

**Oracle® ZFS Storage Appliance
Administration Guide, Release OS8.7.x**



Part No: E79457-03
September 2017

Part No: E79457-03

Copyright © 2009, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Référence: E79457-03

Copyright © 2009, 2017, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

- About Oracle ZFS Storage Appliance** 19
 - Oracle ZFS Storage Appliance Key Features 20
 - Supported Protocols 20
 - Appliance Data Services 21
 - Data Availability 21
 - Browser User Interface (BUI) 22
 - Network Icons 30
 - Dashboard Icons 30
 - Analytics Toolbar Icons 31
 - Identity Mapping Icons 32
 - Supported Browsers 33
 - Command Line Interface (CLI) 33
 - CLI Contexts 35
 - CLI Properties 42
 - Working with CLI Scripting 45
 - Using Batch Commands 45
 - Understanding the CLI Scripting Commands 46
 - ▼ Accessing the CLI Script Environment 46
 - Understanding the Built-in CLI Functions 47
 - ▼ Using the Run Function 48
 - ▼ Using the Get Function 48
 - ▼ Using the List Function 49
 - ▼ Using the Children Function 52
 - ▼ Using the Choices Function 53
 - Using the Functions for Generating Output 54
 - Understanding CLI Scripting Errors 55
- Configuring the Appliance** 57

Initial Appliance Configuration	57
Appliance Cluster Configuration	58
Cluster Configuration BUI View	58
▼ Upgrading a Standalone Appliance to a Clustered Configuration (BUI)	60
▼ Shutting Down a Clustered Configuration (BUI)	62
▼ Shutting Down a Clustered Configuration (CLI)	64
Cluster Terminology	66
Understanding Clustering	66
Cluster Advantages and Disadvantages	68
Cluster Interconnect I/O	70
Cluster Resource Management	71
Cluster Takeover and Failback	74
Configuration Changes in a Clustered Environment	76
Clustering Considerations for Storage	77
Clustering Considerations for Networking	79
Private Local IP Interfaces	81
Clustering Considerations for InfiniBand	82
Preventing Split-Brain Conditions	85
Estimating and Reducing Takeover Impact	87
Network Configuration	89
Network Configuration (BUI)	90
Network Configuration (CLI)	102
Working with Network Configuration	111
Configuring Management Interfaces	113
Configuring Network Datalinks	113
Configuring Network Interfaces	116
Configuring Network IP MultiPathing (IPMP)	117
Configuring Network Performance and Availability	118
Configuring Network Routing	119
Configuring Storage	122
▼ Creating a Storage Pool (BUI)	122
▼ Creating a Storage Pool (CLI)	125
▼ Importing an Existing Storage Pool (BUI)	128
▼ Importing an Existing Storage Pool (CLI)	128
▼ Configuring an All-Flash Storage Pool (BUI)	130
▼ Configuring an All-Flash Storage Pool (CLI)	131
▼ Adding a Disk Shelf to an Existing Storage Pool (BUI)	133

▼ Adding a Disk Shelf to an Existing Storage Pool (CLI)	135
▼ Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI)	137
▼ Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI)	138
▼ Removing a Cache or Log Device from an Existing Storage Pool (BUI)	141
▼ Removing a Cache or Log Device from an Existing Storage Pool (CLI)	142
▼ Unconfiguring a Storage Pool (BUI)	143
▼ Unconfiguring a Storage Pool (CLI)	144
▼ Renaming a Storage Pool (BUI)	145
▼ Renaming a Storage Pool (CLI)	146
▼ Scrubbing a Storage Pool (BUI)	147
▼ Scrubbing a Storage Pool (CLI)	148
▼ Viewing Pool and Device Status (BUI)	149
Storage Pool Concepts	150
Data Profiles for Storage Pools	152
Understanding the Appliance Status	155
Status Dashboard	155
Summary of Pool Usage	161
Summary of Memory Usage	161
Disk Activity Dashboard	162
Dashboard CLI	164
▼ Running the Dashboard Continuously	165
Status Dashboard Settings	166
▼ Changing the Displayed Activity Statistics	168
▼ Changing the Activity Thresholds	168
NDMP Status	169
NDMP States	170
Configuring Storage Area Network (SAN)	170
▼ Configuring FC Port Modes (BUI)	171
▼ Discovering FC Ports (BUI)	173
▼ Creating FC Initiator Groups (BUI)	174
▼ Associating a LUN with an FC Initiator Group (BUI)	176
▼ Changing FC Port Modes (CLI)	177
▼ Discovering FC Ports (CLI)	177
▼ Creating FC Initiator Groups (CLI)	178
▼ Associating a LUN with an FC Initiator Group (CLI)	179

▼ Scripting Aliases for Initiators and Initiator Groups (CLI)	179
▼ Creating an Analytics Worksheet (BUI)	181
▼ Configuring SAN iSER Targets	182
▼ Adding an iSCSI Target with an Auto-generated IQN (CLI)	185
▼ Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)	185
▼ Adding an iSCSI Initiator with CHAP Authentication (CLI)	186
▼ Adding an iSCSI Target Group (CLI)	187
▼ Adding an iSCSI Initiator Group (CLI)	187
▼ Configuring SRP Target (BUI)	188
▼ Configuring SRP Targets (CLI)	189
Understanding SAN	190
SAN Fibre Channel Configuration	192
SAN iSCSI Configuration	196
SAN iSCSI Initiator Configuration	197
SAN SRP Configuration	199
SAN Terminology	200
Configuring Users	202
▼ Adding an Administrator or User (BUI)	203
▼ Adding an Administrator or User (CLI)	205
▼ Changing a User Password (BUI)	207
▼ Changing a User Password (CLI)	207
▼ Editing Exceptions for a User (BUI)	208
▼ Editing Exceptions for a User (CLI)	209
▼ Deleting Exceptions for a User (BUI)	211
▼ Deleting Exceptions for a User (CLI)	212
▼ Adding a Role (BUI)	213
▼ Adding a Role (CLI)	213
▼ Editing Authorizations for a Role (BUI)	214
▼ Editing Authorizations for a Role (CLI)	215
▼ Deleting Authorizations from a Role (BUI)	217
▼ Deleting Authorizations from a Role (CLI)	217
▼ Adding a User Who Can View the Dashboard	218
Understanding Users and Roles	218
User Authorizations	219
Managing User Properties	222
Setting Appliance Preferences	224

▼ Setting Preferences (BUI)	224
▼ Setting Preferences (CLI)	225
▼ Setting SSH Public Keys (BUI)	226
▼ Setting SSH Public Keys (CLI)	226
Preference Properties	228
Configuring Alerts	228
▼ Adding an Alert Action (BUI)	229
▼ Adding an Alert Action (CLI)	230
Sending Email Alerts (CLI)	231
Sending an SNMP Trap (CLI)	231
Alert Categories	232
Threshold Alerts	233
Resuming/Suspending Analytics Datasets and Worksheets	234
Configuring Certificates	235
▼ Creating a New Server Certificate (BUI)	236
▼ Creating a New Server Certificate (CLI)	237
▼ Uploading CA Certificates from Non-root CAs (BUI)	238
▼ Uploading CA Certificates from Non-root CAs (CLI)	239
▼ Viewing CSR and Certificate Details (BUI)	240
▼ Viewing CSR and Certificate Details (CLI)	240
▼ Destroying a CSR or Certificate (BUI)	241
▼ Destroying a CSR or Certificate (CLI)	241
▼ Setting the Appliance Certificate (BUI)	242
▼ Setting the Appliance Certificate (CLI)	242
▼ Uploading Trusted Certificates (BUI)	243
▼ Uploading Trusted Certificates (CLI)	243
▼ Viewing Trusted Certificate Details (BUI)	244
▼ Viewing Trusted Certificate Details (CLI)	244
▼ Destroying a Trusted Certificate (BUI)	245
▼ Destroying a Trusted Certificate (CLI)	245
▼ Assigning a Certificate to a Service (BUI)	246
▼ Assigning a Certificate to a Service (CLI)	246
Appliance Services	247
Managing Services	248
▼ Viewing a Service in the BUI	248
▼ Selecting a Service in the CLI	249

▼ Enabling a Service (BUI)	249
▼ Enabling a Service (CLI)	250
▼ Disabling a Service (BUI)	250
▼ Disabling a Service (CLI)	250
▼ Viewing Service States in the CLI	251
▼ Viewing Service Help in the CLI	251
▼ Setting Service Properties (BUI)	252
▼ Setting Service Properties (CLI)	253
▼ Viewing Service Logs (BUI)	254
▼ Viewing Service Logs (CLI)	255
List of Available Appliance Services	256
Required Service Ports	258
Configuring Services	258
Active Directory Configuration	259
DNS Configuration	265
Dynamic Routing Configuration	272
FTP Configuration	273
HTTP Configuration	275
HTTPS Configuration	279
Identity Mapping Configuration	280
IPMP Configuration	291
iSCSI Configuration	292
Kerberos Configuration	293
LDAP Configuration	310
NDMP Configuration	319
NFS Configuration	328
NIS Configuration	333
NTP Configuration	335
Phone Home Configuration	339
RESTful API Configuration	342
Service Tags Configuration	342
SFTP Configuration	342
Shadow Migration Configuration	346
SMB Configuration	346
SMTP Configuration	365
SNMP Configuration	366
SRP Configuration	370

SSH Configuration	370
Syslog Configuration	371
System Identity Configuration	377
TFTP Configuration	378
Virus Scan Configuration	378
Shares and Projects	383
▼ Creating a Project (BUI)	384
▼ Creating a Project (CLI)	384
▼ Editing a Project (BUI)	386
▼ Editing a Project (CLI)	386
▼ Renaming a Project (BUI)	388
▼ Renaming a Project (CLI)	388
▼ Deleting a Project (BUI)	389
▼ Deleting a Project (CLI)	389
▼ Creating a Filesystem or LUN in a Project (BUI)	390
▼ Creating a Filesystem or LUN in a Project (CLI)	391
▼ Editing a Filesystem or LUN (BUI)	393
▼ Editing a Filesystem or LUN (CLI)	394
▼ Renaming a Filesystem or LUN (BUI)	396
▼ Renaming a Filesystem or LUN (CLI)	396
▼ Moving a Filesystem or LUN to a Different Project (BUI)	397
▼ Moving a Filesystem or LUN to a Different Project (CLI)	397
▼ Deleting a Filesystem or LUN (BUI)	398
▼ Deleting a Filesystem or LUN (CLI)	398
▼ Setting User or Group Quotas (BUI)	399
▼ Setting User or Group Quotas (CLI)	400
About Storage Pools, Projects, and Shares	402
Project and Share Properties	403
Inherited Properties	404
LUN Local Properties	412
Other Properties	413
Static Properties	414
Project Properties	417
Filesystem Properties	423
LUN Properties	431
Space Management for Shares	436

Shares Terminology	436
Managing Filesystem and Project Space	437
Setting User or Group Quotas	438
Working with Identity Management	439
Working with Filesystem Namespace	439
Share Usage Statistics	441
Share and Project Protocols	442
NFS Protocol	442
SMB Protocol	449
HTTP Protocol	455
FTP Protocol	456
SFTP Protocol	456
TFTP Protocol	456
Access Control Lists for Filesystems	456
Root Directory Access	457
ACL Behavior on Mode Change	458
ACL Inheritance Behavior	459
Root Directory ACL	460
Working with Schemas	464
▼ Creating a Schema (BUI)	464
▼ Creating a Schema (CLI)	464
Schema Properties	466
Shadow Migration	469
Understanding Shadow Migration	470
Creating a Shadow Filesystem	472
Managing Background Migration	473
Handling Migration Errors	473
Monitoring Migration Progress	474
▼ Monitoring Migration Progress and Errors (BUI)	474
▼ Monitoring Migration Progress and Errors (CLI)	475
Canceling Migration	477
Snapshotting Shadow File Systems	477
Backing Up Shadow File Systems	478
Replicating Shadow File Systems	478
Migrating Local File Systems	478
Using Shadow Migration Analytics	479

▼ Testing Potential Shadow Migration using the CLI	479
▼ Migrating Data from an Active NFS Server using the CLI	480
Snapshots and Clones	481
Snapshot Space Management	482
▼ Taking a Snapshot (BUI)	484
▼ Taking a Snapshot (CLI)	485
▼ Scheduling Snapshots (BUI)	485
▼ Scheduling Snapshots (CLI)	487
▼ Setting a Scheduled Snapshot Label (BUI)	489
▼ Setting a Scheduled Snapshot Label (CLI)	489
▼ Viewing Snapshots and Schedules (BUI)	490
▼ Viewing Snapshots and Schedules (CLI)	491
▼ Editing a Snapshot Retention Policy (BUI)	492
▼ Editing a Snapshot Retention Policy (CLI)	493
▼ Removing a Snapshot Schedule (BUI)	494
▼ Removing a Snapshot Schedule (CLI)	495
▼ Making a Filesystem Snapshot Directory Visible (BUI)	496
▼ Making a Filesystem Snapshot Directory Visible (CLI)	497
▼ Accessing a Hidden Filesystem Snapshot Directory (CLI)	498
▼ Accessing a Visible Filesystem Snapshot Directory (CLI)	498
▼ Renaming a Snapshot (BUI)	499
▼ Renaming a Snapshot (CLI)	500
▼ Rolling Back to a Snapshot (BUI)	501
▼ Rolling Back to a Snapshot (CLI)	501
▼ Destroying a Snapshot (BUI)	502
▼ Destroying a Snapshot (CLI)	503
▼ Cloning a Snapshot (BUI)	504
▼ Cloning a Snapshot (CLI)	506
▼ Cloning a Clone	508
▼ Viewing Clones of a Snapshot (BUI)	508
▼ Viewing Clones of a Snapshot (CLI)	509
▼ Viewing a Clone Origin (BUI)	509
▼ Viewing a Clone Origin (CLI)	510
Remote Replication	511
▼ Remote Replication Workflow	511

Configuring Remote Replication	512
▼ Checking Source and Target Compatibility	513
▼ Setting Up Network Interfaces and Static Routing (BUI)	513
▼ Setting Up Network Interfaces and Static Routing (CLI)	515
▼ Creating a Replication Target (BUI)	516
▼ Creating a Replication Target (CLI)	517
▼ Creating a Replication Action (BUI)	518
▼ Creating a Replication Action (CLI)	520
▼ Configuring Automatic Snapshot Retention on a Target (BUI)	522
▼ Configuring Automatic Snapshot Retention on a Target (CLI)	523
▼ Manually Sending a Replication Update (BUI)	525
▼ Manually Sending a Replication Update (CLI)	526
▼ Configuring Replication for a Clustered Configuration	526
Configuring Offline Replication (BUI)	527
Configuring Offline Replication (CLI)	532
▼ Disabling Replication Compression (BUI)	544
▼ Disabling Replication Compression (CLI)	545
▼ Editing a Replication Target (BUI)	546
▼ Editing a Replication Target (CLI)	546
▼ Editing a Replication Action (BUI)	547
▼ Editing a Replication Action (CLI)	547
Monitoring Remote Replication	548
▼ Monitoring Replication Progress (BUI)	548
▼ Monitoring Replication Progress (CLI)	549
▼ Setting Replication Alerts	551
Replication Audit Actions	552
▼ Monitoring Replication Delays and RPO (BUI)	553
▼ Monitoring Replication Delays and RPO (CLI)	554
Using Replication Analytics	556
Managing Replication Packages	556
Managing User-Generated Snapshots	557
▼ Canceling a Replication Update (BUI)	558
▼ Canceling a Replication Update (CLI)	559
▼ Cloning a Replication Package (BUI)	560
▼ Cloning a Replication Package (CLI)	563
▼ Cloning a Snapshot in a Replication Package (BUI)	567
▼ Cloning a Snapshot in a Replication Package (CLI)	570

▼ Severing a Replication Package (BUI)	575
▼ Severing a Replication Package (CLI)	575
▼ Editing a Replication Package (BUI)	576
▼ Editing a Replication Package (CLI)	577
▼ Disabling a Replication Package (BUI)	579
▼ Disabling a Replication Package (CLI)	579
Disaster Recovery with Remote Replication	580
▼ Setting Up a Replication Target at a Recovery Site (BUI)	580
▼ Switching Operations to the Recovery Site (BUI)	581
▼ Updating the Production Site (BUI)	582
▼ Reversing Replication Back to the Production Site (BUI)	582
▼ Setting Up a Replication Target at a Recovery Site (CLI)	583
▼ Switching Operations to the Recovery Site (CLI)	584
▼ Updating the Production Site (CLI)	585
▼ Reversing Replication Back to the Production Site (CLI)	586
Remote Replication Concepts	588
Replication Terminology	590
Replication Targets	590
Replication Actions and Packages	591
Replication Action Properties	594
Replication Package Properties	597
Replication Storage Pools	598
Project vs. Share Replication	598
Replication Authorizations	600
Deduplicated Replication	600
Replication Configuration for Clustered Appliances	603
Example: Replication Configuration for Clustered Appliances	604
Replication Snapshots and Data Consistency	611
Replication Snapshot Management	612
iSCSI Configurations and Replication	615
Resumable Replication	615
Replication Alerts	616
Replication Failures	617
Compressed Replication	620
Replication Packages	621
Cloning a Replication Package or Share	621
Exporting Replicated Filesystems	623

Severing Replication	624
Reverse the Direction of Replication	625
Destroying a Replication Package	628
Target Replica Backups	628
Data Encryption	629
▼ Data Encryption Workflow	630
▼ Configuring LOCAL Keystore Encryption (BUI)	630
▼ Configuring LOCAL Keystore Encryption (CLI)	633
▼ Configuring OKM Keystore Encryption (BUI)	634
▼ Configuring OKM Keystore Encryption (CLI)	635
▼ Creating an Encrypted Project (BUI)	636
▼ Creating an Encrypted Project (CLI)	637
▼ Changing a Project Encryption Key (BUI)	638
▼ Changing a Project Encryption Key (CLI)	640
▼ Creating an Encrypted Filesystem or LUN (BUI)	640
▼ Creating an Encrypted Filesystem or LUN (CLI)	641
▼ Changing a Share Encryption Key (BUI)	642
▼ Changing a Share Encryption Key (CLI)	644
▼ Backing Up a LOCAL Key (BUI)	644
▼ Backing Up a LOCAL Key (CLI)	645
▼ Deleting an Encryption Key (BUI)	645
▼ Deleting an Encryption Key (CLI)	648
▼ Restoring a LOCAL Key (BUI)	649
▼ Restoring a LOCAL Key (CLI)	650
Encryption Properties	651
Managing Encryption Keys	652
Maintaining Keys	653
Understanding Encryption Key Values	653
Performance Impact of Encryption	654
Encryption Key Life Cycle	655
Backing up and Restoring Encrypted Data	655
Replicating an Encrypted Share	656
Maintenance Workflows	657
Understanding Workflows	658
Understanding Workflow Parameters	659

Constrained Workflow Parameters	660
Optional Workflow Parameters	661
Workflow Error Handling	662
Workflow Input Validation	663
Workflow Execution Auditing and Reporting	664
Understanding Workflow Versioning	667
Using Workflows for Alert Actions	668
Using Scheduled Workflows	670
Using a Scheduled Workflow	671
Coding Workflow Schedules	671
Creating a Worksheet Based on a Specified Drive Type	673
Uploading and Executing Workflows Using the BUI	676
▼ Downloading Workflows using the CLI	677
▼ Listing Workflows using the CLI	678
▼ Executing Workflows using the CLI	679
▼ Auditing Workflows using the CLI	680
Integration	681
Configuring Oracle ZFS Storage Appliance for Oracle Database Clients	682
Plug-ins for Oracle Products	682
Oracle Enterprise Manager Plug-in for Oracle ZFS Storage Appliance	683
Oracle VM Storage Connect Plug-in for Oracle ZFS Storage Appliance	683
Oracle ZFS Storage Appliance Network File System Plug-in for Oracle Solaris Cluster	684
Oracle ZFS Storage Appliance Plug-in for Oracle Solaris Cluster Geographic Edition	684
Plug-ins for Non-Oracle Products	684
Oracle ZFS Storage Appliance Virtual Storage Manager Plug-ins for VMware vSphere and VMware vSphere Web Client	685
Oracle ZFS Storage Appliance Storage Replication Adapter for VMware Site Recovery Manager	686
Oracle ZFS Storage Appliance Plug-in for VMware vSphere Storage APIs for Array Integration – NAS	686
Oracle ZFS Storage Appliance Provider for VMware vSphere APIs for Storage Awareness	686
Oracle ZFS Storage Appliance Provider for Volume Shadow Copy Service Software	687
Oracle ZFS Storage Appliance Plug-in for Veritas NetBackup OpenStorage	687

Oracle ZFS Storage Appliance Plug-in for CommVault Simpana IntelliSnap ...	688
Oracle Intelligent Storage Protocol	688
Database Record Size	688
Synchronous Write Bias Hint	689
Analytics Breakdown by Database Name	689
Caching Hints	689
OISP-Capable Protocols and Clients	689
Fibre Channel and iSCSI Support with Veritas Dynamic Multi-Pathing and Storage Foundation/InfoScale Foundation	690


About Oracle ZFS Storage Appliance

The Oracle ZFS Storage Appliance (appliance) family of products provides efficient file and block data services to clients over a network, and a rich set of data services that can be applied to the data stored on the system.

For information about configuring and working with the Oracle ZFS Storage Appliance product, see the following sections:

- [Oracle ZFS Storage Appliance Key Features](#)
- [“Supported Protocols” on page 20](#)
- [Oracle ZFS Storage Appliance Data Services](#)
- [“Data Availability” on page 21](#)
- [Browser User Interface \(BUI\)](#)
- [Network Icons](#)
- [Dashboard Icons](#)
- [Analytics Icons](#)
- [Identity Mapping Icons](#)
- [Supported Browsers](#)
- [Command Line Interface \(CLI\)](#)
- [Working with CLI Scripting](#)

Oracle ZFS Storage Appliance Key Features



Oracle ZFS Storage Appliance includes technologies to deliver the best storage price/performance and unprecedented observability of your workloads in production, including:

- Analytics, a system for dynamically observing the behavior of your system in real-time and viewing data graphically
- The ZFS Hybrid Storage Pool, composed of optional Flash-memory devices for acceleration of reads and writes, low-power, high-capacity disks, and DRAM memory, all managed transparently as a single data hierarchy
- Support for a variety of hardware

For more information about Analytics and Hardware, refer to the documentation on the [Oracle Technology Network \(http://www.oracle.com/technetwork/indexes/documentation/index.html\)](http://www.oracle.com/technetwork/indexes/documentation/index.html)

Supported Protocols

Oracle ZFS Storage Appliance supports a variety of industry-standard client protocols, including NFS, iSCSI, SMB, FTP, HTTP, NDMP, Fibre Channel, SRP, iSER, and SFTP.

For information on these protocols, see the following:

- [“SAN Fibre Channel Configuration” on page 192](#)
- [“Configuring SAN iSER Targets” on page 182](#)
- [“NFS Configuration” on page 328](#)
- [“iSCSI Configuration” on page 292](#)
- [“SMB Configuration” on page 346](#)

- [“FTP Configuration” on page 273](#)
- [“HTTP Configuration” on page 275](#)
- [“NDMP Configuration” on page 319](#)
- [“SFTP Configuration” on page 342](#)
- [“SRP Configuration” on page 370](#)

Appliance Data Services

To manage the data that you export using these protocols, you can configure the appliance using the built-in collection of advanced data services, including:

- RAID-Z (RAID-5 and RAID-6), mirrored, and striped disk configurations (See [“Configuring Storage” on page 122](#))
- Unlimited read-only and read-write snapshots, with snapshot schedules (See [“Snapshots and Clones” on page 481](#))
- Controlling the elimination of duplicate copies of data (See [“Data Deduplication” on page 406](#))
- Built-in data compression (See [“Data Compression” on page 408](#))
- Remote replication of data for disaster recovery (See [“Remote Replication” on page 511](#))
- Active-active clustering for high availability (See [“Appliance Cluster Configuration” on page 58](#))
- Thin provisioning of iSCSI LUNs (See [“iSCSI Configuration” on page 292](#))
- Virus scanning and quarantine (See [“Virus Scan Configuration” on page 378](#))
- NDMP backup and restore (See [“NDMP Configuration” on page 319](#))

Note - Replication and Cloning are licensed features for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Data Availability

To maximize the availability of your data in production, the appliance includes a complete end-to-end architecture for data integrity, including redundancies at every level of the stack. Key features include:

- Predictive self-healing and diagnosis of all system hardware failures: CPUs, DRAM, I/O cards, disks, fans, power supplies
- ZFS end-to-end data checksums of all data and metadata, protecting data throughout the stack
- RAID-6 (double- and triple-parity) and optional RAID-6 across disk shelves
- Active-active clustering for high availability (See [“Appliance Cluster Configuration” on page 58](#))
- Link aggregations and IP multipathing for network failure protection (See [“Network Configuration” on page 89](#))
- I/O Multipathing between the controller and disk shelves
- Integrated software restart of all system software services (See [“Appliance Services” on page 247](#))
- Phone Home of telemetry for all software and hardware issues (See [“Phone Home Configuration” on page 339](#))
- Lights-out Management of each system for remote power control and console access


Browser User Interface (BUI)

The Oracle ZFS Storage Appliance Browser User Interface (BUI) is the graphical tool for administration of the appliance. The BUI provides an intuitive environment for administration tasks, visualizing concepts, and analyzing performance data. The BUI provides an uncluttered environment for visualizing system behavior and identifying performance issues with the appliance.



Direct your browser to the system using either the *IP address* or *host name* you assigned to the NET-0 port during initial configuration as follows: `https://ipaddress:215` or `https://hostname:215`. The login screen appears.

The online help linked in the top right of the BUI is context-sensitive. For every top-level and second-level screen in the BUI, the associated help page appears when you click the Help button.



Changing a filesystem's properties by moving it into another project using the Projects side panel.

The masthead contains several interface elements for navigation and notification, as well as primary functionality. At left, from top to bottom, are the Sun/Oracle logo, a hardware model badge, and hardware power off/restart button. Across the right, again from top to bottom: login identification, logout, help, main navigation, and subnavigation.




System alerts appear in the Masthead as they are triggered. If multiple alerts are triggered sequentially, refer to the list of recent alerts found on the Dashboard screen or the full log available on the Logs screen.

Use the main navigation links to view between the Configuration, Maintenance, Shares, Status, and Analytics areas of the BUI. Use sub-navigation links to access features and functions within each area.

If you provide a session annotation, it appears beneath your login ID and the logout control. To change your session annotation for subsequent administrative actions without logging out, click on the text link. For details about session annotations, see [“Configuring Users” on page 202](#).


The title bar appears below the Masthead and provides local navigation and functions that vary depending on the current view.



For example, the Identity mapping service title bar enables the following:

- Navigation to the full list of services through the side panel
- Controls to enable or disable the Identity Mapping service
- A view of Identity Mapping uptime
- Navigation to the Properties, Rules and Logs screens for your Identity Mapping service
- Button to Apply configuration changes made on the current screen
- Button to Revert configuration changes applied on the current screen

To quickly navigate between Service and Project views, open and close the side panel by clicking the title or the reveal  arrow.



To add projects, click the Add... link in the sidebar.

To move Shares between Projects, click the move  icon and drag a filesystem Share to the appropriate Project in the side panel.

Note that dragging a share into another project will change its properties if they are set to be inherited from its parent project.

Most BUI controls use standard web form inputs; however, there are a few key exceptions worth noting:

TABLE 1 Key Web Form Exceptions

Summary of BUI Controls	
Modify a property	Click the edit  icon and complete the dialog
Add a list item or property entry	Click the add  icon
Remove a list item or property entry	Click the remove  icon
Save changes	Click the Apply button
Undo saved changes	Click the Revert button
Delete an item from a list	Click the trash  icon (hover the mouse over the item row to see the icon)
Search for an item in a list	Click the search  icon at the top right of the list
Sort by list headings	Click on the bold sub-headings to re-sort the list
Move or drag an item	Click the move  icon


Summary of BUI Controls

Rename an item	Click the rename  icon
View details about your system	Oracle logo or click the model badge to go to the oracle.com web page for your model
Automatically open side panel	Drag an item to the side panel


When setting permissions, the RWX boxes are clickable targets. Clicking on the access group label (User, Group, Other) toggles all permissions for that label on and off.



To edit Share properties, deselect the Inherit from project checkbox.




