).
3. Press **Enter** to connect to the Web Management Interface.



4. In the login dialog box, enter **admin** as the user name and **admin** as the password (the system defaults), then click **Login**.
5. Complete the **Initial Setup Wizard** to either create a new SnapScale cluster or join an existing cluster.

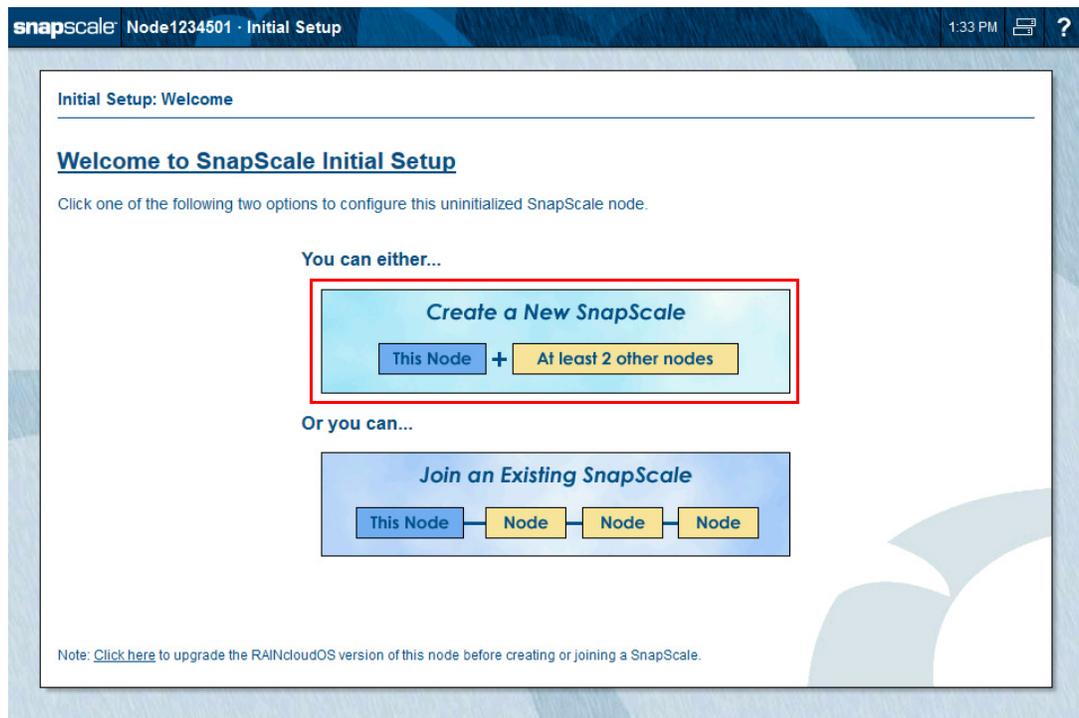
## Connect Using SSM

1. Launch **SnapServer Manager (SSM)**.  
SSM discovers all SnapServers, SnapScale clusters, and Uninitialized nodes on its local network segment and displays their Server names, IP addresses, and other information in the main console. If you do not have a DHCP server, there might be a delay before the node appears on the network.  
**NOTE:** To distinguish multiple SnapServers and SnapScale nodes, you may need to find their default names as explained in [Connect Using the Node Name on page 21](#).
2. If using a DHCP server, proceed to [Step 3](#); otherwise, assign an **IP address** to one of the nodes to be configured in the cluster.  
**NOTE:** Only one node needs to be configured with an IP address in order to create the cluster.
  - a. In SSM, right-click the **node name**.
  - b. Select **Set IP Address**.
  - c. Enter an IP address and a subnet mask, then click **OK**.
3. In SSM, right-click the node name and select **Launch Web Administration**.
4. Log into the **Web Management Interface**.  
In the login dialog box, enter **admin** as the user name and **admin** as the password (the system defaults), then click **OK**.

5. Complete the **Initial Setup Wizard** to either create a new SnapScale cluster or join an existing cluster.

## Setup a New SnapScale Cluster (via Wizard)

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



From the Initial Setup Wizard, you can either use this node to create a new SnapScale cluster by connecting it to two or more other nodes or add the node to an existing cluster (see [Join an Existing SnapScale Cluster \(via Wizard\) on page 36](#)). To start the creation of a new cluster, click the **Create a New SnapScale** box. The Initial Setup Wizard consists of several steps:

**Step 1 – Select SnapScale Nodes** (select the nodes to be included in the cluster)

**Step 2 – Client Network Configuration Overview** (review the Client network information)

**Step 3 – Choose Client Network Static TCP/IP Settings** (choose the static TCP/IP settings for the Client network)

**Step 4 – Configure Node Static IP Addresses** (populate the Static IP addresses for the nodes)

**Step 5 – Basic SnapScale Properties** (enter the basic SnapScale properties)

**Step 6 – Set Date and Time** (set the date and time)

**Step 7 – Summary Verification & Cluster Creation** (verify the settings and create a SnapScale cluster)

**NOTE:** After the cluster is created, you are asked to configure the Administrator's password. It is highly recommended for security to set it to something other than the default setting.

## Step 1 – Select SnapScale Nodes

Select the nodes you want to use from the list of eligible nodes.

**NOTE:** At least three nodes are required to create a SnapScale clustered network. All nodes must have the identical version of RAINcloudOS (ROS) and be on a subnet that does not contain an existing cluster. The Client network interfaces for all the nodes must be located on the same public network subnet, and the Storage network interfaces for all nodes must be located on the same private Storage network subnet. The nodes cannot have any expansion units attached.

Any combination of node types (X4 and X2) can be used to create a cluster.

**Initial Setup: Create SnapScale - Select SnapScale Nodes**

Step 1 | Step 2 | Step 3 | Step 4 | Step 5 | Step 6 | Step 7

Select the nodes below that you want to add to this new SnapScale and click Next. (All eligible nodes are selected by default.)

**Note:** 3 nodes are required as a minimum to create a SnapScale. All nodes in a SnapScale must have identical RAINcloudOS (ROS) versions, and the client network interface for all nodes must be located on the same subnet.

13 Eligible Nodes. (This node: ROS version=4.1.119, IP address=192.168.199.101, Subnet mask=255.255.254.0)

Node	Model	ROS Version	Disks	Add to SnapScale
<input checked="" type="checkbox"/> Node1234501 (This Node)	X2	4.1.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Node1234503	X2	4.1.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Node1234507	X2	4.1.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Node1234510	X4	4.1.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: 1.95 TB 11: 3.91 TB 12: 3.91 TB 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB Rear Disks 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk)	<input checked="" type="checkbox"/>

Back Re-Detect Available Nodes Next

Verify that the boxes in the Add to SnapScale column for the nodes you want to use are checked. Click **Re-Detect Available Nodes** to refresh the list. When ready, click **Next**.

**NOTE:** If you deselect one or more of the detected nodes, when you click **Next** a message page is displayed recommending that you add all the nodes at once.

snapscale Node2413872 · Initial Setup 3:58 PM

### Initial Setup: Create SnapScale - Select SnapScale Nodes

**Important:** You have selected 3 out of 6 eligible nodes for your new SnapScale. Overland recommends that you select all eligible nodes now for addition to the SnapScale, rather than adding them after the SnapScale is created.

Click Keep My Node Selections if you wish to ignore the above concern(s) and continue Initial Setup using the nodes you have selected.

Back Keep My Node Selections

## Step 2 – Client Network Configuration Overview

Review the information about setting up your Client network. Click **Next** to continue.

snapscale Node2413872 · Initial Setup 4:02 PM

### Initial Setup: Create SnapScale - Configure Client Network : Overview

Step 1 Step 2 Step 3 Step 4 Step 5 Step 6 Step 7

Your SnapScale nodes are connected together via a *storage network* that allows both user data and meta-data to be communicated between nodes. This storage network is automatically configured for you as part of this Initial Setup process. Your SnapScale also includes a *client network* which will be used by users to access (read & write) user data stored on the SnapScale. This client network is configured and managed by you using static IP addresses.

Back Next

(Click Next to configure the client network settings for your SnapScale.)

## Step 3 – Choose Client Network Static TCP/IP Settings

Use this step to specify the static TCP/IP settings that will be common to all nodes in the cluster. Then click **Next** to continue to the next page to set the actual node static IP addresses.

The screenshot shows the SnapScale initial setup wizard interface. At the top, the title bar reads "snap scale Node2413872 · Initial Setup" with a timestamp of "4:04 PM" and a help icon. Below the title bar is a progress indicator with seven steps: Step 1, Step 2, Step 3 (highlighted), Step 4, Step 5, Step 6, and Step 7. The main content area is titled "Initial Setup: Create SnapScale - Configure Client Network : Static TCP/IP Settings". It contains the following text: "Specify the static TCP/IP settings that are common to all nodes in the SnapScale. In the next step you will specify a list of static IP addresses to be used by the nodes." Below this is the section "Client network static TCP/IP settings." with the following fields:

Subnet Mask	<input type="text" value="255.255.0.0"/>	
WINS Servers:	<input type="text"/>	(optional)
	<input type="text"/>	(optional)
	<input type="text"/>	(optional)
	<input type="text"/>	(optional)
Default Gateway	<input type="text" value="10.25.1.1"/>	(optional)
DNS Domain Name	<input type="text" value="devnet.myoverland.net"/>	(optional)
Domain Name Servers:	<input type="text" value="10.6.8.34"/>	(optional)
	<input type="text" value="10.6.8.35"/>	(optional)
	<input type="text"/>	(optional)

At the bottom right of the form area, there are two buttons: "Back" and "Next".

## Step 4 – Configure Node Static IP Addresses

A SnapScale cluster requires a set of static IP addresses: one for each node, and one for the Management IP. Use this page to specify the static IP addresses for each of your nodes and for the SnapScale Management IP address used to access the Web Management Interface for this cluster.

The screenshot shows the SnapScale web interface during the initial setup phase. The browser address bar indicates the URL is `Node2413872 · Initial Setup`. The page title is "Initial Setup: Create SnapScale - Configure Client Network : Static Node IP Addresses". A progress bar at the top shows seven steps, with Step 4 being the current active step. The main content area contains the following text:

Specify 4 static IP addresses: one for each of your SnapScale nodes and one for the SnapScale Management IP address (the Management IP address is used to access the Web Management Interface for this SnapScale). These IP addresses must all be located on the same subnet, and they will be automatically assigned to your nodes when the SnapScale is created.

**Optional:** Enter a starting IP address and click the "Populate" button to populate the list below with sequential static IP addresses. You can then review or change the IP addresses before clicking Next.

Starting IP Address   (optional)

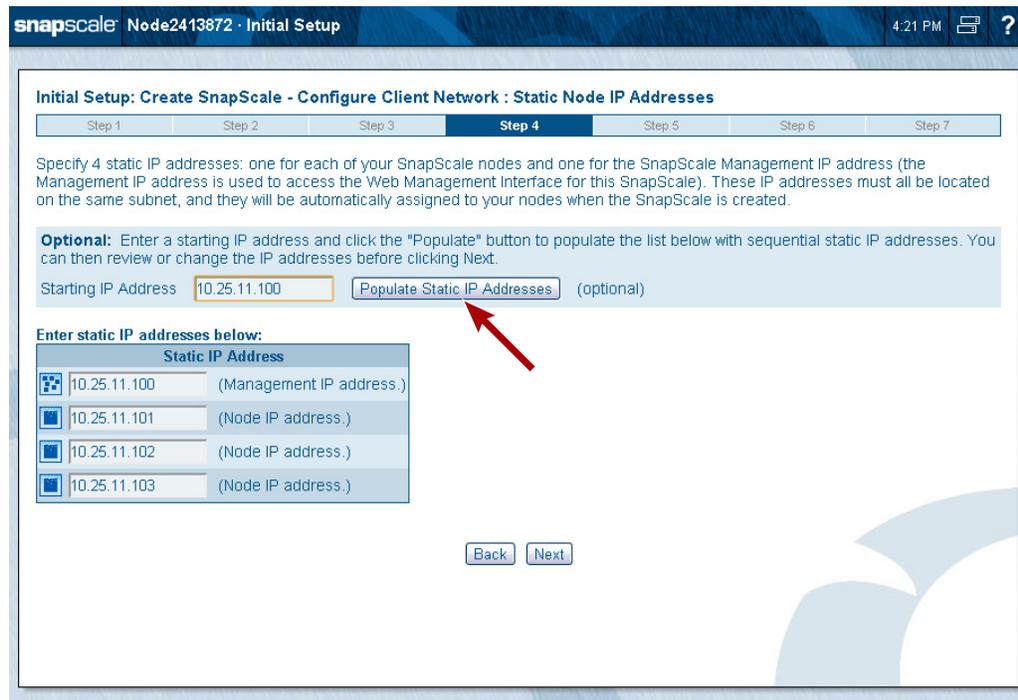
**Enter static IP addresses below:**

Static IP Address	
<input type="text"/>	(Management IP address.)
<input type="text"/>	(Node IP address.)
<input type="text"/>	(Node IP address.)
<input type="text"/>	(Node IP address.)

At the bottom of the form, there are "Back" and "Next" buttons.

These IP addresses must all be located on the same subnet. They are automatically assigned to your nodes when the SnapScale cluster is created.

The **Populate Static IP Addresses** button can be used to automatically enter a sequential list of static IP addresses. Just enter an IP address on the subnet and click **Populate Static IP Addresses**. The fields below it are automatically populated.



The screenshot shows the SnapScale web interface during the initial setup phase. The page title is "Initial Setup: Create SnapScale - Configure Client Network : Static Node IP Addresses". A progress bar at the top indicates that Step 4 is the current step. The main content area contains instructions for specifying 4 static IP addresses. An optional section allows the user to enter a starting IP address (10.25.11.100) and click the "Populate Static IP Addresses" button, which is highlighted with a red arrow. Below this, a table titled "Enter static IP addresses below:" shows the resulting IP addresses: 10.25.11.100 (Management IP address.), 10.25.11.101 (Node IP address.), 10.25.11.102 (Node IP address.), and 10.25.11.103 (Node IP address.). At the bottom of the form, there are "Back" and "Next" buttons.

Initial Setup: Create SnapScale - Configure Client Network : Static Node IP Addresses

Step 1 Step 2 Step 3 **Step 4** Step 5 Step 6 Step 7

Specify 4 static IP addresses: one for each of your SnapScale nodes and one for the SnapScale Management IP address (the Management IP address is used to access the Web Management Interface for this SnapScale). These IP addresses must all be located on the same subnet, and they will be automatically assigned to your nodes when the SnapScale is created.

**Optional:** Enter a starting IP address and click the "Populate" button to populate the list below with sequential static IP addresses. You can then review or change the IP addresses before clicking Next.

Starting IP Address   (optional)

**Enter static IP addresses below:**

Static IP Address	
<input type="text" value="10.25.11.100"/>	(Management IP address.)
<input type="text" value="10.25.11.101"/>	(Node IP address.)
<input type="text" value="10.25.11.102"/>	(Node IP address.)
<input type="text" value="10.25.11.103"/>	(Node IP address.)

Click **Next** to continue.

## Step 5 – Basic SnapScale Properties

Use this step to enter the basic properties for your new SnapScale cluster, then click **Next**.

**Initial Setup: Create SnapScale - Basic SnapScale Properties**

Step 1 Step 2 Step 3 Step 4 **Step 5** Step 6 Step 7

Enter basic properties for the new SnapScale and click Next.

**SnapScale name and description.**  
 SnapScale Name   
 SnapScale Description  (optional)

**Data replication count.**  
 The data replication count specifies how many copies of each data file or folder to maintain. A replication count of 3x offers higher data protection yet uses more disk space.  
 Data Replication Count  (Note: Once the SnapScale is created, the replication count can be changed only from 3x to 2x.)

**Spare disks.**  
 Spare disks specifies the number of available disks in the SnapScale to reserve for spares. A spare disk is used to automatically replace a failed peer set member.  
 Allocate spare disks  
 Spare Disks

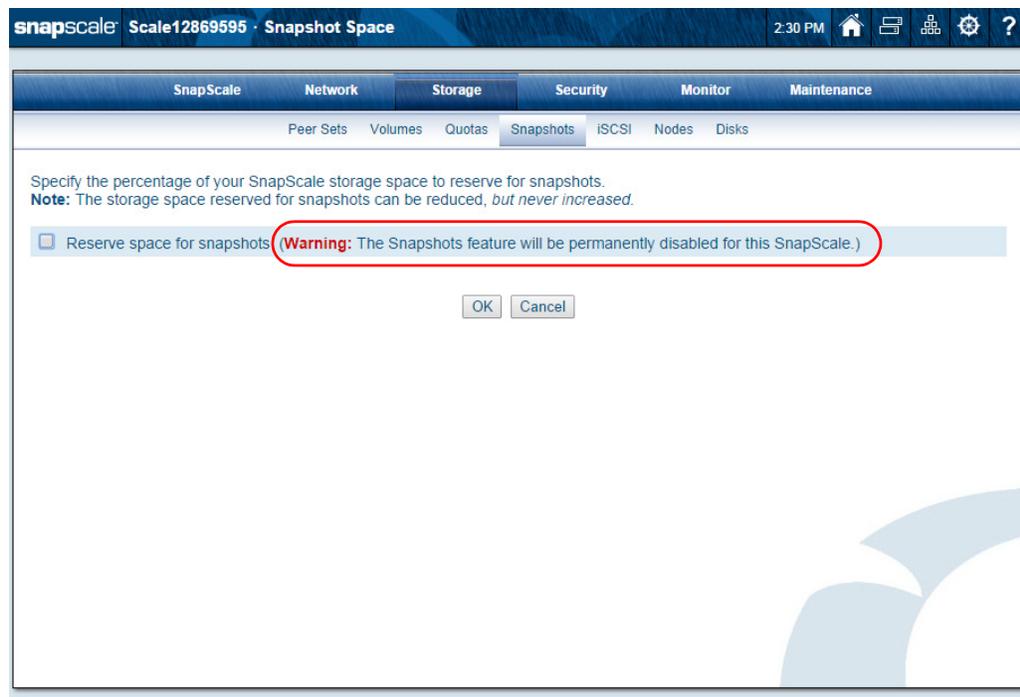
**Snapshots.**  
 If you plan on using snapshots, it is recommended that you reserve at least 20% of your SnapScale storage space for snapshots.  
**Note:** Once the SnapScale is created, the storage space reserved for snapshots can be decreased, *but never increased*.  
 Reserve space for snapshots  
 Percentage of SnapScale storage to reserve for snapshots

This table lists and describes the basic options:

Option	Description
<b>SnapScale Name</b>	Either accept the default name or enter an alphanumeric name up to 15 characters in length. Network clients use this name with round robin DNS name resolution to connect to the cluster.  The default name is "Nodennnnnn" (where nnnnnn is the appliance number of the node used to create the cluster).
<b>SnapScale Description</b>	This optional field provides a place to define the cluster in the overall scheme of your network and better identify the cluster on a LAN.
<b>Data Replication Count</b>	The data replication count establishes the level of data redundancy in the cluster. The setting specifies how many disks are in a peer set and as a result how many copies of each data file or folder to maintain. A count of 3x offers higher data protection but uses more disk space.  Once the cluster is created, the count can only be decreased from 3x to 2x. It cannot be increased from 2x to 3x.

Option	Description
<b>Spare Disks Allocation</b>	<p>Check the box and select the number of spare disks you want to reserve. A spare disk is used to automatically replace a failed Peer Set member.</p> <p>If there are unused drives remaining after allocating the number of spares requested, they are used for other peer sets. If there is an insufficient number of drives left to create a final peer set, the drives are configured as additional spares.</p>
<b>Reserve Space for Snapshots</b>	<p>Check the box and select the percentage of the storage space you want to reserve for snapshots. It is recommended that at least 20% of your SnapScale storage space be set aside for snapshots.</p> <p><b>NOTE:</b> Once the SnapScale cluster is created, the storage space reserved for snapshots can only be decreased. It can never be increased.</p>

**NOTE:** If you uncheck the box for reserving space for snapshots, an alert is displayed to remind you that the feature will be **permanently disabled** for the cluster.



## Step 6 – Set Date and Time

At the **Registration** page:

1. Enter the **four required items** in the appropriate fields.
2. Click **Download Registration File**.  
The information, including all the node data, is incorporated into a CSV file.
3. Depending on your browser settings, make sure that you **save the CSV file** to your local computer.

4. Email the downloaded CSV file to **warranty@overlandstorage.com**.

Use the subject line **SnapScale Registration Request** for the email.

Nodes automatically synchronize time with one another. You can either manually set the date and time to specific values, or you can use NTP (Network Time Protocol) servers to automatically synchronize the date and time. Visit [www.ntp.org](http://www.ntp.org) for a list of public NTP primary and secondary servers, or simply use the default NTP servers below.

**Initial Setup: Create SnapScale - Date & Time**

Step 1 Step 2 Step 3 Step 4 Step 5 **Step 6** Step 7

You can set the date and time of this SnapScale to specific values, or you can use NTP (Network Time Protocol) servers to automatically synchronize the date and time of this SnapScale. Visit [www.ntp.org](http://www.ntp.org) for a list of public NTP primary and secondary servers, or simply use the default NTP servers below.

**Please select one of the following two options:**

Set the date and time of this SnapScale to the following:

Date: 2012 - 12 - 12 (YYYY-MM-DD)

Time: 12 : 12 : 12 PM

Automatically synchronize the date and time of this SnapScale to the following NTP servers.

NTP Servers: 0.pool.ntp.org  
1.pool.ntp.org (optional)

Time Zone: (UTC) Coordinated Universal Time

Back Next

If you intend to join the cluster to a Windows domain, configure the cluster using the manual settings to set the date and time. Otherwise, configure the cluster to synchronize with up to two NTP servers.

**NOTE:** NTP cannot be used if you are joining a Windows Active Directory domain.

Default NTP servers automatically populate the server fields. The Time Zone is set automatically to UTC time but can be changed using the drop-down list.

Click **Next** to continue.

## Step 7 – Summary Verification & Cluster Creation

At this step, review the current settings and go back if you need to make changes.

**NOTE:** Make note of the Management IP address for later use. Also, both the Client and Storage network bond types can be changed after the cluster is created.

**snapScale** Node1234501 - Initial Setup 2:18 PM  ?

---

**Initial Setup: Create SnapScale - Create SnapScale Summary**

Step 1 | Step 2 | Step 3 | Step 4 | Step 5 | Step 6 | **Step 7**

Please review your settings below and click Create New SnapScale to complete the creation of this SnapScale.

**Important Security Note:** You will be asked to change your administrator password after the SnapScale has been successfully created. The administrator password is used to access this Web Management Interface.

**SnapScale settings.**

SnapScale Name	MyCluster
Data Replication Count	2x
Spare Disks	2
Snapshot Space Reserved	20%
Management IP Address	192.168.199.100 (Please make note of this IP address for later use.) 
Subnet Mask	255.255.0.0
Default Gateway	192.168.192.1
Domain Name Servers	192.168.192.22
WINS Servers	-
DNS Domain Name	devnet.myoverland.net
Time Zone	(UTC-08:00) Pacific Time (US & Canada)

**3 SnapScale Nodes.** (Note: The IP addresses displayed below will not necessarily be assigned to their associated node.)

Node 	IP Address	Model	ROS Version	Disks											
<input checked="" type="checkbox"/> Node1234501 (This Node)	192.168.199.101	X2	4.0.119	1: 1.95 TB 	2: 1.95 TB 	3: 3.91 TB 	4: 3.91 TB 	5: 1.95 TB 	6: 1.95 TB 	7: 3.91 TB 	8: 3.91 TB 	9: 1.95 TB 	10: (No Disk) 	11: (No Disk) 	12: (No Disk) 
<input checked="" type="checkbox"/> Node1234503	192.168.199.102	X2	4.0.119	1: 1.95 TB 	2: 1.95 TB 	3: 3.91 TB 	4: 3.91 TB 	5: 1.95 TB 	6: 1.95 TB 	7: 3.91 TB 	8: 3.91 TB 	9: 1.95 TB 	10: (No Disk) 	11: (No Disk) 	12: (No Disk) 
<input checked="" type="checkbox"/> Node1234507	192.168.199.103	X2	4.0.119	1: 1.95 TB 	2: 1.95 TB 	3: 3.91 TB 	4: 3.91 TB 	5: 1.95 TB 	6: 1.95 TB 	7: 3.91 TB 	8: 3.91 TB 	9: 1.95 TB 	10: (No Disk) 	11: (No Disk) 	12: (No Disk) 

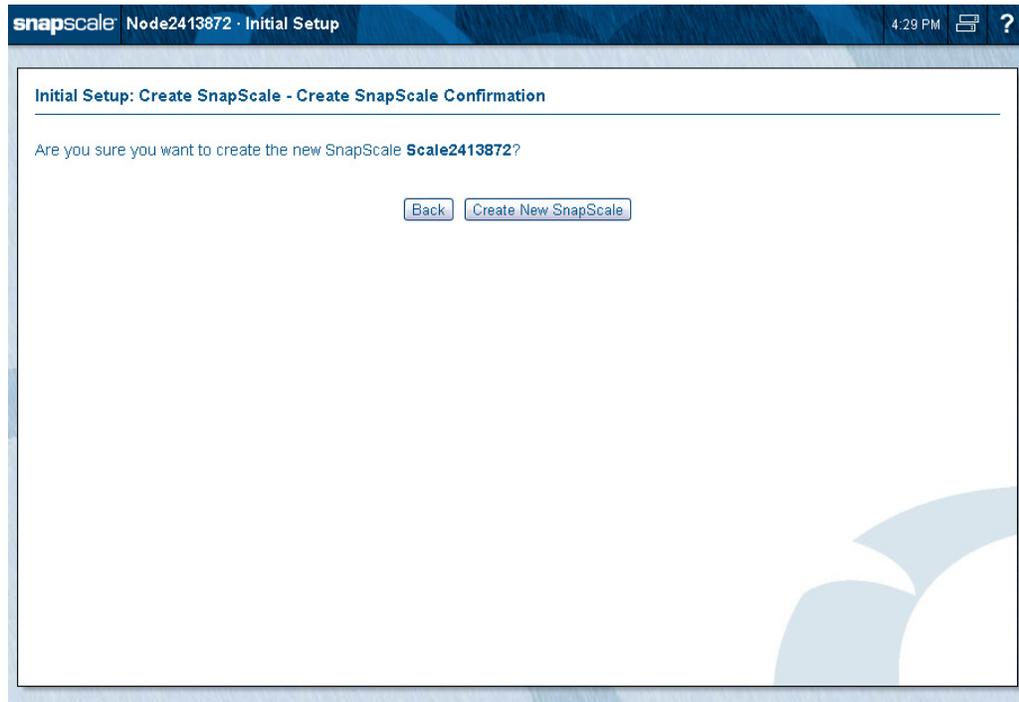
**NOTE:** If you uncheck the box for reserving space for snapshots, an alert is displayed to remind you that the feature will be **permanently disabled** for the cluster.

**SnapScale settings.**

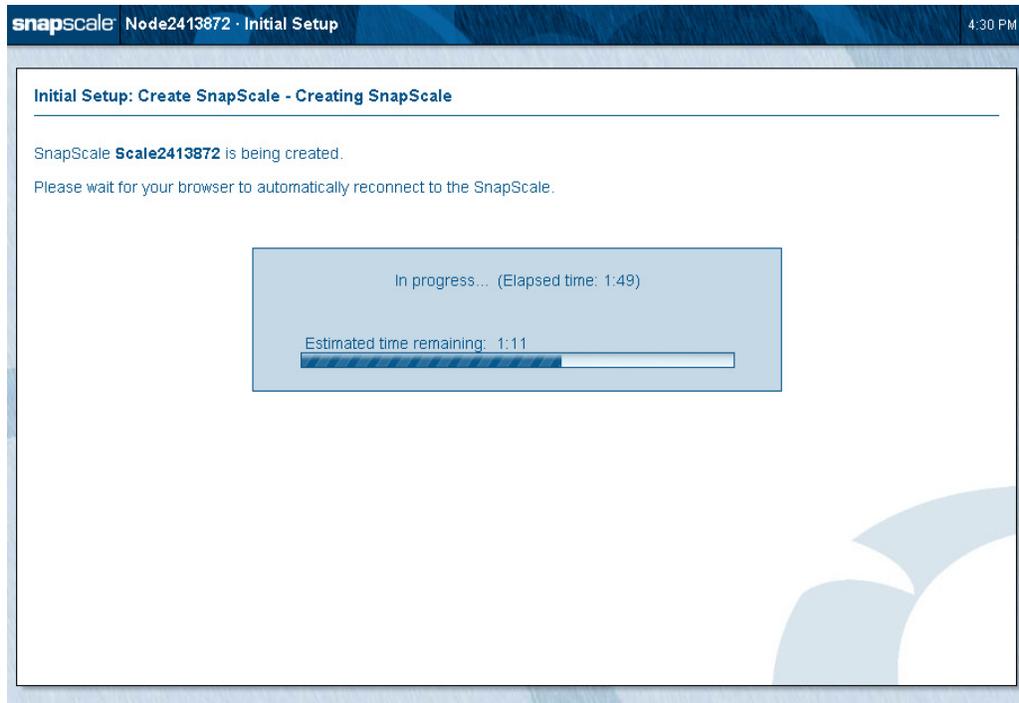
SnapScale Name	SimScale
Data Replication Count	2x
Spare Disks	2
Snapshot Space Reserved	<b>Warning:</b> No space is being reserved for snapshots. The Snapshots feature will be permanently disabled for this SnapScale.
Management IP Address	192.168.199.100 (Please make note of this IP address for later use.)
Subnet Mask	255.255.254.0
Default Gateway	192.168.192.1



Click **Create New SnapScale** to complete the process. A confirmation page is shown.



Click **Create New SnapScale** again to create the cluster. A progress bar is displayed as the SnapScale cluster is created.



Once the cluster is created and the system changes the uninitialized node IP addresses from DHCP to the configured static IP address, a completion page is displayed stating that a share was created and suggesting UPS units be enabled. To enhance security, you are asked to change the default administrator password after the cluster has been successfully created:

The screenshot shows the SnapScale web interface for cluster Scale2413872. The page title is "Initial Setup: Create SnapScale - Complete". The main content area contains the following text:

SnapScale **Scale2413872** has been successfully created and is ready to use. A public share, **SHARE1**, has been created that allows full read/write access to the root of your SnapScale.

You can manage shares under Security > Shares.

**Important:** It is highly recommended that you configure your SnapScale to automatically respond to a low power warning sent from an uninterruptable power supply (UPS). You can do this from the SnapScale > UPS page.

Please change your administrator password now, in order to protect your SnapScale from unauthorized changes.

New Administrator Password  (Leave blank to keep existing administrator password.)

Confirm New Administrator Password

OK

It is highly recommended that you use the password fields at the bottom of the page to change the Administrator's password for the cluster.

After changing the Administrator's password and clicking **OK**, a success page is shown:

The screenshot shows the SnapScale web interface for cluster Scale2302216. The page title is "Create SnapScale - Complete". The main content area contains the following text:

Your administrator password has been successfully changed. Note that you can change your administrator password at any time by changing the password for the "admin" user from the [Security > Local Users](#) page.

Click OK to continue administering your SnapScale.

OK

Click **OK** to continue. The **Login** page is shown. Log in using the new password.

After changing the password and logging back in, the **Registration** page is displayed to facilitate activating your warranty:

**Register your SnapScale to activate your warranty and stay informed of important software and product updates.**  
(Note: You must register your SnapScale whenever you add new nodes.)

Please enter your contact information below and then click the Download button to download the registration (text) file to your computer. You can then email the registration file to [warranty@overlandstorage.com](mailto:warranty@overlandstorage.com) with the subject, **SnapScale Registration Request**.

Name

Email Address

Company Name

Shipping Address  (For RMA purposes.)

You can view the Overland Storage privacy policy by visiting [www.snapserver.com/privacy](http://www.snapserver.com/privacy) (a new browser window will open).

Complete the registration fields and then click **Download Registration File**. Email that file (SnapScaleRegistration.csv) to Overland Storage Service ([warranty@overlandstorage.com](mailto:warranty@overlandstorage.com)) using the subject line “SnapScale Registration Request” to initiate your warranty coverage.

Click **Close**. You will receive a confirmation email to confirm and complete the registration.

When you close that page, the **Administration** page is displayed:



It is recommended that you configure your DNS in your network so clients can resolve the cluster using round-robin name resolution:

- Add a host record for the cluster management name (<clustername>-mgt) to resolve to the Management IP address.
- Add multiple host records for the cluster name resolving to each of the node IP addresses. The DNS resolves lookups for the cluster name via round robin.

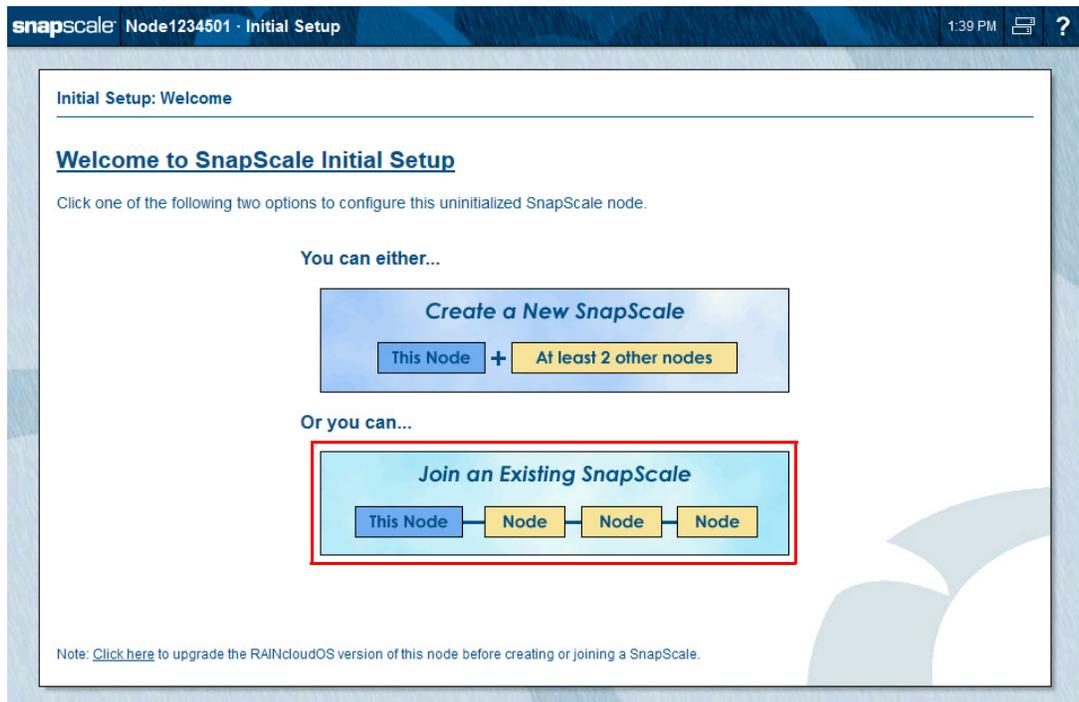
## Join an Existing SnapScale Cluster (via Wizard)

**IMPORTANT:** While the Initial Setup Wizard can be used to add one or more new nodes to an existing cluster, it is recommended that you log into the existing cluster's Web Management Interface and add the nodes using the Add Nodes function (**Storage > Nodes > Add Nodes**). Refer to [Adding Nodes on page 127](#) for more information.

At any time, one or more new nodes can be added to the cluster to expand the storage pool.

**NOTE:** To create new peer sets to expand cluster storage, it is recommended that the number of new nodes you add is equal to the Data Replication Count being used (2x or 3x) and they all be added at the same time.

When you log into any of the new, uninitialized nodes, the Initial Setup Wizard launches displaying the **Welcome** page and its two options. To add this and other nodes to an existing SnapScale cluster, click **Join an Existing SnapScale**.



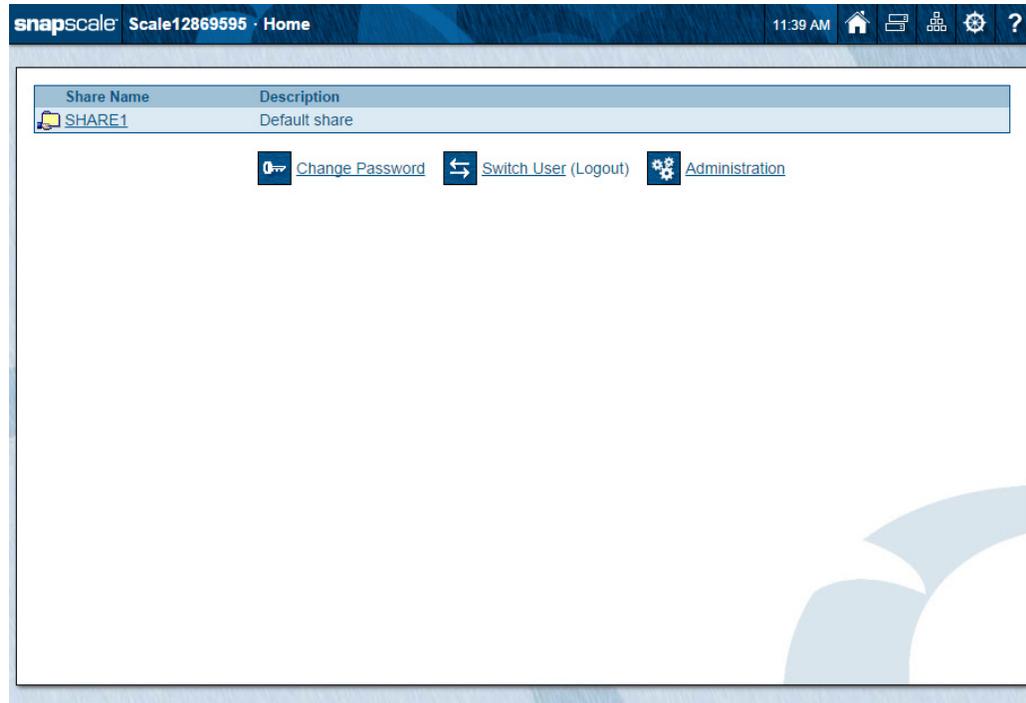
The Initial Setup Wizard then redirects you to the **Add Nodes** page in the Web Management Interface where this node (and all other discovered/new nodes) can be easily added to the cluster. (See [Adding Nodes on page 127](#) for more information.) You are then directed to select the nodes to add, set the static IP addresses, and confirm the settings.

**NOTE:** If no existing SnapScale cluster is detected, a warning is displayed. Verify that the node is on the same Storage network as the other nodes in the cluster, then click **Re-Detect SnapScale**.

## Web Management Interface

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

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



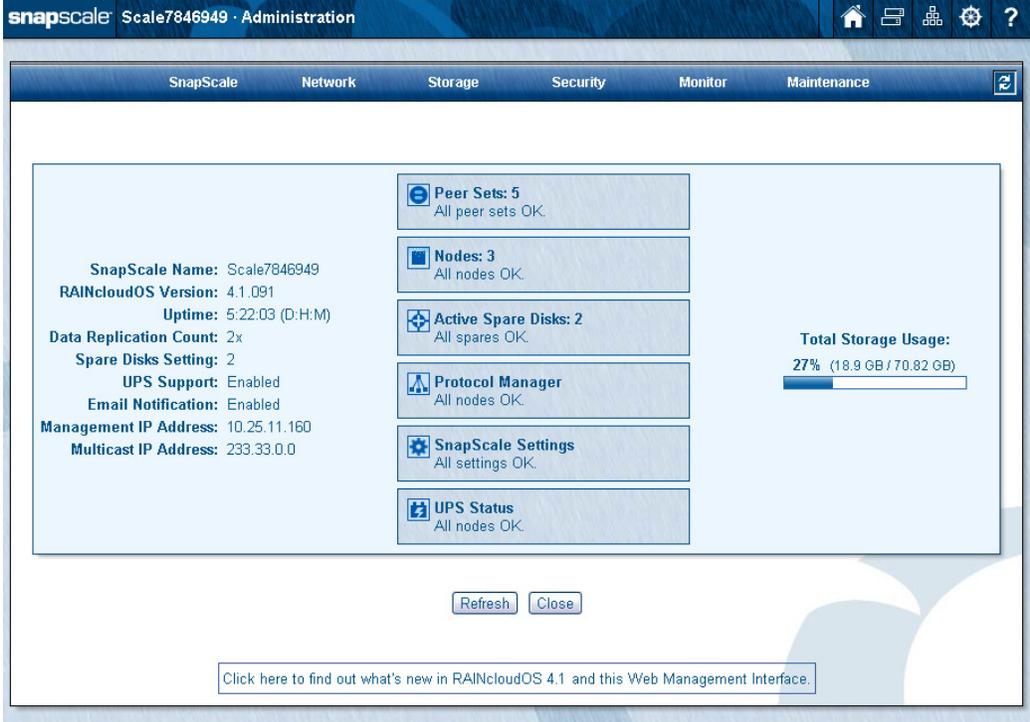
**NOTE:** If you have not gone through the initial setup or authentication is required, you may be prompted to log in when you first access the Web Management Interface.

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

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

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

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



The screenshot shows the SnapScale Administration interface for cluster Scale7846949. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The main content area displays cluster health metrics:

- Peer Sets:** 5 (All peer sets OK)
- Nodes:** 3 (All nodes OK)
- Active Spare Disks:** 2 (All spares OK)
- Protocol Manager:** All nodes OK
- SnapScale Settings:** All settings OK
- UPS Status:** All nodes OK

Cluster details on the left include:

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

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

Buttons for Refresh and Close are located at the bottom of the health summary. A link at the bottom of the page reads: "Click here to find out what's new in RAINcloudOS 4.1 and this Web Management Interface."

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

## Alert Messages

Alert messages are displayed on Administrator-level Web Management Interface pages that display a menu. Some alerts (such as Spare Distributor and Data Balancer) have clickable options:

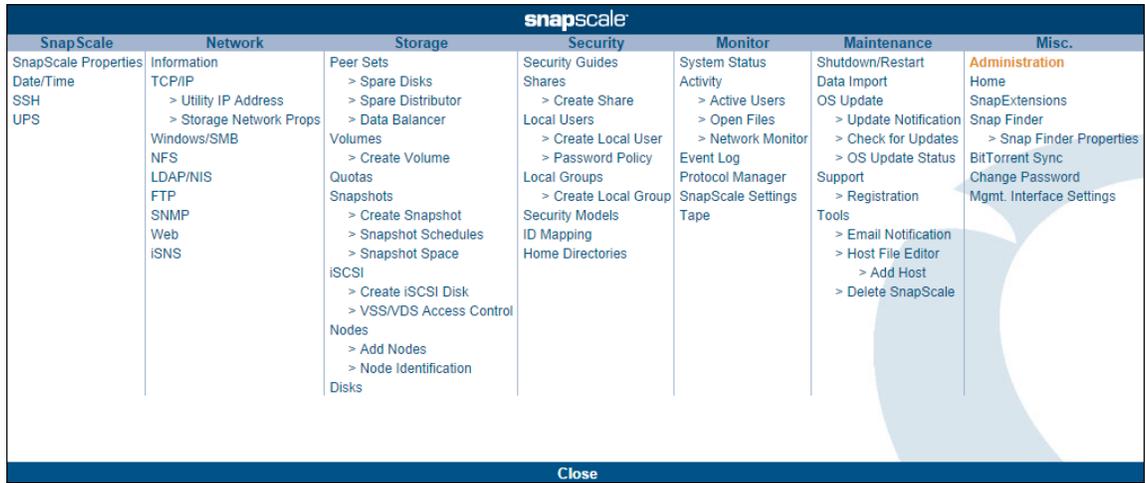
- **[Later]** - Hides the alert for 24 hours or until after feature is run, whichever is first.
- **[Hide]** - Suppresses the alert. It will not be shown again until after the feature called out in the alert is run and a new alert for that feature is generated.

When a cluster is restarted, the Web Management Interface shows the status while the cluster is booting. Because some components are not immediately available, an alert message is displayed showing the percent done and as a reminder that the process is not complete, some nodes may appear offline, and so forth. Some of the status boxes may show warnings.

The screenshot displays the SnapScale Administration interface for cluster Scale1234501. At the top, a navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. A prominent yellow alert banner states: "Important: Your SnapScale is currently starting up. 58% complete. Warning: During this time, status information displayed in this Web Management Interface may be incomplete. (Nodes may appear offline, peer sets may be forming, IP addresses may not be assigned, etc.) Please wait before using or making any changes to your SnapScale." Below the alert, a central panel provides system details and status boxes. System details include SnapScale Name (Scale12869595), RAINcloudOS Version (4.1.056), Uptime (0:15:15), Data Replication Count (2x), Spare Disks Setting (2), UPS Support (Disabled), Email Notification (Disabled), Management IP Address (10.25.11.160), and Multicast IP Address (233.33.0.0). Status boxes show: Peer Sets (80, All OK), Nodes (11, 1 critical, 3 warnings), Active Spare Disks (5, 6 too small), Protocol Manager (10 warnings), SnapScale Settings (1 failed, 1 pending), and UPS Status (6 critical, 2 warnings). A Total Storage Usage bar shows 15% (29.94 TB / 200 TB). Refresh and Close buttons are at the bottom, along with a link to find out what's new in RAINcloudOS 4.1.

## Site Map

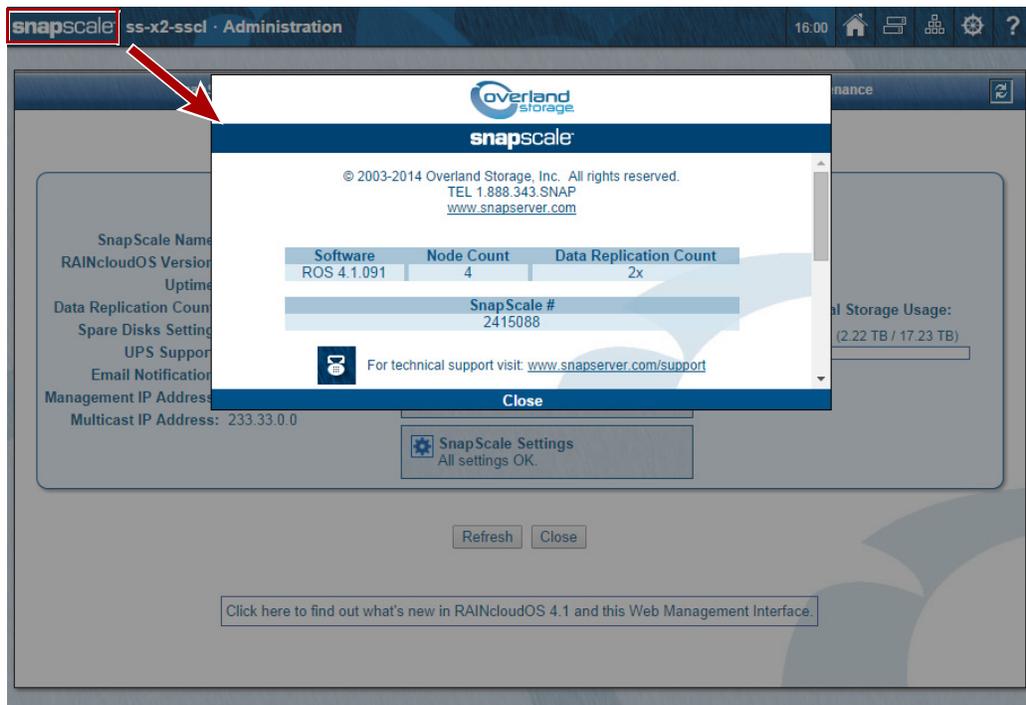
The RAINcloud OS site map (⚙️) provides links to all the web pages that make up the Web Management Interface. All the pages are each covered in detail in the following chapters.



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

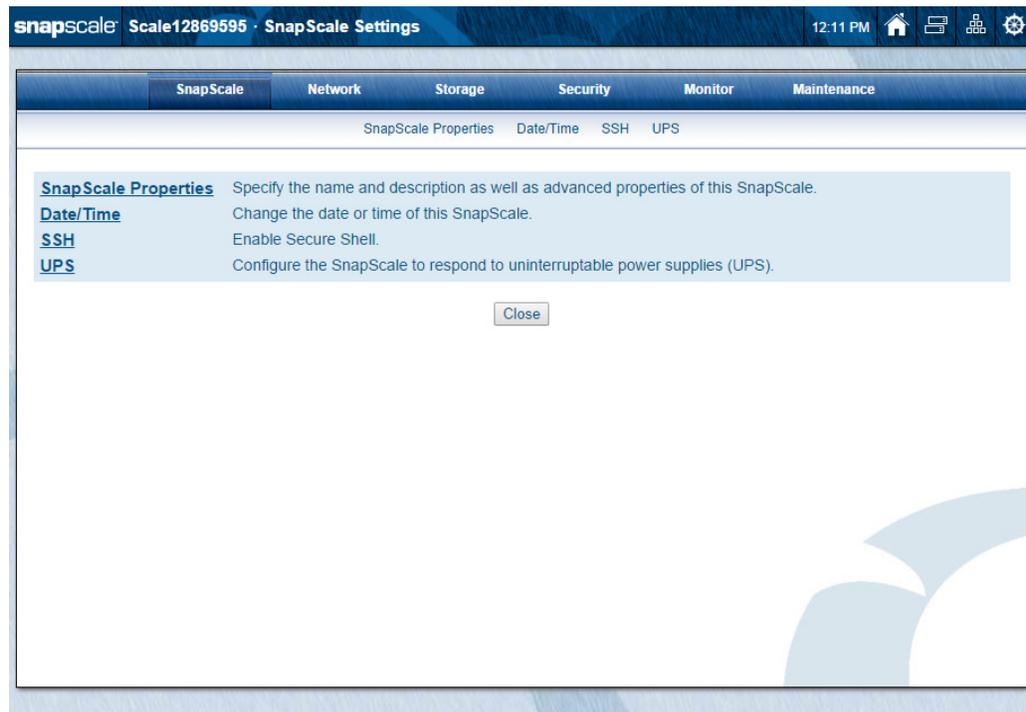
## Contact, Hardware & Software Information

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



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

This section covers the configuration options for a SnapScale cluster of three or more nodes. The four options for cluster settings are found under the SnapScale tab. They can also be accessed using the site map icon (⚙️).



**Topics in SnapScale Settings:**

- [SnapScale Properties](#)
- [Date/Time](#)
- [Secure Shell](#)
- [UPS Protection](#)

## SnapScale Properties

These basic options are found under **SnapScale Properties**:

The screenshot shows the SnapScale Properties configuration page. The interface includes a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. The SnapScale Properties tab is active, showing sub-tabs for SnapScale Properties, Date/Time, SSH, and UPS. The main content area contains the following fields and options:

- SnapScale Name:** Scale12869595
- Description:** (optional)
- Data Replication Count:** 2x (Note: The replication count cannot be changed.)
- Spare disks:** Specifies the number of available disks in the SnapScale to reserve for spares. A spare disk is used to automatically replace a failed peer set member.
  - Allocate spare disks (Active spares: 2)
  - Spare Disks:** 2
- Storage Usage Warning Percentage:** 80
- Storage Usage Critical Percentage:** 95

Buttons for OK and Cancel are located at the bottom of the form.

This table details the options on the **SnapScale Properties** page:

Option	Description
<b>SnapScale Name and Description</b>	<p>Either accept the default cluster name or enter an alphanumeric name up to 15 characters in length. Network clients can use this name along with round robin DNS name resolution to connect to the cluster.</p> <p>The default name is "Scalennnnnnnn" (where nnnnnnnn is the appliance number of the node used to create the cluster).</p>
<b>Description</b>	<p>This optional field provides a place to define the cluster in the overall scheme of your network and better identify the cluster on a LAN.</p>
<b>Data Replication Count</b>	<p>The data replication count establishes the level of data redundancy in the cluster. The setting specifies how many copies of each data file or folder to maintain. A count of 3x offers higher data protection but uses more disk space.</p> <p>Once the cluster is created, the count can only be decreased from 3x to 2x. It cannot be increased from 2x to 3x.</p>
<b>Spare Disks</b>	<p>Check the box and select the number of spare disks you want to reserve. A spare disk is used to automatically replace a failed Peer Set member.</p> <p>If there are unused drives remaining after allocating the number of spares requested, they are used for other peer sets. If there is an insufficient number of drives left to create a final peer set, the drives are configured as additional spares.</p>

Option	Description
Storage Utilization	Use the two drop-down lists to select the percentage of storage used before a warning or critical notice is sent. If not done already, use the link in this section to set up email notification. See <a href="#">Email Notification</a> .

## Date/Time

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

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

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

The screenshot displays the 'Date/Time' configuration page in the SnapScale web management interface. The page title is 'Scale12869595 - Date/Time'. The breadcrumb trail shows 'SnapScale Properties > Date/Time > SSH > UPS'. The main content area contains the following elements:

- A header bar with navigation tabs: SnapScale, Network, Storage, Security, Monitor, Maintenance.
- A sub-header: SnapScale Properties, Date/Time, SSH, UPS.
- Introductory text: "You can set the date and time of this SnapScale to specific values, or you can use NTP\* (Network Time Protocol) servers to automatically synchronize the date and time of this SnapScale. Visit [www.ntp.org](http://www.ntp.org) for a list of public NTP primary and secondary servers."
- Two radio button options:
  - Manual:** "Set the date and time of this SnapScale to the following:" with input fields for Date (2014-09-19) and Time (12:20:07 PM).
  - Automatic:** "Automatically synchronize the date and time of this SnapScale to the following NTP servers." with input fields for NTP Servers (0.pool.ntp.org and 1.pool.ntp.org).
- A note: "\*Note: Due to security restrictions, NTP cannot be used when the SnapScale is joined to an Active Directory domain."
- A Time Zone dropdown menu set to "(UTC-08:00) Pacific Time (US & Canada)".
- Buttons for "OK" and "Cancel".
- A footer note: "Note: Changing the date, time or time zone may require you to re-login to this Web Management Interface."

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

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

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

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

4. Click **OK**.

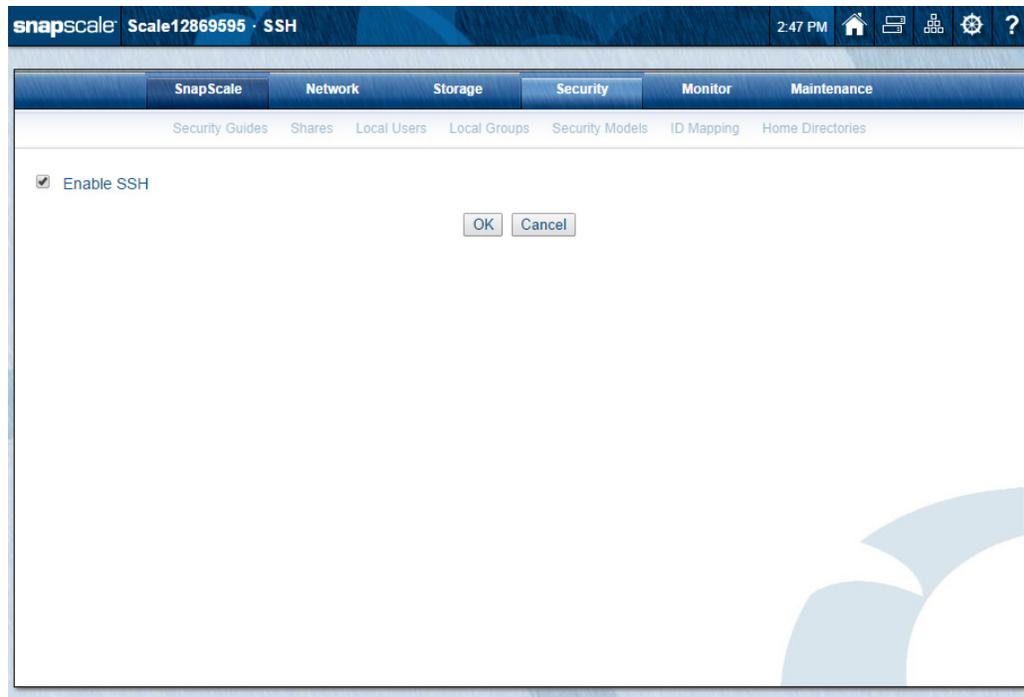
## Secure Shell

Secure Shell (SSH) is a service that provides a remote console to access a command line shell that allows the user to perform basic management and update functions outside the Web Management Interface. See [Command Line Interface on page 267](#) for more information. The SSH implementation requires SSH v2.

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

### Disable SSH

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



## Connect to the CLI using SSH

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

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

You will automatically be placed in the CLI shell.

## UPS Protection

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

**NOTE:** If UPS devices have not been configured, the first time you select this option, you are automatically shown the **UPS Properties** page. See [Edit UPS Properties on page 47](#).

The screenshot shows the SnapScale web interface with the 'UPS Status' page selected. The page title is 'Scale2302216 · UPS Status'. The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'UPS' sub-menu is active. Below the navigation, there are tabs for SnapScale Properties, Date/Time, SSH, and UPS. The main content area shows the 'UPS device status for 3 nodes' with the following table:

Node	Description	UPS Type	Primary UPS $\Delta$	Secondary UPS
VM-Node13710160 (Mgmt. Node)	SnapScale Vi...	Network/Single	Online 10.25.2.52	-
VM-Node5268314	SnapScale Vi...	Network/Single	Online 10.25.2.52	-
VM-Node5898369	SnapScale Vi...	Network/Single	Online 10.25.2.52	-

Below the table are three buttons: 'UPS Properties', 'Refresh', and 'Close'.

An APC Smart-UPS series device allows the SnapScale cluster to shut down gracefully in the event of an unexpected power interruption. You can configure the cluster to automatically shut down when a low power warning is sent from one or more APC network-enabled or USB-based UPS devices (some serial-only APC UPS devices are also supported by using the IOGear

GUC232A USB to Serial Adapter Cable). To do this, you must enable UPS support on the cluster (as described in this section) to listen to the IP address of one or more APC UPS devices and you must supply the proper authentication phrase configured on the UPS devices.

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

## Edit UPS Properties

To manage the network UPS devices, click **UPS Properties**:

**NOTE:** If UPS devices have not been configured, the first time you select that option, you are automatically shown the **UPS Properties** page.

Scale12869595 · UPS Properties 3:30 PM

SnapScale Properties Date/Time SSH UPS

You can configure this SnapScale to automatically respond when a low power warning is sent from a USB-based (if supported on node) or network-based uninterruptable power supply (UPS). Currently, only APC® brand UPS devices\* are supported.

Enable UPS Support

Low battery response requires a low battery message from both primary and secondary UPS devices, if applicable.

Network UPS Devices (0):

(none)

Add Add a network UPS device.  
Change Modify a network UPS device.  
Delete Delete network UPS devices.

UPS devices for 3 nodes.

Node	Description	UPS Type	Primary UPS ▲	Secondary UPS
VM-Node 12869595	-	Please select...	-	(Not supported.)
VM-Node 14122451	-	Please select...	-	(Not supported.)
VM-Node 9715283 (Mgmt. Node)	-	Please select...	-	(Not supported.)

OK Cancel

\*UPS technology copyright © 2001 American Power Conversion Corporation.

UPS Properties page options:

Option	Description
<b>Enable UPS Support</b>	Check the <b>Enable UPS Support</b> box to enable support.
<b>Low battery response message</b>	Check the box to initiate a graceful shutdown only when both the primary and secondary UPS devices for a node send a low battery message.
<b>Network UPS Devices (#)</b>	This field shows a list of UPS devices that are used with the cluster. Use the <b>Add</b> , <b>Change</b> , and <b>Delete</b> buttons to manage the list.

Option	Description
<b>UPS Type</b> (Third column in Node table)	Use the drop-down list in the third column of the Node table to select which UPS device is used: <ul style="list-style-type: none"> <li>• <b>USB</b> – Select this option to use a direct-attached (USB) device.</li> <li>• <b>Network/Single</b> – Use this option to select a network UPS device.</li> <li>• <b>Network/Dual</b> – Use this option to activate the option of a secondary network UPS device.</li> </ul>
<b>Primary UPS</b> (Fourth column in Node table)	Selecting the Network/Single option under UPS Type causes a drop-down list to be displayed in this column. Select the primary UPS to associate with the node from the list (which is based on the Network UPS Devices table).
<b>Secondary UPS</b> (Fifth column in Node table)	If supported, selecting the Network/Dual option (under UPS Type) causes a drop-down list to be displayed in this column. Select the secondary UPS to associate with the node from the list (which is based on the Network UPS Devices table).

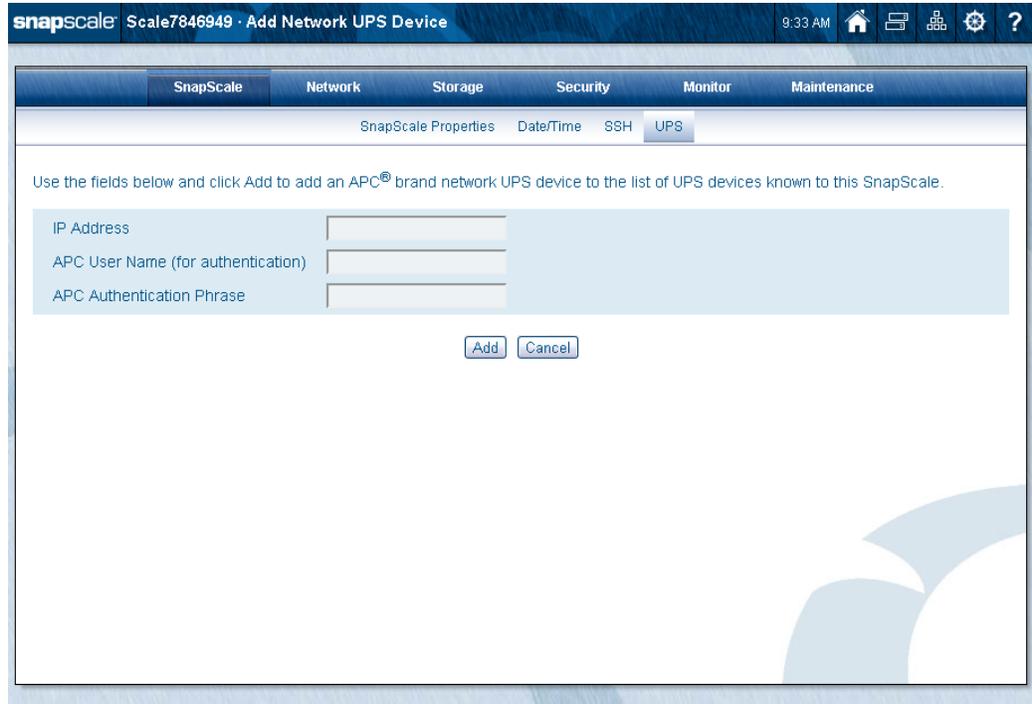
## Configure UPS Protection

1. Check **Enable UPS Support**.
2. If desired, check the **low battery message** option.  
Both Primary and Secondary UPS devices need to have low battery levels before the notice is sent to initiate a graceful shutdown.
3. If necessary, **add** network UPS devices.  
See [Add Network UPS Device](#).
4. Select or change the following from the drop-down lists in the **UPS device table**:
  - UPS Type
  - Primary UPS
  - Secondary UPS
5. Click **OK** to finish.

## Add Network UPS Device

Devices need to be added to the Network UPS Devices table on the **UPS** page for the nodes to be associated with them.

1. To the right of the Network UPS Devices table, click **Add**.
2. At the **Add Network UPS Device** page, enter:
  - IP Address of the device
  - APC User Name (usually the UPS administrator name, default is **apc**)
  - APC Authentication Phrase (found under low battery shutdown configuration in the APC UPS interface; it is **NOT** the Administrator password)



The screenshot shows the SnapScale web interface for adding a network UPS device. The page title is "Scale7846949 · Add Network UPS Device". The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The current page is "UPS". The form contains three input fields: "IP Address", "APC User Name (for authentication)", and "APC Authentication Phrase". Below the fields are "Add" and "Cancel" buttons. A message above the fields reads: "Use the fields below and click Add to add an APC® brand network UPS device to the list of UPS devices known to this SnapScale."

3. Click **Add**.

You are returned to the **UPS** page and the device is shown in the Network UPS Devices table. The table title UPS count is increased by one. Repeat the process for additional devices.

## Change Network UPS Device

To change the settings of a network UPS device:

1. Select a **device** in the Network UPS Devices field to change.
2. Click **Change**.
3. Edit any of the **three options** for the device.
4. Click **Change** again.

Any changes you make are applied to all nodes that are currently using this device.

## Delete Network UPS Device

To delete a network UPS device:

1. If the device is still connected to **any nodes**, deselect the device from the nodes.
2. Highlight the **device** in the Network UPS Devices field.
3. Click **Delete**.

The device is deleted from the list.

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

Network bonding options allow you to configure the SnapScale Client or Storage networks for load balancing, failover, Switch Trunking, and Link Aggregation (802.3ad). Network file protocols control how network clients can access the cluster. Access to the cluster's storage space is provided via Windows (SMB), Unix (NFS), FTP/FTPS, and the Web (HTTP/HTTPS).

**NOTE:** Uninitialized nodes are configured to use DHCP until they are added to a cluster when they switch to the static IP addresses used by the cluster.



**Topics in Network Access:**

- [Network Information](#)
- [TCP/IP Networking](#)
- [Windows/SMB Networking](#)
- [NFS Access](#)
- [LDAP/NIS Domains](#)
- [FTP/FTPS Access](#)
- [SNMP Configuration](#)

- [Web Access](#)
- [iSNS Configuration](#)

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

## Network Information

Browse to **Network > Information** to access the **Network Information** page that displays either the Client or Storage network settings for SnapScale and identifies the node currently serving as the management node. The information is broken into two parts displaying the common and node-specific network information. Use the **View Network** drop-down menu on the upper right side to select either the Client or Storage network details. Error messages are also shown in this area.

### Client Network Information

This page shows the information on the public **Client** network:

The screenshot shows the SnapScale web interface for Network Information. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Network' tab is active, and the 'Information' sub-tab is selected. The page displays SnapScale client network information and node-specific client network information.

**SnapScale client network information:**

Subnet Mask	255.255.0.0
Default Gateway	10.25.1.1
Domain Name	devnet.myoverland.net
Domain Name Servers	10.6.8.34, 10.6.8.35
WINS Servers	-
Bonding Status	Load Balance (ALB)
Management IP Address	10.25.11.160

**Node-specific client network information:**

Node	Ethernet Port Status	IP Address	Speed/Duplex Status	Ethernet Address
<b>Node2470036</b> (Mgmt. Node)	Eth 7: OK	192.168.51.168	10000 Mbps (Auto) / Full Duplex (Auto)	90:E2:BA:2C:42:D0
	Eth 8: OK			00:C0:B6:25:B0:95
	Eth 2: Disabled			-
	Eth 3: Disabled			-
	Eth 4: Disabled			-
<b>Node2470040</b>	Eth 7: OK	192.168.51.167	10000 Mbps (Auto) / Full Duplex (Auto)	90:E2:BA:2F:91:5C
	Eth 8: OK			00:C0:B6:25:B0:99
	Eth 2: Disabled			-
	Eth 3: Disabled			-
	Eth 4: Disabled			-
<b>Node2470146</b>	Eth 7: OK	192.168.51.166	10000 Mbps (Auto) / Full Duplex (Auto)	90:E2:BA:2F:95:FC
	Eth 8: OK			00:C0:B6:25:B1:00
	Eth 2: Disabled			-
	Eth 3: Disabled			-
	Eth 4: Disabled			-

Buttons: Refresh, Close

Field definitions are detailed in the following table:

SnapScale Client Network Information Section	
<b>Subnet Mask</b>	Combines with the IP address to identify the subnet on which the cluster's Client network interfaces are located.
<b>Default Gateway</b>	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.
<b>Domain Name</b>	The ASCII name that identifies the DNS domain name that is added to the cluster name to form the fully-qualified host name of the cluster. Additional space-separated domain names are added to the cluster's domain search suffix list.
<b>Domain Name Servers</b>	The IP address of up to three servers that maintain a mapping of all host names and IP addresses for translating domain names into IP addresses.
<b>WINS Servers</b>	The IP address of up to four Windows Internet Naming Service (WINS) servers which locate network resources in a TCP/IP-based Windows network by automatically configuring and maintaining name and IP address mapping tables.
<b>Bonding Status</b>	Shows Load Balance (ALB), Failover, Switch Trunking, or Link Aggregation (802.3ad) as the selected bonding.
<b>Management IP Address</b>	The IP address configured to access and manage the SnapScale cluster through the Web Management Interface.
Node-specific Client Network Information Section	
<b>Show/Hide Disabled Ethernet Ports</b>	Use the link on the right above the table to toggle the display of Ethernet ports that have been disabled. The link is only available when ports have been disabled (such as, by the addition of a 10Gb PCIe card disabling the built-in 1Gb ports).
<b>Node</b>	The name of the specific node. The node designated as the Management node is so noted.
<b>Ethernet Port Status</b>	Shows abbreviated references of the Ethernet ports of the node and their statuses. <ul style="list-style-type: none"> <li>• <b>OK</b> – A blue icon () indicates a healthy connection.</li> <li>• <b>No Link</b> – A yellow icon () indicates no link for that port.</li> <li>• <b>Failed</b> – A red icon () indicates that the port has failed.</li> </ul>
<b>IP Address</b>	The unique 32-bit value that identifies the node on a network subnet. This is automatically assigned to each node from the pool of IP addresses configured on the cluster.
<b>Speed/Duplex Status</b>	<b>Speed</b> – Ethernet link speed. <b>Duplex Status</b> – Full-duplex; simultaneous two-way data flow.
<b>Ethernet Address</b>	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port (xx:xx:xx:xx:xx:xx).

## Storage Network Information

This page shows the information on the private **Storage** network:

The screenshot shows the SnapScale Network Information page. The top navigation bar includes SnapScale, Scale12869595, Network Information, and a clock showing 11:08 AM. Below the navigation bar, there are tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Network tab, there are sub-tabs for Information, TCP/IP, Windows/SMB, NFS, LDAP/NIS, FTP, SNMP, Web, and iSNS. The main content area displays 'SnapScale storage network information' with a 'View Network' dropdown set to 'Storage'. It shows a table with Subnet Mask (255.255.0.0), Bonding Status (Failover), and Multicast IP Address (233.33.0.0). Below this is a section for 'Node-specific storage network information' with a table listing three nodes: VM-Node12869595, VM-Node14122451, and VM-Node9715283 (Mgmt. Node). Each node row includes Ethernet Port Status (Eth 3 and Eth 4), IP Address, Speed/Duplex Status, and Ethernet Address. At the bottom of the table are 'Refresh' and 'Close' buttons.

Field definitions are detailed in the following table:

SnapScale Storage Network Information	
<b>Subnet Mask</b>	Combines with the IP address to identify the subnet on which the cluster's Storage network interfaces are located.
<b>Bonding Status</b>	Shows Load Balance (ALB), Failover, Switch Trunking, or Link Aggregation (802.3ad) as the selected bonding.
<b>Multicast IP Address</b>	Multicast address used for inter-node cluster messaging.

Node-specific Storage Network Information Section	
<b>Show/Hide Disabled Ethernet Ports</b>	Use the link on the right above the table to toggle the display of Ethernet ports that have been disabled. The link is only available when ports have been disabled (such as, by the addition of a 10Gb PCIe card disabling the built-in 1Gb ports).
<b>Node</b>	The name of the specific node. The node designated as the Management node is so noted.
<b>Ethernet Port Status</b>	Shows abbreviated references of the Ethernet ports of the node and their statuses. <ul style="list-style-type: none"> <li>• <b>OK</b> – A blue icon () indicates a healthy connection.</li> <li>• <b>No Link</b> – A yellow icon () indicates no link for that port.</li> <li>• <b>Failed</b> – A red icon () indicates that the port has failed.</li> </ul>

Node-specific Storage Network Information Section	
<b>IP Address</b>	The unique 32-bit value that identifies the node on a network subnet. This is automatically assigned to each node from the pool of IP addresses configured on the cluster.
<b>Speed/Duplex Status</b>	<b>Speed</b> – Ethernet link speed. <b>Duplex Status</b> – Full-duplex; simultaneous two-way data flow.
<b>Ethernet Address</b>	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port (xx:xx:xx:xx:xx:xx).

## TCP/IP Networking

SnapScale nodes ship with either four 1GbE or 10GbE ports at the rear for network connections. The Storage network ports are always bonded using Failover mode. The Client network ports are bonded by default using Load Balance (ALB), but can be changed after the cluster is created to one of the other bonding modes:

- Failover
- Switch Trunking
- Link Aggregation (802.3ad)

See [Bonding Options on page 54](#) for descriptions.

The **TCP/IP Networking** page provides configuration of the common cluster network settings, the static Management IP address, and the pool of static IP addresses to automatically assign to cluster nodes.

**NOTE:** If the Client network runs a DHCP server, be sure the static IP addresses assigned to the nodes and Management IP are excluded from DHCP assignment.