To view controls for an item in a list, hover the mouse over the row.



NAME	SIZE	MOUNTPOINT
admin / accounts	78.5K	/export/accounts
admin / exports	78.5K	/export/exports
admin / primary	78.5K	/export/primary
admin / traffic	78.5K	/export/traffic
admin / workflow	78.5K	/export/workflow

All modal dialogs have titles and buttons that identify and commit or cancel the current action at top, and content below. The modal content area follows the same interface conventions as the main content area, but are different in that they must be dismissed using the buttons in the title bar before other actions can be performed.



Add Threshold Alert [CANCEL] [APPLY]

Threshold

CPU: percent utilization exceeds 95 percent

Timing

for at least 5 minutes ☐ only between 00:00 and 00:00 ☐ only during weekdays

☒ Repost alert every 5 minutes while this condition persists.

☒ Also post alert when this condition clears for at least 5 minutes

Alert actions

Send email [TEST] Send to admin@hostname.com Subject: Alert! CPU has exceeded threshold

Icons indicate system status and provide access to functionality, and in most cases serve as buttons to perform actions when clicked. It is useful to hover your mouse over interface icons to view the tooltip. The tables below provide a key to the conventions of the user interface.

The status lights are basic indicators of system health and service state:

TABLE 2 Status Indicators

Icon	Description	Icon	Description
	on		warning
	off		disabled

The following icons are found throughout the user interface and cover most of the basic functionality:

TABLE 3 BUI Icons

Icon*		Description	Icon*		Description
--		rename (edit text)	--		sever
--		move	--		clone
		edit	--		rollback
		destroy	--		appliance power
		add	--		apply
		remove	--		revert
		cancel/close	--		info
--		error	--		sort list column (down)
--		alert	--		sort list column (up)
		on/off toggle			first page
		restart			previous page
--		locate			next page
		disable/offline			last page
		lock	--		search
--		wait spinner			menu
--		reverse direction			panel

* Disabled icons are shown at left.

The following icons are used to distinguish different types of objects and provide information of secondary importance.

TABLE 4 Miscellaneous Icons

Icon	Description	Icon	Description
	allow		SAS
	deny		SAS port
	storage pool		

Network Icons

These icons indicate the state of network devices and type of network datalinks:

TABLE 5 Network Icons

Icon	Description	Icon	Description
	active network device		active InfiniBand port
	inactive network device		inactive InfiniBand port
	network datalink		network datalink for an InfiniBand partition
	network datalink VLAN		
	network datalink aggregation		
	network datalink aggregation VLAN		

Dashboard Icons

The following icons indicate the current state of monitored statistics with respect to user-configurable thresholds set from within Settings.

TABLE 6 Dashboard Icons

Icon	Description	Icon	Description
	sunny		hurricane
	partly cloudy		hurricane class 2
	cloudy		hurricane class 3
	rainy		hurricane class 4
	stormy		hurricane class 5

Analytics Toolbar Icons

This set of icons is used in a toolbar to manipulate display of information within Analytics worksheets.

TABLE 7 Analytics Toolbar Icons

Icon	Description	Icon	Description
	back		show minimum
	forward		show maximum
	forward to now		show line graph
	pause		show mountain graph
	zoom out		crop outliers
	zoom in		sync worksheet to this statistic
	show one minute		unsync worksheet statistics
	show one hour		drilldown

Icon	Description	Icon	Description
	show one day		export statistical data (download to client)
	show one week		save statistical data
	show one month		archive dataset
			send worksheet with support bundle

For more information about Analytics, refer to the documentation on the [Oracle Technology Network](http://www.oracle.com/technetwork/indexes/documentation/index.html) (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)

Identity Mapping Icons

These icons indicate the type of role being applied when mapping users and groups between Windows and Unix.

TABLE 8 Identity Mapping Icons

Icon*	Description	Icon*	Description
	allow Windows to Unix		allow Unix to Windows
	deny Windows to Unix		deny Unix to Windows
	allow bidirectional		

**Disabled icons shown at left.*

Related Topics

- [“Understanding the Appliance Status” on page 155](#)
- [“Network Configuration” on page 89](#)
- [“Configuring Storage” on page 122](#)
- [“Configuring Alerts” on page 228](#)
- [“Appliance Services” on page 247](#)

For more information about Analytics, refer to the documentation on the [Oracle Technology Network](http://www.oracle.com/technetwork/indexes/documentation/index.html) (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)

Supported Browsers

This section defines BUI browser support.

The BUI is fully featured and functional on the following browsers:

- Firefox 10 and newer
- Internet Explorer 9 and newer
- Safari 5 and newer
- Google Chrome 31 and newer

BUI elements may be cosmetically imperfect on the following browsers, and some functionality may not be available, although all necessary features work correctly. A warning message appears during login if you are using one of the following browsers:

- Firefox 6 to 9
- Internet Explorer 7 and 8
- Google Chrome 21 to 30
- Opera 23 and older

The following browsers are incompatible, unsupported, and known to have issues; login will not complete.

- Firefox 5 & older
- Internet Explorer 6 & older
- Google Chrome 20 & older
- Safari 4 & older
- Opera 22 & older

Related Topics

- [“Configuring Users” on page 202](#)
- [“Setting Appliance Preferences” on page 224](#)

Command Line Interface (CLI)

The CLI is designed to imitate the capabilities of the BUI, while also providing a powerful scripting environment for performing repetitive tasks. The command line is an efficient and powerful tool for repetitive administrative tasks. The appliance presents a CLI available through

either the Console or SSH. There are several situations in which the preferred interaction with the system is the CLI:

- **Network unavailability** - If the network is unavailable, browser-based management is impossible; the only vector for management is the Console, which can only accommodate a text-based interface
- **Expediency** - Starting a browser may be prohibitively time-consuming, especially if you only want to examine a particular aspect of the system or make a quick configuration change
- **Precision** - In some situations, the information provided by the browser may be more qualitative than quantitative in nature, and you need a more precise answer
- **Automation** - Browser-based interaction cannot be easily automated; if you have repetitive or rigidly defined tasks, script the tasks
- **Accessibility** - The CLI is an alternative, and equivalent, way to access the BUI features and functionality. Because the operating systems that run on Oracle ZFS Storage Appliance systems support assistive technologies to read the content of the screen, you can use the CLI as an equivalent means to access the color-based, mouse-based, and other visual-based utilities that are part of the BUI. For example, you can use a keyboard to enter CLI commands to identify faulted hardware components, check system status, and monitor system health.

Oracle strives to make its products, services, and supporting documentation usable and accessible to the disabled community. To that end, products, services, and documentation include features that make the product accessible to users of assistive technology. For more information about Oracle's commitment to accessibility, visit the [Oracle Accessibility Program website \(http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc\)](http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc).

When navigating through the CLI, there are two principles to be aware of:

- **Tab completion is used extensively** - if you are not sure what to type in any given context, pressing the Tab key will provide you with possible options. Throughout the documentation, pressing Tab is presented as the word "tab" in bold italics.
- **Help is always available** - the help command provides context-specific help. Help on a particular topic is available by specifying the topic as an argument to help, for example `help commands`. Available topics are displayed by tab-completing the help command, or by typing `help topics`.

You can combine these two principles as follows:

```
dory:> help tab
builtins  commands  general  help      properties  script
```

To log in remotely using the CLI, use an ssh client. If you have not followed the instructions in [“Configuring Users” on page 202](#) to administer the appliance, you will need to log in as

root. When you log in, the CLI will present you with a prompt that consists of the hostname, followed by a colon, followed by a greater-than sign:

```
% ssh root@dory
Password:
Last login: Mon Oct 13 15:43:05 2009 from kiowa.sf.fishpo
dory:>
```

Related Topics

- [“Browser User Interface \(BUI\)” on page 22](#)
- [“CLI Contexts” on page 35](#)
- [“CLI Properties” on page 42](#)

CLI Contexts

A central principle in the CLI is the *context* in which commands are executed. The context dictates which elements of the system can be managed and which commands are available. Contexts have a tree structure in which contexts may themselves contain nested contexts and the structure generally mirrors that of the views in the BUI.

The initial context upon login is the *root context*, and serves as the parent or ancestor of all contexts. To navigate to a context, execute the name of the context as a command. For example, the functionality available in the Configuration view in the browser is available in the configuration context of the CLI. From the root context, this can be accessed by typing it directly:

```
dory:> configuration
dory:configuration>
```

Note that the prompt changes to reflect the context, with the context provided between the colon and the greater-than sign in the prompt.

The show command shows child contexts. For example, from the configuration context:

```
dory:configuration> show
Children:
    net => Configure networking
    services => Configure services
    version => Display system version
    users => Configure administrative users
    roles => Configure administrative roles
    preferences => Configure user preferences
    alerts => Configure alerts
    storage => Configure Storage
```

These child contexts correspond to the views available under the Configuration view in the browser, including Network, Services, Users, Preferences, and so on. To select one of these child contexts, type its name:

```
dory:configuration> preferences
dory:configuration preferences>
```

Navigate to a descendant context directly from an ancestor by specifying the intermediate contexts separated with spaces. For example, to navigate directly to configuration preferences from the root context, simply type it:

```
dory:> configuration preferences
dory:configuration preferences>
```

Some child contexts are *dynamic* in that they correspond not to fixed views in the browser, but rather to dynamic entities that have been created by either the user or the system. There are two ways to navigate to these contexts: You can use the `select` command followed by the name of the dynamic context, or surround the name of the dynamic context with double quotes. The names of the dynamic contexts contained within a given context are shown using the `list` command. For example, the users context is a static context, but each user is its own dynamic context.

```
dory:> configuration users
dory:configuration users> list
```

NAME	USERNAME	UID	TYPE
John Doe	bmc	12345	Dir
Super-User	root	0	Loc

To select the user named `bmc`, issue the command `select bmc` or `"bmc"`:

```
dory:configuration users> "bmc"
dory:configuration users bmc>
```

Alternately, double quotes, `select` and `destroy` can in some contexts be used to select an entity based on its properties. For example, one could select log entries issued by the `reboot` module in the maintenance `logs` system context by issuing the following command:

```
dory:maintenance logs system> select module=reboot
dory:maintenance logs system entry-034> show
```

Properties:

```
timestamp = 2016-8-14 06:24:41
module = reboot
priority = crit
text = initiated by root on /dev/console syslogd: going down on signal 15
```

As with other commands, `select` or double quotes may be appended to a context-changing command. For example, to select the user named `bmc` from the root context:


```
dory:> configuration users select bmc
dory:configuration users bmc>
```

Use the `last` command to navigate to a previously selected or created context. The following example creates a replication action, and then uses the `last` and `get id` commands to retrieve the replication action ID. Then a different action is selected, and the `last` and `get id` commands are used to retrieve the ID of the last-visited replication action.

Using `last`, you can return to the last-visited node:

```
dory:configuration net interfaces> "igb4"
dory:configuration net interfaces igb4> done
dory:configuration net interfaces> last
net:configuration net interfaces igb4>
```

The `last` command is also useful to retrieve values that have been automatically set by the appliance during the creation of a dynamic node. For example, each replication action is assigned an ID by the appliance when it is created. Using the `last` command with the `get id` command, you can retrieve the ID without using the name of the replication action:

```
dory:shares p1/share replication> create
dory:shares p1/share action (uncommitted)> set target=dory
                                     target = dory (uncommitted)
dory:shares p1/share action (uncommitted)> set pool=p0
                                     pool = p0 (uncommitted)
dory:shares p1/share action (uncommitted)> commit
dory:shares p1/share replication> last get id
                                     id = 7034367a-d4d8-e26f-fa93-c3b454e3b595
dory:shares p1/share replication>
```

Note that when `last` is combined with another command (in this case, `get id`), the command is run in the context of the last-visited node, but the current node remains unchanged.

Because `last` allows you to retrieve the last-visited node and its values without specifying the name of the node, this command is particularly convenient for scripting:

```
script
    project = 'myproj';
    target = 'mytarget';
    target_pool = 'notmypool';

    run('cd /');
    run('shares select ' + project);
    run('replication');
    run('create');
    set('target', target);
    set('pool', target_pool);
    run('commit');
```

```
run('last');
id = get('id');
printf("Sending update for replication action id %s ...", id);
run('sendupdate');
while (get('state') != 'idle') {
    printf(".");
    run('sleep 1');
}
printf("done\n");
```

To return to the previous context, use the `done` command:

```
dory:configuration> done
dory:>
```

This returns to the previous context, which is not necessarily the parent context, as follows:

```
dory:> configuration users select bmc
dory:configuration users bmc> done
dory:>
```

The `done` command can be used multiple times to backtrack to earlier contexts:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> done
dory:configuration users> done
dory:configuration> done
dory:>
```

To navigate to a parent context, use the `cd` command. Inspired by the classic UNIX command, `cd` takes an argument of `..` to denote moving to the parent context:

```
dory:> configuration users select bmc
dory:configuration users bmc> cd ..
dory:configuration users>
```

And as with the UNIX command, `"cd /"` moves to the root context:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> cd /
dory:>
```

And as with its UNIX analogue, `"cd ../../"` may be used to navigate to the grandparent context:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> cd ../../
dory:configuration>
```

Note that the `cd /` and `cd ..` commands support limited variations. For more versatility, use the `top` command and the `up` command.

Use the `top` command to navigate to the root context:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> top
dory:>
```

Use the `top` command followed by a context name to directly navigate to the specified context relative to the root context. For example, to directly navigate from context `configuration users` to context `configuration services`, use the `top configuration services` command:

```
dory:> configuration
dory:configuration> users
dory:configuration users> top configuration services
dory:configuration services>
```

When the `top` command is used in conjunction with a specific context, the `done` command can be used to navigate back to the context before the `top` command was executed. In the following example, the first `done` command returns to the previous context. The second `done` command returns to the context before the `top` command. The third `done` command returns to the context two nodes before the `top` command.

```
dory:> maintenance system
dory:maintenance system> updates
dory:maintenance system updates> top configuration services
dory:configuration services> ftp
dory:configuration services ftp> done
dory:configuration services> done
dory:maintenance system updates> done
dory:>
```

Like the `cd ..` command, the `up` command can be used to navigate to the parent context:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> up
dory:configuration users>
```

Additionally, you can go to a context *n* nodes up from the current context by repeating the up command *n* times:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> up up
dory:configuration>
```

To go back to a specific context relative to the current parent context, enter the context name after the up command. Likewise, use the up up command followed by a context name to go back to a specific context relative to the current grandparent context. For example, to go from context configuration users bmc to context configuration services, use the command up up services:

```
dory:> configuration
dory:configuration> users
dory:configuration users> select bmc
dory:configuration users bmc> up up services
dory:configuration services>
```

When the up command is used in conjunction with a specific context, the done command can be used to navigate back to the context before the up command was executed. In the following example, the first done command returns to the context before the up command. The second done command returns to the context two nodes before the up command, and the third done command returns to the context three nodes before the up command.

```
dory:> configuration
dory:configuration> services
dory:configuration services> ftp
dory:configuration services ftp> up http
dory:configuration services http> done
dory:configuration services ftp> done
dory:configuration services> done
dory:configuration> done
dory:>
```

Context names will tab complete, be they static contexts (via normal command completion) or dynamic contexts (via command completion of the select command). Following is an example of selecting the user named bmc from the root context with just fifteen keystrokes, instead of the thirty-one that would be required without tab completion:

```
dory:> configtab
dory:> configuration utab
dory:> configuration users setab
dory:> configuration users select tab
bmc    root
```

```
dory:> configuration users select btab
dory:> configuration users select bmcenter
dory:configuration users bmc>
```

Once in a context, execute context-specific commands. For example, to get the current user's preferences, execute the get command from the configuration preferences context:

```
dory:configuration preferences> get
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    session_annotation =
    advanced_analytics = false
```

If there is input following a command that changes context, that command will be executed in the target context, but control will return to the calling context. For example, to get preferences from the root context without changing context, append the get command to the context navigation commands:

```
dory:> configuration preferences get
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    session_annotation =
    advanced_analytics = false
```

When creating a new entity in the system, the context associated with the new entity will often be created in an *uncommitted* state. For example, create a threshold alert by executing the create command from the configuration alerts threshold context:

```
dory:> configuration alerts thresholds create
dory:configuration alerts threshold (uncommitted)>
```

The (uncommitted) in the prompt denotes that this is an uncommitted context. An uncommitted entity is committed via the commit command; any attempt to navigate away from the uncommitted context will prompt for confirmation:

```
dory:configuration alerts threshold (uncommitted)> cd /
Leaving will abort creation of "threshold". Are you sure? (Y/N)
```

When committing an uncommitted entity, the properties associated with the new entity will be validated, and an error will be generated if the entity cannot be created. For example, the creation of a new threshold alert requires the specification of a statistic name; failure to set this name results in an error:

```
dory:configuration alerts threshold (uncommitted)> commit
error: missing value for property "statname"
```

To resolve the problem, address the error and reattempt the commit:

```
dory:configuration alerts threshold (uncommitted)> set statname=cpu.utilization
      statname = cpu.utilization (uncommitted)
dory:configuration alerts threshold (uncommitted)> commit
error: missing value for property "limit"
dory:configuration alerts threshold (uncommitted)> set limit=90
      limit = 90 (uncommitted)
dory:configuration alerts threshold (uncommitted)> commit
dory:configuration alerts thresholds> list
THRESHOLD      LIMIT      TYPE  STATNAME
threshold-000      90      normal cpu.utilization
```

Related Topics

- [“Command Line Interface \(CLI\)” on page 33](#)
- [“CLI Properties” on page 42](#)

CLI Properties

Properties are typed name/value pairs that are associated with a context. Properties for a given context can be ascertained by running the "help properties" command. Following is an example of retrieving the properties associated with a user's preferences:

```
dory:configuration preferences> help properties
Properties that are valid in this context:

locale           => Locality

login_screen     => Initial login screen

session_timeout  => Session timeout

session_annotation => Current session annotation

advanced_analytics => Make available advanced analytics statistics
```

The properties of a given context can be retrieved with the get command. Following is an example of using the get command to retrieve a user's preferences:

```
dory:configuration preferences> get
      locale = C
      login_screen = status/dashboard
      session_timeout = 15
      session_annotation =
      advanced_analytics = false
```

The `get` command will return any properties provided to it as arguments. For example, to get the value of the `login_screen` property:

```
dory:configuration preferences> get login_screen
login_screen = status/dashboard
```

The `get` command will tab complete with the names of the available properties. For example, to see a list of available properties for the iSCSI service:

```
dory:> configuration services iscsi get tab
<status>          isns_server          radius_secret      target_chap_name
isns_access        radius_access        radius_server      target_chap_secret
```

The `select` command, or a command surrounded by double quotes, will select a dynamic node by property. For example, to select `key-000` by user:

```
hostname:configuration services sftp keys> show
Keys:

NAME      MODIFIED          CIPHER  USER  COMMENT
key-000  2015-6-5 19:48:23  RSA     u1     1

hostname:configuration services sftp keys> "user=u1"
hostname:configuration services sftp key-000>
```

The `set` command will set a property to a specified value, with the property name and its value separated by an equals sign. For example, to set the `login_screen` property to be `"shares"`:

```
dory:configuration preferences> set login_screen=shares
login_screen = shares (uncommitted)
```

Note that in the case of properties that constitute state on the appliance, setting the property does *not* change the value, but rather records the set value and indicates that the value of the property is uncommitted.

To force set property values to take effect, they must be explicitly committed, allowing multiple values to be changed as a single, coherent change. To commit any uncommitted property values, use the `commit` command:

```
dory:configuration preferences> get login_screen
login_screen = shares (uncommitted)
dory:configuration preferences> commit
dory:configuration preferences> get login_screen
login_screen = shares
```

If you attempt to leave a context that contains uncommitted properties, you will be warned that leaving will abandon the set property values, and will be prompted to confirm that you wish to leave. For example:

```
dory:configuration preferences> set login_screen=maintenance/hardware
      login_screen = maintenance/hardware (uncommitted)
dory:configuration preferences> done
You have uncommitted changes that will be discarded. Are you sure? (Y/N)
```

If a property in a context is set from a different context -- that is, if the set command has been appended to a command that changes context -- the commit is *implied*, and happens before control is returned to the originating context. For example:

```
dory:> configuration preferences set login_screen=analytics/worksheets
      login_screen = analytics/worksheets
dory:>
```

Some properties take a list of values. For these properties, the list elements should be separated by a comma. For example, the NTP servers property may be set to a list of NTP servers:

```
dory:configuration services ntp> set servers=0.pool.ntp.org,1.pool.ntp.org
      servers = 0.pool.ntp.org,1.pool.ntp.org (uncommitted)
dory:configuration services ntp> commit
```

If a property value contains a comma, an equals sign, a quote or a space, the entire value must be double quoted. For example, the sharenfs shares property for the default project may be set to read-only, but provide read/write access to host kiowa. For more information, see [“Shares and Projects” on page 383](#).

```
dory:> shares select default
dory:shares default> set sharenfs="ro,rw=kiowa"
      sharenfs = ro,rw=kiowa (uncommitted)
dory:shares default> commit
```

Some properties are immutable; you can get their values, but you cannot set them. Attempts to set an immutable property results in an error. For example, attempting to set the immutable space_available property of the default project. For more information, see [“Shares and Projects” on page 383](#).

```
dory:> shares select default
dory:shares default> get space_available
      space_available = 1.15T
dory:shares default> set space_available=100P
error: cannot set immutable property "space_available"
```

Some other properties are only immutable in certain conditions. For these properties, the set command is not valid. For example, if the user named bmc is a network user, the fullname property will be immutable:

```
dory:> configuration users select bmc set fullname="Rembrandt Q. Einstein"
error: cannot set immutable property "fullname"
```

Related Topics

- [“Browser User Interface \(BUI\)” on page 22](#)
- [“Command Line Interface \(CLI\)” on page 33](#)

Working with CLI Scripting

The CLI is designed to provide a powerful scripting environment for performing repetitive tasks.

You can use Batching Commands or Scripting Commands (or some combination), but in any case the automated infrastructure requires automated access to the appliance. This must be done by user configuration, user authorizations, and setting SSH public keys using the CLI.

For information about configuring users, see the following:

- [“Configuring Users” on page 202](#)
- [“User Authorizations” on page 219](#)
- [“Setting SSH Public Keys \(CLI\)” on page 226](#)

To use CLI scripting, use the following sections:

- [Using Batch Commands](#)
- [Understanding the CLI Scripting Commands](#)
- [Accessing the CLI Script Environment](#)
- [Understanding the Built-in CLI Functions](#)
- [Using the Run Function](#)
- [Using the Get Function](#)
- [Using the List Function](#)
- [Using the Children Function](#)
- [Using the Choices Function](#)
- [Using the Functions for Generating Output](#)
- [Understanding CLI Scripting Errors](#)

Using Batch Commands

The simplest scripting mechanism is to batch appliance shell commands. For example, to automatically take a snapshot called "newsnap" in the project "myproj" and the filesystem "myfs", put the following commands in a file:

```
shares
select myproj
```

```
select myfs
snapshots snapshot newsnap
```

Then ssh onto the appliance, redirecting standard input to be the file:

```
% ssh root@dory < myfile.txt
```

In many shells, you can abbreviate this by using a "here file", where input up to a token is sent to standard input. Following is the above example in terms of a here file:

```
% '''ssh root@dory << EOF
shares
select myproj
select myfs
snapshots snapshot newsnap
EOF'''
```

This mechanism is sufficient for the simplest kind of automation, and may be sufficient if wrapped in programmatic logic in a higher-level shell scripting language on a client, but it generally leaves much to be desired.

Understanding the CLI Scripting Commands

While batching commands is sufficient for the simplest of operations, it can be tedious to wrap in programmatic logic. For example, if you want to get information on the space usage for every share, you must have many different invocations of the CLI, wrapped in a higher level language on the client that parsed the output of specific commands. This results in slow, brittle automation infrastructure. To allow for faster and most robust automation, the appliance has a rich *scripting environment* based on ECMAScript 3. An ECMAScript tutorial is beyond the scope of this document, but it is a dynamically typed language with a C-like syntax that allows for:

- Conditional code flow (if/else)
- Iterative code flow (while, for, etc.)
- Structural and array data manipulation via first-class Object and Array types
- Perl-like regular expressions and string manipulation (split(), join(), etc.)
- Exceptions
- Sophisticated functional language features like closures

▼ Accessing the CLI Script Environment

1. In the CLI, enter the script environment using the `script` command:

```
dory:> script
("." to run)>
```

2. **At the script environment prompt, you can input your script, finally entering "." alone on a line to execute it:**

```
dory:> script
("." to run)> for (i = 10; i > 0; i--)
("." to run)>     printf("%d... ", i);
("." to run)> printf("Blastoff!\n");
("." to run)> .
10... 9... 8... 7... 6... 5... 4... 3... 2... 1... Blastoff!
```

3. **If your script is a single line, you can simply provide it as an argument to the script command, making for an easy way to explore scripting:**

```
dory:> script print("It is now " + new Date())
It is now Tue Oct 14 2009 05:33:01 GMT+0000 (UTC)
```

Understanding the Built-in CLI Functions

Of course, scripts are of little utility unless they can interact with the system at large. There are several built-in functions that allow your scripts to interact with the system:

TABLE 9 Built-in Functions to Support System Interactions

Function	Description
get	Gets the value of the specified property. Note that this function returns the value in native form, e.g. dates are returned as Date objects.
list	Returns an array of tokens corresponding to the dynamic children of the current context.
run	Runs the specified command in the shell, returning any output as a string. Note that if the output contains multiple lines, the returned string will contain embedded newlines.
props	Returns an array of the property names for the current node.
set	Takes two string arguments, setting the specified property to the specified value.
choices	Returns an array of the valid property values for any property for which the set of values is known and enumerable.

▼ Using the Run Function

1. **The simplest way for scripts to interact with the larger system is to use the "run" function: it takes a command to run, and returns the output of that command as a string. For example:**

```
dory:> configuration version script dump(run('get boot_time'))
'
          boot_time = 2009-10-12 07:02:17\n'
```

2. **The built-in dump function dumps the argument out, without expanding any embedded newlines. ECMAScript's string handling facilities can be used to take apart output. For example, splitting the above based on whitespace:**

```
dory:> configuration version script dump(run('get boot_time').split(/\s+/))
[&#39;', 'boot_time', '=', '2009-10-12', '07:02:17', &#39;]
```

▼ Using the Get Function

The run function is sufficiently powerful that it may be tempting to rely exclusively on parsing output to get information about the system -- but this has the decided disadvantage that it leaves scripts parsing human-readable output that may or may not change in the future. To more robustly gather information about the system, use the built-in "get" function. In the case of the boot_time property, this will return not the string but rather the ECMAScript Date object, allowing the property value to be manipulated programmatically.

1. **For example, you might want to use the boot_time property in conjunction with the current time to determine the time since boot:**

```
script
run('configuration version');
now = new Date();
uptime = (now.valueOf() - get('boot_time').valueOf()) / 1000;
printf('up %d day%s, %d hour%s, %d minute%s, %d second%s\n',
      d = uptime / 86400, d < 1 || d >= 2 ? 's' : '',
      h = (uptime / 3600) % 24, h < 1 || h >= 2 ? 's': '',
      m = (uptime / 60) % 60, m < 1 || m >= 2 ? 's': '',
      s = uptime % 60, s < 1 || s >= 2 ? 's': '');
```

2. **Assuming the above is saved as a "uptime.aksh", you could run it this way:**

```
% ssh root@dory < uptime.aksh
Pseudo-terminal will not be allocated because stdin is not a terminal.
Password:
up 2 days, 10 hours, 47 minutes, 48 seconds
```

The message about pseudo-terminal allocation is due to the ssh client; the issue that this message refers to can be dealt with by specifying the "-T" option to ssh.

▼ Using the List Function

In a context with dynamic children, it can be very useful to iterate over those children programmatically. This can be done by using the `list` function, which returns an array of dynamic children.

1. **The following example script iterates over every share in every project, printing out the amount of space consumed and space available:**

```
script
  run('shares');
  projects = list();

  for (i = 0; i < projects.length; i++) {
    run('select ' + projects[i]);
    shares = list();

    for (j = 0; j < shares.length; j++) {
      run('select ' + shares[j]);
      printf("%s/%s %1.64g %1.64g\n", projects[i], shares[j],
        get('space_data'), get('space_available'));
      run('cd ..');
    }

    run('cd ..');
  }
}
```

2. **Here is the output of running the script, assuming it were saved to a file named "space.aksh":**

```
% ssh root@koi < space.aksh
Password:
admin/accounts 18432 266617007104
admin/exports 18432 266617007104
admin/primary 18432 266617007104
admin/traffic 18432 266617007104
admin/workflow 18432 266617007104
aleventhal/hw_eng 18432 266617007104
bcantrill/analytx 1073964032 266617007104
bgregg/dashbd 18432 266617007104
bgregg/filesys01 26112 107374156288
```

```
bpijewski/access_ctrl 18432 266617007104
...
```

3. **If one would rather a "pretty printed" (though more difficult to handle programmatically) variant of this, one could directly parse the output of the `get` command:**

```
script
run('shares');
projects = list();

printf('%-40s %-10s %-10s\n', 'SHARE', 'USED', 'AVAILABLE');

for (i = 0; i < projects.length; i++) {
    run('select ' + projects[i]);
    shares = list();

    for (j = 0; j < shares.length; j++) {
        run('select ' + shares[j]);

        share = projects[i] + '/' + shares[j];
        used = run('get space_data').split(/\s+/)[3];
        avail = run('get space_available').split(/\s+/)[3];

        printf('%-40s %-10s %-10s\n', share, used, avail);
        run('cd ..');
    }

    run('cd ..');
}
```

4. **Here is the output of running this new script, assuming it were named "prettyspace.aksh":**

```
% ssh root@koi < prettyspace.aksh
Password:
SHARE                                USED      AVAILABLE
admin/accounts                      18K       248G
admin/exports                       18K       248G
admin/primary                       18K       248G
admin/traffic                       18K       248G
admin/workflow                      18K       248G
aleventhal/hw_eng                   18K       248G
bcantrill/analytx                   1.00G     248G
bgregg/dashbd                      18K       248G
bgregg/filesys01                   25.5K     100G
bpijewski/access_ctrl              18K       248G
...
```

5. The list function supports optional arguments depth and filter.

The format is: `list ([depth, [filter]])`. The argument `depth` can be defined by a number. The greater number of `depth`, the more details will be returned. The argument `filter` is formatted as `{<prop1>:<val1>, <prop2>:<val2> ...}`. If `filter` is specified, `depth` must also be specified.

Usage and input behavior:

- `list()` - Returns only node names.
- `list(0)` - Return properties of node and only children names.
- `list(0, {kiosk_mode: true})` - Return a filtered list for `kiosk_mode` is `true` with names of children.
- `list(1)` - Return properties of node, names and properties of children, only names of grandchildren.
- `list(1, {kiosk_mode: true})` - Return a filtered list for `kiosk_mode` is `true` with details up to `depth=1`.
- `list(2)` - Return properties of node, names and properties of children and `list(0)` output of grandchildren.
- `list(2, {fullname: 'Super*', kiosk_mode: true})` - Return a filtered list for `fullname` containing `Super` and `kiosk_mode` is `true` with details up to `depth=2`.

6. This is an example output for a list with depth=2:

The label `name` shows the name of the list item (that is, a node). The label `properties` shows the properties of the list item. The label `children` shows static children of the list item. The label `list` shows dynamic children of the list item.

```
script
("." to run)> dump(list(2));
("." to run)> .

[
  {
    name: 'restuser',
    properties: {
      kiosk_screen: 'status/dashboard',
      kiosk_mode: false,
      roles: ['basic'],
      require_annotation: false,
      initial_password: 'DummyPassword',
      fullname: 'REST User',
      logname: 'restuser'
    },
    children: [
      {
        name: 'preferences',
```

```
        properties: {
            advanced_analytics: false,
            session_timeout: 15,
            login_screen: 'status/dashboard',
            locale: 'C'
        }
    }, {
        name: 'exceptions',
        list: [{
            name: 'auth-000',
            properties: {
                allow_configure: false,
                scope: 'alert'
            }
        }], {
            name: 'auth-001',
            properties: {
                allow_workgroup: false,
                allow_domain: false,
                name: '*',
                scope: 'ad'
            }
        }
    ]
}]
}]
```

▼ Using the Children Function

Even in a context with static children, it can be useful to iterate over those children programmatically. This can be done by using the `children` function, which returns an array of static children.

1. **For example, here's a script that iterates over every service, printing out the status of the service:**

```
configuration services
script
    var svcs = children();
    for (var i = 0; i < svcs.length; ++i) {
        run(svcs[i]);
        try {
            printf("%-10s %s\n", svcs[i], get('<status>'));
        } catch (err) { }
        run("done");
    }
```



```
}
```

2. **Here's the output of running the script, assuming it were saved to a file named "svcinfo.aksh":**

```
% ssh root@koi < space.aksh
Password:
cifs      disabled
dns       online
ftp       disabled
http      disabled
identity  online
idmap     online
ipmp      online
iscsi     online
ldap      disabled
ndmp      online
nfs       online
nis       online
ntp       online
scrk      online
sftp      disabled
smtp      online
snmp      disabled
ssh       online
tags      online
vscan     disabled
```

▼ Using the Choices Function

The choices function returns an array of the valid property values for any property for which the set of values is known and enumerable. For example, the following script retrieves the list of all pools on the shares node using the choices function and then iterates all pools to list projects and shares along with the available space.

1. **For example, the following script retrieves the list of all pools on the shares node using the choices function and then iterates all pools to list projects and shares along with the available space.**

```
fmt = '%-40s %-15s %-15s\n';
printf(fmt, 'SHARE', 'USED', 'AVAILABLE');
run('cd /');
run('shares');
pools = choices('pool');
for (p = 0; p < pools.length; p++) {
```

```
set('pool', pools[p]);
projects = list();
for (i = 0; i < projects.length; i++) {
    run('select ' + projects[i]);
    shares = list();
    for (j = 0; j < shares.length; j++) {
        run('select ' + shares[j]);
        share = pools[p] + ':' + projects[i] + '/' + shares[j];
        printf(fmt, share, get('space_data'),
            get('space_available'));
        run('cd ..');
    }
    run('cd ..');
}
```

2. Here is the output of running the script:

SHARE	USED	AVAILABLE
pond:projectA/fs1	31744	566196178944
pond:projectA/fs2	31744	566196178944
pond:projectB/lun1	21474836480	587670999040
puddle:deptA/share1	238475	467539219283
puddle:deptB/share1	129564	467539219283
puddle:deptB/share2	19283747	467539219283

Using the Functions for Generating Output

Reporting state on the system requires generating output. Scripts have several built-in functions made available to them to generate output:

TABLE 10 Built-in Functions for Generating Output

Function	Description
dump	Dumps the specified argument to the terminal, without expanding embedded newlines. Objects will be displayed in a JSON-like format. Useful for debugging.
print	Prints the specified object as a string, followed by a newline. If the object does not have a toString method, it will be printed opaquely.
printf	Like C's printf(3C), prints the specified arguments according to the specified formatting string.

Understanding CLI Scripting Errors

When an error is generated, an exception is thrown. The exception is generally an object that contains the following members:

- `code` - a numeric code associated with the error
- `message` - a human-readable message associated with the error

Exceptions can be caught and handled, or they may be thrown out of the script environment. If a script environment has an uncaught exception, the CLI will display the details. For example:

```
dory:> script run('not a cmd')
error: uncaught error exception (code EAKSH_BADCMD) in script: invalid command
      "not a cmd" (encountered while attempting to run command "not a cmd")
```

You could see more details about the exception by catching it and dumping it out:

```
dory:> script try { run('not a cmd') } catch (err) { dump(err); }
{
  toString: <function>,
  code: 10004,
  message: 'invalid command "not a cmd" (encountered while attempting to
           run command "not a cmd")'
}
```

This also allows you to have rich error handling, for example:

```
#!/usr/bin/ksh -p

ssh -T root@dory <<EOF
script
  try {
    run('shares select default select $1');
  } catch (err) {
    if (err.code == EAKSH_ENTITY_BADSELECT) {
      printf('error: "$1" is not a share in the ' +
            'default project\n');
      exit(1);
    }

    throw (err);
  }

  printf('default/$1: compression is %s\n', get('compression'));
  exit(0);
EOF
```

If this script is named "share.ksh" and run with an invalid share name, a rich error message will be generated:

```
% ksh ./share.ksh bogus
error: "bogus" is not a share in the default project
```

Configuring the Appliance

To configure the appliance, use the following sections:

- [“Initial Appliance Configuration” on page 57](#)
- [“Appliance Cluster Configuration” on page 58](#)
- [“Network Configuration” on page 89](#)
- [“Configuring Storage” on page 122](#)
- [“Understanding the Appliance Status” on page 155](#)
- [“Configuring Storage Area Network \(SAN\)” on page 170](#)
- [“Configuring Users” on page 202](#)
- [“Setting Appliance Preferences” on page 224](#)
- [“Configuring Alerts” on page 228](#)
- [“Configuring Certificates” on page 235](#)

Initial Appliance Configuration

If you are setting up a new appliance, follow the initial configuration steps in [“Configuring the Appliance for the First Time” in *Oracle ZFS Storage Appliance Installation Guide*](#).

You can repeat initial configuration at a later time by clicking the INITIAL SETUP button on the Maintenance > System screen, or by entering the maintenance system setup context in the CLI.

Related Topics

- [“Appliance Cluster Configuration” on page 58](#)
- [“Network Configuration” on page 89](#)
- [“Configuring Storage” on page 122](#)

Appliance Cluster Configuration

The Oracle ZFS Storage Appliance product supports cooperative clustering of appliances. This strategy can be part of an integrated approach to availability enhancement that may also include client-side load balancing, proper site planning, proactive and reactive maintenance and repair, and the single-appliance hardware redundancy built into all appliances.

Note - If you are configuring clustering for two new controllers, follow the procedure [“Configuring the Appliance for the First Time” in *Oracle ZFS Storage Appliance Installation Guide*](#).

For tasks related to appliance clustering, see:


- [“Connecting Cluster Cables” in *Oracle ZFS Storage Appliance Cabling Guide*](#)
- [“Cluster Configuration BUI View” on page 58](#)
- [“Upgrading a Standalone Appliance to a Clustered Configuration \(BUI\)” on page 60](#)
- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

For a better understanding about appliance clustering, see:

- [“Cluster Terminology” on page 66](#)
- [“Understanding Clustering” on page 66](#)
- [“Cluster Advantages and Disadvantages” on page 68](#)
- [“Cluster Interconnect I/O” on page 70](#)
- [“Cluster Resource Management” on page 71](#)
- [“Cluster Takeover and Failback” on page 74](#)
- [“Configuration Changes in a Clustered Environment” on page 76](#)
- [“Clustering Considerations for Storage” on page 77](#)
- [“Clustering Considerations for Networking” on page 79](#)
- [“Private Local IP Interfaces” on page 81](#)
- [“Clustering Considerations for InfiniBand” on page 82](#)
- [“Preventing Split-Brain Conditions” on page 85](#)
- [“Estimating and Reducing Takeover Impact” on page 87](#)

Cluster Configuration BUI View

The Configuration > Cluster view provides a graphical overview of the status of the cluster card, the cluster controller node states, and all of the resources.

FIGURE 1 Cluster Connections

Note - Cluster cables must be connected between the two controllers to see the three solid line connections in the BUI. For cluster cabling details, see [“Connecting Cluster Cables” in Oracle ZFS Storage Appliance Cabling Guide](#).

The interface contains the following objects:

- A thumbnail picture of each system, with the system whose administrative interface is being accessed shown at left. Each thumbnail is labeled with the canonical appliance name, and its current cluster state (the icon above, and a descriptive label).
- A thumbnail of each cluster card connection that dynamically updates with the hardware: a solid line connects a link when that link is connected and active, and the line disappears if that connection is broken or while the other system is restarting/rebooting.
- A list of the PRIVATE and SINGLETON resources currently assigned to each system, shown in lists below the thumbnail of each cluster node, along with various attributes of the resources.
- For each resource, the appliance to which that resource is assigned (that is, the appliance that will provide the resource when both are in the CLUSTERED state). When the current appliance is in the OWNER state, the owner field is shown as a pop-up menu that can be edited and then committed by clicking Apply.
- For each resource, a lock icon  indicating whether or not the resource is PRIVATE. When the current appliance is in either of the OWNER or CLUSTERED states, a resource can be locked to it (made PRIVATE) or unlocked (made a SINGLETON) by clicking the lock icon  and then clicking Apply. Note that PRIVATE resources belonging to the remote peer will not be displayed on either resource list.

The BUI contains the following buttons:

TABLE 11 Cluster Interface Buttons

Button	Description
Setup	If the cluster is not yet configured, execute the cluster setup guided task, and then return to the current screen.

Button	Description
Unconfig	Upgrade a node to standalone operation by unconfiguring the cluster.
Apply	If resource modifications are pending (rows highlighted in yellow), commit those changes to the cluster.
Revert	If resource modifications are pending (rows highlighted in yellow), revert those changes and show the current cluster configuration.
Failback	If the current appliance (left-hand side) is the OWNER, fail-back resources owned by the other appliance to it, leaving both nodes in the CLUSTERED state (active/active).
Takeover	If the current appliance (left-hand side) is either CLUSTERED or STRIPPED, force the other appliance to reboot, and take-over its resources, making the current appliance the OWNER.

Related Topics

- [“Performing Initial Configuration \(BUI\)” in *Oracle ZFS Storage Appliance Installation Guide*](#)
- [“Upgrading a Standalone Appliance to a Clustered Configuration \(BUI\)” on page 60](#)
- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

▼ Upgrading a Standalone Appliance to a Clustered Configuration (BUI)

Use this procedure to upgrade a standalone appliance to a clustered configuration.

Note - It is strongly recommended that you use the BUI to configure clustered controllers.

Before You Begin Check for the following:

- The second controller is a new controller or a controller that has been reset to the factory settings. See [“Performing a Factory Reset” in *Oracle ZFS Storage Appliance Customer Service Manual*](#).
- Both controllers must be the same model. Note that the 7420 (with 2GHz or 2.40GHz CPUs) is based on the same platform and can be clustered with the 7420 (with 1.86GHz or 2.00 GHz CPUs).
- The standalone appliance is powered on. There is no need to power down the standalone appliance during this procedure.

1. **Connect the cluster cables between the standalone appliance and second controller.**
For cluster cabling details, see [“Connecting Cluster Cables” in Oracle ZFS Storage Appliance Cabling Guide](#).
2. **On the second controller, connect the power cables into power supply 0 and power supply 1. Then connect each cable to the external power source.**
The second controller powers on automatically.
3. **Connect the second controller to the disk shelves.**
See the documentation that came with your appliance or refer to [“Getting Started with Cabling” in Oracle ZFS Storage Appliance Cabling Guide](#).
4. **On the standalone controller, go to Configuration > Cluster.**
5. **Confirm that the communication links between the two controllers are connected and active.**
If three solid lines are not shown, ensure that the three cluster cables are properly connected and secure in their connectors.
6. **Click SETUP.**
7. **Enter the host name for the second controller and the same root password that is set on the first controller.**

Note - An initial cluster configuration setup can take several minutes to complete.

8. **On the standalone controller, go to Configuration > Cluster and click the lock icon  for the management interface.**
Locking the management interface to the controller will prevent a transfer of resources when a failback occurs.
9. **From the standalone controller, configure the management interface for the second controller.**
 - a. **Go to Configuration > Network and click the add icon  next to Interfaces.**
 - b. **Enter a name for the management interface, and check the boxes for Enable Interface and Allow Administration.**
 - c. **Select an IP address and click APPLY.**

10. **Go to Configuration > Cluster and click FAILBACK to bring the cluster to Active: Active mode.**
The two controllers are now configured as clustered peers.
11. **On the second controller, go to Configuration > Cluster and click the lock icon  for the management interface.**

Related Topics


- [“Clustering Considerations for Storage” on page 77](#)
- [“Cluster Configuration BUI View” on page 58](#)

▼ Shutting Down a Clustered Configuration (BUI)

Use this procedure to shut down a clustered configuration.

1. **From one of the peer controllers, go to Configuration > Cluster.**
2. **Determine the state of both controllers.**

In the following figure, the active controller is controller-a, and the standby controller is controller-b.



Use the following table to determine the state of each controller.

controller-a	controller-b	Condition
 Active	 Active	Both controllers are running in a normal clustered condition.

controller-a	controller-b	Condition
 Active (takeover completed)	 Ready (waiting for failback)	controller-a owns all of the resources and is the active controller. controller-b, is in standby mode and has no resources. To limit the number of times a pool is moved, shut down the standby controller first.
 Active (takeover completed)	 Rejoining cluster ...	controller-b is rebooting and controller-a has all resources.
 Active (takeover completed)	Unknown (disconnected or restarting)	controller-b is powered off or rebooting, all of its cluster interconnect links are down, or clustering has not yet been configured.

3. **Log in to the BUI of Controller B and click the power icon  on the left side under the masthead.**

Note - To limit the number of times a pool is moved, shut down the standby controller first.

4. **From the BUI of Controller A, go to Configuration > Cluster to confirm that Controller B is powered off, with the cluster state: Unknown (disconnected or restarting).**
5. **From the BUI of Controller A, click the power icon  on the left side under the masthead.**
6. **(Optional) To confirm that both controllers are powered off, log into the Oracle ILOM and enter :**

```
->show /SYS power_state
```

For information about accessing ILOM, see [“Logging in to Oracle ILOM Remotely Using a Command Line Interface” in Oracle ZFS Storage Appliance Customer Service Manual.](#)

7. **Power off the disk shelves.**
 - a. **Place the power supply on/off switches to the "O" off position.**
 - b. **Disconnect the power cords from the external power source for the cabinet.**

Note - All power cords must be disconnected to completely remove power from the disk shelf.

For more information, see [“Powering Off a Disk Shelf” in Oracle ZFS Storage Appliance Installation Guide](#).

▼ Shutting Down a Clustered Configuration (CLI)

Use this procedure to shut down a clustered configuration.

Note - For the purpose of this procedure, the clustered controllers are referred to as controller-a and controller-b.

1. Verify the cluster state of each controller, using the following commands:

In the following example, controller-a is the owner and in the active state. Its peer, controller-b, is the standby controller and in the stripped state.

```
controller-a:>configuration cluster
controller-a:configuration cluster> show
state = AKCS_OWNER
description = Active (takeover completed)
peer_asn = 365ed33c-3b9d-c533-9349-8014e9da0408
peer_hostname = controller-b
peer_state = AKCS_STRIPPED
peer_description = Ready (waiting for failback)
```

2. Use the following table to verify the status of each controller:

controller-a	controller-b	Condition
AKCS_CLUSTERED	AKCS_CLUSTERED	Both controllers are running in a normal clustered condition.
AKCS_OWNER	AKCS_STRIPPED	controller-a owns all of the resources and is the active controller. controller-b, is in standby mode and has no resources. To limit the number of times a pool is moved, shut down the STRIPPED controller first.
AKCS_OWNER	rebooting	controller-b is rebooting and controller-a has all resources.
AKCS_OWNER	unknown	controller-b is powered off or rebooting, all of its cluster interconnect links are down,

controller-a	controller-b	Condition
		or clustering has not yet been configured.

Note - If the status of each controller does **not** agree, the cluster may be experiencing a problem. Contact Oracle Support before proceeding.

3. Shut down the controller-b, using the following commands:

```
controller-b:configuration cluster> cd /
controller-b:> maintenance system poweroff
This will turn off power to the appliance. Are you sure? (Y/N)Y
```

Note - If both controllers have a status of AKCS_CLUSTERED, a takeover of the surviving controller begins automatically.

4. From controller-a, use the show command to verify that controller-b has been powered off and is in state OWNER/unknown.

```
controller-a:configuration cluster> show
state = AKCS_OWNER
description = Active (takeover completed)
peer_asn = 365ed33c-3b9d-c533-9349-8014e9da0408
peer_hostname = controller-b
peer_state = OWNER/unknown
peer_description =
```

5. Shut down controller-a using the following commands:

```
controller-a:configuration cluster> cd /
controller-a:> maintenance system poweroff
This will turn off power to the appliance. Are you sure? (Y/N) Y
```

6. (Optional) To confirm that both controllers are powered off, log into the Oracle ILOM and enter:

```
->show /SYS power_state
```

For information about accessing ILOM, see [“Logging in to Oracle ILOM Remotely Using a Command Line Interface”](#) in *Oracle ZFS Storage Appliance Customer Service Manual*.

7. Power off the disk shelves.

- a. Place the power supply on/off switches to the "O" off position.

- b. **Disconnect the power cords from the external power source for the cabinet.**

Note - All power cords must be disconnected to completely remove power from the disk shelf.

For more information, see [“Powering Off a Disk Shelf” in Oracle ZFS Storage Appliance Installation Guide](#).

Related Topics

- [“Understanding Clustering” on page 66](#)
- [“Cluster Terminology” on page 66](#)
- [“Cluster Resource Management” on page 71](#)
- [“Cluster Takeover and Failback” on page 74](#)
- [“Configuration Changes in a Clustered Environment” on page 76](#)

Cluster Terminology

The terms defined here are used throughout the document. In most cases, they are explained in greater context and detail along with the broader concepts involved. The cluster states and resource types are described in the next section. Refer back to this section for reference as needed.


- **Export:** the process of making a resource inactive on a particular controller
- **Failback:** the process of moving from AKCS_OWNER state to AKCS_CLUSTERED, in which all foreign resources (those assigned to the peer) are exported, then imported by the peer
- **Import:** the process of making a resource active on a particular controller
- **Peer:** the other appliance in a cluster
- **Rejoin:** to retrieve and resynchronize the resource map from the peer
- **Resource:** a physical or virtual object present, and possibly active, on one or both controllers
- **Takeover:** the process of moving from AKCS_CLUSTERED or AKCS_STRIPPED state to AKCS_OWNER, in which all resources are imported

Understanding Clustering

The clustering subsystem incorporated into the series consists of three main building blocks (see the following figure). The cluster I/O subsystem and the hardware device provide

a transport for inter-controller communication within the cluster and are responsible for monitoring the peer's state. This transport is used by the resource manager, which allows data service providers and other management subsystems to interface with the clustering system. Finally, the cluster management user interfaces provide the setup task, resource allocation and assignment, monitoring, and takeover and failback operations. Each of these building blocks is described in detail in the following sections.

FIGURE 2 Clustering Subsystem



Unconfiguring Clustering

Unconfiguring clustering is a destructive operation that returns the clustered controllers to standalone controllers. There are two reasons to unconfigure clustering:

- You no longer wish to use clustering; instead, you wish to configure two independent storage appliances.
- You are replacing a failed storage controller with new hardware or a storage controller with factory-fresh appliance software (typically this replacement is performed by your service provider).

The peer node must be turned off before unconfiguration can occur. The peer node must be factory reset before being used again in the same clustered configuration.



Caution - Because unconfiguring a cluster may result in data loss, contact Oracle support.

Cluster Advantages and Disadvantages

It is important to understand the scope of the Oracle ZFS Storage Appliance clustering implementation. The term 'cluster' is used in the industry to refer to many different technologies with a variety of purposes. We use it here to mean a metasystem comprised of two appliance controllers and shared storage, used to provide improved availability in the case in which one of the controllers succumbs to certain hardware or software failures. A cluster contains exactly two appliances or storage controllers, referred to for brevity throughout this document as *controllers*. Each controller may be assigned a collection of storage, networking, and other resources from the set available to the cluster, which allows the construction of either of two major topologies. Many people use the terms *active-active* to describe a cluster in which there are two (or more) storage pools, one of which is assigned to each controller along with network resources used by clients to reach the data stored in that pool, and *active-passive* to refer to which a single storage pool is assigned to the controller designated as *active* along with its associated network interfaces. Both topologies are supported by the appliance. The distinction between these is artificial; there is no software or hardware difference between them and one can switch at will simply by adding or destroying a storage pool. In both cases, if a controller fails, the other (its *peer*) will take control of all known resources and provide the services associated with those resources.

As an alternative to incurring hours or days of downtime while the controller is repaired, clustering allows a peer appliance to provide service while repair or replacement is performed. In addition, clusters support rolling upgrade of software, which can reduce the business disruption associated with migrating to newer software. Some clustering technologies have certain additional capabilities beyond availability enhancement; the Oracle ZFS Storage Appliance clustering subsystem was not designed to provide these. In particular, it does not provide for load balancing among multiple controllers, improve availability in the face of storage failure, offer clients a unified filesystem namespace across multiple appliances, or divide service responsibility across a wide geographic area for disaster recovery purposes. These functions are likewise outside the scope of this document; however, the appliance and the data protocols it offers support numerous other features and strategies that can improve availability:

- Replication of data, which can be used for disaster recovery at one or more geographically remote sites
- Client-side mirroring of data, which can be done using redundant iSCSI LUNs provided by multiple arbitrarily located storage servers
- Load balancing, which is built into the NFS protocol and can be provided for some other protocols by external hardware or software (applies to read-only data)

- Redundant hardware components including power supplies, network devices, and storage controllers
- Fault management software that can identify failed components, remove them from service, and guide technicians to repair or replace the correct hardware
- Network fabric redundancy provided by LACP and IPMP functionality
- Redundant storage devices (RAID)

Additional information about other availability features can be found in the appropriate sections of this document.

When deciding between a clustered and standalone Oracle ZFS Storage Appliance configuration, it is important to weigh the costs and benefits of clustered operation. It is common practice throughout the IT industry to view clustering as an automatic architectural decision, but this thinking reflects an idealized view of clustering risks and rewards promulgated by some vendors in this space. In addition to the obvious higher up-front and ongoing hardware and support costs associated with the second controller, clustering also imposes additional technical and operational risks. Some of these risks can be mitigated by ensuring that all personnel are thoroughly trained in cluster operations; others are intrinsic to the concept of clustered operation. Such risks include:

- The potential for application intolerance of protocol-dependent behaviors during takeover,
- The possibility that the cluster software itself will fail or induce a failure in another subsystem that would not have occurred in standalone operation,
- Increased management complexity and a higher likelihood of operator error when performing management tasks,
- The possibility that multiple failures or a severe operator error will induce data loss or corruption that would not have occurred in a standalone configuration, and
- Increased difficulty of recovering from unanticipated software and/or hardware states.

These costs and risks are fundamental, apply in one form or another to all clustered or cluster-capable products on the market (including the Oracle ZFS Storage Appliance product), and cannot be entirely eliminated or mitigated. Storage architects must weigh them against the primary benefit of clustering: the opportunity to reduce periods of unavailability from hours or days to minutes or less in the rare event of catastrophic hardware or software failure. Whether that cost/benefit analysis will favor the use of clustering in an Oracle ZFS Storage Appliance deployment will depend on local factors such as SLA terms, available support personnel and their qualifications, budget constraints, the perceived likelihood of various possible failures, and the appropriateness of alternative strategies for enhancing availability. These factors are highly site-, application-, and business-dependent and must be assessed on a case-by-case basis. Understanding the material in the rest of this section will help you make appropriate choices during the design and implementation of your unified storage infrastructure.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Cluster Interconnect I/O

All inter-controller communication consists of one or more messages transmitted over one of the three cluster I/O links provided by the CLUSTRON hardware (see [“Controller Cluster I/O Ports” in Oracle ZFS Storage Appliance Cabling Guide](#)). This device offers two low-speed serial links and one Ethernet link. The use of serial links allows for greater reliability; Ethernet links may not be serviced quickly enough by a system under extremely heavy load. False failure detection and unwanted takeover are the worst way for a clustered system to respond to load; during takeover, requests will not be serviced and will instead be enqueued by clients, leading to a flood of delayed requests after takeover in addition to already heavy load. The serial links used by the appliances are not susceptible to this failure mode. The Ethernet link provides a higher-performance transport for non-heartbeat messages such as rejoin synchronization and provides a backup heartbeat.

All three links are formed using ordinary straight-through EIA/TIA-568B (8-wire, Gigabit Ethernet) cables. To allow for the use of straight-through cables between two identical controllers, the cables must be used to connect opposing sockets on the two connectors as shown in [“Connecting Cluster Cables” in Oracle ZFS Storage Appliance Cabling Guide](#).

Clustered controllers only communicate with each other over the secure private network established by the cluster interconnects, and never over network interfaces intended for service or administration. Messages fall into two general categories: regular heartbeats used to detect the failure of a remote controller, and higher-level traffic associated with the resource manager and the cluster management subsystem. Heartbeats are sent, and expected, on all three links; they are transmitted continuously at fixed intervals and are never acknowledged or retransmitted as all heartbeats are identical and contain no unique information. Other traffic may be sent over any link, normally the fastest available at the time of transmission, and this traffic is acknowledged, verified, and retransmitted as required to maintain a reliable transport for higher-level software.