SnapScale client network settings.  
 Subnet Mask: 255.255.0.0  
 WINS Servers: (optional)  
 Default Gateway: 10.25.1.1 (optional)  
 DNS Domain Name: devnet.myoverland.net (optional)  
 Domain Name Servers: 10.6.8.34 (optional), 10.6.8.35 (optional)  
 Bond Type: Load Balance (ALB)

SnapScale management and node client network static IP addresses.  
 Static IP Address (Management IP address): 10.25.11.160  
 Static IP Address (Node IP address): 10.25.11.161  
 Static IP Address (Node IP address): 10.25.11.162  
 Static IP Address (Node IP address): 10.25.11.163

Optional: Enter a starting IP address and click "Populate" to populate the list on the left with sequential IP addresses.  
 Starting IP Address:   (optional)

The following table describes the configuration options found on the **TCP/IP Networking** page:

Column	Description
<b>Subnet Mask</b>	Combines with the IP address to identify the subnet on which the cluster's Client network interfaces are located.
<b>WINS Servers</b>	The IP address of up to four Windows Internet Naming Service (WINS) servers which locate network resources in a TCP/IP-based Windows network by automatically configuring and maintaining name and IP address mapping tables.
<b>Default Gateway</b>	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.
<b>DNS Domain Name</b>	The ASCII name that identifies the DNS domain name that is added to the cluster name to form the fully-qualified host name of the cluster, and also serves as the primary DNS search suffix. Additional space-separated domain names can be specified to extend the domain search suffix list.
<b>Domain Name Servers</b>	The IP address of up to three servers that maintain a mapping of all host names and IP addresses for translating domain names into IP addresses.
<b>Bond Type</b>	Use the drop-down list to select one of the four bonding modes for the Client network interface on all nodes. See <a href="#">Bonding Options on page 54</a> for more details.

Column	Description
Static IP Address	<p>This table shows the SnapScale Management IP address and the pool of Client network static IP addresses to be automatically assigned by the cluster to the different nodes.</p> <p>To change or populate the list with a contiguous range of IP addresses, in the area to the right, enter a starting IP address and click <b>Populate Static IP Addresses</b>.</p>

## Bonding Options

### The bonding options available for SnapScale nodes:

- **Failover** – This is the default mode for Storage networks. It uses one Ethernet port as the primary network interface and one port 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 on a node should be connected to different switches (though this is not required).

**NOTE:** Failover mode provides switch fault tolerance, as long as ports are connected to different switches.

- **(Automatic) Load Balance (ALB)** – This is the default mode for Client networks. An intelligent software adaptive agent repeatedly analyzes the traffic flow from the node and distributes the packets based on destination addresses, evenly distributing network traffic for optimal network performance. Both ports of the bond need to be connected to the same switch or logical switch.
- **Switch Trunking** – This mode groups multiple physical Ethernet links to create one logical interface. Provides high fault tolerance and fast performance between switches, routers, and servers. Both ports of the bond need to be connected to the same physical or logical switch, and the switch ports must be configured for static link aggregation.
- **Link Aggregation (802.3ad)** – This method of combining or aggregating multiple network connections in parallel is used to increase throughput beyond what a single connection could handle. It also provides a level of redundancy in case one of the links fails. It uses Link Aggregation Control Protocol (LACP), also called dynamic link aggregation, to autonegotiate trunk settings. Both ports of the bond need to be connected to the same switch or logical switch.

## Guidelines in TCP/IP Configuration

Consider the following guidelines when connecting a SnapScale cluster to the network.

### Configure the DNS for Name Resolution and Round Robin Load Distribution

To evenly distribute client access loads to the cluster nodes, add a DNS A record for the cluster name for each IP in the node IP address pool. The DNS server then rotates through the node IP addresses in a round-robin basis when serving name resolution requests for the cluster name.

Do not add an A record for the cluster name pointing to the Management IP address. If desired, or if using Snap EDR, add an A record for the cluster name followed by “-MGT” for the Management IP address. For example, if the cluster name is Scale1234567, create an A record for hostname “Scale1234567-MGT.”

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

All Ethernet ports on the cluster nodes are set to autonegotiate speed and duplex settings with the Ethernet switch. The switch to which the SnapScale is connected *must* be set to autonegotiate; otherwise, network throughput or connectivity to the node may be seriously impacted.

### Cluster Restart Required when Switching the Storage Network to or from Switch Trunking or Link Aggregation (802.3ad)

To prevent the interruption of communication on the Storage network during reconfiguration of the Storage switch, the cluster must be shutdown before changing the Storage network bond setting to or from Switch Trunking or Link Aggregation (802.3ad). After all the Storage switches have been reconfigured, restart the cluster normally by turning the nodes back on.

### Configure the Client Switch for Load Balancing

If you select either Switch Trunking or Link Aggregation (802.3ad) network bonding configuration for the Client or Storage network bond, be sure the switch is configured correctly for that bonding method **after** configuring the bond on the cluster. No switch configuration is required for Adaptive Load Balancing (ALB).

## Edit Storage Network Properties



**IMPORTANT:** Changing the bond type for your SnapScale's storage network may require changes to your network switch.

---

The bond type for the Storage network of a SnapScale cluster can be changed as needed.



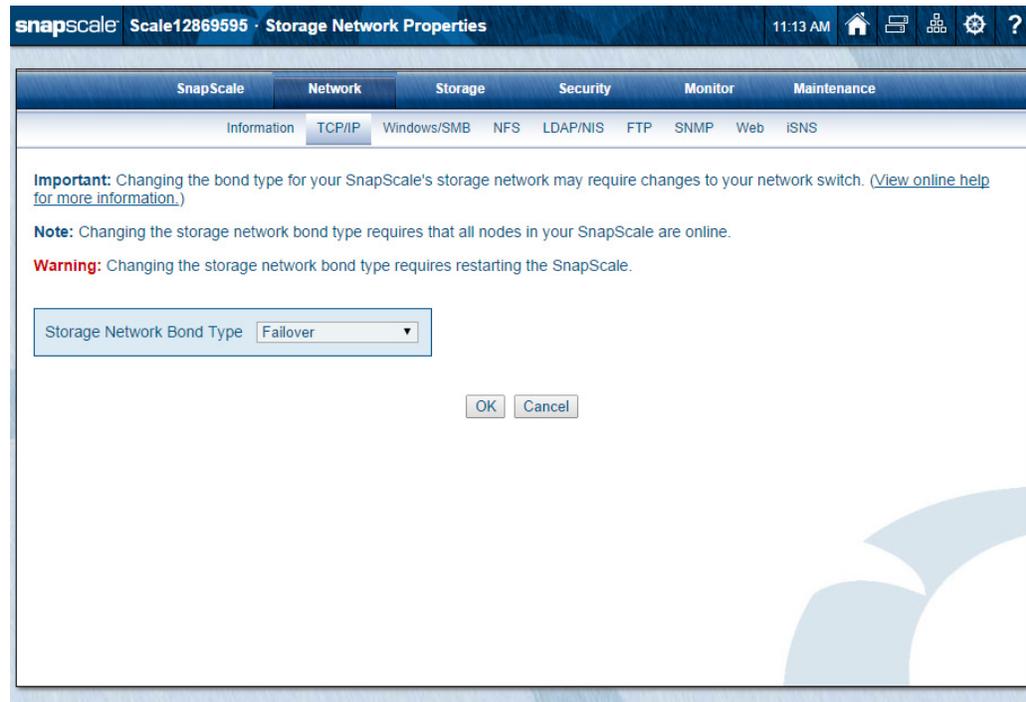
**CAUTION:** All cluster nodes must be online when their bond type is changed. After changing the bond type, the cluster must be restarted. If the switch is being reconfigured, the cluster must be shut down completely, the Storage network switches reconfigured to the new bond type, and then all nodes restarted.

---

The following bond types are supported:

- Failover
- Load Balance (ALB)
- Switch Trunking
- Link Aggregation (802.3ad)

The following page shows the **bonding options** available from the drop-down list:



When changing the bond type, depending on the type of change, the following requirements must be met:

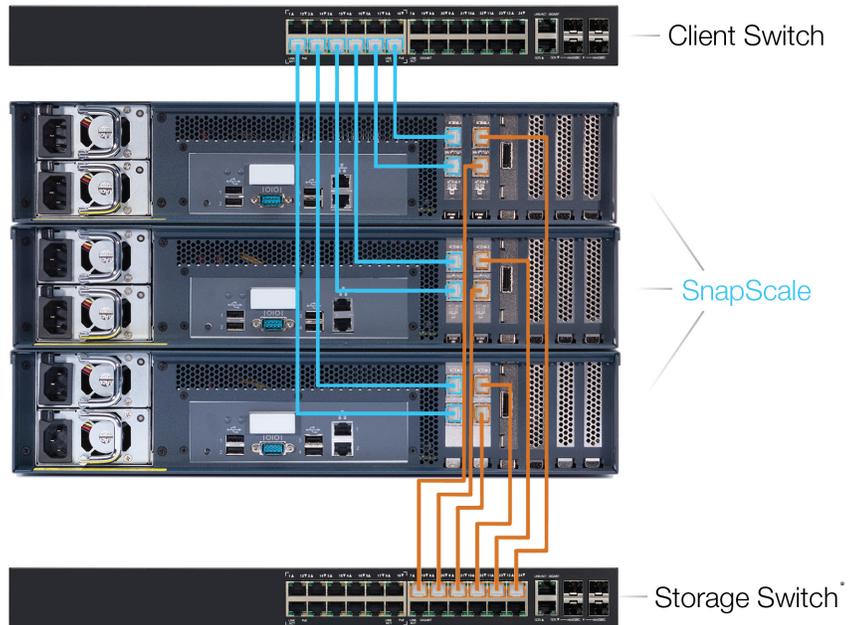
- If changing the Storage network between Failover and ALB, the cluster must reboot.
- If changing the Storage network to or from Switch Trunking or Link Aggregation (802.3ad), the cluster must be shut down completely, the Storage network switches reconfigured to the new bond type, and then all nodes restarted.

---

**CAUTION:** If you change the bonding mode from the default Failover to ALB, Switch Trunking, or Link Aggregation (802.3ad), you **MUST** re-cable the Storage network ports on each node to the same switch. You **CANNOT** straddle them across two Storage network switches like you do for Failover.

---

**Example of Cabling for ALB, Switch Trunking, or Link Aggregation (802.3ad)**



**Utility IP Address**

To assign an additional static IP address to a specific node, click **Utility IP Address** on the **TCP/IP Networking** page.



The Utility IP address can be used to reliably access a specific node by a known IP address, and is particularly useful for backup agents and media server installations. The Utility IP address is assigned to the node in addition to its static IP address automatically assigned from the cluster IP address pool, as well as the Management IP if the node serves as the Management node.

**NOTE:** The Utility IP address must be located on the same subnet as the SnapScale Client network. The address should be assigned BEFORE installing a backup agent or media server on a node. Once the Utility IP address has been assigned, you must add a host record to the DNS server for the node name pointing to the Utility IP address (do NOT add it as another host record for the cluster name).

Only one Utility IP address can exist on a cluster. The Web Management Interface will not allow a new Utility IP address to be created if a Utility IP address currently exists, or when an address does not exist but there are one or more offline nodes (which may have an address already configured on them). The Utility IP address also must not be the same as the Management IP address or any existing address in the cluster IP address pool.

#### Configure a Utility IP Address:

1. On the **Utility IP Address** page, in the empty field, enter a **static IP address** on the same subnet as the IP address pool on the Client network.
2. Using the drop-down list, select the **cluster node** to which the Utility IP address will be assigned.
3. Click **OK**.
4. At the confirmation page, click **Save Changes**.

The Utility address is displayed on the **Network Information** page (**Network > Information**) beneath the static address of the node on which it was configured. The Utility IP address remains with the node, even if the node is restarted or goes offline.

## Windows/SMB Networking

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

The top screenshot shows the SnapScale Windows/SMB configuration page. The 'Member of' dropdown is set to 'Workgroup', which is circled in red. Below it, the 'Workgroup Name' is 'Workgroup'. The 'Enable Windows Networking (SMB)' checkbox is checked. Other checked options include 'Enable Opportunistic Locking', 'Enable this SnapScale as the Master Browser', and 'Enable SMB2'. There are 'OK' and 'Cancel' buttons at the bottom.

The bottom screenshot shows the same page but with the 'Member of' dropdown set to 'Active Directory Domain', which is circled in red. The 'Domain Name' is 'Workgroup'. Below it, there are input fields for 'Administrator Name', 'Administrator Password', and 'Organizational Unit', each with a '(Required to join a domain)' note. The 'LDAP Signing' dropdown is set to 'Plain'. The same checked options and buttons are present as in the top screenshot.

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

## Support for Windows/SMB Networking

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

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

### Support for Microsoft Name Resolution Servers

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

### ShareName\$ Support

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

**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, use the hidden share option under **Advanced Share Properties** when creating a new share or modifying an existing share. This hides a share from SMB, AFP, HTTP, HTTPS, and FTP clients. However, shares are not hidden from NFS, which cannot connect to shares that aren't visible. To hide shares from NFS clients, consider disabling NFS access on hidden shares.

## Support for Windows Network Authentication

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

**NOTE:** When a SnapScale cluster joins a domain, it does so under its cluster name (*Scalennnnnnn*). When a domain user is authenticated on a node, the cluster name is used. As such, a user can use any node of the cluster to be authenticated and log on.

### Windows Networking Options

Windows environments operate in either workgroup mode, where each cluster 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 cluster in the workgroup.

Option	Description
<b>Active Directory Service (ADS)</b>	<p>When operating in a Windows Active Directory domain environment, the SnapScale is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log into the domain through their Windows-based client machines. Active Directory domains resolve user authentication and group membership through the domain controller.</p> <p>Once joined to a Windows Active Directory domain, the SnapScale cluster 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 SnapScale cluster.</p> <p><b>NOTE:</b> 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 or cluster. The user credentials are always encrypted before they are transmitted over the network.

The SnapScale cluster supports the Microsoft Windows implementation of Kerberos. In Windows ADS, the domain controller is also the directory server, the Kerberos Key Distribution Center (KDC), and the origin of group policies that are applied to the domain.

**NOTE:** Kerberos requires the cluster's time to be closely synchronized to the domain controller's time. This means that (1) the cluster 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 SnapScale supports the Microsoft Windows family of servers that run in ADS mode. Any SnapScale cluster can join Active Directory domains as a member server. References to the SnapScale 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 SnapScale cluster

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

## Connect from a Windows Client

Windows clients can connect to the SnapScale using either the cluster name or any IP address in the node IP address pool. However, if possible, clients should use the cluster name to benefit from round robin DNS resolution (see [Configure the DNS for Name Resolution and Round Robin Load Distribution on page 56](#)).

To navigate to the cluster using Windows Explorer, use one of these procedures:

- For Microsoft Windows Vista, 2008, and 7 clients, navigate to **Network > server\_name**.
- For Microsoft Windows XP, 2000, or 2003 clients, navigate to **My Network Places > workgroup\_name > server\_name**.

## Connect a Mac OS X Client Using SMB

Mac OS X clients can connect using SMB. Specify the cluster name (or an IP address from the node IP address pool) in the Connect to Server window (from **Finder** press **Cmd + K**, or select **Finder > Go > Connect to Server**) as one of the following:

**NOTE:** If possible, clients should use the cluster name to benefit from round robin DNS resolution (see [Configure the DNS for Name Resolution and Round Robin Load Distribution on page 56](#)).

- `smb://cluster_name`
- `smb://node_ip_address`

**Tip:** To disconnect from the SnapScale cluster, drag its icon into the Trash.

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

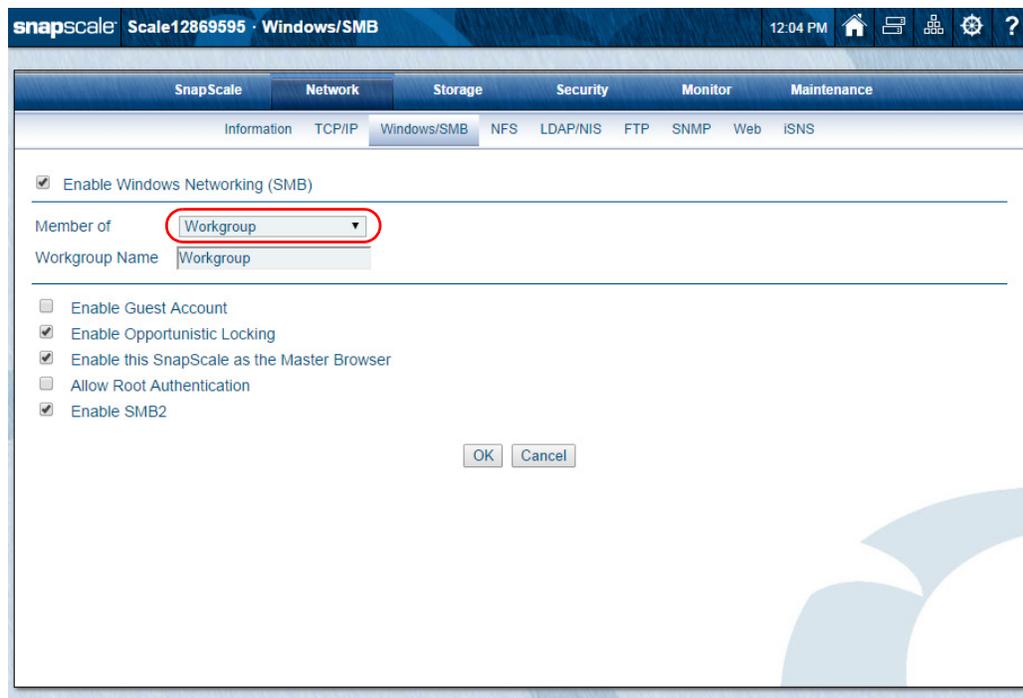
## Configure Windows/SMB Networking

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

Before performing the configuration procedures provided here, be sure you are familiar with the information provided in [Support for Windows/SMB Networking on page 62](#) and [Support for Windows Network Authentication on page 62](#).

### To Join a Workgroup

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



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

Option	Settings
<b>Enable Windows Networking (SMB)</b>	Check the box to enable SMB and activate the options. Clear the box to disable.
<b>Member Of</b>	Verify that it is set to <b>Workgroup</b> .  <b>NOTE:</b> For the Active Directory Domain option, see <a href="#">To Join an Active Directory Domain on page 65</a> .
<b>Workgroup Name</b>	The default settings make the SnapScale available in the workgroup named <i>Workgroup</i> . Enter the workgroup name to which the cluster belongs.
<b>Enable Guest Account</b>	Check the box to allow unknown users (or users explicitly logging in as <b>Guest</b> ) to access the SnapScale using the guest account. Clear the box to disable this feature.
<b>Enable Opportunistic Locking</b>	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the box to disable this feature.
<b>Enable this SnapScale as the Master Browser</b>	Enabled by default. The SnapScale can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the box if you plan to install this cluster in a Windows environment and you want it to be able to serve as the Master Browser for a workgroup. Clear the box to disable this feature.
<b>Allow Root Authentication</b>	Check the box to allow root login to the cluster; clear the box to disable this feature.  <b>NOTE:</b> The root password is synchronized with the cluster's admin password.
<b>Enable SMB2</b>	Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).

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

### To Join an Active Directory Domain

When the cluster joins a domain, it does so as a single unit under the cluster name, and all nodes operate equally under the cluster name to authenticate against the domain. This provides multi-point access to the domain through each node.

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

- From the drop-down Member list, select **Active Directory Domain** to view the configuration page.

The screenshot shows the SnapScale configuration interface for Windows/SMB. The 'Member of' dropdown menu is highlighted with a red circle and contains the text 'Active Directory Domain'. Below it, the 'Domain Name' field contains 'Workgroup'. There are input fields for 'Administrator Name' and 'Administrator Password', both with a note '(Required to join a domain)'. Below these is an 'Organizational Unit' field. A second red circle highlights the 'LDAP Signing' dropdown menu, which is currently set to 'Plain'. Below this are several checkboxes: 'Enable Guest Account' (unchecked), 'Enable Opportunistic Locking' (checked), 'Enable this SnapScale as the Master Browser' (checked), 'Allow Root Authentication' (unchecked), 'Disable NetBIOS over TCP/IP' (unchecked), 'Enable Trusted Domains' (unchecked), and 'Enable SMB2' (checked). At the bottom right, there are 'OK' and 'Cancel' buttons.

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

- Edit the **fields** shown in the following table:

Option	Description
<b>Enable Windows Networking (SMB)</b>	Check the box to enable SMB and activate the options. Clear the box to disable.
<b>Member Of</b>	Verify it shows <i>Active Directory Domain</i> .
<b>Domain Name</b>	The default settings make the SnapScale available in the workgroup named <i>Workgroup</i> . Enter the domain name to which the cluster belongs. <b>NOTE:</b> Windows 2000 domain controllers must run SP2 or later.
<b>Administrator Name / Administrator Password</b>	If joining a domain, enter the user name and password of a user with domain join privileges (typically an administrative user).
<b>Organizational Unit</b>	To create a machine account at a different location than the default, enter a name in the field. By default, this field is blank, signaling the domain controller to use a default defined within the controller. <b>NOTE:</b> Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([org_unit]/[sub-unit1]/[sub-unit1a])
<b>LDAP Signing</b>	Use the drop-down list to set ADS domain LDAP signing to <b>Plain</b> (no signing), <b>Sign</b> , or <b>Seal</b> , as appropriate for your domain. Default setting is <b>Plain</b> .

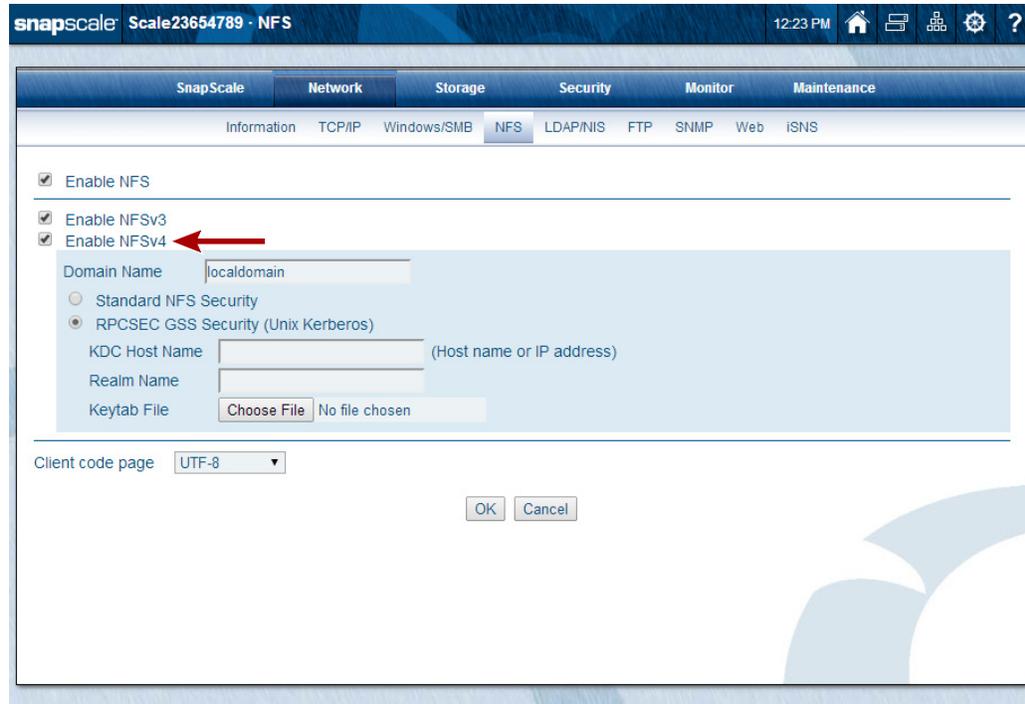
Option	Description
<b>Enable Guest Account</b>	Check the box to allow unknown users (or users explicitly logging in as <b>Guest</b> ) to access the SnapScale using the guest account. Clear the box to disable this feature.
<b>Enable Opportunistic Locking</b>	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the box to disable this feature.
<b>Enable this SnapScale as the Master Browser</b>	Enabled by default. The SnapScale can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the box if you plan to install this cluster in a Windows environment and you want it to be able to serve as the Master Browser for a workgroup. Clear the box to disable this feature.
<b>Allow Root Authentication</b>	Check the box to allow root login to the cluster. Clear the box to disable this feature.  <b>NOTE:</b> The root password is synchronized with the cluster's admin password.
<b>Disable NetBIOS over TCP/IP</b>	Some administrators may wish to disable NetBIOS over TCP/IP. Check the box to disable NetBIOS; clear the box to leave NetBIOS enabled.  <b>NOTE:</b> If you disable NetBIOS and you are joining a domain, you must enter the domain name as a fully qualified domain name (such as, "actdirdomainname.companyname.com"). A short form such as ActDirDomName does not work.
<b>Enable Trusted Domains</b>	SnapScale clusters recognize trust relationships established between the domain to which the SnapScale is joined and other domains in a Windows environment by default. Check the box to enable this feature; clear the box to disable this feature.  <b>NOTE:</b> SnapScale clusters 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, is retained.
<b>Enable SMB2</b>	Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).

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

## NFS Access

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

**NOTE:** Only NFSv3 is enabled by default. If you wish to enable NFSv4, check the **Enable NFSv4** box on the **Network > NFS** page to view additional options.



NFS client access to shares can be specified by navigating to the **Security > Shares** page and clicking the **NFS Access** link next to the share. To ensure proper Unicode representation on the file system, set the client code page to indicate the code page used by NFS clients to represent characters in filenames (usually UTF-8 on modern Unix/Linux-based operating systems).

These versions of the NFS protocol are supported:

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

\*NFSv4 ACLs are not supported.

## Assigning Share Access to NFS Users

The NFSv3 protocol does not support user-level share access control, but rather supports host- and subnet-based access control. NFSv4 supports user-level access control via Kerberos configuration, but otherwise uses the same form of host-based access control. On a standard Unix server, share access is configured in an “exports” file. On SnapScale clusters, the exports for each share are configured on the **NFS Share Access** page independently of user-based share access for other protocols.

## Enable NFS Access to the Cluster

1. Go to **Network > NFS**.
2. Check the **Enable NFS** box.
3. Check the **versions** you want to enable.  
Select one or more from **NFSv3** and **NFSv4**.

4. Choose the desired **Client code page** from the drop-down list.  
Select **UTF-8**, **ISO-8859-1**, **ISO-8859-15**, or **EUC-JP**.
5. Click **OK**.

## Configure NFSv4 Access

1. Go to **Network > NFS**.
2. Check the **Enable NFS** and **Enable NFSv4** boxes.  
A new set of security options are displayed below the **Enable NFSv4** option.

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

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

4. If you selected **RPSEC GSS Security (Unix Kerberos)** security, complete the new options displayed using the table below. Note the following:
- The service will not start unless the TCP/IP domain name is set up exactly the same as the keytab.
  - You must create the NFS and host service entries in the keytab with the fully qualified domain name of the SnapScale cluster.

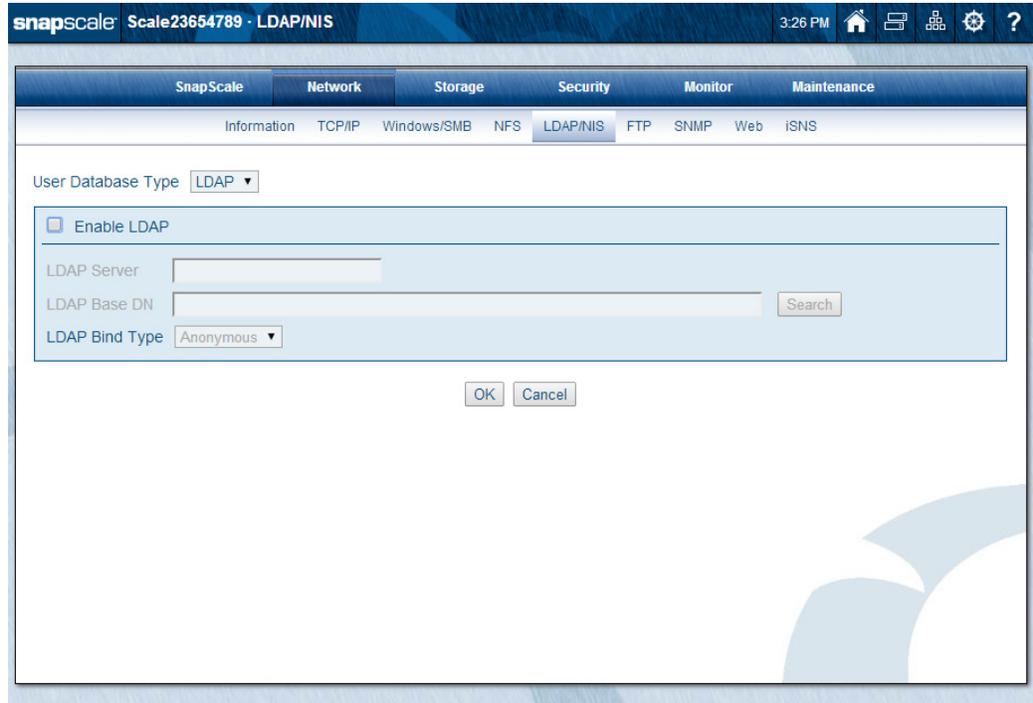
Option	Description
<b>KDC Host Name</b>	Enter the host name of the Kerberos server (for example, "kerberos-2000.mit.edu").
<b>Realm Name</b>	Enter the Kerberos realm name (For example, "ATHENA.MIT.EDU"). <b>NOTE:</b> Realm names are conventionally specified in all CAPITAL letters, but this is not required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

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

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

## LDAP/NIS Domains

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



### LDAP vs. NIS Overview

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

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

NIS (Network Information Service) is a client-server directory service protocol for distributing system configuration data such as user and host names between computers on a computer network. The SnapScale cluster can join an NIS domain and function as an NIS client. It can then read the users and groups maintained by the NIS domain to translate user/group names to UIDs/GIDs for configuration of quotas, ID mapping, and home directories. As such, you must use the NIS server to make modifications.

**NOTE:** Changes you make on the NIS server do not immediately appear on the SnapScale cluster. It may take up to 10 minutes for changes to be replicated.

### Configuring LDAP

Use this procedure to configure LDAP on your SnapScale cluster:

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

The default LDAP database type page is shown.

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

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

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

## Configuring NIS

**NOTE:** Unless UID/GID assignments are properly handled, NIS users and groups may fail to display properly. For guidelines on integrating compatible SnapScale cluster UIDs, see [User and Group ID Assignments on page 143](#).

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

1. Go to **Network > LDAP/NIS**.  
The default LDAP database type page is shown.
2. From the **User Database Type** drop-down list, select **NIS**.

The screenshot shows the SnapScale web interface for configuring LDAP/NIS. The top navigation bar includes 'SnapScale', 'Scale23654789', and 'LDAP/NIS'. The main menu has tabs for 'Information', 'TCP/IP', 'Windows/SMB', 'NFS', 'LDAP/NIS', 'FTP', 'SNMP', 'Web', and 'iSNS'. The 'User Database Type' dropdown is set to 'NIS'. Below this, there is a section for 'Enable NIS' with a checked checkbox. Underneath, there is a text input field for 'NIS Domain Name'. Two radio button options are available: 'Broadcast and bind to any NIS server' (which is selected) and 'Broadcast and bind to the following NIS server' (which is unselected). At the bottom of this section are 'OK' and 'Cancel' buttons.

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

Options	Description
<b>NIS Domain Name</b>	Enter the NIS domain name.
<b>NIS Server</b>	To bind to an NIS server, select either: <ul style="list-style-type: none"> <li>• <b>Broadcast and Bind to Any NIS server</b> to bind to any available NIS servers.</li> <li>• <b>Broadcast and Bind to the following NIS server</b> and enter the IP address for a specific NIS server in the field provided.</li> </ul>

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

## FTP/FTPS Access

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

For FTPS, it is recommended that your FTPS client application use explicit FTPS (such as, FTPES or Auth TLS).

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

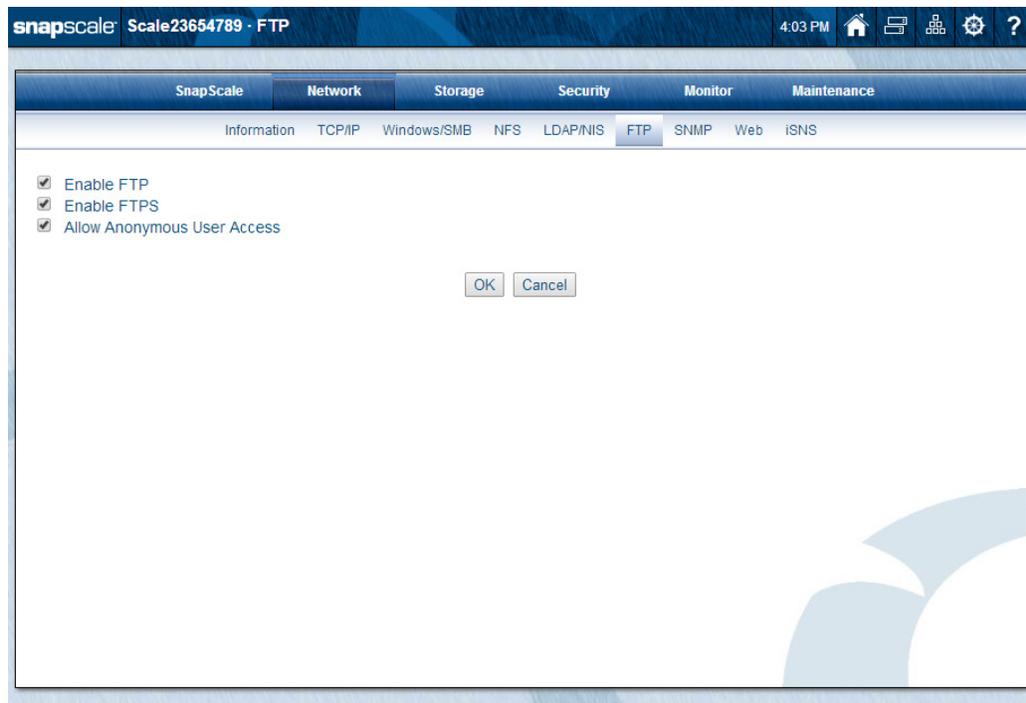
### Supported FTP Clients

SnapScale clusters have been tested with the most common FTP clients and work as expected based on the commands required by RFC 959. SnapScale clusters have been proven to work with these products for standard FTP.

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

### Configure FTP/FTPS Access

1. Go to **Network > FTP**.



2. Edit the **settings** as shown in the following table:

Option	Settings
<b>Enable FTP</b>	Check the box to enable standard FTP services; leave the box blank to disable access to this cluster via standard FTP.
<b>Enable FTPS</b>	Check the box to enable FTPS services; leave the box blank to disable access to this cluster via FTPS.
<b>Allow Anonymous User Access</b>	<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 SnapScale users can access the cluster.</p> <ul style="list-style-type: none"> <li>• <b>Check the box</b> to allow users to connect to the cluster using the anonymous user account. The anonymous user is mapped to the local guest user account. You can set share access for anonymous FTP/FTPS users by granting either read-write (the default access) or read-only access to the guest account on a share-by-share basis.</li> <li>• <b>Leave the box blank</b> so users cannot log in anonymously but must instead log in via a locally created user name and password.</li> </ul>

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

## Connect via FTP/FTPS

1. To connect to the SnapScale cluster:
  - For **standard FTP**, enter the name of the cluster 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 cluster_name`
    - To connect via a **Web browser**, enter:  
`ftp://cluster_name`

(where *cluster\_name* is the name or IP address of the cluster)
  - For **secure FTPS**, configure your FTPS client application to use explicit FTPS (such as, FTPES or “Auth TLS”) and enter the name of the cluster or IP address.

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

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

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

## SNMP Configuration

The SnapScale cluster 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

cluster communicates with SNMP managers in the same read-only community. A community name is a password that authorizes managers and agents to interact. The cluster only responds to managers that belong to the same read-only community.

## Default Traps

A *trap* is a signal from the SnapScale cluster or any individual node informing an SNMP manager program that an event has occurred. SnapScale supports the default traps shown in this table:

Trap	Initiating Action
<b>coldStart</b>	Whenever SNMP is enabled and a node boots.
<b>linkDown</b>	A node's Ethernet interface has gone offline.
<b>linkUp</b>	A node's Ethernet interface has come back online.
<b>authenticationFailure</b>	An attempt to query the SNMP agent using an incorrect read-only community string was made, and resulted in a failure.
<b>enterpriseSpecific</b>	<p>SnapScale-generated traps that correspond to the error-level, warning-level, and fatal-error-level traps of RAINcloudOS. These traps contain a descriptive message that helps to diagnose a problem using the following OIDs:</p> <ul style="list-style-type: none"> <li>1.3.6.1.4.1.6411.2000.1000.1:loglevel 0 syslog messages (<i>emergency</i>)</li> <li>1.3.6.1.4.1.6411.2000.1001.1:loglevel 1 syslog messages (<i>alert</i>)</li> <li>1.3.6.1.4.1.6411.2000.1002.1:loglevel 2 syslog messages (<i>critical</i>)</li> <li>1.3.6.1.4.1.6411.2000.1003.1:loglevel 3 syslog messages (<i>error</i>)</li> </ul> <p><b>NOTE:</b> There is no specific MIB that defines traps sent by SnapScale clusters or nodes.</p>

## Supported Network Manager Applications and MIBs

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

## Configure SNMP

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

The screenshot shows the SnapScale web interface for configuring SNMP. The top navigation bar includes 'SnapScale', 'Scale12869595', and 'SNMP'. The main menu has tabs for 'SnapScale', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Network', there are sub-tabs for 'Information', 'TCP/IP', 'Windows/SMB', 'NFS', 'LDAP/NIS', 'FTP', 'SNMP', 'Web', and 'ISNS'. The 'SNMP' configuration page is active, showing the following settings:

- Enable SNMP
- Read-Only Community:
- Location:  (optional)
- Contact:  (optional)
- Enable SNMP traps
- IP Address 1:
- IP Address 2:  (optional)
- IP Address 3:  (optional)
- IP Address 4:  (optional)
- Send a test trap to listed IP addresses upon saving settings.

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

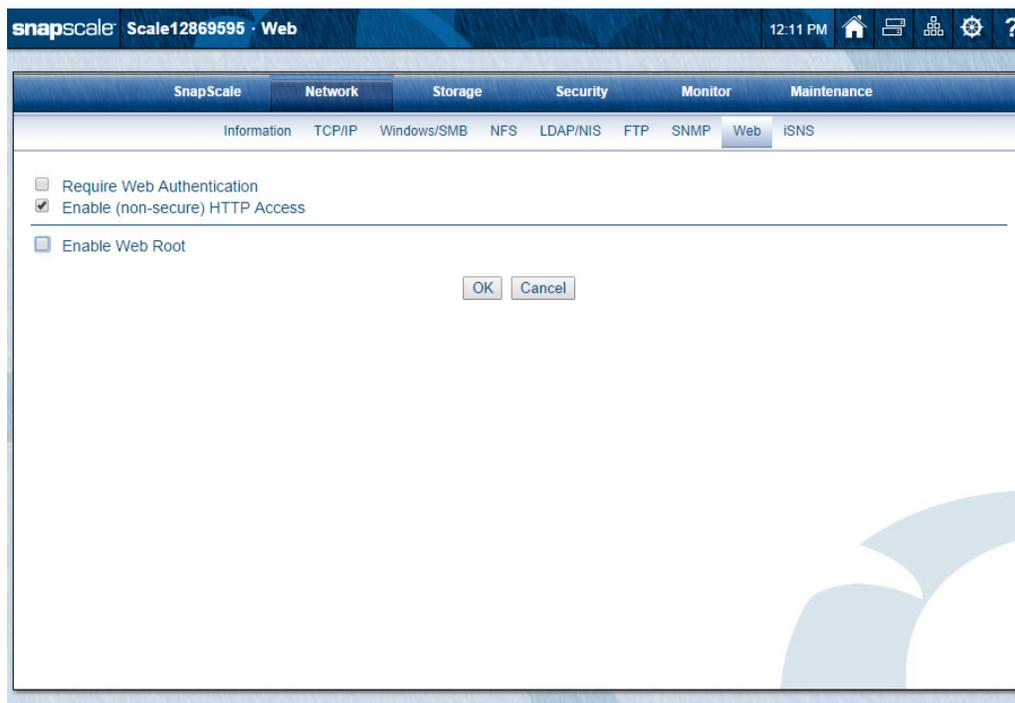
Edit settings as described in the following table, and then click **OK**. Once enabled, SNMP managers can access MIB-II and Host Resources MIBs management data on the cluster.

Option	Description
<b>Enable SNMP</b>	To enable SNMP, check the <b>Enable SNMP</b> box. Leave the box blank to disable SNMP.
<b>Read-Only Community</b>	To allow SNMP managers to read data from this cluster, enter a read-only community string or accept the default <i>snap_public</i> . <b>NOTE:</b> As a precaution against unauthorized access, Overland Storage recommends that you create your own community string.
<b>Location</b>	Optionally enter information that helps a user identify the physical location of the cluster. 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> .
<b>Contact</b>	Optionally enter information that helps a user report problems with the cluster. For example, you might include the name and title of the system administrator, a telephone number, pager number, or email address.
<b>Enable SNMP Traps</b>	Check the <b>Enable SNMP Traps</b> box to enable traps. Clear the box to disable SNMP traps.

Option	Description
IP Address 1-4	Only available when SNMP traps are enabled. Enter the IP address of at least one SNMP manager in the first field as a trap destination. Optionally, you can enter up to three additional IP addresses in fields 2-4.
Send a Test Trap	Only available when SNMP traps are enabled. To verify your settings, check the <b>Send a test trap</b> box, then click <b>OK</b> .

## Web Access

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



### Configuring HTTP/HTTPS

You can require web authentication, disable HTTP (non-secure) access, and enable the Web Root feature. All HTTP access is made via the root node and the Management IP address.

#### To Require Web Authentication

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

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

### To Enable HTTP Access to the SnapScale cluster

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

Option	Description
<b>Enable (non-secure) HTTP Access</b>	Check the <b>Enable HTTP Access</b> box to enable non-secure HTTP access. Leave the box blank to disable access to the cluster via HTTP.  <b>NOTE:</b> This option applies to both Web View and Web Root modes.

### To Connect via HTTPS or HTTP

1. Enter the **cluster name** (or Management IP address or any IP address from the node IP address pool) 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 SnapScale cluster name or URL, as shown in the following examples:

```
https://SnapNode2302216/Share1/my_files
https://10.10.5.23/Share1/my_files
```

2. Press **Enter**.

The **Web View** page opens.

## Configure the SnapScale Cluster as a Simple Web Server Using Web Root

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

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

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

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

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

**NOTE:** SnapScale 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 SnapScale cluster.

## Configuring Web Root

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

The screenshot shows the SnapScale Web Management Interface. The top navigation bar includes tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Web' sub-tab is active, showing options for authentication and web root configuration. The 'Enable Web Root' checkbox is checked. Below it, 'Allow Directory Listings' is unchecked. Under 'Automatically create and configure a WebRoot share', the radio button is selected. The 'Web Root Share' dropdown is set to 'SHARE1'. There is a list for 'Default Index File Names (0)' with '(none)' in it, and 'Add' and 'Delete' buttons. A note states 'index.html' will be used if no index files are defined. An important note at the bottom provides the URL <http://10.25.11.160/config>.

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

Option	Description
<b>Allow Directory Listings</b>	If <b>Allow Directory Listings</b> is checked and no user-defined index pages are configured or present, the browser opens to a page allowing browsing of all directories underneath the Web Root.  <b>NOTE:</b> Checking or unchecking this option only affects directory browsing in Web Root. It does not affect access to Web View directory browsing.
<b>Create and configure a Web Root share</b>	Select one of the following: <ul style="list-style-type: none"> <li>• <b>Automatically create and configure a Web Root share:</b> A share named “WebRoot” is automatically created. By default, the share is hidden from network browsing and has all network access protocols except HTTP/HTTPS enabled (as such, it can be accessed from a browser as the Web Root but can not be accessed via Web View). You can change these settings at <b>Security &gt; Shares</b>.</li> <li>• <b>Use existing share:</b> From the drop-down list of existing shares for selection, select a share and click the <b>Properties</b> button to edit the selected share's properties (see <b>Security &gt; Shares</b>).</li> </ul>

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

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

### Accessing the Web Management Interface when Web Root is Enabled

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

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

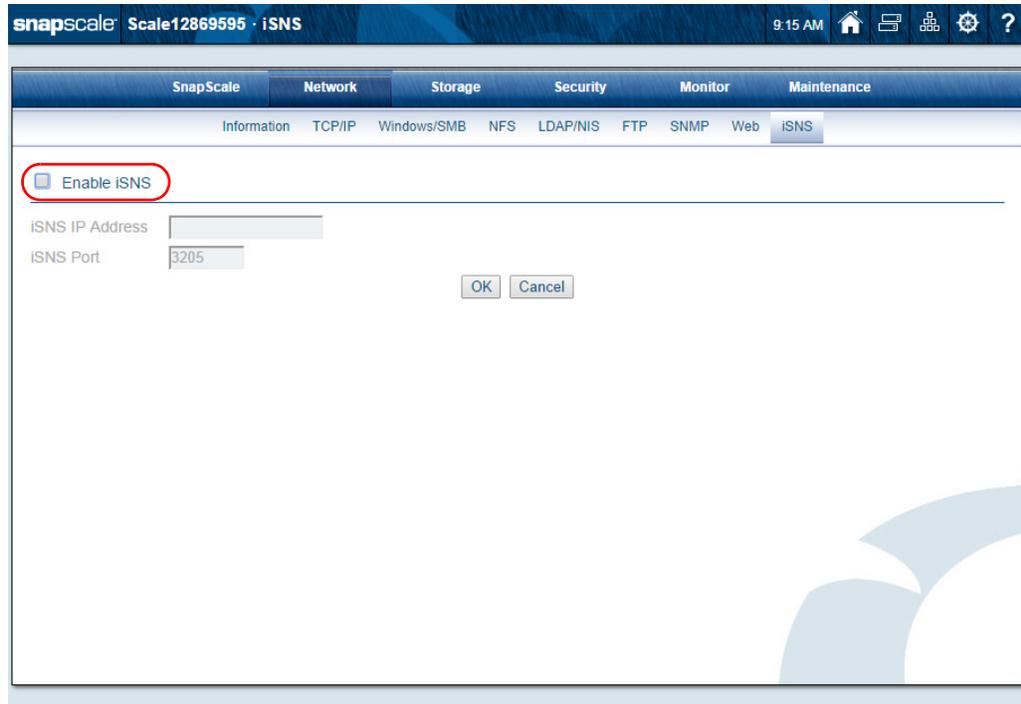
You are prompted for your User ID and password, then you are placed into the Web Management Interface.

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

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

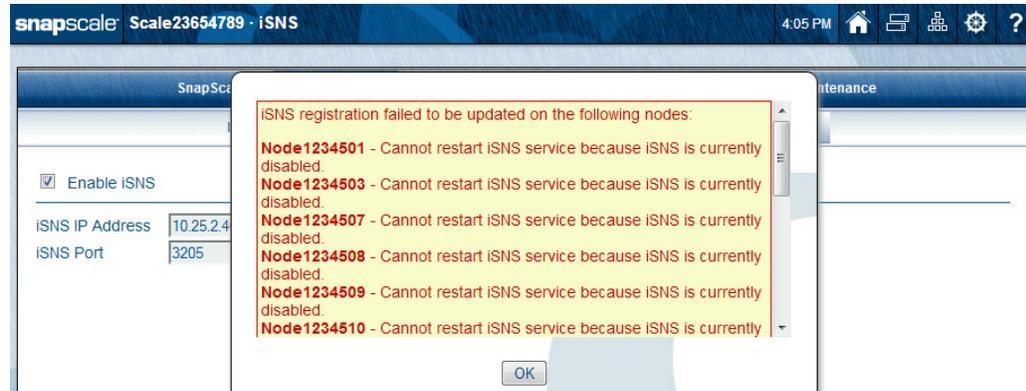
## iSNS Configuration

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



### Configure the iSNS Settings

1. If not already installed, install the iSNS service on a **Windows server**.  
Note the IP address of the server or workstation on which the iSNS service is installed.
2. Configure iSNS on the **SnapScale cluster**:
  - a. On the **Network > iSNS** page, check the **Enable iSNS** box.
  - b. Enter the **IP address** of the iSNS server.  
If the iSNS server does not use the default port, the iSNS port default value of 3205 can be changed on this page as well.
  - c. Click **OK**.  
A pop-up window shows the status and success/failure of the process. If there are missed registrations, see [Update iSNS Registration on page 83](#).

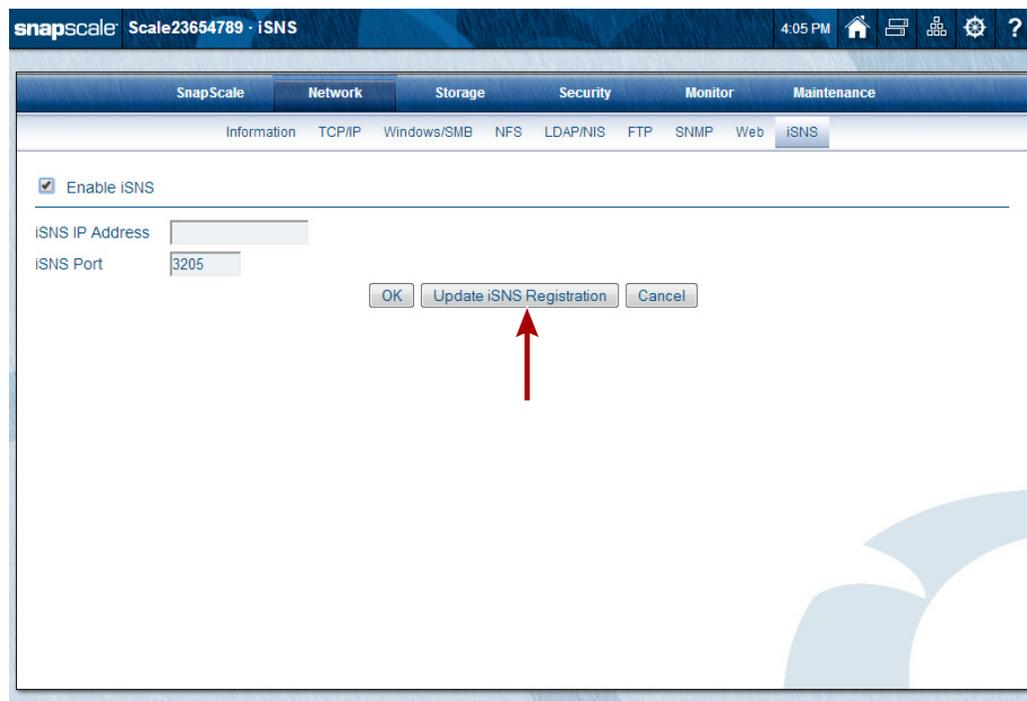


3. Configure the **iSCSI initiator** to discover iSCSI targets via the iSNS server.

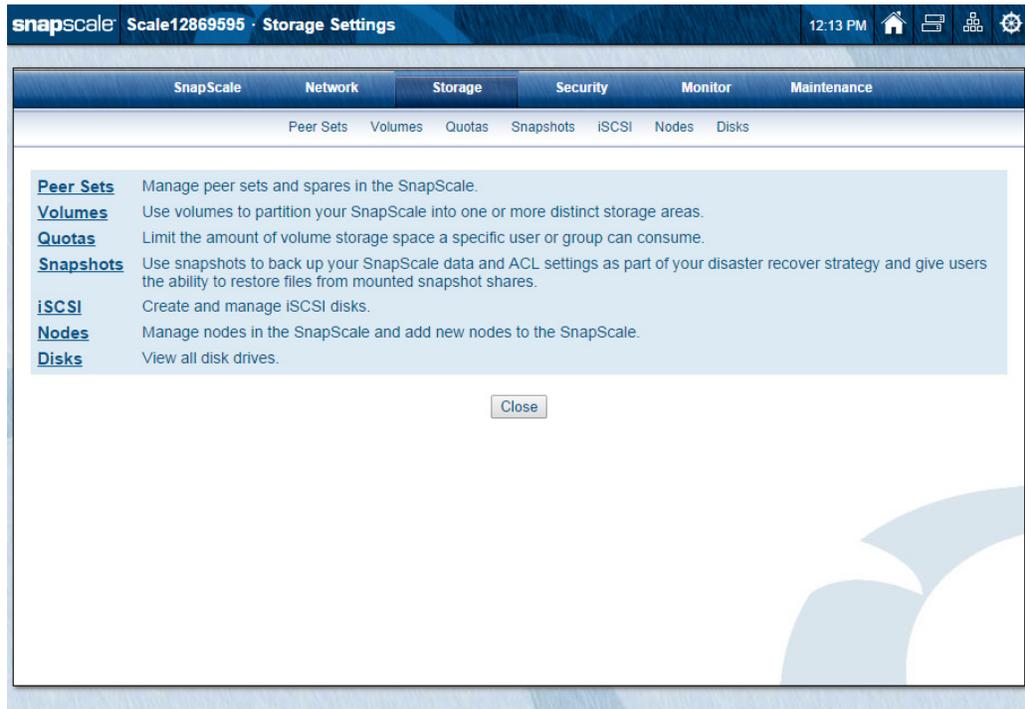
**NOTE:** After you have completed this procedure, all the iSCSI targets on the SnapScale cluster automatically appear in the Microsoft Initiators target list.

## Update iSNS Registration

Once iSNS is enabled, the **Update iSNS Registration** button is shown. It can be used to update any missed registrations that might occur.



From the storage default page (**Storage Settings**), you can access and configure the storage options for your SnapScale cluster including nodes and drives.



**Topics in Storage Options:**

- [Peer Sets](#)
- [Volumes](#)
- [Quotas](#)
- [Snapshots](#)
- [iSCSI Disks](#)
- [Nodes](#)
- [Disks](#)

## Peer Sets

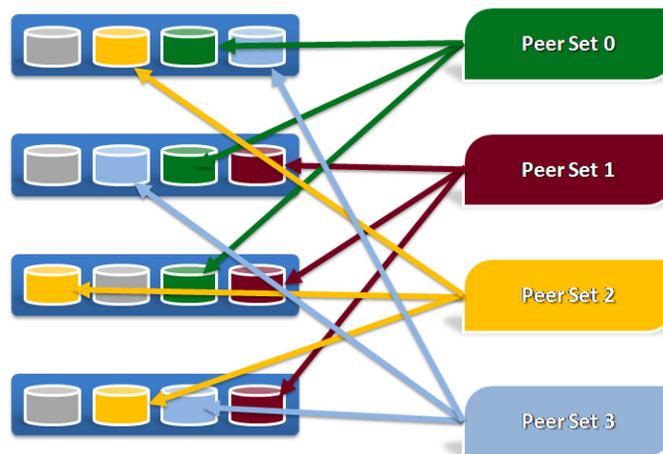
In a cluster, a node is a file server working in tandem with other nodes. The drives on every node are grouped into peer sets or hot spares. Each peer set contains two or three drives, depending on the Data Replication Count, that mirror the same data. To ensure availability, each drive in a peer set resides in a different node.

SnapScale aggregates all the storage on the peer sets in the cluster to form a unified data storage space for network client access. Data access is transparent between the cluster storage space and the peer sets so that users never directly access the peer sets.

When you create a cluster or add new nodes to an existing cluster, SnapScale automatically creates peer sets with the available drives. By distributing peer set members throughout the cluster, the system ensures that content is protected from failure of either individual drives or entire nodes. When they are created, peer sets are assigned a unique peer set ID.

Nodes can be added to expand cluster storage at any time. Based on the configuration settings, the additional drives are either used to create more peer sets or left as hot spares. Nodes can be removed from a cluster for replacement with a new node, and the drives in the replacement node are automatically synchronized with the existing peer sets.

On a four-node cluster configured for 3x replication count, four hot spares, and four drives per node, the peer set formation might look something like this:



Each peer set has members on three different nodes, shown below as peer set 0, 1, 2, and 3. Hot spares are automatically distributed throughout the cluster in order to replace any failed peer set member. When a peer set member fails, a hot spare is assigned from a node on which the peer set does not already have an active member.

The example above uses a 3x Data Replication Count, which means that each peer set contains three members, and as a result all data is replicated three times. The cluster can also be configured for a 2x Data Replication Count, in which case the distribution of two-member peer sets would be different. The system automatically determines which drives are used to form each peer set; you cannot choose them.

**NOTE:** The Data Replication Count can be decreased from 3x to 2x to increase cluster storage, but cannot be increased from 2x to 3x once the cluster is created.

## Peer Sets and Recovery

Though data on peer sets is served indirectly by the unified cluster storage space, access to files stored on a given peer set is dependent on the health of that peer set. When a drive in a peer set fails, data is served from the remaining peer set member drives. If there is a spare reserved for the cluster that does not exist on the same node as another active member of the peer set and is not smaller than other members, the peer set can claim the drive and rebuild the data (using the integrated RapidRebuild feature) onto that spare without administrator intervention.

Peer Set	Status	Member 1	Member 2	DSM
PeerSet16	RapidRebuild: 77% complete.	Node2414532: Disk 12 - 931.51 GB	Node2414528: Disk 2 - 1.82 TB	931.5 GB
PeerSet10	RapidRebuild: 69% complete.	Node2414532: Disk 8 - 931.51 GB	Node2414528: Disk 8 - 931.51 GB	OK
PeerSet4	RapidRebuild: 0% complete.	Node2414532: Disk 4 - 931.51 GB	Node2414528: Disk 4 - 931.51 GB	OK
PeerSet15	RapidRebuild: 0% complete.	Node2414538: Disk 12 - 931.51 GB	Node2414528: Disk 12 - 931.51 GB	OK
PeerSet0	OK	Node2414532: Disk 1 - 931.51 GB	Node2414538: Disk 1 - 931.51 GB	OK
PeerSet1	OK	Node2414532: Disk 2 - 931.51 GB	Node2414538: Disk 2 - 931.51 GB	OK
PeerSet2	OK	Node2414532: Disk 3 - 931.51 GB	Node2414538: Disk 3 - 931.51 GB	OK
PeerSet3	OK	Node2414538: Disk 4 - 931.51 GB	Node2414528: Disk 3 - 931.51 GB	OK
PeerSet5	OK	Node2414532: Disk 5 - 931.51 GB	Node2414538: Disk 5 - 931.51 GB	OK
PeerSet6	OK	Node2414528: Disk 5 - 931.51 GB	Node2414538: Disk 6 - 931.51 GB	OK
PeerSet7	OK	Node2414528: Disk 6 - 931.51 GB	Node2414532: Disk 6 - 931.51 GB	OK
PeerSet8	OK	Node2414532: Disk 7 - 931.51 GB	Node2414538: Disk 7 - 931.51 GB	OK
PeerSet9	OK	Node2414538: Disk 8 - 931.51 GB	Node2414528: Disk 7 - 931.51 GB	OK
PeerSet11	OK	Node2414532: Disk 9 - 931.51 GB	Node2414538: Disk 9 - 931.51 GB	OK
PeerSet12	OK	Node2414528: Disk 9 - 931.51 GB	Node2414538: Disk 10 - 931.51 GB	OK
PeerSet13	OK	Node2414528: Disk 10 - 931.51 GB	Node2414532: Disk 10 - 931.51 GB	OK
PeerSet14	OK	Node2414532: Disk 11 - 931.51 GB	Node2414538: Disk 11 - 931.51 GB	OK

If a peer set is missing one drive but at least one other drive is available, the peer set continues to be accessible but is in degraded mode. This table shows the different peer set statuses:

Peer Set Status	Failure Type	Data Availability
OK	The peer set drives are healthy and connected.	Data is fully available for read and write.
RapidRebuild	Spare made available to rebuild the peer set using RapidRebuild.	Data is fully available for read and write.
Degraded	One drive missing from the peer set	Data is fully available for read and write.
Degraded – Cannot repair; no spares	The peer set cannot be repaired because there are no spare drives.	Data is fully available for read and write.
Degraded – Cannot repair; spares too small	The peer set cannot be repaired because all eligible spares are too small.	Data is fully available for read and write.

Peer Set Status	Failure Type	Data Availability
<b>Degraded – Cannot repair; spares on same node</b>	The peer set cannot be repaired because the only eligible spares are located on the same node as an active member of the peer set.	Data is fully available for read and write.
<b>Failed</b>	All drives in peer set have failed.	No availability. Contact Overland Support.
<b>Initializing</b>	The peer set is being created or initialized.	Data is not yet available.
<b>Inconsistent</b>	The peer set has more members than the data replication count.	Contact Overland Support.

## Peer Set Utilization

Each file's data is spread across multiple peer sets, and the cluster automatically distributes data for different files throughout the peer sets in the cluster. Metadata for files and directories is independently distributed among different peer sets using a hash algorithm for optimum performance and protection.

## Peer Set Basics

New drives are initially configured automatically as spare drives. Subsequently, if enough spare drives exist on different nodes to construct new peer sets but still satisfy the spare count setting, the SnapScale automatically creates new peer sets and expand cluster storage space.

Drives in a cluster do not all need to have the same capacity, but drives in a given peer set should have the same capacity or space is wasted on the larger drives.

The following points must be observed in regards to drives used in the cluster:

- The drives in a cluster must all be the same type of drive (such as SAS) and the same rotational speed.
- The storage capacity of a peer set is limited to the smallest capacity drive in the peer set.

In case of peer set drive failure, RAINcloudOS continues to serve data reads and writes to that peer set from another member of the peer set as long as the peer set is not offline. If clients are currently using data on the peer set, it continues to operate as-is.

**Data Replication Count** is an administrator specified, cluster-wide count of the degree of redundancy of data on the cluster. The Data Replication Count can be either 2x or 3x and determines the number of drives (2 or 3) that make up each peer set.

## Hot Spares

Each node can have a number of hot spares in the event of a drive failure. The total number of hot spares for the cluster is user selectable. A suggested number of hot spares for various node sizes is provided. If a peer set member drive fails, data from a healthy peer set drive on another node is re-synced onto an available spare on any node that doesn't have another active member of that peer set, and the spare then becomes a member of that peer set.

Drives added to nodes as additions or as replacements to failed drives are automatically configured as spares. If enough spares exist across different nodes to satisfy the Data Replication Count and the spare drives count, the cluster automatically creates a new peer set out of available spare drives.

## Snapshot Limitations

- All snapshots are deleted when:
  - New peer sets are automatically created when new drives are installed.
  - One or more new nodes are added to a cluster.
  - If a complete peer set fails.
- A Snapshot may be deleted if:
  - Any peer set member drive runs out of snapshot space.
  - A second member of a peer set (containing *unique* snapshot data) fails, even though the main file system data may still be healthy.

## Peer Sets Page

The screenshot shows the SnapScale interface for the Peer Sets page. The page title is "Scale12869595 · Peer Sets". The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The Storage menu is expanded, showing Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. The main content area displays "5 peer sets. Data replication count: 2x" and "Active spares: 2". Below this is a table with the following data:

Peer Set	Status	Member 1	Member 2	DSM
PeerSet0	OK	VM-Node12869595: Disk 1 - 50 GB	VM-Node9715283: Disk 1 - 50 GB	OK
PeerSet1	OK	VM-Node12869595: Disk 3 - 50 GB	VM-Node14122451: Disk 3 - 50 GB	OK
PeerSet2	OK	VM-Node12869595: Disk 2 - 50 GB	VM-Node14122451: Disk 2 - 50 GB	OK
PeerSet3	OK	VM-Node9715283: Disk 3 - 50 GB	VM-Node12869595: Disk 4 - 50 GB	OK
PeerSet4	OK	VM-Node9715283: Disk 4 - 50 GB	VM-Node14122451: Disk 1 - 50 GB	OK

Below the table are buttons for Spare Disks, Spare Distributor, Data Balancer, Refresh, and Close.

The following table details the items listed on this page:

Option	Description
<b># Peer Sets</b> (above table, left)	Displays the total number (#) of peer sets configured and shown in the table.
<b>Data Replication Count</b> (above table, left)	Displays the cluster-wide replication count and links to the <b>SnapScale Properties</b> page.

Option	Description
<b>Active Spare Disks</b> (above table, right)	Shows the number of drives allocated as spares. They are broken out to show the status of spares and the number of spares with that status. <ul style="list-style-type: none"> <li>• <b>OK</b> – A blue icon (🟢) indicates spares are active and available.</li> <li>• <b>Too Small</b> – A yellow icon (🟡) indicates spares are too small to be used in some of the peer sets. A yellow icon with an “X” (🟡X) indicates a spare is too small to use with any available peer set.</li> <li>• <b>Failed</b> – A red icon (🔴) indicates spares have failed.</li> </ul> Clicking this link opens the <b>Spare Disks</b> page. This is the same as clicking the <b>Spare Disks</b> button at the bottom of the page.
<b>Peer Set</b>	Lists the peer set name and shows a usage bar. Position the cursor over a name (or usage bar) to show the percentage and actual amount of storage space used.
<b>Status</b>	Shows the current status. Refer to <a href="#">Peer Sets and Recovery on page 86</a> for complete details.
<b>Member 1</b>	Shows the node, drive/slot number, and the size of the first member of this peer set. Click to view the <b>Disks</b> page and identify the specific disk drive's location.
<b>Member 2</b>	Shows the node, drive/slot number, and the size of the second member of this peer set. Click to view the <b>Disks</b> page and identify the specific disk drive's location.
<b>Member 3 (if shown)</b>	Shows the node, drive/slot number, and the size of the third member of this peer set when the <b>Data Replication Count</b> is set to “3x.” Click to view the <b>Disks</b> page and identify the specific disk drive's location.
<b>DSM</b> (Drive Size Mismatch)	Shows either <b>OK</b> or mismatch size difference. If the member drives are not the exact same size, then capacity is limited to the smallest drive in the peer set, and extra space on larger drives is wasted. In this case, the size displayed reflects the unutilized capacity of the peer set. Position the cursor over a name (or usage bar) to show the unutilized capacity of the peer set.
<b>Spare Disks (button)</b>	Launches the <b>Spare Disks</b> page. See <a href="#">Spare Disks Page on page 90</a> .
<b>Spare Distributor (button)</b>	Launches the <b>Spare Distributor</b> page. See <a href="#">Spare Distributor on page 91</a> .
<b>Data Balancer (button)</b>	Launches the <b>Data Balancer</b> page. See <a href="#">Data Balancer on page 93</a> .
<b>Refresh (🔄 button)</b>	Refreshes the page when clicked.
<b>Close (button)</b>	Closes this page and returns to the <b>Storage Settings</b> page.

## Spare Disks Page

When you click the **Spare Disks** button (or the **Active spares** link on the upper right above the table on the **Peer Sets** page), the **Spare Disks** page opens.

2 spare disks. Active spares usable by any existing peer set: 2 Spare disks setting: 2

Spare Disk	Node	Slot	Spare Status ▲
<a href="#">50 GB SAS</a>	<a href="#">VM-Node14122451</a>	4	OK
<a href="#">50 GB SAS</a>	<a href="#">VM-Node9715283</a>	2	OK

Note: The minimum size required for a spare disk to be able to repair any existing peer set is 50 GB.

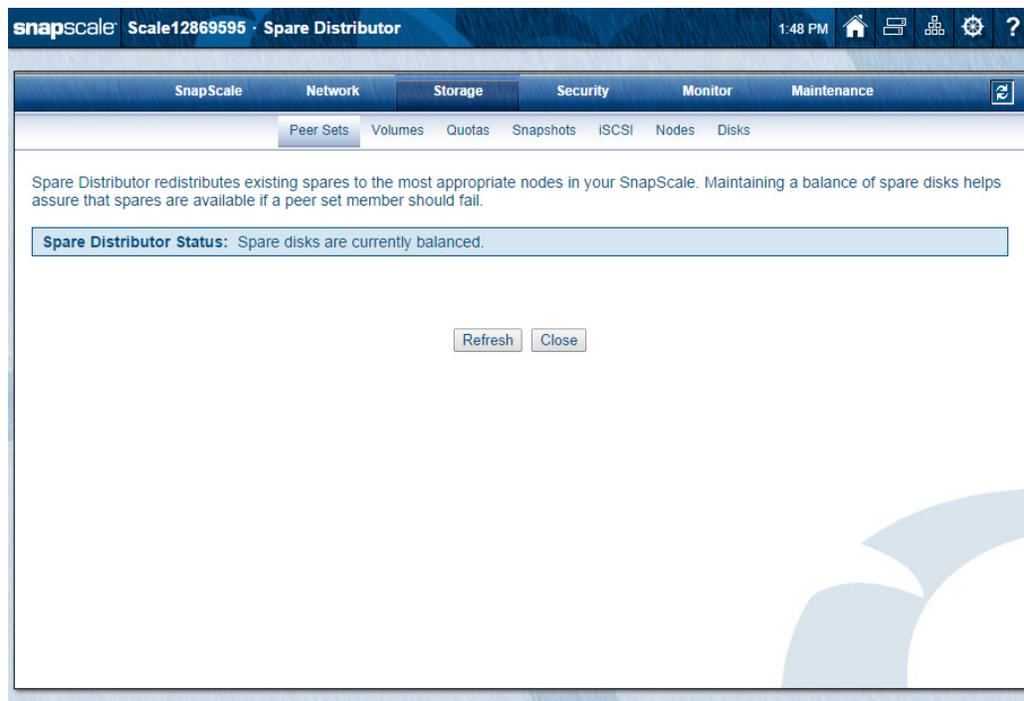
The following table details the items listed on this page:

Option	Description
<b># Spare Disks</b> (above table, left)	Displays the total number (#) of spare drives configured and shown in the table.
<b>Active Spares</b> (above table, left)	Displays the cluster-wide number of spares usable by all peer sets.
<b>Spare Disks Setting</b> (above table, right)	Displays the quantity set for spare drives. Clicking this link takes you to the <b>SnapScale Properties</b> page to edit the setting.  <b>NOTE:</b> This setting may not equal the number of spare drives currently displayed if there are fewer spare drives available than the setting specifies, or if there is an insufficient number of extra drives to automatically create a new peer set and satisfy the cluster's Data Replication Count.
<b>Spare Disk</b>	Displays disk drive capacity and type. Click a name in the column to open the <b>Disks</b> page and identify a specific disk drive's location.
<b>Node</b>	Displays the name of the node in which the drive resides. Click a name in the column to open the <b>Node Properties</b> page for the specific node.
<b>Slot</b>	Displays the slot number of the listed node where this spare drive is located.

Option	Description
<b>Spare Status</b>	Shows the current status: <ul style="list-style-type: none"> <li>• <b>OK</b> – Spare drive is healthy and can be used by all peer sets.</li> <li>• <b>Spare Too Small</b> – Spare is too small to either repair any existing peer sets or repair <i>n</i> existing peer sets.</li> <li>• <b>Failed</b> – Spare drive has failed.</li> </ul>
<b>Refresh</b> (  button)	Refreshes the page when clicked.
<b>Close</b> (button)	Closes this page and returns to the main <b>Peer Sets</b> page.

## Spare Distributor

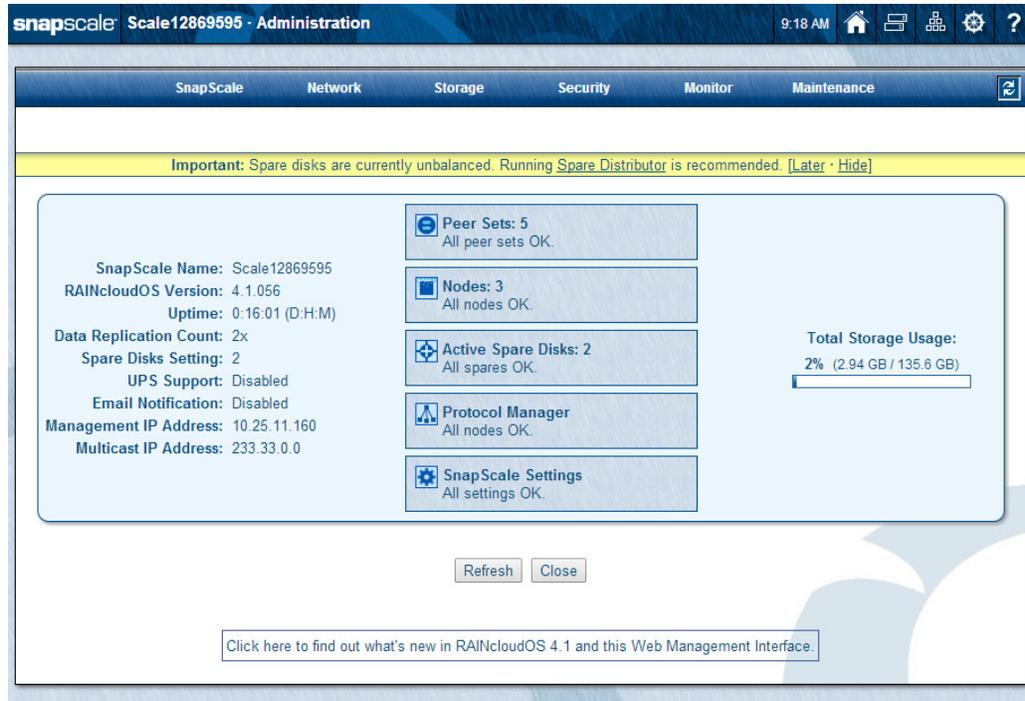
Spare Distributor (formerly the Spare Disk Balancer) evenly redistributes spares and peer set members across the cluster nodes. Maintaining a balance of spare drives helps ensure that spares are available if a peer set member should fail.