Regardless of its type or origin, every message is sent as a single 128-byte packet and contains a data payload of 1 to 68 bytes and a 20-byte verification hash to ensure data integrity. The serial links run at 115200 bps with 9 data bits and a single start and stop bit; the Ethernet link runs at 1Gbps. Therefore the effective message latency on the serial links is approximately 12.2 ms. Ethernet latency varies greatly; while typical latencies are on the order of microseconds, effective latencies to the appliance management software can be much higher due to system load.

Normally, heartbeat messages are sent by each controller on all three cluster I/O links at 50ms intervals. Failure to receive any message is considered link failure after 200ms (serial links) or 500ms (Ethernet links). If all three links have failed, the peer is assumed to have failed;

takeover arbitration will be performed. In the case of a panic, the panicking controller will transmit a single notification message over each of the serial links; its peer will immediately begin takeover regardless of the state of any other links. Given these characteristics, the clustering subsystem normally can detect that its peer has failed within:

- 550ms, if the peer has stopped responding or lost power, or
- 30ms, if the peer has encountered a fatal software error that triggered an operating system panic.

All of the values described in this section are fixed; the appliance does not offer the ability (nor is there any need) to tune these parameters. They are considered implementation details and are provided here for informational purposes only. They may be changed without notice at any time.

Note - To avoid data corruption after a physical re-location of a cluster, verify that all cluster cabling is installed correctly in the new location. For more information, see [“Preventing Split-Brain Conditions” on page 85](#).

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Cluster Resource Management

The resource manager is responsible for ensuring that the correct set of network interfaces is plumbed up, the correct storage pools are active, and the numerous configuration parameters remain in sync between two clustered controllers. Most of this subsystem's activities are invisible to administrators; however, one important aspect is exposed. Resources are classified into several types that govern when and whether the resource is imported (made active). Note that the definition of active varies by resource class; for example, a network interface belongs to the net class and is active when the interface is brought up.

The three most important resource types are singleton, private, and replica.

- **Replica resources** - Replicas are simplest: they are never exposed to administrators and do not appear on the cluster configuration screen. Replicas always exist and are always active on both controllers. Typically, these resources simply act as containers for service properties that must be synchronized between the two controllers.
- **Singleton resources** - Like replicas, singleton resources provide synchronization of state; however, singletons are always active on exactly one controller. Administrators can choose the controller on which each singleton should normally be active; if that controller has failed, its peer will import the singleton. Singletons are the key to clustering's availability characteristics; they are the resources one typically imagines moving from a

failed controller to its surviving peer and include network interfaces and storage pools. Because a network interface is a collection of IP addresses used by clients to find a known set of storage services, it is critical that each interface be assigned to the same controller as the storage pool clients will expect to see when accessing that interface's address(es). In Illustration 4, all of the addresses associated with the PrimaryA interface will always be provided by the controller that has imported pool-0, while the addresses associated with PrimaryB will always be provided by the same controller as pool-1.

- **Private resources** - Private resources are known only to the controller to which they are assigned, and are never taken over upon failure. This is typically useful only for network interfaces; see the following discussion of specific use cases.

FIGURE 3 ZS3-2 Clustering Example



Several other resource types exist; these are implementation details that are not exposed to administrators. One such type is the symbiote, which allows one resource to follow another as it is imported and exported. The most important use of this resource type is in representing the disks and flash devices in the storage pool. These resources are known as disksets and must always be imported before the ZFS pool they contain. Each diskset consists of half the disks in an external storage enclosure; a clustered storage system may have any number of disksets attached (depending on hardware support), and each ZFS pool is formed from the storage devices in one or more disksets. Because disksets may contain ATA devices, they must be explicitly imported and exported to avoid certain affiliation-related behaviors specific to ATA devices used in multipathed environments. Representing disks as resources provides a simple way to perform these activities at the right time. When an administrator sets or changes the ownership of a storage pool, the ownership assignment of the disksets associated with it is transparently changed at the same time. Like all symbiotes, diskset resources do not appear in the cluster configuration user interface.

TABLE 12 Cluster Resource Management

Resource	Icon	Omnipresent	Taken over on failure
SINGLETON		No	Yes
REPLICA	None	Yes	N/A
PRIVATE		No	No
SYMBIOTE	None	Same as parent type	Same as parent type

When a new resource is created, it is initially assigned to the controller on which it is being created. This ownership cannot be changed unless that controller is in the AKCS_OWNER state; it is therefore necessary either to create resources on the controller which should own them normally or to take over before changing resource ownership. It is generally possible to destroy resources from either controller, although destroying storage pools that are exported is not possible. Best results will usually be obtained by destroying resources on the controller which currently controls them, regardless of which controller is the assigned owner.

Most configuration settings, including service properties, users, roles, identity mapping rules, SMB autoshare rules, and iSCSI initiator definitions are replicated on both controllers automatically. Therefore it is never necessary to configure these settings on both controllers, regardless of the cluster state. If one appliance is down when the configuration change is made, it will be replicated to the other when it rejoins the cluster on next boot, prior to providing any service. There are a small number of exceptions:

- Share and LUN definitions and options may be set only on the controller which has control of the underlying pool, regardless of the controller to which that pool is ordinarily assigned.
- The "Identity" service's configuration (i.e., the appliance name and location) is not replicated.
- Names given to chassis are visible only on the controller on which they were assigned.
- Each network route is bound to a specific interface. If each controller is assigned an interface with an address in a particular subnet, and that subnet contains a router to which the appliances should direct traffic, a route must be created for each such interface, even if the same gateway address is used. This allows each route to become active individually as control of the underlying network resources shifts between the two controllers. For more information, see [“Clustering Considerations for Networking” on page 79](#).
- SSH host keys are not replicated and are never shared. Therefore if no private administrative interface has been configured, you may expect key mismatches when attempting to log into the CLI using an address assigned to a node that has failed. The same limitations apply to the SSL certificates used to access the BUI.

The basic model, then, is that common configuration is transparently replicated, and administrators will assign a collection of resources to each appliance controller. Those resource

assignments in turn form the binding of network addresses to storage resources that clients expect to see. Regardless of which appliance controls the collection of resources, clients are able to access the storage they require at the network locations they expect.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Cluster Takeover and Failback

Clustered controller nodes are in one of a small set of states at any given time:

TABLE 13 Cluster States

State	Icon	CLI/BUI Expression	Description
UNCONFIGURED		Clustering is not configured	A system that has no clustering at all is in this state. The system is either being set up or the cluster setup task has never been completed.
OWNER		Active (takeover completed)	Clustering is configured, and this node has taken control of all shared resources in the cluster. A system enters this state immediately after cluster setup is completed from its user interface, and when it detects that its peer has failed (i.e. after a takeover). It remains in this state until an administrator manually executes a failback operation.
STRIPPED		Ready (waiting for failback)	Clustering is configured, and this node does not control any shared resources. A system is STRIPPED immediately after cluster setup is completed from the user interface of the other node, or following a reboot, power disconnect, or other failure. A node remains in this state until an administrator manually

State	Icon	CLI/BUI Expression	Description
CLUSTERED		Active	executes a fail-back operation. Clustering is configured, and both nodes own shared resources according to their resource assignments. If each node owns a ZFS pool and is in the CLUSTERED state, then the two nodes form what is commonly called an active-active cluster.
-		Rejoining cluster ...	The appliance has recently rebooted, or the appliance management software is restarting after an internal failure. Resource state is being resynchronized.
-		Unknown (disconnected or restarting)	The peer appliance is powered off or rebooting, all its cluster interconnect links are down, or clustering has not yet been configured.

Transitions among these states take place as part of two operations: takeover and failback.

Takeover can occur at any time, and is attempted whenever peer failure is detected. It can also be triggered manually using the cluster configuration CLI or BUI, which can be useful for testing purposes. Finally, takeover will occur when a controller boots and detects that its peer is absent. This allows service to resume normally when one controller has failed permanently or when both controllers have temporarily lost power.

Failback never occurs automatically. When a failed controller is repaired and booted, it will rejoin the cluster (resynchronizing its view of all resources, their properties, and their ownership) and proceed to wait for an administrator to perform a failback operation. Until then, the original surviving controller will continue to provide all services. This allows for a full investigation of the problem that originally triggered the takeover, validation of a new software revision, or other administrative tasks prior to the controller returning to production service. Because failback is disruptive to clients, it should be scheduled according to business-specific needs and processes. There is one exception: Suppose that controller A has failed and controller B has taken over. When controller A rejoins the cluster, it becomes eligible to take over if it detects that controller B is absent or has failed. The principle is that it is always better to provide service than not, even if there has not yet been an opportunity to investigate the original problem. So while failback to a previously-failed controller will never occur automatically, it may still perform takeover at any time.

After you set up a cluster, the initial state consists of the node that initiated the setup in the OWNER state and the other node in the STRIPPED state. After performing an initial failback operation to hand the STRIPPED node its portion of the shared resources, both nodes are CLUSTERED. If both cluster nodes fail or are powered off, then upon simultaneous startup they will arbitrate and one of them will become the OWNER and the other STRIPPED.

During failback all foreign resources (those assigned to the peer) are exported, then imported by the peer. A pool that cannot be imported because it is faulted will trigger reboot of the STRIPPED node. An attempt to failback with a faulted pool can reboot the STRIPPED node as a result of the import failure.

To minimize service downtime, statistics and datasets are not available during failback and takeover operations. Data is not collected, and any attempt to suspend or resume statistics is delayed until failback and takeover operations have completed and data collection automatically resumes.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Configuration Changes in a Clustered Environment

The vast majority of appliance configuration is represented as either service properties or share/LUN properties. While share and LUN properties are stored with the user data on the storage pool itself (and thus are always accessible to the current owner of that storage resource), service configuration is stored within each controller. To ensure that both controllers provide coherent service, all service properties must be synchronized when a change occurs or a controller that was previously down rejoins with its peer. Since all services are represented by replica resources, this synchronization is performed automatically by the appliance software any time a property is changed on either controller.

It is therefore unnecessary and redundant for administrators to replicate configuration changes. Standard operating procedures should reflect this attribute and call for making changes to only one of the two controllers once initial cluster configuration has been completed. Note as well that the process of initial cluster configuration will replicate all existing configuration onto the newly-configured peer. Generally, then, we derive two best practices for clustered configuration changes:

- Make all storage- and network-related configuration changes on the controller that currently controls (or will control, if a new resource is being created) the underlying storage or network interface resources.

- Make all other changes on either controller, but not both. Site policy should specify which controller is to be considered the *master* for this purpose, and should in turn depend on which of the controllers is functioning and the number of storage pools that have been configured. Note that the appliance software does not make this distinction.

The problem of *amnesia*, in which disjoint configuration changes are made and subsequently lost on each controller while its peer is not functioning, is largely overstated. This is especially true of Oracle ZFS Storage Appliance, in which no mechanism exists for making independent changes to system configuration on each controller. This simplification largely alleviates the need for centralized configuration repositories and argues for a simpler approach: whichever controller is currently operating is assumed to have the correct configuration, and its peer will be synchronized to it when booting. While future product enhancements may allow for selection of an alternate policy for resolving configuration divergence, this basic approach offers simplicity and ease of understanding: the second controller will adopt a set of configuration parameters that are already in use by an existing production system (and are therefore highly likely to be correct). To ensure that this remains true, administrators should ensure that a failed controller rejoins the cluster as soon as it is repaired.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Clustering Considerations for Storage

When sizing an Oracle ZFS Storage Appliance for use in a cluster configuration, two considerations are very important:

- Whether all pools are owned by the same controller, or split between the two controllers.
- Whether you want pools with no single point of failure (NSPF).

Assigning storage pool ownership - Perhaps the most important decision is whether all storage pools will be assigned ownership to the same controller, or split between them. There are several trade-offs to consider, as shown in [Table 14, “Clustering Considerations for Storage Pools,” on page 78](#).

Generally, pools should be configured on a single controller except when optimizing for throughput during nominal operation or when failed-over performance is not a consideration. The exact changes in performance characteristics when in the failed-over state will depend to a great deal on the nature and size of the workload(s). Generally, the closer a controller is to providing maximum performance on any particular axis, the greater the performance degradation along that axis when the workload is taken over by that controller's peer. Of course, in the multiple pool case, this degradation will apply to both workloads.

Read cache devices are located in the controller or disk shelf, depending on your configuration.

Read cache devices, located in a controller slot (internal L2ARC), do not follow data pools in takeover or failback situations. A read cache device is only active in a particular cluster node when the pool that is assigned to the read cache device is imported on the node where the device resides. Absent additional configuration steps, read cache will not be available for a pool that has migrated due to a failover event. In order to enable a read cache device for a pool that is not owned by the cluster peer, take over the pool on the non-owning node, and then add storage and select the cache devices for configuration. Read cache devices in a cluster node should be configured as described in the [“Configuring Storage” on page 122](#). Write-optimized log devices are located in the storage fabric and are always accessible to whichever controller has imported the pool.

If read cache devices are located in a disk shelf (external L2ARC), read cache is always available. During a failback or takeover operation, read cache remains sharable between controllers. In this case, read performance is sustained. For external read cache configuration details, see [“Disk Shelf Configurations” in Oracle ZFS Storage Appliance Customer Service Manual](#).

Configuring NSPF - A second important consideration for storage is the use of pool configurations with no single point of failure (NSPF). Since the use of clustering implies that the application places a very high premium on availability, there is seldom a good reason to configure storage pools in a way that allows the failure of a single disk shelf to cause loss of availability. The downside to this approach is that NSPF configurations require a greater number of disk shelves than do configurations with a single point of failure; when the required capacity is very small, installation of enough disk shelves to provide for NSPF at the desired RAID level may not be economical.

The following table describes storage pool ownership for cluster configurations.

TABLE 14 Clustering Considerations for Storage Pools

Variable	Single controller pool ownership	Multiple pools owned by different controllers
Total throughput (nominal operation)	Up to 50% of total CPU resources, 50% of DRAM, and 50% of total network connectivity can be used to provide service at any one time. This is straightforward: only a single controller is ever servicing client requests, so the other is idle.	All CPU and DRAM resources can be used to provide service at any one time. Up to 50% of all network connectivity can be used at any one time (dark network devices are required on each controller to support failover).
Total throughput (failed over)	No change in throughput relative to nominal operation.	100% of the surviving controller's resources will be used to provide service. Total throughput relative to nominal operation may range from approximately 40% to 100%, depending on utilization during nominal operation.

Variable	Single controller pool ownership	Multiple pools owned by different controllers
I/O latency	Internal read cache is not available during a failback or takeover operation, which can significantly increase latencies for read-heavy workloads that fit into available read cache. Latency of write operations is unaffected. With external read cache configurations (EL2ARC), read performance is unaffected. Read cache is shared between cluster peers during a failback or takeover operation, resulting in no read latency.	Internal read cache is not available during a failback or takeover operation, which can significantly increase latencies for read-heavy workloads that fit into available read cache. Latency of both read and write operations may be increased due to greater contention for controller resources. This is caused by running two workloads on the surviving controller instead of the usual one. When nominal workloads on each controller approach the controller's maximum capabilities, latencies in the failed-over state may be extremely high. With external read cache configurations (EL2ARC), read performance is unaffected. Read cache is shared between cluster peers during a failback or takeover operation, resulting in no read latency.
Storage flexibility	All available physical storage can be used by shares and LUNs.	Only the storage allocated to a particular pool can be used by that pool's shares and LUNs. Storage is not shared across pools, so if one pool fills up while the other has free space, some storage may be wasted.
Network connectivity	All network devices in each controller can be used while that controller is providing service.	Only half of all network devices in each controller can be used while that controller is providing service. Therefore each pool can be connected to only half as many physically disjoint networks.

Related Topics


- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Clustering Considerations for Networking

Network device, datalink, and interface failures do not cause a clustered subsystem controller to fail. To protect against network failures inside or outside of the appliance, IPMP and/or LACP

should be used. A comprehensive approach to availability requires the correct configuration of the network and a network-wide plan for redundancy.

FIGURE 4 Clustering for Networking



Network interfaces can be configured as either singleton or private resources, provided they have a static IP configuration. Interfaces configured using DHCP must be private and using DHCP in clusters is discouraged. When configured as a singleton resource, all datalinks and devices used to construct an interface can be active on only one controller at a time. Likewise, corresponding devices on each controller must be attached to the same networks in order for service to be provided in a failed-over state. An example of this is shown in the previous diagram.

For a cluster to operate correctly when you construct network interfaces from devices and datalinks, it is essential that each singleton interface has a device using the same identifier and capabilities available on both controllers. Since device identifiers depend on the device type and the order in which they are first detected by the appliance, clustered controllers **MUST** have identical hardware installed. Each slot in both controllers must be populated with identical hardware and slots must be populated in the same order on both controllers. Your qualified Oracle reseller or service representative can assist in planning hardware upgrades that meet these requirements.

A route is always bound explicitly to a single network interface. Routes are represented within the resource manager as symbiotes and can become active only when the interfaces to which they are bound are operational. Therefore, a route bound to an interface which is currently in standby mode (exported) has no effect until the interface is activated during the takeover process. This is important when two pools are configured and are made available to a common subnet. If a subnet is home to a router that is used by the appliances to reach one or more other networks, a separate route (for example, a second default route), must be configured and bound to each of the active and standby interfaces attached to that subnet.

Example:

- Interface e1000g3 is assigned to 'alice' and e1000g4 is assigned to 'bob'.
- Each interface has an address in the 172.16.27.0/24 network and can be used to provide service to clients in the 172.16.64.0/22 network, reachable via 172.16.27.1.
- Two routes should be created to 172.16.64.0/22 via 172.16.27.1; one should be bound to e1000g3 and the other to e1000g4.

It is a good idea to assign each clustered controller an IP address used only for administration (most likely on a dedicated management network) and to designate the interface as a private resource. This ensures that it is possible to reach a functioning controller from the management network even if it is in a AKCS_STRIPPED state and awaiting failback. This is important if services such as LDAP and Active Directory are in use and require access to other network resources when the controller is not providing service. If this is not practical, the service processor should be attached to a reliable network and/or serial terminal concentrator so that the controller can be managed using the system console.

If neither of these actions is taken, it is impossible to manage or monitor a newly-booted controller until failback is completed. You may want to monitor or manage the controller that is providing service for a particular storage pool. This is likely to be useful when you want to modify some aspect of the storage itself such as modifying a share property or create a new LUN. This can be done by using one of the service interfaces to perform administrative tasks or by allocating a separate singleton interface to be used only for managing the pool to which it is matched. In either case, the interface should be assigned to the same controller as the pool it is used to manage.

Impact to NFSv4.1 clients - Certain networking changes in a cluster configuration can adversely affect the servicing of requests for NFSv4.1 clients. If the relationship between an IP address and its owner changes, then the best practice is to remount the filesystems from the client. Unlike NFSv4.0, NFSv4.1 protocol enables client connections over multiple IP addresses to be associated with the same NFSv4.1 protocol lease. When the relationship between an IP address and its owner changes, the group of IP addresses that failover together is no longer the same, forcing the client to re-establish the lease relationships by remounting the filesystems.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Private Local IP Interfaces

Use the following guidelines when creating private local IP interfaces:

- Creating an IP interface with the same name as a private IP interface on cluster peer, results in the local creation of a private IP interface.

- Datalinks in use by the peer's private interfaces can not be deleted and the delete button is greyed out.
- IP interfaces that belong to an IPMP group must all be of the same type and belong to the same controller. To create an IPMP group you must use either all singleton or all private IP interfaces and your cluster node must be the owner of these interfaces.
- The IPMP group type is set only at creation, and is determined by the type of underlying links.
- IP interfaces that belong to IPMP groups do not appear on the Cluster:Resources page because IP interface ownership cannot be modified independently of the IPMP group ownership.
- Private IPMP groups do not appear in the Cluster:Resources page because this type or ownership cannot be modified.


Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Clustering Considerations for InfiniBand

Like a network built on top of Ethernet devices, an InfiniBand network needs to be part of a redundant fabric topology in order to guard against network failures inside and outside of the appliance. The network topology should include IPMP to protect against network failures at the link level with a broader plan for redundancy for HCAs, switches and subnet managers.

FIGURE 5 Clustering Considerations for InfiniBand




To ensure proper cluster configuration, each controller must be populated with identical HCAs in identical slots. Furthermore, each corresponding HCA port must be configured into the same partition (pkey) on the subnet manager with identical membership privileges and attached to

the same network. To reduce complexity and ensure proper redundancy, it is recommended that each port belong to only one partition in the InfiniBand sub-network. Network interfaces may be configured as either singleton or private resources, provided they have static IP configuration. When configured as a singleton resource, all of the IB partition datalinks and devices used to construct an interface may be active on only one controller at any given time. A concrete example of this is shown in the illustration above. Changes to partition membership for corresponding ports must happen at the same time and in a manner consistent with the clustering rules described earlier. Your qualified Oracle reseller or service representative can assist in planning hardware upgrades that will meet these requirements.

The following illustration shows cluster configuration for subnet manager redundancy. Greater redundancy is achieved by connecting two dual-port HCAs to a redundant pair of server switches.

FIGURE 6 Cluster Configuration for Subnet Manager Redundancy



Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Preventing Split-Brain Conditions

A common failure mode in clustered systems is known as *split-brain*; in this condition, each of the clustered controllers believes its peer has failed and attempts takeover. Absent additional logic, this condition can cause a broad spectrum of unexpected and destructive behavior that can be difficult to diagnose or correct. The canonical trigger for this condition is the failure of the communication medium shared by the controllers; in the case of Oracle ZFS Storage Appliance, this would occur if the cluster I/O links fail. In addition to the built-in triple-link redundancy (only a single link is required to avoid triggering takeover), the appliance software will also perform an arbitration procedure to determine which controller should continue with takeover.

A number of arbitration mechanisms are employed by similar products; typically they entail the use of *quorum disks* (using SCSI reservations) or *quorum servers*. To support the use of ATA disks without the need for additional hardware, the appliance uses a different approach relying on the storage fabric itself to provide the required mutual exclusivity. The arbitration process consists of attempting to perform a SAS ZONE LOCK command on each of the visible SAS expanders in the storage fabric, in a predefined order. Whichever appliance is successful in its attempts to obtain all such locks will proceed with takeover; the other will reset itself. Since a clustered appliance that boots and detects that its peer is unreachable will attempt takeover and enter the same arbitration process, it will reset in a continuous loop until at least one cluster I/O link is restored. This ensures that the subsequent failure of the other controller will not result in an extended outage. These SAS zone locks are released when failback is performed or approximately 10 seconds has elapsed since the controller in the AKCS_OWNER state most recently renewed its own access to the storage fabric.

This arbitration mechanism is simple, inexpensive, and requires no additional hardware, but it relies on the clustered appliances both having access to at least one common SAS expander in the storage fabric. Under normal conditions, each appliance has access to all expanders, and arbitration will consist of taking at least two SAS zone locks. It is possible, however, to construct multiple-failure scenarios in which the appliances do not have access to any common expander. For example, if two of the SAS cables are removed or a disk shelf is powered down, each appliance will have access to disjoint subsets of expanders. In this case, each appliance will successfully lock all reachable expanders, conclude that its peer has failed, and attempt to proceed with takeover. This can cause unrecoverable hangs due to disk affiliation conflicts and/or severe data corruption.

Note that while the consequences of this condition are severe, it can arise only in the case of multiple failures (often only in the case of 4 or more failures). The clustering solution embedded in Oracle ZFS Storage Appliance is designed to ensure that there is no single point of failure, and to protect both data and availability against any plausible failure without adding undue cost or complexity to the system. It is still possible that massive multiple failures will cause loss of service and/or data, in much the same way that no RAID layout can protect against an unlimited number of disk failures.

FIGURE 7 Preventing Split-Brain

Fortunately, most such failure scenarios arise from human error and are completely preventable by installing the hardware properly and training staff in cluster setup and management best practices. Administrators should always ensure that all three cluster I/O links are connected and functional (see illustration), and that all storage cabling is connected as shown in the setup poster delivered with your appliances. It is particularly important that two paths are detected to each disk shelf (see illustration) before placing the cluster into production and at all times afterward, with the obvious exception of temporary cabling changes to support capacity increases or replacement of faulty components. Administrators should use alerts to monitor the state of cluster interconnect links and disk shelf paths and correct any failures promptly. Ensuring that proper connectivity is maintained will protect both availability and data integrity if a hardware or software component fails.

FIGURE 8 Cluster Two Paths


Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Estimating and Reducing Takeover Impact

There is an interval during takeover and failback during which access to storage cannot be provided to clients. The length of this interval varies by configuration, and the exact effects on clients depends on the protocol(s) they are using to access data. Understanding and mitigating these effects can make the difference between a successful cluster deployment and a costly failure at the worst possible time.

NFS (all versions) clients typically hide outages from application software, causing I/O operations to be delayed while a server is unavailable. NFSv2 and NFSv3 are stateless protocols that recover almost immediately upon service restoration. NFSv4.0 and NFSv4.1 incorporate a client grace period at startup, during which I/O typically cannot be performed. The duration of this grace period can be tuned in Oracle ZFS Storage Appliance; reducing it will reduce the apparent impact of takeover and/or failback. For planned outages, the appliance provides grace-less recovery for NFSv4.0 and NFSv4.1 clients, which avoids the grace period delay. For more information about grace-less recovery, see the Grace period property in [“NFS Service Properties” on page 328](#).

FIGURE 9 Cluster Grace Period

iSCSI behavior during service interruptions is initiator-dependent, but initiators will typically recover if service is restored within a client-specific timeout period. Check your initiator's documentation for additional details. The iSCSI target will typically be able to provide service as soon as takeover is complete, with no additional delays.

SMB, FTP, and HTTP/WebDAV are connection-oriented protocols. Because the session states associated with these services cannot be transferred along with the underlying storage and network connectivity, all clients using one of these protocols will be disconnected during a takeover or failback, and must reconnect after the operation completes.

While several factors affect takeover time (and its close relative, failback time), in most configurations these times will be dominated by the time required to import the diskset resource (s). Typical import times for each diskset range from 15 to 20 seconds, linear in the number of disksets. Recall that a diskset consists of one half of one disk shelf, provided the disk bays in that half-disk shelf have been populated and allocated to a storage pool. Unallocated disks and empty disk bays have no effect on takeover time. The time taken to import diskset resources is unaffected by any parameters that can be tuned or altered by administrators, so administrators planning clustered deployments should either:

- Limit installed storage so that clients can tolerate the related takeover times, or
- Adjust client-side timeout values above the maximum expected takeover time.

Note that while diskset import usually comprises the bulk of takeover time, it is not the only factor. During the pool import process, any intent log records must be replayed, and each share

and LUN must be shared via the appropriate service(s). The amount of time required to perform these activities for a single share or LUN is very small - on the order of tens of milliseconds - but with very large share counts this can contribute significantly to takeover times. Keeping the number of shares relatively small - a few thousand or fewer - can therefore reduce these times considerably.

Failback time is normally greater than takeover time for any given configuration. This is because failback is a two-step operation: first, the source appliance exports all resources of which it is not the assigned owner, then the target appliance performs the standard takeover procedure on its own assigned resources only. Therefore it will always take longer to failback from controller A to controller B than it will take for controller A to take over from controller B in case of failure. This additional failback time is much less dependent upon the number of disksets being exported than is the takeover time, so keeping the number of shares and LUNs small can have a greater impact on failback than on takeover. It is also important to keep in mind that failback is always initiated by an administrator, so the longer service interruption it causes can be scheduled for a time when it will cause the lowest level of business disruption.

Note - Estimated times cited in this section refer to software/firmware version 2009.04.10,1-0. Other versions may perform differently, and actual performance may vary. It is important to test takeover and its exact impact on client applications prior to deploying a clustered appliance in a production environment.

Related Topics

- [“Shutting Down a Clustered Configuration \(CLI\)” on page 64](#)

Network Configuration

The Networking Configuration features lets you create a variety of advanced networking setups using your physical network ports, including link-aggregations, virtual NICs (VNICs), virtual LANs (VLANs), and multipathing groups. You can then define any number of IPv4 and IPv6 addresses for these abstractions, for use in connecting to the various data services on the system.