### Using Spare Distributor

When the cluster detects an uneven distribution of spare drives, an alert banner is displayed in the Web Management Interface and the **Spare Distributor** page is enabled.

**NOTE:** You can click [Later](#) to turn off the alert for 24 hours or [Hide](#) to dismiss the alert.

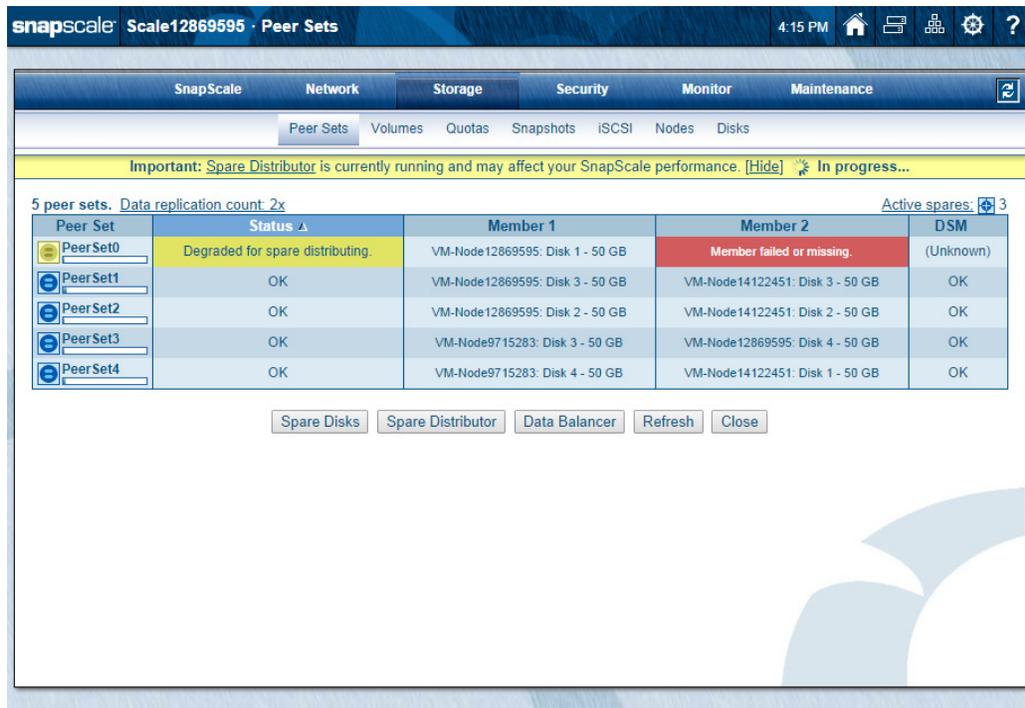


1. Go to **Storage > Peer Sets > Spare Distributor**.

If responding to an alert, you can click the Spare Distributor link in the alert to go directly to the page.

2. Click the **Start Spare Distributor** button to start the process.

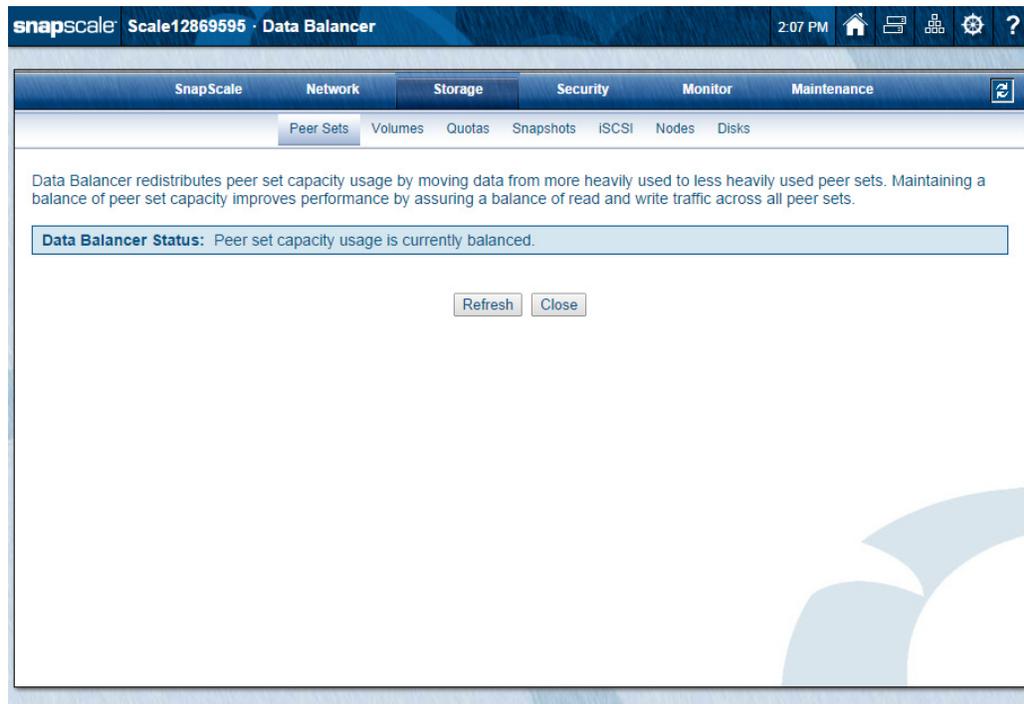
The Spare Distributor redistributes spares and peer set members across the cluster nodes to provide spares on different nodes for better spare availability. Go to the **Peer Sets** page to view the status of the balancer.



**NOTE:** If needed, click **Stop Spare Distributor** on the **Spare Distributor** page to stop the operation. Any peer sets currently degraded and being rebuilt by the Spare Distributor will continue with the rebuilding process until completed.

## Data Balancer

Data Balancer (formerly Capacity Balancer) redistributes peer set utilization by moving data from more to less heavily used peer sets. Maintaining a balance of peer set capacity improves performance by assuring a balance of read and write traffic across all peer sets.



### Using Data Balancer

If the peer set data utilization becomes unbalanced, an alert banner is displayed in the Web Management Interface.

**NOTE:** You can click **Later** to turn off the alert for 24 hours or **Hide** to dismiss the alert.

The screenshot shows the SnapScale Administration page for instance Scale12869595. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. A yellow alert banner at the top states: "Important: Spare disks are currently unbalanced. Running Spare Distributor is recommended. [Later - Hide]". The main content area is divided into two columns. The left column lists system information: SnapScale Name: Scale12869595, RAINcloudOS Version: 4.1.056, Uptime: 0:16:01 (D:H:M), Data Replication Count: 2x, Spare Disks Setting: 2, UPS Support: Disabled, Email Notification: Disabled, Management IP Address: 10.25.11.160, and Multicast IP Address: 233.33.0.0. The right column contains five status boxes: Peer Sets: 5 (All peer sets OK), Nodes: 3 (All nodes OK), Active Spare Disks: 2 (All spares OK), Protocol Manager (All nodes OK), and SnapScale Settings (All settings OK). A "Total Storage Usage" section shows 2% (2.94 GB / 135.6 GB) with a progress bar. At the bottom, there are "Refresh" and "Close" buttons, and a link: "Click here to find out what's new in RAINcloudOS 4.1 and this Web Management Interface."

1. Go to **Storage > Peer Sets > Data Balancer**.  
If responding to an alert, you can click the **Data Balancer** link in the alert to go directly to the page.
2. Review the default **File Size Limit** and change it, if needed.

The screenshot shows the SnapScale Data Balancer configuration page for instance Scale1234501. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The Storage section is active, with sub-tabs for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. A yellow alert banner states: "Data Balancer Status: Peer set capacity usage is currently unbalanced. Running Data Balancer is recommended." Below the alert, a text box explains: "Data Balancer will only move files that are not larger than the File Size Limit. If peer set capacity usage remains unbalanced after running Data Balancer, you can increase the file size limit to accommodate your storage environment and then re-run Data Balancer. (View online help for more information.)" A "File Size Limit" input field is set to "2" GB, with a red circle around it. A note next to it says: "Max. file size limit is 98.01 GB, based upon existing peer set capacity." At the bottom, there are "Start Data Balancer", "Refresh", and "Close" buttons.

The **File Size Limit** represents the maximum size of a file the Data Balancer will attempt to move to rebalance peer set consumption. The default is 2GB.

- Click the **Start Data Balancer** button to start the process.

The Data Balancer moves files between peer sets to improve performance and usability. A table is displayed showing that the Data Balancer is running and the percent completed. If needed, click **Stop Data Balancer** to end the operation.

The screenshot shows the SnapScale Data Balancer status page. The page title is "Scale1234501 · Data Balancer". The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The Storage section is active, with sub-tabs for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. The main content area displays the Data Balancer status as "Data Balancer is currently running." Below this is a table titled "Status of currently running Data Balancer operation:" with the following data:

Status of currently running Data Balancer operation:	
Operation status:	<b>In progress.</b>
Size processed:	512.08 GB of 2.08 TB
File size limit:	2 GB
Start time:	2013-09-17 4:57:08 PM
Elapsed time:	47 seconds
Est. time remaining:	2 minutes (23% complete)

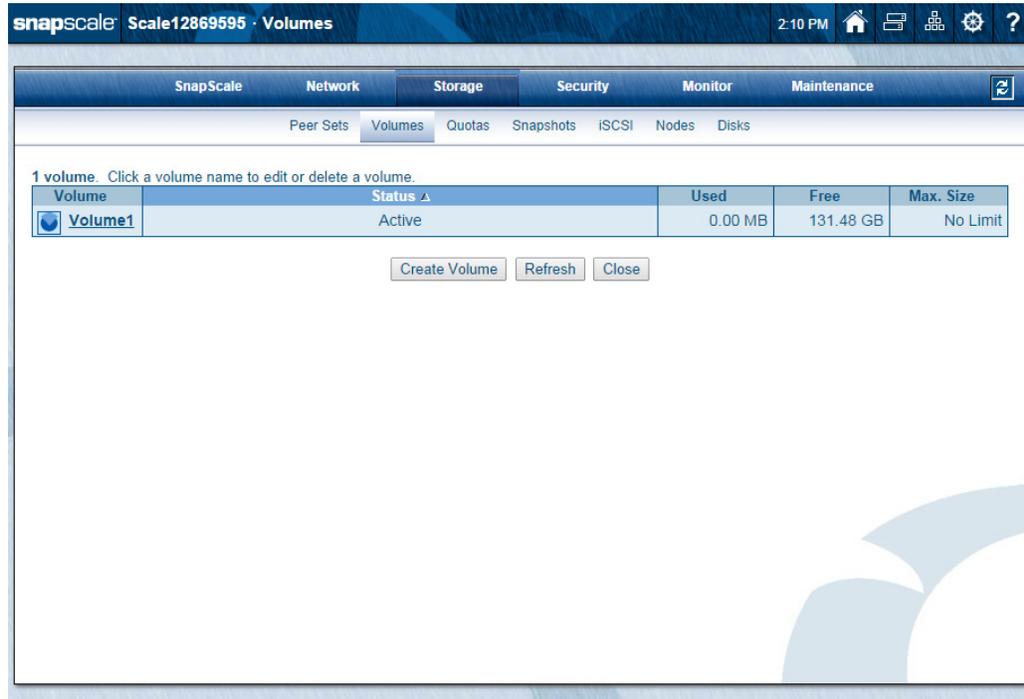
At the bottom of the page, there are three buttons: "Stop Data Balancer", "Refresh", and "Close".

**NOTE:** The cluster continues to be available for client access during the process. The Data Balancer will skip any file that is currently opened by clients, and will abort moving a file if a client opens it during the move.

An alert banner is displayed on any Administration-level page showing the progress.

## Volumes

Use the **Volumes** page (**Storage > Volumes**) to manage the volumes that have been created.



From this page, you can:

- Create a new volume.
- Edit or delete the volume (by clicking the name to access the **Properties** page).

All the peer sets are unified into a single cluster storage space that can be accessed from any node thus providing multiple access points. One or more volumes can be created to provision the cluster storage:

- All volumes share the same cluster storage space and are thinly provisioned to provide better utilization rates of the space.
- Volumes can be configured with a maximum size setting (quota) to prevent one volume from consuming too much shared cluster storage space. See [Creating Volumes on page 96](#).

### Creating Volumes

By default, the full cluster storage space is accessible as one large storage space. However, the storage space can be divided into multiple volumes in order to thinly provision space for specific projects, departments, or roles. Volumes can be constrained to use no more than a certain amount of space available in the clustered storage space.

1. At **Storage > Volumes**, click the **Create Volume** button to open the **Create Volume** page.

snapscale Scale12869595 · Create Volume 2:13 PM

SnapScale Network **Storage** Security Monitor Maintenance

Peer Sets **Volumes** Quotas Snapshots iSCSI Nodes Disks

Specify the name of the new volume and click Create Volume.

Volume Name

Specify the maximum size for this volume:

No Limit (allow volume to consume space on SnapScale as needed)

Limit volume to  GB (Max. size is 135.60 GB)

2. Make any necessary changes to the **options**.
  - It is recommended to enter a **Volume Name** to easily identify the specific volume.
  - If desired, keep the default of **No Limit** to allow the volume to consume an unlimited amount of cluster storage.
  - Otherwise, enter a **size**, changing the measurement units if needed.
3. Click the **Create Volume** button again. After a few moments, a confirmation page is shown:

snapscale Scale12869595 · Create Volume 2:15 PM

SnapScale Network **Storage** Security Monitor Maintenance

Peer Sets **Volumes** Quotas Snapshots iSCSI Nodes Disks

Volume **Volume2** has been successfully created.

Click Create Share to create a share that users will connect to when accessing the volume.  
Click Close to return to the main Volumes page without creating a share.

- At the confirmation page, click **Create Share** to create a share pointing to this volume (takes you to the **Shares** page).

The screenshot shows the 'Create Share' page in the SnapScale web interface. The page title is 'SnapScale Scale12869595 · Create Share'. The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Security' menu is expanded, showing Security Guides, Shares, Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. The main content area contains the following fields and options:

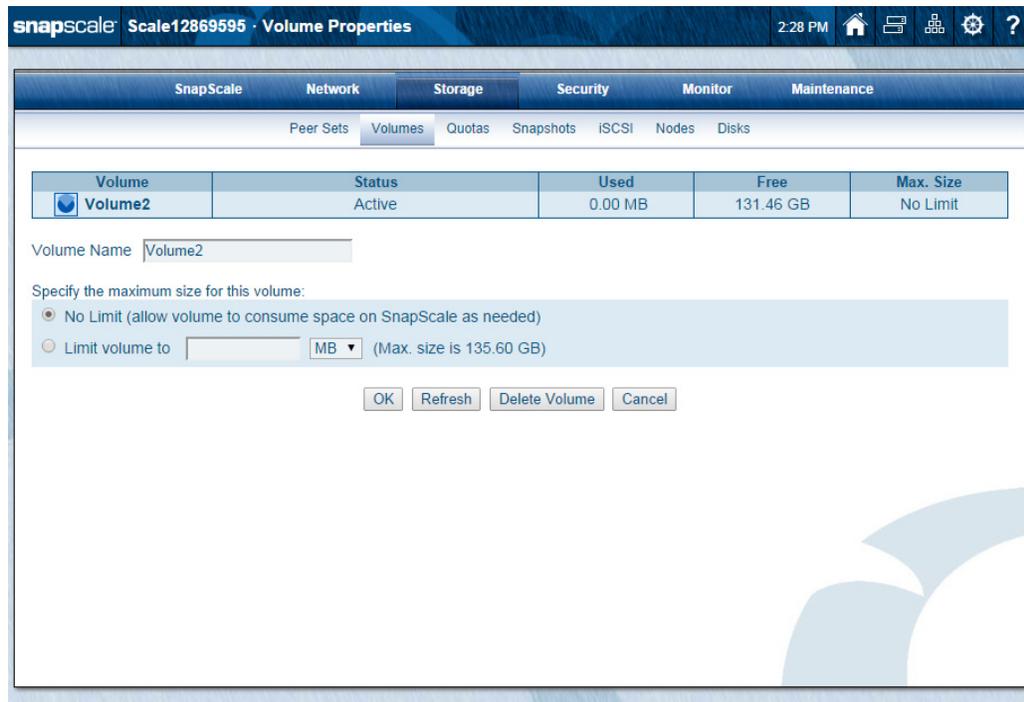
- Name:** SHARE3
- Volume:** Volume1
- Path:** / (with a 'Browse' button)
- Description:** (optional)
- Security model for path:** Windows/Unix
- Create share with full read and write access for all users
- Create share with Admin-only access and proceed to Share Access page
- [Advanced Share Properties >>](#) (highlighted with a red circle and a red arrow pointing to it with the text 'Click to display options shown below')
- 

**NOTE:** The snapshot options at the bottom are only shown if snapshot space has been reserved.

- Enter the appropriate **data** and select the necessary options, then click **Create Share**. Additional options can be accessed by clicking the **Advanced Share Properties** link at the bottom. See [Shares on page 147](#) for complete details. When share creation is completed, the **Security > Shares** page is shown with the new share listed in the table.
- Click **Close** to go to the **Security** menu page.

## Edit Volume Properties

By clicking a volume's name on the **Storage > Volumes** page, details of that particular volume are shown on the **Volume Properties** page.



From this secondary page, you can:

- Change the volume name.
- Set maximum volume size (specific limit or no limit).
- Delete the entire volume.

### Rename a Volume

In the **Volume Name** field, enter a unique volume name of 32 alphanumeric characters and spaces, then click **OK**.

### Specify Maximum Volume Size

There are two options controlling the maximum size of a volume:

- **No Limit** – This is the recommended option because it allows the volume to consume space as needed.
- **Limit Volume to** – Establish a maximum volume size limit by entering the amount and selecting a unit of measure (MB, GB, TB, or PB). The volume then grows in size until it reaches its maximum. If email notification has been enabled, alerts are sent as the maximum is approached. (To enable email notification, see [Email Notification on page 226](#))

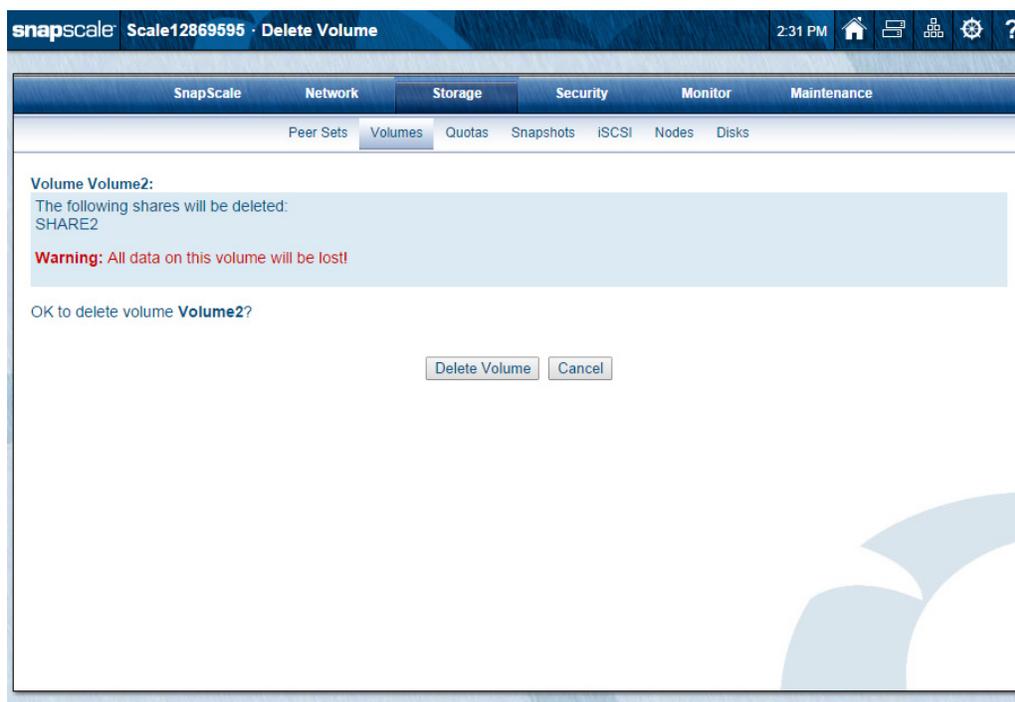
**NOTE:** If you reset the maximum size of a volume to less than its current size, the volume is treated as full and no more data can be written to it until the actual space consumed drops below the maximum size again. When done, click **OK**.

## Deleting Volumes

To delete a volume, at the **Volume Properties** page, click the **Delete Volume** button. At the confirmation page, click the **Delete Volume** button again. You are returned to the **Volumes** page and the volume is deleted in the background.

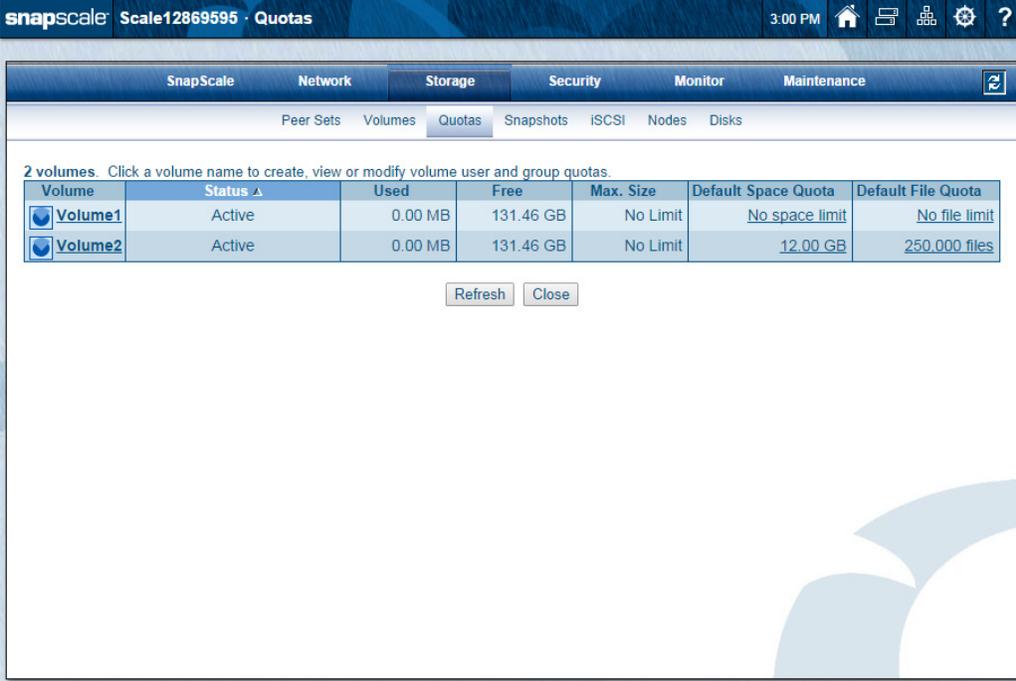


**CAUTION:** Deleting a volume deletes all the shares and data on the volume.



## Quotas

Quotas are configured by accessing the **Storage > Quotas** page of the Web Management Interface. This default page shows all volumes on the cluster and their space/file quotas.



The screenshot shows the SnapScale Quotas page. At the top, there is a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below this, there are sub-tabs for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. The main content area displays a table with 2 volumes. The table has columns for Volume, Status, Used, Free, Max. Size, Default Space Quota, and Default File Quota. Below the table are Refresh and Close buttons.

Volume	Status <span>Δ</span>	Used	Free	Max. Size	Default Space Quota	Default File Quota
<input checked="" type="checkbox"/> Volume1	Active	0.00 MB	131.46 GB	No Limit	No space limit	No file limit
<input checked="" type="checkbox"/> Volume2	Active	0.00 MB	131.46 GB	No Limit	12.00 GB	250,000 files

Assigning quotas ensures that no one user or group consumes a disproportionate amount of volume capacity measured by either space consumed or number of files created. The **Quotas** page also displays space consumed and files created by each user or LDAP/NIS group regardless of whether a quota is applied to them, allowing for precise tracking of usage patterns. You can set individual quotas for any LDAP, NIS, Windows domain, or local user known to the SnapScale. Group quotas are available only for LDAP/NIS groups.

For users and groups, there are no pre-assigned default quotas on the SnapScale. When quotas are assigned, you can assign a default space or file 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.

In calculating usage, the SnapScale looks at all the files on a volume 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 files are copied to the cluster, their size and count are applied against both the applicable user and LDAP/NIS group quotas.

### Default Quotas

On the main **Quotas** page (**Storage > Quotas**), the last two columns of the table show the default quotas for disk space and number of files. To change these settings, click the number (or no space/file limit text) in the row under the default space or file quota column. A page is shown for the appropriate quota type options:

## Default Space Quota Page

Specify the default space quota for users on volume **Volume2**.

**Note:** By default, if you do not assign a space quota to a user, that user will automatically be assigned the default space quota.

Default space quota for users:

- Allow user to consume the entire disk (no space limit)
- Limit user to  GB (Max. size is 135.6 GB)

OK Cancel

To make changes, choose to either use the entire disk or a space of a specific size. For a specific size, enter the maximum amount and select the units. Click **OK** to accept.

## Default File Quota Page

Specify the default file quota for users on volume **Volume2**.

**Note:** By default, if you do not assign a file quota to a user, that user will automatically be assigned the default file quota.

Default file quota for users:

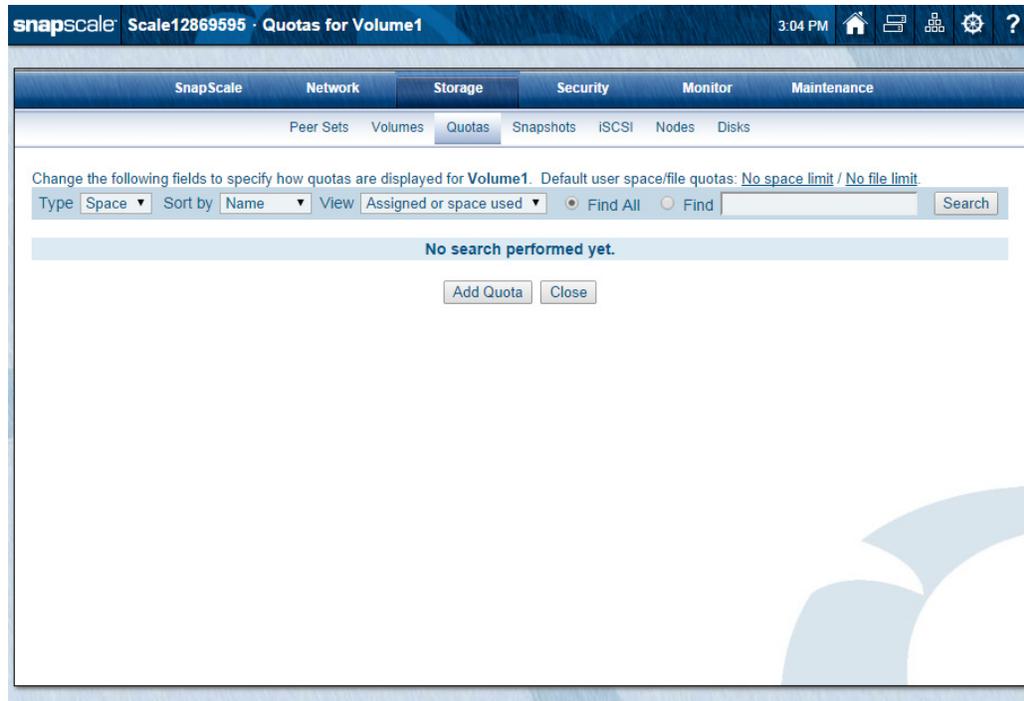
- Allow user to create any number of files (no file limit)
- Limit user to  files. (Max. assignable limit is 1,000,000,000 files)

OK Cancel

To make changes, choose either to have no limit or a specific number of files. For a specific limit, enter the maximum number of files. Click **OK** to accept.

## Quotas for Volume Page

From the **Quotas** page, you can create, view, or modify user and group quotas for a volume by clicking the volume's name in the **Volume** column on the far left. A **Quotas for Volume** page is displayed:



The page shows the available search and view options for the selected volume and the **Default user space/file quotas**. The two defaults shown can be either an amount or a text string:

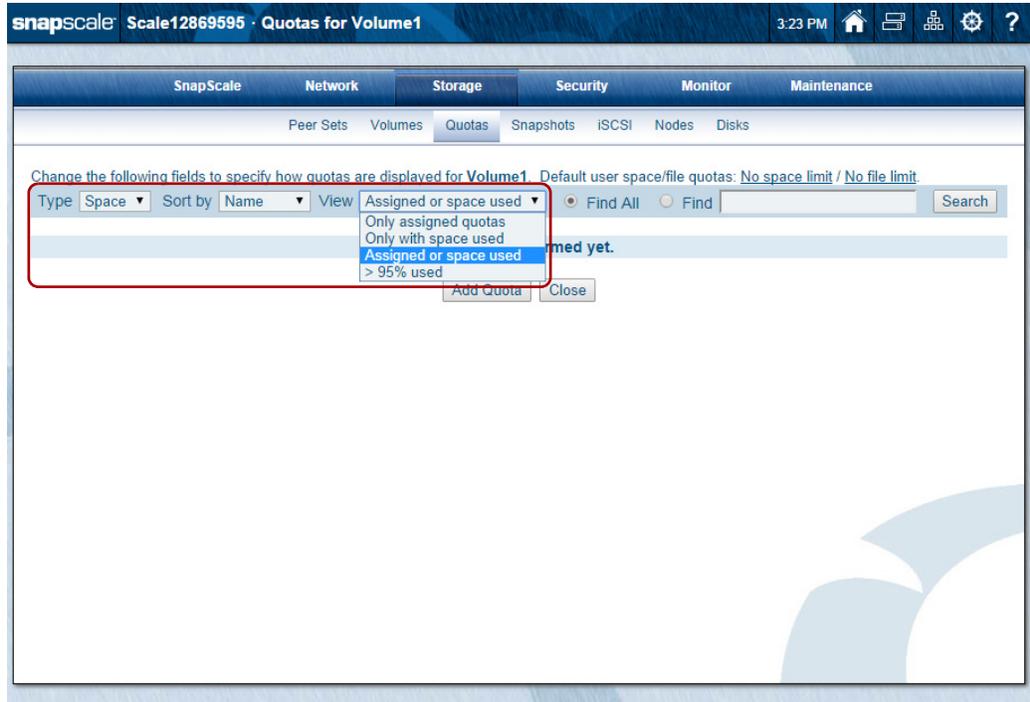
- **An amount** – the default quota size or file count assigned to users in that volume who do not have a specific quota assigned to them.
- **A text string** – the text strings **No space limit** and/or **No file limit** are displayed when quotas are enabled but the default space size and/or file count of no limits are configured for users in that volume. This means the users can consume the entire disk, or as many files as desired, respectively.

The space and file count limits also double as links to access the Web Management Interface pages where default space and file quotas can be configured.

### Search for Quotas or Space Consumed by a User or LDAP/NIS Group

To narrow down the list shown on this page or find a specific user or LDAP/NIS group, first use the search bar just under **Default user space/file quotas**.

1. Select the **Type**, **Sort By**, and **View** parameters.



- **Type** – Choose **Space** or **File**.
  - **Sort by** – Select **Name**, **Limit**, **Used**, or **Used (%)**.
  - **View** – Choose one of these view options:
    - **Only assigned quotas**
    - **Only with files used / Only with space used** (depends on Type setting)
    - **Assigned or files used**
    - **> 95% used**
2. Select **Find All** or **Find**.  
When entering a search string for **Find**:
    - Returned results will include all users and groups whose name **contains** the string entered.
    - To search a specific Windows or NIS domain, enter the domain name followed by a slash (/) or backslash (\) before the search string.
    - To search only local users and groups, enter “local” followed by a backslash (\) before the search string.
  3. Click **Search**.  
A detailed list of users or LDAP/NIS groups that match the parameters is displayed including the quota and space used numbers:

Change the following fields to specify how quotas are displayed for **Volume1**. Default user space/file quotas: No space limit / No file limit.

Type  Sort by  View   Find All  Find  Search

Quotas: 2 found. Note: A quota with a default limit is displayed as "(No space limit)".

User or Group (click to edit)	Domain	Quota	Used	Used (%)
AI*	Local Users	15.00 GB	0.00 MB	0%
Freddy	Local Users	10.00 GB	0.00 MB	0%

\*User or group is ID mapped.

**NOTE:** The search results returned may be automatically limited. Fine tune your search by using a more specific string to return a shorter list or the name desired.

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

No parentheses around the quota amount indicates a specific quota has been assigned. If the default quota limit is set to "no limit" but a particular user's or group's quota is set to 750GB, then "750 GB" is shown instead of the default "(no limit)."

The one exception to this is LDAP/NIS groups. They don't use a volume default quota, so "no limit" (without parentheses) is shown.

If Windows domain users or groups have been ID mapped to local, LDAP, or NIS users or groups, the **Quotas** page may display either of the following:

- **ID Mapped User/Group** - The consumption of the current ID-mapped user/group pair is identified by the name of the Windows domain user/group with an asterisk and footnote (for example, a footnote saying **\*User or group is ID mapped.**)
- **Unknown UID/GID** - If ID mapping changes are not applied to the file system, consumption (if any) by the former UID/GID of the Windows domain user/group is represented with the UID/GID and the name of the domain user/group in parentheses. Quotas for these unknown UIDs/GIDs cannot be edited.

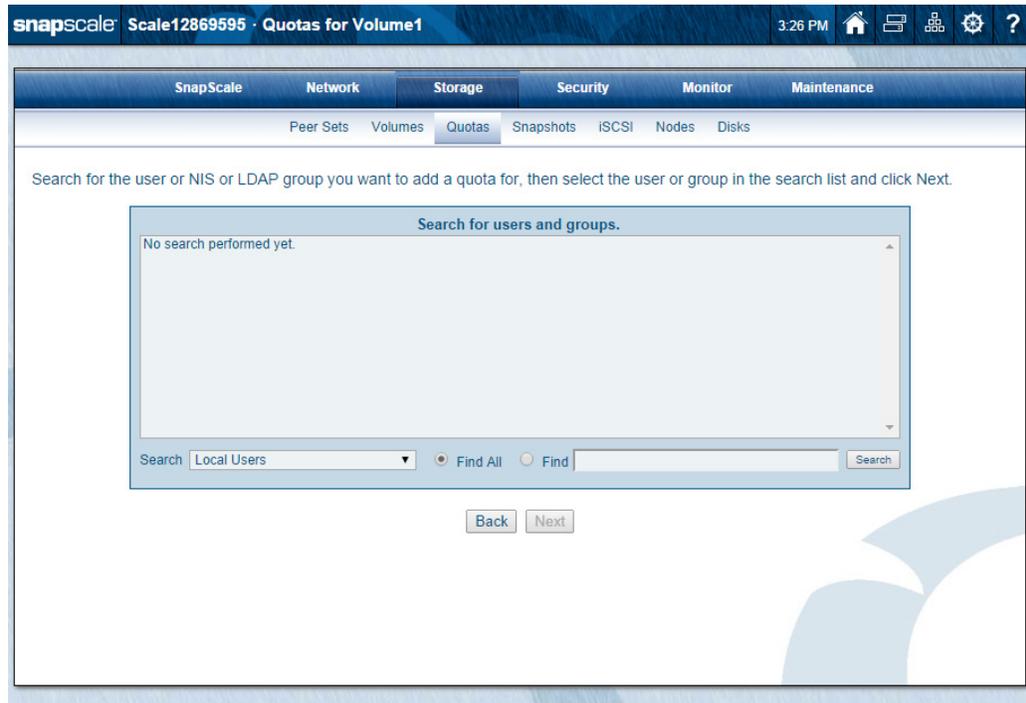
4. To make **changes**, click the user or LDAP/NIS group name.

See [Editing or Removing Quotas on page 108](#).

## Add Quota Wizard

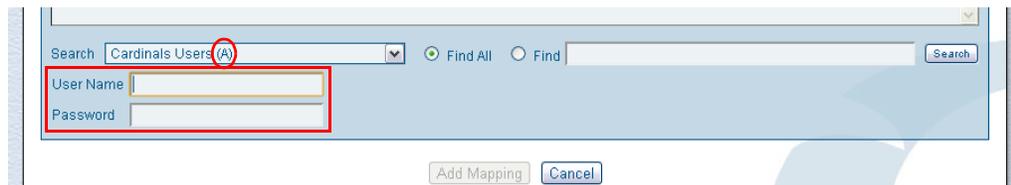
1. Click the **Volume name** link on the **Quotas** initial page to open the **Quotas for Volume** page for that volume.

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

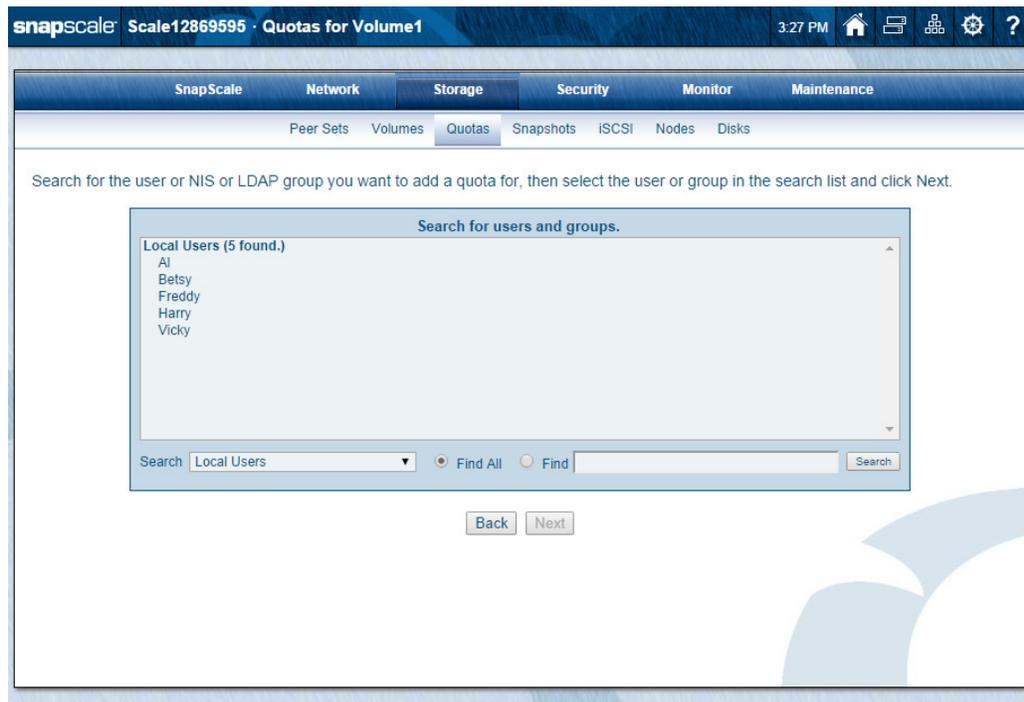


3. To search for a user or LDAP/NIS group, select the local or domain option from the Search drop-down list, enter the search string (or select Find All), and click **Search**.

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



Returned results will include all users and LDAP/NIS groups whose name **begins** with the string entered in the Search field.



Note the following caveats:

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

- From the search results, select the appropriate **user or LDAP/NIS group**, and click **Next** to show the configuration page for that user or group.

Quota for user **Local User/Freddy** on volume **Volume1** (Volume1 maximum size is 135.6 GB).

Quota	Used	Used (%)
(No space limit)	0.00 MB	0%
(No file limit)	0 files	0%

**User space quota:**

- No space limit (user can consume space up to the maximum size of the volume)
- Limit to  PB (Max. size is 135.6 GB)
- Use default user space quota (No space limit)

**User file quota:**

- No file limit (user can create any number of files)
- Limit to  files. (Max. assignable limit is 1,000,000,000 files)
- Use default user file quota (No file limit)

OK Cancel

- Select or enter the desired **space and file quota** amounts, and click **OK**.

**NOTE:** LDAP/NIS groups do not display the third option for using the default user space or file quota.

## Editing or Removing Quotas

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

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

- Click the **Volume name** link on the **Quotas** initial page to open the **Quotas for Volume** page for that volume.
- If necessary, **search** for a specific user to narrow the list to a more reasonable number. See [Search for Quotas or Space Consumed by a User or LDAP/NIS Group](#) on page 103.

- To edit or remove the quota, from the search results, select the appropriate **user** or **LDAP/NIS group** in the left column to open the settings page for that user.

Quota for user **Local User/Freddy** on volume **Volume1** (Volume1 maximum size is 135.6 GB).

Quota	Used	Used (%)
10 GB	0.00 MB	0%
123,456 files	0 files	0%

User space quota:

- No space limit (user can consume space up to the maximum size of the volume)
- Limit to   (Max. size is 135.6 GB)
- Use default user space quota (No space limit)

User file quota:

- No file limit (user can create any number of files)
- Limit to  files. (Max. assignable limit is 1,000,000,000 files)
- Use default user file quota (No file limit)

OK Cancel

- Select or enter the new **quota** options:
  - When **editing**, enter a limit or select the default quota options.
  - To **remove** a specific quota limit, set the space or file quotas to no limit.

**NOTE:** LDAP/NIS groups do not display the third option for the default space or file quotas.

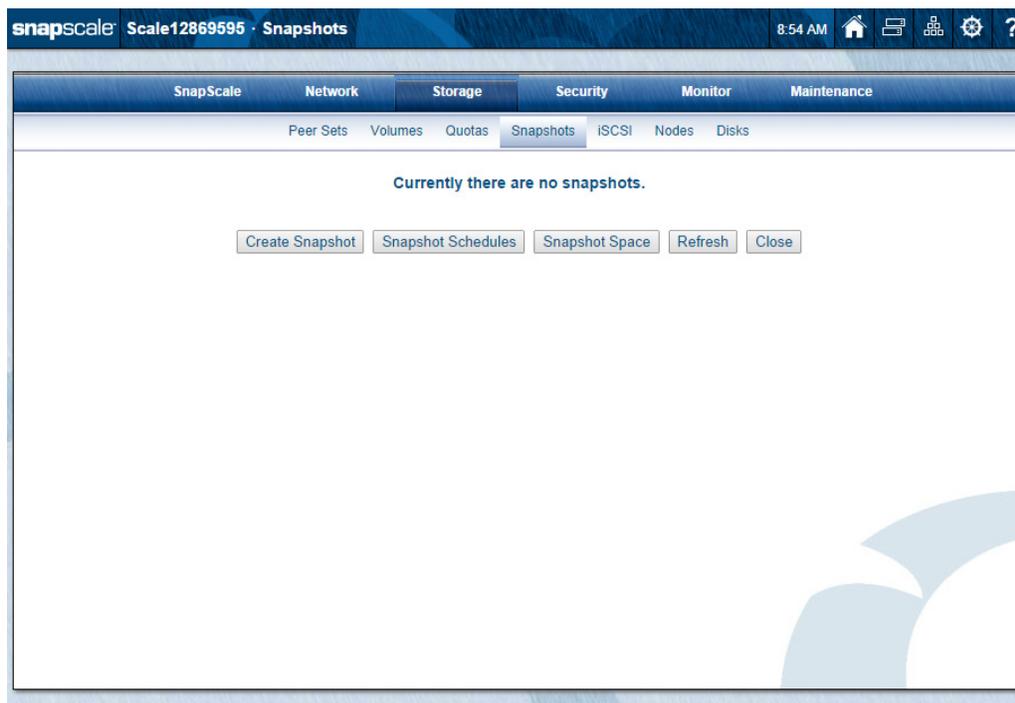
- Click **OK**.

## Snapshots

A *snapshot* is a consistent, stable, point-in-time image of the cluster storage space that can be backed up independent of activity on the cluster storage. Snapshots can also satisfy short-term backup situations such as recovering a file deleted in error 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:** To preserve your cluster configuration and protect your data from loss or corruption, it is critical to schedule backups and snapshots.

To create or schedule snapshots, navigate to **Storage > Snapshots** in the Web Management Interface:



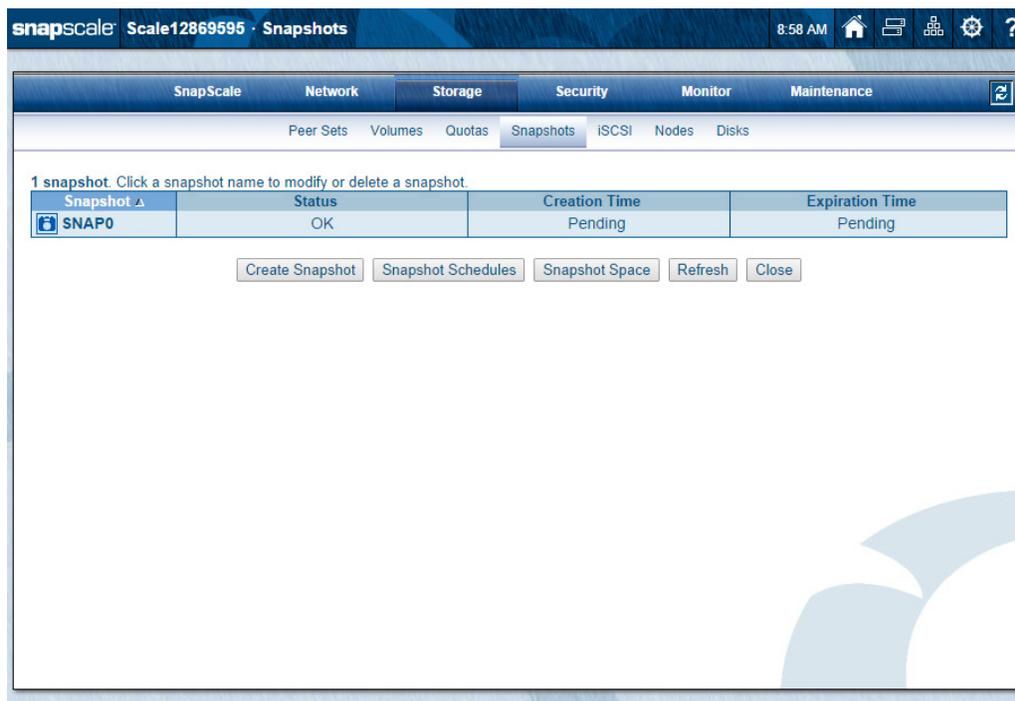
When working with snapshots, consider the following caveats:

- It is recommended that snapshots be taken when the system is idle or under low data traffic to minimize conflicts.
- Snapshots for the cluster storage space use snapshot space reserved on each peer set member drive. If no space is reserved (by unchecking the option box), snapshots are permanently disabled on the cluster.
- While 1% to 90% of the space can be reserved for snapshots, it is recommended that snapshot space be set to 20% of the cluster storage space during setup. Once set, the snapshot space can only be reduced, never increased.
- Snapshot space reserved from each peer set member drive is not necessarily identical to snapshot space of other drives in the same peer set. (This is most likely to occur if two or more drives in the same peer set have recently failed, even if they've been replaced with spares.) As a result, failure of a drive with unique snapshot data may cause one or more snapshots to be automatically deleted.
- Addition of a peer set to the cluster (including automatic peer set creation using new drives inserted into nodes or the addition of new nodes to the cluster) deletes all existing snapshots.
- Failure of a peer set deletes all snapshots.

## Creating Snapshots

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

If no snapshots are currently configured, you only see an empty page. Once a snapshot is created, the page displays a table of snapshots:



These options are available for snapshots:

Action	Procedure
<b>Create a New Snapshot</b>	Click <b>Create Snapshot</b> . The process involves defining snapshot parameters and scheduling.  <b>NOTE:</b> Do not take more snapshots than your system can store, or more than 250 snapshots. Under normal circumstances, nine or ten snapshots are sufficient to safely back up any system.
<b>Edit a Snapshot Schedule</b>	Click the <b>Snapshot Schedules</b> button, and then click the snapshot name. You can modify all snapshot parameters.
<b>Adjust Snapshot Space</b>	Specify the percentage of your SnapScale storage space to reserve for snapshots.  <b>NOTE:</b> The storage space reserved for snapshots can be reduced, but it can <b>never</b> be increased once it is created.
<b>Edit and Delete</b>	Click the snapshot's name in the <b>Snapshot</b> column to open the <b>Snapshot Properties</b> page. You can edit the snapshot's name and duration, or delete the snapshot.
<b>Refresh the Page</b>	Clicking the <b>Refresh</b> button updates the page. This is helpful when waiting for a snapshot to complete.

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

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

**NOTE:** The presence of one or more snapshots on a cluster can impact write performance. Additional snapshots do not have additional impact; in other words, the write performance impact of one snapshot on a cluster is the same as the impact of 100 snapshots.

## 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.

### To Create a Snapshot

Using the **Snapshots** page in the Web Management Interface, you can create a snapshot now, later, or on a recurring schedule. When you select the **Create Snapshot Later** option, additional options are displayed.

Follow these steps to create a snapshot:

1. Go to **Storage > Snapshots**, and click **Create Snapshot**.
2. Enter or select the **options** for the snapshot:
  - a. Type in the **Snapshot Name** (20 character maximum).
  - b. Specify **when** to create the snapshot.
    - Click **Create Snapshot Now** to run the snapshot immediately.
    - Click **Create Snapshot Later** to schedule the snapshot for a later time.

When you select the **Create Snapshot Later** button, a new input section appears below the option. Enter the Start Date and Start Time. Select either to create the snapshot only once (**One Time**) or to have it recurring periodically (**Recurring**) using an interval in hours, days, weeks, or months.

- c. Specify the **duration** of the snapshot.

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

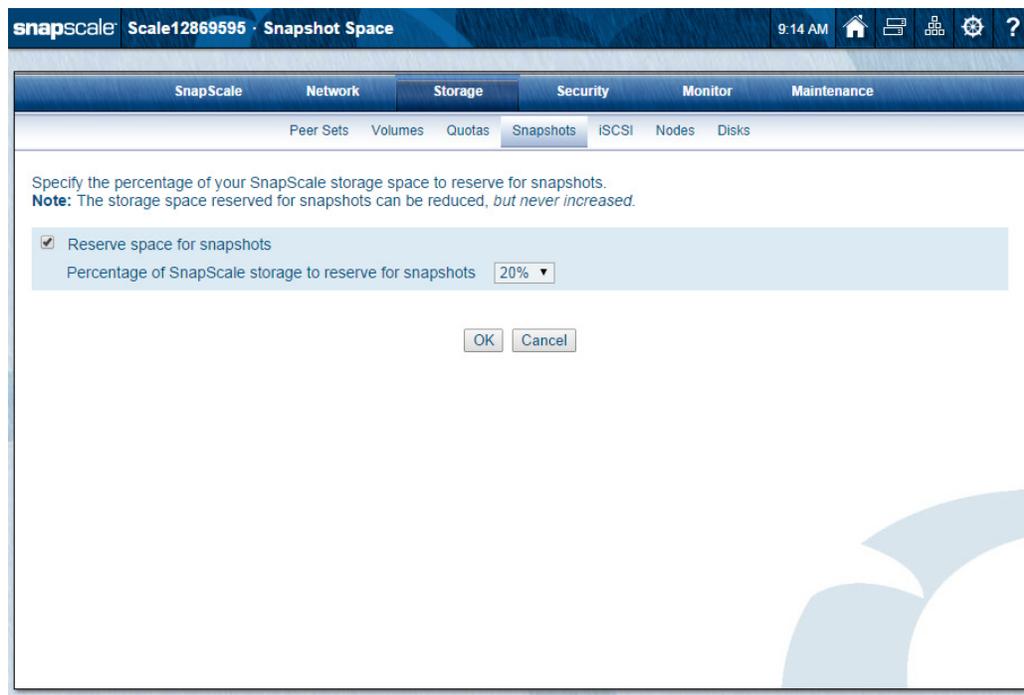
3. Create the snapshot by clicking **Create Snapshot**.

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.

## Adjusting Snapshot Space

**NOTE:** Once the SnapScale cluster is created, the storage space reserved for snapshots can only be decreased. It can never be increased.

If you already reserved storage space for snapshots during the setup of your cluster but now want to decrease the size of that space, use the **Snapshot Space** button to make that change.



1. Go to **Storage > Snapshots**, and click **Snapshot Space**.
2. To reduce or remove the **reserved space**, do one of the following:
  - To decrease the space size, choose a **smaller percentage** of reserved space using the drop-down list.
  - To release all reserved space, uncheck the **reserve space for snapshots** box.

---

**CAUTION:** Unchecking the reserve space box causes all the reserved space to be released, deletes all existing snapshots, and, because the space can never be increased (re-added), **permanently** disables snapshots on the cluster.

---

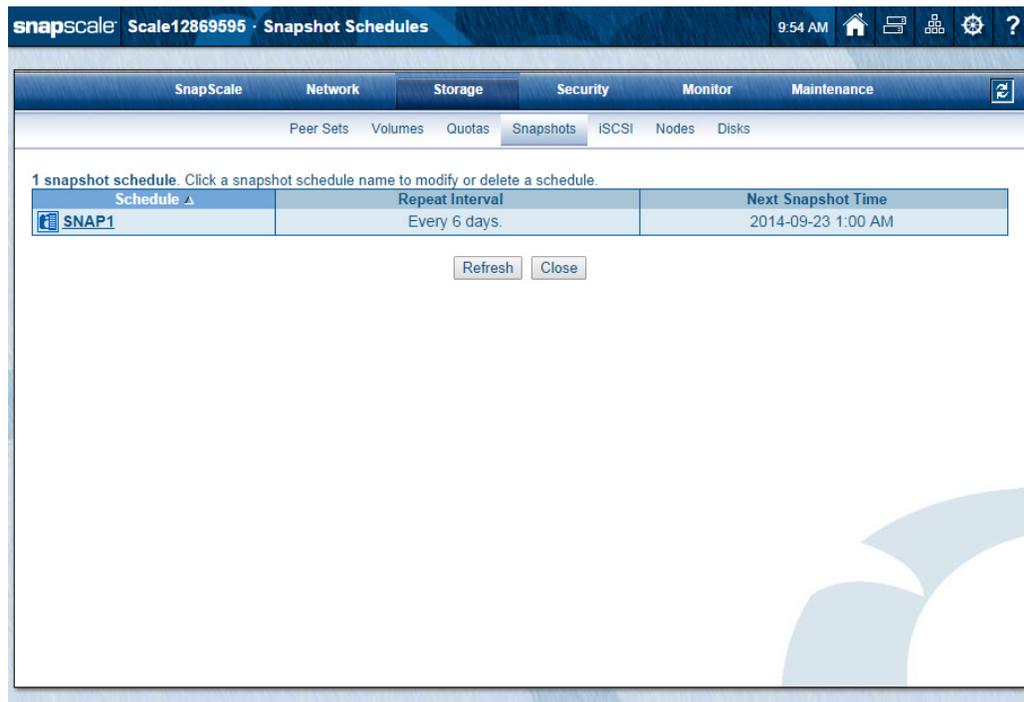
3. Click **OK** to complete the process.

## Accessing Snapshots

After snapshots are created, they can be accessed via a snapshot share. Just as a share provides access to a portion of a live volume, a snapshot share provides access to the same portion of the filesystem on all current snapshots of the volume. The snapshot share's path into snapshots mimics the original share's path into the live volume. The snapshot share is created in the **Shares** section under the **Security** tab. See [Shares on page 147](#) for details.

## Scheduling Snapshots

To view when snapshots are currently scheduled to occur, click **Snapshot Schedules**:



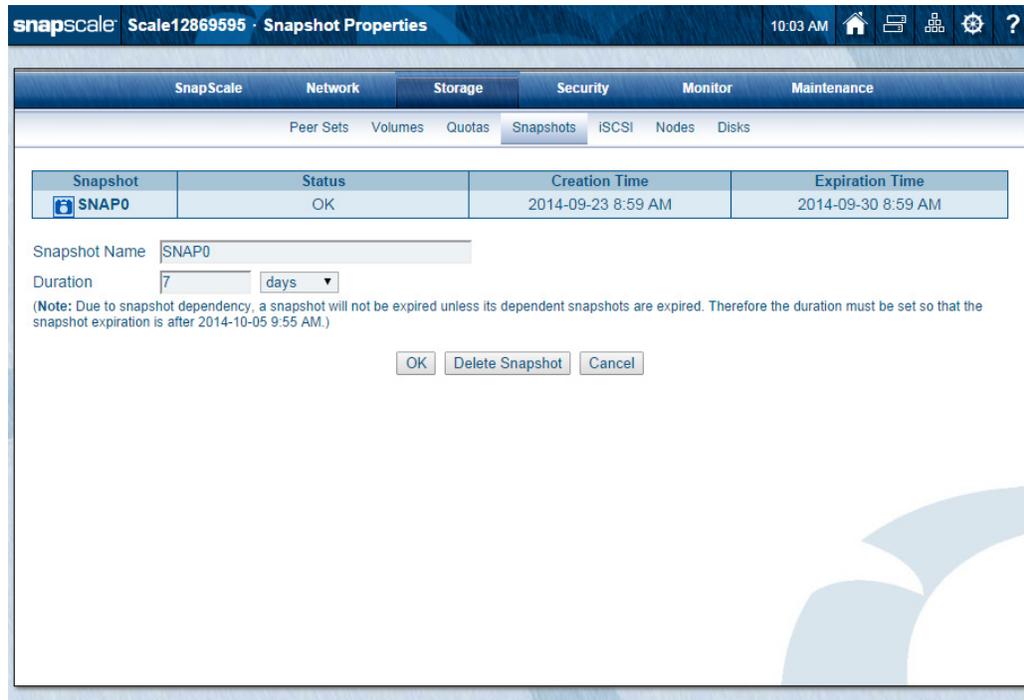
The **Snapshot Schedules** page shows a list of any scheduled snapshots pending. **Repeat Interval** and **Next Snapshot Time** shows the details of when snapshots are scheduled to be taken.

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.

The default snapshot name is `SNAP $n$`  (where  $n$  is a number starting at 0). If it is a recurring, scheduled snapshot, an “ $n$ ” is added to the end of the name.

## Edit Snapshot Properties

From the **Snapshot** primary page table, you can click a snapshot name to access the **Snapshot Properties** page. There you can edit the name and duration, or delete the snapshot:



### Edit a Snapshot

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

### Delete a Snapshot

Click **Delete Snapshot** and then click it again on the confirmation page. The snapshot is deleted.

## iSCSI Disks

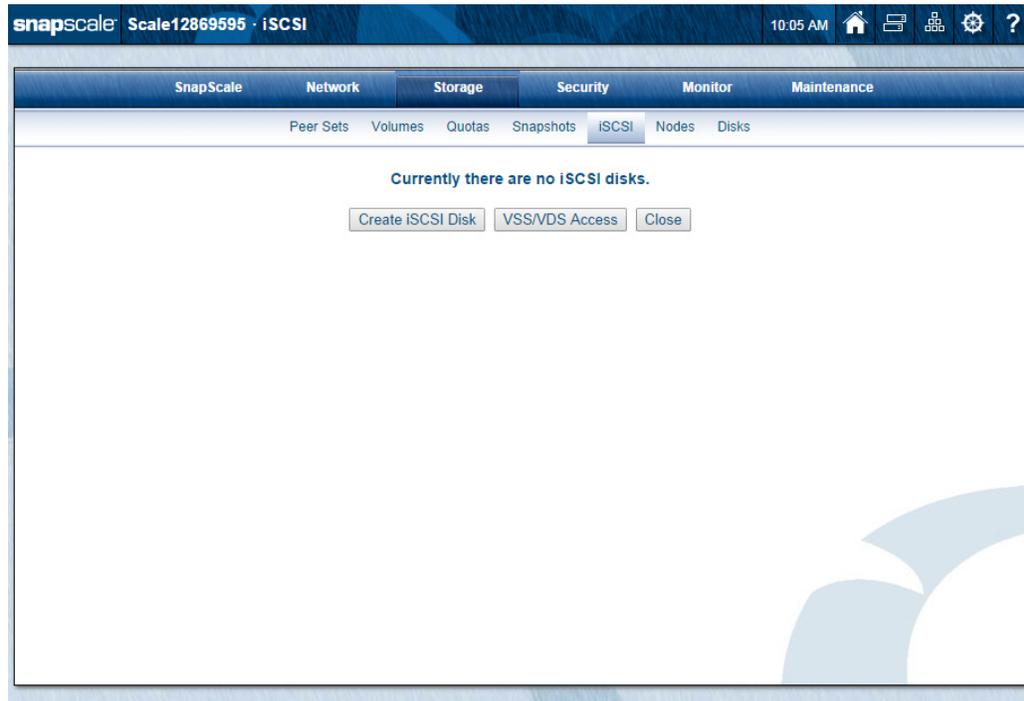
Internet SCSI (iSCSI) is a standard that defines the encapsulation of SCSI packets in Transmission Control Protocol (TCP) and their transmission via IP. On SnapScale clusters, an iSCSI disk consumes cluster storage space as a single large file, but appears to a client machine as a local SCSI drive. This 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 volumes, SnapScale cluster iSCSI disks can be formatted by the iSCSI client to accommodate different application requirements.

### Configuring iSCSI Initiators

Overland Storage has qualified a number of software initiators, PCI cards, and drivers to interoperate with SnapScale clusters. Refer to the vendor's documentation to properly install and configure you initiator to connect to the SnapScale iSCSI disks.

## iSCSI Configuration on the SnapScale

iSCSI disks are created on the **Storage > iSCSI** page of the Web Management Interface.



Before setting up iSCSI disks on your SnapScale cluster, 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.
- iSCSI 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.

### Back up an iSCSI Disk from the Client, not the SnapScale

An iSCSI disk is not accessible from a share and thus cannot be backed up from the SnapScale cluster. 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 SnapScale cluster, the result would be inconsistent or corrupted backup data if any clients are connected during the operation. Only the client can maintain the filesystem embedded on the iSCSI disk in the consistent state that is required for data integrity.

## iSCSI Multi-Initiator Support

Check the **Support Multiple Initiators** box to allow two or more initiators to simultaneously access a single iSCSI target. Multiple initiator support is designed for use with applications or environments in which clients coordinate with one another to properly write and store data on the target disk. Data corruption becomes possible when multiple initiators write to the same disk in an uncontrolled fashion.

**NOTE:** RAINcloudOS supports Windows 2003 and Windows 2008 Server failover clustering.

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

**Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable Multi-Initiator Support if your environment or application supports it.**

It functions as a reminder that data corruption is possible if this option is used when creating an iSCSI disk.

## Disconnect iSCSI Disk Initiators before Shutting Down the Cluster

Shutting down the cluster 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 cluster nodes.

## iSCSI Disk Naming Conventions

iSCSI disks are assigned formal iSCSI Qualified Names (IQNs). These are used when connecting an initiator to an iSCSI target, and differ from the **iSCSI Disk Name** (alias) assigned when the iSCSI disk is created in the Web Management Interface. The full IQN is displayed for each iSCSI disk on the **iSCSI Disk Properties** page.

The screenshot shows the SnapScale Web Management Interface. The top navigation bar includes 'SnapScale', 'Scale12869595', and 'iSCSI Disk Properties'. The main content area has tabs for 'Peer Sets', 'Volumes', 'Quotas', 'Snapshots', 'iSCSI', 'Nodes', and 'Disks'. The 'iSCSI' tab is active, displaying a table with the following data:

iSCSI Disk	Status	Active Clients	IP Address	iSCSI Disk EUI	Device	Size
iscsi0	OK	0	10.25.11.161	00c0b67dd3000001	/hd/cfs/blockstor/snapbd0	21.56 GB

Below the table, the 'iSCSI Disk IQN' is displayed as `iqn.1997-10.com.snapscale.scale12869595.snapbd0`. The 'iSCSI Disk Size' is set to 21.56 GB. There are two checkboxes: 'Support Multiple Initiators' (unchecked) and 'Enable CHAP Logon' (unchecked). A warning message is visible below the 'Support Multiple Initiators' checkbox: 'Warning: Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable multi-initiator support if your environment or application supports it.' At the bottom, there are buttons for 'OK', 'Delete iSCSI Disk', 'Refresh', and 'Cancel'.

The full IQN is also shown when you mouseover the iSCSI name in the disk table.