There are four components to a system's network configuration:

- **Devices** - Physical network ports. These correspond to your physical network connections or IP on InfiniBand (IPoIB) partitions.
- **Datalinks** - The basic construct for sending and receiving packets. Datalinks may correspond 1:1 with a device (that is, with a physical network port) or IB Partition, or you may define Aggregation, VLAN and VNIC datalinks composed of other devices and datalinks.

- **Interface** - The basic construct for IP configuration and addressing. Each IP interface is associated with a single datalink, or is defined to be an IP MultiPathing (IPMP) group comprised of other interfaces.
- **Routing** - IP routing configuration. This controls how the system will direct IP packets.

To configure the network for the appliance, use the following sections:

- [“Network Configuration \(BUI\)” on page 90](#)
- [“Network Configuration \(CLI\)” on page 102](#)
- [“Working with Network Configuration” on page 111](#)
- [“Configuring Management Interfaces” on page 113](#)
- [“Configuring Network Datalinks” on page 113](#)
- [“Configuring Network Interfaces” on page 116](#)
- [“Configuring Network IP MultiPathing \(IPMP\)” on page 117](#)
- [“Configuring Network Performance and Availability” on page 118](#)
- [“Configuring Network Routing” on page 119](#)

Network Configuration (BUI)

When using the BUI to reconfigure networking, the system makes every effort to preserve the current networking connection to your browser. However, some network configuration changes such as deleting the specific address to which your browser is connected, will unavoidably cause the browser to lose its connection. For this reason it is recommended that you assign a particular IP address and network device for use by administrators and always leave the address configured. You can also perform particularly complex network reconfiguration tasks from the CLI over the serial console if necessary.

The following icons are used in the Configuration > Network section:

TABLE 15 Network Configuration Icons

Icon	Description
	Add new datalink, interface, or route
	Edit datalink, interface, or route settings
	Editing disabled
	Destroy datalink, interface, route
	Destruction disabled

Icon	Description
	Drag-and-drop icon
	Connected network port
	Connected network port with I/O activity
	Disconnected network port (link down, cable problem)
	Active InfiniBand port
	Active InfiniBand port with I/O activity
	Inactive InfiniBand port (down, init, or arm state)
	InfiniBand partition device is up
	InfiniBand partition device is down (subnet manager problem)
	Network datalink
	Network datalink VLAN or VNIC
	Network datalink aggregation
	Network datalink aggregation VLAN or VNIC
	Network datalink IB partition
	Interface is being used to send and receive packets (either up or degraded)
	Interface has been disabled by the user
	Interface is offline (owned by the cluster peer)
	Interface has failed or has been configured with a duplicate IP address

At top right is local navigation for Configuration, Addresses and Routing, which display alternate configuration views.

The Configuration page is shown by default, and lists Devices, Datalinks, and Interfaces, along with buttons for administration. Hover over an entry to expose an additional  icon, and click on any entry to highlight other components that are associated with it.

The Devices list shows links status on the right, as well as an icon to reflect the state of the network port. If ports appear disconnected, check that they are plugged into the network properly.

To configure an IP address on a network devices, first create a datalink, and then create an interface to use that datalink. The  icon may be used to do both, which will display dialogs for the Datalink and Interface properties.

There is more than one way to configure a network interface. Try clicking on the move icon  for a device, then dragging it to the datalink table. Then drag the datalink over to the interfaces table. Other moves are possible. This can be helpful for complex configurations, where valid moves are highlighted.

This page shows a summary table of the current network configuration, with fields:

TABLE 16 Summary of the Current Network Configuration

Field	Description	Example
Network Datalink	Datalink name and detail summary	datalink1 (via igb0)
Network Interface	Interface name and details summary	IPv4 DHCP, via datalink1
Network Addresses	Addresses hosted by this interface	192.168.2.80/22
Host Names	Resolved host names for the network addresses	caji.sf.example.com

This page provides configuration of the IP routing table and associated properties, as discussed above. By default, all entries in the routing table are shown, but the table can be filtered by type by using the subnavigation bar.

To check a specific route, in the CLI use `traceroute`.

```
zfssa-source:> traceroute 10.80.198.102
traceroute: Warning: Multiple interfaces found; using 10.80.198.101 @ igb3
traceroute to 10.80.198.102 (10.80.198.102), 30 hops max, 40 byte packets
 1 10.80.198.1 (10.80.198.1) 6.490 ms 0.924 ms 0.834 ms
 2 10.80.198.102 (10.80.198.102) 0.152 ms 0.118 ms 0.099 ms
zfssa-target:> traceroute 10.80.198.101
traceroute: Warning: Multiple interfaces found; using 10.80.198.102 @ igb3
traceroute to 10.80.198.101 (10.80.198.101), 30 hops max, 40 byte packets
 1 10.80.198.1 (10.80.198.1) 1.031 ms 0.905 ms 0.769 ms
 2 10.80.198.101 (10.80.198.101) 0.158 ms 0.111 ms 0.109 ms
```

▼ Configuring Management Interfaces (BUI)

Use the following procedure to configure management interfaces.

1. **Go to Configuration > Network > Configuration.**


2. Click the add icon  next to Datalinks.
 3. Set the following minimum datalink properties and click **APPLY**.
 - **VNIC** - Select this check box.
 - **Name** - Type a name for the datalink.
 4. Drag the resulting datalink to the **Interfaces** column.
 5. In the **Network Interface** dialog box, set the following minimum interface properties and click **APPLY**:
 - **Name** - Type a name for interface.
 - **Enable Interface** - Select this check box to enable the interface.
 - **Allow Administration** - Select this check box to make this a management interface, which enables BUI connections on port 215 and CLI connections on ssh port 22.
-
- Note** - The Allow Administration option makes it a management interface, enabling BUI connections on port 215 and CLI connections on ssh port 22.
-
- **Use IPv4 Protocol** or **Use IPv6 Protocol** - Select a protocol, its type of address, and enter one or more IP addresses in CIDR notation.
6. If clustered controllers, repeat steps 1-5 on the second controller.
 7. Click the trash icon  next to **Untitled Interface**, which is the default interface, to destroy it.

Note - When an interface is deleted, all routes associated with the interface are also removed.

8. In the **Update Default Route** dialog box, type the **Default Gateway** and select an **Interface** from the drop-down menu. Click **COMMIT WITH ROUTE**.

The default gateway is the default router IP address. For the interface, select the datalink that you assigned to the first management interface.

Note - It is strongly recommended to set a route because it enables communication with the appliance via the BUI and CLI. Without a route, the only means of communication with the appliance is through an Oracle ILOM connection to the SP.



Related Topics

- For an overview of network interface configuration, see [“Working with Network Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.x.](#)
- For further configuration, see [“Configuring the Appliance” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.x.](#)
- To upgrade the software on standalone controller, see [“Upgrading the Software” in Oracle ZFS Storage Appliance Customer Service Manual.](#)
- To lock the cluster management interfaces, see [“Locking Cluster Management Interfaces \(BUI\)” on page 94.](#)

▼ Locking Cluster Management Interfaces (BUI)

After initial configuration, clustered controllers are in an active-active state. When a failover occurs, an active controller takes over all non-private interfaces, and the peer controller becomes passive and inaccessible by its BUI and CLI. To maintain access to a controller regardless of its state, lock its management interface to make it private. The following procedure locks the management interface on each clustered controller.



Caution - Failure to configure locked management interfaces on clustered controllers may lead to longer than necessary fault diagnosis and resolution times.

1. In the BUI of the first controller, navigate to **Configuration > Cluster**.
2. In the BUI of the second controller, navigate to **Configuration > Cluster**.
3. From the BUI of the first controller, choose the management interface for the first controller from the **Resource list**.
4. Click the padlock icon  to lock the management interface to this controller.

The interface displays a locked icon  next to its name in the **Resource list**.

5. From the BUI of the second controller, choose the management interface for the second controller from the Resource list.
6. Click the padlock icon  to lock the management interface to this controller.

The interface displays a locked icon  next to its name in the Resource list.

Related Topics

- To upgrade the software, see [“Upgrading the Software” in Oracle ZFS Storage Appliance Customer Service Manual](#).

▼ Creating a Single Port Interface (BUI)

1. Go to Configuration > Network > Configuration.
2. Click the Datalinks add item icon .
3. Optionally, type a name and select the Custom MTU button then type 9000 in the text box.
4. Choose a device from the Devices list.
5. Click **APPLY**.
The datalink will appear in the Datalinks list.
6. Click the Interface add item icon .
7. Set desired properties, and choose the datalink previously created.
8. Click **APPLY**.
The interface will appear in the Interfaces list.
9. The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click **APPLY** at the top to commit the configuration.

▼ Modifying an Interface (BUI)

1. Go to Configuration > Network > Configuration.

2. Click the edit icon  for a network datalink or network interface.
3. Change settings to desired values.
4. Click APPLY.
5. Click APPLY at the top of the page to commit the configuration.

▼ Unlocking Data Interfaces (BUI)

1. Go to Configuration > Cluster.
2. Verify that the lock icons for data interfaces are gray, indicating that it is unlocked .
3. If the lock icons are black , click each icon to unlock the interfaces.
4. Click APPLY to save your changes.

▼ Creating an LACP Aggregated Link Interface (BUI)

1. Go to Configuration > Network > Configuration.
2. Click the Datalinks add item icon .
3. Optionally set the datalink name.
4. Select LACP Aggregation.
5. Select two or more devices from the Devices list, and click APPLY.
6. Click the Interfaces add item icon .
7. Set desired properties, choose the aggregated link from the Datalinks list, and click APPLY.
8. Click APPLY at the top to commit the configuration.

▼ Creating an IPMP Group Using Probe-Based and Link-State Failure Detection (BUI)

Create one or more "underlying" IP interfaces that will be used as components of the IPMP group. Each interface must have an IP address to be used as the probe source (see [“Creating a Single Port Interface \(BUI\)” on page 95](#)).

Do not use probe-based failure detection when there no systems (other than the cluster peer) on the same subnet as the IPMP test addresses that are configured to answer ICMP echo requests.

1. **Go to Configuration > Network > Configuration.**
2. **Click the Interface add item icon .**
3. **Optionally, change the name of the interface.**
4. **Click the IP MultiPathing Group check box.**
5. **Click the Use IPv4 Protocol and/or the Use IPv6 Protocol, and specify the IP addresses for the IPMP interface.**
6. **Choose the interfaces created in the first step from the Interfaces list.**
7. **Set each chosen interface to be either Active or Standby, as desired.**
8. **Click APPLY.**

▼ Creating an IPMP Group Using Link-State Only Failure Detection (BUI)

Create one or more "underlying" IP interfaces with the IP address 0.0.0.0/8 to be used as the components of the IPMP group (see [“Creating a Single Port Interface \(BUI\)” on page 95](#)).

1. **Go to Configuration > Network > Configuration.**
2. **Click the Interface add item icon .**
3. **Optionally, change the name of the interface.**
4. **Click the IP MultiPathing Group check box.**

5. Click the Use IPv4 Protocol or/and the Use IPv6 Protocol and specify the IP addresses for the IPMP interface.
6. Choose the interfaces created in the first step from the Interfaces list.
7. Set each chosen interface to be either Active or Standby, as desired.
8. Click APPLY.

▼ Extending an LACP Aggregation (BUI)

1. Go to Configuration > Network > Configuration.
2. Hover over a device in the Devices list.
3. Click the move icon , then drag and drop the device onto an aggregation datalink.
4. Click APPLY at the top of the page to commit this configuration.

▼ Extending an IPMP Group (BUI)

1. Go to Configuration > Network > Configuration.
2. Hover over an interface in the Interfaces list.
3. Click the move icon , then drag and drop the device onto an IPMP interface.
4. Click APPLY at the top of the page to commit this configuration.

▼ Creating an InfiniBand Partition Datalink and Interface (BUI)

1. Go to Configuration > Network > Configuration.
2. Click the Datalinks add item icon .
3. Optionally, set a name.
4. Click the IB Partition checkbox.
5. Choose a device from the Partition Devices list.

6. **Enter a four-digit hexadecimal number for the partition key, which must match what was configured on the InfiniBand subnet manager.**
7. **Choose link mode from the drop down menu.**
8. **Click APPLY. The new partition datalink will appear in the Datalinks list.**
9. **Click the Interface add item icon .**
10. **Set desired properties, and choose the datalink previously created.**
11. **Click APPLY.**
The interface will appear in the Interfaces list.
12. **The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click APPLY at the top to commit the configuration.**

▼ **Creating a VNIC Without a VLAN ID for Clustered Controllers (BUI)**

This example is for an active-active configuration with half of the network ports on standby. This task creates an IP interface over a device datalink and assigns it to a head. A VNIC is built on top of the same datalink, and an IP interface is configured on top of the VNIC and assigned to the other head. Configuring one instead of multiple VNICs over a given datalink ensures peak performance. Traffic flows over the cable associated with the underlying active port on one head, as well as the underlying standby port on the other head. Thus, the otherwise idle standby port can be used with VNICs.

1. **Go to Configuration > Network > Configuration.**
2. **When the cluster is in state AKCS_CLUSTERED, click the Datalinks add item icon .**
3. **Optionally, set a name and MTU value.**
4. **Choose a device from the Devices list and click APPLY.**
The datalink appears in the Datalinks list.
5. **Click the Interface add item .**

6. **Set desired properties, choose the datalink previously created, and click APPLY.**
The interface appears in the Interfaces list.
7. **Click the Datalinks add item icon .**
8. **Select the VNIC checkbox, optionally set name and MTU (equal to or less than the value in step 2), and click APPLY.**
The new VNIC datalink appears in the Datalinks list.
9. **Click the Interface add item icon .**
10. **Set desired properties, choose the VNIC datalink previously created, and click APPLY.**
The interface appears in the Interfaces list.
11. **The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click APPLY at the top to commit the configuration.**
12. **Click the Cluster tab.**
The two newly created interfaces appear in the Resource section with default owners.
13. **Use the Owner pull-down list to assign one of the two interfaces to the other head and click APPLY.**

▼ **Creating VNICs with the Same VLAN ID for Clustered Controllers (BUI)**

This example is for an active-active configuration with half of the network ports on standby. This task creates two VNICs with identical VLAN IDs on top of the same device datalink. Each VNIC is configured with an interface, and each interface is assigned to a different head. Traffic flows over the cable associated with the underlying active port on one head, as well as the underlying standby port on the other head. Thus, the otherwise idle standby port can be used with VNICs.

1. **Go to Configuration > Network > Configuration.**
2. **When the cluster is in state AKCS_CLUSTERED, click the Datalinks add item icon .**

3. **Select the VNIC checkbox, optionally set a name and MTU, set the VLAN ID, choose a device from the Devices list, and click APPLY.**
The new VNIC datalink appears in the Datalinks list.
4. **Click the Interface add item icon .**
5. **Set desired properties, choose the VNIC datalink previously created, and click APPLY.**
The interface appears in the Interfaces list.
6. **Create another VNIC as described in steps 2 and 3 with the same Device and VLAN ID, and create an interface for it as described in steps 4 and 5.**
7. **The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click APPLY at the top to commit the configuration.**
8. **Click the Cluster tab.**
The two newly created interfaces appear in the Resource section with default owners.
9. **Use the Owner pull-down list to assign one of the two interfaces to the other head and click APPLY.**

▼ Adding a Static Route (BUI)

1. **Go to Configuration > Network > Routing.**
2. **Click the add item icon  next to Routing Table Entries.**
3. **Select a family protocol and type. Specify a destination address, gateway address, and an interface.**
4. **Click ADD.**
The new route will appear in the table.

▼ Deleting a Static Route (BUI)

1. **Go to Configuration > Network > Routing.**
2. **Hover over the route entry, then click the trash icon  on the right.**

Network Configuration (CLI)

Network configuration is under the configuration `net`, which has sub commands for devices, datalinks, interfaces, and routing. The `show` command can be used with each to show the current configuration:

```
caji:> configuration net
```

```
caji:configuration net> devices show
```

Devices:

DEVICE	UP	SPEED	MAC
igb0	true	1000 Mbit/s	0:14:4f:9a:b9:0
igb1	true	1000 Mbit/s	0:14:4f:9a:b9:1
igb2	true	1000 Mbit/s	0:14:4f:9a:b8:fe
igb3	true	1000 Mbit/s	0:14:4f:9a:b8:ff

```
caji:configuration net> datalinks show
```

Datalinks:

DATALINK	CLASS	LINKS	LABEL
igb0	device	igb0	datalink1

```
caji:configuration net> interfaces show
```

Interfaces:

INTERFACE	STATE	CLASS	LINKS	ADDRS	LABEL
igb0	up	ip	igb0	192.168.2.80/22	caji

```
caji:configuration net> routing show
```

Properties:

multihoming = loose

Routes:

ROUTE	DESTINATION	GATEWAY	INTERFACE	TYPE
route-000	0.0.0.0/0	192.168.1.1	igb0	dhcp
route-001	192.168.0.0/22	192.168.2.142	igb0	system

Type `help` in each section to see the relevant commands for creating and configuring datalinks, interfaces, and routes. Subcommands that are valid in this context:

`help [topic]` => Get context-sensitive help. If `[topic]` is specified, it must be one of "builtins", "commands", "general", "help", "script" or "properties".

`show` => Show information pertinent to the current context

<code>commit</code>	=> Commit current state, including any changes
<code>abort</code>	=> Abort creation of "vnic"
<code>done</code>	=> Finish operating on "vnic"
<code>get [prop]</code>	=> Get value for property [prop]. ("help properties" for valid properties.) If [prop] is not specified, returns values for all properties.
<code>set [prop]</code>	=> Set property [prop] to [value]. ("help properties" for valid properties.) For properties taking list values, [value] should be a comma-separated list of values.
<code>available</code>	=> Get values that can be assigned to the links parameter when creating a network component.

The `available` command is used to see what values can be assigned to the `links` parameter when creating a network component. The following shows the output from the CLI command `available`:

```
caji:configuration net datalinks> device
caji:configuration net datalinks device (uncommitted)> available
igb7,igb6

caji:configuration net datalinks> vnic
caji:configuration net datalinks vnic (uncommitted)> available
igb5,igb4,aggr2,aggr1

caji:configuration net datalinks> vlan
caji:configuration net datalinks vlan (uncommitted)> available
igb5,igb4,aggr2,aggr1

caji:configuration net datalinks> aggregation
caji:configuration net datalinks aggregation (uncommitted)> available
igb7,igb6

caji:configuration net interfaces> ip
caji:configuration net interfaces ip (uncommitted)> available
aggr2,aggr1

caji:configuration net interfaces> ipmp
caji:configuration net interfaces ipmp (uncommitted)> available
vnic4,vnic3,igb5,igb4
```

The following demonstrates creating a datalink using the `device` command, and interface using the `ip` command:

```

caji:configuration net> datalinks
caji:configuration net datalinks> device
caji:configuration net datalinks device (uncommitted)> set links=igb1
    links = igb1 (uncommitted)
caji:configuration net datalinks device (uncommitted)> set label=datalink2
    label = datalink2 (uncommitted)
caji:configuration net datalinks device (uncommitted)> set mtu=9000
    mtu = 9000 (uncommitted)
caji:configuration net datalinks device (uncommitted)> commit
caji:configuration net datalinks> show
Datalinks:

```

DATALINK	CLASS	LINKS	LABEL
igb0	device	igb0	datalink1
igb1	device	igb1	datalink2

```

caji:configuration net datalinks> cd ..
caji:configuration net> interfaces
caji:configuration net interfaces> ip
caji:configuration net interfaces ip (uncommitted)> set label="caji2"
    label = caji2 (uncommitted)
caji:configuration net interfaces ip (uncommitted)> set links=igb1
    links = igb1 (uncommitted)
caji:configuration net interfaces ip (uncommitted)> set v4addrs=10.0.1.1/8
    v4addrs = 10.0.1.1/8 (uncommitted)
caji:configuration net interfaces ip (uncommitted)> commit
caji:configuration net interfaces> show
Interfaces:

```

INTERFACE	STATE	CLASS	LINKS	ADDRS	LABEL
igb0	up	ip	igb0	192.168.2.80/22	caji
igb1	up	ip	igb1	10.0.1.1/8	caji2

The following demonstrates creating a default route via 10.0.1.2 over the new igb1 IP interface:

```

caji:configuration net routing> create
caji:configuration net route (uncommitted)> set family=IPv4
    family = IPv4 (uncommitted)
caji:configuration net route (uncommitted)> set destination=0.0.0.0
    destination = 0.0.0.0 (uncommitted)
caji:configuration net route (uncommitted)> set mask=0
    mask = 0 (uncommitted)
caji:configuration net route (uncommitted)> set interface=igb1
    interface = igb1 (uncommitted)
caji:configuration net route (uncommitted)> set gateway=10.0.1.2
    gateway = 10.0.1.2 (uncommitted)
caji:configuration net route (uncommitted)> commit

```

▼ Configuring Management Interfaces (CLI)

Use the following procedure to configure management interfaces.

1. **Go to configuration net, and then enter datalinks.**

```
hostname:> configuration net
hostname:configuration net> datalinks
```

2. **Enter show to view the datalink(s).**

```
hostname:configuration net datalinks> show
Datalinks:
```

DATALINK	CLASS	LINKS	STATE	ID	LABEL
igb0	device	igb0	up	-	Untitled Datalink

3. **Create a VNIC for the management datalink by going to the vnic context and setting its label to indicate that the link is for management, and optionally assign a VLAN ID. Enter commit and then enter cd .. to return to the correct context for the next step.**

```
hostname:configuration net datalinks> vnic
hostname:configuration net datalinks vnic (uncommitted)> set links=igb0
links = igb0 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> set label=management-datalink-1
label = management-datalink-1 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> commit
hostname:configuration net datalinks vnic> cd ..
```

To assign a VLAN ID:

```
hostname:configuration net datalinks vnic> set id=100
id = 100 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> commit
hostname:configuration net datalinks vnic> cd ..
```

4. **If clustered controllers, create a VNIC for the second management datalink by going to the vnic context and setting a unique label, and optionally assign a VLAN ID. Enter commit and then enter cd .. to return to the correct context for the next step.**

```
hostname:configuration net datalinks> vnic
hostname:configuration net datalinks vnic (uncommitted)> set links=igb0
links = igb0 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> set label=management-datalink-2
```

```

label = management-datalink-2 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> commit
hostname:configuration net datalinks vnic> cd ..

```

To assign a VLAN ID:

```

hostname:configuration net datalinks vnic (uncommitted)> set id=100
id = 100 (uncommitted)
hostname:configuration net datalinks vnic> commit
hostname:configuration net datalinks vnic> cd ..

```

5. Create an IP interface for the controller and assign it to the VNIC.

```

hostname:configuration net> interfaces
hostname:configuration net interfaces> ip
hostname:configuration net interfaces ip (uncommitted)> set v4addrs=192.168.1.101/24
v4addrs = 192.168.1.101/24 (uncommitted)
hostname:configuration net interfaces ip (uncommitted)> set label=management-
controller-1
label = management-controller-1 (uncommitted)
hostname:configuration net interfaces ip (uncommitted)> set links=vnic1
links = vnic1 (uncommitted)
hostname:configuration net interfaces ip (uncommitted)> commit

```

6. If clustered controllers, create an IP interface and assign it to the VNIC for the second controller. After entering commit, enter done.

```

hostname:configuration net interfaces> ip
hostname:configuration net interfaces ip (uncommitted)> set v4addrs=192.168.1.102/24
v4addrs = 192.168.1.102/24 (uncommitted)
hostname:configuration net interfaces ip (uncommitted)> set label=management-
controller-2
label = management-controller-2 (uncommitted)
hostname:configuration net interfaces ip (uncommitted)> set links=vnic2
links = vnic2 (uncommitted)
hostname:configuration net datalinks vnic (uncommitted)> commit
hostname:configuration net datalinks> done

```

7. Configure routing for the first controller. If clustered controllers, configure routing for the second controller. After entering commit, enter done.

First controller:

```

hostname:configuration net> routing
hostname:configuration net routing> create
hostname:configuration net route (uncommitted)> set destination=0.0.0.0
destination = 0.0.0.0 (uncommitted)
hostname:configuration net route (uncommitted)> set mask=0
mask = 0 (uncommitted)

```

```

hostname:configuration net route (uncommitted)> set interface=vnic1
                               interface = vnic1 (uncommitted)
hostname:configuration net route (uncommitted)> set gateway=192.168.1.1
                               gateway = 192.168.1.1 (uncommitted)
hostname:configuration net route (uncommitted)> set family=IPv4
                               family = IPv4
hostname:configuration net route (uncommitted)> commit
hostname:configuration net route> done

```

Second controller:

```

hostname:configuration net> routing
hostname:configuration net routing> create
hostname:configuration net route (uncommitted)> set destination=0.0.0.0
                               destination = 0.0.0.0 (uncommitted)
hostname:configuration net route (uncommitted)> set mask=0
                               mask = 0 (uncommitted)
hostname:configuration net route (uncommitted)> set interface=vnic2
                               interface = vnic2 (uncommitted)
hostname:configuration net route (uncommitted)> set gateway=192.168.1.1
                               gateway = 192.168.1.1 (uncommitted)
hostname:configuration net route (uncommitted)> set family=IPv4
                               family = IPv4
hostname:configuration net route (uncommitted)> commit
hostname:configuration net route> done

```

8. **Destroy the default interface, which is named Untitled Interface, enter `cd ..` and then enter `done`.**

Note - When an interface is deleted, all routes associated with the interface are also removed.

```

hostname:configuration net> interfaces
hostname:configuration net interfaces> show
Interfaces:

INTERFACE  STATE   CLASS  LINKS  ADDRS                LABEL
igb0       up      ip      igb0   192.168.1.101/24     Untitled Interface
vnic1      duplicate ip    vnic1  192.168.1.101/24     management-controller-1
vnic2      duplicate ip    vnic2  192.168.1.102/24     management-controller-2
hostname:configuration net interfaces> destroy igb0
This will destroy "igb0" and any networking objects exclusively built over it.

Are you sure? (Y/N) y

hostname:configuration net interfaces> cd ..
hostname:configuration net> done

```

Related Topics

- For further configuration, see “Configuring the Appliance” in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.x*.
- To upgrade the software on a standalone controller, see “Upgrading the Software” in *Oracle ZFS Storage Appliance Customer Service Manual*.
- To lock the management interfaces, see “Locking Cluster Management Interfaces (CLI)” on page 109.

▼ Configuring Network Interfaces (CLI)

Before You Begin To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure static /32 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done. The two source cluster nodes can replicate to the same target node only if the target node provides two IP addresses, one for use by each node in the source cluster. Replicating to the same target IP address from both nodes of a source cluster is not supported.

1. Navigate to configuration services routing on the source appliance.

Use a static /32 (host-specific) route to the target system IP address via the dedicated network interface. In the following example, mask=32 means this is a host-specific route.

```
host_source:configuration services routing> create

host_source:configuration services route (uncommitted)> get
    family = (unset)
    destination = (unset)
    mask = (unset)
    gateway = (unset)
    interface = (unset)
host_source:configuration services route (uncommitted)> set family=IPv4
host_source:configuration services route (uncommitted)> set destination=203.34.56.78
host_source:configuration services route (uncommitted)> set mask=32
host_source:configuration services route (uncommitted)> set gateway=203.34.56.254
host_source:configuration services route (uncommitted)> set interface=nge3
host_source:configuration services route (uncommitted)> commit
host_source:configuration services routing> show
route-000  0.0.0.0/0                203.24.30.254    nge0            static
route-001  203.24.30.0/32                  203.24.30.28     nge0            dynamic
route-002  203.24.150.0/32                 203.24.150.10    ibd0            dynamic
route-003  203.24.101.65/32                203.24.30.254    nge1            inactive
route-005  203.34.56.78/32                 203.34.56.254    nge3            static
```


2. **After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.**
3. **To ensure traffic is routed through the correct source and target interfaces, use the `traceroute` command.**

For information about using `traceroute`, see [“Configuring Network Routing” on page 119](#).

Note - When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Locking Cluster Management Interfaces (CLI)

After initial configuration, clustered controllers are in an active-active state. When a failover occurs, an active controller takes over all non-private interfaces, and the peer controller becomes passive and inaccessible by its BUI and CLI. To maintain access to a controller regardless of its state, lock its management interface to make it private. The following procedure locks the management interface on each clustered controller.



Caution - Failure to configure locked management interfaces on clustered controllers may lead to longer than necessary fault diagnosis and resolution times.

1. **On the first controller, go to configuration cluster resources and select the management interface for the first controller, prefacing it with `net/`.**


```
controller-a:> configuration cluster resources select net/igb0
```
2. **Lock the interface by setting the type to private:**


```
configuration cluster resources (uncommitted)> set type=private
configuration cluster resources (uncommitted)> commit
```
3. **On the second controller, go to configuration cluster resources and select the management interface for the second controller, prefacing it with `net/`.**


```
controller-b:> configuration cluster resources select net/igb1
```
4. **Lock the interface by setting the type to private:**

```
configuration cluster resources (uncommitted)> set type=private
configuration cluster resources (uncommitted)> commit
```

Related Topics

- To upgrade the software, see [“Upgrading the Software” in Oracle ZFS Storage Appliance Customer Service Manual](#).

▼ Adding a Static Route (CLI)

1. Go to configuration net routing.
2. Enter create.
3. Type show to list required properties, and set each.
4. Enter commit.

▼ Deleting a Static Route (CLI)

1. Go to configuration net routing.
2. Type show to list routes, and route names (e.g., route-002).
3. Enter destroy route name.

▼ Unlocking Data Interfaces (CLI)

1. Go to configuration cluster resources.

```
hostname:> configuration cluster resources
```

2. Enter show.

All data interfaces should have a TYPE set to singleton.

```
hostname:configuration cluster resources> show
Resources:
```

RESOURCE	OWNER	TYPE	LABEL	CHANGES	DETAILS

net/vnic2	zs34-02	private	mgmt-02	no	10.80.218.170
net/vnic3	zs34-02	private	data-01	no	10.80.216.46
net/vnic4	zs34-02	singleton	data-02	no	10.80.216.47
zfs/pool-01	zs34-01	singleton		no	
zfs/pool-02	zs34-02	singleton		no	53.5T

3. If a data interface has a type set to private, select the resource and set the type to singleton.

```
hostname:configuration cluster resources> select net/vnic3
hostname:configuration cluster resources net/vnic3> set type=singleton
type = singleton
```

4. Enter commit.


```
hostname:configuration cluster resources> commit
```

▼ Changing the Multihoming Property to Strict (CLI)

1. Go to configuration net routing.
2. Enter set multihoming=strict.
3. Enter commit.

Working with Network Configuration

In the appliance model, network devices are created by the system to represent the available network or InfiniBand ports; they have no configurable settings. Datalinks are a layer 2 entity and must be created to apply settings such as LACP to these network devices. Interfaces are a layer 3 entity containing the IP settings, which they make available via a datalink. This model has separated network interface settings into two parts: datalinks for layer 2 settings and interfaces for layer 3 settings.



The Devices column corresponds to the physical network interface card (NIC) ports on the controller and are typically labeled igb0, igb1, igb2, and igb3. Port NET-0 corresponds to device igb0, port NET-1 to igb1, and so on. It is strongly recommended to use one NIC port per controller as a *management interface*. This column also contains the physical InfiniBand ports on the controller, and are typically labeled ibp0, ibp1, ibp2, and ibp3.

The Datalinks column corresponds to the construct for sending and receiving packets for a specific network device. They support VLANs, VNICS, IB partitions, and LACP aggregation. Datalinks are required to complete network configuration, even if they do not apply specific settings to network devices.

The datalink entity (which we named "aggr1") groups the network devices in a configurable way (LACP aggregation policy). The interface entity (which we named "phobos") provides configurable IP address settings, which it makes available on the network via the datalink. The network devices (named "igb1", "igb2", ..., by the system) have no direct settings.

The Interfaces column corresponds to configurable IP address settings and other properties for datalinks. Interfaces can be available via a single datalink or as defined in an IP MultiPathing (IPMP) group comprising a pool of datalinks, which allows automatic migration of IP addresses from failed to working datalinks.

An example of a single IP address on a single port (common configuration) is:

- Devices - igb0
- Datalink - datalink1

- Interface - deimos (192.168.2.80/22)

The following configuration is for a 3-way link aggregation:

- Devices - igb1, igb2, igb3
- Datalink - aggr1 (LACP aggregation)
- Interface - phobos (192.168.2.81/22)

Configuring Management Interfaces

If you did not set a management interface during initial configuration, use the following procedures to configure a network interface card (NIC) port as a management interface. A management interface is a network interface with administrative access.

The physical ports correspond to the interfaces, from which you can select to make a management interface. All standalone controllers should have at least one NIC port configured as a management interface. All cluster installations should have at least one NIC port on each controller configured as a management interface. In addition, the NIC instance number must be unique on each clustered controller.

For clustered controllers, it is recommended that you lock the management interfaces.

To configure management interfaces and lock cluster management interfaces, use the following procedures:

- [“Configuring Management Interfaces \(BUI\)” on page 92](#)
- [“Configuring Management Interfaces \(CLI\)” on page 105](#)
- [“Locking Cluster Management Interfaces \(BUI\)” on page 94](#)
- [“Locking Cluster Management Interfaces \(CLI\)” on page 109](#)

Configuring Network Datalinks

Network datalinks manage devices, and are used by interfaces. They support:

- **Link Aggregation Control Protocol (LACP)** - LACP is used to bundle multiple network devices such that they behave as one. This improves performance (by increasing bandwidth) and reliability (by protecting from network port failure); however, the appliance must be connected to a switch that supports LACP and has it enabled for those ports.
- **InfiniBand (IB) Partitions** - InfiniBand partitions connect to logically isolated IB fabric domains.
- **Virtual LANs (VLANs)** - VLANs are used to improve local network security and isolation. VLANs are recommended for administering the appliance; otherwise, use VNICs.

- **Virtual Network Interface Cards (VNICs)** - VNICs allow single or aggregated Ethernet datalinks to be split into multiple virtual (Ethernet) datalinks. VNICs can be optionally tagged with VLAN IDs, and can allow physical network port sharing in a cluster. Step-by-step instructions can be found in [“Clustering Considerations for Networking” on page 79](#) below.

Note - VNIC-based and VLAN-based datalinks cannot share the same VLAN ID.

The IEEE802.3ad (link aggregation) standard does not explicitly support aggregations across multiple switches, but some vendors provide multi-switch support via proprietary extensions. If a switch configured with those extensions conforms to the IEEE standard and the extensions are transparent to the end-nodes, its use is supported with the appliance. If an issue is encountered, Oracle support may require it to be reproduced on a single-switch configuration.

The following datalink settings are available:

TABLE 17 Datalink Settings

Property	Description
Name	Use the defined custom name. For example: "internal", "external", "adminnet", etc.
Speed	Use the defined speed. Valid values are auto, 10, 100, 1000 and 10000, representing autonegotiation, forced 10Mbit/sec, forced 100Mbit/sec, forced 1Gbit/sec and forced 10Gbit/sec. Speed and duplex must be either both forced to specific values or both set to autonegotiate. Not all networking devices support forcing to all possible speed/duplex combinations. Disabling autonegotiation is strongly discouraged. However, if the switch has autonegotiation disabled, it may be necessary to force speed (and duplex) to ensure the datalink runs at the expected speed and duplex.
Duplex	Use the defined transmission direction. Valid CLI values are auto, half, and full, representing autonegotiation, half- and full-duplex respectively. Speed and duplex must be either both forced to specific values or both set to autonegotiate.
VLAN	Use VLAN headers.
VLAN ID	Use the defined VLAN identifier; optional for VNICs.
VNIC	Use a VNIC.
MTU	Use the defined maximum transmission unit (MTU) size. The default MTU is 1500 bytes. Specify a lower MTU (minimum 1280) to leave packet headroom (for example, for tunneling protocols). Specify a larger MTU (maximum 9000) to improve network performance. All systems and switches on the same LAN must be configured with the chosen MTU. After the MTU value

Property	Description
	is set and the new network configuration is committed to the system, you can return to the network screen and view the datalink status to see the exact MTU value in bytes that was selected. Note that a VLAN or VNIC cannot be configured with an MTU value larger than that of the underlying datalink.
LACP Aggregation	Use multiple network device LACP aggregation.
LACP Policy	Use the defined LACP policy for selecting an outbound port. L2 hashes the source and destination MAC address; L3 uses the source and destination IP address; L4 uses the source and destination transport level port
LACP Mode	Use the defined LACP communication mode. Active mode will send and receive LACP messages to negotiate connections and monitor the link status. Passive mode will listen for LACP messages only. Off mode will use the aggregated link but not detect link failure or switch configuration changes. Some network switch configurations, including Cisco Etherchannel, do not use the LACP protocol; the LACP mode should be set to "off" when using non-LACP aggregation in your network.
LACP Timer	Use the defined interval between LACP messages for Active mode.
IB Partition	Use IB Partitions.
Partition Key	Use the partition (fabric domain) in which the underlying port device is a member. The partition key (pkey) is found on and configured by the subnet manager. The pkey may be defined before configuring the subnet manager but the datalink will remain "down" until the subnet partition has been properly configured with the port GUID as a member. It is important to keep partition membership for HCA ports consistent with "IPMP Configuration" on page 291 and "Appliance Cluster Configuration" on page 58 rules on the subnet manager.
IB Link Mode	Use the defined IB Link Mode. IPoIB provides two link modes: Connected (the default) and Unreliable Datagram. Connected mode provides higher throughput and is recommended over Unreliable Datagram. Use Unreliable Datagram only if technically required. Connected mode uses IB queue pairs and dedicates a local queue pair to communicate with a dedicated remote queue pair. Connected mode uses an MTU of 65520 and provides higher throughput than Unreliable Datagram. Unreliable Datagram lets a local queue pair communicate with multiple other queue pairs on any host and messages are communicated unacknowledged at the IB layer. Unreliable Datagram mode uses an MTU of 2044 and yields a lower throughput rate.

Configuring Network Interfaces

Interfaces configure IP addresses via datalinks. They support the following:

- IPv4 and IPv6 protocols.
- IPMP - IP MultiPathing, to improve network reliability by allowing IP addresses to automatically migrate from failed to working datalinks.

For information on how to configure network interfaces, see [“Configuring Network Interfaces \(CLI\)” on page 108](#).

The following interface settings are available:

TABLE 18 Interface Settings

Property	Description
Name	Custom name for the interface
Enable Interface	Enable this interface to be used for IP traffic. If an interface is disabled, the appliance will no longer send or receive IP traffic over it, or make use of any IP addresses configured on it. At present, disabling an active IP interface in an IPMP group will not trigger activation of a standby interface.
Allow Administration	Allow connections to the appliance administration BUI or CLI over this interface. If your network environment included a separate administration network, this could be enabled for the administration network only to improve security
IPv4 Configure with	Either "Static Address List" manually entered, or "DHCP" for dynamically requested
IPv4 Address/Mask	One or more IPv4 addresses in CIDR notation (192.168.1.1/24)
IPv6 Configure with	Either "Static Address List" manually entered, or "IPv6 AutoConfiguration" to use automatically generated link-local address (and site-local if an IPv6 router responds)
IPv6 Address/Mask	One or more IPv6 addresses in CIDR notation (1080::8:800:200C:417A/32)
Directly Reachable Network(s)	Directly reachable subnet(s), expressed as an IP address and mask in CIDR notation, that the local IP address is not a member of, but to which the datalink of its interface is physically connected. This improves scalability by conserving IP addresses, and could ease traffic congestion through core switches and routers.
IP MultiPathing Group	Configure IP multipathing, where a pool of datalinks can be used for redundancy

Configuring Network IP MultiPathing (IPMP)

IP MultiPathing groups are used to provide IP addresses that will remain available in the event of an IP interface failure (such as a physical wire disconnection or a failure of the connection between a network device and its switch) or in the event of a path failure between the system and its network gateways. The system detects failures by monitoring the IP interface's underlying datalink for link-up and link-down notifications, and optionally by probing using test addresses that can be assigned to each IP interface in the group, described below. Any number of IP interfaces can be placed into an IPMP group so long as they are all on the same link (LAN, IB partition, or VLAN), and any number of highly-available addresses can be assigned to an IPMP group.

Each IP interface in an IPMP group is designated either *active* or *standby*:

- **Active** - The IP interface will be used to send and receive data so long as IPMP has determined it is functioning correctly.
- **Standby** - The IP interface will only be used to send and receive data if an active interface (or a previously activated standby) stops functioning.

Multiple active and standby IP interfaces can be configured, but each IPMP group must be configured with at least one active IP interface. IPMP will strive to activate as many standbys as necessary to preserve the configured number of active interfaces. For example, if an IPMP group is configured with two active interfaces and two standby interfaces and all interfaces are functioning correctly, only the two active interfaces will be used to send and receive data. If an active interface fails, one of the standby interfaces will be activated. If the other active interface fails (or the activated standby fails), the second standby interface will be activated. If the active interfaces are subsequently repaired, the standby interfaces will again be deactivated.

IP interface failures can be discovered by either link-based detection or probe-based detection (i.e., a test address is configured).

If probe-based failure detection is enabled on an IP interface, the system will determine which target systems to probe dynamically. First, the routing table will be scanned for gateways (routers) on the same subnet as the IP interface's test address and up to five will be selected. If no gateways on the same subnet were found, the system will send a multicast ICMP probe (to 224.0.0.1 for IPv4 or ff02::1 for IPv6) and select the first five systems on the same subnet that respond. Therefore, for network failure detection and repair using IPMP, you should be sure that at least one neighbor on each link or the default gateway responds to ICMP echo requests. IPMP works with both IPv4 and IPv6 address configurations. In the case of IPv6, the interface's link-local address is used as the test address.

Note - Do not use probe-based failure detection when there are no systems (other than the cluster peer) on the same subnet as the IPMP test addresses that are configured to answer ICMP echo requests.

The system will probe selected target systems in round-robin fashion. If five consecutive probes are unanswered, the IP interface will be considered failed. Conversely, if ten consecutive probes are answered, the system will consider a previously failed IP interface as repaired. You can set the system's IPMP probe failure detection time from the IPMP screen. This time indirectly controls the probing rate and the repair interval -- for instance, a failure detection time of 10 seconds means that the system will send probes at roughly two second intervals and that the system will need 20 seconds to detect a probe-based interface repair. You cannot directly control the system's selected targeted systems, though it can be indirectly controlled through the routing table.

The system will monitor the routing table and automatically adjust its selected target systems as necessary. For instance, if the system using multicast-discovered targets but a route is subsequently added that has a gateway on the same subnet as the IP interface's test address, the system will automatically switch to probing the gateway. Similarly, if multicast-discovered targets are being probed, the system will periodically refresh its set of chosen targets (e.g., because some previously selected targets have become unresponsive).

For step-by-step instructions on building IPMP groups, see: [“IPMP Configuration” on page 291](#).

For information about private local interfaces, see [“Appliance Cluster Configuration” on page 58](#).

Configuring Network Performance and Availability

IPMP and link aggregation are different technologies available in the appliance to achieve improved network performance as well as maintain network availability. In general, you deploy link aggregation to obtain better network performance, while you use IPMP to ensure high availability. The two technologies complement each other and can be deployed together to provide the combined benefits of network performance and availability.

In link aggregations, incoming traffic is spread over the multiple links that comprise the aggregation. Thus, networking performance is enhanced as more NICs are installed to add links to the aggregation. IPMP's traffic uses the IPMP interface's data addresses as they are bound to the available active interfaces. If, for example, all the data traffic is flowing between only two IP addresses but not necessarily over the same connection, then adding more NICs will not improve performance with IPMP because only two IP addresses remain usable.

Performance can be affected by the number of VNICS/VLANs configured on a datalink for a given device, as well as by using a VLAN ID. Configuring multiple VNICS over a given device may impact the performance of all datalinks over that device by up to five percent, even when VNICS are not in use. If more than eight VNICS/VLANs are configured over a given datalink, performance may degrade significantly. Also, if a datalink uses a VLAN ID, all datalink performance for that device may be impacted by an additional five percent.

Configuring Network Routing

The system provides a single IP routing table, consisting of a collection of routing table entries. When an IP packet needs to be sent to a given destination, the system selects the routing entry whose destination most closely matches the packet's destination address (subject to the system's multihoming policy; see below). It then uses the information in the routing entry to determine what IP interface to send the packet on and, if the destination is not directly reachable, the next-hop gateway to use. If no routing entries match the destination, the packet will be dropped. If multiple routing entries tie for closest match (and are not otherwise prioritized by multihoming policy), the system will load-spread across those entries on a per-connection basis.

The system does not act as a router.

The routing table is comprised of routing entries, each of which has the following fields:

TABLE 19 Routing Entry Fields

Field	Description	Examples
Destination	Range of IP destination addresses (in CIDR notation) that can match the route	192.168.0.0/22
Gateway	Next hop (IP address) to send the packet to (except for "system" routes -- see below)	192.168.2.80
Family	Internet protocol	IPv4, IPv6
Type	Origin of the route	dhcp, direct, static, system
Status	Route status	active, inactive (static or direct route associated with a disabled or offline IP interface)
Interface	IP interface the packet will be sent on	igb0

A routing entry with a "destination" field of `0.0.0.0/0` matches any packet (if no other route matches more precisely), and is thus known as a 'default' route. In the BUI, default routes are distinguished from non-default routes by an additional property:

TABLE 20 Distinguishing Default from Non-default Routes

Kind	Route kind	Default, Network

As above, a given packet will be sent on the IP interface specified in the routing entry's "interface" field. If an IPMP interface is specified, then one of the active IP interfaces in the

IPMP group will be chosen randomly on a per-connection basis and automatically refreshed if the chosen IP interface subsequently becomes unusable. Conversely, if a given IP interface is part of an IPMP group, it cannot be specified in the "interface" field because such a route would not be highly-available.

Routing entries come from a number of different origins, as identified by the "type" field. Although the origin of a routing entry has no bearing on how it is used by the system, its origin does control if and how it can be edited or deleted. The system supports the following types of routes:

TABLE 21 Supported Route Types

Type	Description
Static	Created and managed by the appliance administrator.
Dynamic	Created automatically by the appliance via the RIP and RIPv6 dynamic routing protocols (if enabled).
DHCP	Created automatically by the appliance part of enabling an IP interface that is configured to use DHCP. A DHCP route will be created for each default route provided by the DHCP server.
System	Created automatically by the appliance as part of enabling an IP interface. A system route will be created for each IP subnet the appliance can directly reach. Since these routes are directly reachable, the "gateway" field instead identifies the appliance's IP address on that subnet.
Direct	Created and managed as a network interface property: Directly Reachable Network(s). Directly reachable subnet that the local IP address is not a member of, but to which the datalink of its interface is physically connected. This improves scalability by conserving IP addresses, and could ease traffic congestion through core switches and routers.

Note that direct routes are configured as network interfaces using either the Configuration > Network BUI screen or the configuration net interfaces CLI context. Direct routes are not managed via the Routing BUI screen nor the routing CLI context.

TABLE 22 Routing Properties

Property	Description
Multihoming model	Controls the system policy for accepting and transmitting IP packets when multiple IP interfaces are simultaneously enabled. Allowed values are "loose" (default), "adaptive", and "strict". See the discussion below.

If a system is configured with more than one IP interface, then there may be multiple equivalent routes to a given destination, forcing the system to choose which IP interface to send a packet on. Similarly, a packet may arrive on one IP interface, but be destined to an IP address that is hosted on another IP interface. The system's behavior in such situations is determined by the selected multihoming policy. Three policies are supported:

TABLE 23 Multihoming Policies

Policy	Description
Loose	Do not enforce any binding between an IP packet and the IP interface used to send or receive it: 1) An IP packet will be accepted on an IP interface so long as its destination IP address is up on the appliance. 2) An IP packet will be transmitted over the IP interface tied to the route that most specifically matches an IP packet's destination address, without any regard for the IP addresses hosted on that IP interface. If no eligible routes exist, drop the packet.
Adaptive	Identical to loose, except prefer routes with a gateway address on the same subnet as the packet's source IP address: 1) An IP packet will be accepted on an IP interface so long as its destination IP address is up on the appliance. 2) An IP packet will be transmitted over the IP interface tied to the route that most specifically matches an IP packet's destination address. If multiple routes are equally specific, prefer routes that have a gateway address on the same subnet as the packet's source address. If no eligible routes exist, drop the packet.
Strict	Require a strict binding between an IP packet and the IP interface used to send or receive it: 1) An IP packet will be accepted on an IP interface so long as its destination IP address is up on that IP interface. 2) An IP packet will only be transmitted over an IP interface if its source IP address is up on that IP interface. To enforce this, when matching against the available routes, the appliance will ignore any routes that have gateway addresses on a different subnet from the packet's source address. If no eligible routes remain, drop the packet.

When selecting the multihoming policy, a key consideration is whether any of the appliance's IP interfaces will be dedicated to administration (for example, for dedicated BUI access) and thus accessed over a separate administration network. In particular, if a default route is created to provide remote access to the administration network, and a separate default route is created to provide remote access to storage protocols, then the default system policy of "loose" may cause the administrative default route to be used for storage traffic. By switching the policy to "adaptive" or "strict", the appliance will consider the IP address associated with the request as part of selecting the route for the reply. If no route can be found on the same IP interface, the

"adaptive" policy will cause the system to use any available route, whereas the "strict" policy will cause the system to drop the packet.

Configuring Storage

The appliance uses storage pools to manage physical storage devices. After configuring these pools based on physical characteristics and the desired level of data redundancy, you can store filesystems and LUNs, collectively known as shares, in these pools. Shares, which are contained in projects, automatically grow within the disk space allocated to the pool, and pools can span multiple storage devices. Although there is no need to statically size shares, you can control space usage using quotas and reservations. For more information, see [“Space Management for Shares” on page 436](#).

To configure and manage storage, use these tasks:

- Creating a Storage Pool - [BUI](#), [CLI](#)
- Importing an Existing Storage Pool - [BUI](#), [CLI](#)
- Configuring an All-Flash Storage Pool - [BUI](#), [CLI](#)
- Adding a Disk Shelf to an Existing Storage Pool - [BUI](#), [CLI](#)
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool - [BUI](#), [CLI](#)
- Removing a Cache or Log Device from an Existing Storage Pool - [BUI](#), [CLI](#)
- Unconfiguring a Storage Pool - [BUI](#), [CLI](#)
- Renaming a Storage Pool - [BUI](#), [CLI](#)
- Scrubbing a Storage Pool - [BUI](#), [CLI](#)
- [Viewing Pool and Device Status](#)

To understand storage basics, use these topics:

- [“Storage Pool Concepts” on page 150](#)
- [“Data Profiles for Storage Pools” on page 152](#)
- [“Space Management for Shares” on page 436](#)

▼ Creating a Storage Pool (BUI)

Storage pools store data and can be created during or after initial configuration. Pools can contain data drives, and log, read cache, and meta devices. The following task assumes that

initial configuration has been completed. Creating and configuring a storage pool is a two-step process. First, the storage devices are verified for presence and minimum functionality, and you assign drives or even entire disk shelves to the pool. Second, you select a profile for the drives based on your storage needs. If for some reason a pool is unconfigured, you can import it as described in [“Importing an Existing Storage Pool \(BUI\)” on page 128](#).

To reduce redundant data, which can be especially prevalent in replication workloads, consider the benefits of using deduplication. Allocate meta devices if deduplication will be enabled for projects or shares in this pool. For more information, see [Data Deduplication](#). There is also an all-flash storage pool, which utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. See [“Configuring an All-Flash Storage Pool \(BUI\)” on page 130](#).

- Before You Begin**
- For recommendations on how many drives to select per pool, see [“Number of Devices per Pool” on page 150](#).
 - To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
 - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.
 - To use the enhanced data depulication feature in a storage pool, upgrade to software release OS8.7.0 or later and accept all deferred updates, including Data Deduplication v2. See [“Data Deduplication v2 Deferred Update” in Oracle ZFS Storage Appliance Customer Service Manual](#).
1. **Go to Configuration > Storage.**
 2. **Next to Available Pools, click the add icon .**
 3. **Type a name for the storage pool and click APPLY.**
 4. **Select the number of data drives for the storage pool for each disk shelf. You can also select available log, cache, and meta devices.**

For more information on log, cache, and meta devices, see [“Data Profiles for Storage Pools” on page 152](#).

Confirm that all devices are present and minimally functional, and allocate them to a storage pool. ABORT COMMIT

Verify and allocate devices Step 1 of 2

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

NAME	MODEL	TYPE	RPM	DATA	LOG	CACHE	META
chena	Sun Netra X4270 M3	System	--	-	-	-	-
1236FMT002	Oracle Storage DE2-24C	HDD	7200	5 (13.6T)	-	0 (1G)	0 (1G)
1236FMT001	Oracle Storage DE2-24C	HDD	7200	24 (65.5T)	-	0 (1G)	0 (1G)
1237FMT00H	Oracle Storage DE2-24C	HDD	7200	20 (54.6T)	2 (137G)	0 (1G)	0 (1G)



Caution - Once a data disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

If all attached disk shelves do not appear, click ABORT, check the disk shelf cabling and power, and begin this procedure again.

- If all drives are the same size or rotational speed, or if one size is selected among multiple sizes, the maximum number of drives available is allocated by default. If the storage device contains drives of different rotational speeds or models, no drives are allocated by default.
- It is strongly recommended that pools include only devices of the same size and rotational speed to provide consistent performance characteristics.
- Monitor or limit space usage because you may experience reduced performance when pools approach full capacity.

5. Click COMMIT.

The drives are allocated to the storage pool, and verified for presence and minimum functionality. If verification fails, click ABORT, fix the problem, and begin this procedure again. If you allocate a pool with missing or failed devices, you will not be able to add the missing or failed devices later.

6. On the Choose Storage Profile screen, select the desired data profile that meets your reliability, availability, serviceability, and performance goals.

For a description of each profile, click on the data profile name, or see [“Data Profiles for Storage Pools” on page 152](#).

7. If you allocated log, cache, or meta devices, select the appropriate profiles.

- For log devices, click Log Profile and select either the mirrored or striped profile. If you allocated an even number of log devices to the pool, select the mirrored profile.



Caution - A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see [“Data Profiles for Storage Pools” on page 152](#).

- For cache devices, the profile is always striped, as shown under Cache Profile.
- For meta devices, click Metadata Profile and select either the mirrored or striped profile.

Note - Once meta devices are added to a storage pool, they cannot be removed from the pool.

8. Click COMMIT.

Related Topics

- [“Data Profiles for Storage Pools” on page 152](#)
- [“Importing an Existing Storage Pool \(BUI\)” on page 128](#)
- [“Adding a Disk Shelf to an Existing Storage Pool \(BUI\)” on page 133](#)
- [“Renaming a Storage Pool \(BUI\)” on page 145](#)
- [“Storage Pool Concepts” on page 150](#)
- [Data Deduplication](#)

▼ Creating a Storage Pool (CLI)

Storage pools store data and can be created during or after initial configuration. Pools can contain data drives, and log, read cache, and meta devices. The following task assumes that initial configuration has been completed. Creating and configuring a storage pool is a two-step process. First, the storage devices are verified for presence and minimum functionality, and you assign drives or even entire disk shelves to the pool. Second, you select a profile for the drives based on your storage needs. If for some reason a pool is unconfigured, you can import it as described in [“Importing an Existing Storage Pool \(CLI\)” on page 128](#).

To reduce redundant data, which can be especially prevalent in replication workloads, consider the benefits of using deduplication. Allocate meta devices if deduplication will be enabled for projects or shares in this pool. For more information, see [Data Deduplication](#). There is also an all-flash storage pool, which utilizes SSDs as data devices and optional log devices,

but does not contain read cache or meta devices. See [“Configuring an All-Flash Storage Pool \(CLI\)” on page 131](#).

- Before You Begin**
- For recommendations on how many drives to select per pool, see [“Number of Devices per Pool” on page 150](#).
 - To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
 - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
 - To use the enhanced data deduplication feature in a storage pool, upgrade to software release OS8.7.0 or later and accept all deferred updates, including Data Deduplication v2. See [“Data Deduplication v2 Deferred Update” in Oracle ZFS Storage Appliance Customer Service Manual](#).

1. **Go to configuration storage.**
2. **Enter `config` and a name for the new storage pool.**

```
hostname: configuration storage> config pool0
hostname: configuration storage (pool0) verify>
```

3. **Enter `show` to see the device information for the pool:**

```
hostname:configuration storage (pool0) verify> show
```

ID	STATUS	ALLOCATION	DATA	LOG	CACHE	META	RPM
0	ok	custom	0	0	0/4	0/4	1.86T
1	ok	custom	0	0/2	34G	0	15000
2	ok	custom	0	0/2	34G	0	15000

4. **Enter `set` and the disk shelf or controller ID, and the number of data drives to use. You can also select available cache, meta, and log devices.**

For more information on log, cache, and meta devices, see [“Data Profiles for Storage Pools” on page 152](#).