The format of IQNs for new SnapScale iSCSI disks is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]
```

where `[clustername]` is the name of the SnapScale cluster, and `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster. All the `[blockdevice]` names are automatically created using the term `snapbd` appended with a sequence number (such as, `snapbd0`, `snapbd1`, and so on).

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0
```

The format of IQNs for **VSS-based iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice].[nnn]
```

where `[clustername]` is the name of the SnapScale cluster, `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster, and `[nnn]` is a sequential number starting from 000. For example:

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0.000
```

The format of IQNs for **VDS-based iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]
```

and the format for IQNs for **snapshots of iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]-snap[n]
```

where, in both cases, `[clustername]` is the name of the SnapScale cluster, `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster, and `[n]` is a sequential number starting from 0. For example:

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0-snap0
```

## Create iSCSI Disks

Navigate to **Storage > iSCSI** and click **Create iSCSI Disk** to create iSCSI disks on the SnapScale cluster. Be sure to read [iSCSI Configuration on the SnapScale on page 116](#) before you begin creating iSCSI Disks.



The screenshot shows the SnapScale web interface for creating an iSCSI disk. The breadcrumb navigation is **Storage > iSCSI**. The page title is **Create iSCSI Disk**. The interface includes a top navigation bar with tabs for **SnapScale**, **Network**, **Storage**, **Security**, **Monitor**, and **Maintenance**. Below this, there are sub-tabs for **Peer Sets**, **Volumes**, **Quotas**, **Snapshots**, **iSCSI**, **Nodes**, and **Disks**. The main form contains the following fields and options:

- iSCSI Disk Name:**
- IP Address:**
- Size:**
- Support Multiple Initiators**  
**Warning:** Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable multi-initiator support if your environment or application supports it.
- Enable CHAP Logon**

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

1. Navigate to **Storage > iSCSI** and click **Create iSCSI Disk**.
2. Enter the **iSCSI settings** for the disk name and size (16GB minimum).  
You can accept the default name or enter a new one using up to 20 alphanumeric, lowercase characters. You can also accept the default size (the remaining cluster space) or enter a specific size.
3. If you want your iSCSI Disk to allow **multiple** initiator connections, check that box.

**NOTE:** Data corruption is possible if this option is checked. See [iSCSI Multi-Initiator Support on page 117](#) for more information.

- If desired, enable CHAP authentication by checking the **Enable CHAP Logon** box to display the hidden options.

The screenshot shows the 'Create iSCSI Disk' configuration page in the SnapScale web interface. The page has a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. The iSCSI tab is active. The form includes the following fields and options:

- ISCSI Disk Name:
- IP Address:
- Size:
- Support Multiple Initiators
- Warning:** Uncontrolled simultaneous access of multiple initiators to the same iSCSI target can result in data corruption. Only enable multi-initiator support if your environment or application supports it.
- Enable CHAP Logon** (highlighted with a red box)
- User Name:
- Target Secret:  (Minimum 12 characters)
- Confirm Target Secret:
- 

Enter a **User Name** and **Target Secret** (password), and then confirm the password. Consider the following:

- Both items 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.
- Click the **Create iSCSI Disk** button.
  - At the confirmation page, verify the settings and click the **Create iSCSI Disk** button again to complete the process.

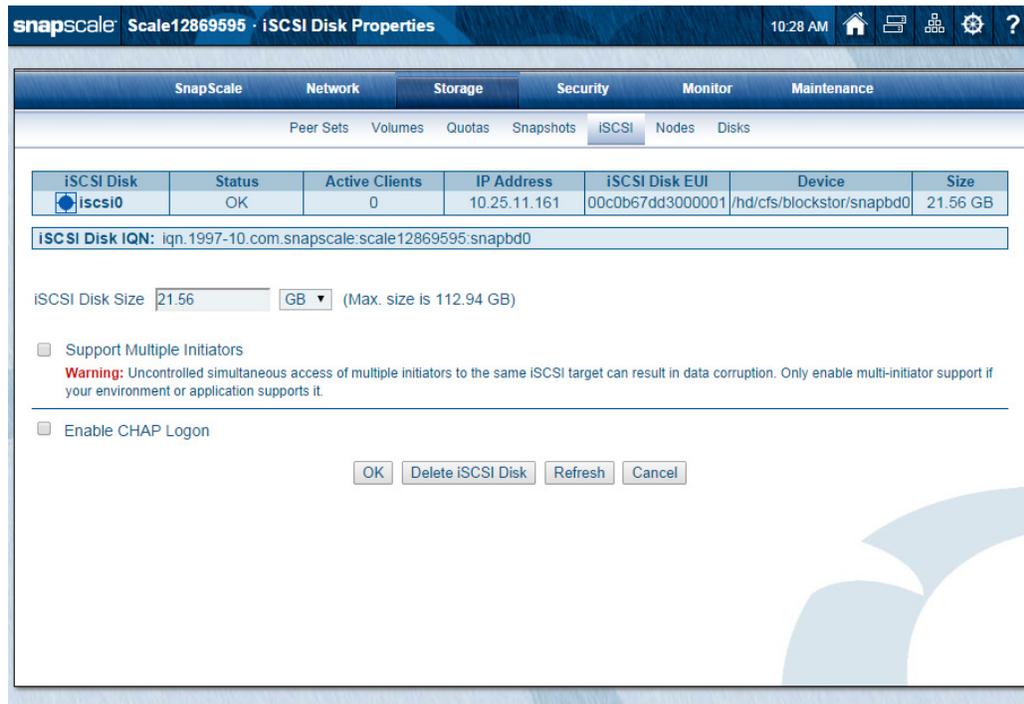
You are returned to the **iSCSI** page and the new iSCSI disk is displayed in the table there with the following information:

Label	Description
<b>iSCSI Disk</b>	The name of the iSCSI disk.
<b>Status</b>	Current condition of the iSCSI disk: <ul style="list-style-type: none"> <li><b>OK</b> – The iSCSI disk is online and accessible.</li> <li><b>Stopped</b> – The iSCSI disk is currently stopped.</li> <li><b>Failed</b> – The iSCSI disk has failed.</li> </ul>
<b>Active Clients</b>	The number of current sessions.
<b>IP Address</b>	The IP address used by the iSCSI disk.
<b>Authentication</b>	Either <b>CHAP</b> or <b>None</b> .
<b>Size</b>	The size of the iSCSI disk.

## Edit iSCSI Disk Properties

**NOTE:** You cannot edit an iSCSI disk until all active clients have been disconnected from that disk. The hostname and IQN name of all connected initiators are displayed in the properties table.

After disconnecting all client initiators, go to **Storage > iSCSI** and click the iSCSI disk name in the table to display the **iSCSI Disk Properties** page.



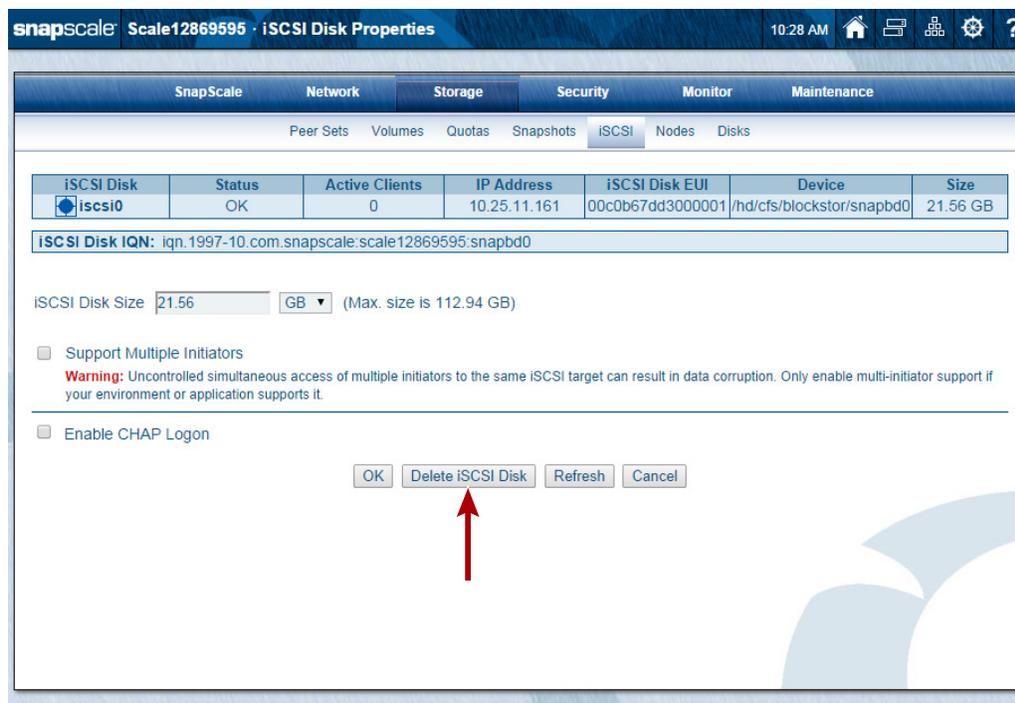
- On this **page**, you can do the following:
  - View the iSCSI Disk **IQN**.
  - Increase (but not decrease) the **size** of the iSCSI disk (if space is available).
  - Enable or disable support for **multiple initiators**.
  - Enable or disable **CHAP logon**.
- Click **OK** to accept the changes.

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

## Delete an iSCSI Disk

**NOTE:** You cannot delete an iSCSI disk until all active clients have been disconnected.

After disconnecting all client initiators, go to **Storage > iSCSI** and click the iSCSI disk name in the table to display the **iSCSI Disk Properties** page.



Click **Delete iSCSI Disk** (which is followed by a confirmation page) to delete the iSCSI disk.

## Configuring VSS/VDS for iSCSI Disks

RAINcloudOS 4.1 provides VSS and VDS hardware providers to support Microsoft Volume Shadow Copy Services (VSS) and Virtual Disk Service (VDS) for iSCSI disks.

- The **VSS** hardware provider provides a mechanism for taking application-consistent Windows-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 on a Windows initiator host, so that the job backs up the snapshot volume rather than the main production volume.

**NOTE:** VSS iSCSI snapshots are managed by the Windows client and represent the iSCSI disk, not the Snap volume on which the iSCSI disk resides. They are not related to RAINcloudOS snapshots as described in [Snapshots on page 109](#). The VSS iSCSI snapshot rollback feature is not currently supported.

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

### Backing up an iSCSI Disk using VSS Snapshots

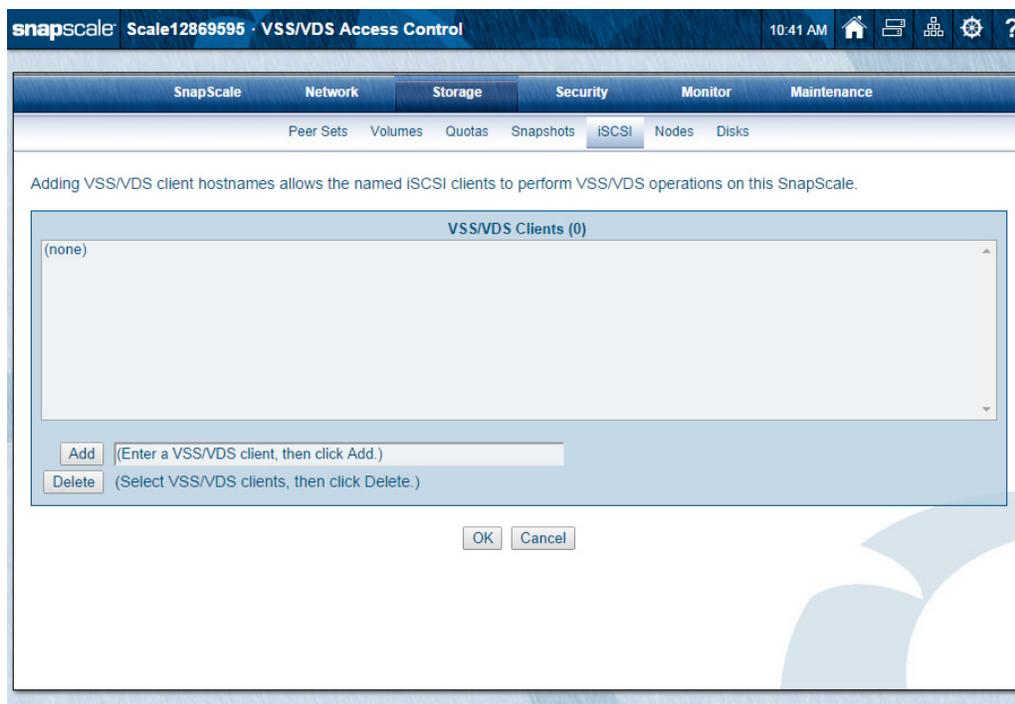
Windows VSS-compatible backup applications can create snapshots of SnapScale cluster iSCSI disks to perform consistent backups of application data without stopping the application, using the snapshot instead of the live volume as the backup source.

Each VSS snapshot of an iSCSI target consumes additional cluster storage space. The required space is 10% of the size of the iSCSI disk per snapshot. If this amount of free space is not available on the pool or volume, the VSS snapshot will not be created and an error will be reported by the SnapScale cluster 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 cluster.

**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. Go to **Storage > iSCSI** and click **VSS/VDS Access**.



2. Add **VSS clients** to the SnapScale cluster:
  - a. At **VSS/VDS Access Control**, enter the **VSS client hostname** to which you want to grant access, and click **Add**.  
The client hostname is not case-sensitive and should appear in the VSS/VDS Clients box.

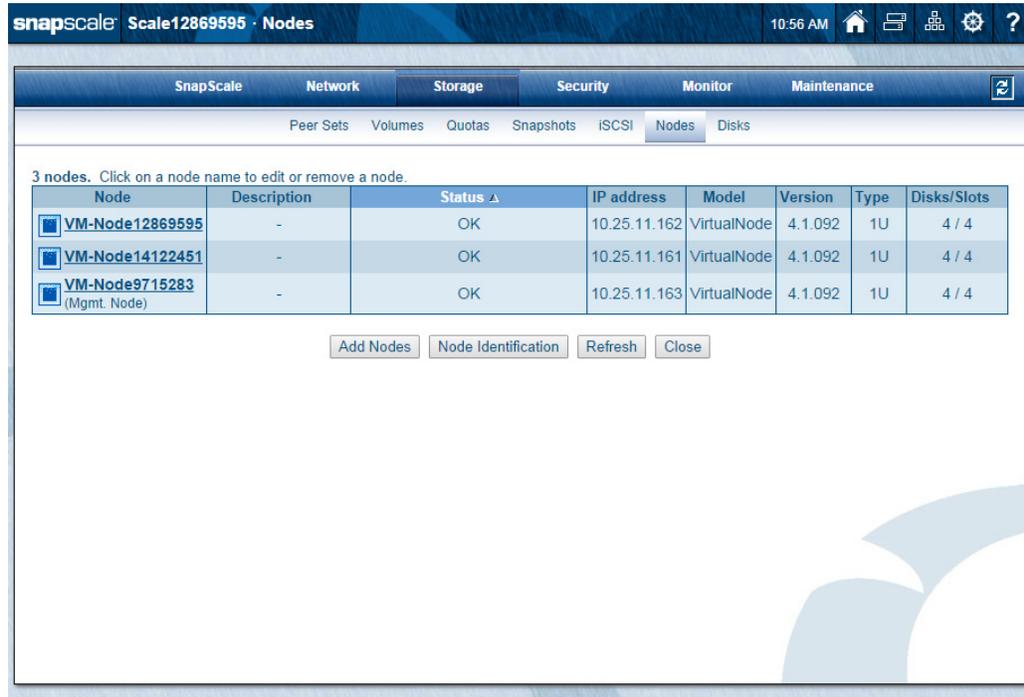
**NOTE:** Use only the short hostname (*myclientname*) of the client only. Do not use the IP address or fully-qualified name (for example, *myclientname.mydomain.com*).

  - b. Repeat [Step a](#) for any other hostnames you want to add.
  - c. When you have finished adding VSS clients, click **OK**.
3. Install the **VSS hardware** provider on the Windows iSCSI client.
  - a. Depending on the Windows client, go to the appropriate heading to locate and download either *SnapServerToolsInstall32.exe* or *SnapServerToolsInstall64.exe* from the Overland website:  
<http://docs.overlandstorage.com/ssm>
  - b. Double-click the **executable** (.exe) to run the Installation Wizard on the VSS client and select the VSS/VDS hardware providers option. This will add the hardware provider to the Windows iSCSI client.
4. 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.



## Nodes

Use the **Nodes** page to manage the nodes that make up the cluster.



From this page, you can:

- Add a new node.
- Access the **Node Properties** page to edit or remove a node from the cluster.
- Identify physical nodes via flashing LEDs.

Some important points about SnapScale nodes:

- Users can access the cluster storage over any of the configured network protocols by connecting to any of the nodes.
- Because the storage space is unified across the cluster, connecting to any of the nodes provides access to the same data as any other cluster node.
- To balance network client access to the nodes, enter an **A** record to the DNS pointing to the cluster name for each IP address in the node IP address range. The DNS then uses round-robin name resolution requests for the cluster name among the node IP addresses. Alternatively, manually distribute clients accessing the cluster to different IP addresses in the node IP address range.
- When a node fails, the IP address it uses is automatically reassigned to another node. Clients connected to that IP address are forwarded to the new node, though this may cause a momentary interruption to storage access.
- Files opened by clients connected to any node are recognized by all nodes, and file locks are respected by all nodes.

## Edit Node Properties

By clicking a node's name in the **Node** column of the table on the **Node** page, the **Node Properties** page is shown with details of that particular node.

The screenshot shows the SnapScale Node Properties page for a node named VM-Node14122451. The page has a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Storage tab, there are sub-tabs for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. The main content area contains a table with the following data:

Node	Description	Status	IP address	Model	Version	Type	Disks/Slots
VM-Node14122451	-	OK	10.25.11.161	VirtualNode	4.1.092	1U	4 / 4

Below the table, there are two buttons: a light blue box with a plus sign and the text "Click here to identify this node by flashing its LEDs for 5 minutes." and a red box with an "X" and the text "Click here to stop flashing this node's LEDs." Below these buttons is a text input field for the Node Description, followed by several action buttons: OK, View Disks, Remove Node from SnapScale, Refresh, and Cancel.

From this page, you can:

- Flash the node drive LEDs to help identify the node.
- Change the node description.
- View the drives in the node.
- Remove the node from the SnapScale cluster.

### Flash the Node LEDs

Click the light-blue box () under the node name to start the LEDs flashing for up to five minutes. Click the box with the red "X" () to stop flashing the LEDs.

## Node Drives

To view the drives that are installed in the node, click the **View Disks** button on the properties page.



When you click the name of the drive in the **Disk** column on the **Node Disks** page, the **Disks** page is displayed with the physical location of the disk drive. See [Disks on page 135](#) for more information.

Clicking the member name in the **Type** column takes you to the **Peer Sets** page with that member highlighted. See [Peer Sets Page on page 88](#) for more information.

When done, click **Close** to return to the **Node Properties** page.

## Adding Nodes

A SnapScale cluster can also be expanded by adding more nodes. Expansion kits are available that consist of either two or three additional nodes and all the necessary cables. Documentation is included with each node that details how to install, cable, and power up the new node.

Once installed in a rack, clicking the **Add Node** button on the **Nodes** page starts a wizard to add one or more nodes to the cluster to expand the storage space. By default, all eligible nodes are pre-selected.

---

**IMPORTANT:** In order to expand storage by adding nodes, the cluster must be able to create peer sets with each member on different nodes. To efficiently increase cluster storage, it is recommended that the number of new nodes you add is equal to or greater than the Data Replication Count (2x or 3x). It is highly recommended that these new nodes all be added in the same operation.

---

When adding nodes, **all** the following must be taken into consideration:

- The nodes must all be running the same version of RAINcloudOS (ROS). See [OS Update on page 213](#).

- All the nodes, those already part of the cluster and those being added, must be attached to the same Client subnet.
- No expansion units can be attached to a node.
- The appropriate ports on the node must be available to create the proper bonding. (See [Client and Storage Networks on page 16.](#))

Using the wizard, follow these steps to add your nodes:

1. At the **Node** page, click **Add Nodes**.



**IMPORTANT:** It is highly recommended that all active iSCSI users be disconnected before continuing.

snapScale Scale12869595 · Add Nodes to SnapScale 11:08 AM

SnapScale Network Storage Security Monitor Maintenance

Peer Sets Volumes Quotas Snapshots iSCSI Nodes Disks

### Add Nodes to SnapScale

This guide will discover available nodes and help you add them to your SnapScale.

**Warning:** In order to avoid possible data loss, it is highly recommended that all active iSCSI users (across all iSCSI disks) be disconnected from this SnapScale before continuing.

**Add Nodes to your SnapScale**

Node Node Node Node

New Node New Node

Back Next

(Click Next to search for available SnapScale nodes.)

2. Click **Next** to display node choices (Wizard Step 1):

**Add Nodes to SnapScale: Select Nodes to Add**

Step 1 | Step 2 | Step 3

Select the nodes below that you want to add to your SnapScale and click Next. (All eligible nodes are selected by default.)

**Note:** All nodes in the SnapScale must have identical RAINcloudOS (ROS) versions, and the client network interface for all nodes must be located on the same subnet.

**Important:** The data replication count is set at 2x. In order to prevent possible peer set synchronization issues you should select a minimum of 2 nodes to add to the SnapScale. Also, it is recommended that you select all nodes to add now, rather than adding nodes incrementally.

**2 Eligible Nodes:** (Note: 4 nodes are not eligible to be added to this SnapScale.)

Node	Model	ROS Version	Disks	Add to SnapScale
Node1234511	X2	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
Node1234512	X4	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: 1.95 TB 11: 3.91 TB 12: 3.91 TB 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB ↓ Rear Disks ↓ 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk) 33: (No Disk) 34: (No Disk) 35: (No Disk) 36: (No Disk)	<input checked="" type="checkbox"/>
Node1234502	X4	4.0.077	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: 1.95 TB 11: 3.91 TB 12: 3.91 TB 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB ↓ Rear Disks ↓ 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk) 33: (No Disk) 34: (No Disk) 35: (No Disk) 36: (No Disk)	<input type="checkbox"/> (Different/older ROS version; click here to upgrade this node.)

Back Re-Detect Available Nodes Next

3. Check the boxes for the **nodes** you want to add to the cluster, and click **Next**.

**NOTE:** By default, all eligible nodes are pre-selected. It is recommended to accept all the nodes to ensure the optimum configuration.

If the node bond type doesn't match the cluster bond type, special informational messages are shown in the wizard based on the existing situation:

- If one or more of the nodes have a different Storage network bond type than the cluster, a note table is displayed in Wizard Step 1 showing the bond issues (or other errors).
- The final wizard page displays a percent completed status while the updated nodes with changed bond types are being rebooted.

If an available node is not eligible to be added to the cluster due to an OS version mismatch, it can be upgraded to make it eligible.

snapScale Scale12869595 · Add Nodes to SnapScale 11:16 AM

SnapScale Network Storage Security Monitor Maintenance

Peer Sets Volumes Quotas Snapshots iSCSI Nodes Disks

**Add Nodes to SnapScale: Select Nodes to Add**

**Warning:** There are available SnapScale nodes, however none are eligible to be added to this SnapScale (see below for details).

**Note:** All nodes in the SnapScale must have identical RAINcloudOS (ROS) versions, and the client network interface for all nodes must be located on the same subnet.

0 Eligible Nodes: (Note: 1 node is not eligible to be added to this SnapScale.)

Node	Model	ROS Version	Disks	Add to SnapScale
VM-Node5381351	VirtualNode	4.1.056	1: 50 GB 2: 50 GB 3: 50 GB 4: 50 GB	(Different/older ROS version; <a href="#">click here to upgrade this node.</a> )

Back Re-Detect Available Nodes

Click the **click here** link in the **Add to SnapScale** column, then login to the node, and perform an OS update (see [OS Update on page 213](#)). When updated, go back to the browser tab showing the cluster and click **Re-Detect Available Nodes**.

- At Wizard Step 2, configure the **static IP addresses** for the nodes, and click **Next**.

snapScale Scale1234501 · Add Nodes to SnapScale 5:39 PM

SnapScale Network Storage Security Monitor Maintenance

Peer Sets Volumes Quotas Snapshots iSCSI Nodes Disks

**Add Nodes to SnapScale: Static Node IP Addresses**

Step 1 Step 2 Step 3

Please specify 2 static IP addresses for the 2 nodes you are adding.

**Important:** The IP addresses you specify must all be located on the same subnet as the SnapScale client network: (Range: 192.168.198.1 - 192.168.199.254).

**Optional:** Enter a starting IP address and click the "Populate" button to populate the list below with sequential static IP addresses. You can then review or change the IP addresses before clicking Next.

Starting IP Address  Populate Static IP Addresses (optional)

Enter 2 static IP addresses below:  
([Click here](#) to auto-populate IP addresses.)

Static IP Address
<input type="text"/>
<input type="text"/>

192.168.199.100 (Mgmt. IP)  
192.168.199.101  
192.168.199.102  
192.168.199.103  
192.168.199.104  
192.168.199.105

These 12 static IP addresses are currently being used by your SnapScale. (Note: The IP addresses you specify for your new nodes cannot already be in use by your SnapScale.)

Back Next

It is recommended that you click the **Click here** option at the lower left to automatically add IP addresses based on the addresses being currently used.

**NOTE:** When more new nodes are being added to the cluster than there are unused IP addresses in the node address pool, more IP addresses must be added to the pool.

You can also enter a starting address in the Populate field based on the static IP addresses (in the list on the right) currently being used by your SnapScale cluster, and then click the **Populate Static IP Addresses** button.

- At Wizard Step 3, verify the data and click **Add Nodes to SnapScale**.

**Add Nodes to SnapScale: Add Nodes Summary**

Step 1 | Step 2 | **Step 3**

Please review your settings below and click Add Nodes to SnapScale when ready.

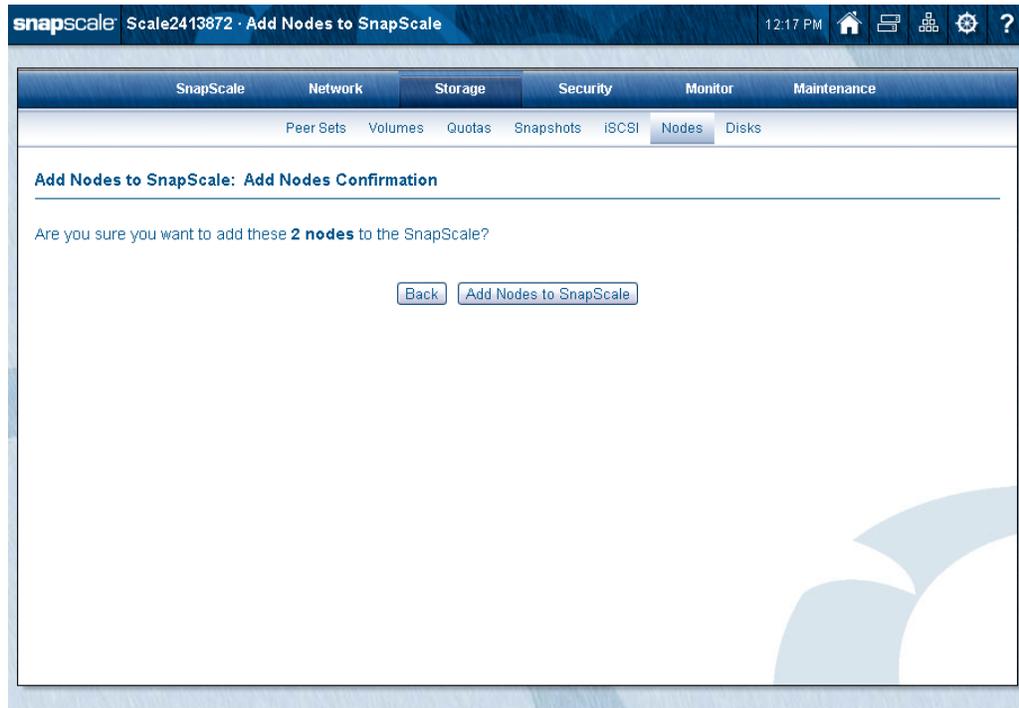
**2 New SnapScale Nodes.**

Node	IP Address	Model	ROS Version	Disks
Node1234511	192.168.199.112*	X2	4.0.119	1: 1.95 TB   2: 1.95 TB   3: 3.91 TB   4: 3.91 TB 5: 1.95 TB   6: 1.95 TB   7: 3.91 TB   8: 3.91 TB 9: 1.95 TB   10: (No Disk)   11: (No Disk)   12: (No Disk)
Node1234512	192.168.199.113*	X4	4.0.119	1: 1.95 TB   2: 1.95 TB   3: 3.91 TB   4: 3.91 TB 5: 1.95 TB   6: 1.95 TB   7: 3.91 TB   8: 3.91 TB 9: 1.95 TB   10: 1.95 TB   11: 3.91 TB   12: 3.91 TB 13: 1.95 TB   14: 1.95 TB   15: 3.91 TB   16: 3.91 TB 17: 1.95 TB   18: 1.95 TB   19: 3.91 TB   20: 3.91 TB 21: 1.95 TB   22: 1.95 TB   23: 3.91 TB   24: 3.91 TB ↓ Rear Disks ↓ 25: 1.95 TB   26: 1.95 TB   27: 3.91 TB   28: (No Disk) 29: (No Disk)   30: (No Disk)   31: (No Disk)   32: (No Disk) 33: (No Disk)   34: (No Disk)   35: (No Disk)   36: (No Disk)

\*Note: This IP address will not necessarily be assigned to its associated node.

Back | Add Nodes to SnapScale

6. At the confirmation page, click **Add Nodes to SnapScale** again.



The new nodes are added to the cluster and the peer sets are built. This process takes several minutes.

---

**IMPORTANT:** If the Storage network bond type on the new nodes needs to be changed to or from Switch Trunking or Link Aggregation (802.3ad) to match the cluster bond type, an additional step is shown after shutting down the node so you can reconfigure the network switch and restart the node. You can also cancel the adding of the nodes if desired. However, if no action is taken in the first hour, after that time the adding of the nodes is automatically canceled.

---

Example of network bond message:

**Scale1234501 - Add Nodes to SnapScale** 4:32 PM ?

**Add Nodes to SnapScale: Adding Nodes**

The following nodes are being added. **Please wait...** (Elapsed time: 0:46)

**3 New SnapScale Nodes.**

Node	IP Address	Status
Node1234511	192.168.199.106*	Storage network bond type changed; node shut down. <b>Waiting for user response.</b>
Node1234512	192.168.199.107*	Storage network bond type changed; node shut down. <b>Waiting for user response.</b>
Node1234513	192.168.199.108*	Storage network bond type changed; node shut down. <b>Waiting for user response.</b>

\*Note: This IP address will not necessarily be assigned to its associated node.

**User Action Required:**

The storage network bond type for the new nodes has been changed from Failover to **Switch Trunking**, and the nodes have been shut down. Please perform the following steps, in order:

1. Configure your network switch for these nodes.
2. **Warning: Data corruption of your SnapScale can occur if you add a node with an improperly configured storage network.** Manually power-on the nodes.
3. Click the Continue Adding Nodes button, below.

Cancel Adding Nodes Continue Adding Nodes

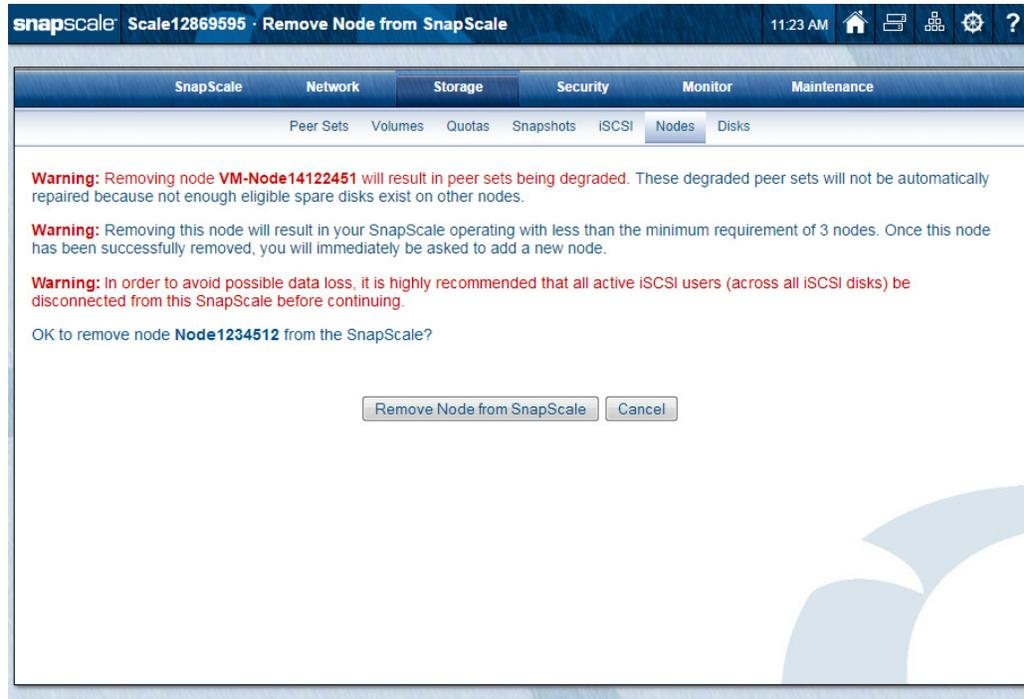
Note: This Add Nodes process will automatically be canceled after 1 hour if you do not click the Continue button.

## Removing Nodes

To remove a node from a SnapScale cluster, go to **Storage > Nodes**, click the node name to view the **Node Properties** page, and then click **Remove Node from SnapScale**. At the confirmation page, click **Remove Node from SnapScale** again.



**IMPORTANT:** It is highly recommended that all active iSCSI users be disconnected before continuing.



**IMPORTANT:** Removing a node may result in one or more peer sets becoming degraded. These degraded peer sets may not be automatically repaired if there are not enough eligible spare drives on other nodes. Removing this node may also result in your SnapScale cluster operating with less than the minimum requirement of 3 nodes.

**NOTE:** If removing the node would destroy one or more peer sets, an error message is returned and the node is not removed.

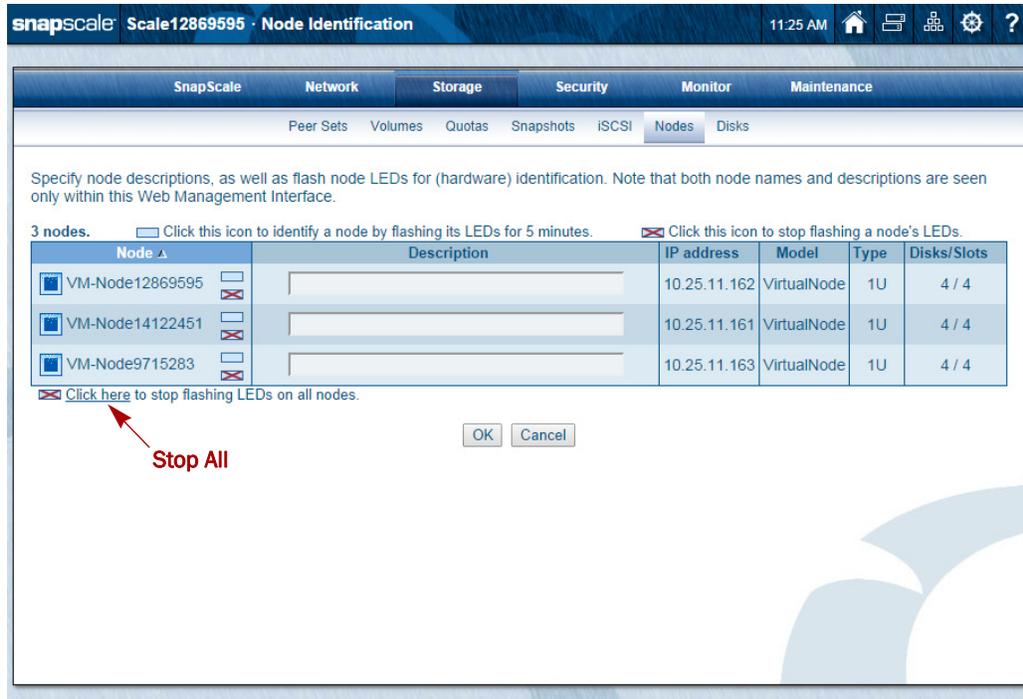
The node itself is no longer associated with the cluster and becomes an Uninitialized node that can be added to another cluster.

## Node Identification

The **Node Identification** page (accessed by the button on the main **Nodes** page) provides a convenient place to check the nodes and optionally change their descriptions for easier identification in the Web Management Interface.

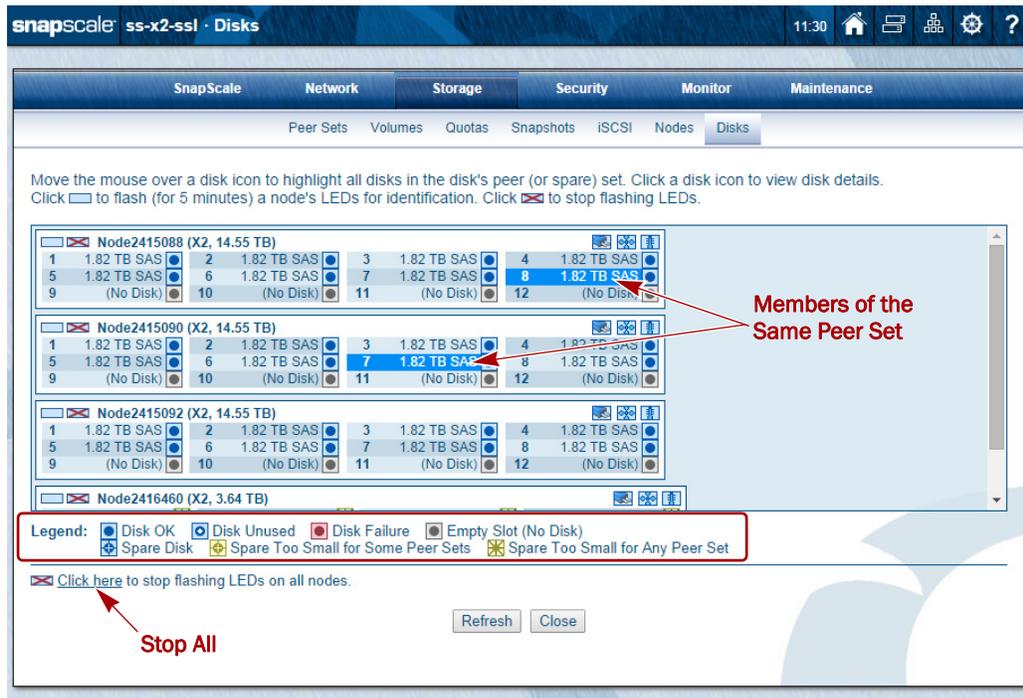
**NOTE:** These options are also accessible from the **Node Properties** page.

Click a light-blue box () next to the node name to start the node's LEDs flashing for up to five minutes. Click the box with the red "X" () to stop flashing the LEDs or click the link next to the same icon below the nodes table to stop all node LEDs flashing.



## Disks

The **Disks** page is a graphic representation of the peer sets and disk drive status on your cluster. The legend on the page explains the meaning of each icon.



- Click a drive icon (  ) to view drive details.
- Hover over a drive icon (  ) to highlight in blue all the other members of the peer set.

- Hover over a spare drive icon () to view other spare drives.
- Click a unit's LED icon () to flash the unit's status and drive status LEDs for identification. The LEDs flash amber. Click the LED stop icon () to stop the unit's LEDs from flashing.

**NOTE:** The LEDs continue to flash for five minutes unless stopped. To stop flashing LEDs for all units, click the link next to the stop icon located below the Legend list.

## Replacing Drives

Should a drive fail (solid red LED), it can be replaced (hot-swapped) without shutting down the SnapScale node. If a spare is available on a node that doesn't already have an active member of the failed drive's peer set, the spare automatically replaces the failed drive and the new drive being installed automatically becomes a spare. If no spares are available, the new drive automatically becomes a member of the failed drive's peer set.

A failed drive can be removed and replaced anytime. When hot-swapping multiple drives, it is recommended to swap one drive at a time to avoid timing issues.

**NOTE:** Hot-removed (non-spare) drives cannot be added directly back into the peer set from which they were removed. When any non-spare drive is physically removed, it is also removed from the peer set to which it belonged. That peer set then becomes degraded and it attempts to incorporate a suitable cluster spare. If the removed drive is reinserted and there is a degraded peer set or not enough spares to satisfy the spare count, the reinserted drive is reconfigured as a spare; otherwise, the reinserted drive becomes unused and available to the user to incorporate manually via the Web Management Interface.

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

## Adding Drives



**CAUTION:** If new peer sets are created when adding new drives to a node in a SnapScale cluster, all existing snapshots will be deleted.

---

**NOTE:** SnapScale only supports SAS (and nearline SAS) drives.

If empty slots are available on one or more nodes, you can add Overland-approved drives to either expand cluster storage space or add hot spares. When adding drives to expand storage space, distribute the new drives evenly across all the cluster nodes.

**NOTE:** Drives should be added without shutting down the node so that the cluster properly recognizes each drive. Note that drives with different rotational speeds cannot be combined in the same cluster.

In order to properly create peer sets with each member on different nodes, if you have the Data Replication Count set at 2x, you must add drives in groups of two, each one in a different node. For a 3x count, add drives in groups of three, each one in a different node.

Once the new drives are added to the nodes, they must be incorporated using the **New Disks Detected** page in the Web Management Interface to enable the cluster to use them properly to create new peer sets and spares.

## Important Considerations

When adding new drives, consider the following:

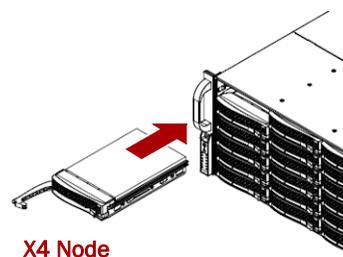
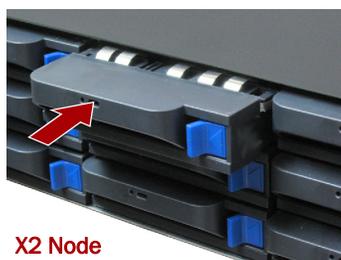
- Only add new drives (hot-insert) when the cluster is up and fully operational. Do not add new drives when the cluster is powered down.
- New drives can be added to Uninitialized nodes either while running or powered down.
- If there are fewer spares in the cluster than the spare count specifies, drives added to any cluster node are automatically configured as hot spares until the spare count is satisfied.
- If there is a degraded peer set on the cluster when adding a new drive and there are no existing spares, the drive will automatically be incorporated into the peer set as long as it is not on one of the nodes containing another active member of the peer set.
- When replacing a failed drive, it is recommended that the new drive be installed in the same slot as the old one to maintain a capacity balance.
- When adding a drive that replaces a failed drive in a peer set, the **Peer Sets** page will display that peer set as rebuilding the new drive into the peer set.

## Drive Installation

**NOTE:** Do not remove the disk drives from their carriers. Doing so voids the drive warranty. Unless adding drives to an Uninitialized nodes, the cluster must be up and fully operational.

Install the drives into the available slots:

1. Remove the **blank drive carriers** from the slots that will be used for the new drives (leaving the remaining blank carriers in place).
2. Positioning a **drive carrier** in front of the appropriate **bay**:
  - For the **X2** node, slide it in until the **latch** clicks, locking the assembly in the bay.
  - For the **X4** node, slide it in until it stops and then close the **latch** to lock it in place.

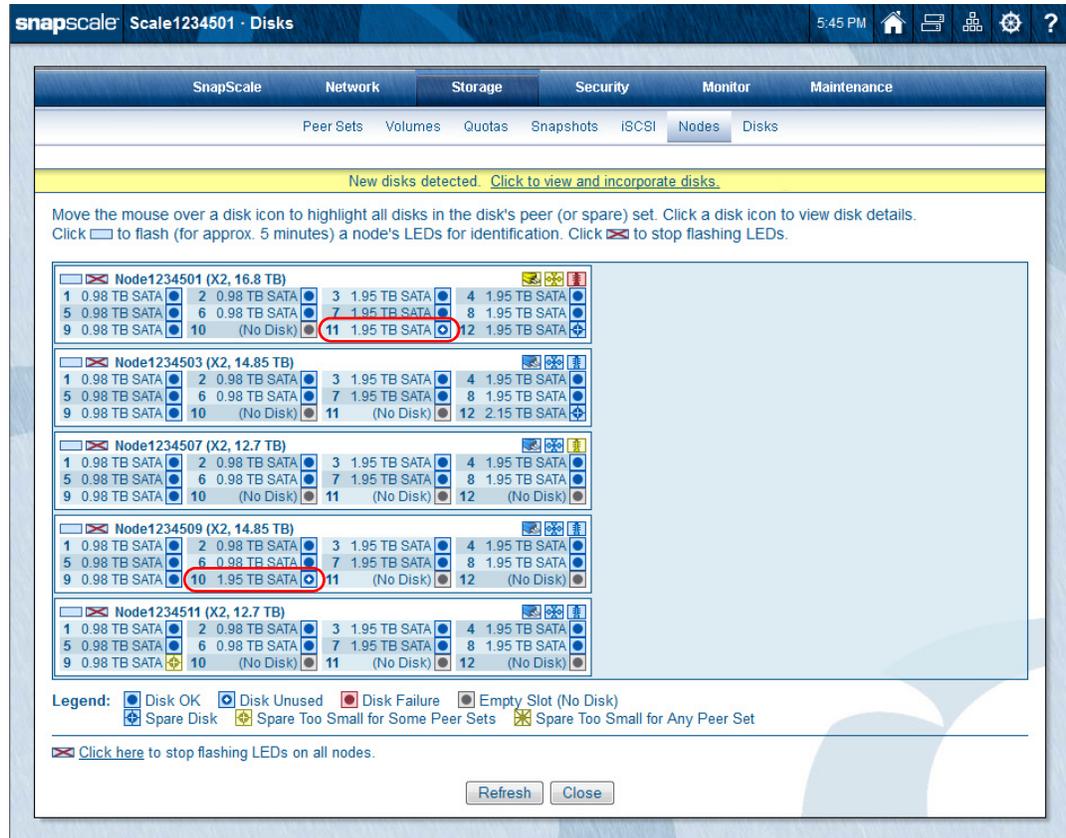


3. Repeat [Step 2](#) for **each** remaining drive carrier being installed.



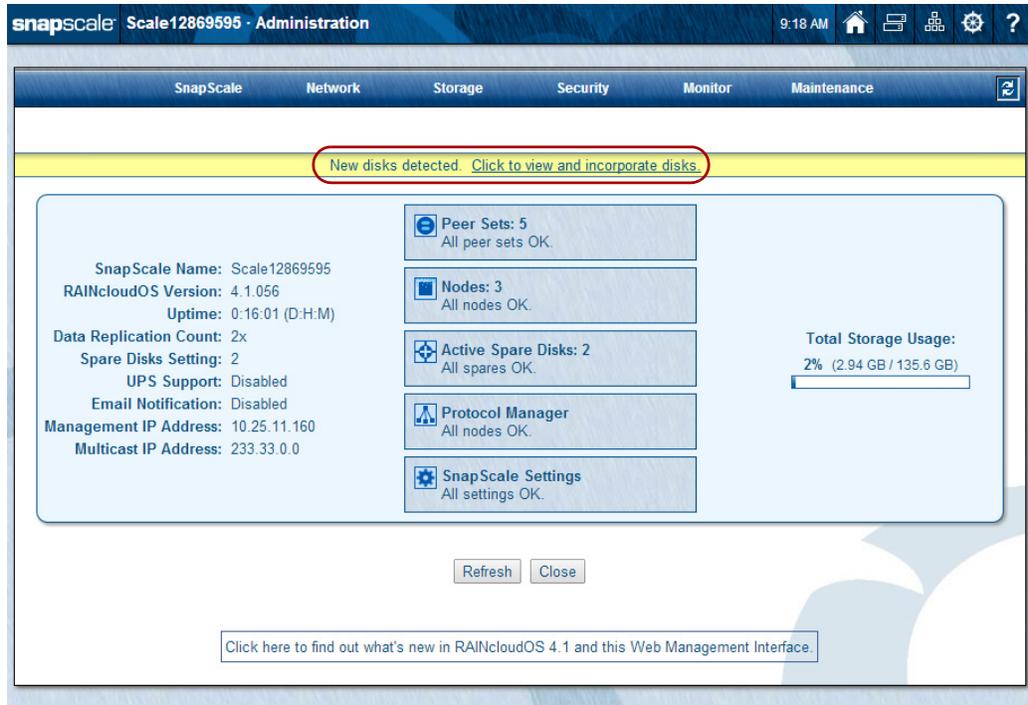
**IMPORTANT:** To maintain proper airflow and cooling, a drive carrier or a blank carrier must be installed in every slot. No empty slots are allowed.

On the **Disks** page, any newly detected drives show a disk unused icon next to the drive. It may take a minute or two before the drives appear as unused and the new disks detected banner is displayed. The alert link in the banner takes you to the **New Disks Detected** page for incorporation.

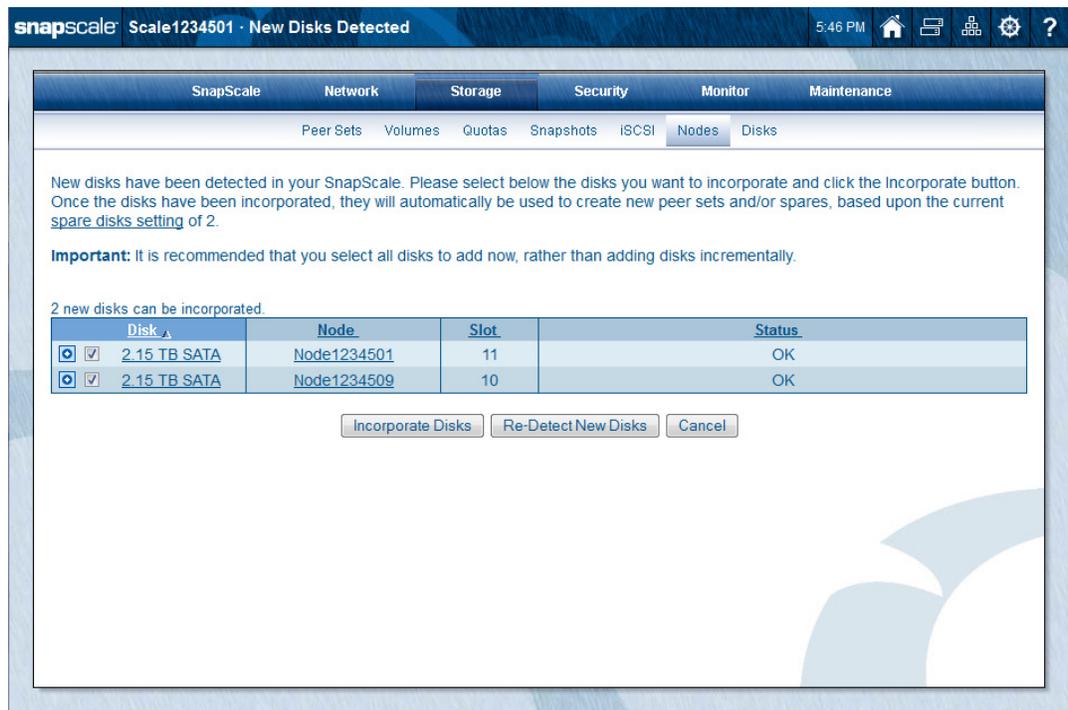


### Peer Set/Hot Spare Incorporation

When newly installed drives are detected, SnapScale first auto-incorporates drives to fix any failed peer set and fulfill any reserved spare count. The Web Management Interface then displays an alert banner about the new drives and the **New Disks Detected** page is activated.



Click the link in the alert to go to the **New Disks Detected** page. The page displays all new drives available to be formed into new peer sets or used as new spares.



The boxes next to the drive name can be unchecked to remove a drive from the incorporation list. However, to balance peer set and spare creation, it is recommended that all drives be incorporated at the same time.

Click **Incorporate Disks** to begin the process.

- If enough new or spare drives exist on different nodes based on your Data Replication Count (2x or 3x), new peer sets are formed as long as the spare count is satisfied.
- If there are not enough drives or they are not on different nodes, the drives are used to create additional hot spares.

After incorporation, the drives are displayed normally as peer sets or hot spares on the **Disks** page.

The Security options control the access to the SnapScale cluster and its data.



SnapScale cluster authentication validates a user’s identity by requiring the user to provide a registered login name (User ID) and corresponding password. The node ships with predefined local users and groups that allow administrative (admin) and guest user access to the cluster via all protocols.

Administrators may choose to join the SnapScale cluster to a Windows Active Directory domain, and CIFS/SMB and AFP clients can then authenticate to the cluster using their domain credentials. To accommodate NFS clients, the SnapScale cluster can also join an LDAP or NIS domain, and the SnapScale can look up user IDs (UIDs) and group IDs (GIDs) maintained by the domain for configuration of quotas and ID mapping. For authentication control beyond the guest account, Mac and FTP client login credentials can be created locally on the cluster. See [User and Group ID Assignments on page 143](#).

### Topics in Security Options

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

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

## Security Considerations

SnapScale cluster default security configuration provides one share to the entire volume. All network protocols for the share are enabled, and all users are granted read-write permission to the share via the guest account. By default, the `guest` user is disabled in SMB but enabled for HTTP and FTP.

Network clients can initially access the cluster using the guest account (where enabled), but if you require a higher degree of control over individual access to the filesystem for these clients, you must create local accounts (or use Windows Active Directory security for CIFS/SMB clients).

Local users or groups are created using **Security > Local Users** or **Security > Local Groups** in the Web Management Interface. Local users are also used for administrative access to the cluster through the Web Management Interface, SSM, or the CLI through SSH.

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

- **admin** – The local user admin account is used to log into the Web Management Interface. The default password for the admin account is also *admin*.
- **guest** – The local user guest account requires no password.
- **admingrp** – The Admin group account includes the default admin user account. Any local user accounts created with admin rights are also automatically added to this group.

### Guidelines for Local Authentication

These password authentication guidelines are for both users and groups.

**Duplicating Client Login Credentials for Local Users and Groups.** To simplify user access for Windows Workgroup, duplicate their local client logon credentials on the SnapScale cluster by creating local accounts on the cluster that match those used to log on to client workstations. This strategy allows users to bypass the login procedure when accessing the cluster.



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

---

**Default Local Users and Groups.** Default users and groups *admin*, *guest*, and *admingrp* appear on the list of users or groups on the user or group management pages, but they cannot be deleted or modified (although the admin password can be changed).

**Changing Local UIDs or GIDs.** The SnapScale cluster 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 cluster makes these fields editable.

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

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

**Local Account Management Tools.** The following tools are available for creating, modifying, and editing local user and group accounts:

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

## User and Group ID Assignments

SnapScale uses the POSIX standard to assign UIDs or GIDs, in which each user and group must have a unique ID. This requirement applies to all users and groups on the cluster, including Windows Active Directory, LDAP, NIS, and local users.

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

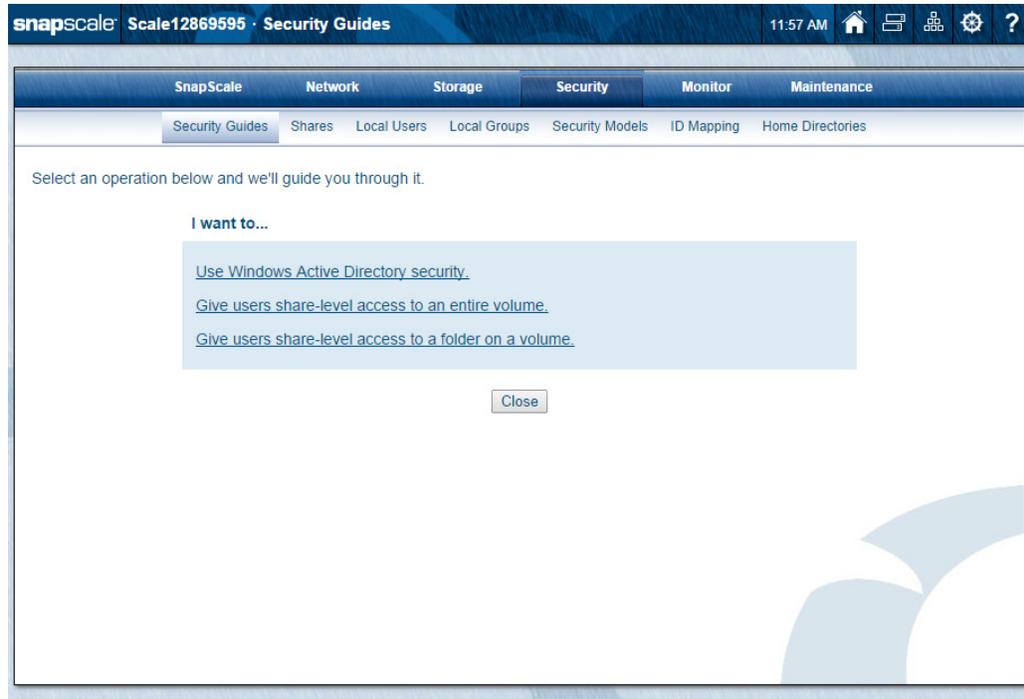
- UIDs and GIDs from 0 to 100 are unavailable for use. If you try to assign a UID or GID that is less than 101 (or in use by Windows, LDAP, or NIS domain), you will get an error message.
- When the cluster automatically generates UIDs or GIDs for imported Windows domain users or groups, UIDs or GIDs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

The `nfsnobody` and `nobody` user IDs (UID 65534 and 65535, respectively) and GIDs are reserved. They are not mappable to other IDs, nor is another ID mappable to `nfsnobody` or `nobody`.

## Security Guides

Security Guides are special wizards to guide you through:

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



## Security Guide for Windows Active Directory

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

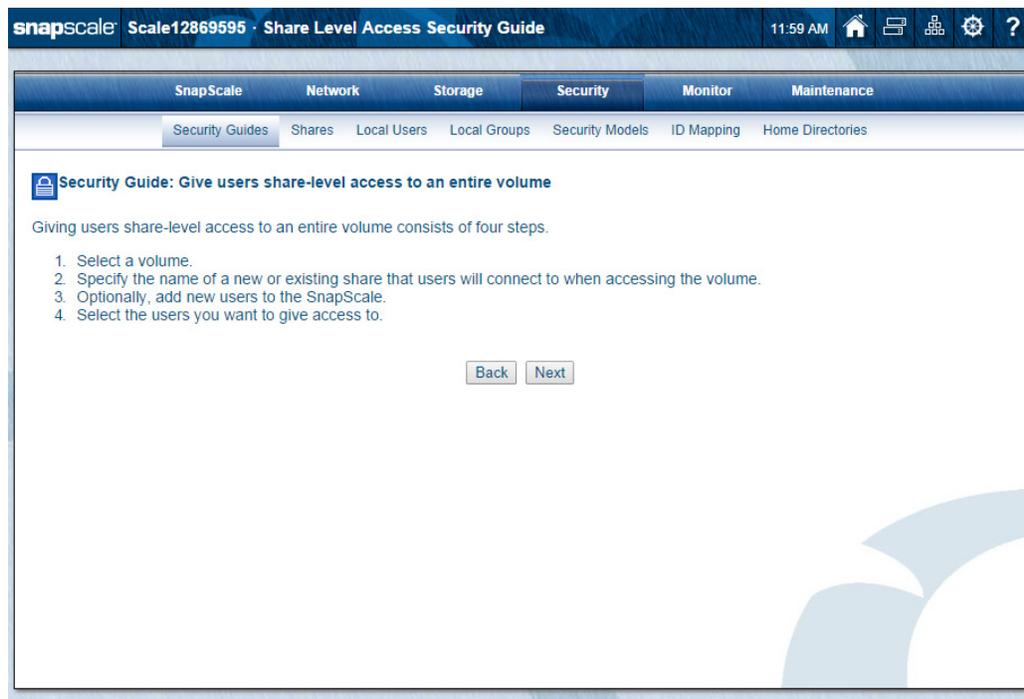
**NOTE:** You cannot join an Active Directory domain if NTP is enabled. If you see such a message, click the NTP link to change your settings.

When the cluster joins a domain, it does so as a single unit under the cluster name, and all nodes operate equally under the cluster name to authenticate against the domain. This provides multi-point domain-authenticated access to the cluster through each node.



## Security Guide for Entire Volume Access

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



## Security Guide for Folder Access on Volume

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



## Shares

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

2 shares. Click a share name to modify a share's properties or to delete a share.

Share	Volume	Path	Access	NFS Access	Protocols	Attributes
SHARE1	Volume1	/	<a href="#">Open</a>	<a href="#">Default</a>	SMB NFS HTTP FTP	-
SHARE2	Volume2	/	<a href="#">Open</a>	<a href="#">Default</a>	SMB NFS HTTP FTP	-

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

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

[Create Share](#) [Refresh](#) [Close](#)

SnapScale supports file access in Windows, and Unix networks, as well as access via HTTP and FTP. New shares are created by default with full read-write access to all users, subject to the filesystem permissions on the share target directory. The first step to securing a cluster is to specify access at the individual share level. Administrators can assign read-write or read-only share access to individual Windows (and local) users and groups.

## Create Shares

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

The screenshot shows the 'Create Share' page in the SnapScale web interface. The page title is 'Scale12869595 · Create Share' and the time is 12:04 PM. The navigation menu includes SnapScale, Network, Storage, Security (selected), Monitor, and Maintenance. The sub-menu includes Security Guides, Shares (selected), Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. The main content area contains the following fields and options:

To create a new share, specify a name, volume, and path to a folder.

Name:

Volume:

Path:

Description:  (optional)

Security model for path:

Create share with full read and write access for all users  
 Create share with Admin-only access and proceed to Share Access page

[Advanced Share Properties >>](#)

By clicking the **Advanced Share Properties** link, additional options are displayed. Use these options to hide the share from network browsing, select the protocols supported and create a snapshot share associated with this share.

The screenshot shows the 'Create Share' page in the SnapScale web interface. The page title is 'Scale12869595 · Create Share'. The navigation menu includes SnapScale, Network, Storage, Security (selected), Monitor, and Maintenance. The sub-menu includes Security Guides, Shares (selected), Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. The main content area has the instruction: 'To create a new share, specify a name, volume, and path to a folder.' The form includes:
 

- Name: SHARE2
- Volume: Volume2
- Path: / (with a 'Browse' button)
- Description: (optional)
- Security model for path: Windows/Unix
- Radio buttons for:
  - Create share with full read and write access for all users
  - Create share with Admin-only access and proceed to Share Access page
- A link 'Advanced Share Properties' circled in red, with a red arrow pointing to it and the text 'Click to display the options shown below'.
- Checkbox: Hide this share (the share will be hidden from network browsing)
- Section: 'Select only the protocols that are needed for network users to access this share.' with checkboxes for:
  - Windows (SMB)
  - Linux/Unix (NFS)
  - Web (HTTP, HTTPS)
  - FTP, FTPS
- Text: 'You can create a "snapshot" share that is associated with this new share. The snapshot share allows access (using identical security) to snapshots of the data that the new share references.'
- Checkbox: Create Snapshot Share
- Checkbox: Hide this share (the snapshot share will be hidden from network browsing)
- Text field: Snapshot Share Name (SHARE2\_SNAP)
- Buttons: Create Share, Cancel

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

## Creating a Share

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

1. Accept the default **share name** or enter a new one.  
To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including spaces).
2. Choose the **volume** you need from the drop-down menu.
3. Select from the following **path options**:
  - **To create a share to the entire volume** – The current Path field defaults to the root path of the volume. Simply leave it blank if this is the desired configuration.
  - **To create a share to a folder on the volume** – Browse to the folder to which you want to point the share, click the folder name, and click **OK**.

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

4. If desired, enter a **description** to clarify the purpose of the share.
5. Choose a **security model for path** by selecting **Windows/Unix**, **Windows**, or **Unix** from the drop-down list.

The option defaults to the current security model at the specified path. If changed to a different security model, the change will propagate to all files and subdirectories underneath. For more information, see [Security Models on page 173](#).

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

Choose either **Create share with full read and write access for all users**, or **Create share with Admin-only access and proceed to Share Access page** to configure the user share access. For more information, see [Share Access Behaviors on page 154](#).

**NOTE:** If selecting the share with Admin-only access option and the share has NFS enabled, be sure to configure the NFS access settings afterward.

7. To further configure the share, click **Advanced Share Properties** and enter any of the following:

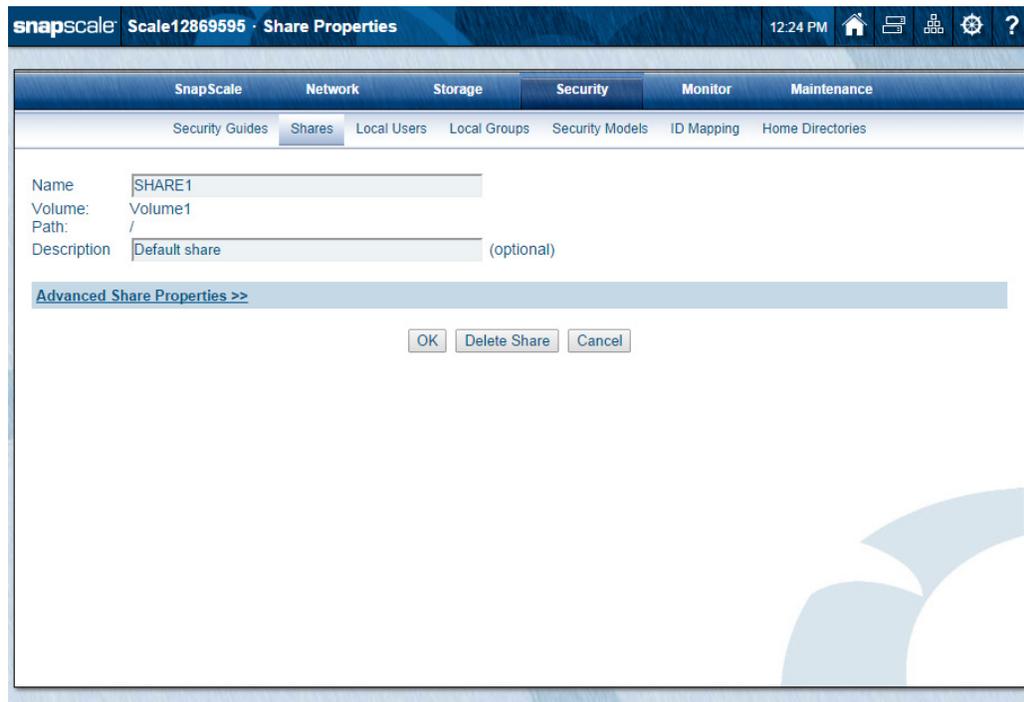
Option	Description
<b>Hide this Share</b>	Select this option if you want the share to be hidden from network browsing using SMB, HTTP/HTTPS, and FTP protocols (but not NFS).
<b>Protocols</b>	Select the access protocols for the share: Windows (SMB), Linux/Unix (NFS), Web (HTTP/HTTPS), and FTP/FTPS. Check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

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

## Edit Share Properties

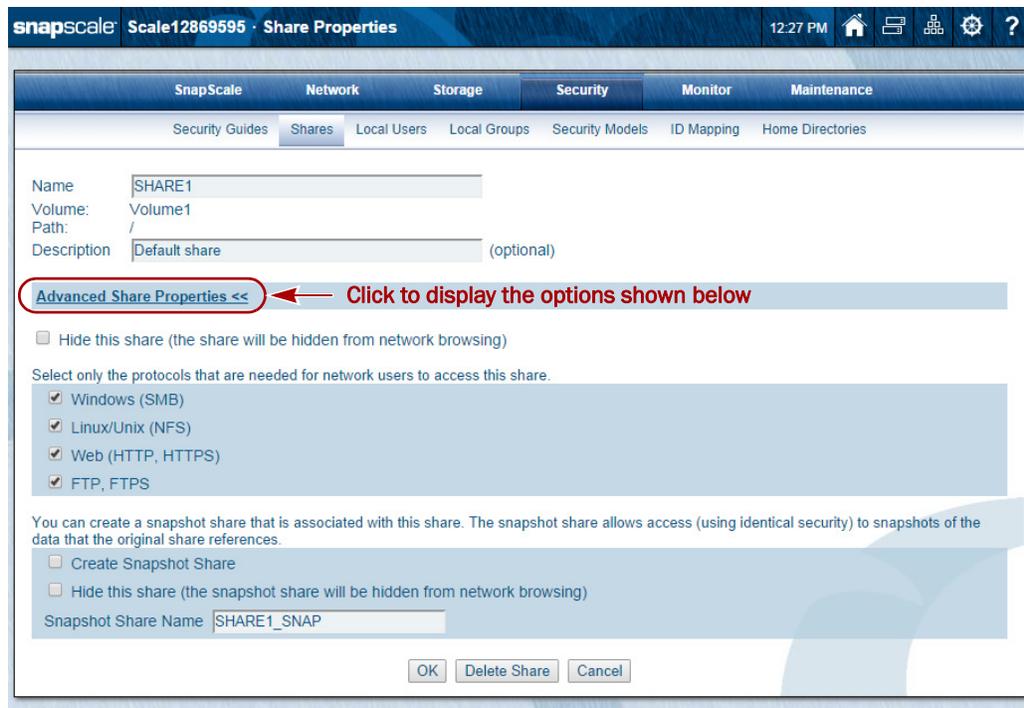
**NOTE:** You cannot change the volume (or path) of a share once it is created. If you need to change the share's volume, you must delete the share and create a new share on the other volume.

Once a share has been created, you can change its name, description and the advanced properties. To edit the properties, go to **Security > Shares > Share Properties** (displayed by clicking the share name in the table).



By clicking the **Advanced Share Properties** link, additional options are displayed. Use these options to hide the share from network browsing, select the protocols supported, and create a snapshot share associated with this share.

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



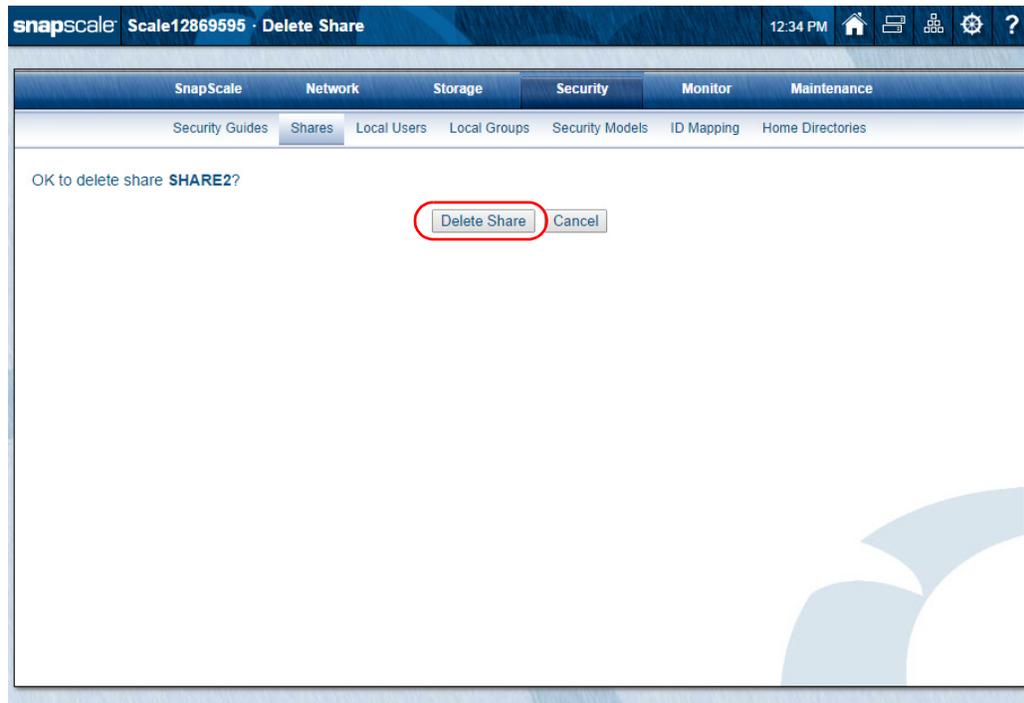
Option	Description
<b>Name</b>	Accept the default share name or enter a new one. If you change the default, observe the following guidelines: <ul style="list-style-type: none"> <li>• Make sure the share name is unique on this cluster.</li> <li>• To ensure compatibility with all protocols, share names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

## Delete Shares

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

1. At the **Delete Share** page, click **Delete Share**.

2. At the confirmation page, click **Delete Share** again.



## Configuring Share Access

In **Security > Shares**, in the **Access** column, click the link next to the share you want to configure. The **Share Access** page is displayed. You can set access levels for the share, as well as grant or deny access to specific users and groups.

**NOTE:** To add a new user to a share, you must first create the user, then add that user to the share. Please see [Local Users](#) on page 161 for information on creating new users.

The top screenshot shows the SnapScale web interface for the 'Shares' page. It displays a table with columns: Share, Volume, Path, Access, NFS Access, Protocols, and Attributes. Two shares are listed: SHARE1 (Volume1, Path /) and SHARE2 (Volume2, Path /). The 'Access' column for both shares has an 'Open' button circled in red. Below the table, there are buttons for 'Create Share', 'Refresh', and 'Close'. An important security note is also present.

The bottom screenshot shows the 'Share Access' page for 'SHARE1'. It features two main sections: 'Users and groups with specific access to share. (1)' and 'Search for users and groups.'. The first section shows 'AllUsers' with a 'Remove' button. The second section shows a list of local users: admin, Al, Betsy, Freddy, Harry, and Vicky. There are buttons for '<- Add', 'Full Access', and a search bar.

### Share Access Behaviors

Administrators tasked with devising security policies for SnapScale cluster will find the following share access behaviors informative:

- **Share access defaults to full control** – The default permission granted to users and groups when they are granted access to the share is full control. You may restrict selected users and groups to read-only access.
- **User-based share access permissions are cumulative** – An SMB, HTTP, and FTP user's effective permissions for a resource are the sum of the permissions that you assign to the individual user account and to all of the groups to which the user belongs in the **Share Access** page. For example, if a user has read-only permission to the share, but is also a member of a group that has been given full-access permission to the share, the user gets full access to the share.

- **NFS access permissions are not cumulative** – An NFS user's access level is based on the permission in the NFS access list that most specifically applies. For example, if a user connects to a share over NFS from IP address 192.168.0.1, and the NFS access for the share gives both read-write access to "\*" (All NFS clients) and read-only access to 192.168.0.1, the user will get read-only access.
- **Interaction between share-level and file-level access permissions** – When both share-level and file-level permissions apply to a user action, the more restrictive of the two applies. Consider the following examples:

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

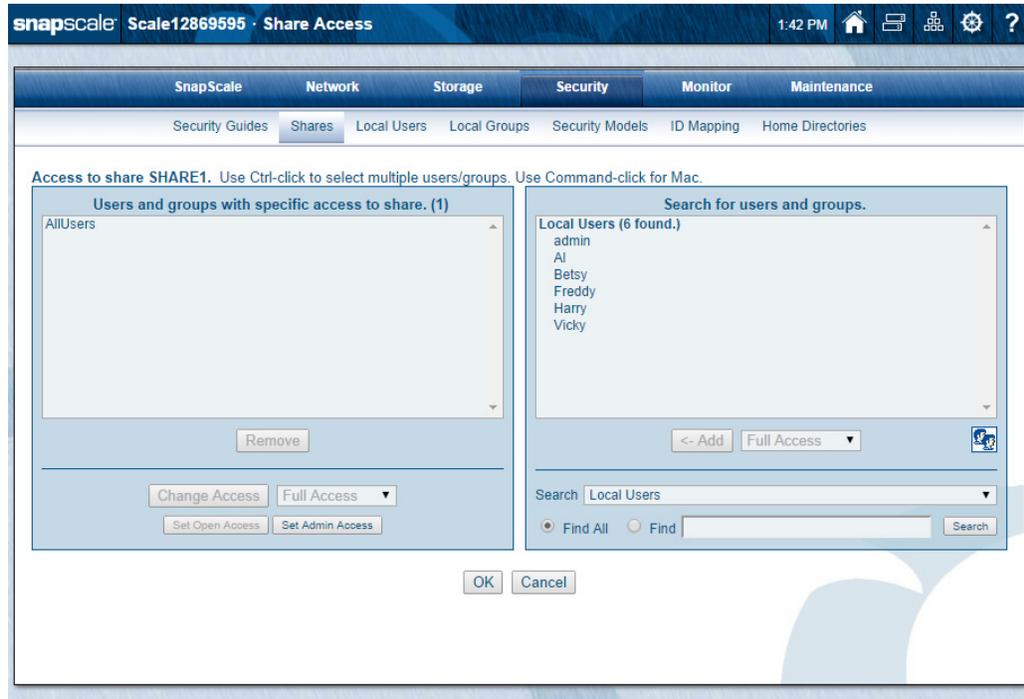
Share Level	File Level	Result
Full control	Read-only to File A	Full control over all directories and files in SHARE1 <i>except</i> where a more restrictive file-level permission applies. The user has read-only access to File A.

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

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

## Setting User-based Share Access Permissions

Share permissions for Windows, HTTP, and FTP users are configured from **Security > Shares** by clicking the link in the **Access** column of the share you want to configure. Share permissions for NFS are configured and enforced independently. See [NFS Access for Shares on page 158](#) for more information.



User-based share access permissions apply to users connecting over SMB, HTTP, or FTP. Users and groups with assigned share access permissions appear in the list on the left (**Users and groups with specific access to share**). To search for those without assigned access, use the box on the right (**Search for users and groups**).

The default permission granted to users and groups when they are granted access to the share is **Full Access**. You may restrict selected users and groups to **Read-only Access**.

Share-Level Access Permissions	
<b>Full</b>	Users can read, write, modify, create, or delete files and folders within the share.
<b>Read-only</b>	Users can navigate the share directory structure and view files.

1. Display the **Share Access** page (**Security > Shares > *access\_link***).
2. To **add** share access permissions for a user or group:
  - a. At the bottom, using the drop-down list, select the **domain** or **local user/group list** to search.

**NOTE:** For domains that require authentication (showing an "(A)" after the name), after selecting the domain name, enter the **User Name** and **Password** for that domain. The user name and password can be for any user in the domain and are used to retrieve basic information (like the user and group lists) from the domain.

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

When entering a search string:

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

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

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

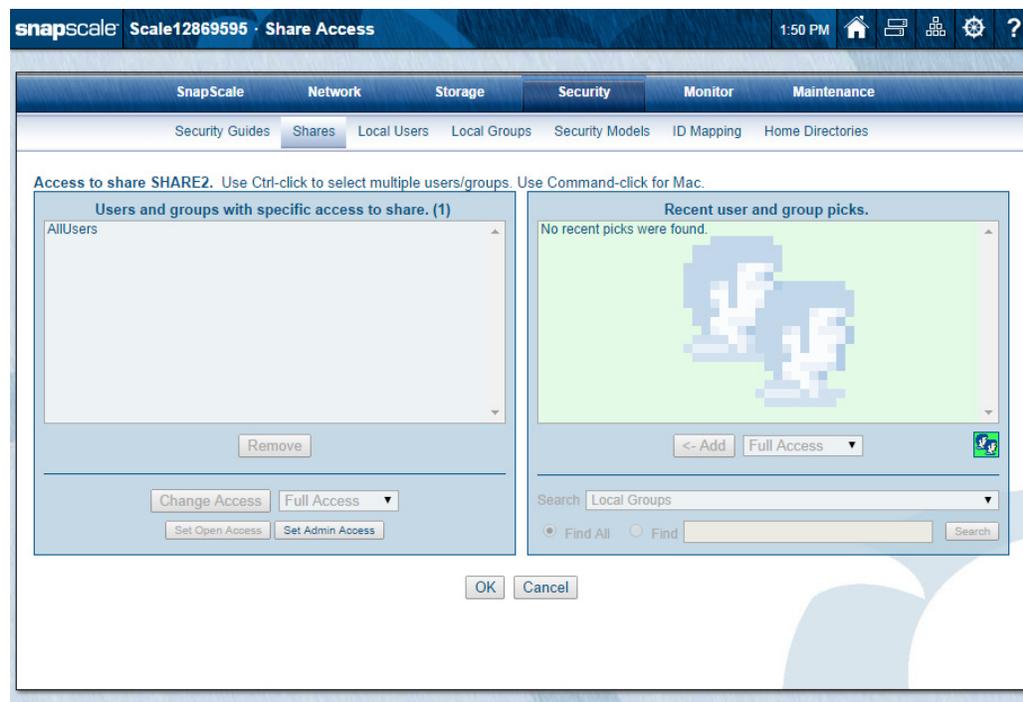
d. Select one or more **names** in the list.

Users that already have access are shown in purple font with a plus sign (+) in front of their name.

e. Choose either **Full Access** or **Read Only** from the drop-down list.

f. Click **Add**.

**NOTE:** To display recent user or group picks, click the **faces** (👤) icon. A list with a green background is displayed. Click the now green icon again to return to the normal search box.



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

- Select one or more **users or groups** in the left box.
- Click **Remove**.

4. To change **access permissions** for a user or group, select one or more users or groups in the left box, then select either **Full Access** or **Read Only** from the drop-down list, and finally click **Change Access**.
5. To quickly specify either Open or Admin-only **access** for the entire share, click either **Set Open Access** or **Set Admin Access**.
6. Click **OK** to save share permissions.

### NFS Access for Shares

**NOTE:** Multiple shares pointing to the same target directory must have the same NFS access settings. The Web Management Interface applies the same NFS access for all shares pointing to the same directory.

To configure NFS access, click the link shown in the **NFS Access** column for the share you want to configure. You can configure NFS access to the share using standard Linux “exports” file syntax.

On the **Shares** page, click the name of the access type listed in the **NFS Access** column to open the **NFS Share Access** page.

The top screenshot shows the SnapScale web interface for the 'Shares' page. It features a navigation bar with 'SnapScale', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Below this is a sub-navigation bar with 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. The main content area displays '2 shares' and a table with columns: Share, Volume, Path, Access, NFS Access, Protocols, and Attributes. The table contains two rows: 'SHARE1' (Volume1, Path /, Access Open, NFS Access Default, Protocols SMB-NFS-HTTP-FTP, Attributes -) and 'SHARE2' (Volume2, Path /, Access Open, NFS Access Default, Protocols SMB-NFS-HTTP-FTP, Attributes -). The 'NFS Access' column for both rows is circled in red. Below the table is an 'Important Security Note' and buttons for 'Create Share', 'Refresh', and 'Close'. A red arrow points from the circled 'Default' text in the 'NFS Access' column to the 'NFS Share Access' page in the bottom screenshot.

The bottom screenshot shows the 'NFS Share Access' page. It has the same navigation bars as the top screenshot. The main content area includes instructions to 'Add a host using the following (basic) permissions to the exports list below.' with options for 'SnapScale default options' (checked) and 'Read-only (ro,async)'. There is an 'NFS Host' input field and an 'Add Host' button. Below this is a text box for 'NFS access (exports) for share SHARE1:' containing the text: `*(rw,insecure,async,root_squash,no_all_squash)`. At the bottom are 'OK' and 'Cancel' buttons.

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

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

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

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

The entry options are explained in the following table:

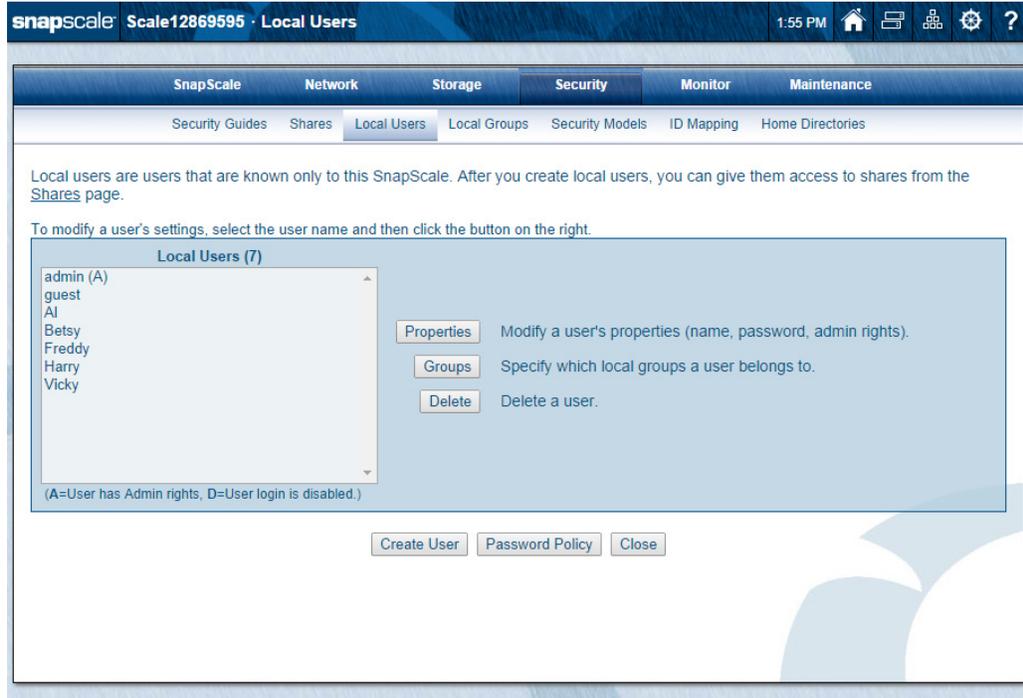
Entry Code	Meaning
<b>Asterisk</b>	All NFS clients
<b>ro</b>	The directory is shared read only (ro).
<b>rw</b>	The client machine will have read and write (rw) access to the directory.
<b>insecure</b>	Turns off the options that require requests to originate on an Internet port less than IPPORT_RESERVED (1024).
<b>root_squash</b>	Forces users connected as root to interact as the “nobody” user (UID 65534). This is the SnapScale default.
<b>no_root_squash</b>	no_root_squash means that if root is logged in on your client machine, it will have root privileges over the exported filesystem. By default, any file request made by user root on the client machine is treated as if it is made by user <b>nobody</b> on the cluster. (Exactly which UID the request is mapped to depends on the UID of user <b>nobody</b> on the cluster, 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 cluster. 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.
<b>async</b>	Tells a client machine that a file write is complete – that is, has been written to stable storage – when NFS has finished handing the write over to the filesystem.
<b>no_all_squash</b>	Allows non-root users to access the nfs export with their own privileges.

**Using the Add Host Option.** Follow these steps:

1. Select **one** of the following options:
  - **SnapScale Default Options** – Inserts the default options as described above.
  - **Read Only** – Inserts the read only option only.
  - **Both** – Inserts default options, but substitutes read only for read/write.
2. Do **one** of the following in the NFS host text box:
  - **To apply the options to all NFS hosts** – Leave this field blank.
  - **To apply the options to specific hosts** – Enter one or more IP addresses.
3. Click **Add Host**.

## Local Users

The **Local Users** page provides all the options to manage local users. Local users are users that are known only to the cluster being accessed. Each SnapScale cluster comes with two predefined users: admin and guest. The admin user has full Administrator rights. Go to **Security > Local Users** to view settings or make changes.



## Create a User

Click **Create** to create a new user on this cluster. Enter the user data, select any special options, and click **Create User** again.

The screenshot shows the 'Create Local User' page in the SnapScale web management interface. The page title is 'Scale12869595 · Create Local User'. The navigation menu includes SnapScale, Network, Storage, Security (selected), Monitor, and Maintenance. The Security sub-menu includes Security Guides, Shares, Local Users (selected), Local Groups, Security Models, ID Mapping, and Home Directories. The main content area contains the following form:

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

Name:

Full Name:  (optional)

Password:

Confirm Password:

User ID (UID):

Disable user login

Grant admin rights to this user (A local user with admin rights will be able to access this Web Management Interface.)

### To Create a Local User

1. On the **Local Users** page, click **Create User**.
2. On the **Create Local User** page that opens, enter the requested **information**:

Option	Description
<b>Name</b>	Use up to 31 alphanumeric characters and the underscore.
<b>Full Name</b>	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
<b>Password</b>	Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
<b>Confirm Password</b>	Type the chosen password again for verification.
<b>User ID (UID)</b>	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see <a href="#">User and Group ID Assignments on page 143</a> .

Option	Description
<b>Disable User Login</b>	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box.  This box can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
<b>Exempt from Password Expiration and Character Requirements</b>	This checkbox is only visible if Password Policy is enabled. Check this box to exempt this user from password expiration and character requirement policies.
<b>Grant Admin Rights To This User</b>	Check this box to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

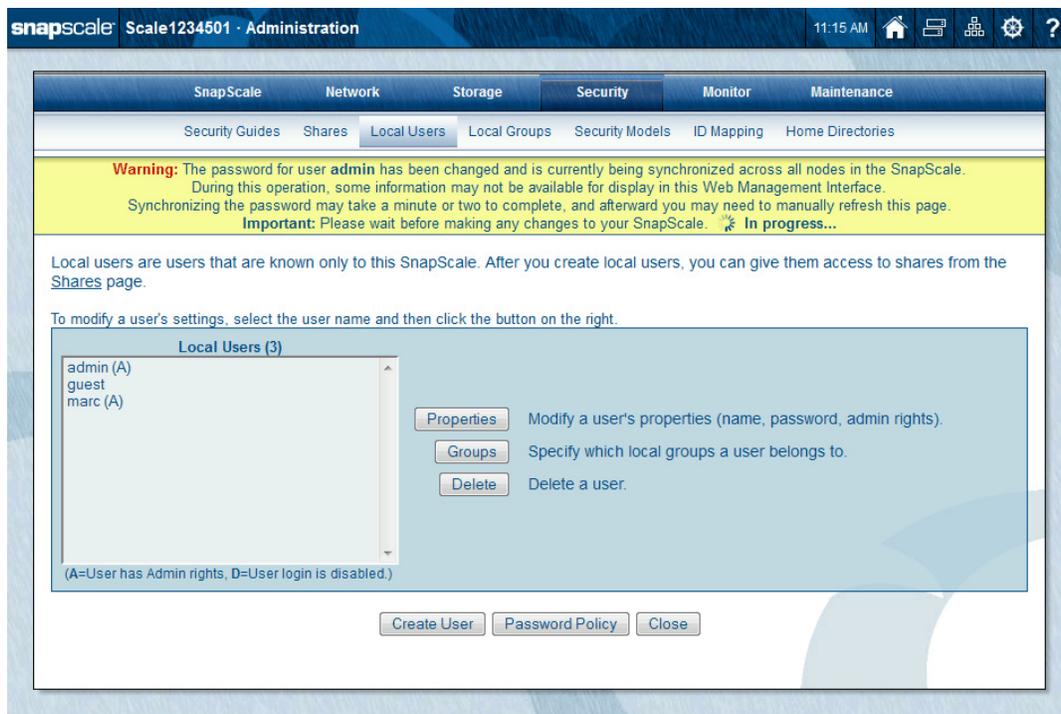
3. Click **Create User** again to create the user account.

## Edit User Properties

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

The screenshot shows the 'Local User Properties' dialog box in the SnapScale web interface. The user being edited is 'Freddy' with the full name 'Fredrick Sandstone' and UID '18001'. The 'Exempt this user from password expiration and character requirements' checkbox is highlighted with a red box and a red arrow, with a note indicating it is only shown if Password Policy is enabled. Other options include 'Disable user login' and 'Grant admin rights to this user'. The dialog has 'OK' and 'Cancel' buttons at the bottom.

**NOTE:** When changing the Admin password, it can take a minute or so to synchronize the new password across all nodes. During this time, a warning message is displayed in the Web Management Interface. While the cluster, all the nodes, and all the data are fully accessible during this synchronization process, you should wait for the message to disappear before making further changes to your SnapScale cluster.



### To Edit Local User Properties

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

Option	Description
<b>Name</b>	<b>NOTE:</b> Cannot be modified. Instead, delete and recreate the user with the same UID if you need to change the user name.
<b>Full Name</b>	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
<b>Password</b>	Passwords are case-sensitive. To keep the existing password, leave this field blank.
<b>Password Verify</b>	Type the chosen password again for verification. To keep the existing password, leave this field blank.
<b>User ID (UID)</b>	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see <a href="#">User and Group ID Assignments on page 143</a> . <b>NOTE:</b> Changing a user's UID may alter filesystem access permissions that apply to that UID. In addition, any existing permissions for a UID previously assigned to a user that are changed to a different UID may become active if another user is created with the same UID. Carefully consider security configuration on existing files and directories before changing the UID of a user.

Option	Description
<b>Disable User Login</b>	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box.  This box can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
<b>Exempt from Password Expiration and Character Requirements</b>	<b>NOTE:</b> This box is only visible if Password Policy is enabled. Check this box to exempt this user from password expiration and character requirement policies.
<b>Grant Admin Rights To This User</b>	Check this box to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **OK**.

## Local User Password Policies

**NOTE:** Local users can be individually exempted from password expiration and character requirements. This may be necessary for some special users, such as users configured to perform backups. See [To Create a Local User on page 162](#) for procedures to set password policy for local users. Also, the built-in *admin* user is automatically exempt from all password policies.

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

### To Set Password Policy for Local Users

1. On the **Local Users** page, click **Password Policy**.
2. On the **Local Users Password Policy** page, check the **Enable Password Policy** box.

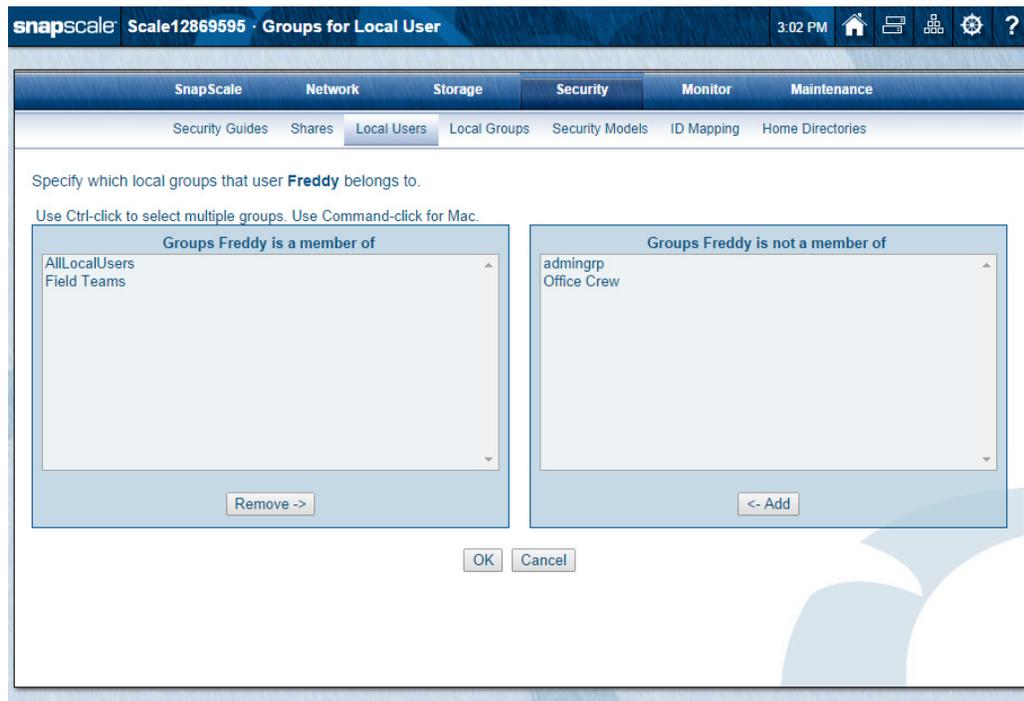
3. Enter the following **information**:

Option	Description
<b>Character Requirements</b>	Select the alpha/numeric/special character requirements for the password from the drop-down list.
<b>Minimum Number of Characters</b>	Check this box to enable the policy, then enter the minimum number of characters required for the password.
<b>Disable Login After <i>n</i> Attempts</b>	Check this box to enable the policy, then enter the number of times a user can fail to login before the system locks the user out. This applies to failed logins when connecting to any node in the cluster.  <b>NOTE:</b> To unlock a user, clear the <b>Disable User Login</b> box for the user in the <b>Local Users</b> page.
<b>Re-enable a Disabled Login After <i>n</i> Minutes</b>	If you have defined a limit to the number of times a user can fail to log in, you can also check this box and enter a time period after which the system will allow the user to log in again. This saves the administrator from having to manually re-enable the user.
<b>Expire Password After <i>n</i> Days</b>	Check this box to enable the policy, then enter the number of days before the password must be changed.  <b>NOTE:</b> Local users with expired passwords can change their passwords at: <a href="http://&lt;cluster_name&gt;/changepassword">http://&lt;cluster_name&gt;/changepassword</a> .

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

## Assign User to Group

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

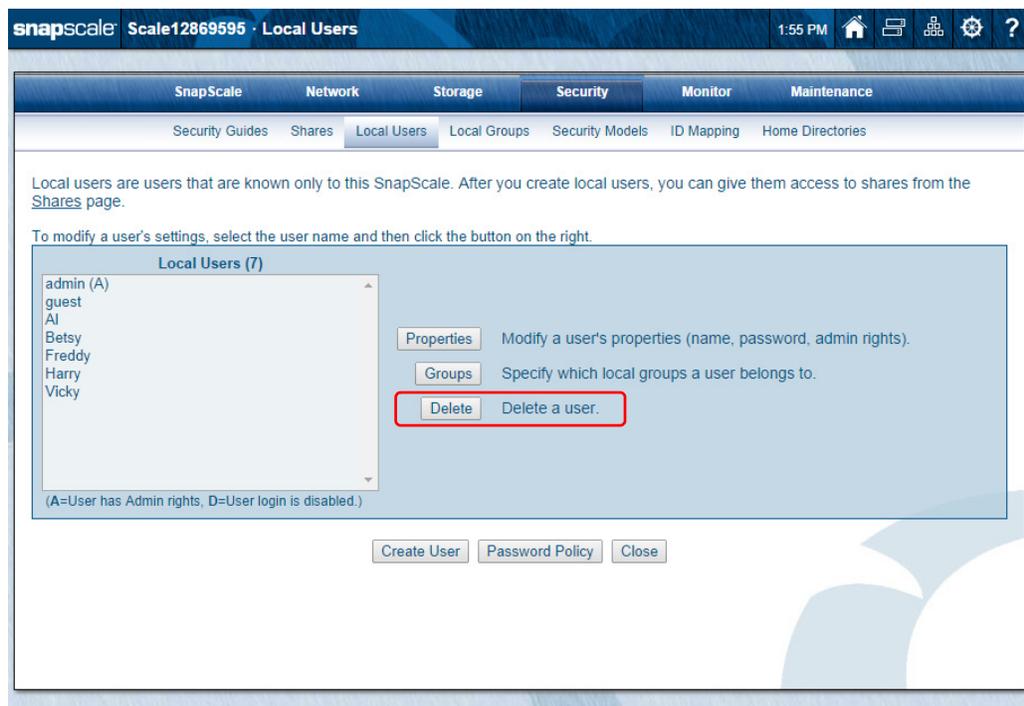


## Add or Remove Users from Groups

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

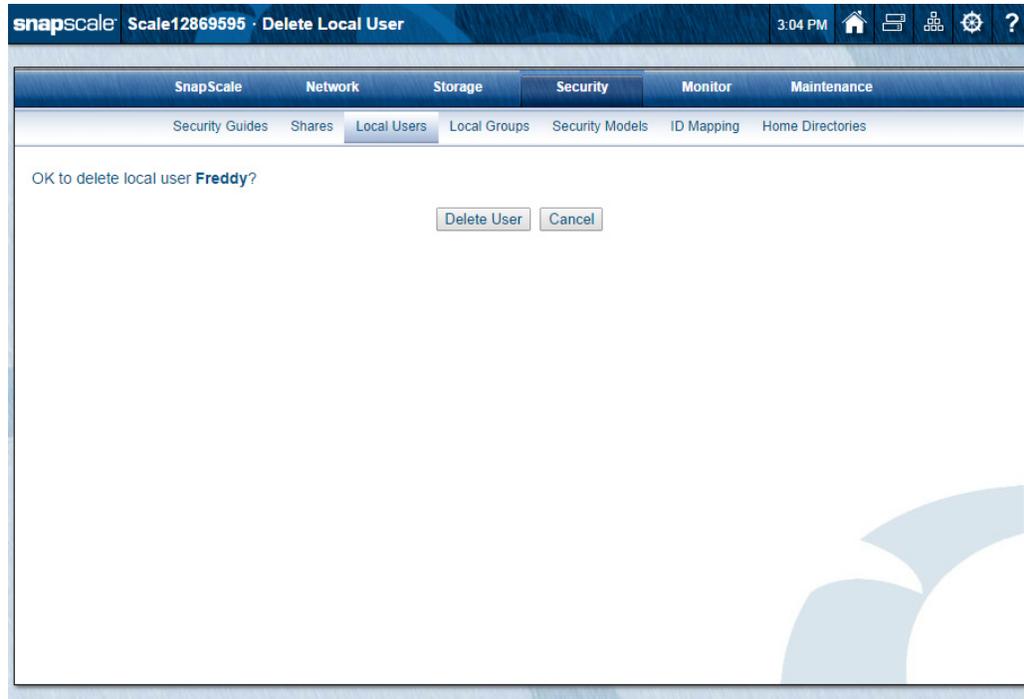
## Delete Local User

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



### To Delete a Local User

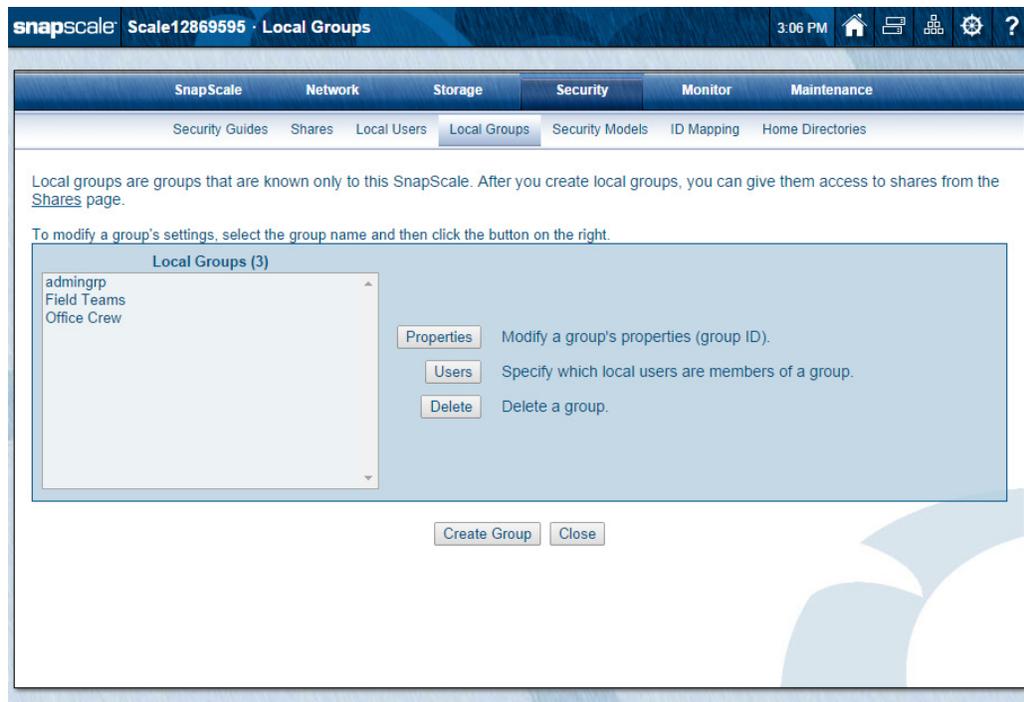
1. On the **Local Users** page, select the user to be deleted.
2. Click **Delete**.  
The confirmation page is displayed.



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

## Local Groups

The **Local Groups** page (**Security > Local Groups**) provides all the options to manage local groups. Local groups are groups of local users that are known only to the cluster being accessed. Each SnapScale cluster comes with one predefined group (**admingrp**).



## Create New Group

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

### To Create a New Local Group

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

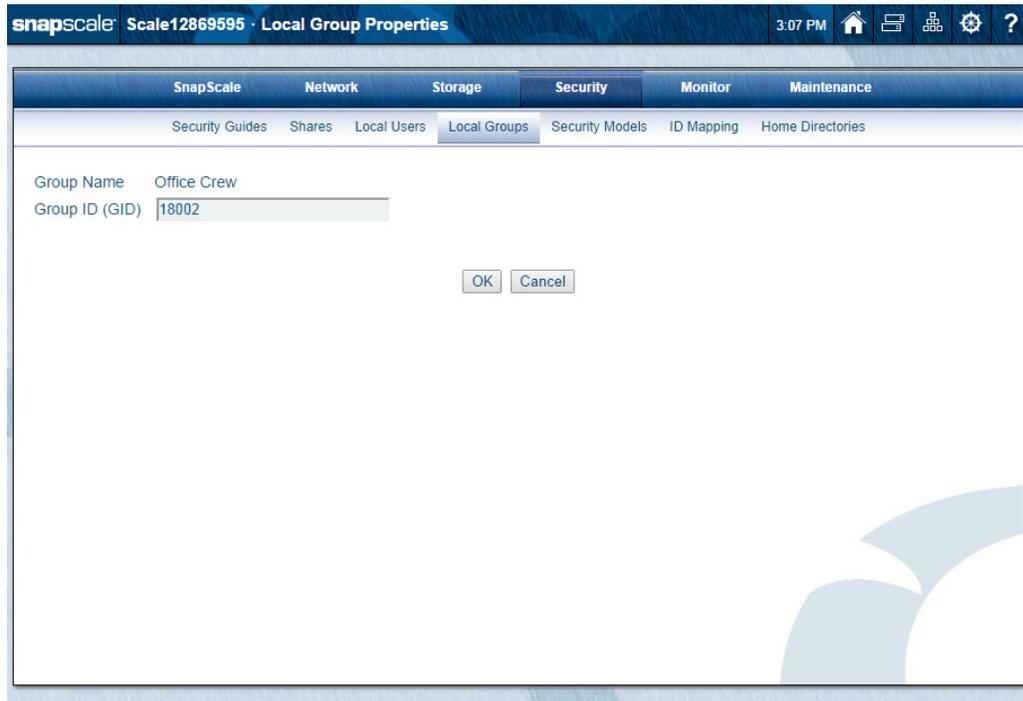
2. Enter the following **information**:

Option	Description
<b>Group Name</b>	Use up to 31 characters (alphanumeric and the underscore only).
<b>Group ID (GID)</b>	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see <a href="#">User and Group ID Assignments on page 143</a> .

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

## Edit Group Properties

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



### To Edit Local Group Properties

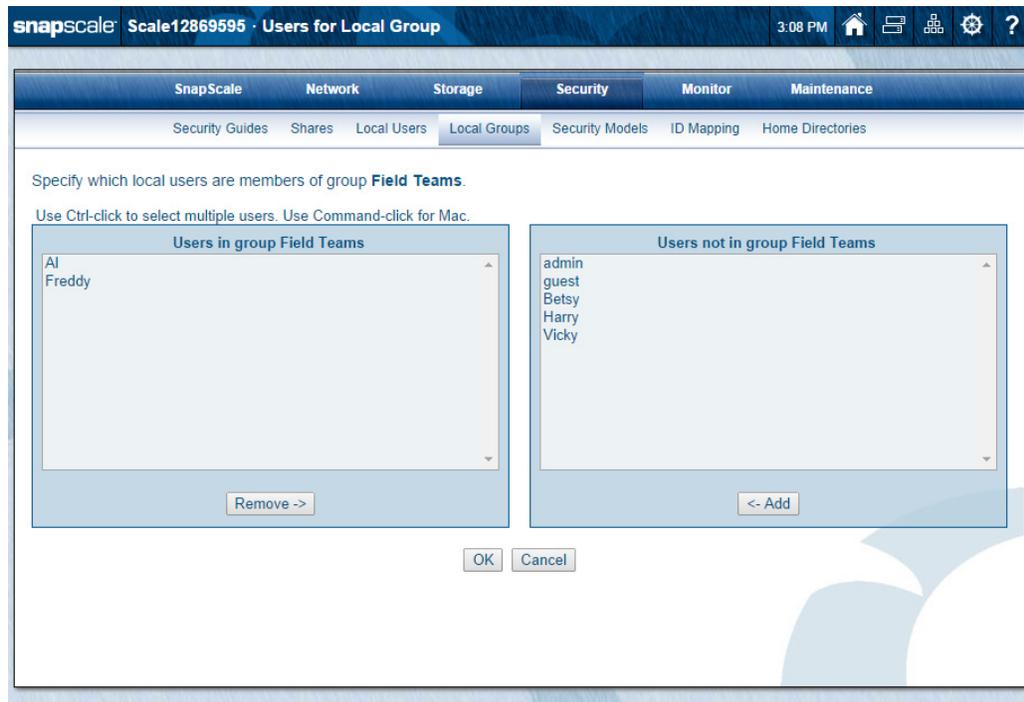
1. On the **Local Groups** page, select the group you want to edit and click **Properties**.
2. On the **Local Groups Properties** page that opens, you can change the **GID**. For information on available UID ranges, see [User and Group ID Assignments on page 143](#).

**NOTE:** Changing a group's GID may alter filesystem access permissions that apply to that GID. In addition, any existing permissions for a GID previously assigned to a group that are changed to a different GID may become active if another group is created with the same GID. Carefully consider security configuration on existing files and directories before changing the GID of a group.

3. Click **OK**.

## Specify Users in Group

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

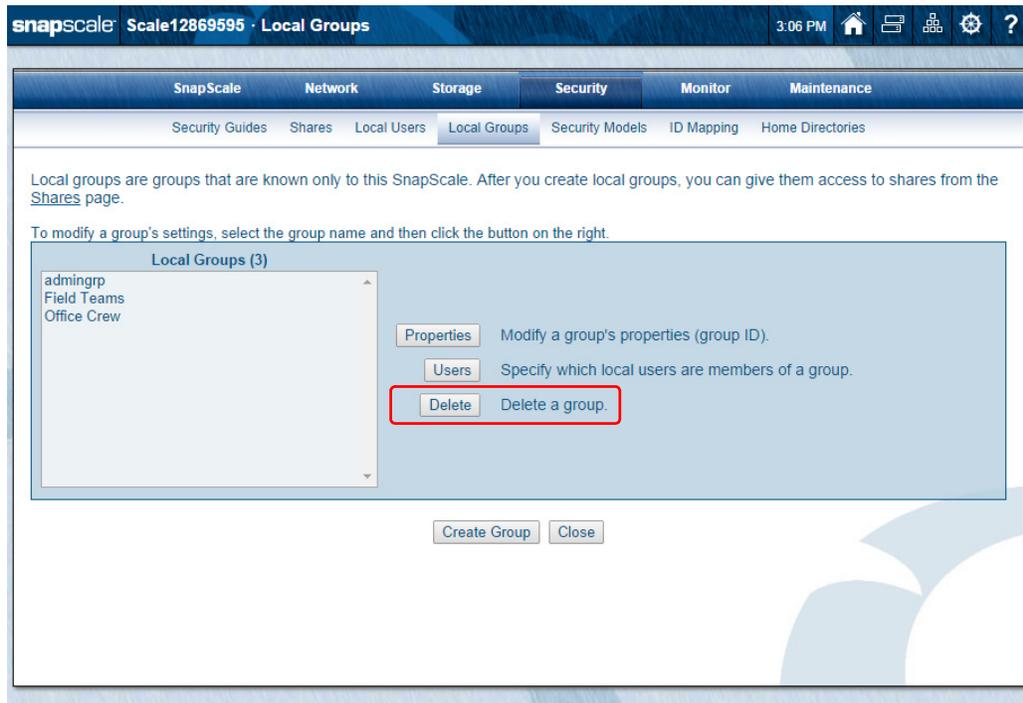


### Add or Remove Group Users

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

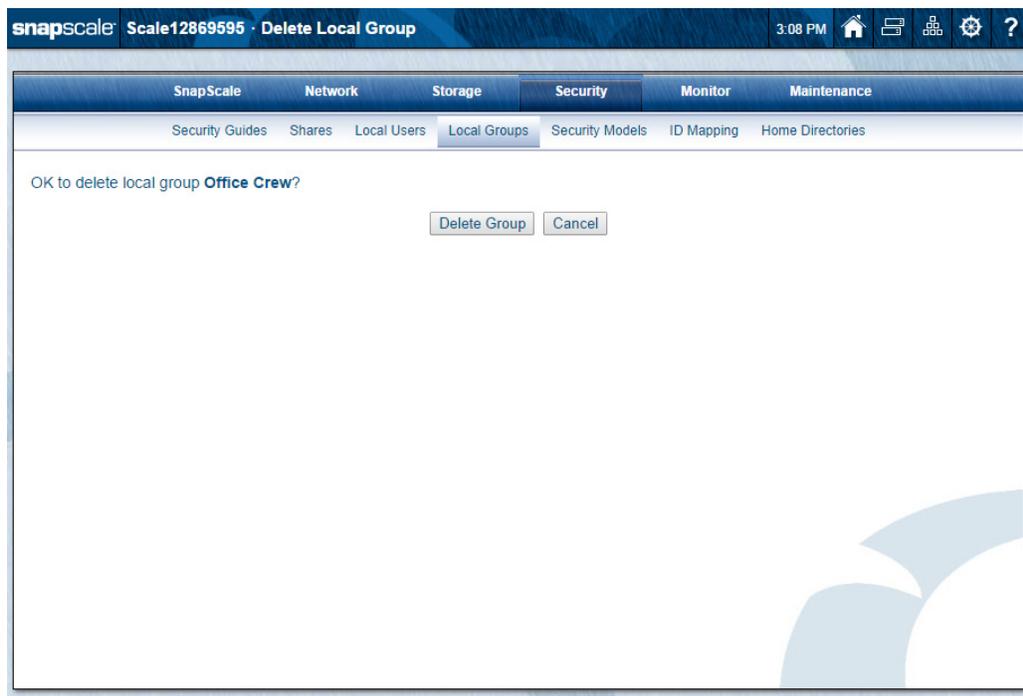
## Delete Group

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



### Delete a Group

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



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

## Security Models

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

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

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

**NOTE:** Folders created in a volume default to the security model of that volume. The folder's security model may differ from the personality of the folders (for example, folders with a **Windows/Unix** security may have a Unix personality).

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

## Managing Volume Security Models

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

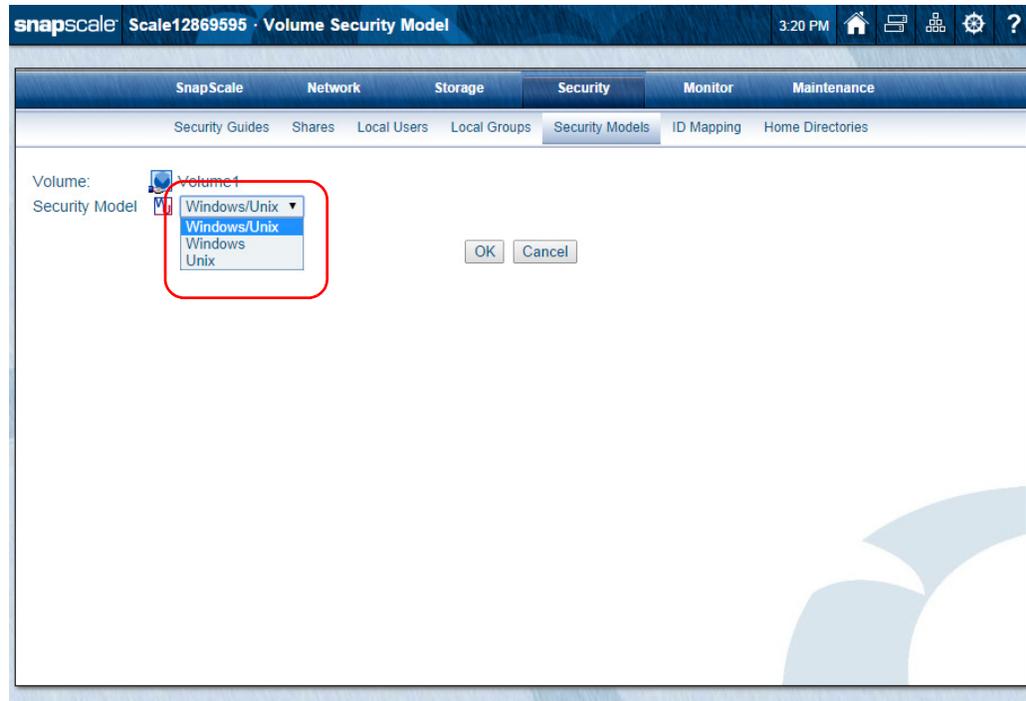
The screenshot displays the SnapScale Security Models management page. The top navigation bar includes SnapScale, Scale12869595, and Security Models. The main content area has a sub-navigation bar with Security Guides, Shares, Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. The main text instructs users to manage the security model for an entire volume. A table lists two volumes: Volume1 with a Windows/Unix security model and Volume2 with a Windows security model. Below the table, a status box shows a successful conversion on 2014-09-23 3:18:11 PM for the path /hd/cfs/shares/project2, with the security model set to Windows and 1 of 1 files/folders converted (100%). Refresh and Close buttons are located at the bottom of the status box.

Volume	Security Model for Volume
Volume1	Windows/Unix
Volume2	Windows

Security model conversion status.  
**Status:** Conversion completed successfully on 2014-09-23 3:18:11 PM.  
 Path converted: /hd/cfs/shares/project2  
 Security model: Windows  
 Apply security model to all files & sub-folders in this folder: No  
 Files/folders converted: 1 of 1 (100%)

Refresh Close

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

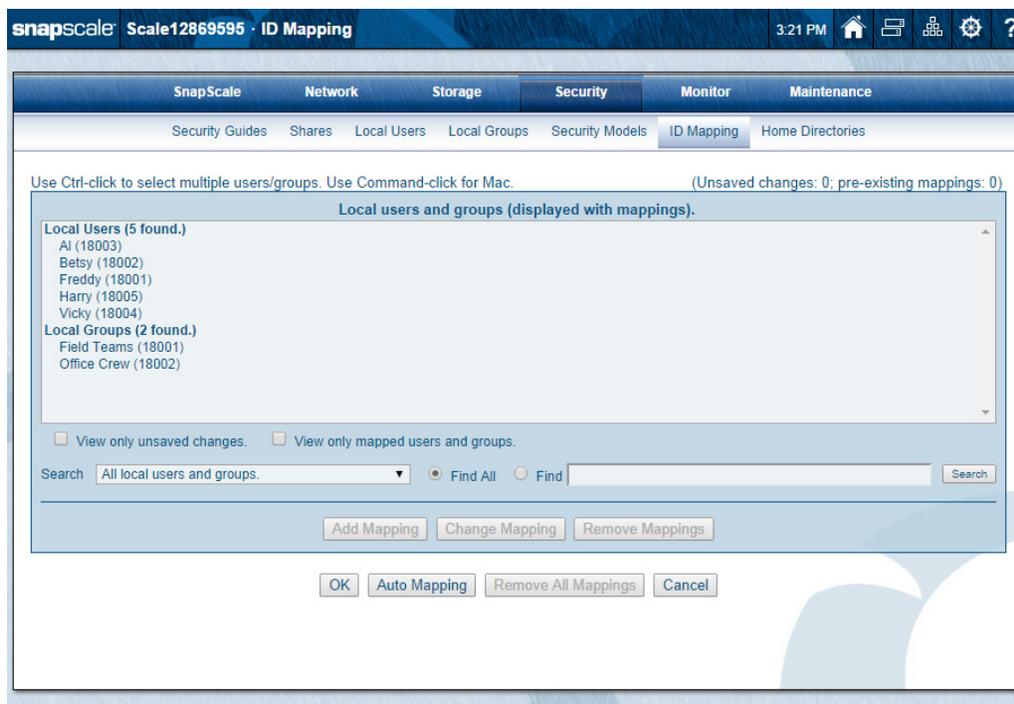


4. At the confirmation message, click **Apply Security Model**.

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

## ID Mapping

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



Select a local, LDAP, or NIS user or group from the displayed list on the default page. You can then use **Add Mapping** to map the user's UID or group's GID to that of a Windows domain user or group. **Change Mapping** is used to change existing mappings. **Remove Mappings** removes one or more mappings while **Remove All Mappings** removes all mappings that had been previously established.

Options to simplify the discovery of a desired user or group to manage their ID mapping search options are presented at the bottom of the selection pages:

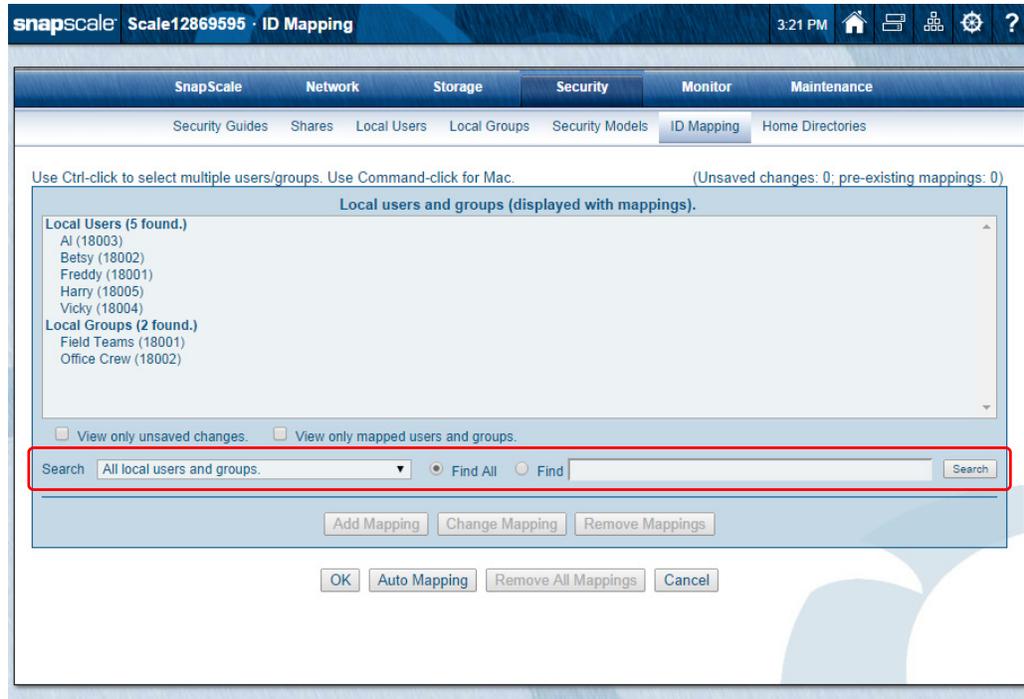
- Check **View only unsaved changes** to display only mapping changes that have not yet been applied.
- Check **View only mapped users and groups** to display only local, LDAP, or NIS users and groups that have been mapped to a Windows domain user or group.

### Add Mapping

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

Follow this procedure to map a user or group:

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



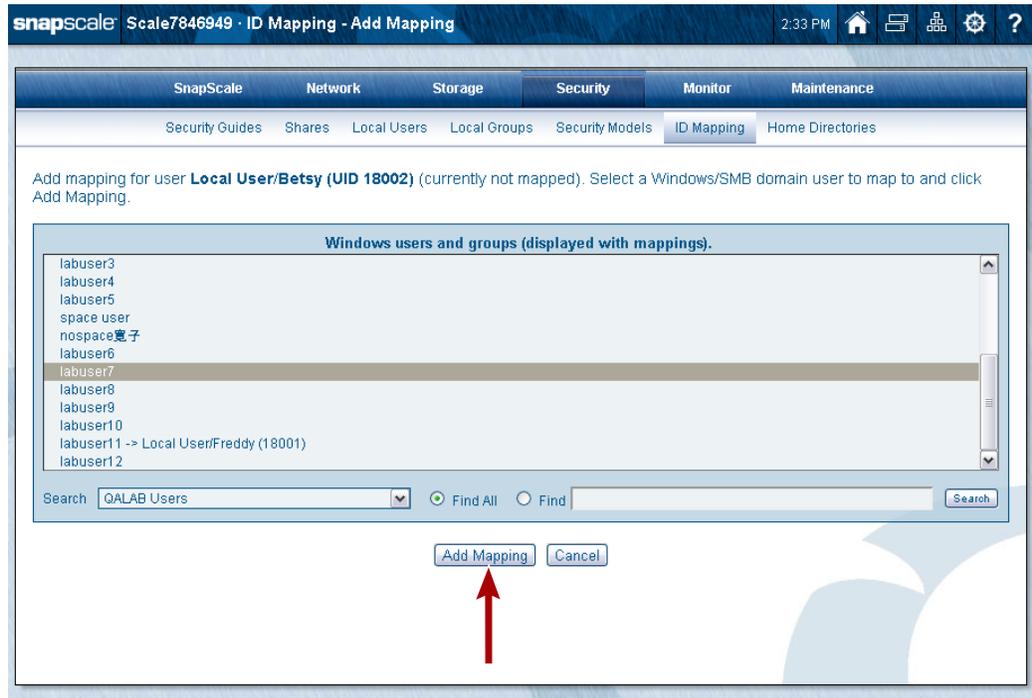
- a. At the bottom of the list, using the **Search** drop-down list, select the local, LDAP, or NIS user or group **list** to be searched.
  - b. Select **Find** and enter the **search string** (or select **Find All**).  
Enter the exact **name** (or a string with a wild card "\*" before or after) as the search string.
  - c. Click **Search** to display any matches.
2. Select a **user or group** from the results list and click **Add Mapping**.
  3. At the **Add Mapping** page, to find the user/group you want to map to, select the Windows domain **user or group list**, the scope of the search, enter a search string if needed, and click **Search**.
    - To search for a specific user or group, use either **Find All** or a **Find** search string (wild card "\*" before or after is allowed).



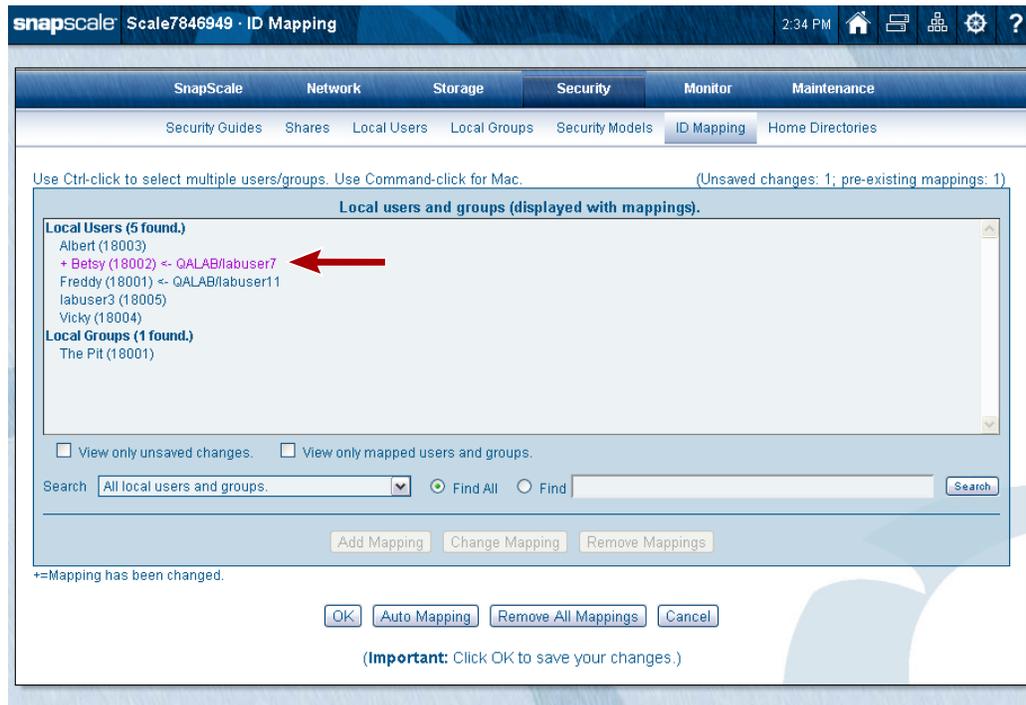
- For domains that **REQUIRE** authentication (showing an **(A)** after the name), select the domain name, enter the user name and password for that domain, and use either **Find All** or a **Find** search string (using the first few letters of the user/group name).



- On the rare occasion you need to search for a Windows domain that's not listed (remote domain), select a Windows domain from the **Search** drop-down list through which to search, then enter in the **Find** box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.
4. From the search results, select the Windows domain **user/group** to which you want to map the local, LDAP, or NIS user, and click **Add Mapping**.



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



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

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



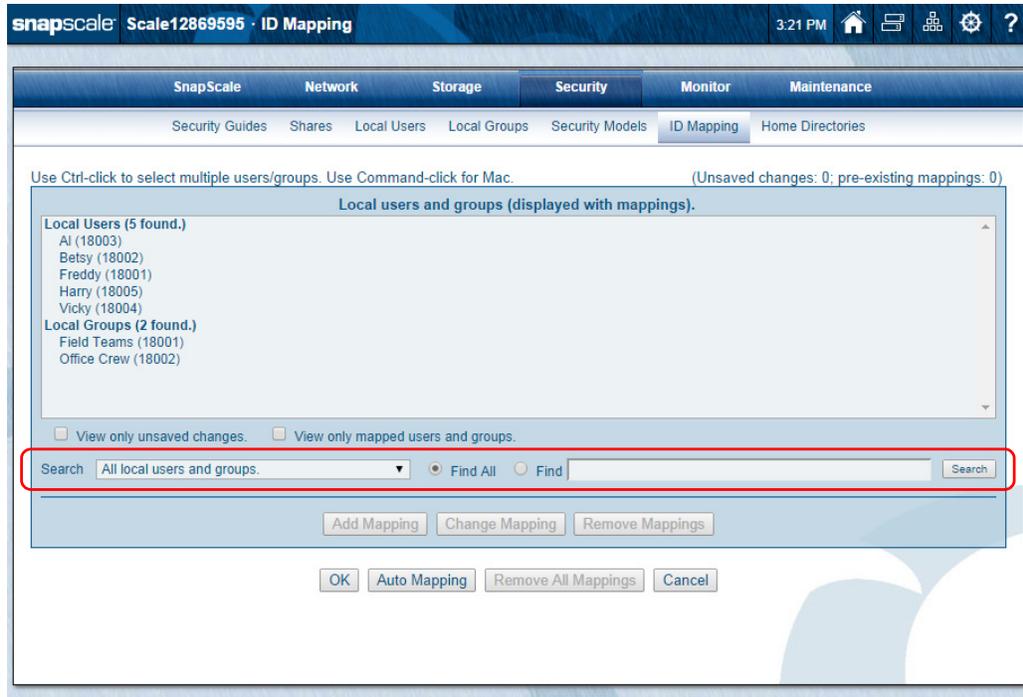
**IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated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

See [Filesystem Updates on page 187](#) for more details.

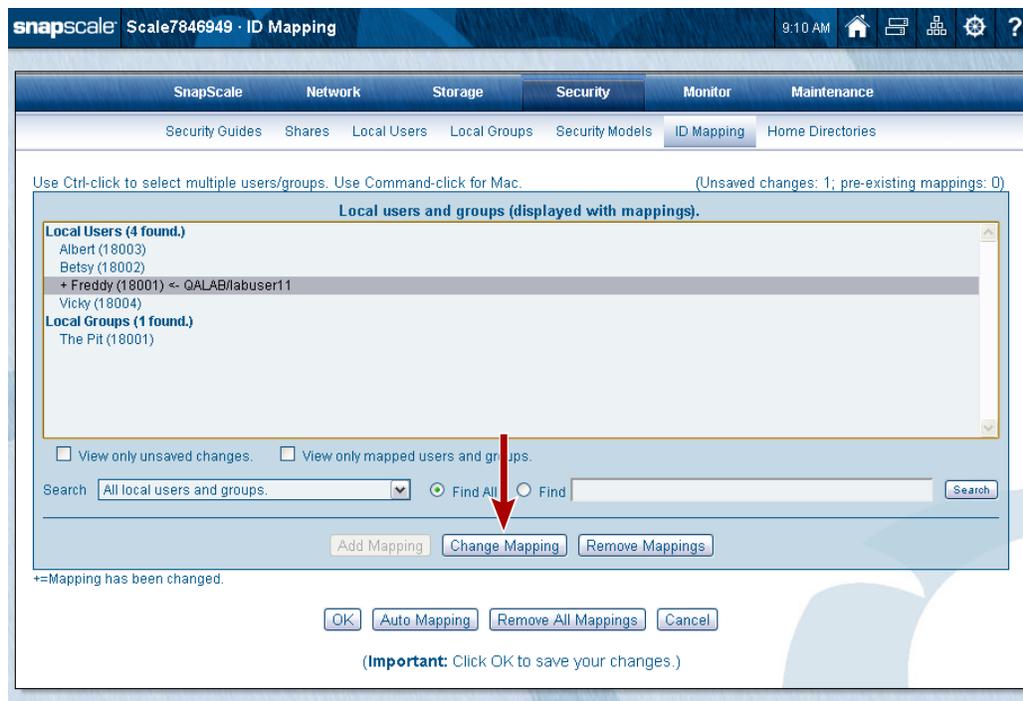
## Change Mapping

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

1. If the desired user or group to be changed does not appear in the default page list, use the **search option** to locate them.



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



- At the **Change Mapping** page, to find the user/group you want to map to, select the Windows domain **user or group list**, the scope of the search, enter a search string if needed, and click **Search**.
  - To search for a specific user or group, use either **Find All** or a **Find** search string (wild card "\*" before or after is allowed).

- For domains that **REQUIRE** authentication (showing an **(A)** after the name), select the domain name, enter the user name and password for that domain, and use either **Find All** or a **Find** search string (using the first few letters of the user/group name).

- On the rare occasion you need to search for a Windows domain that's not listed (remote domain), select a Windows domain from the **Search** drop-down list through which to search, then enter in the **Find** box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, **remote\_domain\user\_name**). After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.
- From the search results, select the Windows domain **user/group** you want to re-map the local, LDAP, or NIS user to and click **Change Mapping**.

- Repeat **Steps 1–4** until all **changes** are made.

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



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

See [Filesystem Updates on page 187](#) for more details.

## Auto Mapping

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

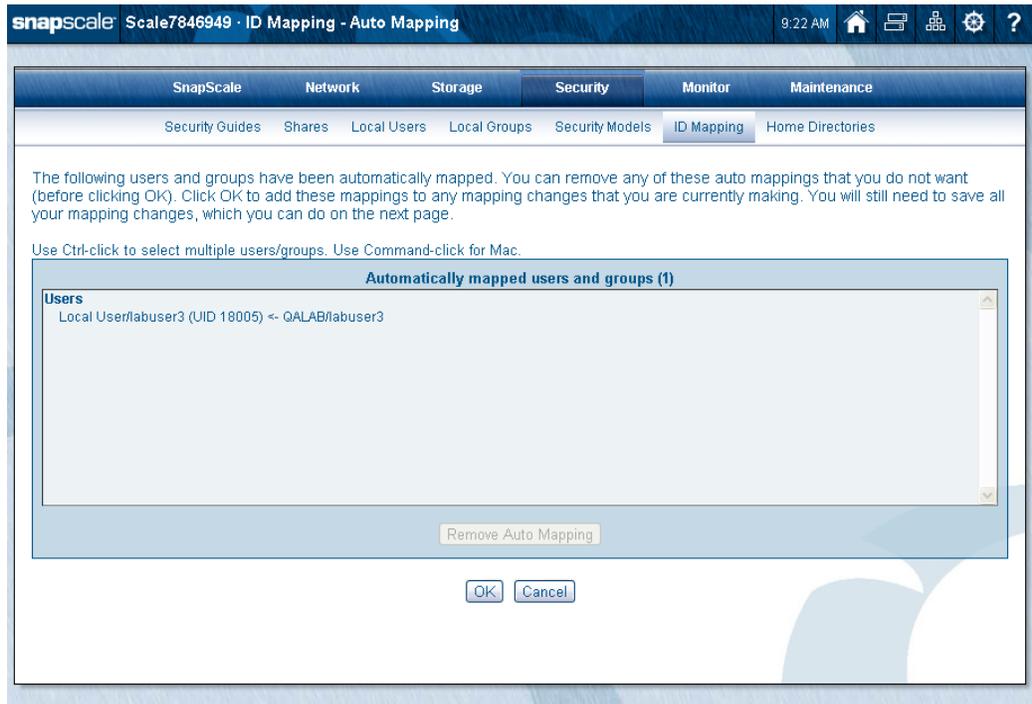
1. Click **Auto Mapping** to generate a **list** of Windows domain users/groups that have the same name as your local, LDAP, or NIS users and groups.

Domain, local, LDAP, and NIS user/group lists are compared. The matches are automatically queued. Users and groups already mapped are not affected.

2. At the **Auto Mapping** confirmation page, click **View Auto Mappings** to display a page summarizing your changes.



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



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



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

See [Filesystem Updates on page 187](#) for more details.

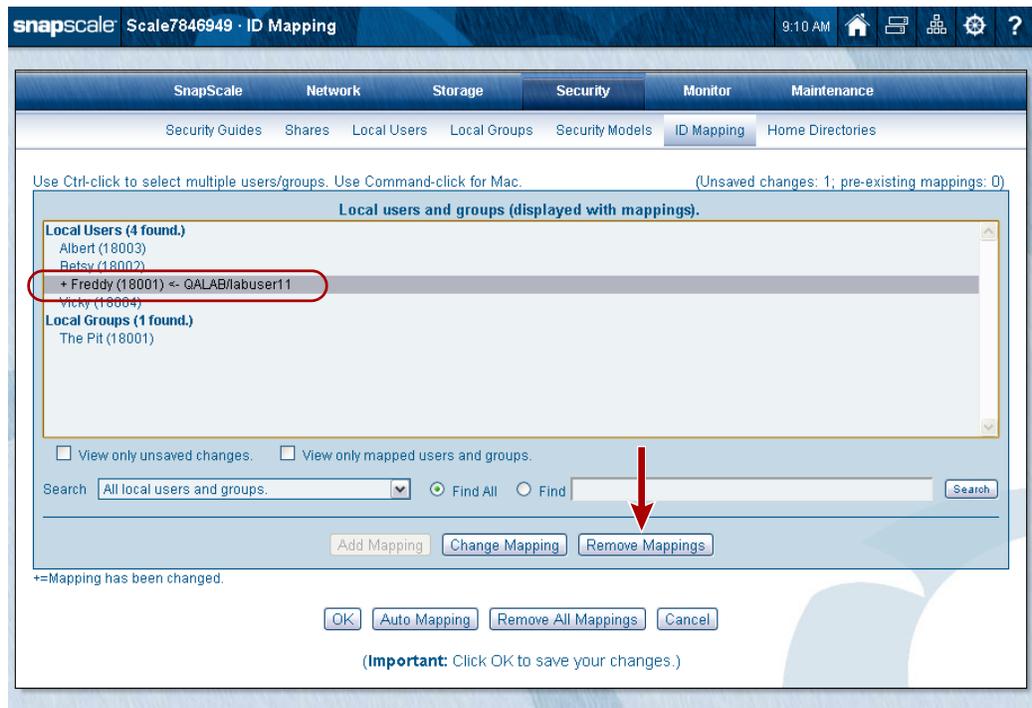
## Remove Mappings

User mappings can be removed individually or all at once. Once removed, they can not be restored but must be added back using [Add Mapping on page 175](#). You also have the option to update the filesystem after removing the ID mappings.

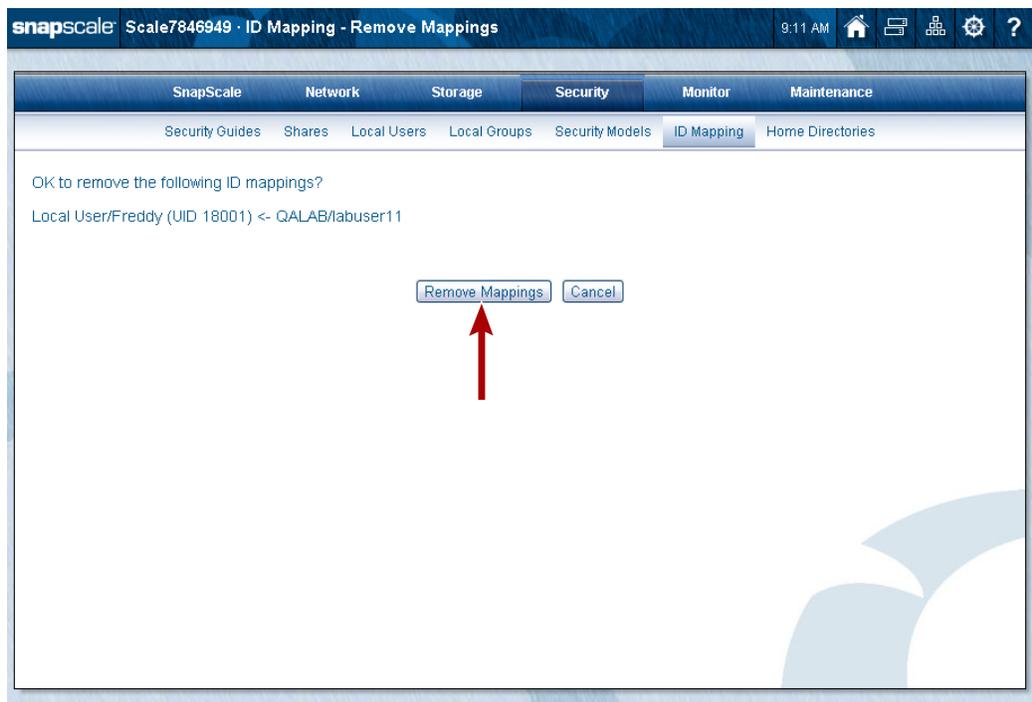
### Remove Individual Mappings

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

2. Click **Remove Mappings**.



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



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

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

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

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

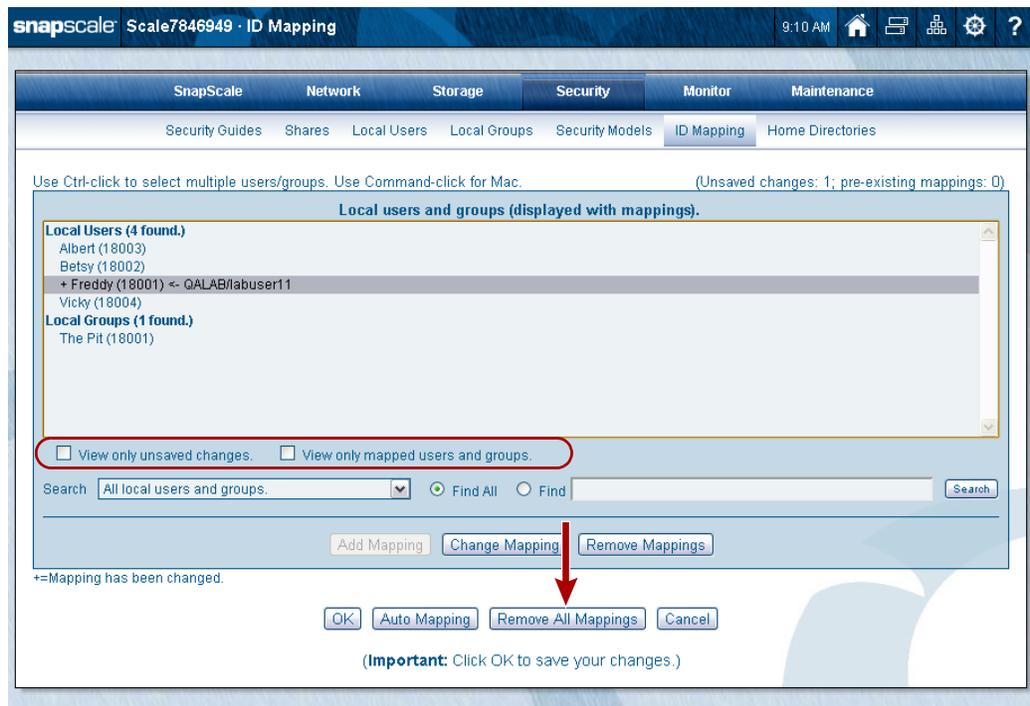
See [Filesystem Updates on page 187](#) for more details.

## Remove All Mappings

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

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

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



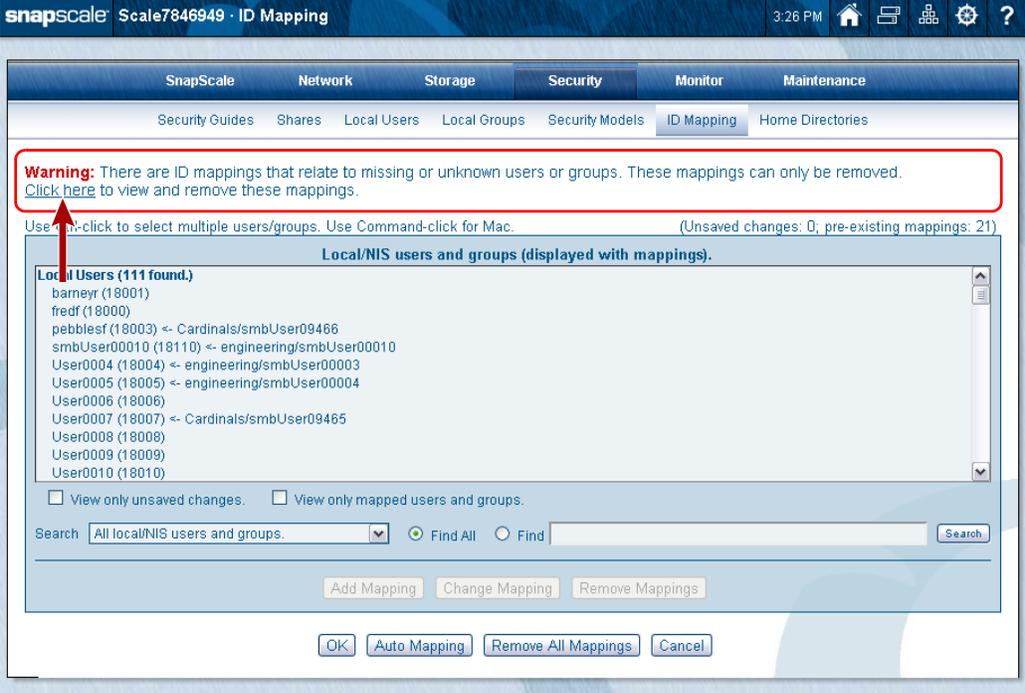
- At the confirmation page, click **Remove Mappings**.  
All the mappings are removed and the default page is displayed with the users/groups that were unmapped in purple with a plus (+) in front of the names.
- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings,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

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

See [Filesystem Updates on page 187](#) for more details.

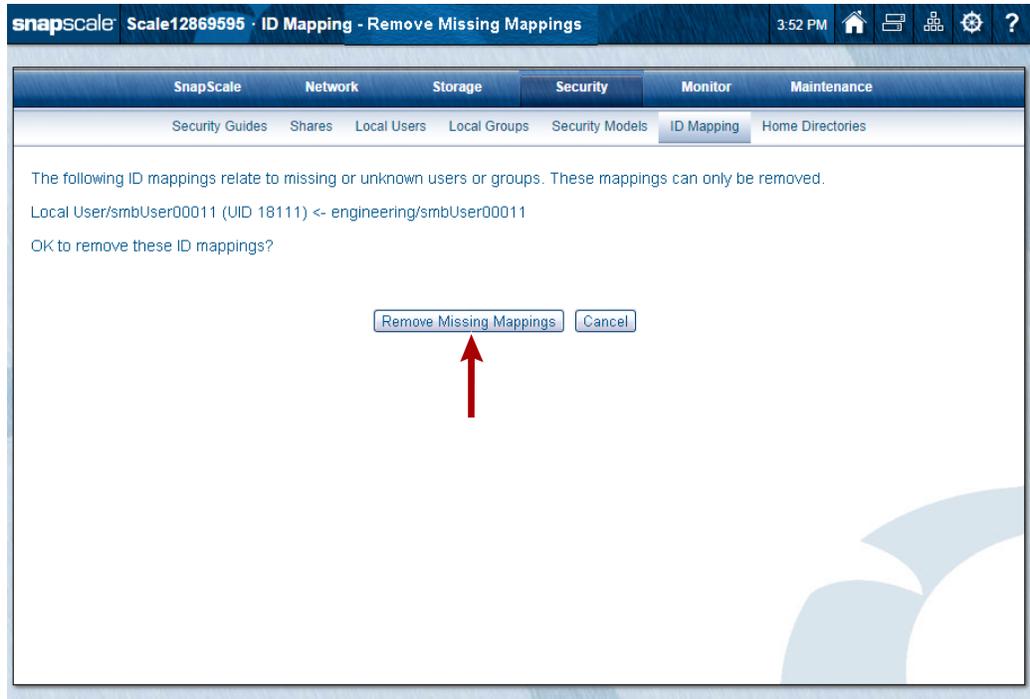
## Remove Missing ID Mappings

If the cluster has mappings for users or groups that no longer exist, the following warning message may be displayed at the top of the main **ID Mappings** page:

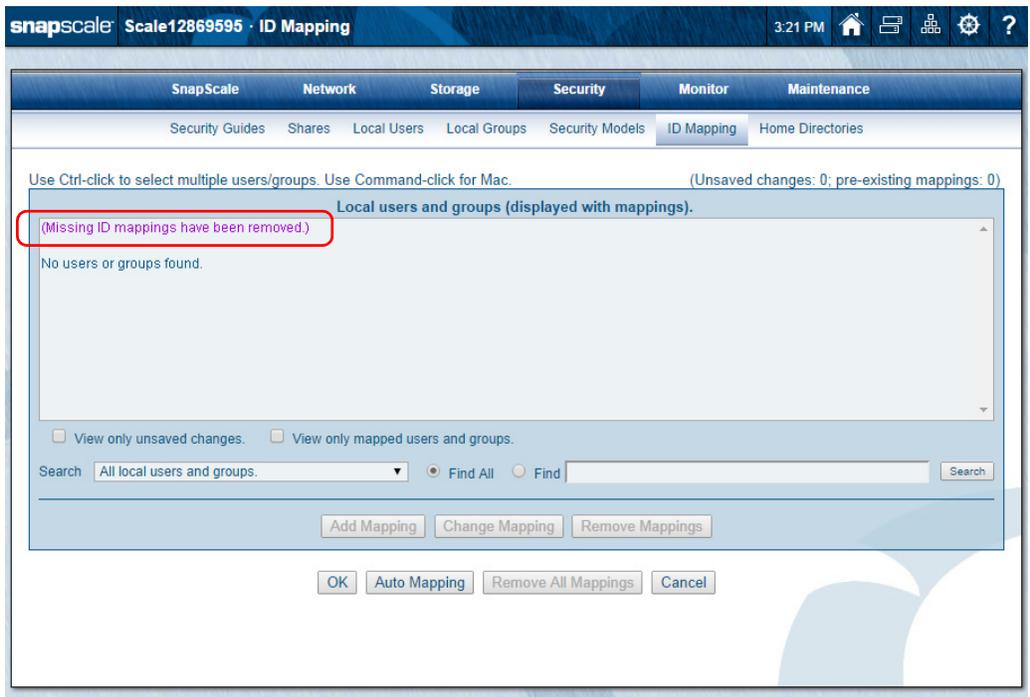


The screenshot shows the SnapScale web interface for ID Mapping. At the top, there is a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below this, there are sub-tabs for Security Guides, Shares, Local Users, Local Groups, Security Models, ID Mapping, and Home Directories. A red box highlights a warning message: "Warning: There are ID mappings that relate to missing or unknown users or groups. These mappings can only be removed. [Click here](#) to view and remove these mappings." Below the warning, there is a search bar and a list of local users and groups. The list is titled "Local/NIS users and groups (displayed with mappings)." and contains the following entries: Local Users (111 found), barneyr (18001), fredf (18000), pebblesf (18003) <- Cardinals/smbUser09466, smbUser00010 (18110) <- engineering/smbUser00010, User0004 (18004) <- engineering/smbUser00003, User0005 (18005) <- engineering/smbUser00004, User0006 (18006), User0007 (18007) <- Cardinals/smbUser09465, User0008 (18008), User0009 (18009), and User0010 (18010). Below the list, there are checkboxes for "View only unsaved changes." and "View only mapped users and groups.", a search bar with the text "All local/NIS users and groups.", and buttons for "Add Mapping", "Change Mapping", and "Remove Mappings". At the bottom, there are buttons for "OK", "Auto Mapping", "Remove All Mappings", and "Cancel".

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



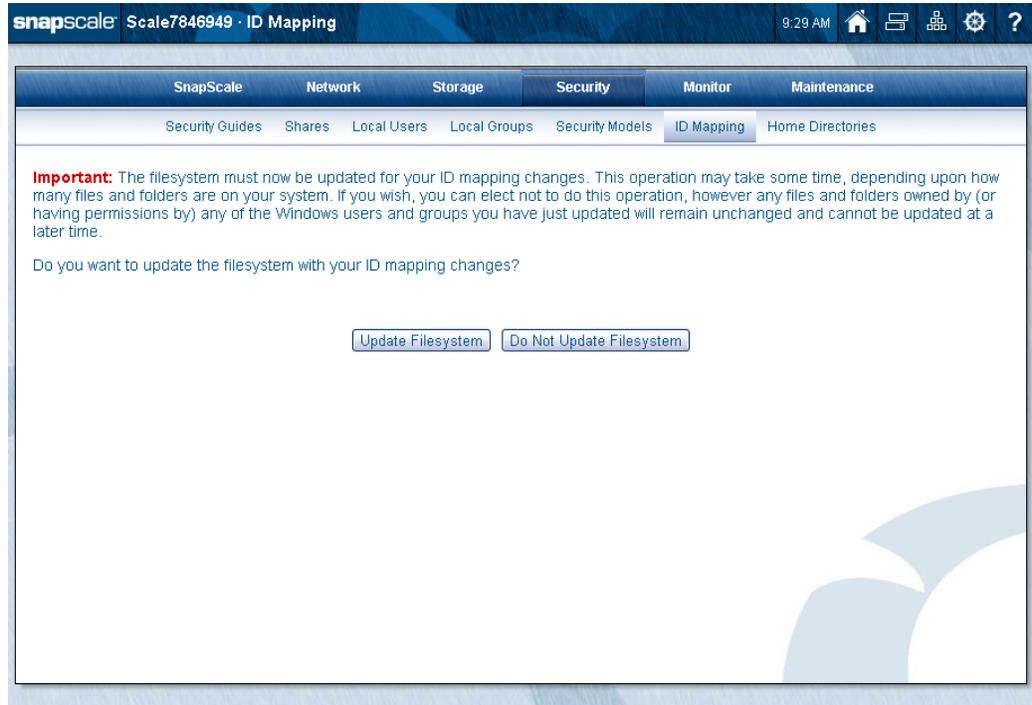
2. Click **Remove Missing Mappings** to clear the missing mappings from the system. A confirmation is shown on the **ID Mapping** main page.



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

## Filesystem Updates

After making any changes to ID mappings, you are presented with a filesystem update option page, where you can choose either **Update Filesystem** or **Do Not Update Filesystem** options.



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

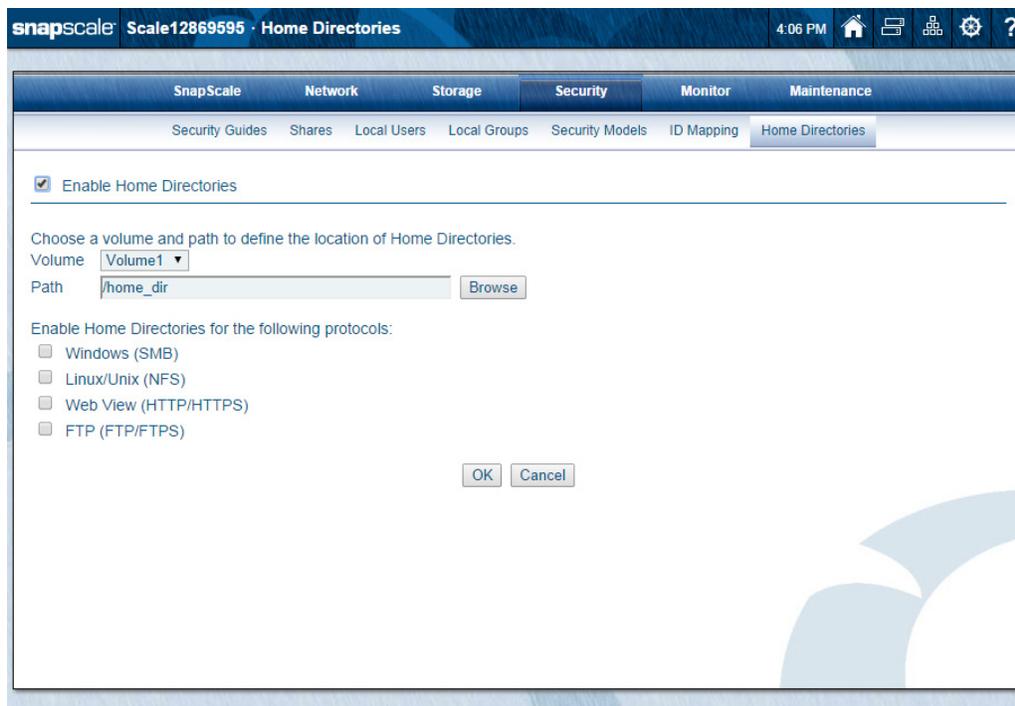
---

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

---

## Home Directories

To enable Home Directories, go to **Security > Home Directories** and check **Enable Home Directories**. Choose the volume, path, and protocols you want.



The Home Directories feature creates a private directory for every local or Windows domain user that accesses the system. When enabling Home Directories (from **Security > Home Directories**), the administrator creates or selects a directory to serve as the home directory root. When a user logs in to the cluster 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, HTTP/HTTPS, and FTP/FTPS. They are accessed by clients in the following manner:

- For SMB 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 cluster continue to be accessible in the usual fashion.
- For NFS, the home directory is exported. When a user mounts the home directory root, all home directories are visible inside the root, but the user's home directory is accessible only by the user and the administrator.

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

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

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

- If a local user home directory accumulates files before the local and domain users are mapped and if the domain user's home directory is empty, the local user's files are copied to the domain user's home directory the first time the local user connects after the users are mapped.
- If both the local and domain user home directories accumulate files before the local and domain users are mapped, the files in the local user's home directory are not copied to the domain user's home directory.

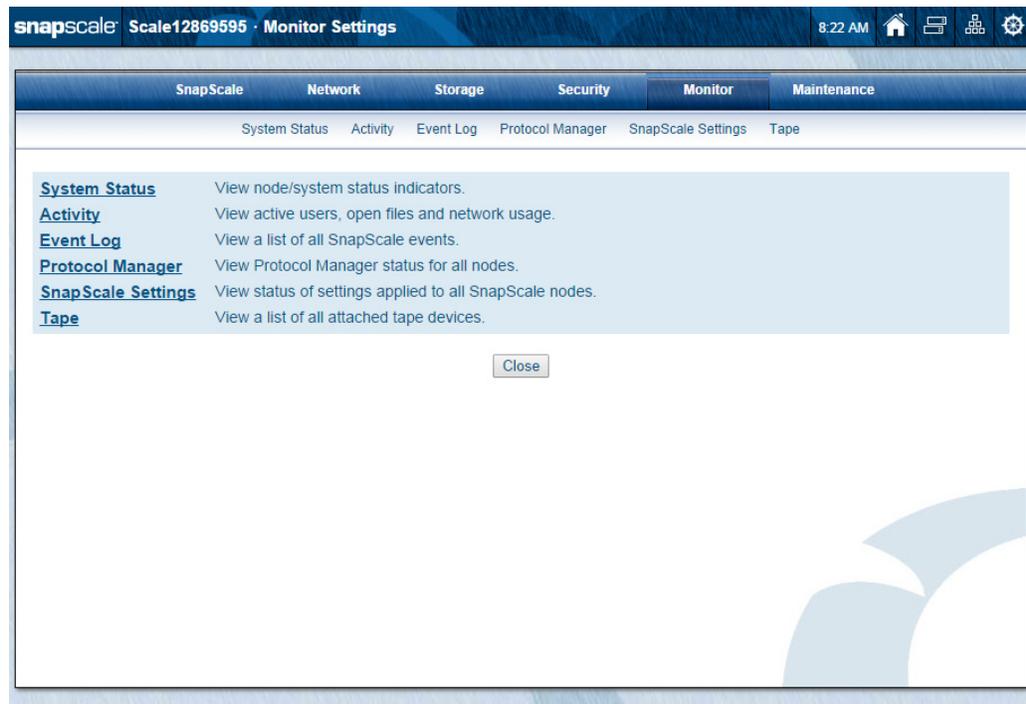
## Configure Home Directories

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

Field	Description
<b>Enable Home Directories</b>	Check to enable Home Directories for local users and activate the options. Remove the check to disable.
<b>Volume</b>	Select the volume where the Home Directories will be located. <b>NOTE:</b> Be sure the volume you select has enough disk space. Once Home Directories are placed, they cannot be moved.
<b>Path</b>	Provide the path to the Home Directories or click <b>Browse</b> to create a new folder. The default path is <code>/home_dir/</code> .
<b>Protocols</b>	Check each of the protocols where Home Directories will be enabled.

**NOTE:** Do not put Home Directories on a volume that might be deleted. If you delete the volume, you will also delete the Home Directories.

This chapter addresses the options for monitoring the SnapScale cluster.



**Topics in System Monitoring:**

- [System Status](#)
- [Activity](#) Submenus:
  - [Active Users](#)
  - [Open Files](#)
  - [Network Monitor](#)
- [Event Log](#)
- [Protocol Manager](#)
- [SnapScale Settings](#)
- [Tape](#)

## System Status

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

The screenshot displays the 'System Status' page in the SnapScale web interface. The page title is 'Scale2302216 · System Status'. The navigation menu includes 'SnapScale', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Under 'Monitor', there are sub-links for 'System Status', 'Activity', 'Event Log', 'Protocol Manager', 'SnapScale Settings', and 'Tape'. The main content area shows two nodes:

- Node13710160 (Mgmt. Node)**: SnapScale X2, RAINcloudOS 4.1.092, Hardware: 04.05.50, Node Number: 13710160, BIOS: r6.0, Serial Number: 56 4d fa 83 2d 2e 54 83, JVM: 1.6.0\_26, Uptime: 1:00:19 (Days:Hours:Minutes), Memory: 1024 MB, CPU: AuthenticAMD Dual-Core AMD Opteron(tm) Processor 2216, Client Network: 1000 Mbps / Full Duplex (Primary), Storage Network: 1000 Mbps / Full Duplex, Ambient Temp.: 35° C / 95° F (Normal), CPU Temp.: 49° C / 120° F (Normal), Power Supply: PS 1: OK, PS 2: OK, Fan Status: 3 Chassis Fans: OK.
- Node5268314**: SnapScale X2, RAINcloudOS 4.1.092, Hardware: 04.05.50, Node Number: 5268314, BIOS: r6.0, Serial Number: 56 4d 95 54 17 e3 12 5c, JVM: 1.6.0\_26, Uptime: 1:00:19 (Days:Hours:Minutes), Memory: 1024 MB.

At the bottom of the node details, there are 'Refresh' and 'Close' buttons.

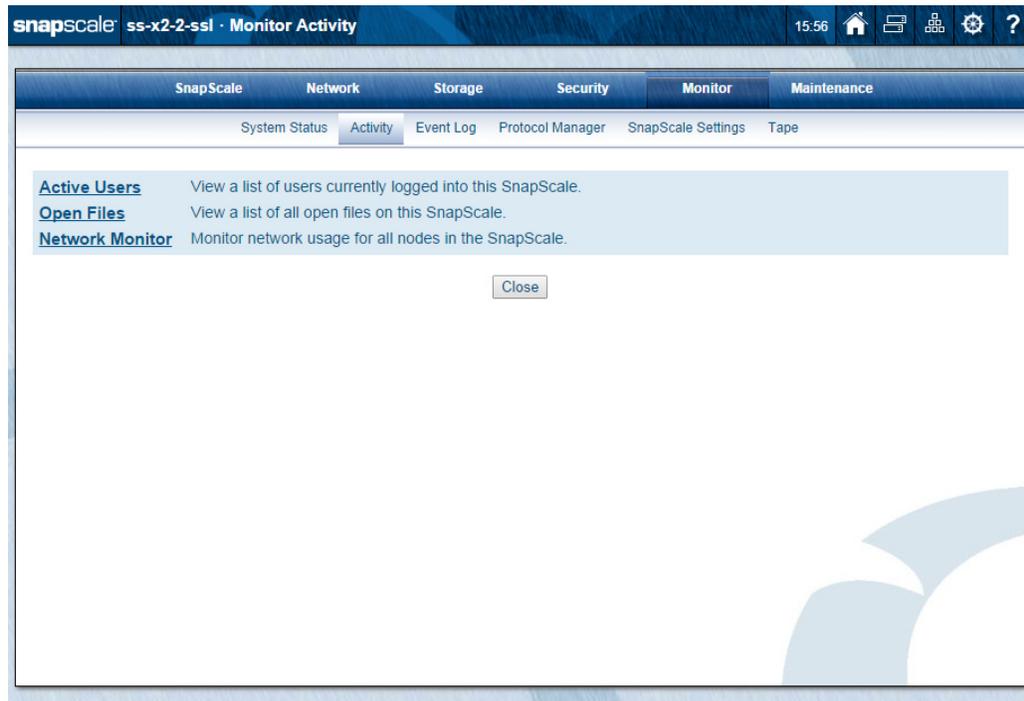
The following status fields are displayed for each node that is part of the SnapScale cluster. Any critical messages are displayed in a **red** font.

Field	Description
<b>Node Name</b>	Name of the node. The default node name is Nodennnnnnn, where <i>nnnnnnn</i> is your node number (for example, Node1123578).
<b>Node Model</b>	Node model name/number.
<b>OS Version</b>	The version of RAINcloudOS currently loaded on the SnapScale node.
<b>Hardware</b>	The node's hardware platform ID.
<b>Node Number</b>	Number derived from the MAC address of the Ethernet 1 port, used as part of the default node name.
<b>BIOS</b>	The BIOS version for the node.
<b>Serial Number</b>	Unique number assigned to the node.
<b>JVM</b>	The Java Virtual Machine version.
<b>Uptime</b>	The amount of time the node has been up (since the last reboot) in "days:hours:minutes" format.
<b>Memory</b>	Amount of system RAM.

Field	Description
CPU	The type of central processing unit (CPU). If more than one CPU exists, each is listed separately.
Client Network	Provides details on the client-side network.
Storage Network	Provides details on the storage-side network.
Ambient Temp.	The temperature of the space inside the chassis.
CPU Temp.	Current CPU temperature.
Power Supply	The status of power supply modules
Fan Status	The status of fan modules.

## Activity

The **Activity** tab provides access to a submenu of options and features for monitoring activity on the cluster.



This submenu is used to access three other pages:

- [Active Users](#)
- [Open Files](#)
- [Network Monitor](#)

## Active Users

This option is used to view read-only details on the active users logged on to each of the nodes on the cluster.

0 active users. 3 nodes online.

Node	User Name	Workstation	Authorized via	Open Files	Protocol	View Node
<input type="checkbox"/> VM-Node12869595	(No active users.)	-	-	-	-	-
<input type="checkbox"/> VM-Node14122451	(No active users.)	-	-	-	-	-
<input type="checkbox"/> VM-Node9715283 (Mgmt. Node)	(No active users.)	-	-	-	-	-

Refresh Close

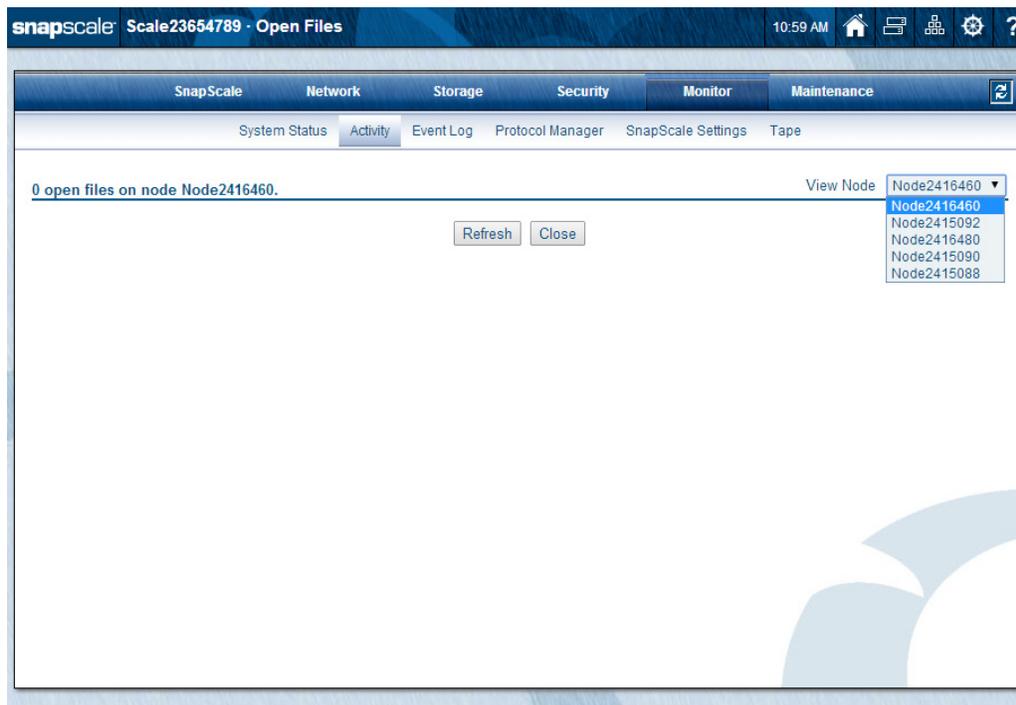
Information available on this page includes user names of all active users, their workstation names, authorization, the number of open files they have on the node, the protocol, and when they logged on. Columns can be sorted in ascending or descending order by clicking the column head.

Use the drop-down list on the upper right to select whether to view all the nodes or individual nodes. Close the page to return to the **Activity** tab.

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

## Open Files

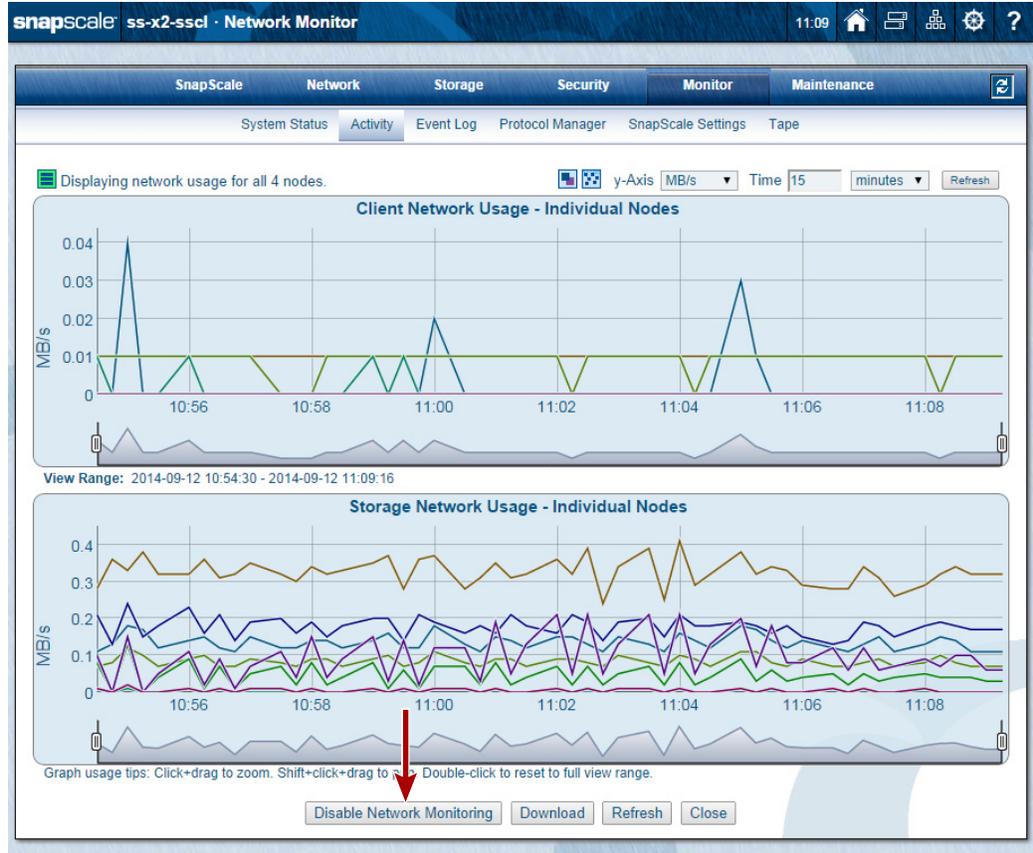
Use this option to view read-only details on the open files on a specific node.



Use the drop-down list on the upper right to choose a different node to view. Click **Close** to return to the **Activity** tab.

## Network Monitor

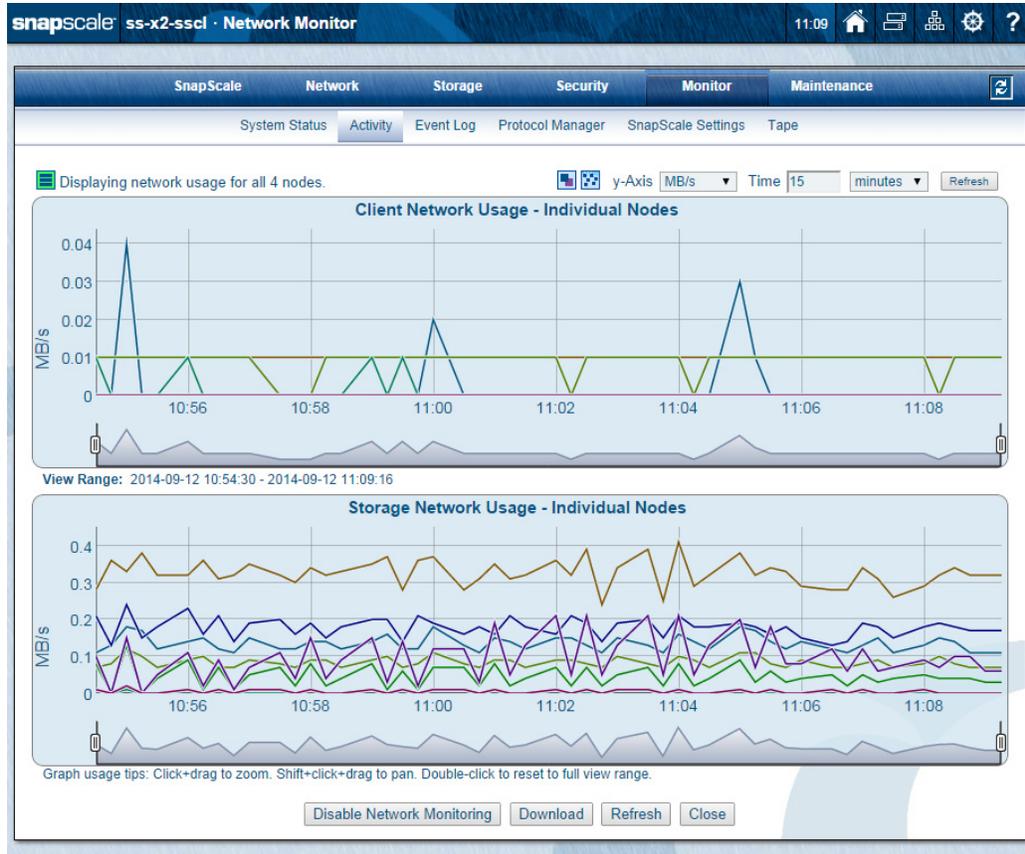
This feature can be used to monitor the network utilization on both the Client and Storage networks. Monitoring is enabled by default. Go to **Monitor > Activity > Network Monitor** and click **Disable Network Monitoring** to turn it off.



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

## View Usage

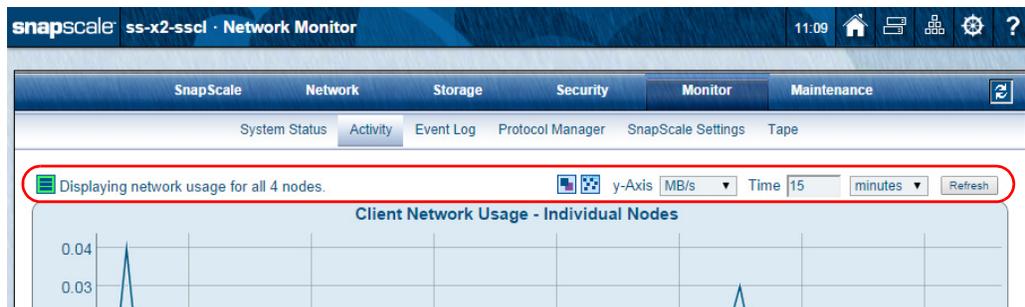
You can go to **Monitor > Activity > Network Monitor** to view graphs showing current usage, total throughput, or combined usage for a user-configurable time period. The data is refreshed every 15 seconds.



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

Icons and options are located above the graphs to modify the views:

- Icons are used to change the information presented in the graphs.
- Units of measure for the Y-axis are selectable (**Percent**, **MB/s**, or **GB/s**).
- Time interval of the data can be adjusted to be shown as a period of time (1 to 999 **minutes**, **hours**, or **days**).



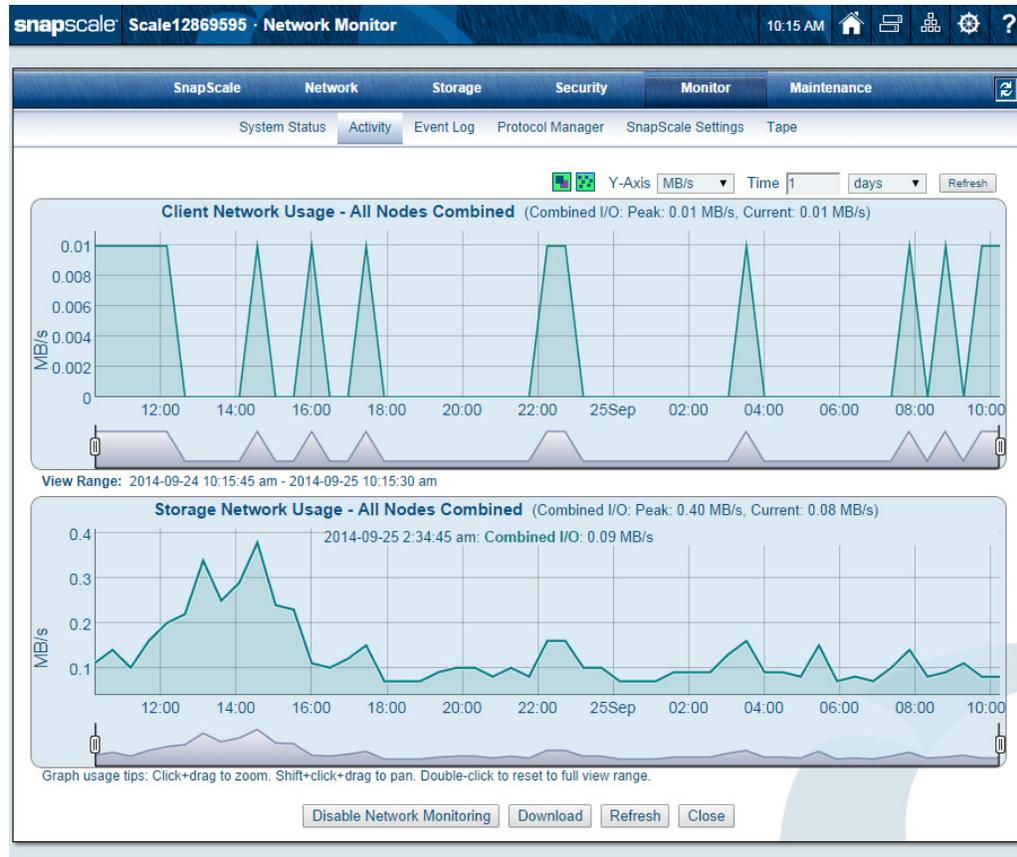
Icons / Options	Description
<b>Select Nodes</b>  Ind All	<p>Click this icon to select the nodes. Choose either individual nodes or <b>Select All Nodes</b>. You can select multiple nodes by using Ctrl+Click.</p> <ul style="list-style-type: none"> <li>When blue, individual nodes are selected.</li> <li>When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

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

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

Different views are displayed by changing the selection of the icons:

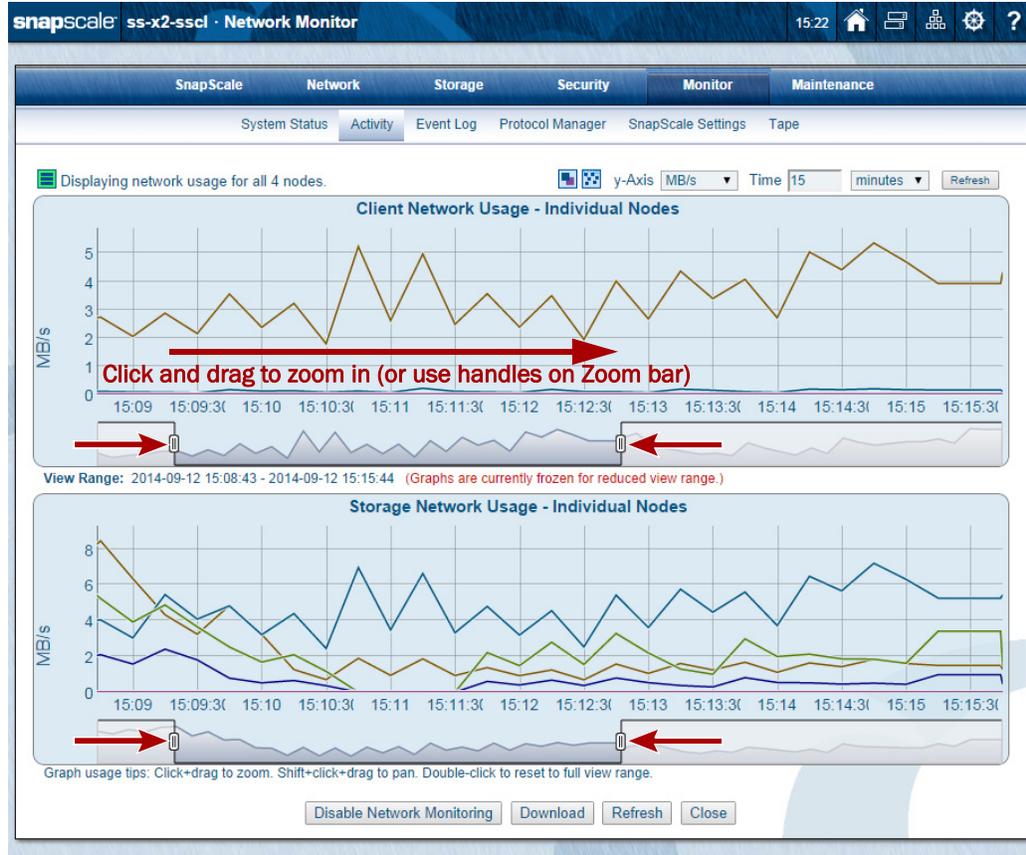
- **Individual activity** – only a single node is selected using Select Nodes.
- **Total activity** – all nodes are selected using Select Nodes.
- **Individual throughput** – a single node (Select Nodes) and the Total Throughput icon are selected.
- **Total throughput** – all nodes (Select Nodes) and the Total Throughput icon are selected.
- **Combined activity** – only the Combined icon is selected.
- **Cluster-wide activity** – both the Total Throughput and Combined icons are selected.



## Graph Options

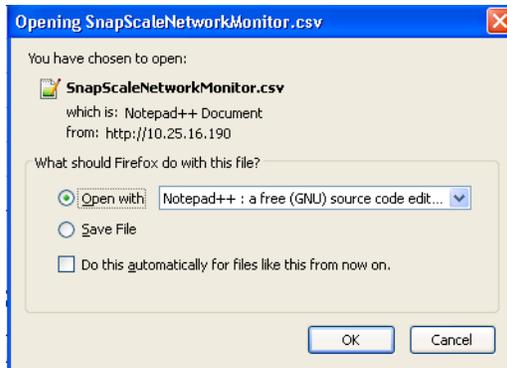
Below each graph is a Zoom Bar that can be used to show a specific time range. When zoomed in, the graphs are frozen and not updated.

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



### Download Usage Records

To download the record displayed as a CSV file, click **Download**. A dialog box asks you to determine the location of the downloaded file.



## Event Log

Use the **Event Log** page to view a log of operations performed on the cluster.

Change the following fields and click Refresh to specify how the log is displayed.

View Log: SnapScale | Severity: Errors, Warnings, and Info | Display Last: 2 Days |  Most Recent First | Refresh

Severity	Time	Message	Source
I	09/23 3:18:13 PM	VM-Node9715283: Security model - set	EventSystem
I	09/23 2:02:47 PM	VM-Node14122451: Password policy - set	EventSystem
I	09/23 2:02:47 PM	VM-Node12869595: Password policy - set	EventSystem
I	09/23 2:02:41 PM	VM-Node9715283: Password policy - set	EventSystem
I	09/23 1:49:19 PM	VM-Node9715283: Accessor recently used - set	EventSystem
I	09/23 1:49:17 PM	VM-Node9715283: Share access set - for 1 accessors	EventSystem
I	09/23 9:55:56 AM	VM-Node12869595: Snapshot schedule SNAP1 - set	EventSystem
I	09/23 9:55:56 AM	VM-Node14122451: Snapshot schedule SNAP1 - set	EventSystem
I	09/23 9:55:56 AM	VM-Node9715283: Snapshot schedule SNAP1 - set	EventSystem
I	09/23 9:55:56 AM	VM-Node12869595: Snapshot space - reclaimed	EventSystem
I	09/23 9:55:53 AM	VM-Node14122451: Snapshot space - reclaimed	EventSystem
I	09/23 9:55:52 AM	VM-Node9715283: Snapshot space - reclaimed	EventSystem
W	09/23 9:55:46 AM	VM-Node9715283: Set cluster to read-write	EventSystem
I	09/23 9:55:46 AM	VM-Node14122451: Snapshot - created SNAP1_0	EventSystem
I	09/23 9:55:46 AM	VM-Node12869595: Snapshot - created SNAP1_0	EventSystem
I	09/23 9:55:43 AM	VM-Node9715283: Snapshot - created SNAP1_0	EventSystem
W	09/23 9:55:39 AM	VM-Node9715283: Set cluster to read-only	EventSystem

Refresh Close

Entries are color coded according to severity as described in the following table:

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

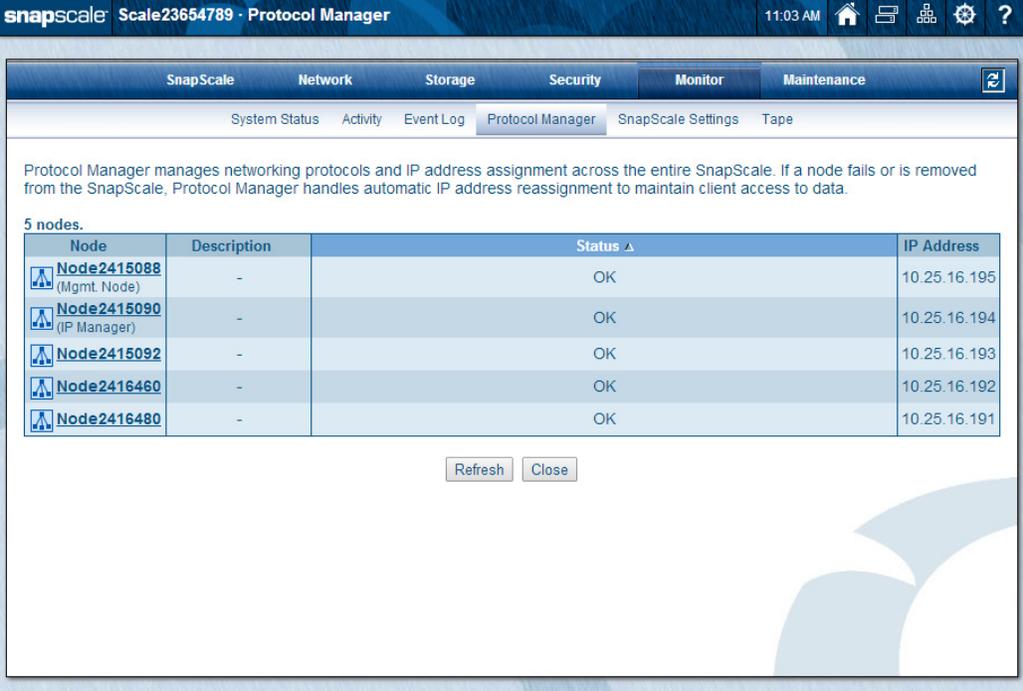
### Filter the Log

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

Option	Description
<b>View Log</b>	Select to view either the SnapScale cluster-wide or node-specific logs. The SnapScale option shows general cluster-related log messages while the node-specific options show log messages specific to the selected node.
<b>Severity</b>	Select the type of alerts and information you want to view.
<b>Display Last <i>n</i> Days</b>	Enter the number of days' worth of entries you want to view.
<b>Most Recent First</b>	Check this box to start the list with the most recent entry; deselect to start the list with the oldest entry.

## Protocol Manager

Protocol Manager manages networking protocols and IP address assignment across the entire SnapScale. If a node fails or is removed from the SnapScale, Protocol Manager handles automatic IP address reassignment to maintain client access to data.



The screenshot shows the SnapScale Protocol Manager interface. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The Monitor tab is active, showing a sub-menu with System Status, Activity, Event Log, Protocol Manager, SnapScale Settings, and Tape. The main content area displays a description of Protocol Manager's function and a table of 5 nodes. Below the table are Refresh and Close buttons.

Node	Description	Status	IP Address