Caution - Once a data disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, `1-data=8` allocates eight data drives from the first disk shelf.

```
hostname:configuration storage (pool1) verify> set 1-data=8
1-data = 8
```

This example allocates one cache device from the controller:

```
hostname:configuration storage (pool1) verify> set 0-cache=1
0-cache = 1
```

This example allocates one meta device from the controller:

```
hostname:configuration storage (pool1) verify> set 0-meta=1
0-meta = 1
```

5. Enter done.

```
hostname:configuration storage (pool1) verify> done
```

6. Enter show to display the profile.

```
hostname:configuration storage (pool1) config> show
PROFILE          CAPACITY  NSPF  DESCRIPTION
log_profile = log_stripe  17G   no   Striped log
```

Note - If you allocated cache devices to the pool, the profile is always striped.

7. If you allocated log devices to the pool, enter set log_profile= and set the log profile to either log_mirror or log_stripe. Use log_mirror if the pool contains an even number of log devices.



Caution - A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see [“Data Profiles for Storage Pools” on page 152](#).

```
hostname:configuration storage (pool1)> set log_profile=log_mirror
```

8. If you allocated meta devices to the pool, enter set meta_profile= and set the meta profile to either meta_mirror or meta_stripe.

```
hostname:configuration storage (pool1)> set meta_profile=meta_mirror
```

9. Enter done to complete the task.

```
hostname:configuration storage (pool1)> done
```

Related Topics

- [“Data Profiles for Storage Pools” on page 152](#)
- [“Importing an Existing Storage Pool \(CLI\)” on page 128](#)
- [“Adding a Disk Shelf to an Existing Storage Pool \(CLI\)” on page 135](#)

- [“Renaming a Storage Pool \(CLI\)” on page 146](#)
- [“Storage Pool Concepts” on page 150](#)
- [Data Deduplication](#)

▼ Importing an Existing Storage Pool (BUI)

The import action allows you to import an unconfigured storage pool. A storage pool can be unconfigured because of an inadvertent action, factory reset, or service operation to recover user data. Importing a storage pool requires scanning all attached storage devices and discovering any existing state. This can take a significant amount of time, during which no other storage configuration activities can take place.

Before You Begin Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to **Maintenance > System**.

1. Go to Configuration > Storage.

A list of storage pools is displayed, including some identifying characteristics. If the storage has been destroyed or is incomplete, the storage pool is not importable. Unlike storage configuration, the storage pool name is not initially shown, but it is shown after selecting the storage pool.

2. Click IMPORT.

3. Select the storage pool you want to import.

By default, the previous storage pool names are displayed.

4. To rename the storage pool, click the pool name and change it.

5. Click COMMIT.

Related Topics

- [“Unconfiguring a Storage Pool \(BUI\)” on page 143](#)
- [“Renaming a Storage Pool \(BUI\)” on page 145](#)

▼ Importing an Existing Storage Pool (CLI)

The import action allows you to import an unconfigured storage pool. A storage pool can be unconfigured because of an inadvertent action, factory reset, or service operation to recover

user data. Importing a storage pool required iterating over all attached storage devices and discovering any existing state. This can take a significant amount of time, during which no other storage configuration activities can take place.

Before You Begin Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

1. Go to configuration storage.

2. Enter `import`.

```
hostname:configuration storage (pool0)> import
```

Search for storage. Begin the process of searching for existing storage pools.

Subcommands that are valid in this context:

```
help [topic]      => Get context-sensitive help. If [topic] is specified,
                    it must be one of "builtins", "commands", "general",
                    "help" or "script".

show              => Show information pertinent to the current context

abort             => Abort this task (potentially resulting in a
                    misconfigured system)

done              => Finish operating on "discover"
```

```
hostname:configuration storage (pool0) discover>
```

3. Enter `done`.

4. Enter `show`.

```
hostname:configuration storage (pool0)> show
```

Pools:

POOL	OWNER	DATA PROFILE	LOG PROFILE	STATUS	ERRORS
-> pool0	hostname	mirror	log_stripe	online	0
pool1	hostname	-	-	exported	-

Properties:

```
pool = pool0
status = online
errors = 0
owner = hostname
profile = mirror
log_profile = log_stripe
```

```
cache_profile = cache_stripe
scrub = none requested
```

5. Enter **set pool=** and the name of the pool you want to import.

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage select> set pool=pool1
pool = pool1
```

A message reminds you to verify that storage is correctly attached and functioning.

6. Enter **done**.

Related Topics

- [“Unconfiguring a Storage Pool \(CLI\)” on page 144](#)
- [“Renaming a Storage Pool \(CLI\)” on page 146](#)

▼ Configuring an All-Flash Storage Pool (BUI)

An all-flash storage pool utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. All-flash pools are suitable for virtualization environments or backup workloads.

Before You Begin

- Follow the cabling guidelines for all-flash shelves described in [“Cabinet and Cabling Guidelines” in Oracle ZFS Storage Appliance Cabling Guide](#).
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).
- To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).

Note - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.

1. Go to **Configuration > Storage**.
2. Click the add icon  above the list of storage pools.
3. From the **Verify and allocate devices** screen, select the storage type **SSD**, and then select a device size.

Verify and allocate devices Step 1 of 2

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

Storage type ☐ HDD ☒ SSD
 Device size ☐ 373G ☒ Any

NAME	MODEL	TYPE	RPM	DATA	LOG	CACHE
sbm-7440b	Oracle ZFS Storage ZS4-4	System	--	-	-	-
1504NMK01M	Oracle Storage DE3-24C	HDD	7200	-	-	-
1407NM400E	Oracle Storage DE3-24P	SSD	--	0 (1.46T)	-	-
1504NMK01W	Oracle Storage DE3-24C	HDD	7200	-	-	-
1242FM4003	Oracle Storage DE2-24P	HDD	10000	-	-	-
1452NMK00T	Oracle Storage DE3-24C	HDD	7200	-	-	-

- For each SSD disk shelf, select the number of drives to include in the pool.

Note - An all-flash pool cannot contain read cache devices or meta devices.

- (Optional) Select log devices to add to the all-flash pool.
- Click **COMMIT**.
- On the **Configure Added Storage** screen, select the data profile appropriate for your workload that balances performance, availability, and capacity.
For a description of available profiles, see [“Data Profiles for Storage Pools”](#) on page 152.
- (Optional) If you allocated log devices, select an appropriate profile.
- Click **COMMIT**.

Related Topics

- All-Flash Storage Configuration
- “Setting a Threshold Alert for SSD Endurance (BUI)” in *Oracle ZFS Storage Appliance Customer Service Manual*

▼ Configuring an All-Flash Storage Pool (CLI)

An all-flash storage pool utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. All-flash pools are suitable for virtualization environments or backup workloads.

- Before You Begin**
- Follow the cabling guidelines for all-flash shelves described in [“Cabinet and Cabling Guidelines”](#) in *Oracle ZFS Storage Appliance Cabling Guide*.

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).
- To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).

Note - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

1. Verify SSDs are correctly attached and functioning.

If any devices are missing or malfunctioning, make the necessary corrections.

Note - An all-flash pool cannot contain read cache devices or meta devices.

2. Go to configuration storage, enter config and a unique name for the pool:

```
hostname:configuration storage> config allflashpool
```

Instructions and subcommands that can be used in this context are displayed.

3. Show the available devices for the pool.

```
hostname:configuration storage verify> show
```

	ID	STATUS	ALLOCATION	DATA	LOG	CACHE	RPM	TYPE
	---	-----	-----	-----	-----	-----	-----	

	0	ok	custom	0	0	0		
system								
	1	ok	custom	0/7 3.46T	0/2 373G	0		ssd
	2	ok	custom	0/24 6.55T	0	0		ssd
				-----	-----	-----		
				0	0	0		

4. List the available properties:

```
hostname:configuration storage verify> help properties
```

```
0                => Chassis 0
1-data           => Chassis 1 data
1-log            => Chassis 1 log
2                => Chassis 2
2-data           => Chassis 2 data
```

5. Assign the devices to the pool, as shown in this example:

```
hostname:configuration storage verify> set 1-data=3 2-data=3
```



```
1-data = 3
2-data = 3
```

This example assigns 3 devices from chassis 1 (1-data=3) and 3 devices from chassis 2 (2-data=3) to the pool.

6. (Optional) Select log devices to add to the all-flash pool.

7. Enter done to close verify.

```
hostname:configuration storage verify> done
```

8. Show the available storage profile types:

```
hostname:configuration storage config>show
```

9. Select the data profile appropriate for your workload, that balances performance, availability, and capacity.

For a description of available profiles, see [“Data Profiles for Storage Pools” on page 152](#).

```
hostname:configuration storage config>set profile=
```

10. (Optional) If you allocated log devices, select an appropriate profile.

11. Enter done.

```
hostname:configuration storage config> done
```

Related Topics

- [All-Flash Storage Configuration](#)
- [“Setting a Threshold Alert for SSD Endurance \(CLI\)” in Oracle ZFS Storage Appliance Customer Service Manual](#)

▼ Adding a Disk Shelf to an Existing Storage Pool (BUI)

Use the following task to add a disk shelf to an existing storage pool.

- Before You Begin**
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).

- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
- If there is insufficient storage to configure the system for the data profile and its options, some attributes may not be supported. For example, it is impossible to preserve NSPF characteristics when adding a single disk shelf to a double parity RAID configuration with the NSPF option. You can add the disk shelf, but cannot use the NSPF option.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.



Caution - Once a disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

1. **Install the new disk shelf using [“Adding a New Disk Shelf” in Oracle ZFS Storage Appliance Customer Service Manual](#).**
2. **Go to Configuration > Storage.**
3. **From the Available Pools list, select an online pool to which to add the disk shelf.**
4. **Click ADD.**
5. **For this disk shelf, select the number of data drives for the storage pool.**

If the new disk shelf does not appear, click ABORT, check the disk shelf cabling and power, and begin this procedure again.

 - If all drives are the same size or rotational speed, or if one size is selected among multiple sizes, the maximum number of drives available is allocated by default. If the storage device contains drives of different rotational speeds or models, no drives are allocated by default.
 - It is strongly recommended that pools include only devices of the same size and rotational speed to provide consistent performance characteristics.
 - Monitor or limit space usage because you may experience reduced performance when pools approach full capacity.
6. **(Optional) Add any cache or log devices from the disk shelf to the pool.**
7. **Click COMMIT.**
8. **For data drives, select the same data profile used in the existing pool.**
9. **If you allocated log or cache devices, select the appropriate profiles.**

- For log devices, click Log Profile and select either the mirrored or striped profile. If you allocated an even number of log devices to the pool, select the mirrored profile.



Caution - A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see [“Data Profiles for Storage Pools” on page 152](#).

- For cache devices, the profile is always striped, as shown under Cache Profile.

10. **Click COMMIT.**

Related Topics

- [“Unconfiguring a Storage Pool \(BUI\)” on page 143](#)
- [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(BUI\)” on page 137](#)

▼ Adding a Disk Shelf to an Existing Storage Pool (CLI)

Use the following task to add a disk shelf to an existing storage pool.

Before You Begin

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).
- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
- If there is insufficient storage to configure the system for the data profile and its options, some attributes may not be supported. For example, it is impossible to preserve NSPF characteristics when adding a single disk shelf to a double parity RAID configuration with the NSPF option. You can add the disk shelf, but cannot use the NSPF option.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.



Caution - Once a disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

1. **Install the new disk shelf using [“Adding a New Disk Shelf” in Oracle ZFS Storage Appliance Customer Service Manual](#).**
2. **Go to configuration storage.**

3. If you have multiple pools, a default pool is selected and displayed. If this is not the pool to which you want to add the device, enter `set pool=` and specify another online pool.

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

A message reminds you to verify that the device is correctly installed. Note that mixing device types and speeds is strongly discouraged.

4. Enter `add`.

```
hostname:configuration storage (pool1)> add
```

5. Enter `show` to see the device information for the pool:

```
hostname:configuration storage (pool1) verify> show
ID STATUS  ALLOCATION  DATA  LOG    CACHE  RPM
0      ok      custom      0      0      0/4    1.86T
1      ok      custom      0      0/2    34G    0      15000
2      ok      custom      0      0/2    34G    0      15000
```

6. Specify which disk shelf or the controller and the number of data drives to use. ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, `1-data=8` allocates eight data drives from the first disk shelf.

```
hostname:configuration storage (pool1) verify> set 1-data=8
1-data = 8
```

7. (Optional) Specify which disk shelf or the controller and the number of log or cache devices to use.

ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, `set 0-cache=1` allocates one cache device from the controller:

```
hostname:configuration storage (pool1) verify> set 0-cache=1
0-cache = 1
```

8. Enter `done`.

```
hostname:configuration storage (pool1) verify> done
```

The storage devices are verified for presence and minimum functionality. If verification fails, fix the problem, and begin this procedure again. If you allocate a pool with missing or failed devices, you will not be able to add the missing or failed devices later.

9. Enter `show` to display the profile.

```
hostname:configuration storage (pool1) config> show
PROFILE  CAPACITY  NSPF  DESCRIPTION
log_profile  17G      no    Striped log
```

10. Enter the same data profile as the remainder of the pool by entering `set profile=` and the profile name.**11. Enter `done`.****12. If you allocated log devices to the pool, enter `set log_profile=` and either `log_mirror` or `log_stripe`. Use `log_mirror` if the pool contains an even number of log devices.**

```
hostname:configuration storage (pool1)> set log_profile=log_mirror
```

Note - If you allocated cache devices to the pool, the profile is always striped.

13. Enter `done`.**Related Topics**

- [“Unconfiguring a Storage Pool \(CLI\)” on page 144](#)
- [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(CLI\)” on page 138](#)

▼ Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI)

Use the following task to add a log, read cache, or meta device to an existing storage pool.

- Before You Begin**
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).
 - You must select the same data profile currently used in the existing pool. To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
 - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.
 - A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

1. **Install the new log, read cache, or meta device into the first available and appropriate slot. To determine the appropriate slot, see [“Disk Shelf Configurations” in Oracle ZFS Storage Appliance Customer Service Manual](#).**
2. **Go to Configuration > Storage.**
3. **From the Available Pools list, select an online pool to which to add the device.**
4. **Click ADD.**
5. **Select the device to add to the pool and click COMMIT.**
6. **Select the appropriate profiles.**
 - For log devices, click Log Profile and select either the mirrored or striped profile. Use the mirrored profile if the pool now contains an even number of log devices.



Caution - A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see [“Data Profiles for Storage Pools” on page 152](#).

- For cache devices, the profile is always striped, as shown under Cache Profile.
 - For meta devices, click Metadata Profile and select either the mirrored or striped profile. The striped profile is recommended for better performance in the event of a meta device failure.
7. **Click COMMIT.**

Related Topics

- [“Removing a Cache or Log Device from an Existing Storage Pool \(BUI\)” on page 141](#)
- [“Adding a Disk Shelf to an Existing Storage Pool \(BUI\)” on page 133](#)
- [“Data Profiles for Storage Pools” on page 152](#)
- [“Storage Pool Concepts” on page 150](#)

▼ Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI)

Use the following task to add a read cache device or log device to an existing storage pool.

- Before You Begin**
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see [“Storage Pool Concepts” on page 150](#).
 - You must select the same data profile currently used in the existing pool. To understand the different data profiles, see [“Data Profiles for Storage Pools” on page 152](#).
 - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
 - A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

1. **Install the new log, read cache, or meta device into the first available and appropriate slot. To determine the appropriate slot, see [“Disk Shelf Configurations” in Oracle ZFS Storage Appliance Customer Service Manual](#).**
2. **Go to configuration storage.**
3. **If you have multiple pools, a default pool is selected and displayed. If this is not the pool to which you want to add a device, enter `set pool=` and specify another online pool.**

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

A message reminds you to verify that storage is correctly attached and functioning.

4. **Enter `add`:**

```
hostname:configuration storage (pool1)> add
```

5. **Enter `show` to display device information for the pool.**

```
hostname:configuration storage (pool1) verify> show
```

ID	STATUS	ALLOCATION	DATA	LOG	CACHE	META	RPM
0	ok	custom	0	0	0/4	0/4	1.86T
1	ok	custom	0	0/2 34G	0	0	15000
2	ok	custom	0	0/2 34G	0/2	0	15000

6. **Enter `set` and use tab completion to see if cache, meta, and log devices are available.**

```
hostname:configuration storage (pool1) verify> set
0-cache 1-data 2-cache 2-meta 2-log
```

7. **Enter `set` and the disk shelf or controller ID, and the number of log, cache, or meta devices to use.**

ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, 2-log=1 allocates one log device from the second disk shelf.

```
hostname:configuration storage (pool1) verify> set 2-log=1
2-log = 1
```

Note - A value of "1-log=2" would allocate two log devices from the first disk shelf.

This example allocates one cache device from the second disk shelf.

```
hostname:configuration storage (pool1) verify> set 2-cache=1
2-cache = 1
```

This example allocates one meta device from the second disk shelf:

```
hostname:configuration storage (pool1) verify> set 2-meta=1
2-meta = 1
```

8. **Enter `done`.**

```
hostname:configuration storage (pool1) verify> done
```

9. **Enter `show` to display the profile.**

```
hostname:configuration storage (pool1) config> show
PROFILE          CAPCTY  NSPF  DESCRIPTION
log_profile = log_stripe    17G    no  Striped log
```

Note - If you allocated cache devices to the pool, the profile is always striped.

10. **If you allocated log devices to the pool, enter `set log_profile=` and set the log profile to either `log_mirror` or `log_stripe`. Use `log_mirror` if the pool now contains an even number of log devices.**



Caution - A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see [“Data Profiles for Storage Pools” on page 152](#).

```
hostname:configuration storage (pool1)> set log_profile=log_mirror
```

11. **If you allocated meta devices to the pool, enter `set meta_profile=` and set the meta profile to either `meta_mirror` or `meta_stripe`.**


```
hostname:configuration storage (pool1)> set meta_profile=meta_mirror
```

12. Enter done to complete the task.

```
hostname:configuration storage (pool1)> done
```

Related Topics

- [“Removing a Cache or Log Device from an Existing Storage Pool \(CLI\)” on page 142](#)
- [“Adding a Disk Shelf to an Existing Storage Pool \(CLI\)” on page 135](#)
- [“Data Profiles for Storage Pools” on page 152](#)
- [“Storage Pool Concepts” on page 150](#)

▼ Removing a Cache or Log Device from an Existing Storage Pool (BUI)

Use the following task to remove a read cache or log device from an existing storage pool. This capability is useful when preparing for a system update that requires the removal of certain cache devices.

Note - Meta devices cannot be removed from a storage pool.

If a pool has cache devices on both controllers of a clustered configuration, you must perform this procedure on each controller.

To add a device to a different, existing storage pool, see [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(BUI\)” on page 137](#).

Before You Begin Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.

- 1. Go to Configuration > Storage.**
- 2. From the Available Pools list, select an online pool from which to remove the device.**
- 3. Click REMOVE.**
- 4. Select the number of log and cache devices to be removed from the storage pool.**

Note - If the log devices use a mirrored profile, a message reminds you to select an even number of log devices to remove. If they use a striped profile, you may remove an even or odd number of devices.

5. **Click COMMIT.**

Related Topics

- [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(BUI\)” on page 137](#)

▼ Removing a Cache or Log Device from an Existing Storage Pool (CLI)

Use the following task to remove a read cache or log device from an existing storage pool. This capability is useful when preparing for a system update that requires the removal of certain cache devices.

Note - Meta devices cannot be removed from a storage pool.

If a pool has cache devices on both controllers of a clustered configuration, you must perform this procedure on each controller.

To add a device to a different, existing storage pool, see [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(CLI\)” on page 138](#).

Before You Begin Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

1. **Go to configuration storage.**
2. **If you have multiple pools, a default pool is displayed and selected. If this is not the pool to which you want to add the device, enter `set pool=` and specify another online pool.**

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

3. Enter show to see the device information for the pool.

```
hostname:configuration storage (pool1) verify> show
ID STATUS ALLOCATION DATA LOG CACHE META RPM
0 ok custom 0 0 0/4 0/4 1.86T
1 ok custom 0 0/2 34G 0 0 15000
2 ok custom 0 0/2 34G 0/2 0 15000
```

4. Enter remove.

```
hostname:configuration storage (pool1)> remove
```

5. Specify the controller or disk shelf, and the number of log or cache devices to remove.

ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, 1-log=2 removes two log devices from the first disk shelf:

```
hostname:configuration storage (pool1) remove> set 1-log=2
1-log = 2
```

This example removes one cache device from the controller:

```
hostname:configuration storage (pool1) remove> set 0-cache=1
0-cache = 1
```

6. Enter done.

```
hostname:configuration storage (pool1) remove> done
```

Note - If the log devices use a mirrored profile, a message reminds you to select an even number of log devices to remove. If the log devices use a striped profile, you may remove an even or odd number of devices.

Related Topics

- [“Adding a Cache, Meta, or Log Device to an Existing Storage Pool \(CLI\)” on page 138](#)

▼ Unconfiguring a Storage Pool (BUI)

Unconfiguring a storage pool removes any active filesystems and LUNs and makes the raw storage available for future storage configuration. This process can be undone by importing the unconfigured storage pool, if the raw storage has not since been used as part of an active storage pool.



Caution - Unconfiguring a pool renders data inaccessible, creates the potential for data loss, and fails inbound replications.

Before You Begin

- Do not unconfigure a pool while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.
- Do not unconfigure a pool while the peer controller is down or unreachable.
- If an error message reports that the target is in use, wait and try the operation again.

1. **Go to Configuration > Storage.**
2. **From the Available Pools list, select an online pool to unconfigure.**
3. **Click UNCONFIG.**

Related Topics

- [“Importing an Existing Storage Pool \(BUI\)” on page 128](#)
- [“Renaming a Storage Pool \(BUI\)” on page 145](#)

▼ Unconfiguring a Storage Pool (CLI)

Unconfiguring a storage pool removes any active filesystems and LUNs and makes the raw storage available for future storage configuration. This process can be undone by importing the unconfigured storage pool, if the raw storage has not since been used as part of an active storage pool.



Caution - Unconfiguring a pool renders data inaccessible, creates the potential for data loss, and fails inbound replications.

Before You Begin

- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to maintenance system updates.
- Do not unconfigure a pool while the peer controller is down or unreachable.
- If an error message reports that the target is in use, wait and try the operation again.

1. **Go to configuration storage.**
2. **If you have multiple pools, a default pool is selected and displayed. If this is not the pool you want to unconfigure, enter `set pool=` and specify another online pool.**

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

3. Enter `unconfig`.

```
hostname:configuration storage (pool1)> unconfig
```

4. Enter `done`.

Related Topics

- [“Importing an Existing Storage Pool \(CLI\)” on page 128](#)
- [“Renaming a Storage Pool \(CLI\)” on page 146](#)

▼ Renaming a Storage Pool (BUI)

To rename a storage pool, you must unconfigure it and then immediately import it with a new name. While storage is unconfigured, data will be inaccessible and there is a potential for data loss. Importing a storage pool can take a considerable amount of time.

- Before You Begin**
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System.
 - Do not rename a pool while the peer controller is down or unreachable.

1. **Go to Configuration > Storage.**
2. **From the Available Pools list, select the online pool to rename.**
3. **Click UNCONFIG, then COMMIT.**
4. **Click IMPORT, then select the storage pool you just unconfigured.**
5. **Click the storage pool name and change it.**
6. **Click COMMIT.**

Related Topics

- [“Unconfiguring a Storage Pool \(BUI\)” on page 143](#)

- [“Importing an Existing Storage Pool \(BUI\)” on page 128](#)

▼ Renaming a Storage Pool (CLI)

To rename a storage pool, you must unconfigure it and then immediately import it with a new name. While storage is unconfigured, data will be inaccessible and there is a potential for data loss. Importing a storage pool can take a considerable amount of time.

- Before You Begin**
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
 - Do not rename a pool while the peer controller is down or unreachable.

1. **Go to configuration storage.**
2. **If you have multiple pools, a default pool is selected and displayed. If this is not the pool you want to rename, enter `set pool=` and specify another online pool.**

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

3. **Enter `unconfig`.**

```
hostname:configuration storage (pool1)> unconfig
```

4. **Enter `done`.**

5. **Enter `import`.**

```
hostname:configuration storage> import
```

Search for storage. Begin the process of searching for existing storage pools.

Subcommands that are valid in this context:

```
help [topic]      => Get context-sensitive help. If [topic] is specified,
                    it must be one of "builtins", "commands", "general",
                    "help" or "script".
```

```
show              => Show information pertinent to the current context
```

abort => Abort this task (potentially resulting in a misconfigured system)

done => Finish operating on "discover"

```
hostname:configuration storage> discover>
```

6. Enter done.

7. To select the storage pool you just unconfigured, enter `set pool=` and the pool name.

```
hostname:configuration storage select> set pool=pool1
pool = pool1
```

8. To rename the storage pool, enter `set name=` and a new name.

```
hostname:configuration storage (pool1)> set name=NewPool
pool = NewPool
```

9. Enter done.

Related Topics

- [“Unconfiguring a Storage Pool \(CLI\)” on page 144](#)
- [“Importing an Existing Storage Pool \(CLI\)” on page 128](#)

▼ Scrubbing a Storage Pool (BUI)

Scrubbing a storage pool verifies the content by checking for errors. If any unrecoverable errors are found, either through a scrub or through normal operation, the BUI displays the affected files.

In general, a scrub should be performed as often as your oldest backup expires, at a minimum. The recommended time frame for performing a scrub is quarterly. A scrub should also be run before performing a software upgrade.

- 1. Go to Configuration > Storage.**
- 2. From the Available Pools list, select the online pool to scrub.**
- 3. Click SCRUB.**

The scrub status is displayed, including the date and time of the scrub, the number of errors, and the filenames with errors.

4. (Optional) To stop the scrub, click STOP.

Clicking SCRUB again resumes the scrub from where it was stopped.

Related Topics

- [“Storage Pool Concepts” on page 150](#)

▼ Scrubbing a Storage Pool (CLI)

Scrubbing a storage pool verifies the content by checking for errors. If any unrecoverable errors are found, either through a scrub or through normal operation, the CLI displays the affected files. If desired, the scrub process can be stopped before completion.

In general, a scrub should be performed as often as your oldest backup expires, at a minimum. The recommended time frame for performing a scrub is quarterly. A scrub should also be run before performing a software upgrade.

- 1. Go to configuration storage.**
- 2. If you have multiple pools, a default pool is selected and displayed. If this is not the pool you want to scrub, enter `set pool=` and specify another online pool.**

Note - If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
pool = pool1
```

3. Enter scrub start.

```
hostname:configuration storage (pool1)> scrub start
```

4. (Optional) Stop the scrub before it has completed by entering scrub stop.

```
hostname:configuration storage (pool1)> scrub stop
```

Entering scrub start again resumes the scrub from where it was stopped.

Related Topics

- [“Storage Pool Concepts” on page 150](#)

▼ Viewing Pool and Device Status (BUI)

You can check the status of pool and component devices. If there is a problem with a pool, details about device status will also be listed.

1. **Navigate to Configuration > Storage.**
2. **Click a pool to select it and see more details.**

Refer to the following table for a description of the pool status.

Pool Status	Description
Online	A pool that has all devices operating normally.
Degraded	A pool with one or more failed devices, but the data is still available due to a redundant configuration.
Faulted	One or more component devices are offline and there are insufficient replicas to continue functioning.
Offline	A pool was explicitly taken offline.
Unavailable	A pool with corrupted metadata, or one or more unavailable devices and insufficient replicas to continue functioning.
Exported	A pool is active on the cluster peer and is ready for a cluster failback to occur.

3. **View the selected pool's device statuses under the Device Status section.**

Refer to the following table for a description of the device status.

Device Status	Description
Online	The device is online and functioning. You may not see this status, and instead see the message "No device faults have been detected in the storage pool."
Degraded	The device is not in an optimal state. Either it is expected to fail soon, or a spare has not finished resilvering yet.
Faulted	The device is faulty; more information can be found in the maintenance logs.
Offline	The device was explicitly taken offline; no reads or writes will occur to this device until it has been onlined.
Removed	The device has been physically removed.
Hot Spare	This spare device is actively being used as a data disk in the pool as a replacement for a device that failed.
Unavailable	The device could not be opened or the pool could not detect this device.