Node2415088 (Mgmt. Node)	-	OK	10.25.16.195
Node2415090 (IP Manager)	-	OK	10.25.16.194
Node2415092	-	OK	10.25.16.193
Node2416460	-	OK	10.25.16.192
Node2416480	-	OK	10.25.16.191

The following table addresses the possible status:

Status	Description
<b>OK</b>	This node is fully functional.
<b>Disconnected</b>	This node could not be connected through the Storage network and is currently not participating in the cluster. If there is a public IP address associated with this node it should have been taken over by a different node. No services are running on this node.
<b>Banned</b>	This node failed too many recovery attempts and has been banned from participating in the cluster temporarily. Any public IP addresses have been taken over by other nodes.
<b>Disabled</b>	This node has been administratively disabled. This node is still functional and participates in the cluster but its IP addresses have been taken over by a different node and no services are currently being hosted.
<b>Unhealthy</b>	A service provided by this node is malfunctioning. The node itself is operational and participates in the cluster, however its public IP addresses have been taken over by a different node and no services are currently being hosted.
<b>Stopped</b>	A node that is stopped does not host any public IP addresses, and does not participate in the cluster.

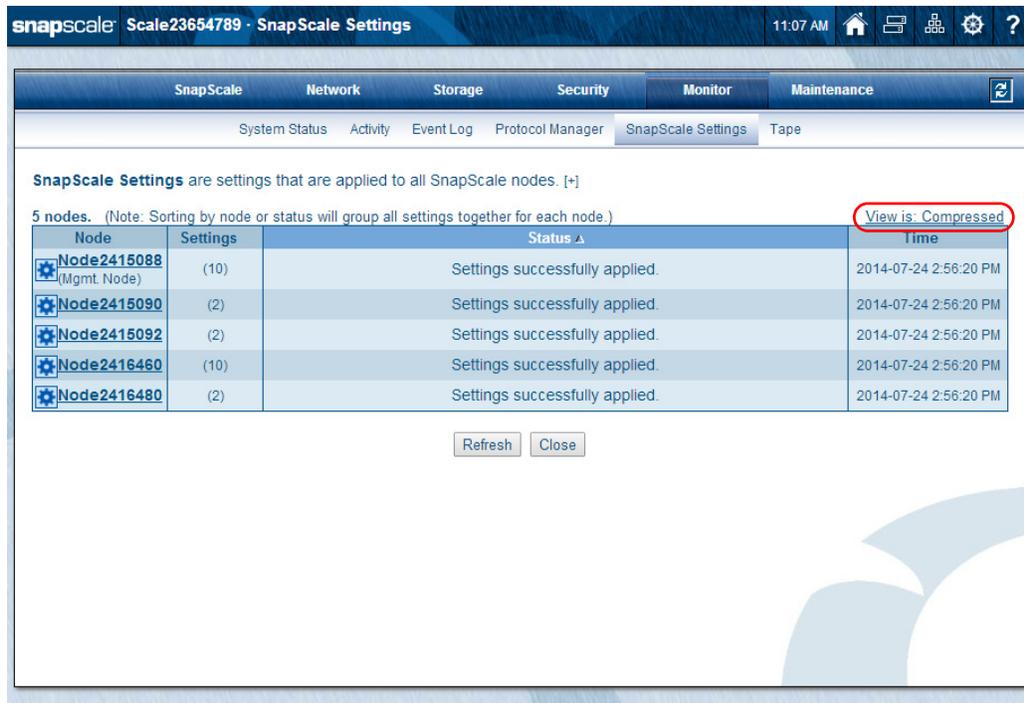
Status	Description
PartiallyOnline	A node that is partially online participates in the cluster like a node that is OK. Some network interfaces which serve public IP addresses are down, but at least one interface is up.

## SnapScale Settings

When cluster settings are configured in the Web Management Interface, success or failure of the operation is determined by the attempt to perform it on the Management node. If successful, the same configuration operation is pushed to all member nodes in the background.

The **SnapScale Settings** page displays a list of settings that have been applied to the nodes in the cluster and the status of each setting. When you make changes to your SnapScale via the Web Management Interface, the settings are applied to the Management node first to determine success or failure of the configuration, then the settings are applied to the other nodes in the background. When a SnapScale setting has not been applied yet, its status is displayed as **Pending**. When a SnapScale setting fails to be applied, its status is displayed in detail and the failed settings are automatically re-applied until they are successful.

The initial view is compressed to show all nodes and a count of the settings:



Click the upper right text that says **View is: Compressed** to switch to the expanded view:

The screenshot shows the SnapScale Settings page. At the top, there are navigation tabs: SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below these are sub-tabs: System Status, Activity, Event Log, Protocol Manager, SnapScale Settings, and Tape. The main content area displays 'SnapScale Settings' and a table of settings for 5 nodes. The table has columns for Node, Settings, Status, and Time. The 'View is: Expanded' text is circled in red.

Node	Settings	Status	Time
Node2415088 (Mgmt. Node)	Profiles	Settings successfully applied.	2014-07-24 2:56:20 PM
	Users & Groups	Settings successfully applied.	2014-07-24 2:53:26 PM
	Third Party Apps	Settings successfully applied.	2014-07-23 4:55:14 PM
	Email	Settings successfully applied.	2014-07-10 3:14:27 PM
	Snap EDR	Settings successfully applied.	2014-06-09 6:20:14 PM
	Server	Settings successfully applied.	2014-06-09 6:18:40 PM
	Share Access	Settings successfully applied.	2014-06-09 6:18:40 PM
	NFS Exports	Settings successfully applied.	2014-06-09 6:18:40 PM
	Shares	Settings successfully applied.	2014-06-09 6:18:39 PM
	Volumes	Settings successfully applied.	2014-06-09 6:18:39 PM
Node2415090	Profiles	Settings successfully applied.	2014-07-24 2:56:20 PM
	Users & Groups	Settings successfully applied.	2014-07-24 2:53:27 PM

Buttons: Refresh, Close

Use the scroll bar on the right side to view all the data. In expanded mode, detailed information on each node is available:

- Each line reports the setting category, status, and date/time the setting was applied.
- If an operation is still in progress on a node, the **Status** will be set to **Pending** with a yellow background.
- If an operation failed on a node, the **Status** will have an error message and a red background.
- Column headers can be clicked to sort by **Node**, **Settings**, **Status**, or **Time**.
- When sorting by **Status** or **Node**, settings are grouped by node.

Click the **View is: Expanded** text to revert back to the compressed view. The displayed view from that point on will be the last view selected. Clicking the column heading resorts the table on that function.

## Tape

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

The screenshot shows the SnapScale Tape Monitor interface. At the top, there's a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below this is a sub-navigation bar with System Status, Activity, Event Log, Protocol Manager, SnapScale Settings, and Tape. The main content area displays '0 tape devices. 5 nodes online.' and a table with the following data:

Node	Device Model	Device Type	Device name	CX	Bus	ID	LUN
Node2415088 (Mgmt. Node)	No tape devices attached. (SCSI card detected.)	-	-	-	-	-	-
Node2415090	No tape devices attached. (SCSI card detected.)	-	-	-	-	-	-
Node2415092	No tape devices attached. (SCSI card detected.)	-	-	-	-	-	-
Node2416460	No tape devices attached. (SCSI card detected.)	-	-	-	-	-	-
Node2416480	No tape devices attached. (SCSI card detected.)	-	-	-	-	-	-

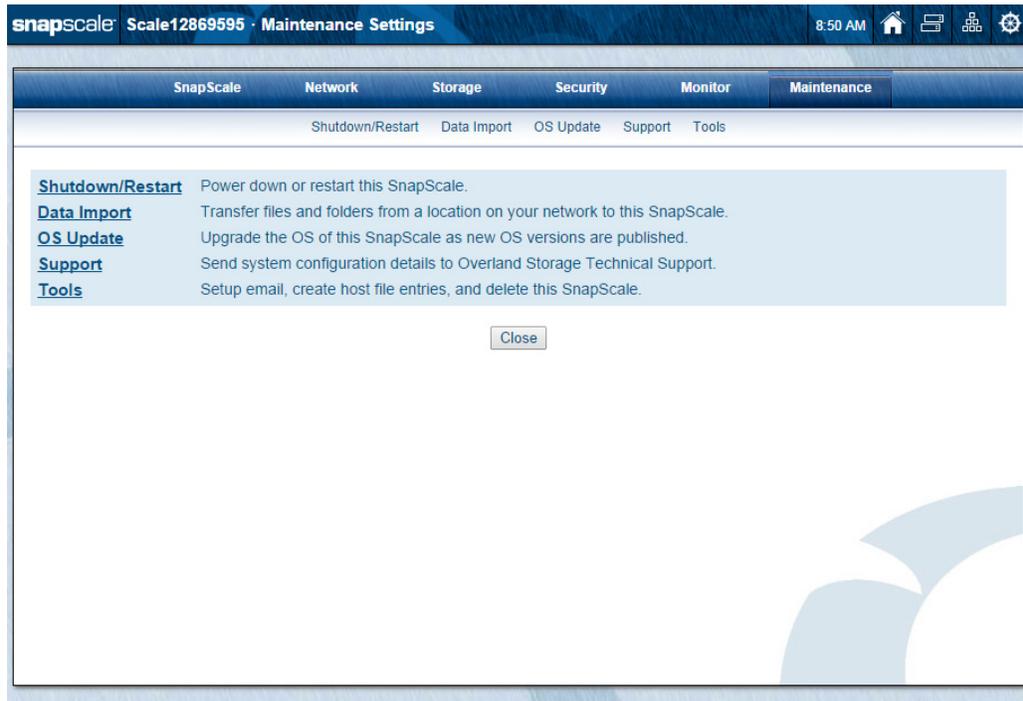
Below the table are 'Refresh' and 'Close' buttons. A 'View Node' dropdown menu is set to 'View all nodes.'.

Use the drop-down list on the upper right to select whether to view all the nodes or individual nodes. Close the page to return to the **Monitor** page.

The following table describes the fields:

Field	Description
<b>Device Model</b>	The manufacturer's model for the device.
<b>Device Type</b>	Type of tape device: either Sequential-Access (tape drive) or Medium-Changer (for example, robotic arm for a tape library).
<b>Device Name</b>	Name of the device node to which the device is bound.
<b>Connection</b>	Identifies the connection type: SCSI or USB.
<b>Bus</b>	Bus number indicating which physical interface (for example, SCSI card) the device is connected to.
<b>ID</b>	ID number (SCSI only)
<b>LUN</b>	LUN identifier (SCSI only)

Clicking the **Maintenance** tab on the Web Management Interface displays options used to maintain this SnapScale cluster. There is also a **Tools** submenu of special, related options.

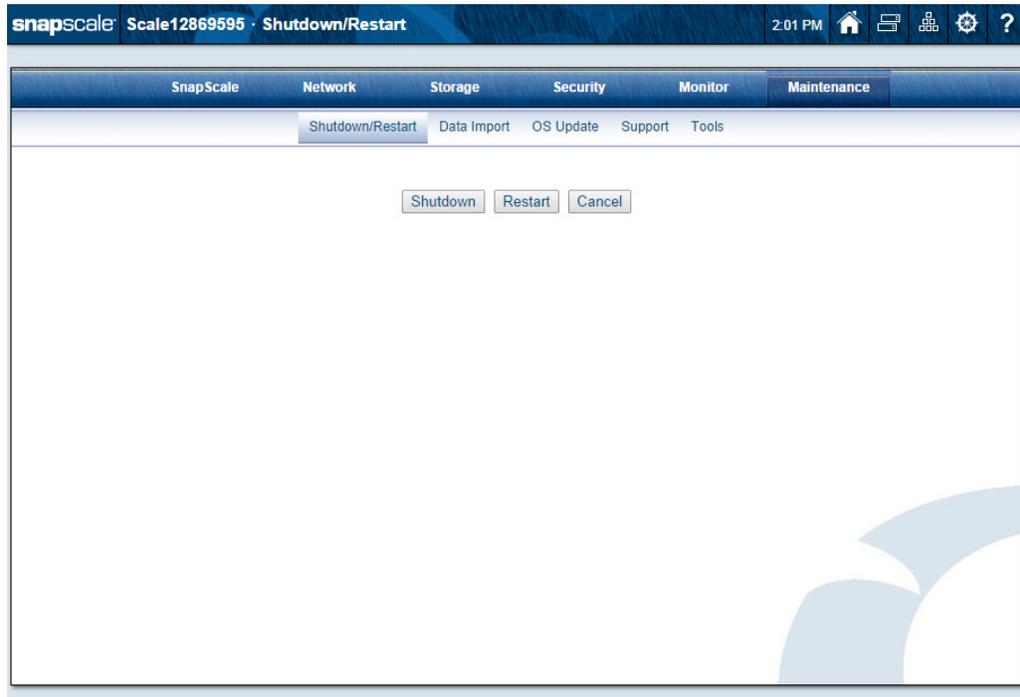


## Topics in Maintenance

- [Shutdown and Restart](#)
- [Data Import](#)
- [OS Update](#)
- [Support](#)
- [Maintenance Tools:](#)
  - [Email Notification](#)
  - [Host File Editor](#)
  - [Delete SnapScale Cluster](#)

## Shutdown and Restart

Use the **Shutdown/Restart** page to reboot or shut down the cluster.



Click one of the following buttons:

- **Shutdown** – Shuts down and powers off all nodes in the cluster.
- **Restart** – Reboots the cluster via a controlled shutdown and restart.

### Manually Powering SnapScale Nodes On and Off

---

**CAUTION:** To prevent possible data corruption or loss, it is NOT recommended to directly power down any nodes that are part of a SnapScale cluster. When powering down a cluster, always use the **Shutdown** button that can be found under **Maintenance > Shutdown/Restart** in the Web Management Interface.

---

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

- To turn the node on, press the Power button on the front of the node.  
The node takes a few minutes to initialize. A green system/status LED indicates that the system is up and running.
- To turn the node off, press and release the Power button to begin the shutdown process. Do not depress this button for more than four seconds.

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

## Data Import

Use the **Data Import** page (**Maintenance > Data Import**) to import (migrate) data from another SnapScale cluster, SnapServer, or other computer that supports CIFS or NFS (v2, v3, or v4) to this cluster.

### Windows/SMB Page:

The screenshot shows the SnapScale Data Import page. The top navigation bar includes 'SnapScale', 'Scale12869595', and 'Data Import'. The main navigation tabs are 'SnapScale', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. The 'Data Import' sub-tab is active, with other options like 'Shutdown/Restart', 'OS Update', 'Support', and 'Tools' visible. The page content includes a description: 'Use Data Import to copy or move files and folders from a location on your network (Source) to this SnapScale (Target)'. The 'Source' section has a 'Network Protocol' dropdown set to 'Windows (SMB)', with a note '(specifies how to communicate with host)'. Below this are input fields for 'Auth. Name', 'Auth. Password', 'Host', 'Share', and 'Path', each with a 'Browse' button. The 'Target (This SnapScale)' section has a 'Volume' dropdown set to 'Volume1' and a 'Path' input field with a 'Browse' button. The 'Options' section includes an 'Import Type' dropdown set to 'Copy (source data is maintained)', and three checked checkboxes: 'Include all sub-folders (if source path specifies a folder)', 'Overwrite existing target files and folders (that have identical names as the source files and folders)', and 'Verify imported data (takes twice as long)'. There are also two unchecked options: 'Preserve file/folder permissions' and 'Verify imported data (takes twice as long)'. A note at the bottom states: 'Note: You can setup [Email Notification](#) (administrative operation event) to be notified when a Data Import operation is complete.' At the bottom right, there are three buttons: 'Start Import', 'View Log', and 'Close'.

**NFS Page:**

The screenshot shows the SnapScale Data Import configuration page. At the top, the breadcrumb navigation includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Data Import' sub-tab is active. The main heading reads: 'Use Data Import to copy or move files and folders from a location on your network (Source) to this SnapScale (Target).'

**Source:**

- Network Protocol: **NFS** (specified how to communicate with host)
- User Name: [ ] (Snap local, NIS or LDAP user)
- Host: [ ]
- Export: [ ]
- Path: [ ]

**Target (This SnapScale):**

- Volume: **Volume1**
- Path: [ ]

**Options:**

- Import Type: **Copy (source data is maintained)**
- Include all sub-folders (if source path specifies a folder)
- Overwrite existing target files and folders (that have identical names as the source files and folders)
- Preserve file/folder permissions
- Verify imported data (takes twice as long)

**Note:** You can setup [Email Notification](#) (administrative operation event) to be notified when a Data Import operation is complete.

Buttons at the bottom:

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

The following import options can be specified:

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

**NOTE:** If you elect to preserve original permissions settings, review [Preserving Permissions on page 212](#).

- Verify imported data

**NOTE:** If you elect to verify imported data, all data is read twice, once for import and once for comparison to the copied data. This could be a lengthy process.

## Setting Up a Data Import Job

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

- For Windows import, specify an administrator or member of the Windows server/domain administrators group.

- For NFS v2/v3 import, consider using the user root and configuring the NFS export on the source to `no_root_squash` for the IP Address of the node for the duration of the import.

**NOTE:** Only one import job can run at a time.

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

- On the **Data Import** page, complete the required **information** for both the source and target.

Option	Description
<i>Source:</i>	
<b>Network Protocol</b>	<p>Protocol that the cluster uses to connect to the source server. Use the drop-down list to select:</p> <p><b>NOTE:</b> If you are importing via SMB, SMB must also be enabled on the target cluster (enable at <a href="#">Network &gt; Windows/SMB</a>).</p> <ul style="list-style-type: none"> <li><b>Windows (SMB)</b> – Uses SMB for Windows with the source data on a Windows root directory (default option).</li> <li><b>NFS</b> – Uses NFS v2/v3 for Unix/Linux-based servers or RAINcloudOS nodes with source data on a Unix root directory.</li> </ul>
<b>Auth. Name &amp; Auth. Password / User Name</b>	<ul style="list-style-type: none"> <li>For the <b>Windows (SMB)</b> network protocol, enter both the <b>Auth. Name</b> and <b>Auth. Password</b> (Windows user name and password to log in to the source server over SMB).</li> <li>For the <b>NFS</b> network protocol, enter the <b>User Name</b> (node local user name or NIS user, representing the UID used to perform the operation over NFS).</li> </ul>
<b>Host</b>	Enter the name or IP address of the source server you are importing data from.
<b>Share/Export</b>	<p>Specify the share (Windows) or export (NFS) on the source server containing the data you want to import.</p> <p><b>NOTE:</b> Wildcards are not supported when specifying the source share to import.</p>
<b>Path</b>	<p>Enter the path to the file or folder you want to import. If you are importing the entire share, you can leave the Path field blank.</p> <p><b>NOTE:</b> Wildcards are not supported when specifying the path to import.</p>
<i>Target (This SnapScale):</i>	
<b>Volume</b>	Specify the volume where you want the data imported.
<b>Path</b>	Specify the path to the directory where you want the data imported.
<i>Options:</i>	
<b>Import Type</b>	<p>Options for the import data are to <b>Copy</b> (source data is maintained) or <b>Move</b> (source data is removed during copy). If <b>Verify Imported Data</b> is enabled, the <b>Move</b> option removes the original data after the verification is complete.</p> <p>The default is <b>Copy</b>.</p> <p><b>NOTE:</b> If you select to <b>Move</b> rather than <b>Copy</b> data, it is strongly recommended that you also select the <b>Verify Imported Data</b> option.</p>

Option	Description
<b>Include All Sub-folders</b>	If the folder you select for import contains subfolders, selecting this option imports all files and folders underneath this folder (default is checked).  <b>NOTE:</b> If disabled, <i>only the files in this folder are imported</i> .
<b>Overwrite Existing Target Files &amp; Folders</b>	If any files/folders on the target have identical names with files/folders on the source, checking this option overwrites those files/folders during import (default is checked).
<b>Preserve File/Folder Permissions</b>	Selecting this option retains the source permissions when the files/folders are imported to the target (unchecked by default).  <b>NOTE:</b> Before selecting this option, review <a href="#">Preserving Permissions on page 212</a> .
<b>Verify Imported Data</b>	Selecting this option causes all source data to be read twice, once to write to the target and once to perform a binary comparison with the data written (default is unchecked).  If enabled and, if the <b>Import Type</b> is <b>Move</b> , files on the source are only removed after verification. Otherwise, files are removed during the process of copying them to the target. If you select to move the files rather than copy them, it is strongly recommended that you enable the <b>Verify Imported Data</b> option.  If a file mismatch occurs during verification, the target file is moved to a <code>data_import_verify_failures</code> directory on the root of the same volume. Check the failed file to determine the problem, then run the import again with <b>Overwrite Existing Target Files &amp; Folders</b> deselected (so you do not re-copy files that have already been copied and verified).  <b>NOTE:</b> Depending upon how much data is being imported, verifying imported data can be a lengthy process.
<b>Email Notification</b>	Clicking the email notification link takes you to the <b>Email Notification</b> page (for more information, see <a href="#">Email Notification on page 226</a> ).  Fill in notification information and check the box next to <b>Administrative Operation Event</b> in order to receive an email when the import operation is complete.

- Once you have completed the import information, click **Start Import** to begin the import. You can see the progress of the import, the estimated time until completion, and the import log on the secondary **Data Import** page.
- When the import is complete, click **View Log** to see details of all errors. Click the **Data Import Error Log** link to download the entire log.

## Stopping an Import Job

To stop the import at any time, click **Stop Import** on the **Data Import** secondary page. If a file was in the process of being copied, the partially-copied file on the target is removed.

## Recreating an Import Job

The **Data Import** log records all errors that occurred during import. You can import just the files and folders that were not imported during the original job due to an error condition (for example, the file was locked).

- Review the **Data Import errors log** and correct all error conditions (such as unlocking a locked file).

2. Reopen the **Data Import** page. All fields (except the password) from the last import will still be visible on the page.

By default, all files will be re-imported. If you want only to import those files that failed to import the first time, you can disable the **Overwrite Existing Target Files** option.

However, make sure that all problematic files from the first import are deleted from the target so they can be re-imported.

**NOTE:** If an import failed, it is strongly recommended that you enable the [Verify imported data option for the re-importation](#).

3. Enter your password and click **Start Import** to run the import again.

## Preserving Permissions

The types of permissions retained will differ, depending on which import scenario is applied.

### Importing from a Windows Security Model to a Windows Personality Directory

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

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

### Importing from a Unix Security Model to a Unix Personality directory

If you are importing from a Unix server to a Unix personality directory, Unix permissions for UIDs/GIDs are copied exactly from source to target; thus, identities of the users and groups are best retained if the SnapScale cluster belongs to the same NIS domain as the Unix server.

### Importing Between Conflicting Security Models

When importing from a Unix source to a Windows security model target, Unix permissions are retained and the security personality on the resulting files and directories will be Unix.

However, when importing from a Windows source to a Unix security model target, permissions cannot be retained (since Unix root directories are required to be Unix personality throughout). Files and directories will inherit the Unix personality and will have a set of default Unix permissions.

### Importing from a SnapServer or SnapScale Cluster

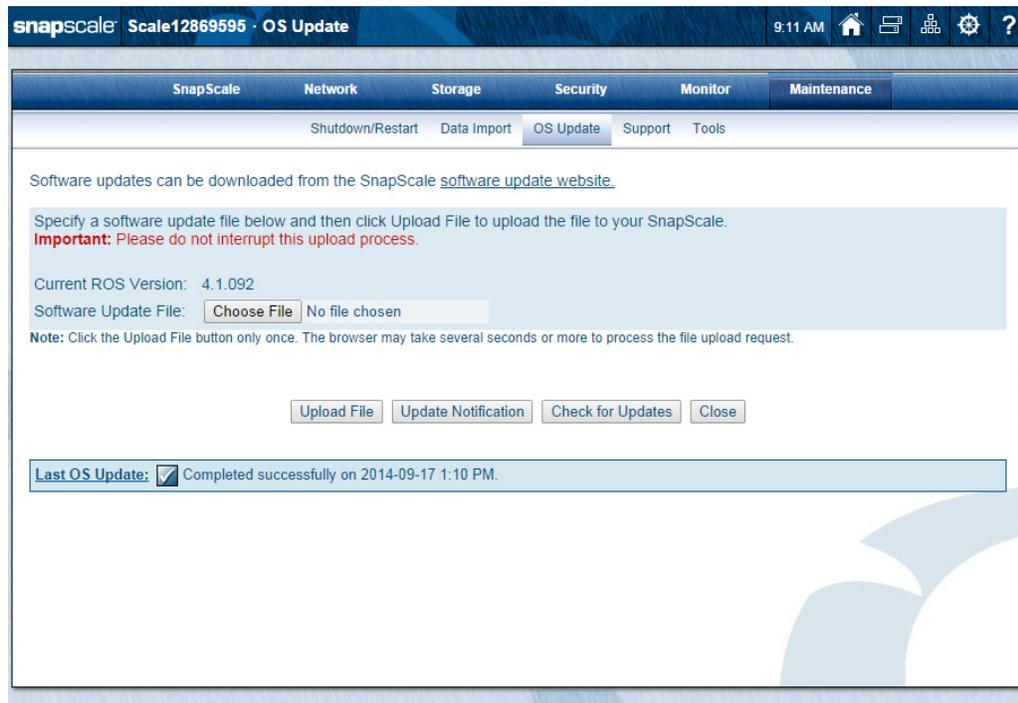
When importing from a different SnapServer or SnapScale cluster, it is recommended that you maintain the same security model on the target that you have on the source.

- If your source uses a Windows security model and has permissions assigned to Windows domain users, use a Windows (SMB) connection for import. Windows permissions are retained exactly as they are on the source, with the same enforcement limitations for unknown users as for importing from Windows servers (see [Importing from a Windows Security Model to a Windows Personality Directory on page 212](#)).
- If your source server or cluster uses a Unix security model and has permissions assigned to local or NIS users, use an NFS connection for import.

**NOTE:** Local users who have Unix permissions on the source are not created on the target with the same UIDs.

## OS Update

Use this page to install updates to RAINcloudOS and other installed software, and to configure your system to automatically check for updates.



Information about the last RAINcloudOS update is listed at the bottom of the page and shows the basic information about the update.



**CAUTION:** Do not interrupt the update process. You may severely damage the cluster if you interrupt a software update operation before it is complete.

**NOTE:** While interrupting an update is not recommended, the **Abort** button provided on the Online/Rolling Update status page can be used to stop the update should there be no alternative. After an abort of a Online/Rolling Update, the user must then complete the update using the Offline method (all nodes restarting).

### Update the RAINcloudOS



**IMPORTANT:** It is highly recommended that all active iSCSI users be disconnected before continuing.

**NOTE:** Snap EDR cannot be installed on a SnapScale cluster using the **OS Update** page (see [Snap EDR on page 239](#)).

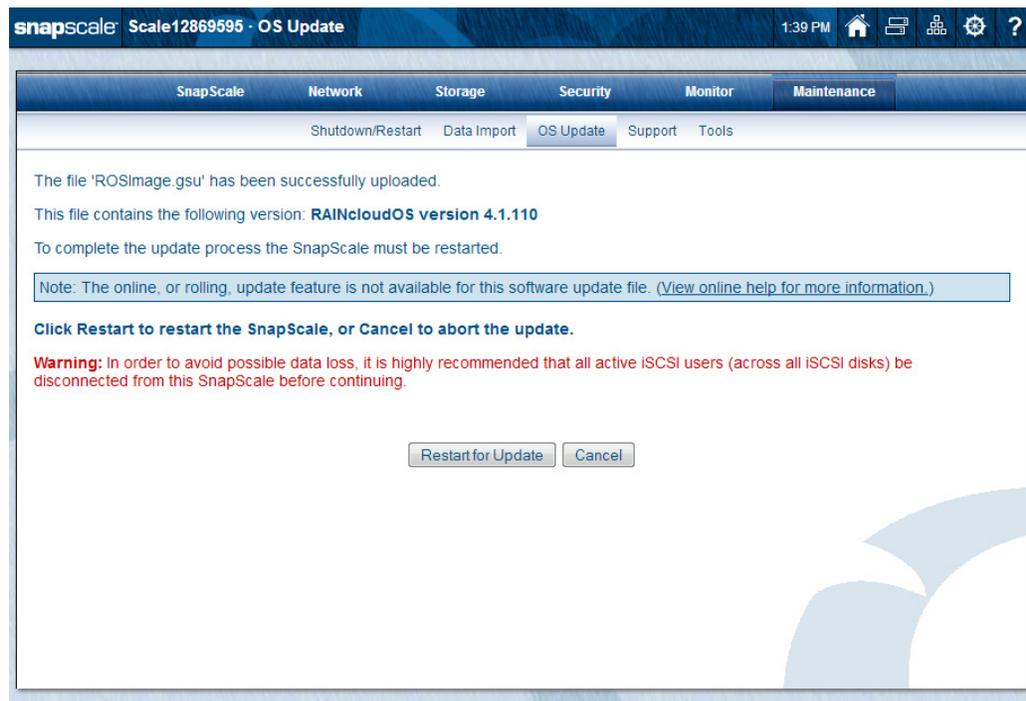
#### 1. Click **Check for Updates**.

If an update is available, follow the instructions on the page to download it.

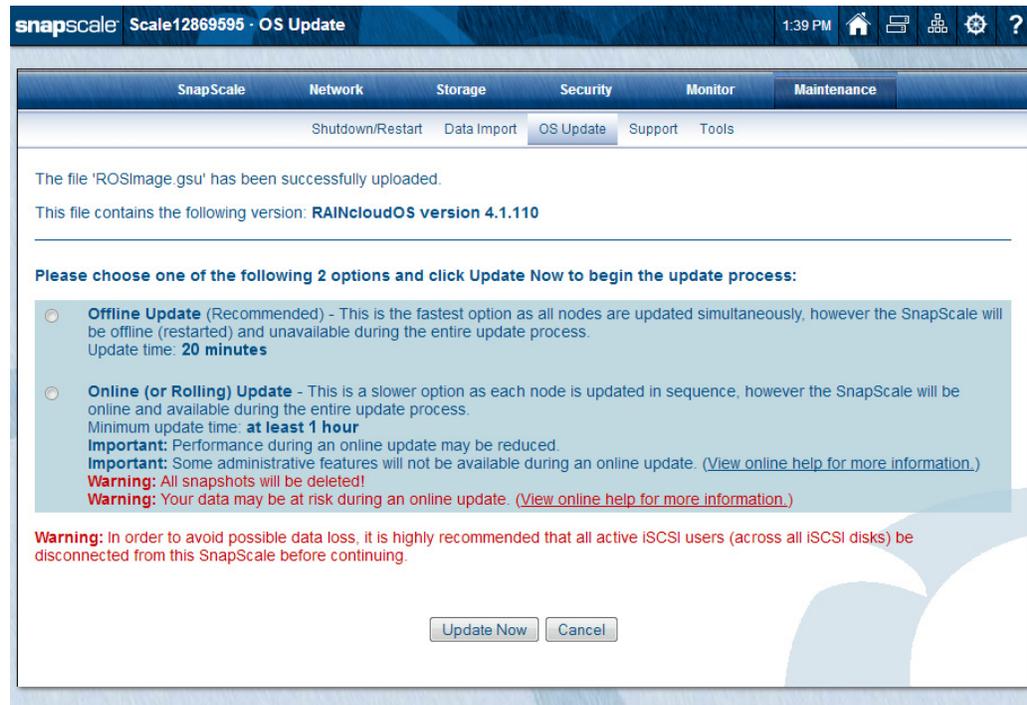
**NOTE:** If the cluster does not have access to the Internet, download the latest RAINcloudOS image (.gsu) or other software package from the [Overland Storage website](#) to a computer on the same network that the cluster can access.

2. On the **OS Update** page, click **Browse** (or **Choose File**, depending upon your browser), locate the file to be uploaded, and select it.
3. Click **Upload File** to start the upload to the cluster.  
Only click the button once. Some browsers show the percent of the upload progress in their bottom status bar. The SnapScale cluster uploads the software package. An update option page is shown.

Under some circumstances, an online/rolling upgrade is not available. You are given the opportunity to either cancel or use the offline upgrade option:



- Choose the update option you want to use and click **Update Now**.



- **Offline Update (Recommended)** – This is the fastest option as all nodes are updated simultaneously. The SnapScale will be offline (restarted) and unavailable during the entire update process.
- **Online (or Rolling) Update** – This is a slower option as each node is updated in sequence. The SnapScale will still be online and available during the entire update process. However, performance will be reduced and some administrative features will not be available. See [Online/Rolling Updates Notes on page 216](#) for caveats.



**CAUTION:** With an Online/Rolling Update, all snapshots are deleted and your data has reduced redundancy during the update.

After an upgrade and reboot, the **OS Update Status** page displays the success or failure of the last update performed for each node.

OS Update Status

Operation status:  Completed successfully on 2014-09-17 1:10 PM.

Operation type: Offline Update

Nodes updated: 3 of 3

3 Nodes: All nodes were updated successfully.

Node	Status	Product	Version	Completion Time
VM-Node12869595	<input checked="" type="checkbox"/> Update complete.	RAINcloudOS Full Upgrade	4.1.092	2014-09-17 1:09 PM
VM-Node14122451	<input checked="" type="checkbox"/> Update complete.	RAINcloudOS Full Upgrade	4.1.092	2014-09-17 1:10 PM
VM-Node9715283 (Mgmt. Node)	<input checked="" type="checkbox"/> Update complete.	RAINcloudOS Full Upgrade	4.1.092	2014-09-17 1:07 PM

Close

## Online/Rolling Updates Notes

The Online/Rolling Update option allows you to update the RAINcloudOS while leaving the SnapScale cluster online and active. Each node is upgraded and rebooted one at a time, and as a result, degrades the peer sets with members in that node during that process. After upgrade/reboot, each node rejoins the cluster and all peer sets are fully rebuilt before the rolling update proceeds to the next node. This repeats until all nodes have been upgraded (or a failure is encountered).

To prevent problems and data loss, consider the following caveats:

- To avoid possible data loss, it is highly recommended that all active iSCSI users (across all iSCSI disks) be disconnected from this SnapScale cluster before updating the OS.
- If upload fails on one or more nodes, the upgrade will abort with an error and list the nodes that had problems. In such a case, an Offline Update is required.
- If you abort the Online/Rolling Update during the update process, the process stops where it is resulting in a mixed-OS state which needs to be resolved immediately. Then proceed using the Offline Update to complete the update of the cluster.
- All nodes must be in a healthy state to perform Online/Rolling Updates (**Storage > Nodes**). If they are not healthy, they need to have any issues addressed before continuing:
  - Using Offline Updates, update the OS on nodes with different OS versions. After updating, verify the node is now healthy, resolving any remaining issues.
  - Replace any unhealthy drives on nodes.
  - Replace any unhealthy nodes.

**Warning:** The OS (Online/Rolling) Update operation was aborted. Please view the [OS Update Status](#) page for details.

**You must perform an OS (Offline) Update now to ensure all nodes are updated to the correct software version.**

**OS Update Status**

Operation status: **Aborted on 2014-02-04 2:07 PM.**

Operation type: Online/Rolling Update

Nodes updated: 3 of 13 (23% complete)

Start time: 2014-02-04 2:06:30 PM

Elapsed time: 52 minutes, 19 seconds

**13 Nodes:** The update process was never started on 10 nodes.

Node	Status	Product	Version	Completion Time
Node1234501 (Mgmt. Node)	(Not started yet.)	-	-	-
Node1234509	(Not started yet.)	-	-	-
Node1234510	(Not started yet.)	-	-	-
Node1234511	(Not started yet.)	-	-	-
Node1234512	(Not started yet.)	-	-	-
Node1234513	(Not started yet.)	-	-	-
Node1234514	(Not started yet.)	-	-	-
Node1234515	(Not started yet.)	-	-	-
Node1234516	(Not started yet.)	-	-	-
Node1234517	(Not started yet.)	-	-	-
Node1234503	OK	RAINcloudOS Full Upgrade	4.0	2014-02-04 2:06 PM (17 minutes)
Node1234507	OK	RAINcloudOS Full Upgrade	4.0	2014-02-04 2:07 PM (21 minutes)
Node1234508	OK	RAINcloudOS Full Upgrade	4.0	2014-02-04 2:07 PM (14 minutes)

Refresh Close

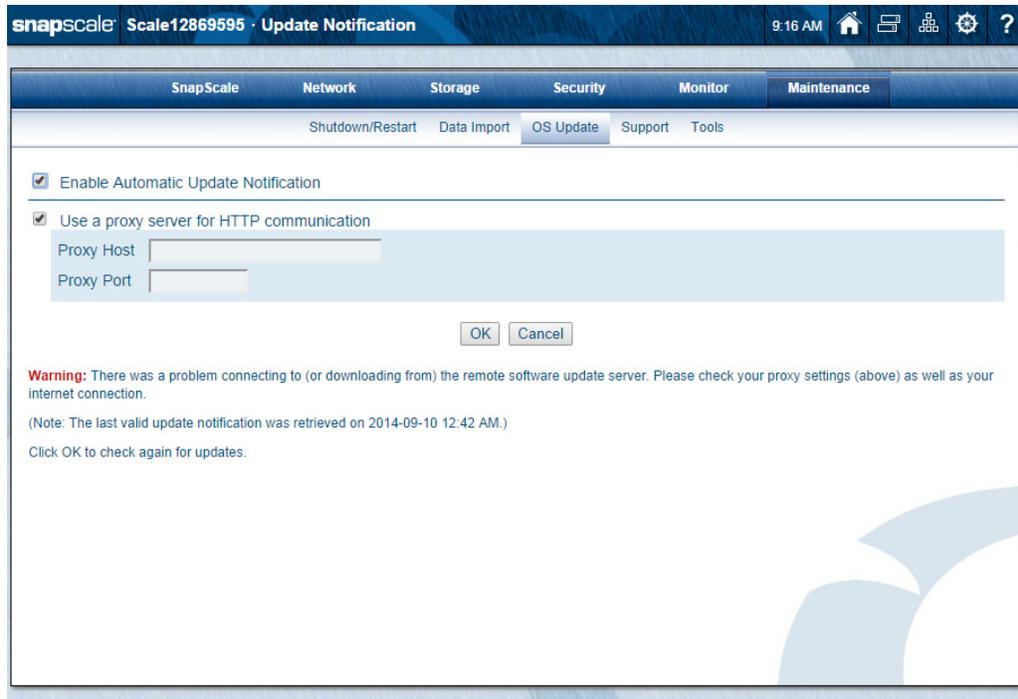
- To minimize possible conflicts, the following Web Management Interface features are blocked during an Online/Rolling Update:
  - Delete SnapScale
  - Add Nodes
  - Remove Node
  - Data Balancer
  - Spare Distributor
  - Update Notification
  - Check for Updates
  - Certain SnapScale Properties:
    - Changing Data Replication Count
    - Decreasing Spare Count
  - Date/Time
  - Create Snapshot
  - Snapshot Space
  - Snapshot Properties
  - Delete Snapshot
  - Snapshot Schedule Properties
  - Delete Snapshot Schedule

- **SnapExtensions – SnapEDR**
  - **Shutdown/Restart SnapScale**
  - **TCP/IP**
  - **Utility IP Address**
  - **Storage Network Properties**
  - **Change Password** (for a user with Admin rights)
  - **Create iSCSI Disk**
  - **Delete iSCSI Disk**
  - **iSCSI Disk Properties** (properties & active clients can still be viewed)
  - **VSS/VDS**
- Under certain circumstances, the Online/Rolling Update feature is not available:

Reason	Solution
Peer sets are degraded or rebuilding.	Address degraded issue and allow peer sets to rebuild before updating.
OS versions are not consistent across all SnapScale nodes.	Update the RAINcloudOS on the nodes individually (using the Offline Update option) so they are all the same.
Peer sets contain disk size mismatches.	Replace the smaller drives in a peer set with ones that match the size of the larger drives in that peer set.
Not supported for this software update file.	Contact Technical Support.
The previously run Online/Rolling Update failed or was aborted.	Contact Technical Support.
Due to an unknown reason.	Contact Technical Support.

## Update Notification Option

You can configure RAINcloudOS to display an alert when updates are available for the cluster.

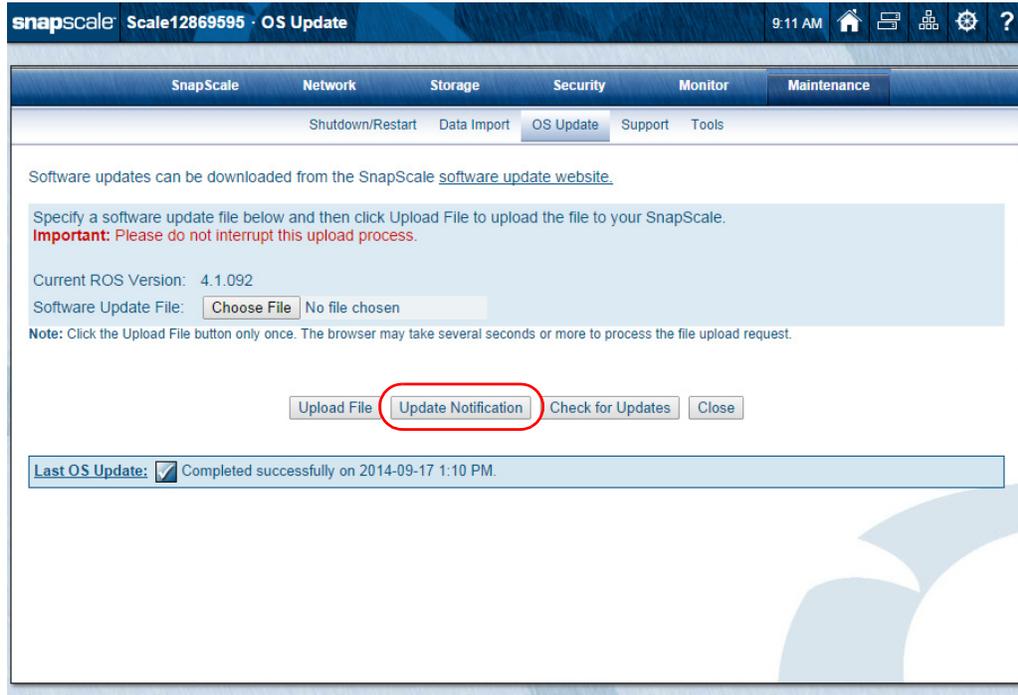


When enabled, **Update Notification** checks weekly for updates that are applicable to the cluster. If updates are available, a banner alert is displayed just below the menu bar on all Web Management Interface pages.

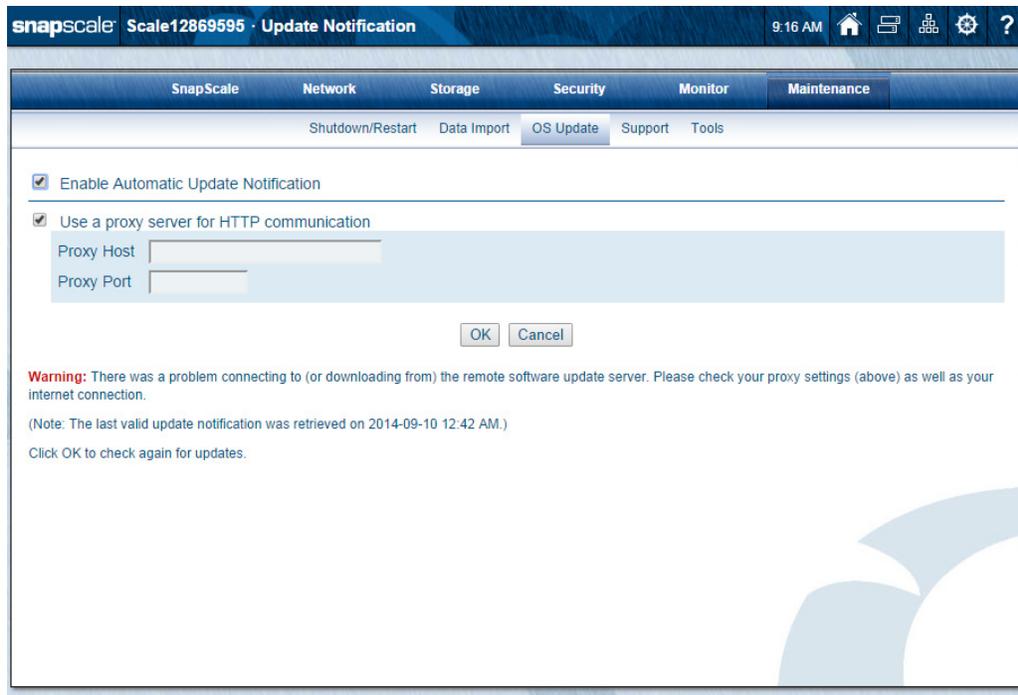
**NOTE:** You can choose to hide the banner by clicking either the *Remind me later* or *Hide this message* link on the banner. For *Remind me later*, the Web Management Interface displays the banner after the next check for updates; for *Hide this message*, the banner is hidden for the update in question until a later version is released.

## Configuring Update Notification

1. Go to **Maintenance > OS Update** and click **Update Notification**:



2. Check the **Enable Automatic Update Notification** box.

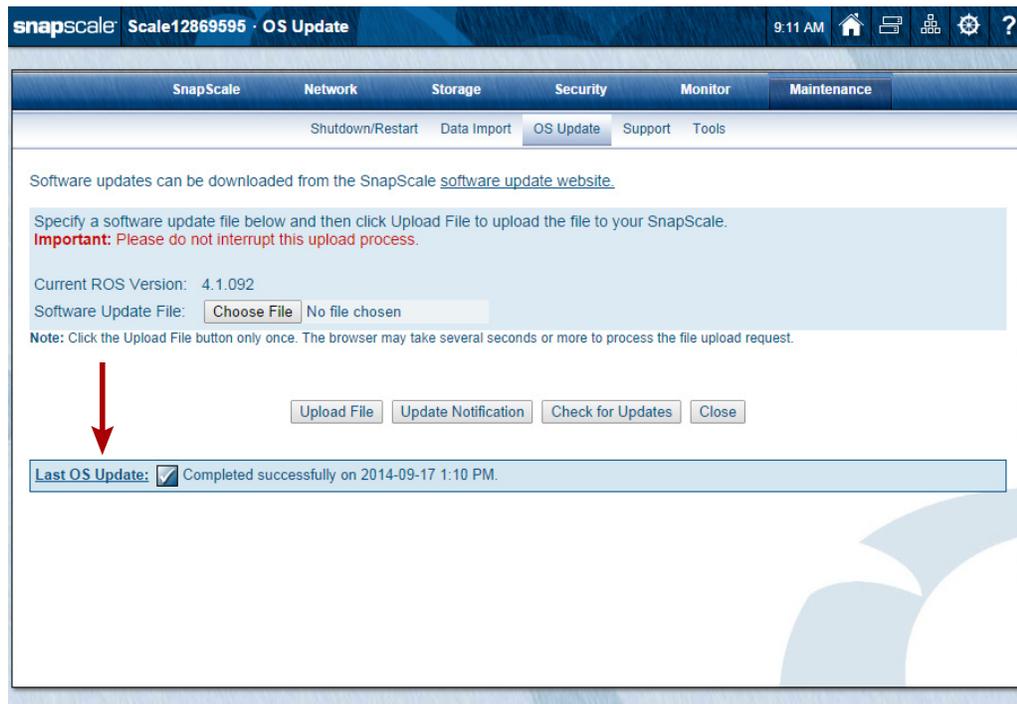


3. If your environment requires using a **proxy server** for external web-based communication:
  - a. Check the **Use a proxy server for HTTP communication** box. Additional proxy options are displayed.

- b. Complete the **Proxy Host** and **Proxy Port** fields.
4. Click **OK**.

## Last OS Update

At the bottom of the **OS Update** page is a **Last OS Update** link and information. Click this link to view a detailed status of the last update applied to the cluster.

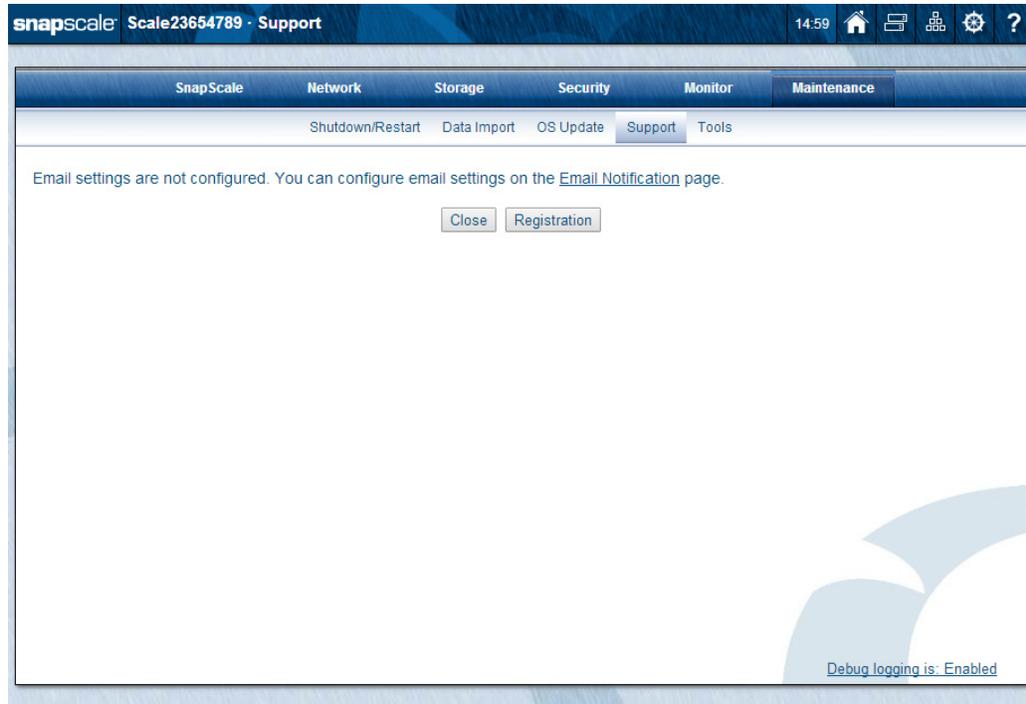


## Support

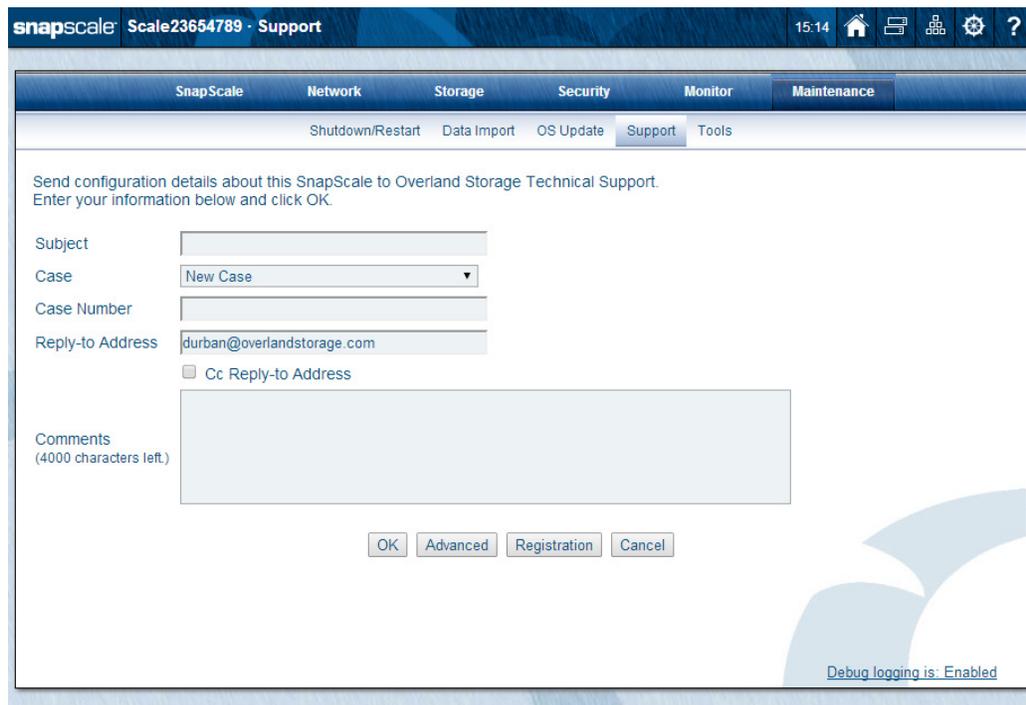
The **Support** page provides an easy way to contact Overland Technical Support, and transmit system logs and files that contain information useful for troubleshooting purposes.



**IMPORTANT:** The **Support** page is not accessible until you have configured **Email Notification** in the **Tools** submenu.



Once email is configured, the **Support** page is available with your contact information entered:

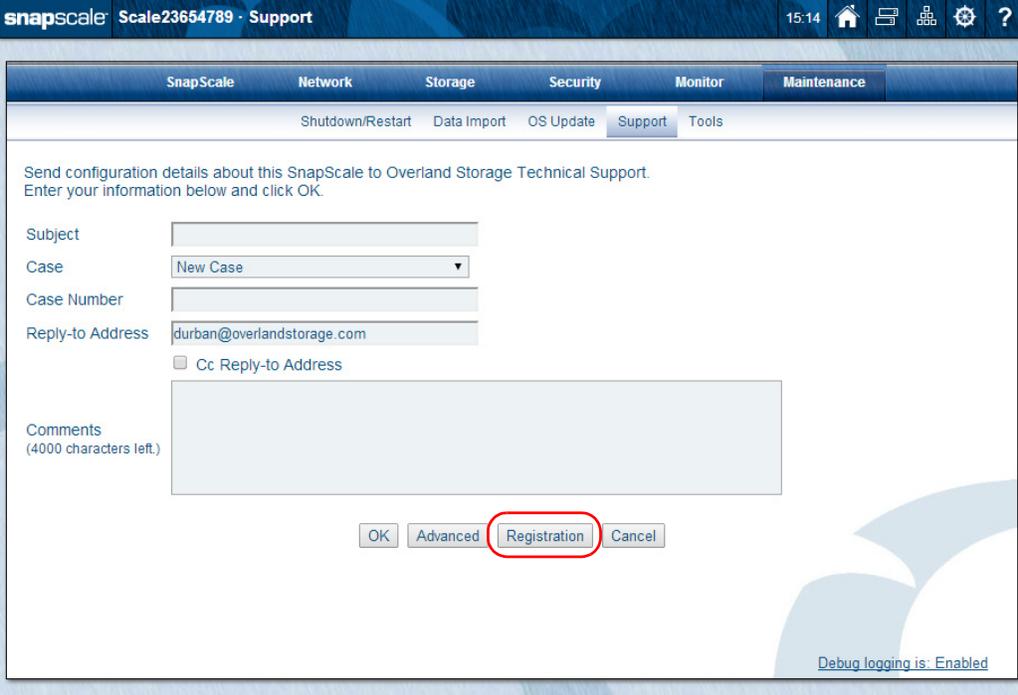


## Registering Your Cluster

The very first time you start your cluster, a **Registration Reminder** page appears. Registering your cluster activates your warranty and allows you to create and track service requests. Registration also provides access to RAINcloudOS upgrades, third-party software, and exclusive promotional offers.

**NOTE:** Warranty information is available at <http://docs.overlandstorage.com/support>.

If you skipped the registration during setup, to register the cluster now, click **Registration** on the **Maintenance > Support** page:



The screenshot shows the SnapScale web interface. At the top, the breadcrumb navigation is "SnapScale > Scale23654789 > Support". The main navigation bar includes "SnapScale", "Network", "Storage", "Security", "Monitor", and "Maintenance". Under "Maintenance", there are sub-tabs for "Shutdown/Restart", "Data Import", "OS Update", "Support", and "Tools". The "Support" tab is active, displaying a form to send configuration details to Overland Storage Technical Support. The form includes fields for "Subject", "Case" (a dropdown menu currently set to "New Case"), "Case Number", and "Reply-to Address" (pre-filled with "durban@overlandstorage.com"). There is a checkbox for "Cc Reply-to Address" which is unchecked. Below these is a large text area for "Comments (4000 characters left.)". At the bottom of the form are four buttons: "OK", "Advanced", "Registration" (which is circled in red), and "Cancel". In the bottom right corner of the page, it says "Debug logging is: Enabled".

### To Register Your Cluster

**NOTE:** To use this feature, access to the Internet is required.

To register your cluster to activate its warranty support, you can either:

- Click the link on the initial **Registration Reminder** page.
- Go to **Maintenance > Support** and click **Registration**.

At the **Registration** page:

**Register your SnapScale to activate your warranty and stay informed of important software and product updates.**  
(Note: You must register your SnapScale whenever you add new nodes.)

Please enter your contact information below and then click the Download button to download the registration (text) file to your computer. You can then email the registration file to [warranty@overlandstorage.com](mailto:warranty@overlandstorage.com) with the subject, **SnapScale Registration Request**.

Name

Email Address

Company Name

Shipping Address  (For RMA purposes.)

You can view the Overland Storage privacy policy by visiting [www.snapserver.com/privacy](http://www.snapserver.com/privacy) (a new browser window will open).

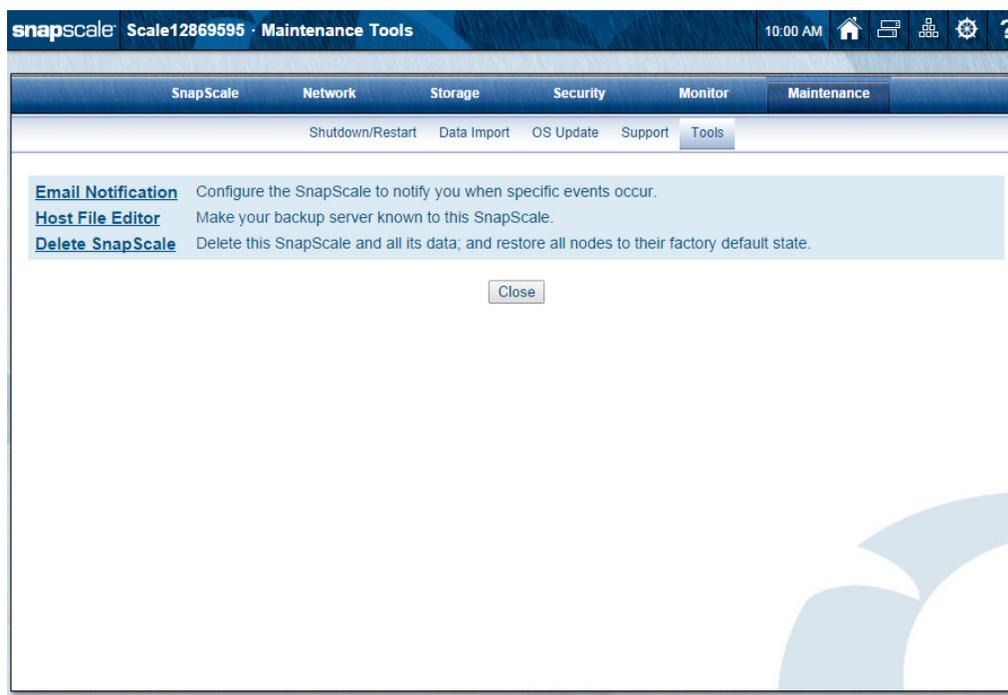
1. Enter the **four required items** in the appropriate fields.
2. Click **Download Registration File**.  
The information, including all the node data, is incorporated into a CSV file.
3. Depending on your browser settings, make sure that you **save the CSV file** to your local computer.
4. Email the downloaded CSV file to **warranty@overlandstorage.com**.  
Use the subject line **SnapScale Registration Request** for the email.

**NOTE:** The same page can also be used to **update** your registration information. For example, when you add new nodes to your cluster, they need to be added to the cluster registration so that they are also covered. Just repeat the steps above.

Once you have registered, you will receive a confirmation email to complete the registration.

## Maintenance Tools

The **Tools** option provides a submenu of general-purpose maintenance options and features.



## Email Notification

To configure the cluster to send email alerts in response to system events or activate Overland support, navigate to **Maintenance > Tools > Email Notification**.

The screenshot shows the 'Email Notification' configuration page in the SnapScale interface. The page is titled 'Scale23654789 - Email Notification'. It features a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Maintenance tab, there are sub-tabs for Shutdown/Restart, Data Import, OS Update, Support, and Tools. The main content area includes the following settings:

- Enable Email Notification
- SMTP Server:  (Host name or IP address)
- SMTP Port:  (Port number for SMTP server)
- Use Authenticated SMTP
- Use Secure Connection
- Email Address of Sender:
  - Use default Scale23654789@devnet.myoverland.net
  - Use specific
- Email Addresses of Recipients:
  - (optional)
  - (optional)
  - (optional)
- Send email notification for the following events:
  - Node shutdown/restart
  - Node hardware event
  - SnapScale system event
  - Administrative operation event
  - Storage usage warning event (See SnapScale Properties)
- Send a test email to listed email addresses upon saving settings.

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

To set up email alerts, you need the SMTP server's IP address and the email address of each recipient (up to four) who is to receive the alert.

### Configuring Email Notification

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

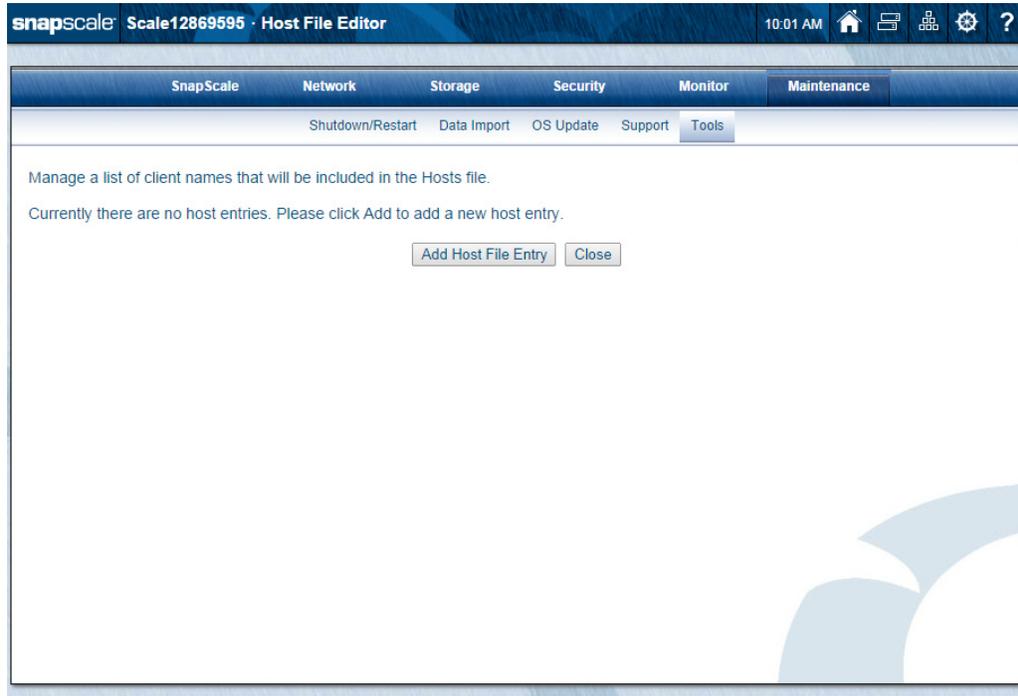
Option	Description
<b>Enable Email Notification</b>	To enable email notification, check the <b>Enable Email Notification</b> box.
<b>SMTP Server</b>	Enter a valid SMTP server IP address or host name.
<b>SMTP Port</b>	Enter a port number for the SMTP server or accept the default. The default is 25.
<b>Use Authenticated SMTP</b>	Check this box to authenticate when an email is sent to the SMTP server by the SnapScale. Provide an authentication User Name and Password in the fields that appear when the feature is enabled. The types of methods supported (in order) are CRAM-MD5, LOGIN, and PLAIN.
<b>Use Secure Connection</b>	Check this box to encrypt emails from the cluster. STARTTLS and TLS/SSL encryption protocols are supported.

Option	Description
<b>Email Address of Sender:</b>	Choose one: <ul style="list-style-type: none"> <li>The default address (<i>cluster_name@domain</i>) where the <i>domain</i> is the DNS domain name. If there is no DNS domain name, then the server's IP address for Eth0 will be used (<i>cluster_name@ipaddress</i>).</li> <li>Specify a specific sender.</li> </ul>
<b>Email Addresses of Recipients</b>	Enter the email addresses to receive the notifications. One address is required but as many as four email addresses can be entered.
<b>Send Email Notification</b>	Check the boxes next to the events you wish to be notified about: <ul style="list-style-type: none"> <li><b>Node shutdown/restart</b> – The node shuts down or reboots due to an automatic or manual process.</li> <li><b>Node hardware event</b> – The internal temperature for a node exceeds its maximum operating temperature or other hardware problems.</li> <li><b>SnapScale system event</b> – A change or error occurs that impacts the entire cluster.</li> <li><b>Administrative operation event</b> – A Data Import operation has finished or experienced an error.</li> <li><b>Storage usage warning event</b> – Storage space usage on a volume reaches either the maximum utilization or the critical utilization setting. See <a href="#">SnapScale Properties on page 43</a>.</li> </ul>
<b>Send a Test Email</b>	To verify your settings, check <b>Send a test email to listed email addresses upon saving settings</b> , then click <b>OK</b> .

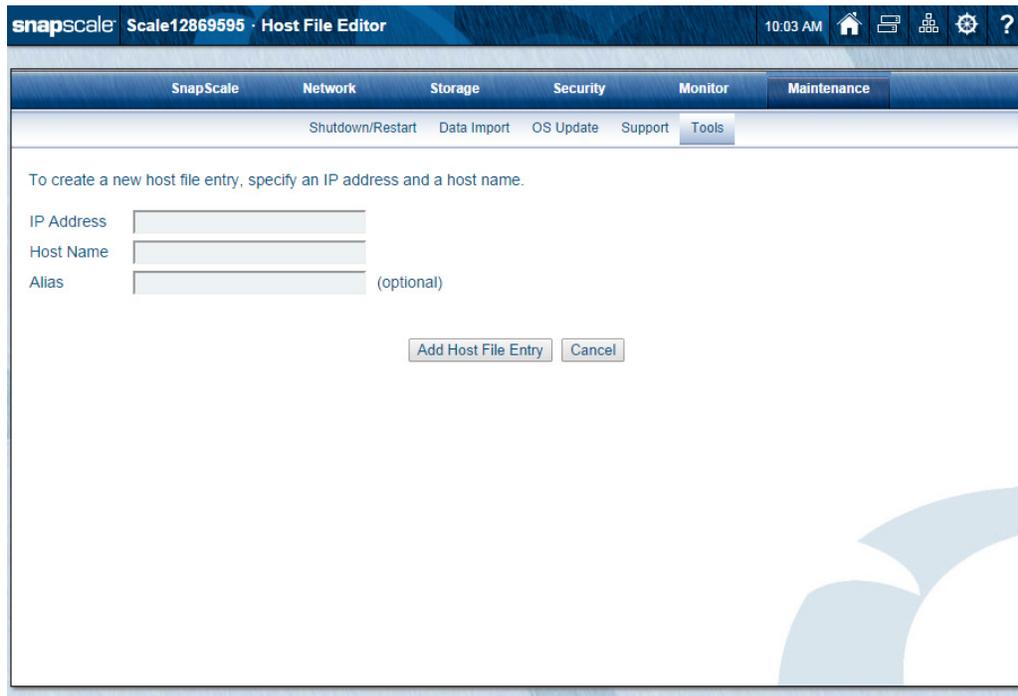
If **Send a Test Email** is checked, when you save your changes, an email is sent to all configured email recipients.

## Host File Editor

Use this page to identify external hosts in the hosts file for the SnapScale cluster. This page allows you to supply a hostname-to-IP address mapping that persists across system reboots.



Click **Add Host File Entry**, complete the fields as described on the table below, and then click **Add Host File Entry** again.



Use this table to complete the options shown:

Option	Description
<b>IP Address</b>	The IP address of the external host.
<b>Host Name</b>	Enter the fully qualified hostname for the external host, using the format: <i>myserver.mydomain.com</i> .  <b>NOTE:</b> Some applications may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.
<b>Alias (optional)</b>	Enter an optional abbreviated address for the external host, using the format: <i>myserver</i> .  <b>NOTE:</b> Some applications may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.

## Delete SnapScale Cluster

This page is used to delete a SnapScale cluster and all its data.



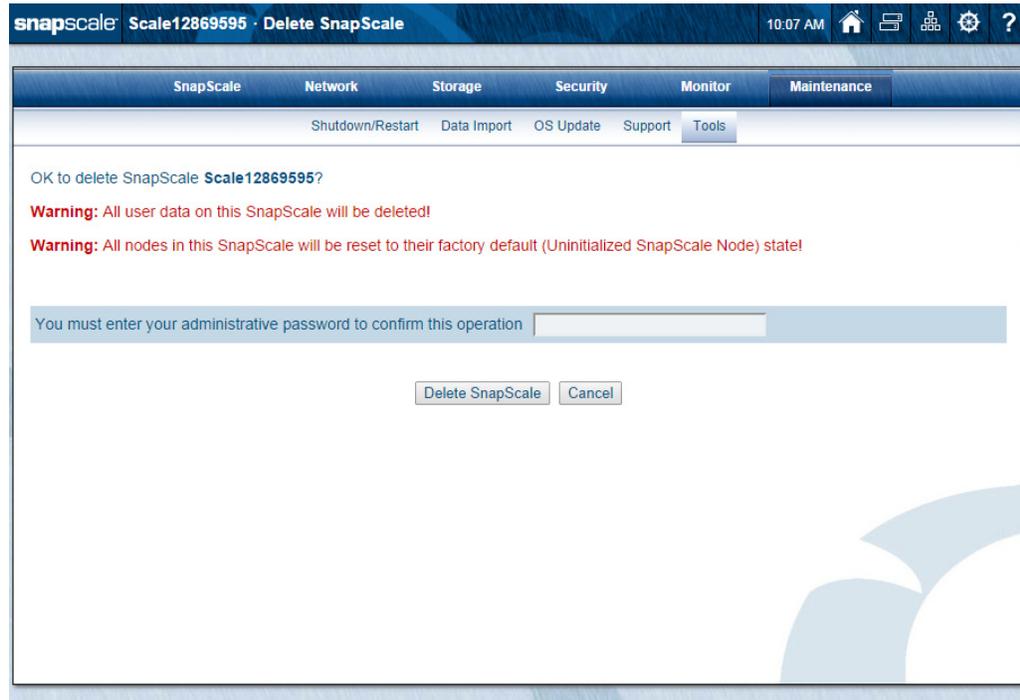
**CAUTION:** All data on all the nodes will be lost and all the nodes will be reset to their original factory default settings. No recovery is possible.

- Select whether to restart or shut down all the nodes after the cluster is deleted.
  - Restart all nodes** – After deleting the SnapScale cluster, the nodes reboot, they automatically perform a fresh install, and then they reboot as Uninitialized nodes.
  - Shut down all nodes** – After deleting the SnapScale cluster, the nodes shut down. The next time the nodes are powered on, they automatically perform a fresh install and then reboot as Uninitialized nodes.

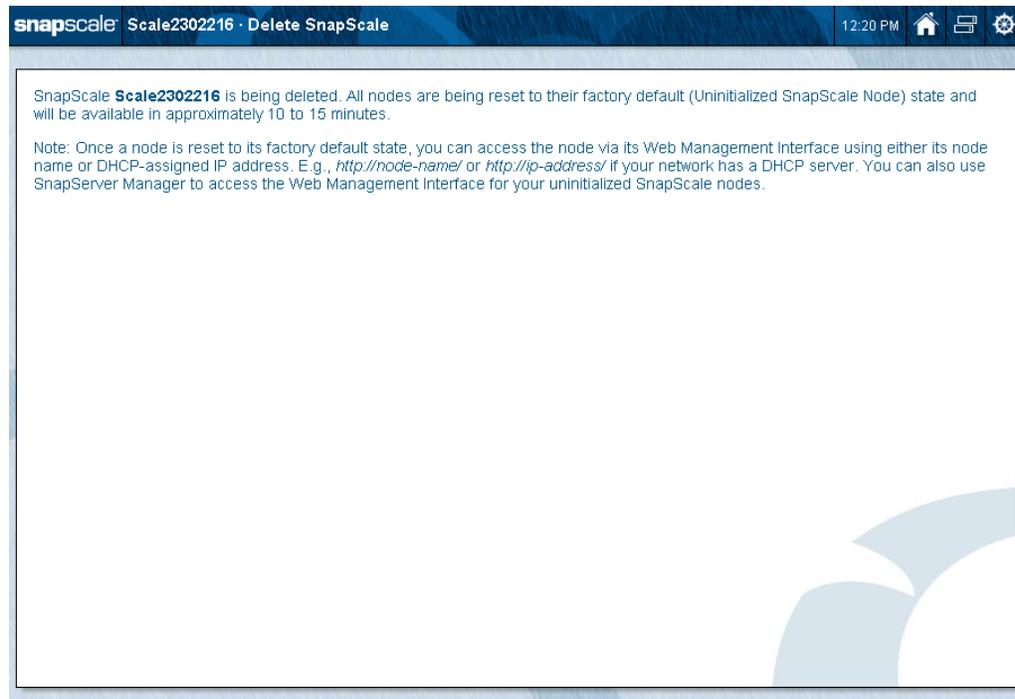


- Click **Delete SnapScale**.

- At the confirmation page, enter your **Admin password** and click **Delete SnapScale** again to start the process.



During the cluster deletion, an information page is shown (such as this one for the Restart option).



The RAINcloudOS site map (⚙️) provides links to a majority of the web pages that make up the Web Management Interface. It also provides, in the last column, special links to higher level options and processes which are the focus of this chapter.

With the exception of **Mgmt. Interface Settings**, these options are also directly navigable from the various menus in the Web Management Interface. Also the **Home**, **Snap Finder**, **SnapExtensions**, **Site Map**, and **Help** options are accessible from any page by clicking their respective icon in the top right corner of the page (see the table in [Web Management Interface on page 37](#)).

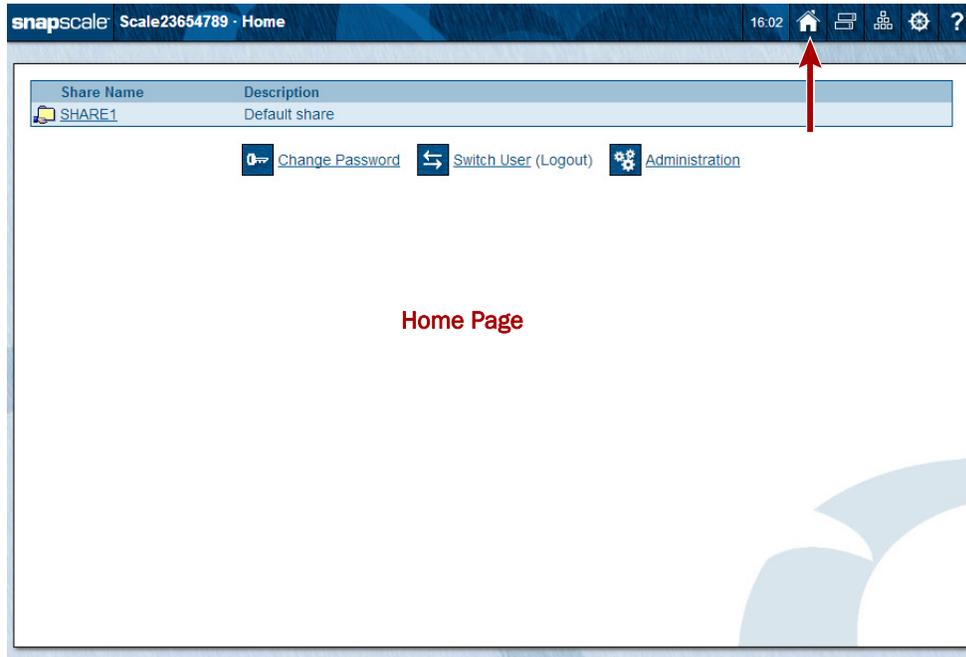
snapscale						
SnapScale	Network	Storage	Security	Monitor	Maintenance	Misc.
SnapScale Properties	Information	Peer Sets	Security Guides	System Status	Shutdown/Restart	Administration
Date/Time	TCP/IP	> Spare Disks	Shares	Activity	Data Import	Home
SSH	> Utility IP Address	> Spare Distributor	> Create Share	> Active Users	OS Update	SnapExtensions
UPS	> Storage Network Props	> Data Balancer	Local Users	> Open Files	> Update Notification	Snap Finder
	Windows/SMB	Volumes	> Create Local User	> Network Monitor	> Check for Updates	> Snap Finder Properties
	NFS	> Create Volume	> Password Policy	Event Log	> OS Update Status	BitTorrent Sync
	LDAP/NIS	Quotas	Local Groups	Protocol Manager	Support	Change Password
	FTP	Snapshots	> Create Local Group	SnapScale Settings	> Registration	Mgmt. Interface Settings
	SNMP	> Create Snapshot	Security Models	Tape	Tools	
	Web	> Snapshot Schedules	ID Mapping		> Email Notification	
	iSNS	> Snapshot Space	Home Directories		> Host File Editor	
		iSCSI			> Add Host	
		> Create iSCSI Disk			> Delete SnapScale	
		> VSS/VDS Access Control				
		Nodes				
		> Add Nodes				
		> Node Identification				
		Disks				

### Topics in Misc. Options

- [Home Pages](#)
  - [Home Page](#)
  - [Administration Page](#)
- [SnapExtensions](#)
  - [BitTorrent Sync](#)
  - [Snap EDR](#)
- [Snap Finder](#)
  - [Edit Snap Finder Properties](#)
- [Change Password](#)
- [Management Interface Settings](#)

## Home Pages

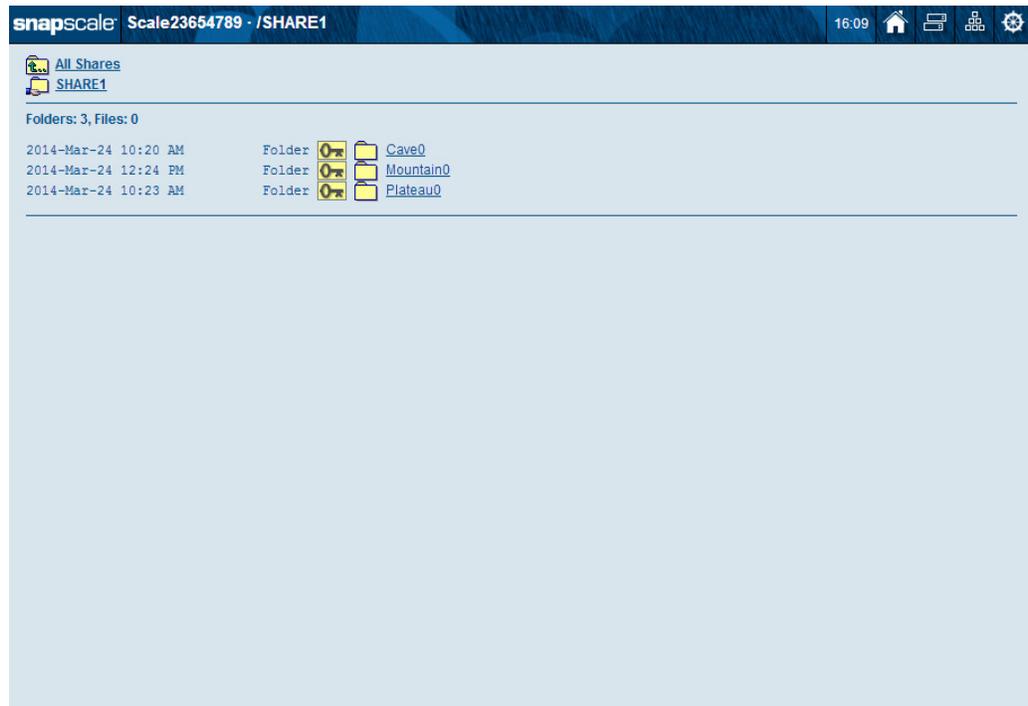
When you first launch the Web Management Interface, the **Home** page is displayed showing any existing shares and three options. Once logged in using the **Administration** link, you can switch between the **Home** page and the **Administration** page using the Home page (🏠) icon on the button bar.



### Home Page

The Web Management Interface **Home** page displays a list of all shares and three basic options. Users can navigate the share structure to locate and view or download files without logging in but they cannot modify or upload files.

Click a share name to see a list of files:

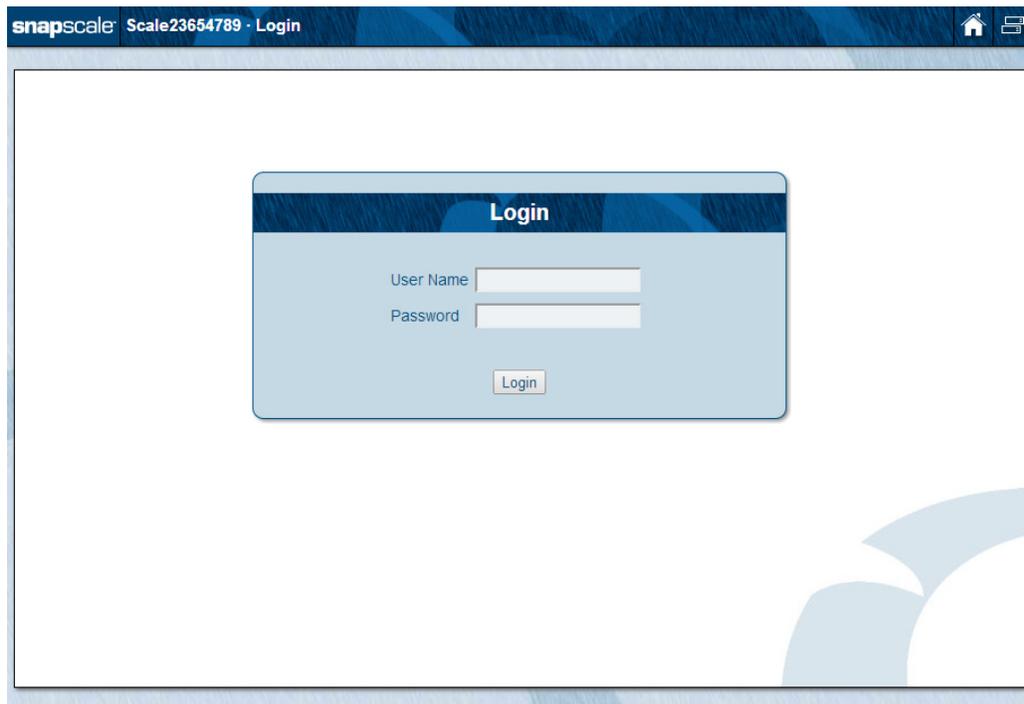


For users with admin rights, a key icon () appears next to the file/folder on the share. Clicking this icon displays a popup box with security information about the file/folder.

This page also provides three key administrative function links:

- **Change Password** () – Takes you to the **Change Password** page where you can change your administration password. Enter your **User Name** and **Current Password** for access. See [Change Password on page 242](#).

- **Switch User (Logout)** (🔒) – Automatically logs out the current user and displays the **Login** page for the new user to gain access to the SnapScale cluster.



- **Administration** (⚙️) – Displays the **Administration** page (see [Administration Page on page 234](#)). You will be prompted to log in if you have not already done so.

If any of the following conditions are present, you may not be able to access the Home page:

- **Require Web Authentication** is enabled (via **Network > Web > Require Web Authentication**) and you do not have a valid user name and password on the cluster.
- The cluster has not completed the **Initial Setup Wizard** (if this is the case, you will not be able to access the **Administration** page of the Web Management Interface either).
- **Web Root** is enabled (via **Network > Web > Enable Web Root**).

## Administration Page

The **Administration** page is accessible by clicking either the **Administration** link in the Site Map or the Administration (⚙️) or Home page (🏠) icons on the **Home** page. If web root is enabled, it can also be accessed directly by entering the address:

```
http://<node_name>/sadmin
```

in a web browser where *<node\_name>* is the unique node name in the format Nodennnnnnnnn. It provides a high-level view of the SnapScale status, the amount of total storage being used, and a link to find out what's new in RAINcloudOS by accessing online help. The tabs at the top provide access to the various functions and features of the RAINcloudOS.

The screenshot shows the SnapScale Administration page for cluster Scale7846949. The page has a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. A help icon (?) is in the top right. Below the navigation bar is a large box containing cluster information and status indicators.

**Cluster Information:**

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

**Status Indicators:**

- Peer Sets: 5 (All peer sets OK)
- Nodes: 3 (All nodes OK)
- Active Spare Disks: 2 (All spares OK)
- Protocol Manager: All nodes OK
- SnapScale Settings: All settings OK
- UPS Status: All nodes OK

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

Buttons: Refresh, Close

Link: Click here to find out what's new in RAINcloudOS 4.1 and this Web Management Interface.

The **Auto-refresh** icon (🔄) is displayed on the right corner just above the Cluster Status Box. Click the icon (or **Refresh** below the Cluster Status Box) to manually refresh the information.

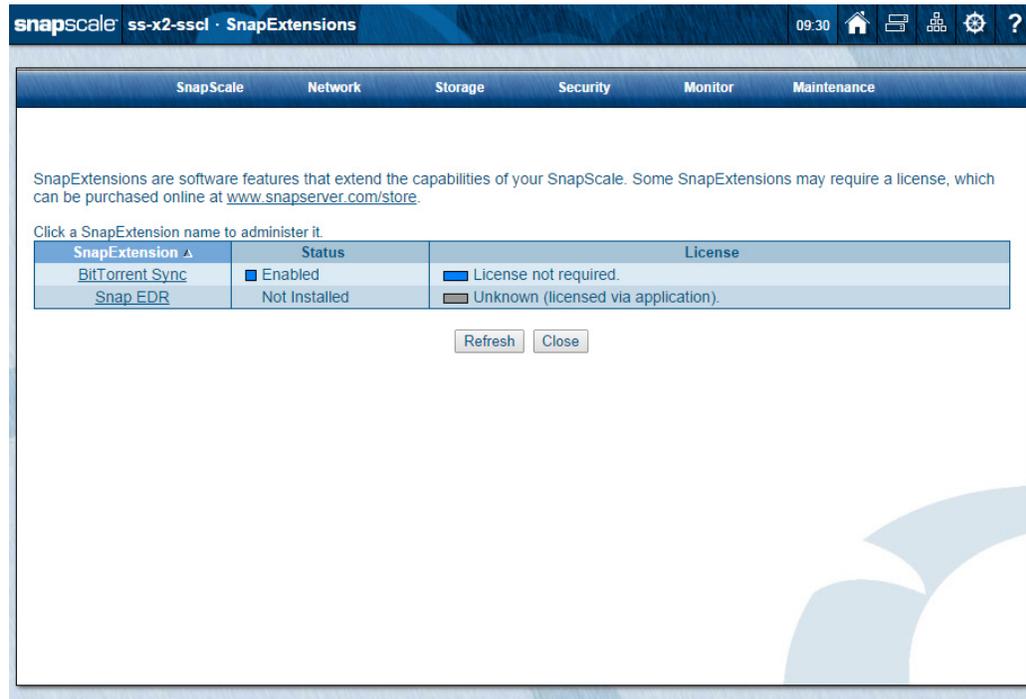
This screenshot is identical to the one above, but with a red arrow pointing to the auto-refresh icon (🔄) located in the top right corner of the main content area, just above the Cluster Status Box.

From the **Administration** page, clicking  takes you to the **Home** page.

## SnapExtensions

The SnapExtensions icon  opens the SnapExtensions page. This page is used to manage the SnapExtensions installed on your SnapScale cluster.

**NOTE:** Mouseover the icon to display a popup menu with direct access to SnapExtensions that are both installed and enabled.



If any SnapExtensions are installed, you can click the SnapExtension name in the left column of the table to display the management page for that extension. Mousing over this icon shows all active extensions and provides a direct access link to them.

### BitTorrent Sync



**CAUTION:** BitTorrent Sync bypasses share and file security. Be sure to only share data that is intended to be accessible by any user with the folder secret.

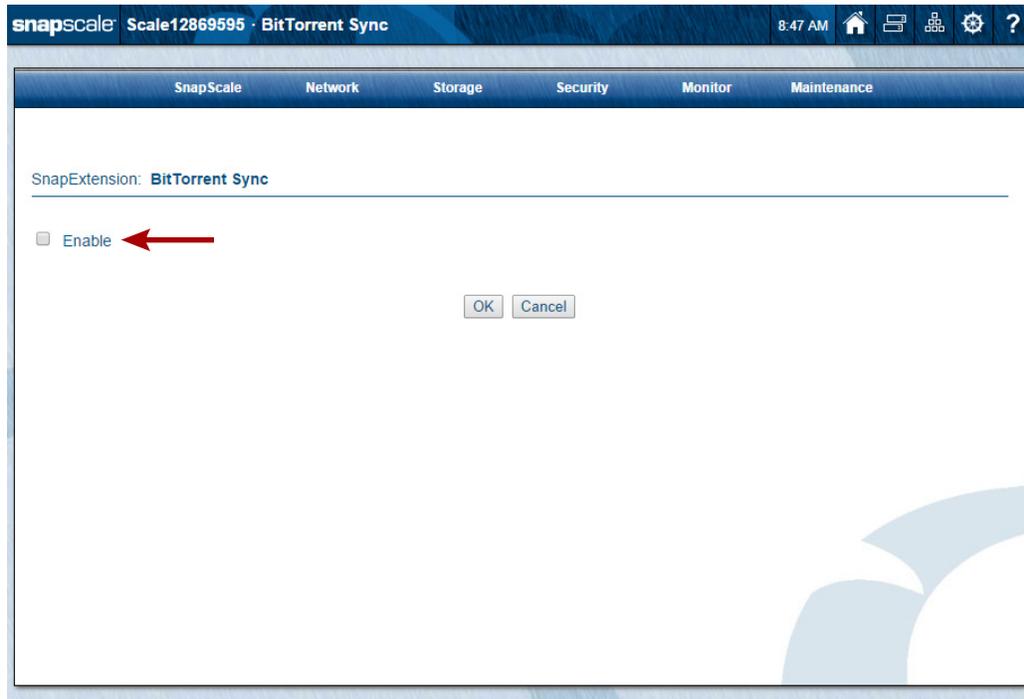
**NOTE:** Cookies must be enabled on your browser for BitTorrent Sync to work.

BitTorrent Sync (BTSync) is a SnapExtension that is preloaded on SnapScale. It lets you share and sync an unlimited number of files and folders of any size across multiple platforms. For more information, visit <http://www.bittorrent.com/sync>.

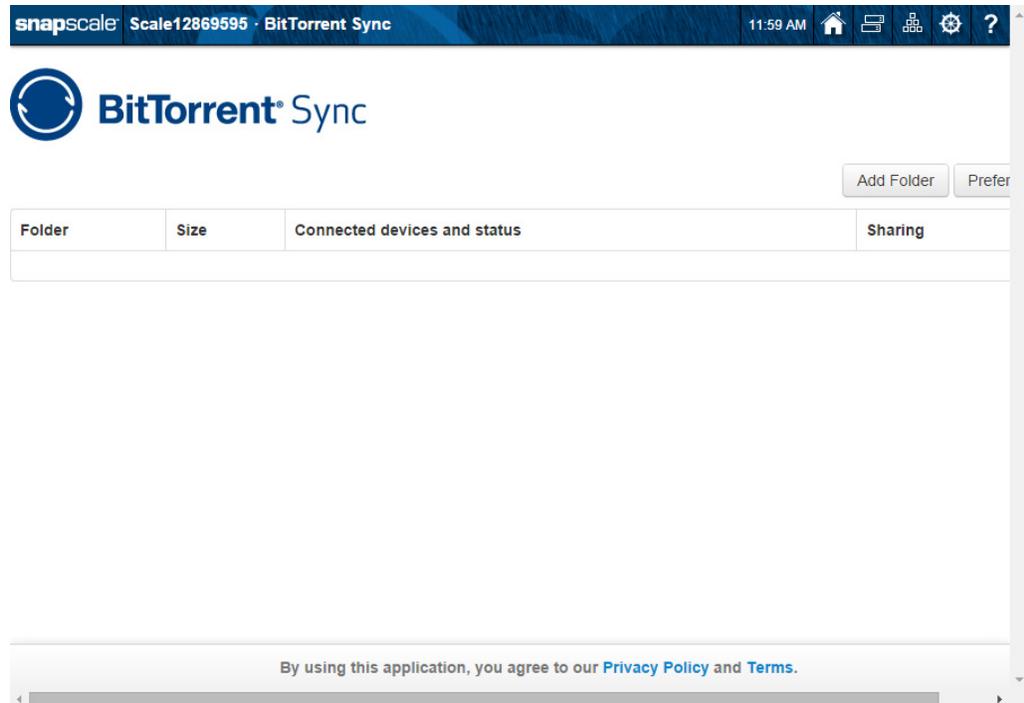
To use BitTorrent Sync, it must first be enabled:

1. On the **SnapExtensions** page, click the **BitTorrent Sync** name in the table to access the configuration page.

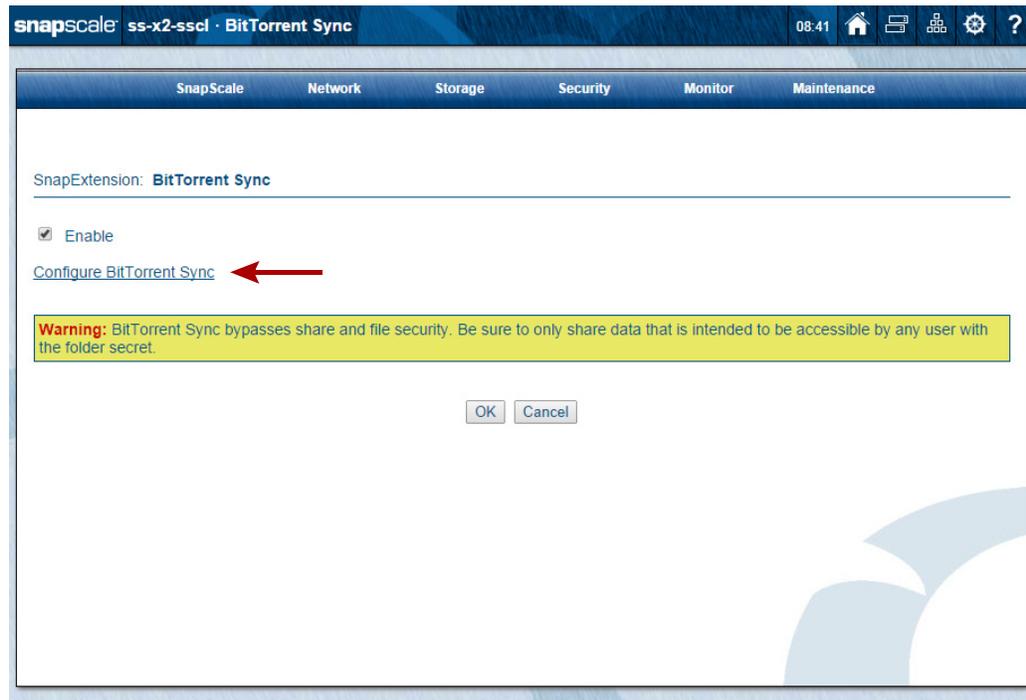
- At the configuration page, check **Enable** and click **OK**.



- At the following page, accept the BitTorrent Privacy Policy and Terms by clicking **I Accept** so you can run BitTorrent Sync.

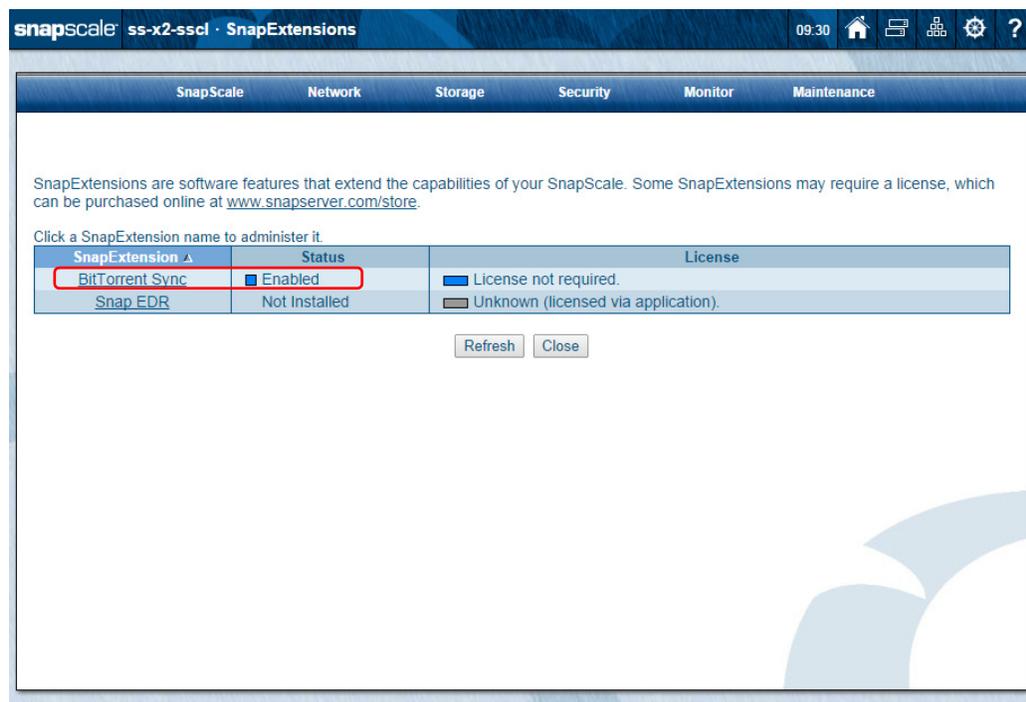


- When returned to the **BitTorrent Sync** page, use the **Configure BitTorrent Sync** link to configure it for your use.



**NOTE:** Once enabled, you can return at any later time to configure or reconfigure BitTorrent Sync by clicking **BitTorrent Sync** on the **SnapExtensions** page or the **Site Map**.

- When configuration is complete, click **OK**.  
You are returned to the **SnapExtensions** page. **BitTorrent Sync** shows **Enabled**.



**NOTE:** To turn off (disable) the BitTorrent Sync feature, uncheck the BitTorrent Sync box.

### BitTorrent Considerations

- For the most recent information and details on BitTorrent Sync configuration and use, refer to the documentation available on their web site.
- BitTorrent Sync cannot be used to replicate snapshots because it requires the ability to write to the sync location and snapshots are read only.
- BitTorrent Sync installs as a hidden directory on a volume. If the volume is deleted or rolled back from a snapshot, the SnapScale attempts to automatically relocate the BitTorrent Sync hidden directory to another volume. If there are no more volumes, or if none can be found that are large enough, BitTorrent Sync is disabled and cannot be re-enabled until a suitable volume becomes available. Once re-enabled, BitTorrent Sync must be completely reconfigured again.
- There are BitTorrent Sync mobile apps that make synced documents available on iOS, Android, Windows Phone 8, and Kindle Fire systems. Refer to their web site for details on configuration and use of BitTorrent Sync on mobile apps.

## Snap EDR

For **SnapEDR**, at the **Configuration** page, select either to configure it as the Management Console or as an agent of another Management Console. If configuring it as an agent, enter the **Name or IP of the Management Console**.

After SnapEDR finishes its configuration, the Management Console screen is shown on the **Configuration** page:

## Snap Finder

**Snap Finder** (☰) is a powerful tool that lists all the SnapServer appliances, SnapScale clusters, and Uninitialized nodes on your network (and on a remote network segment if so configured), and shows the current status of each. Click the unit name (if you have name resolution) or IP address of a cluster, node, or server to access it through the Web Management Interface.

**NOTE:** You can sort the columns (ascending or descending order) by clicking the column heading.

78 SnapServers. 17 SnapScale Clusters. 7 Uninitialized SnapScale Nodes.  
Click a server name (if your network has name resolution) or IP address to connect to a server.

Server	Status	IP Address	OS Version	Model	Number	Avail Cap.	Total Cap.
at-ros41Ekanite	Online	10.25.12.65	ROS 4.1.0.autobuild_Sep23@16.02	-	-	3.32 GB	3.99 G
beryl	OK	10.25.3.27	GOS 7.5.048	DX1	2300028	4.34 TB	4.34 T
BlueBottle	OK	10.25.2.18	GOS 7.6.0.philtonsled11	DX2	2415534	735.27 GB	735.28 G
bm-gos-72	OK	10.25.2.17	GOS 7.2.131	VirtualSnap	12938986	6.03 GB	6.09 G
bmgos4	OK	10.25.2.40	GOS 7.6.0.briansled11	VirtualSnap	15284780	2.26 GB	4.36 G
bmgos5	OK	10.25.2.49	GOS 7.6.0.briansled11	VirtualSnap	15459346	6.07 GB	6.09 G
bmgos8	OK	10.25.2.230	GOS 7.2.130	VirtualSnap	15922671	5.59 GB	6.09 G
bmros41	Online	10.25.12.230	ROS 4.1.0.briansled11	-	-	24.49 GB	27.81 G
bobbert	OK	10.25.3.38	GOS 5.0.133	4400	1723986	118.84 GB	280.71 G
CB-Meadowhawk	OK	10.25.15.62	GOS 7.5.047	DX2	2415100	638.90 GB	719.75 G
CB-Sundragon	OK	10.25.15.60	GOS 7.5.047	DX2	2413126	1.31 TB	1.46 T
CCCloudDX1	Online	10.25.17.237	ROS 4.1.088	-	-	4.26 TB	4.27 T
CCCloudDX2	Online	10.25.17.230	ROS 4.1.088	-	-	28.65 TB	32.40 T
CCDX1SNAP	OK	10.25.3.88	GOS 7.5.047	DX1	2301028	3.52 TB	6.52 T
CCSundragon3	OK	10.25.3.156	GOS 7.6.064	DX2	2411108	27.34 TB	27.34 T
CCWAVE410	OK	10.25.2.152	GOS 5.2.067	410	2250681	517.44 GB	836.12 G
CCWave412	OK	10.25.17.87	GOS 6.5.029	410	2277760	2.97 TB	4.36 T
CCXSR120EXP	OK	10.25.3.35	GOS 7.6.071	XSR 120	2425094	4.23 TB	4.31 T
CLW240	Online	10.25.12.240	ROS 4.1.0.charissa-devel	-	-	8.54 TB	8.54 T
daedalus	OK	10.25.10.32	GOS 7.2.117	DX1	2302760	1.29 TB	2.00 T
devqa	OK	10.25.11.2	GOS 6.5.029	N2000	730062	2.51 TB	2.70 T
DX1-SJSE-CB-RDX	OK	10.25.2.88	GOS 7.6.028	DX1	2303856	2.15 TB	2.15 T
DX1-VGOS-CB	OK	10.25.3.148	GOS 7.5.047	VirtualSnap	9532834	18.44 GB	18.45 G
DX1-VGOS1-CB	OK	10.25.3.149	GOS 7.6.028	VirtualSnap	9525848	0.00 MB	0.00 M
DX2-VGOS-CB	OK	10.25.2.28	GOS 7.6.028	VirtualSnap	9506780	61.77 GB	61.80 G
DX2-VGOS1-CB	OK	10.25.3.135	GOS 7.6.028	VirtualSnap	9531579	61.78 GB	61.80 G
excess	OK	10.25.3.51	GOS 7.6.071	XSR 40	2312726	7.49 TB	7.50 T

The following table describes the columns in the table:

Identification	Description
<b>Server</b>	Name of the SnapServer appliance, SnapScale cluster, or Uninitialized node. The default name is <i>Snapnnnnnnnn</i> , <i>Scalennnnnnnn</i> , or <i>Nodennnnnnnn</i> , where <i>nnnnnnnn</i> is the number of the server, node originally used to create the cluster, or the Uninitialized node. For example, "Scale23022161."
<b>Status</b>	<ul style="list-style-type: none"> <li>The status of the SnapServer or Uninitialized node (for example, <b>OK</b> or <b>Fan Failure</b>).</li> <li>The status of a SnapScale cluster is always <b>Online</b>.</li> </ul>
<b>IP Address</b>	The IP address of the SnapServer, Uninitialized node, or the Management IP address of the SnapScale cluster.
<b>OS Version</b>	The OS version currently installed on the SnapServer, Uninitialized node, or SnapScale cluster.
<b>Model</b>	The hardware model number of the SnapServer or Uninitialized node. This field is not applicable to a SnapScale cluster.
<b>Number</b>	The server or node number derived from the MAC address of the primary Ethernet port, used as part of the default name. This field is not applicable to a SnapScale cluster.
<b>Avail Cap.</b>	The available capacity on the SnapServer or SnapScale cluster. This field is not applicable to an Uninitialized node.
<b>Total Cap.</b>	The total capacity on the SnapServer or SnapScale cluster. This field is not applicable to an Uninitialized node.

To enable remote discovery of clusters, nodes, or servers on a different subnet or to display a warning icon for SnapServers or Uninitialized nodes with an enabled Ethernet port that has no link, click **Properties** at the bottom of the page to open the **Snap Finder Properties** page.

## Edit Snap Finder Properties

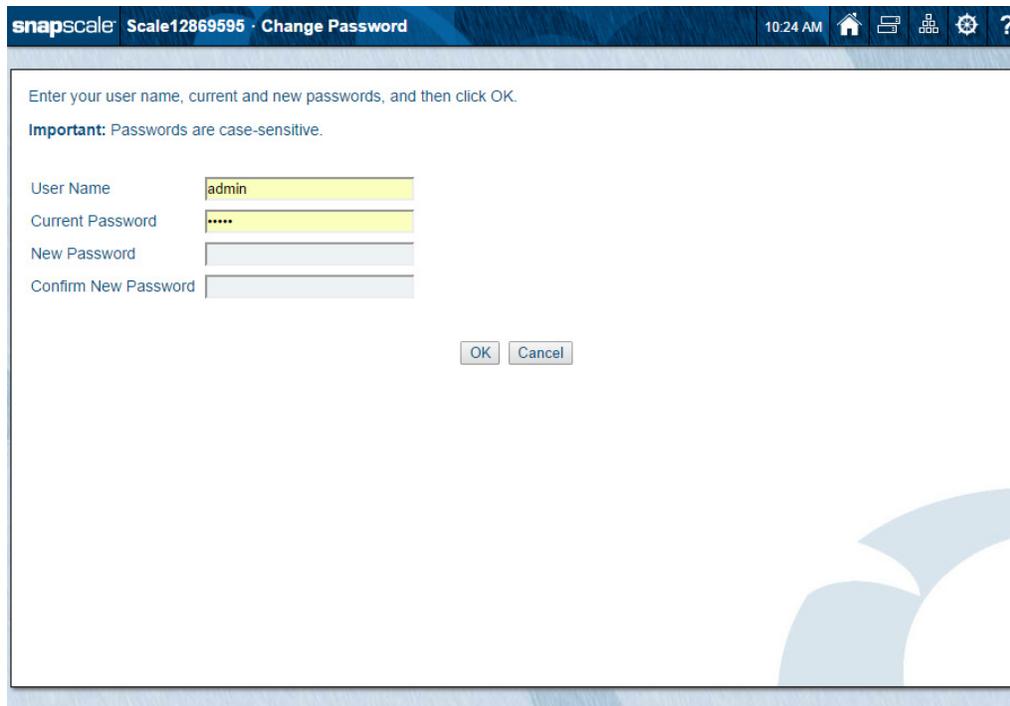
Anyone with administrative privileges can view or edit the Snap Finder properties. Click **Properties** to access the page.

From this page you can choose to display a warning icon for Uninitialized nodes or SnapServers with an enabled Ethernet port that has no link and enable remote discovery of units on a different subnet. Complete the following fields and then click **OK** to return to the **Snap Finder** page:

Option	Description
<b>Display warning if any of a server's Ethernet ports have no link</b>	Check to display a warning icon in the <b>Status</b> column for any nodes or SnapServers that have an enabled Ethernet port with no link. By default, this box is unchecked.
<b>Enable Remote Server Discovery</b>	Check to enable remote discovery of clusters, nodes, or SnapServers on a different subnet.
<b>Add Server</b>	Enter the host name or IP Address of a cluster, node, or server in the field to the right of the <b>Add</b> button, and click <b>Add</b> to incorporate it into the list of Remote Discovery Servers. Remote Discovery Servers send information about themselves as well as all other servers they've discovered on the remote network.
<b>Delete Server</b>	Select a cluster, node, or server, in the Remote Discovery Servers field and click <b>Delete</b> . When asked to confirm the deletion, click <b>Delete</b> again.

## Change Password

To enhance the security of your SnapScale cluster, it is recommended that users change their passwords regularly. This is done using the **Change Password** page.



Enter your user name, current and new passwords, and then click OK.

**Important:** Passwords are case-sensitive.

User Name

Current Password

New Password

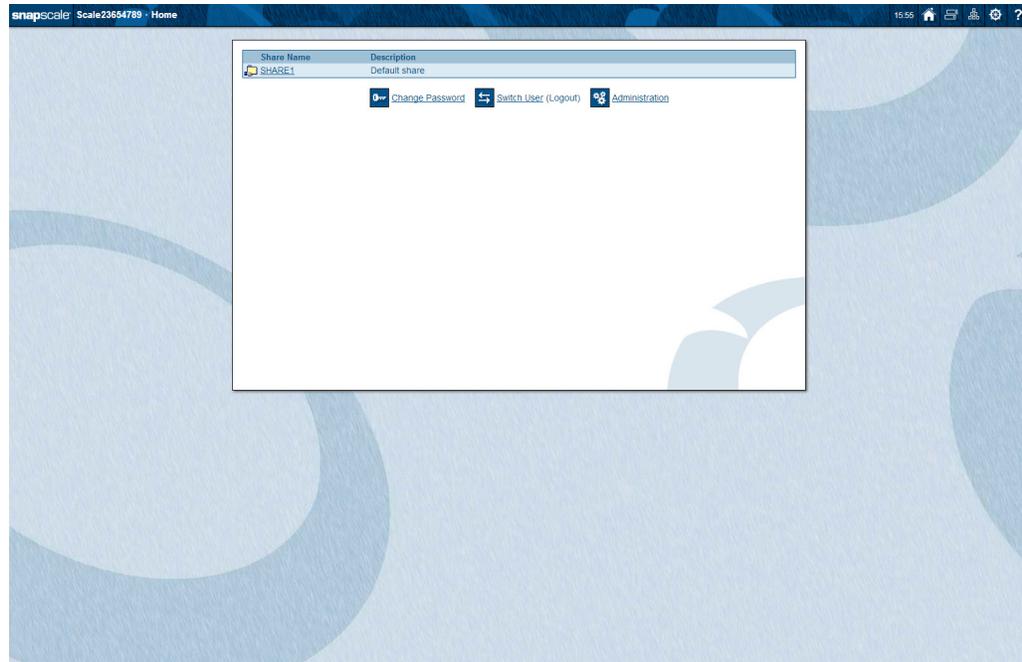
Confirm New Password

### Changing Your Password

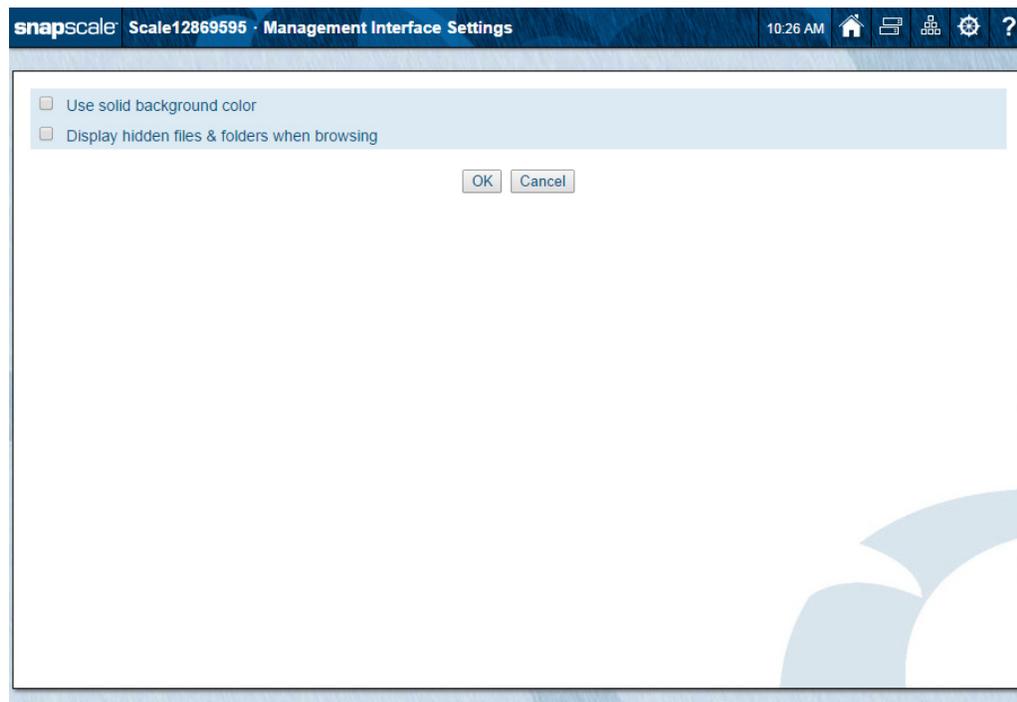
1. On the **Home** page, click the **Change Password** link (🔑).
2. At the **Change Password** page, enter your **User Name** and **Current Password**.
3. Enter and confirm your **new password**.  
Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
4. Click **OK**.
5. At the confirmation page, click **OK** again.  
You are returned to the **Home** page.

## Management Interface Settings

The Web Management Interface default background is light blue with the stylized “O” symbols on a textured blue background:



This can be changed to a solid blue background on the Web Management Interface Settings page by clicking the Site Map icon (⚙️) to access **Management Interface Settings**.



Check the first box to use a solid color background (or clear the box to return to the standard background). Check the second box if you want to display hidden files and folders when browsing volumes for administrative configuration in the Web Management Interface.



This appendix provides a brief description of the supported backup solutions and the Snap Enterprise Data Replicator (Snap EDR) software.

## Topics in Backup Solutions:

- [Backup and Replication Solutions](#)
- [Snap Enterprise Data Replicator](#)
- [Backup via SMB or NFS](#)
- [Off-the-Shelf Backup Solutions](#)

## Backup and Replication Solutions

RAINcloudOS supports several backup methods, including third-party off-the-shelf backup applications and applications that have been customized and integrated with RAINcloudOS on the SnapScale cluster:

- Data and security metadata backup and replication can be performed using the built-in Snap EDR.
- Backup over network file protocols can be performed using various backup packages that can access the cluster via SMB or NFS.
- Backup from the cluster or to a tape attached to a node can be performed using supported backup agents and media libraries installed on a node.

## Snap Enterprise Data Replicator

Snap EDR provides server-to-server synchronization by moving, copying, or replicating the contents of a share from one cluster or server to another share on one or more different clusters or servers. It comes preinstalled on SnapScale clusters and activates a 45-day free trial if configured as a Management Console.

Snap EDR consists of a Management Console and a collection of Agents. The Management Console is installed on a central system. It coordinates and logs the following data transfer activities carried out by the distributed Agents:

- Replicates files between any two systems including SnapServers, SnapScale clusters, and Windows, Linux, or Mac Agents.
- Transfers files from one source host to one or more target hosts
- Transfers files from multiple hosts to a single target host, and stores the files on a local disk or locally attached storage device.
- Backs up data from remote hosts to a central host with locally-attached storage.

- Restores data from a central storage location to the remote hosts from which the data was originally retrieved.

## Snap EDR Usage

The Snap Enterprise Data Replicator software distribution comes preinstalled on the SnapScale but must first be installed in SnapExtensions and then configured before it's available for use. It operates under the cluster management name (in the style of `<clustername>-mgt`) and the Management IP address.

All other Snap EDR installations (including another machine running as the Management Console that the server registers to, other Agents that register to a Snap EDR Management Console running on the server, or other Agents replicating to/from the server) need to be able to resolve the SnapScale cluster name to its IP address in order to interoperate properly with the cluster. This can be accomplished via a DNS host record, local hosts file entries, or other name resolution services in the environment.

When installing EDR to the cluster for the first time, it installs to all nodes and runs on whichever node currently serves as the Management node. When adding a node to an existing cluster with Snap EDR installed, it is automatically installed on that node.

## Configuring Snap EDR

To configure the cluster as a Snap EDR Management Console or an Agent:

1. Click the **SnapExtensions** icon located in the upper right corner of the Web Management Interface.
2. If necessary, install the **software package**:
  - a. Run the **installation routine** from SnapExtensions.  
SnapExtensions displays a Snap EDR link and the status **Not Installed**.
  - b. Click the **Snap EDR link** and confirm the installation.  
Wait for the installation to complete. The **SnapExtensions** page then displays the **Snap EDR Configuration** link.
3. Click the **link** to launch the **Management Console/Agent** configuration page.
4. Select either **Configure as the Management Console** or **Configure as the Agent**.

**NOTE:** If you are configuring a cluster as an Agent, you must provide the server name (for a SnapServer) or cluster management name (for a SnapScale) of the Management Console. The cluster must be able to resolve the Management Console server name to the correct IP address.

5. Once the cluster is configured, select the following **options** from the page that appears:

Option	Description
<a href="#">Click here to configure jobs</a>	Opens the Management Console where jobs can be scheduled.
<b>Stop Service</b>	Stops all services.
<b>Restart Service</b>	Restarts all services.



**CAUTION:** Use only if you have encountered a problem and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service.

## Scheduling Jobs in Snap EDR

To schedule jobs, click the **Snap EDR** link in the **Site Map** (under **Misc.**).

For complete information on scheduling jobs in Snap EDR, see the *Snap EDR Administrator's Guide*.

## Backup via SMB or NFS

A SnapScale cluster can be backed up via standard file server access.

In this configuration, the backup server is set up to use SMB or NFS to connect to the cluster, examine the file system, and then back up the data onto itself. No special agents or media servers are needed.

## Off-the-Shelf Backup Solutions

Special Application Notes for installing the backup agent or media servers can be found on the Overland SnapScale Support website (<http://docs.overlandstorage.com/snapscale>).

**NOTE:** The backup packages shown in the Application Notes do not support the backup of Windows ACLs. If Windows ACL backup is critical, Overland Storage strongly recommends you create a disaster recovery image before you perform a backup.

A SnapScale can be backed up using a third-party Linux agent or media server installed on one cluster node. A tape drive or library can be attached to the node to provide local backup of the cluster.

### Utility IP Address

To provide continuous access to a specific cluster node with a backup agent or media server installed even when a public IP address changes, a special Utility IP address can be assigned that operates in addition to other cluster-assigned IP addresses and is always associated with the specific node. Then, if the public IP address of that node changes, backups continue to function without the need for the administrator to take action.

This IP address setting works for both the node as a backup source (backup from the node to the backup server) and the node as a backup target (backup from any machine including the node to an attached tape drive).

This appendix provides additional information and configuration options about securing and accessing shares and files on the SnapScale. The RAINcloudOS supports share-, file-, and directory-level permissions for all local and Windows domain users and groups.

File and directory security can be configured using either Windows NTFS-style security or classic Unix-style security. The type of security present on a file or directory is its *security personality*.

Files and directories are stored on the cluster on volumes (or the directories underneath) with a configured *security model*. The security model on the volume governs the permitted security personalities, the default personalities, and the ability to change personalities on child files and directories.

The default security model on newly-created volumes is always Windows/Unix. It can be changed to either a Windows or Unix security model.

#### Topics in Shares and File Access:

- [Security Model Rules](#)
- [Security Model Management](#)
- [Special Share Options](#)
- [File and Share Access](#)
- [File-level Security](#)

## Security Model Rules

Files and directories created inside security models acquire the security personality and permissions according to the rules of the chosen security model.

#### Windows/Unix Security Model:

- Files and directories created by SMB clients will have the Windows security personality. Permissions will either be inherited according to the ACL of the parent directory (if Windows) or will receive a default ACL that grants the user full access only (if the parent is Unix or has no inheritable permissions).
- Files and directories created by non-SMB clients will have the Unix personality. Unix permissions will be as set by the client (per the user's local umask on the client).
- The security personality of a file or directory can be changed by any user with sufficient rights to change permissions or ownership. If a client of one security personality changes permissions or ownership of a file or directory of a different personality, the personality will change to match the personality of the client protocol (for example, if an NFS client changes Unix permissions on a Windows file, the file will change to the Unix personality).

#### Windows Security Model:

- All files and directories will have the Windows security personality. Permissions will be inherited according to the ACL of the parent directory.
- The permissions of a file or directory can be changed by any Windows SMB user with sufficient rights to change permissions or ownership. Permissions cannot be changed by NFS or FTP clients.
- The personality of files and directories cannot be changed on a Windows security model. All files and directories always have the Windows personality with a Windows ACL. Standard Unix permissions will appear as 777 (rwxrwxrwx), but only the permissions in the Windows ACL will be enforced.

#### Unix Security Model:

- Files and directories created by non-SMB clients will have the Unix personality. Unix permissions will be as set by the client (per the user's local umask on the client).
- Files and directories created by SMB clients will have the Unix personality. Unix permissions will be set to a default.
- The personality of files and directories cannot be changed on a Unix security model. All files and directories always have the Unix personality.

## Security Model Management

Changes to a security model can optionally be propagated with the corresponding personality and default permission to all files and directories underneath the security model.

When **setting** the security model, only a single security model can be set on the entire volume at the root level but not the directories immediately underneath the volume as they inherit the security model from the top level.

When **changing** the security model:

- If changing from Windows to Unix, all files and directories will be changed to be owned by *admin* and *admingrp*, with Unix permissions of 777(rwxrwxrwx).
- If changing from Unix to Windows, files and directories will be changed to default permissions that allow all users the ability to create and manage their own files and directories and to access other users' files and directories.

## Special Share Options

The basic setup and configuration of shares on a SnapScale cluster is handled on the **Security > Shares** page. This section covers more details about the special options and features of share security in these subsections:

- [Hiding Shares](#)
- [Share Level Permissions](#)
- [Where to Place Shares](#)

## Hiding Shares

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

- Name the share with a dollar-sign (\$) at the end. This is the traditional Windows method of hiding shares; however, it does not truly hide the share since Windows clients themselves filter the shares from share lists. Other protocols can still see dollar-sign shares.
- Hide the share from all protocols (except NFS) by one of these two procedures:
  - While creating a share, navigate to **Security > Shares > Create Share > Advanced Share Properties** and check the **Hide this Share** box.
  - Edit a share by selecting the share, clicking to expand **Advanced Share Properties**, and checking the **Hide this Share** box.

When a share is hidden this way, the share is invisible to clients and must be explicitly specified to gain access.

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

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

## Share Level Permissions

Share-level permissions on GuardianOS are applied cumulatively. For example, if the user *jdoe* has Read-Only share access and belongs to the group *sales*, which has Read/Write share access, the result is that the user *jdoe* will have Read/Write share access.

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

## Where to Place Shares

For security and backup purposes, it is recommended that administrators restrict access to shares at the root of a volume to administrators only. After initialization, all SnapScale clusters have a default share named *SHARE1* that points to the root of the default volume *Volume1*. The share to the root of the volume should only be used by administrators as a “door” into the rest of the directory structure so that, in the event that permissions on a child directory are inadvertently altered to disallow administrative access, access from the root share is not affected. This also allows one root share to be targeted when performing backups. If it is necessary to have the root of the volume accessible, using the Hidden option helps ensure only those that need access to that share can access it.

## File and Share Access

The shares feature also controls access by other users and groups. This section provides information on setting up the shares options to allow proper access to the files.

## NFS Share Access

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



**CAUTION:** If there are multiple shares to the same directory on the disk, and those shares permit access via NFS, they must all have the same NFS export configuration. This is enforced when configuring NFS access to the overlapping shares.

## Snapshot Access

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

## Snapshot Shares and On Demand File Recovery

A *snapshot share* is a read-only copy of a live share that provides users with direct access to versions of their files archived locally on the SnapScale cluster via a snapshot. Users who wish to view or recover an earlier version of a file can retrieve it on demand without administrator intervention.

Snapshot shares are created during the course of creating a share, or thereafter by navigating to the Snapshots page and clicking the name of a snapshot. For instructions on accessing snapshot shares, see [Chapter 6, Security Options](#).

## Creating a Snapshot Share

You create a snapshot share by selecting the **Create Snapshot Share** option on the **Security > Shares > (share\_name) > Share Properties** page, under the **Advanced Share Properties** link.

For example, assume you create a share to a directory called *sales* and you select the **Create Snapshot Share** option. When you connect to the cluster via a file browser or use the **Misc. > Home** link in the Site Map, two shares are displayed:

```
SALES
SALES_SNAP
```

The first share provides access to the live volume and the second share provides access to any archived snapshots. Other than read-write settings (snapshots are read-only), a snapshot share inherits access privileges from its associated live-volume share.

**NOTE:** The same share folders appear on the Home page when you connect to the SnapScale cluster using a Web browser. However, the snapshot share folder does not provide access to the snapshot; it always appears to be empty. You can prevent the snapshot share from displaying on this Home page by selecting the **Hide Snapshot Share** option when creating or editing a share.

## Accessing Snapshots Within the Snapshot Share

A snapshot share contains a series of directories. Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created.

For example, assume the snapshot share named *Sales\_SNAP* contains the following four directories:

```
latest
2014-02-25.120000
2014-03-01.000100
2014-03-07.020200
```

The *latest* directory always points to the most recent snapshot (in this case, **2014-03-07.020200**, or March 7th, 2014, at 2:02 a.m.). A user may view an individual file as it existed at a previous point in time or even roll back to a previous version of the file by creating a file copy to the current live volume.

**NOTE:** The *latest* subdirectory is very useful for setting up backup jobs, as the name of the directory is always the same and always points to the latest available snapshot.

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

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

## File-level Security

RAINcloudOS supports two “personalities” of filesystem security on files and directories:

- **Windows ACLs:** Windows NTFS-style filesystem permissions. Windows ACLs fully support the semantics of NTFS ACLs, including configuration, enforcement, and inheritance models (not including the behavior of some built-in Windows users and groups).
- **Unix:** Traditional Unix permissions (rwx) for owner, group owner, and other.

By default, volumes are created with the Windows/Unix security model (Windows-style ACLs for files created by SMB clients and Unix-style permissions for files created by other protocols and processes), and allow all users to create, delete, and configure permissions on their own files and to access files and directories created by other users.

### Security Personalities and Security Models

The security personality of a file or directory is dependent on the security model of the root directory or volume in which the file or directory exists.

Files and directories in a Windows/Unix security model can have either a Windows or Unix security personality, depending on the network protocol used to create the file or change permissions on it. Files in a Windows security model always have the Windows security personality and permissions can only be set by Windows SMB clients. Files in a Unix security model always have the Unix security personality and permissions can only be set by non-SMB clients.

### Windows ACLs

RAINcloudOS fully supports Windows NTFS-style filesystem ACLs, including configuration, enforcement, and inheritance models. Inside Windows/Unix and Windows security models, files created and managed by Windows clients have the Windows security personality and

behave just as they would on a Windows server. Clients can use the standard Windows Explorer interface to set directory and file permissions for local and Windows domain users and groups on the SnapScale.

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

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

### Default File and Folder Permissions

When a file or directory is created by an SMB client, the owner of the file is the user who created the file (except for files created by local or domain administrators, in which case the owner will be the **Administrators** group, mapped to the local **admingrp**). The ACL is inherited per the inheritance ACEs on the parent directory's ACL. The owner of a file or directory always implicitly has the ability to change permissions, regardless of the permissions established in the ACL. In addition, members of the SnapScale local admin group, as well as members of Domain Admins (if the cluster is configured to belong to a domain) always implicitly have *take ownership* and *change ownership* permissions.

### Setting File and Directory Access Permissions and Inheritance (Windows)

Access permissions for files and directories with the Windows security personality are set using the standard Windows Explorer interface. RAINcloudOS supports:

- All standard generic and advanced access permissions that can be assigned by Windows clients.
- All levels of inheritance that can be assigned to an ACE in a directory ACL from a Windows client.
- Automatic inheritance from parent directories, as well as the ability to disable automatic inheritance from parents.
- Special assignment and inheritance of the CREATOR OWNER, CREATOR GROUP, Users, Authenticated Users, and Administrators built-in users and groups.

Procedure to set file and directory access permissions and inheritance in Windows:

1. Using a Windows client, **map a drive** to the SnapScale cluster, 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.

Basic techniques for identifying and resolving common hardware and networking issues are described here.

### Topics in Troubleshooting SnapScale Nodes

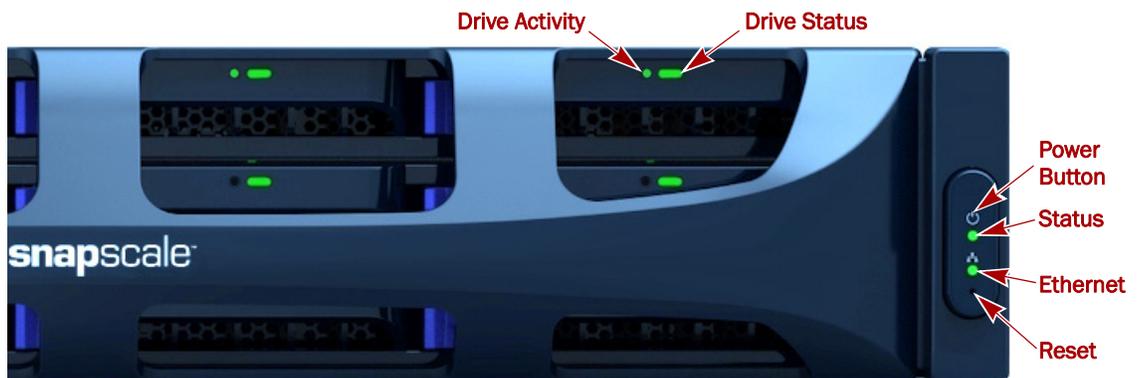
- [LED Indicators](#)
- [Network Reset](#)
- [Networking Issues](#)
- [Miscellaneous Issues](#)
- [Phone Home Support](#)

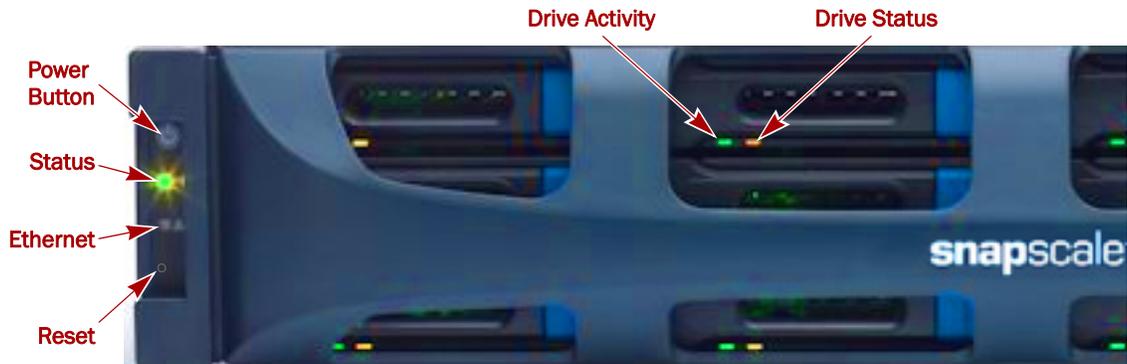
## LED Indicators

LED indicators provide information on the status of basic connectivity, disk drives, fan modules, and power supply modules.

### SnapScale X2 Node LEDs

The SnapScale X2 has one network LED (Ethernet) and one system status LED on the Power Panel located on the right flange, along with a Power button and a Reset button. Each drive has two disk LEDs (Drive Activity and Drive Status) as shown in the following illustration:





The following tables describe the various states that may occur.

## X2 LED States

### Drive Status LED

This is the oblong LED located on the center-right of each disk carrier.

Device State	LED State
No Disk Drive in Carrier	Off
Normal Operation	Solid <b>green</b>
Unit Identification Indicator	Flashing <b>amber</b>
Failed	Flashing <b>red</b>

### Drive Activity LED

This is the round LED located on the center-left of each disk carrier.

Device State	LED State
Powered OFF / No Activity	Off
Drive Activity	Flashing <b>green</b>

### Node Status LED

This is the round LED located just below the Power button on the right flange.

Device State	LED State
Powered OFF	Off
Booting	Solid <b>amber</b>
Normal Operation	Solid <b>green</b>
Shutting down	Flashing <b>green</b>
Maintenance Mode	Flashing <b>green/amber</b>

### Client Network LED

This is the round LED located just below the Status LED on the right flange.

Device State	LED State
Powered OFF	Off
Link Up (SnapScale Powered ON)	Solid <b>green</b>
Link Down	Off

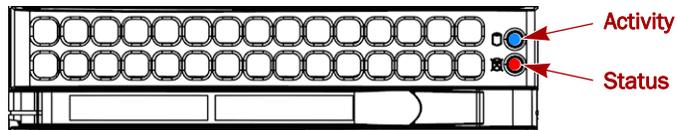
### Power Supply Status LED

This is the LED located just above and to the right of the socket.

Device State	LED State
Normal Operation	Solid <b>green</b>
Standby	Solid <b>red</b>
Power Failure	Off
Fan Failure	Blinking <b>red</b>

### SnapScale X4 Node LEDs

Each drive has two disk LEDs (top is Drive Activity and bottom is Drive Status):



LED	Description
<b>Activity</b>	Solid <b>blue</b> when a drive is present in the carrier. Blinks when there is drive activity.
<b>Status</b>	Red LED that: <ul style="list-style-type: none"> <li>• Solid <b>red</b> when a drive fails.</li> <li>• Blinks <b>red</b> with all the other drive Status LEDs when node ID is activated.</li> <li>• Off if drive is OK or not present.</li> </ul>

The Control Panel on the left flange has a series of LEDs as shown here:



## Performing System Resets Without Network Access

Should Web Management Interface access to the cluster be lost, the **Reset** button can be used to reset cluster settings and reestablish connectivity. This option is also useful should you forget the admin password.

Verify that the cluster is fully booted (as indicated by the system/status LED). Using the end of a straightened paper clip or the fine point of an instrument, press the **Reset** button.

The system will go through a normal shutdown process and then reboot. As a part of the reset and reboot process, the SnapScale does the following:

- Clears user-defined TCP/IP settings such as DHCP configuration.
- Resets the cluster name to its default setting (**Scale**<server\_number>).
- Resets network speed and bonding settings to their defaults.
- Resets the Administrator password to the default (**admin**).
- Resets the web server to allow HTTP access.

## Networking Issues

These are some of the networking issues you may encounter when using your SnapScale cluster.

### The Server Cannot Be Accessed Over 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 cluster and try to connect again.
- Verify that the LED for the primary Ethernet port is lit. (This light indicates network connectivity.) If the light is not lit, perform the following:
  - The most likely cause is the physical connection. Check for a loose or damaged cable, or poor connections in the port connector.
  - This problem may also be caused by a mismatch between the settings on the switch or hub and the settings on the SnapScale Ethernet port. These settings must match. To resolve the problem, make sure the port settings on the hub or switch match the settings for the primary port as configured on the **Network > TCP/IP** page of the Web Management Interface. Use the autonegotiate setting on both the switch and the server port.

### You Have No Access to SnapScale via HTTP

When trying to access the SnapScale via HTTP, the Web browser times out. However, the cluster can be accessed using the ping command or Windows Explorer.

- HTTP and HTTPS are both enabled by default. Try typing HTTPS in the Web address rather than HTTP. If you are able to access the cluster via HTTPS, you can re-enable HTTP on the **Network > Web** page.
- If you cannot access the cluster via HTTPS, try resetting the cluster as described on [Performing System Resets Without Network Access on page 258](#).

### 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 SnapScale will authenticate the users as local SnapScale 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 SnapScale Does Not Operate Properly on a Network Running Gigabit-Full-Duplex

For Gigabit Ethernet to operate properly, both the switch and the SnapScale 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 SnapScale IP Address Is Unknown

Install SnapServer Manager (available from the Overland Storage website) onto a client workstation on the same subnet as the SnapScale. You can then use the utility to discover all SnapServers, SnapScale clusters, and Uninitialized nodes on that network segment, and to assign static IP addresses as necessary.

### Problems Occur with Domain Controller Authentication

You are receiving the following errors in your error log:

```
SMB: Domain Controller unavailable
```

```
SMB: Username not connected to Domain Controller
```

This means that either your Domain Controller is down, or the SnapScale 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 SnapScale but Cannot See It on Network

Ensure that the Ethernet cable is connected securely to both the network port and the primary Ethernet port. Also, check to see that the Link light on the front of the SnapScale 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 SnapScale. To resolve this problem, verify that all settings (if using multiple Ethernet ports) on the switch/hub match the setting on the cluster. When a node 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 cluster.

A SnapScale cluster is configured by default to acquire an IP address from a DHCP server. If no DHCP server is found on the network, the SnapScale 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 SnapScale using either the default cluster name or running the SSM utility.

### You Try to Mount to a Share on Your SnapScale 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 SnapScale from a Linux Workstation

If you are logged in as **root** on your workstation and NFS is enabled on your SnapScale, 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 following the IP address before the share. This returns an Access Denied message. The correct syntax should be the following (added slash shown in **red**):

```
mount 192.168.32.124: /SHARE1 /mnt
```

**NOTE:** The share name is case-sensitive.

### You Cannot Log in as Root to SnapScale

RAINcloudOS allows you to log in as root over SMB. If this operation has failed or you have trouble logging in, be sure that you have enabled root login in the **Network > Windows/SMB** page. Also note that the root account password is tied to the admin account password. If you cannot log in as root, change the password for the admin account on the **Network > Windows/SMB** page. Use the admin password to log in as root.

### You Are Unable to See Your Domain Users When Trying to Set Up Windows Security Permissions on File Folders

The SnapScale has joined the Active Directory domain properly but you cannot see the domain users when you set share permissions from the 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 SnapScale.

### You Backed Up Your Snapshot Share, Are Now Attempting to Restore It, and Operation Fails

A snapshot share is read-only. You can only restore the data to a read-write accessible share.

### Power to the SnapScale Is Unexpectedly Cut Off Due to a Power Outage

Overland Storage recommends that you use an uninterruptible power supply (UPS) with the SnapScale. If you did not have a UPS attached to the cluster at the time of the power outage, do the following:

1. Remove the power cables.
2. Once the power is restored and stabilized, turn the power supplies back on and reboot the cluster.

Once the SnapScale boots, it begins resynchronizing the Peer Sets if necessary. You can use the cluster during the resynchronization, but performance will be a little slower than normal. Do not remove drives, however, while the cluster is resynchronizing the Peer Sets.

### The Server Is Not Responding to File Requests or Configuration Commands

Call your technical support representative.

### You Have Problems Seeing the Tape Library Tape Device, Not the Robotic Arm

When you have problems seeing the actual tape device rather than the robotic arm, it is most likely due to the Tape Loader being configured for Sequential Access. Change the Tape Loader to Random or Mixed Mode.

### The Admin Password to the Web Management Interface Is Not Available

You can perform a limited reset to defaults, which includes the admin password, then use the Web Management Interface to set a new password. See [Performing System Resets Without Network Access on page 258](#).

## Phone Home Support

Once your SnapScale has been registered, Phone Home Support becomes available for use. Phone Home Support emails system logs and files that contain information useful for troubleshooting purposes to Overland Storage Technical Support. You can use the **Maintenance > Support** page to open a new case with technical support; or, in the course of working to resolve an issue, a technical support representative may ask you to fill out and submit this page. If a case is already in progress, you will need to enter the case number provided by the technical support representative.

**NOTE:** Phone Home Support interacts with two fields on the **Maintenance > Tools > Email Notification** page. To use Phone Home Support, you must first enter a valid SMTP server IP address on the Email Notification page. The first email address listed in the Recipients field automatically populates the Reply-to Address field on the **Support** page.

Complete the following fields as appropriate, then click **OK**:

Text Field	Description
<b>Subject*</b>	Enter a concise description that identifies the issue.
<b>Case*</b>	<ul style="list-style-type: none"> <li>Select <b>New Case</b> if you are emailing technical support for the first time.</li> <li>Select <b>Existing Case</b> if you have previously contacted technical support concerning the issue.</li> </ul>
<b>Case Number</b>	If you selected <b>Existing Case</b> above, enter the case number provided by technical support.
<b>Reply-to Address*</b>	<p>This field defaults to the first email address entered as a recipient on the <b>Maintenance &gt; Tools &gt; Email Notification</b> page. If necessary, enter at least one email address that will serve as the contact email address for this issue.</p> <p>To receive a copy of the email and system information attachment, check the <b>Cc Reply-to Address</b> box.</p>
<b>Comments*</b>	Enter additional information that will assist in the resolution of the problem.

\*Required option.

## Advanced Help

If you have an open case and have entered the **Case Number**, clicking **Advanced Support Properties** opens additional options for the phone home feature. These options provide the ability to upload specific log files via FTP, which is sometimes necessary for the large logs the cluster can generate, and tech support may direct a user to use this for a particular case.

Send configuration details about this SnapScale to Overland Storage Technical Support. Enter your information below and click OK.

Subject:

Case:  ▼

Case Number:

Reply-to Address:

Cc Reply-to Address

Comments (3951 characters left):

### Advanced Support Properties

Advanced support properties allow you to upload detailed information about specific nodes in your SnapScale to Overland's Technical Support FTP server. You can either select specific nodes or you can specify only those nodes that relate to a specific file or directory.

**Overland Technical Support FTP settings:**

FTP Server:

FTP Path:

FTP User:

FTP Password:

Upload only information about nodes related to a specific file or directory.

Upload information about specific nodes.

Please select at least one node. Use Ctrl-click to select multiple nodes. Use Command-click for Mac. [Select All](#) [Select None](#)

VM-Node13157556 - 10.25.11.163 - (No description.)

VM-Node179284 - 10.25.11.162 - (No description.)

VM-Node7846949 - 10.25.11.161 - (No description.)

The additional options include:

- **FTP Server** – Name of the server to upload to.
- **FTP Path** – FTP server path used to upload to.
- **FTP User** – Name of the user to log in to the FTP server.
- **FTP Password** – FTP user's password.
- Click an **upload option**:
  - **Upload only information about nodes related to a specific file or directory**  
Use the **Share** drop-down list to select the share and **File or Directory** path field to enter the path to a file or directory to gather logs. Only select (**Use absolute path.**) from the drop-down list when directed by Overland Support.

Upload only information about nodes related to a specific file or directory.  
Please specify a share and path to a file or directory under the share. Select "Use absolute path" only as directed by Overland Support.  
Share   
File or directory   
 Upload information about specific nodes.

- **Upload information about specific nodes**

Select one or more nodes to upload logs from them. Use the **Select All** and **Select None** options on the right as needed.

Upload only information about nodes related to a specific file or directory.  
 Upload information about specific nodes.  
Please select at least one node. Use Ctrl-click to select multiple nodes. Use Command-click for Mac. [Select All](#) [Select None](#)  
VM-Node13157556 - 10.25.11.163 - (No description.)  
VM-Node179284 - 10.25.11.162 - (No description.)  
VM-Node7846949 - 10.25.11.161 - (No description.)

# RAINcloudOS Ports

The following table lists the ports used by RAINcloudOS. The **ROS Feature** column lists access to the feature such as **Storage > iSCSI** (which you can access under the **Storage** tab in the **iSCSI** subsection of the Web Management Interface).

Port #	Layer	ROS Feature	Name	Comment
1	DDP		rtmp	Routing Table Management Protocol
1	TCP & UDP		tcpmux	TCP port service multiplexer
2	DDP		nbp	Name Binding Protocol
22	TCP & UDP	SnapScale > SSH	ssh	Secure Shell (SSH) service
25	TCP & UDP	SnapScale > 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
111	TCP & UDP	<ul style="list-style-type: none"> <li>• Networking &gt; NFS</li> <li>• Assist</li> <li>• SnapServer Manager</li> </ul>	sunrpc	Remote Procedure Call (RPC) Protocol for remote command execution, used by Network Filesystem (NFS) and SnapServer Manager
123	TCP & UDP	SnapScale > Date/Time > Advanced	ntp	Network Time Protocol (NTP)
137	TCP & UDP	Network > Windows/SMB	netbios-ns	NETBIOS Name Services used in Red Hat Enterprise Linux by Samba
138	TCP & UDP	Network > Windows/SMB	netbios-dgm	NETBIOS Datagram Services used in Red Hat Enterprise Linux by Samba
139	TCP & UDP	Network > Windows/SMB	netbios-ssn	NETBIOS Session Services used in Red Hat Enterprise Linux by Samba
161	TCP & UDP	Network > Windows/SMB	snmp	Simple Network Management Protocol (SNMP)
162	TCP & UDP	Network > Windows/SMB	snmptrap	Traps for SNMP
389	TCP & UDP	Network > Windows/SMB	ldap	Lightweight Directory Access Protocol (LDAP)

Port #	Layer	ROS Feature	Name	Comment
443	TCP & UDP	<ul style="list-style-type: none"> <li>• Web Management Interface</li> <li>• SnapServer Manager</li> <li>• SnapExtensions &gt; Snap EDR</li> </ul>	https	Secure Hypertext Transfer Protocol (HTTP).
445	TCP & UDP	Network > Windows/SMB	microsoft-ds	Server Message Block (SMB) over TCP/IP
852	TCP	Network > NFS		Used by rpc.mountd
882	UDP	<ul style="list-style-type: none"> <li>• Snap Finder</li> <li>• SnapServer Manager</li> </ul>	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
2005	TCP	SnapExtensions	SnapExtensions	Bridge from Servlet to Snap Extension framework
2049	TCP & UDP	Network > NFS	nfs [nfsd]	Network Filesystem (NFS)
2050	UDP	Network > NFS	mountd	
2051	UDP	Network > NFS	lockd	
2599	UDP	<ul style="list-style-type: none"> <li>• Snap Finder</li> <li>• SnapServer Manager</li> </ul>	Sysbroker	Multicast Discovery
3052	TCP	SnapScale > UPS		Port for monitoring UPS status
3205	TCP	Network > iSNS	iSNS	iSNS port
3260	TCP	Storage > iSCSI	iSCSI	iSCSI port
8001	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8002	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8003	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8005	TCP	Web Management Interface	tomcat	Tomcat Shutdown port
8008	TCP & UDP	Web Management Interface	http-alt	Tomcat - Apache Bridge
9049	TCP	Sysbroker		Sysbroker Shutdown Port
9050	TCP	Sysbroker		Sysbroker RPC Port
10001	TCP	Snap Extension	Snap Extension	Shutdown Port
32780	TCP	Web Management Interface	tomcat	Random Port
32781	TCP	Web Management Interface	tomcat	Random Port
49221	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port
49229	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port

Port #	Layer	ROS Feature	Name	Comment
1024 - 65535	TCP & UDP	Network > NFS Network > FTP	NFS FTP (Passive)	Dynamically allocated in runtime for user connections

# Command Line Interface

RAINcloudOS includes a command line interface (SnapCLI) that is accessible through SSH. Using the CLI, users can access information about most of the SnapScale configuration parameters and perform configuration and maintenance functions without using the Web Management Interface or SSM.

---

 **IMPORTANT:** Some administrative tasks must still be performed using the Web Management Interface. The CLI is intended as a convenient way to perform some functions; it is not intended as an alternative to using the Web Management Interface.

---

## 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	Result
<b>?</b>	See an overview of the CLI, with a list of available commands and a description of command syntax.
<i>{command}</i> <b>help</b>	See a description of that particular command's function and a list of options available for the command.
<b>tab</b>	Finish the command you have started to type (such as, tab-complete).
<i>{command}</i> <b>tab</b>	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 <code>date timezones</code> to get a list of timezones)

In this instance, to set the date to October 27, 2011, enter:

```
date set day=27 month=10 year=2011
```

**NOTE:** If, instead of typing the word `date`, you had typed `d + [tab]`, the word would have been completed for you. If you entered `d + [tab] + [tab]`, the word would have been completed and the available options displayed.

Suppose, instead of `date`, you typed the command `web`. Two arguments would be available, one with options:

Command	Arguments and Options	Descriptions
web	get	- get WEB properties
	set [OPTIONS]	- set WEB properties
	- require-webview-auth=(yes no)	- require HTTP/HTTPS clients to authenticate in order to access the server
	- non-secure-http=(yes no)	-enable/disable non-secure HTTP access

Thus, the following command string:

```
web set require-webview-auth=yes non-secure-http=no
```

sets HTTP/HTTPS properties on the SnapScale cluster to require clients to authenticate in order to access the cluster and to disable non-secure HTTP access.

## SnapCLI Procedures

Use these procedures to access and exit SnapCLI.

### Logging into SnapCLI

1. Make sure your client has an SSH v2 client application installed.

**NOTE:** Free or low-cost SSH applications are available from the Internet.

2. Connect to a node using its IP address and log in as *admin* (or any other member of *admingrp*).

You will automatically be placed in the CLI shell.

**NOTE:** SSH v2 is required. If you fail to connect to the node, 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
clear	Clear the page
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)
domain get	Get the domains and their properties known to the SnapScale cluster
domain list	List the domains known to the SnapScale cluster
email get	Get email notification settings
email set	Set email notification settings
event clear	Clear all events in the System Event Log
event get	Display the System Event Log
exit	Quit the command line, log off, and exit ssh/bash session. <b>NOTE:</b> If user has started another shell, this command returns them to the command line shell.
globalspares add	Add a disk to the global spares list
globalspares list	List global spares
globalspares remove	Remove a disk from the global spares list
group create	Create a local group
group delete	Delete a local group
group get	Get available groups with their associated information
group list	List available groups
group 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

Command	Description
group set	Change the properties of a local group
history	Print the history of commands typed into the SnapScale command line
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 list	List all host file entries
hostfile set	Set information for a specific host file entry
idmap auto map	View/Save auto-generated ID mappings
idmap count	Count number of ID mappings
idmap group get	Get ID mapping for a (windows domain) group
idmap group remove	Remove ID mapping for a (windows domain) group
idmap group set	Set ID mapping for a (windows domain) group to a local or NIS group
idmap list	List all ID mappings
idmap remove all	Remove all ID mappings
idmap update files	Update filesystem for ID mapping changes
idmap update status	View status of ID mapping update filesystem operation
idmap user get	Get ID mapping for a (windows domain) user
idmap user remove	Remove ID mapping for a (windows domain) user
idmap user set	Set ID mapping for a (windows domain) user to a local or NIS user
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
ldap get	Displays LDAP network settings
less	With a file name, this command allows the user to view any file on the system. It should only be used for text files.
name get	Get the name of the SnapScale cluster
name set	Set the name of the SnapScale cluster
netinfo	Get information about the Ethernet interface
nfs get	Get SnapScale cluster NFS Properties
nfs set	Set SnapScale cluster 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

Command	Description
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
proxy get	Display the HTTP proxy properties
proxy set	Set the HTTP proxy properties
quit	Quit the command line, log off, and exit the ssh/bash session
quota get	Get 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 list	List user or group quotas for a volume
quota set	Set quota settings for a volume
quota user get	Get volume quota limit & usage for a specific user
quota user set	Set volume quota limit & usage for a specific user
reboot	Reboot the SnapScale cluster
securitymodel get	Get the security model on a SnapScale Volume
securitymodel set	Set the security model on a SnapScale Volume
share access delete	Delete access permission of the specified user/group for the share
share access get	Get access list for the share
share access set	Set access list for the share
share create	Create a share
share delete	Delete a share
share get	View a share
share list	List available shares
share nfsaccess delete	Delete NFS access permission of the host for the specified 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 rename	Rename a share
share set	Modify a share
slidingwindow get	Get sliding window settings for a specific interface
slidingwindow list	List sliding window settings for all interfaces
slidingwindow set	Set sliding window settings for a specific interface
snapex	Perform a control operation on the snap extension

Command	Description
snapshot create later	Create a new snapshot schedule
snapshot create now	Create a new one time snapshot to be run immediately
snapshot delete	Delete specified snapshot
snapshot get	Get snapshot properties
snapshot list	Get list of snapshots
snapshot pool get	Get snapshot pool properties
snapshot pool list	List current snapshot pools
snapshot pool set	Set snapshot pool properties
snapshot rollback	Start a rollback for the specified snapshot
snapshot sched delete	Delete specified snapshot schedule
snapshot sched get	Get specified snapshot schedule
snapshot sched list	List current snapshot schedules
snapshot sched set	Set specified snapshot schedule
snapshot set	Set properties for the specified snapshot
syslog all	Create a tar file of syswrapper and all third-party logs
syslog edr	Create a tar file of Snap EDR logs
syslog s2s	Create a tar file of S2Sv2 logs
syslog syswrapper	Create a tar file of syswrapper only
tcpip create bond	Create a bond and set TCP/IP properties.
tcpip delete bond	Remove a TCP/IP bond.
tcpip get	Get TCP/IP parameters
tcpip set	Set TCP/IP parameters.  <b>NOTE:</b> 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.
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 lock	Lock the specified user.
user set	Change the properties of a local user
user unlock	Unlock the specified user.
version	Display current version information, including the Server Number.  <b>NOTE:</b> This is the same information displayed in the Web Management Interface "About" box
volume create	Create a new logical volume

Command	Description
volume delete	Delete a logical volume
volume edit	Edit an existing logical volume
volume get	Get a specific volume's properties
volume list	List of the volumes defined on the SnapScale cluster
volume write-cache	Enable or disable write cache on a volume.
web get	Get current HTTP Web access settings
web set	Enable or Disable HTTP access to Web Management Interface
windows domain-filter	Allows the user to limit the number of domain users and groups that the server imports from the Active Directory domain, by restricting the import to a specific organizational unit as well as to specific domain users and groups.

## Scripts in SnapCLI

Administrative tasks can be automated with shell scripts that call SnapCLI commands.

### Running a SnapCLI Script

1. Create the script and put it in a share on the local server.
  - Be sure to use an application that is compatible with the standard Unix text file format (for example, vi). Avoid using Windows clients to create or edit scripts.
  - Place the script in a share that will never be part of a delete script.
2. Log in to the SnapCLI (see [Logging into SnapCLI on page 268](#) for instructions).
3. Type `osshell` to get a bash prompt (#).
4. At the prompt, make sure the script is executable by typing the following and pressing **Enter**:

```
chmod +x/shares/[sharename]/[scriptname]
```

 where `sharename` is the name of the share where you put the script and `scriptname` is the name of the script.
5. To run the script, type the path again and press **Enter**:

```
/shares/[sharename]/[scriptname]
```

### Sample Script

Following is an example script that can be used to create and remove users, groups, and shares:

```
#!/bin/sh

#####
# Copyright 2003-2007 Overland Storage, Inc. All rights reserved. #
# Permission is granted to use this code provided that it #
# retains the above copyright notice. ##
#####
CLI=/bin/cli
USER=myuser
PASSWORD=myuserpass
GROUP=mygroup
```

```
SHARE=myshare  
VOLUME=VOL0
```

```
# usage: 'mkuser <user_name> <password>'  
mkuser()  
{
```

### Create a User

```
# if the user does not exist then create it  
if ! $CLI user get user-name="$1" > /dev/null 2>&1; then  
echo "Creating user '$1' ..."  
$CLI user create user-name="$1" password="$2" > /dev/null 2>&1  
if [ $? -ne 0 ]; then  
echo "Creation of user '$1' failed."  
return 1  
fi  
else  
echo "User '$1' already exists."  
fi  
  
return 0  
}
```

```
# usage: 'mgroup <group_name>'  
mgroup()  
{
```

## Create a Group

```
# if the group does not exist then create it
if ! $CLI group get group-name="$1" > /dev/null 2>&1; then
    echo "Creating group '$1' ..."
    $CLI group create group-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Creation of group '$1' failed."
    fi
else
    echo "Group '$1' already exists."
fi

return 0
}
```

```
# usage: 'adduser2group <user_name> <group_name>'
adduser2group()
{
```

## Add the User to the Group

```
# if both the user and the group exist add the user as a member of this group
if $CLI user get user-name="$1" > /dev/null 2>&1; then
if $CLI group get group-name="$2" > /dev/null 2>&1; then
    echo "Adding user '$1' to group '$2' ..."
    $CLI group member add user-name="$1" group-name="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Adding user '$1' to group '$2' failed."
    fi
else
    echo "Group '$2' does not exist."
fi
else
    echo "User '$1' does not exist."
fi

return 0
}
```

```
# usage: 'mkshare <share_name> <share_volume>'
mkshare()
{
```

## Create a Share

```
# if the share does not exist create it
if ! $CLI share get share-name="$1" > /dev/null 2>&1; then
    echo "Creating share '$1' ..."
    $CLI share create share-name="$1" share-volume="$2" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Creating share '$1' failed."
    fi
else
    echo "Share '$1' already exists."
fi

return 0
}
```

```
# usage: 'rmuser <user_name>'
rmuser()
{
```

## Delete the User

```
# if the user exists then delete it
if $CLI user get user-name="$1" > /dev/null 2>&1; then
    echo "Deleting user '$1' ..."
    $CLI user delete user-name="$1" > /dev/null 2>&1
else
    echo "User '$1' does not exist."
fi
```

```

        if [ $? -ne 0 ]; then
            echo "Deletion of user '$1' failed."
        fi
    return 1
    fi
else
    echo "User '$1' does not exist."
fi
return 0
}

# usage: 'rmgroup <group_name>'
rmgroup()
{

```

### Delete the Group

```

# if the group exists then delete it
if $CLI group get group-name="$1" > /dev/null 2>&1; then
    echo "Deleting group '$1' ..."
    $CLI group delete group-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of group '$1' failed."
    fi
return 1
fi
else
    echo "Group '$1P' does not exist."
fi
return 0
}

# usage: 'rmshare <share_name>'
rmshare()
{

```

### Delete the Share

```

# if the share exists delete it
if $CLI share get share-name="$1" > /dev/null 2>&1; then
    echo "Deleting share '$1' ..."
    $CLI share delete share-name="$1" > /dev/null 2>&1
    if [ $? -ne 0 ]; then
        echo "Deletion of share '$1' failed."
    fi
return 1
fi
else
    echo "Share '$1' does not exist."
fi
return 0
}

```

### Create a User, Group, and Share; Then Add the User to the Group

```

#####
#   Main   #
#####

# create a user, a group and a share and add the user to the group
mkuser "$USER" "$PASSWORD"
mkgroup "$GROUP"
adduser2group "$USER" "$GROUP"
mkshare "$SHARE" "$VOLUME"

#remove the group, the user and the share
rmgroup "$GROUP"
rmuser "$USER"
rmshare "$SHARE"

```

# Master Glossary & Acronym List

---

**NOTE:** This is a general Overland Storage glossary and acronym list. Not all items may be found in this document or be used by this product.

## 1000BASE-T

1000BASE-T (also known as IEEE 802.3ab) is a standard for gigabit Ethernet over copper wiring. It requires, at a minimum, Category 5 cable (the same as 100BASE-TX), but Category 5e (Category 5 enhanced) and Category 6 cable may also be used and are often recommended. 1000BASE-T requires all four pairs to be present and is far less tolerant of poorly installed wiring than 100BASE-TX.

## Address

An address is a data structure or logical convention used to identify a unique entity, such as a particular process or network device.

## Algorithm

A sequence of steps designed to solve a problem or execute a process.

## ATA

Short for *Advanced Technology Attachment*. A standard interface for connecting storage devices to a PC.

## Authentication

The validation of a user's identity by requiring the user to provide a registered login name and corresponding password.

## Autonegotiation

An Ethernet feature that automatically negotiates the fastest Ethernet speed and duplex setting between a port and a hub or switch. This is the default setting and is recommended.

## Autosensing

An Ethernet feature that automatically senses the current Ethernet speed setting.

## Bar Code

The machine-readable representation of a product code. Bar codes are read by a scanner that passes over the code and registers the product code. The width of black lines and white spaces between varies. Combinations of lines and spaces represent characters. Overland uses 3-of-9 code (Code 39) where each character is represented by 9 bars, 3 of which are wide.

## Bus or Channel

A common physical path composed of wires or other media, across which signals are sent from one part of a computer to another. A channel is a means of transferring data between modules and adapters, or between an adapter and SCSI devices. A channel topology network consists of a single cable trunk that connects one workstation to the next in a daisy-chain configuration. All nodes share the same medium, and only one node can broadcast messages at a time.

## CA

Short for *Certificate Authority*. A trusted third-party in a network that issues and manages security credentials.

## Cat 5 Cable

Short for *Category 5*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors. CAT 5 cabling supports frequencies up to 100 MHz and speeds up to 100 Mbps. It can be used for ATM, token ring, 100BASE-T, and 10BASE-T networking.

Cat 5 is based on the EIA/TIA 568 Commercial Building Telecommunications Wiring Standard developed by the Electronics Industries Association as requested by the Computer Communications Industry Association in 1985.

## Cat 6 Cable

Short for *Category 6*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors made to higher standards that help reduce noise caused by crosstalk and system noise. The ANSI/TIA-568-B.2-1 specification states the cable may be made with 22 to 24 AWG gauge wire, so long as the cable meets the specified testing standards.

It is designed for Gigabit Ethernet that is backward compatible with the Category 5/5e and Category 3 cable standards. Cat 6 features more stringent specifications for crosstalk and system noise. The cable standard provides performance of up to 250 MHz and is suitable for 10BASE-T, 100BASE-TX, and 1000BASE-T (Gigabit Ethernet).

## Channel

A communications path between two computers or devices.

## Checksum

The result of adding a group of data items that are used for checking the group. The data items can be either numerals or other character strings treated as numerals during the checksum calculation. The checksum value verifies that communication between two devices is successful.

## CIFS

Short for *Common Internet Filesystem*. Also known as [SMB](#). The default Windows protocol for communication between computers. A specification for an Internet file access protocol that complements HTTP and FTP.

## daemon

A process that runs in the background.

## default gateway

The router used when there is otherwise no known route to a given subnet.

**DHCP**

Short for *Dynamic Host Configuration Protocol*. A communications protocol that lets network administrators centrally manage and automate the assignment of IP addresses on a computer network. Each system that connects to the Internet/intranet needs a unique IP address.

**Disaster Recovery**

A strategy that allows a company to return to normal activities after a catastrophic interruption. Through failover to a parallel system or by restoration of the failed system, disaster recovery restores the system to its normal operating mode.

**DNS**

Short for *Domain Name Service*. A network service that translates domain names into IP addresses using a server that maintains a mapping of all host names and IP addresses. Normally, this mapping is maintained by the system administrator, but some servers support dynamic mappings.

**Domain**

A set of network resources in Windows 2000/2003/2008, such as users and groups of users. A domain may also include multiple servers on the network. To gain access to these network resources, the user logs into the domain.

**Domain Name**

The ASCII name that identifies the domain for a group of computers within a network.

**Ethernet**

The most widely installed LAN technology. 100BASE-T Ethernet provides transmission speeds of up to 100 Mbps. Fast Ethernet or 1000BASE-T provides transmission speeds up to 1000 Mbps and is typically used for LAN backbone systems, supporting workstations with 100BASE-T cards. Gigabit Ethernet (GbE) provides an even higher level of backbone support at 1000 Mbps (one Gigabit or one billion bits per second).

**Ethernet Address**

The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet interface.

**Ethernet Port**

The port on a network card to provide Ethernet access to the computer.

**Event**

Any significant occurrence or error in the system that may require notifying a system administrator or adding an entry to a log.

**Expansion Slot**

Area in a computer that accepts additional input/output boards to increase the capability of the computer.

**Failover**

A strategy that enables one Ethernet port to assume the role of another port if the first port fails. When the port comes back online, the original identities are restored. Failover is possible only in a multi-Ethernet configuration.

**Failover/Failback**

A combination of Failover and Failback. When a preferred path becomes unavailable, another path is used to route I/O until the preferred path is restored. In this case I/O will “fail back” to the preferred path once it is available again.

**Fibre Channel**

Fibre Channel (FC) is a gigabit-speed network technology which transports SCSI commands over Fibre Channel networks. Fibre Channel was primarily concerned with simplifying the connections and increasing distances, but later designers added the goals of connecting SCSI disk storage, providing higher speeds and far greater numbers of connected devices.

**Firmware**

Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first switched on.

**FTP**

Short for *File Transfer Protocol*. A standard Internet protocol that provides a way to exchange files between computers on the Internet.

**Full-duplex**

A type of transmission that allows communicating systems to both transmit and receive data simultaneously.

**Gateway**

The hardware or software that bridges the gap between two network subnets. It allows data to be transferred among computers that are on different subnets.

**Gigabit Ethernet**

Also known as GigE or GbE, this Ethernet standard uses a one Gigahertz (1000 Hz) clock rate to move data.

**HBA**

Short for *Host Bus Adapter*. An HBA is an I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer of information between the two channels. In order to minimize the impact on host processor performance, the HBA performs many low-level interface functions automatically or with minimal processor involvement.

**Half-duplex**

A type of transmission that transfers data in one way at a time.

**Hidden Share**

A share that restricts the display of the share via the Windows (SMB), Web Home (HTTP/HTTPS), FTP, and AFP protocols. See also [SMB](#).

**Host Name**

The unique name by which a computer is known on a network. It is used to identify the computer in electronic information interchange.

## Hot Swapping

The ability to remove and add disk drives to a system without the need to power down or interrupt client access to filesystems. Not all components are hot-swappable. Please read installation and maintenance instructions carefully.

## HTTP

Short for *Hypertext Transfer Protocol*. An application protocol for transferring files (text, graphic images, sound, video, and other multimedia files) over TCP/IP on the World Wide Web.

## HTTPS

Short for *Hypertext Transfer Protocol Secure*. The HTTP protocol using a Secure Sockets Layer (SSL). SSL provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

## Inheritance

In Windows permissions, inheritance is the concept that when permissions for a folder are defined, any subfolders within the defined folder inherit its permissions. This means an administrator need not assign permissions for subfolders as long as identical permissions are desired. Inheritance greatly reduces administrative overhead and also results in greater consistency in access permission management.

## Initiator Device

An iSCSI system component that originates an I/O command over an I/O bus or network. An initiator issues the commands; a *target* receives them.

An initiator normally runs on a host computer. It may be either a software driver or a hardware plug-in card, often called a Host Bus Adapter (HBA). A software initiator uses one of the computer's Ethernet ports for its physical connection, whereas the HBA will have its own dedicated port.

Software initiators are readily available for most host operating systems. Hardware initiators are not widely used, although they may be useful in very high performance applications or if 10 Gigabit Ethernet support is required.

## I/O (Input/Output)

The operation of transferring data to or from a device, typically through an interface protocol like CIFS, NFS, or HTTP.

## IP

Short for *Internet Protocol*. The unique 32-bit value that identifies the location of the server. This address consists of a network address, optional subnetwork address, and host address. It displays as four addresses ranging from 1 to 255 separated by periods.

## IQN

Short for *iSCSI Qualified Name*. A name format used in the iSCSI protocol. Initiators and targets have IP addresses, just like any other network entity. They are also identified using an iSCSI name, called the iSCSI Qualified Name (IQN). The IQN should be unique worldwide. It is made up of a number of components, specifying the date, identifying the vendor in reverse format, and then uniquely identifying the initiator or target. An example of an IQN is:

```
iqn.2001-04.com.example:storage:diskarray-sn-123456789
```

Since these IQNs are rather unwieldy, initiators and targets also use short, user friendly names (sometimes called alias names or just aliases).

## iSCSI

Short for *Internet SCSI*. iSCSI is an IP-based storage networking standard for linking data storage facilities. iSCSI is a standard that defines the encapsulation of SCSI packets in TCP and then routing it using IP. It allows block-level storage data to be transported over widely used IP networks.

## iSNS Server

Short for *Internet Storage Name Service Server*. A protocol enabling the automatic discovery, configuration, and management of iSCSI devices on a TCP/IP network.

## Kerberos

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

In Windows 2000/XP, the domain controller is the Kerberos server. The Kerberos key distribution center (KDC) and the origin of group policies are applied to the domain.

## LAN

Short for *Local Area Network*. A network connecting computers in a relatively small area such as a building.

## LCD

Short for *Liquid Crystal Display*. An electronic device that uses liquid crystal to display messages.

## LED

Short for *Light-Emitting Diode*. An LED is a type of diode that emits light when current passes through it. Visible LEDs are used as indicator lights on electronic devices.

## Linux

A UNIX-like operating system that was designed to provide personal computer users a free or very low-cost operating system comparable to traditional and usually more expensive UNIX systems.

## Load Balancing

A process available only in multi-Ethernet configurations. The Ethernet port transmission load is distributed among two or more network ports (assuming the cards are configured for load balancing). An intelligent software adaptive agent repeatedly analyzes the traffic flow from the server and distributes the packets based on destination addresses.

## MAC Address

Short for *Media Access Control address*, a hardware address that uniquely identifies each node of a network. In the Open Systems Interconnection (OSI) model, one of two sublayers of the Data Link Control layer concerned with sharing the physical connection to the network among several computers. Each Ethernet port has a unique MAC address.

**MD5 Algorithm**

MD5 is a way to verify data integrity, and is much more reliable than checksum and many other commonly used methods.

**MIB**

Short for *Management Information Base*. A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP.

**Mirroring**

Used in RAID 1 and 10, a process of storing data on one disk and copying it to one or more disks, creating a redundant storage solution. RAID 1 is the most secure method of storing mission-critical data.

**Mounted**

A filesystem that is available.

**MPIO**

Short for *Multipath Input/Output*. A multipath solution built into Microsoft server-grade iSCSI operating systems.

**MTU**

Short for *Maximum Transfer Unit*. It is the largest size packet or frame, specified in octets (eight-bit bytes), that can be sent in a packet- or frame-based network.

**NIC**

Short for *Network Interface Card*. A board that provides network communication capabilities to and from a computer.

**NIS**

Short for *Network Information Service*. It is a client–server directory service protocol for distributing system configuration data such as user and host names between computers on a computer network. Sun Microsystems developed the NIS; the technology is licensed to virtually all other Unix vendors.

**NTFS**

Short for *New Technology File System*. The standard file system used by Windows NT and later versions of the Windows operating system.

**NTP**

Short for *Network Time Protocol*. A protocol for synchronizing the system clocks of computers over a packet-switched network.

**NVRAM**

Abbreviation of *Non-Volatile Random Access Memory*, a type of memory that retains its contents when power is turned off.

**Permissions**

A security category, such as no access, read-only, or read-write, that determines what operations a user or group can perform on folders or files.

**PoP**

Short for *Proof of Purchase*. The number used to obtain a license key for an upgrade to third-party applications.

**Portal**

A target's IP address together with its TCP port number used in iSCSI systems.

**POSIX**

Short for *Portable Operating System Interface*. A set of standard operating system interfaces based on the UNIX operating system. The need for standardization arose because enterprises using computers wanted to develop programs that could run on multiple platforms without the need to recode.

**Protocol**

A standardized set of rules that specifies the format, timing, sequencing, and/or error checking for data transmissions.

**PTP**

Short for *Point-to-Point*. PTP is the common mode of attachment to a single host. PTP is sometimes used to attach to a Fibre Channel switch for [SAN](#) connectivity.

**Quota**

A limit on the amount of storage space on a volume that a specific user or NIS group can consume.

**Router**

A router is a device that enables connectivity between Ethernet network segments.

**SAN**

Short for *Storage Area Network*. Data storage connected to a network that provides network clients access to data using block level protocols. To the clients, the data storage devices appear local rather than remote. An iSCSI SAN is sometimes referred to as an IP-SAN.

**SAS**

Short for *Serial Attached SCSI*. It is a point-to-point serial protocol that replaces parallel SCSI bus technology (multidrop) and uses the standard SCSI command set. It has no termination issues, supports up to 16,384 devices (using expanders), and eliminates clock skew. It consists of an Initiator that originates device service requests, a Target containing logical units that receives device service requests, and a Service Delivery Subsystem that transmits information between the Initiator and the Target.

**Session**

When an initiator wants to establish a connection with a target, it establishes what is known as an iSCSI session. A session consists of one or more TCP/IP connections between an initiator and a target. Sessions are normally established (or re-established) automatically when the host computer starts up, although they also can be established (and broken) manually.

**SMB**

Short for *Server Message Block*. A protocol for Windows clients. SMB uses the TCP/IP protocol. It is viewed as a complement to the existing Internet application protocols such as FTP and HTTP. With SMB, you can access local server files, obtain read-write privileges to local server files, share files with other clients, and restore connections automatically if the network fails.

**SMTP**

Short for *Simple Mail Transfer Protocol*. A TCP/IP protocol used for sending and receiving email.

**SNMP**

Short for *Simple Network Management Protocol*. A system to monitor and manage network devices such as computers, routers, bridges, and hubs. SNMP views a network as a collection of cooperating, communicating devices, consisting of managers and agents.

**SSH**

Short for *Secure Shell*. A service that provides a remote console for special system administration and customer support access to the server. SSH is similar to telnet but more secure, providing strong encryption so that no passwords cross the network in clear text.

**SSL**

Short for *Secure Sockets Layer*. A protocol for managing the security of a message sent on the Internet. It is a type of technology that provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

**Standalone**

A network bonding mode which treats each port as a separate interface. This configuration should be used only in multihomed environments in which network storage resources must reside on two separate subnets.

**Static IP Address**

An IP address defined by the system administrator rather than by an automated system, such as DHCP.

**Storage Area Network**

See [SAN](#).

**Subnet Mask**

A portion of a network that shares a common address component. On TCP/IP networks, subnets are all devices with IP addresses that have the same prefix.

**Target**

A target is a device (peripheral) that responds to an operation requested by an initiator (host system). Although peripherals are generally targets, a peripheral may be required to act temporarily as an initiator for some commands (for example, SCSI COPY command).

Targets are embedded in iSCSI storage controllers. They are the software that makes the RAID storage available to host computers, making it appear just like any other sort of disk drive.

**TCP/IP**

Short for *Transmission Control Protocol/Internet Protocol*. The basic protocol used for data transmission over the Internet.

**Trap**

A signal from a device informing an SNMP management program that an event has occurred.

**U**

A standard unit of measure for designating the height in computer enclosures and rack cabinets. One U equals 1.75 inches. For example, a 3U server chassis is 5.25 inches high.

**UDP**

Short for *User Datagram Protocol*. A communications protocol for sending messages between computers in a network that uses the Internet Protocol (IP). UDP is an alternative to the Transmission Control Protocol but, unlike TCP, does not guarantee reliability or ordering of data packets.

**UPS**

Short for *Uninterruptible Power Supply*. A device that allows a computer to keep running for a short time when the primary power source is lost. It also provides protection from power surges. A UPS device contains a battery that starts when the device senses a loss of power from the primary source.

**URL**

Short for *Uniform Resource Locator*. A Web address.

**USB Port**

USB is short for *Universal Serial Bus*. A USB port is a hardware interface for low-speed peripherals such as the keyboard, mouse, joystick, scanner, printer, and telephony devices.

**Web Management Interface**

A Web-based utility used for configuration and ongoing maintenance, such as monitoring server conditions, configuring email alerts for key events, or for SNMP management.

**Windows Domain Authentication**

Windows-based networks use a domain controller to store user credentials. The domain controller can validate all authentication requests on behalf of other systems in the domain. The domain controller can also generate encrypted challenges to test the validity of user credentials. Other systems use encrypted challenges to respond to CIFS/SMB clients that request access to a share.

**WINS**

Short for *Windows Internet Naming Service*. The server that locates network resources in a TCP/IP-based Windows network by automatically configuring and maintaining the name and IP address mapping tables.

**Workgroup**

A collection of computers that are grouped for sharing resources such as data and peripherals over a LAN. Each workgroup is identified by a unique name.

## Symbols

> (menu flow indicator) 4

## Numerics

802.3ad link aggregation 57

## A

A record (DNS) 56

Abort button 213

access

problems with 258

Access Denied message 259

ACLs

setting file-level permissions (Windows) 253

Active Directory

and name resolution servers 62

joining AD domain 65

SnapScale interoperability with 63

Active Users page 193, 194

add new drives 139

adding nodes 127

adjusting snapshot space 113

admin password

changing 242

default 142

resetting forgotten 261

Admin password synchronization warning. 163

Administration page 234

administration password 233

Advanced Share Properties 149, 151

Advanced Support Properties 262

ALB 56

alert definitions 4

alert messages 40

Application Notes 247

Authentication

default settings 142

HTTPS/HTTP 78

Kerberos 63

LDAP domain 71

NIS domain 71

Automatic Load Balancing (ALB) 56

automatic shutdown 46

automatic update checking 219

## B

backup

off-the-shelf solutions 247

BitTorrent Sync 236

bond type

change process 58

changing 57

definitions 55

## C

CA Unicenter TNg 76

capacity balance 137

Capacity Balancer 16

Capacity Balancer, see *Data Balancer*

change password 233, 242

CLI connection via SSH 45

client access, configuring

FTP 74

HTTPS/HTTP 78

Windows SMB 64

Client network 12, 16, 51

cluster management name 12

cluster name 12

Command Line Interface 267

running scripts 273

syntax 267

connecting

a Mac OS X client 64

- a Windows client **63**
- contact information pop-up **41**
- conventions, typographical **4**
- customer support **3**

## D

- Data Balancer **16, 93**
- data import **208**
- data protection tasks **109**
- Data Replication Count **12, 29, 43, 87, 127**
- date and time settings **44**
- defaults
  - admin password **142**
  - TCP/IP **54**
- directories, home **188**
- disabling snapshots on cluster **110**
- disk drives
  - LED indicators **255**
- DNS A record **56**
- domain search
  - authentication required **106, 156, 176, 180**
- domains
  - joining ADS **65, 144**
- drives
  - adding **136**
  - considerations **137**
  - failed **137**
  - hot swap **136**
  - installing **137**
  - replacing **136**

## E

- Email Notification **226**
- Ethernet, see *Gigabit Ethernet*
- Event Log page **201**
- expansion kits **127**
- exports file, NFS **147**

## F

- failed drive **137**
- Failover
  - failover **16**
  - file/folder security information **233**
  - files, setting permissions for **252**
- FTP

- connecting via **75**

## G

- GID **143**
- Groups
  - creating local **169**
  - file-level access for **253**
- GuardianOS specifications **13**

## H

- hardware information pop-up **41**
- home directories **188**
- Home page **232**
- home pages **232, 234**
- hot spares **87**
- hot swap drives **136**
- HP Open View **76**
- HTTPS/HTTP, configuring **78**

## I

- ID mapping **175**
- incorporate new drives **139**
- Initial Setup Wizard **23, 37**
- internal temperature, e-mail notification of **227**
- IQNs for iSCSI disks **117**
- iSCSI disks **115**
  - backing up **122**
  - configuring iSNS **82**
  - creating **119**
  - LUNs **124**
  - multi-initiator support **117**
  - name (alias) **117**
- iSNS
  - configuring **82**
  - Update iSNS Registration **83**

## K

- Kerberos **63**
- key icon **233**

## L

- LDAP domains **71**
- LEDs

- disk drive indicators **255**
- network **256**
- power/unit status **256**
- system/status **255**
- understanding **254**
- Link Aggregation (802.3ad) **56, 57**
- Load Balance (ALB) **56**
- local groups **168**

## M

- Macintosh, supported OS versions **14**
- maintenance
  - data import **208**
  - OS update **213**
  - shutdown and restart **207**
  - support **221**
  - tools **225**
- Management IP **12**
- Management node **12**
- mapping, ID **175**
- menu flow indicator **4**
- message in Web Management Interface **40**
- monitor
  - active users **193**
  - event log **201**
  - hardware status **192**
  - open files **193, 195**
  - options **191**
  - tape devices **205**
- monitor network traffic **196**

## N

- network
  - access **50**
  - current settings **51**
  - LED indicators **256**
  - problems with access **258**
- network bonding, see *Failover*
- Network Monitor page **196**
- network monitoring
  - download usage records **200**
  - graphing options **199**
  - viewing usage **197**
- new drives detected **139**
- NFS
  - access **67**

- exports file **147**
- read-only share access **68**
- share-level permissions **158**
- supported versions **68**
- NFS backup **247**
- NIS domains **71**
- Node Number **192**
- Node Properties **126**
- nodes
  - adding **127**
  - default page **125**
  - Properties page for nodes **126**

## O

- Open Files page **193, 195**
- operating system **3**
- OS update
  - checking for **213**
  - Online/Rolling update blockers **218**
- OS Updates
  - online/rolling updates **216**
- OS updates
  - options **215**
- Overland technical support **3**

## P

- password
  - changing **242**
  - default for admin account **142**
  - unlock **166**
- paths
  - connecting via web browser **79**
- peer sets
  - basics **87**
  - data recovery **86**
  - definition **12**
  - formation **85**
  - options **88**
  - overview **85**
  - page overview **88**
- permissions
  - share- and file-level interaction **155**
  - file-level, default behavior **253**
- Phone home support **261**
- power/status LED **256**
- product documentation **3**

## Q

### Quotas

- defaults **101**
- deleting **108**
- editing **108**
- initial page **103**
- search for usage **103**
- usage calculation **101**

## R

- RAINcloudOS specifications **13**
- RapidRebuild **86**
- reboot, setting up alert for **227**
- reducing snapshot space **110, 113**
- Registration page **222**
- remote discovery **241**
- replacing disks **136**
- replacing drives **136**
- replication **245**
- reset options **257**
- restart **207**

## S

- Secure Shell (SSH) **45**
- security
  - guides **143**
  - models **173**
  - shares **147**
  - Windows ACLs **252**
- server events notification **226**
- server registration, online **222**
- Shares **147**
  - delete **152**
  - edit properties **151**
- shares with Admin rights **233**
- shutdown **207**
- Simple Network Management Protocol, see *SNMP*
- site map **41, 231**
  - server links **42**
- SMB **61**
- SMB backup **247**
- SMB2 **65, 67**
- SMTP methods supported **226**
- Snap EDR **239, 245**
- Snap Finder **239**

- SnapCLI **267**
  - running scripts **273**
  - syntax **267**
- SnapExtensions
  - BitTorrent **236**
  - main page **236**
  - mouseover icon **236**
  - Snap EDR **239**
- snapshot
  - access **251**
  - shares **251**
- snapshot share **251**
- snapshot space **110**
- snapshots
  - create **112**
  - default page **110**
  - overview **110**
  - scheduled **114**
- SNMP configuration **77**
- SNMP configuration page **77**
- software information pop-up **41**
- software update **213**
- Spare Disk Balancer, see *Spare Distributor*
- Spare Disks page **90**
- Spare Distributor **16, 91**
- specifications, GuardianOS **13**
- specifications, RAINcloudOS **13**
- storage
  - Nodes default page **125**
  - Volumes default page **96**
- Storage network **12, 16, 53**
- support **221**
- Switch Trunking **56, 57**
- switch user (logout) **234**
- synchronize Admin password warning **163**
- system monitoring **191**
- system reset **257**
- System Status page **192**
- system/status LED **255**

## T

- Tape page **205**
- TCP/IP
  - initial configuration **26**
  - options **54**
- technical support **3**

Tivoli NetView **76**  
tools **225**  
troubleshooting **254**  
typographical conventions **4**

## U

UID **143**  
Uninitialized node **12**  
Uninterruptible Power Supplies (UPS) **46**  
unlock a user password **166**  
updates  
    online/rolling updates **216**  
    overview **213**  
    procedure **213**  
UPS  
    configuring **46**  
    enabling support for **46**  
    low-power warning **46**  
user names, active **194**  
users  
    creating local **162**  
    file-level access for **253**  
Utility IP address **60, 247**

## V

VDS-based iSCSI disks **118**  
version of OS **3**  
volumes  
    capacity reached alert **227**  
    default page **96**  
    expanding capacity of **99**  
    Volume Properties **99**  
VSS-based iSCSI disks **118**

## W

warranty activation **222**  
Web Management Interface  
    alert messages **40**  
Web Management Interface, overview **37**  
Web Root **79**  
Web Server **79**  
Windows  
    connecting from a client **63**  
    enabling guest account access **65, 67**  
    guest account access **63**

issues with PDC **259**  
name resolution server support **62**  
networking (SMB) **61**  
security, joining  
    active directory domain **65**  
    see also *Active Directory*  
    see also *Authentication*  
Windows Active Directory **63**  
    setup **65, 144**  
    Shares **147**  
workgroup environment **62**  
workgroup, joining **64**