4. **To see more detailed pool and device error information, navigate to Maintenance > Problems for active errors, or Maintenance > Logs for a history of all problems.**

Storage Pool Concepts

Storage is configured in pools that are characterized by their underlying data redundancy, and provide space that is shared across all filesystems and LUNs. More information about how storage pools relate to individual filesystems or LUNs can be found in [“About Storage Pools, Projects, and Shares” on page 402](#).

Storage Pool Configuration

Pools can be created by configuring a new pool, or importing an existing pool. Importing an existing pool is only used to import pools previously configured on an Oracle ZFS Storage Appliance, and is useful in case of accidental reconfiguration, such as when moving pools between controllers, or due to catastrophic controller failure.

Multiple Pools

Each controller can have any number of pools, and each pool can be assigned ownership independently in a cluster. With the ability to control access to log and cache devices on a per-share basis, the recommended mode of operation is a single pool. While arbitrary number of pools are supported, creating multiple pools with the same redundancy characteristics owned by the same cluster head is not advised. Doing so will result in poor performance, suboptimal allocation of resources, artificial partitioning of storage, and additional administrative complexity. Configuring multiple pools on the same host is only recommended when drastically different redundancy or performance characteristics are desired, for example a mirrored pool for databases and a RAID-Z pool for streaming workloads.

Number of Devices per Pool

Drives within all of the chassis can be allocated individually; however, care should be taken when allocating disks from disk shelves to ensure optimal pool configurations. In general, fewer pools with more disks per pool are preferred because they simplify management and provide a higher percentage of overall usable capacity.

While the system can allocate storage in any increment desired, it is recommended that each allocation include a minimum of 8 disks across all disk shelves and ideally many more.

Drive Characteristics and Performance

Follow these restrictions when configuring storage pools:

- All data disks contained within a head node or disk shelf must have the same rotational speed (media rotation rate). The appliance software will detect misconfigurations and generate a fault for the condition.
- Due to unpredictable performance issues, avoid mixing different disk rotational speeds within the same pool.
- For optimal performance, do not combine disk shelves with different disk rotational speeds on the same SAS fabric (HBA connection). Such a mixture operates correctly, but likely results in slower performance of the faster devices.
- When creating a new pool, avoid mixing different data disk capacities because all disks are then limited to the smallest capacity disk in the pool. When adding a higher capacity disk to an existing pool, the larger disk's capacity is maintained. However, the system preferentially writes to new disks until they begin to reach the same capacity utilization as the old disks. To maintain performance, add as many new higher capacity disks as the total number of disks in the original pool.
- A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

Storage Pool Capacity

When allocating raw storage to pools, keep in mind that filling pools completely will result in significantly reduced performance, especially when writing to shares or LUNs. These effects become more noticeable as the pool reaches full capacity.

All-Flash Storage Configuration

The Oracle Storage Drive Enclosure DE3-24P can be configured as all-flash storage with fully populated flash-based SSD data devices and optional log devices. All-flash storage provides low-latency I/O that increases workload performance.

An all-flash storage pool contains data SSDs and optional log devices. Read flash cache and meta devices cannot be part of an all-flash pool. The remaining lifetime of SSDs can be monitored using threshold alerts.

Related Topics

- Configuring an All-Flash Storage Pool [BUI](#), [CLI](#)
- “Disk Shelf Configurations” in *Oracle ZFS Storage Appliance Customer Service Manual*
- “SSD Endurance” in *Oracle ZFS Storage Appliance Customer Service Manual*

Data Profiles for Storage Pools

After storage devices are physically verified and resources are allocated for a storage pool, the next step is to choose a storage profile that reflects your reliability, availability, serviceability (RAS), and performance goals. The set of possible profiles presented depends on your available storage. The following table lists all possible profiles and their descriptions.

TABLE 24 Data Profiles

Data Profile	Description
Dual Parity Options	
Triple mirrored	Data is triply mirrored, yielding a very highly reliable and high-performing system (for example, storage for a critical database). This configuration is intended for situations in which maximum performance and availability are required. Compared with a two-way mirror, a three-way mirror adds additional IOPS per stored block and higher level protection against failures. Note: A controller without expansion storage should not be configured with triple mirroring.
Double parity RAID	RAID in which each stripe contains two parity disks. As with triple mirroring, this yields high availability, as data remains available with the failure of any two disks. Double parity RAID is a higher capacity option than the mirroring options and is intended either for high-throughput sequential-access workloads (such as backup) or for storing large amounts of data with low random-read component.
Single Parity Options	
Mirrored	Data is mirrored, reducing capacity by half, but yielding a highly reliable and high-performing system. Recommended when space is considered ample, but performance is at a premium (for example, database storage).
Single parity RAID, narrow stripes	RAID in which each stripe is kept to three data disks and a single parity disk. For situations in which single parity protection is acceptable, single parity RAID offers a much higher capacity option than simple mirroring. This higher capacity needs to be balanced against a lower random read capability than mirrored options. Single parity RAID can be considered for non-critical applications with a moderate random read component. For pure streaming workloads, give preference to the Double parity RAID option which has higher capacity and more throughput.
Other	
Striped	Data is striped across disks, with no redundancy. While this maximizes both performance and capacity, a single

Data Profile	Description
	disk failure will result in data loss. This configuration is not recommended. For pure streaming workloads, consider using Double parity RAID. Due to non-redundancy, disks configured in a striped profile will not receive firmware updates, unless the configured storage pools are in an exported state.
Triple parity RAID, wide stripes	RAID in which each stripe has three disks for parity. This is the highest capacity option apart from Striped Data. Resilvering data after one or more drive failures can take significantly longer due to the wide stripes and low random I/O performance. As with other RAID configurations, the presence of cache can mitigate the effects on read performance. This configuration is not generally recommended.

Note - Earlier software versions supported double parity with wide stripes. This has been supplanted by triple parity with wide stripes, as it adds significantly better reliability. Pools configured as double parity with wide stripes under a previous software version continue to be supported, but newly-configured or reconfigured pools cannot select that option.

NSPF Option

For expandable systems, some profiles may be available with an 'NSPF' option. This stands for 'no single point of failure' and indicates that data is arranged in mirrors or RAID stripes such that a pathological disk shelf failure will not result in data loss. Note that systems are already configured with redundancy across nearly all components. Each disk shelf has redundant paths, redundant controllers, and redundant power supplies and fans. The only failure that NSPF protects against is disk backplane failure (a mostly passive component), or gross administrative misconduct (detaching both paths to one disk shelf). In general, adopting NSPF will result in lower capacity, as it has more stringent requirements on stripe width.

Log Devices

Log devices can be configured using only striped or mirrored profiles. Log devices are only used in the event of node failure. For data to be lost with unmirrored logs, it is necessary for both the device to fail and the node to reboot immediately after. This a highly-unlikely event, however mirroring log devices can make this effectively impossible, requiring two simultaneous device failures and node failure within a very small time window.

Note - When different sized log devices are in different chassis, only striped log profiles can be created.

Cache Devices

In a cluster configuration, cache devices installed in controller slots are available only to the controller which has the storage pool imported. In a cluster, it is possible to configure cache devices on both controllers to be part of the same pool. To do this, take over the pool on the passive node, then add storage, and select the cache devices. This has the effect of having half the global cache devices configured at any one time. While the data on the cache devices will be lost on failover, the new cache devices can be used on the new controller.

Cache devices installed in disk shelf slots, when added to a pool, are automatically imported during a cluster failback or takeover. No additional pool configuration is required.

Meta Devices

A meta device is a cache device used to store deduplicated metadata and other metadata for projects and shares. Meta devices can be allocated to a storage pool, but not an all-flash storage pool, during and after storage pool creation. However, they cannot be re-configured as normal cache devices for a pool, nor can they be removed from a pool. A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 (2013.1.7.0) or later.

Before using meta devices and the deduplication feature for new and existing storage pools, accept the deferred software update for Data Deduplication v2, introduced with software version OS8.7.0 (2013.1.7.0). If replicating to other systems, both the replication source and targets must have this deferred update. For more information, see [Data Deduplication](#), and “Data Deduplication v2 Deferred Update” in *Oracle ZFS Storage Appliance Customer Service Manual*.

Hot Spares

Hot spares are allocated as a percentage of total pool size and are independent of the profile chosen (with the exception of striped, which doesn't support hot spares). Because hot spares are allocated for each storage configuration step, it is much more efficient to configure storage as a whole than it is to add storage in small increments.

Related Topics:

- Creating a Storage Pool ([BUI](#), [CLI](#)).
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool ([BUI](#), [CLI](#)).


Understanding the Appliance Status

The Status section provides a summary of appliance status and configuration options. Use the following sections for conceptual and procedural information about appliance status views and related service configuration:

- [“About Oracle ZFS Storage Appliance” on page 19](#)
- [“Status Dashboard” on page 155](#)
- [“Summary of Pool Usage” on page 161](#)
- [“Summary of Memory Usage” on page 161](#)
- [“Disk Activity Dashboard” on page 162](#)
- [“Dashboard CLI” on page 164](#)
- [“Running the Dashboard Continuously” on page 165](#)
- [“Status Dashboard Settings” on page 166](#)
- [“Changing the Displayed Activity Statistics” on page 168](#)
- [“Changing the Activity Thresholds” on page 168](#)
- [“NDMP Status” on page 169](#)
- [“NDMP States” on page 170](#)

Status Dashboard


The dashboard summarizes appliance status.



The status dashboard provides links to all main screens of the browser user interface (BUI). Over 100 visible items on the dashboard link to associated BUI screens indicated by a border or highlighted text that appears when hovered over. The sections that follow describe the areas of the dashboard in detail.

Usage Dashboard

The Usage area of the dashboard provides a summary of your storage pool and main memory usage. The name of the pool appears at the top right of the Usage area. If multiple pools are configured, use the pull-down list to select the pool you want to display.


FIGURE 10 Status Dashboard: Usage

The total pool capacity is displayed to the right of the storage usage pie chart. The storage pie chart details the used and available space. To go to the Shares screen for the pool, click the storage pie chart.

The total system physical memory is displayed to the right of the memory pie chart. To the left is a pie chart showing memory usage by component. To go to the Analytics worksheet for dynamic memory usage broken down by application name, click the Memory pie chart.

Services Dashboard

The services area of the dashboard shows the status of services on the appliance, with a light icon to show the state of each service.

FIGURE 11 Services Dashboard


Most services are green to indicate that the service is online, or grey to indicate that the service is disabled. For a reference of all possible states and icon colors, see [“Browser User Interface \(BUI\)” on page 22](#).

To go to the associated configuration screen, click on a service name. The Properties screen appears with configurable fields, restart, enable, and disable icons, and a link to the associated Logs screen for the service.

Hardware Dashboard

The Hardware area of the dashboard shows an overview of hardware on the appliance.

FIGURE 12 Hardware Dashboard




If there is a known fault, the amber fault  icon appears.

To go to the Hardware screen for a detailed look at hardware state, click the name of a hardware component.

Activity Dashboard

The activity area of the dashboard shows graphs of eight performance statistics by default. The example in this section shows Disk operations/second. The statistical average is plotted in blue and the maximum appears in light grey.

FIGURE 13 Disk Activity Dashboard



To go to the Analytics worksheet for an activity, click one of the four graphs (day, hour, minute, second) for the statistic you want to evaluate.

To view the average for each graph, mouse-over a graph and the average appears in the tooltip. The weather icon in the upper-left provides a report of activity according to thresholds you can customize for each statistic on the [“Status Dashboard Settings” on page 166](#) screen.

TABLE 25 Summary of Statistic Graphs

Summary of Statistic Graphs	
7-day graph (7d)	A bar chart, with each bar representing one day.
24-hour graph (24h)	A bar chart, with each bar representing one hour.
60-minute graph (60m)	A line plot, representing activity over one hour (also visible as the first one-hour bar in the 24-hour graph).
1-second graph	A line plot, representing instantaneous activity reporting.

The average for the selected plot is shown numerically above the graph. To change the average that appears, select the average you want, either 7d, 24h, or 60m.

The vertical scale of all graphs is printed on the top right, and all graphs are scaled to this same height. The height is calculated from the selected graph (plus a margin). The height will rescale based on activity in the selected graph, with the exception of utilization graphs which have a fixed height of 100 percent.

Since the height can rescale, 60 minutes of idle activity may look similar to 60 minutes of busy activity. Always check the height of the graphs before trying to interpret what they mean.

Understanding some statistics may not be obvious - you might wonder, for a particular appliance in your environment, whether 1000 NFSv3 ops/sec is considered busy or idle. This is where the 24-hour and 7-day plots can help, to provide historic data next to the current activity for comparison.

The plot height is calculated from the selected plot. By default, the 60-minute plot is selected. So, the height is the maximum activity during that 60-minute interval (plus a margin). To rescale all plots to span the highest activity during the previous 7 days, select 7d. This makes it easy to see how current activity compares to the last day or week.

The weather icon is intended to grab your attention when something is unusually busy or idle. To go to the weather threshold configuration page, click the weather icon. There is no good or bad threshold, rather the BUI provides a gradient of levels for each activity statistic. The statistics on which weather icons are based provide an *approximate* understanding for appliance performance that you should customize to your workload, as follows:


- Different environments have different acceptable levels for performance (latency), and so there is no one-size-fits-all threshold.

- The statistics on the Dashboard are based on operations/sec and bytes/sec, so you should use Analytics worksheets for an accurate understanding of system performance.

Recent Alerts

This section shows the last four appliance alerts. Click the box to go to the Logs screen to examine all recent alerts in detail.

FIGURE 14 Recent Alerts



RECENT ALERTS

- 2010-2-22 16:53:51 Replication of 'default' to 'tuna' failed.
- 2010-2-22 16:29:23 Finished replicating 'default' to appliance 'tuna'.
- 2010-2-22 16:29 Began replicating 'default' to appliance 'tuna'.
- 2010-2-22 15:59:28 Finished replicating 'default' to appliance 'tuna'.

Summary of Pool Usage

The following table describes the pool usage properties.

TABLE 26 Summary of Pool Usage

Summary Pool Usage	
Used	Space used by this pool including data and snapshots.
Avail	Amount of remaining disk space available to the user. Refers to the amount of available space, excluding unused space that is reserved by projects and shares within a pool.
Compression	Current compression ratio achieved by this pool. Ratio will display 1x if compression is disabled.
Dedup	Current data deduplication ratio achieved by this pool. Ratio will display 1x if data deduplication is disabled.

Summary of Memory Usage

The following table describes the memory usage properties.


TABLE 27 Summary of Main Memory Usage

Summary of Main Memory (RAM) Usage	
Cache	Bytes in use by the filesystem cache to improve performance.
Unused	Bytes not currently in use. After booting, this value will decrease as space is used by the filesystem cache.
Mgmt	Bytes in use by the appliance management software.
Other	Bytes in use by miscellaneous operating system software.
Kernel	Bytes in use by the operating system kernel.

Note that users need the `analytics/component create+read` authorization to view the memory usage. Without this authorization, the memory details do not appear on the dashboard.

Disk Activity Dashboard

The activity area of the dashboard shows graphs of eight performance statistics by default. The example in this section shows Disk operations/second. The statistical average is plotted in blue and the maximum appears in light grey.

FIGURE 15 Disk Activity Dashboard

To go to the Analytics worksheet for an activity, click one of the four graphs (day, hour, minute, second) for the statistic you want to evaluate.

To view the average for each graph, mouse-over a graph and the average appears in the tooltip. The weather icon in the upper-left provides a report of activity according to thresholds you can customize for each statistic on the Status Settings screen.

TABLE 28 Summary of Statistic Graphs

Summary of Statistic Graphs	
7-day graph (7d)	A bar chart, with each bar representing one day.
24-hour graph (24h)	A bar chart, with each bar representing one hour.
60-minute graph (60m)	A line plot, representing activity over one hour (also visible as the first one-hour bar in the 24-hour graph).
1-second graph	A line plot, representing instantaneous activity reporting.

The average for the selected plot is shown numerically above the graph. To change the average that appears, select the average you want, either 7d, 24h, or 60m.

The vertical scale of all graphs is printed on the top right, and all graphs are scaled to this same height. The height is calculated from the selected graph (plus a margin). The height will rescale based on activity in the selected graph, with the exception of utilization graphs which have a fixed height of 100 percent.

Since the height can rescale, 60 minutes of idle activity may look similar to 60 minutes of busy activity. Always check the height of the graphs before trying to interpret what they mean.

Understanding some statistics may not be obvious - you might wonder, for a particular appliance in your environment, whether 1000 NFSv3 ops/sec is considered busy or idle. This is where the 24-hour and 7-day plots can help, to provide historic data next to the current activity for comparison.

The plot height is calculated from the selected plot. By default, the 60-minute plot is selected. So, the height is the maximum activity during that 60-minute interval (plus a margin). To rescale all plots to span the highest activity during the previous 7 days, select 7d. This makes it easy to see how current activity compares to the last day or week.

The weather icon is intended to grab your attention when something is unusually busy or idle. To go to the weather threshold configuration page, click the weather icon. There is no good or bad threshold, rather the BUI provides a gradient of levels for each activity statistic. The statistics on which weather icons are based provide an *approximate* understanding for appliance performance that you should customize to your workload, as follows:

- Different environments have different acceptable levels for performance (latency), and so there is no one-size-fits-all threshold.
- The statistics on the Dashboard are based on operations/sec and bytes/sec, so you should use Analytics worksheets for an accurate understanding of system performance.

Dashboard CLI

A text version of the Status > Dashboard screen is available from the CLI by typing `status dashboard`:

```
hostname:> status dashboard
```

Data:

```
pool_0:
  Used          497G bytes
  Avail         8.43T bytes
  State         online
  Compression   1x
```

Memory:

```
Cache          30.1G bytes
Unused         2.18G bytes
Mgmt           343M bytes
Other          474M bytes
Kernel         38.9G bytes
```

Services:

ad	disabled	smb	disabled
dns	online	ftp	disabled
http	online	identity	online
idmap	online	ipmp	online
iscsi	online	ldap	disabled
ndmp	online	nfs	online
nis	online	ntp	online
routing	online	scrk	maintenance
snmp	online	ssh	online
tags	online	vscan	online

Hardware:

CPU	online	Cards	online
Disks	faulted	Fans	online
Memory	online	PSU	online

Activity:

CPU	1 %util	Sunny
Disk	32 ops/sec	Sunny
iSCSI	0 ops/sec	Sunny
NDMP	0 bytes/sec	Sunny
NFSv3	0 ops/sec	Sunny
NFSv4	0 ops/sec	Sunny
Network	13K bytes/sec	Sunny
SMB	0 ops/sec	Sunny

Recent Alerts:

2013-6-15 07:46: A cluster interconnect link has been restored.

The previous descriptions in this section apply, with the following differences:

- The activity plots are not rendered in text (although we have thought about using aalib).
- The storage usage section will list details for all available pools in the CLI, whereas the BUI only has room to summarize one.

Separate views are available, for example `status activity show`:

```
hostname:> status activity show
Activity:
  CPU          10 %util          Sunny
  Disk         478 ops/sec       Partly Cloudy
  iSCSI         0 ops/sec        Sunny
  NDMP          0 bytes/sec       Sunny
  NFSv3        681 ops/sec       Partly Cloudy
  NFSv4         0 ops/sec        Sunny
  Network     22.8M bytes/sec    Partly Cloudy
  SMB           0 ops/sec        Sunny
hostname:>
```

▼ Running the Dashboard Continuously

You might experience browser memory issues if you leave the Dashboard screen open in a browser continuously (24x7). The browser will increase in size (memory leaks), and need to be closed and reopened. Browsers are fairly good at managing memory when browsing through different websites (and opening and closing tabs). The issue is that the Dashboard screen is left running and not closed, which opens and reopens images for the activity plots, thus degrading image rendering performance.

If you experience this problem while using Firefox, disable the memory cache as follows:

1. **Open `about:config`**
2. **Filter on "memory"**
3. **Set `browser.cache.memory.enable = false`.**

Status Dashboard Settings

The Status > Settings screen enables you to customize the Status Dashboard, including the statistics that appear and thresholds that indicate activity through the weather icons.

Use the layout tab to select the graphs that appear in the dashboard activity area, as defined in the following table.

TABLE 29 Status Layout Settings

Name	Units	Description
<empty>	-	No graph will be displayed in this location.
CPU	utilization	Average cycles the appliance CPUs are busy. CPU cycles includes memory wait cycles.
ARC Ratio	utilization	Average ARC hit/miss percentage. A drop in the hit rate indicates a potential performance problem.
HTTP	operations/sec	Average number of HTTP operations.
Disk	operations/sec	Average number of operations to the physical storage devices.
iSCSI	operations/sec	Average number of iSCSI operations.
FC	operations/sec	Average number of Fibre Channel operations.
NDMP	bytes/sec	Average NDMP network bytes.
NFSv2	operations/sec	Average number of NFSv2 operations.
NFSv3	operations/sec	Average number of NFSv3 operations.

Name	Units	Description
NFSv4.0	operations/sec	Average number of NFSv4.0 operations.
NFSv4.1	operations/sec	Average number of NFSv4.1 operations.
Network	bytes/sec	Average bytes/sec across all physical network interfaces.
SMB	operations/sec	Average number of SMB operations.
SMB2	operations/sec	Average number of SMB2 operations.
SMB3	operations/sec	Average number of SMB3 operations.
FTP	bytes/sec	Average number of FTP bytes.
SFTP	bytes/sec	Average number of SFTP bytes.

Note that to reduce the network traffic required to refresh the Dashboard, configure some of the activity graphs as "<empty>".

Use the Thresholds screen to configure the dashboard activity weather icons. The defaults provided are based on heavy workloads and may not be suitable for your environment.

FIGURE 16 Dashboard Activity Settings

Configure Dashboard Layout **Thresholds**

REVERT **APPLY**

About
Configure the thresholds used for the activity icons on the status dashboard. A threshold is crossed when that statistic has averaged over the threshold for 60 seconds. Weather icons have been used to express activity level, a metaphor suitable considering the nature of these high level summaries.

Dashboard statistic: **CPU**

Use threshold settings: ☐ Custom ☒ **Default**

Sunny		0	%util
Partly Cloudy		20	%util
Cloudy		40	%util
Rainy		60	%util
Stormy		85	%util
Hurricane		92	%util
Cat-2 Hurricane		95	%util
Cat-3 Hurricane		97	%util
Cat-4 Hurricane		98	%util
Cat-5 Hurricane		99	%util


The weather icon that appears on the Dashboard is closest to the threshold value setting for the current activity - measured as a 60 second average. For example, if CPU utilization was at 41%,

by default, the Cloudy weather icon would appear because its threshold is 40% (closest to the actual activity). Select the Custom radio button to configure thresholds and be sure to configure them in the order they appear on the screen.

The dashboard currently cannot be configured from the CLI. Settings saved in the BUI will apply to the dashboard that is visible from the CLI.

▼ Changing the Displayed Activity Statistics

1. Go to the **Status > Settings > Layout** screen.
2. Choose the statistics you want to display on the Dashboard from the drop-down menus.
3. To save your choices, click the **Apply** button.



▼ Changing the Activity Thresholds

1. Go to the **Status > Settings > Thresholds** screen.
2. Choose the statistic to configure from the drop-down menu.
3. Click the **Custom** radio button.
4. Customize the values in the list, in the order they appear. Some statistics will provide a **Units** drop-down, so that Kilo/Mega/Giga can be selected.

5. To save your configuration, click the **Apply** button.

NDMP Status

When the NDMP service has been configured and is active, the Status=>NDMP page shows the NDMP devices and recent client activity. A green indicator shows that the device is online and a gray indicator shows that the device is offline.

To resort the NDMP Device list, click on the Devices column headings. To display details about a device, double click on the device.

NDMP status is not available from the CLI.

FIGURE 17 NDMP Status BUI

The screenshot shows the NDMP Status BUI interface. At the top, there are tabs for Configuration, Maintenance, Shares, Status (selected), and Analytics. Below the tabs, there are sub-tabs for DASHBOARD, SETTINGS, and NDMP. The main content area is titled 'Devices' and shows a list of 20 total devices. The list has columns for TYPE, PATH, VENDOR, MODEL, WWN, and SERIAL. All devices are 'Tape drive' type. The list includes devices from STK, HP, and IBM vendors. Below the device list, there is a 'Recent activity' section showing 0 total sessions and a message 'No active NDMP sessions.'

TYPE	PATH	VENDOR	MODEL	WWN	SERIAL
Tape drive	/dev/rmt/14bn	STK	T1000C	50 01 04 F0 00 AC BB 27	576001000203
Tape drive	/dev/rmt/15bn	HP	Ultrium 5-SCSI	50 01 04 F0 00 79 1C 45	HU1036C6N4
Tape drive	/dev/rmt/16bn	HP	Ultrium 5-SCSI	50 01 04 F0 00 79 1C 4E	HU1036C6N2
Tape drive	/dev/rmt/17bn	HP	Ultrium 5-SCSI	50 01 04 F0 00 79 1C 51	HU1036C6MV
Tape drive	/dev/rmt/20bn	IBM	ULTRIUM-TD5	50 01 04 F0 00 79 1C 2A	1068010909
Tape drive	/dev/rmt/19bn	IBM	ULTRIUM-TD5	50 01 04 F0 00 79 1C 2D	1068011364
Tape drive	/dev/rmt/18bn	HP	Ultrium 5-SCSI	50 01 04 F0 00 AC BB 2A	HU1233PP32
Tape drive	/dev/rmt/21bn	HP	Ultrium 5-SCSI	50 01 04 F0 00 79 1C 42	HU1036C6N0
Tape drive	/dev/rmt/23bn	IBM	ULTRIUM-TD6	50 01 04 F0 00 AC BB 2D	1068000648
Tape drive	/dev/rmt/22bn	IBM	ULTRIUM-TD5	50 01 04 F0 00 AC BB 30	1068025702

TABLE 30 NDMP Status - Devices

Field	Description	Examples
Type	Type of NDMP device	Robot, Tape drive
Path	Path of the NDMP device	/dev/rmt/14bn
Vendor	Device vendor name	STK
Model	Device model name	T1000C
WWN	World Wide Name	50:01:04:F0:00:AC:BB:27

Field	Description	Examples
Serial	Device serial number	576001000203

TABLE 31 NDMP Status - Recent Activity

Field	Description	Examples
ID	NDMP backup ID	49
Active	Backup currently active	No
Remote Client	NDMP client address and port	192.168.1.219:4760
Authenticated	Shows if the client has completed authentication yet	Yes, No
Data State	See Data State	Active, Idle, ...
Mover State	See Mover State	Active, Idle, ...
Current Operation	Current NDMP operation	Backup, Restore, None
Progress	A progress bar for this backup	

NDMP States

The NDMP Data State shows the state of the backup or restore operation. Possible values are:

- **Active** - The data is being backed up or restored.
- **Idle** - Backup or restore has not yet started or has already finished.
- **Connected** - Connection is established, but backup or restore has not yet begun.
- **Halted** - Backup or restore has finished successfully or has failed or aborted.
- **Listen** - Operation is waiting to receive a remote connection.

The NDMP Mover State shows the state of the NDMP device subsystem. Examples for tape devices are:

- **Active** - Data is being read from or written to the tape.
- **Idle** - Tape operation has not yet started or has already finished.
- **Paused** - Tape has reached the end or is waiting to be changed.
- **Halted** - Read/write operation has finished successfully or has failed or aborted.
- **Listen** - Operation is waiting to receive a remote connection.

Configuring Storage Area Network (SAN)

The SAN configuration page lets you connect your appliance to your Storage Area Network (SAN). A SAN is made up of three basic components:

- A client that will access the storage on the network
- A storage appliance that will provide the storage on the network
- A network that will link the client to the storage

To configure SAN, use the following sections:

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)
- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)


To learn more about SAN, see the following:

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ **Configuring FC Port Modes (BUI)**

1. To use FC ports, set them to Target mode on the Configuration > SAN screen of the BUI, using the drop-down menu shown in the following image. You

must have root permissions to perform this action. Note that in a cluster configuration, you set ports to Target mode on each server separately.




2. After setting desired ports to Target, click the Apply button. A confirmation message will appear notifying you that the appliance will reboot immediately. Confirm that you want to reboot.
3. When the appliance boots, the active FC targets appear with the  icon and, on mouse-over, the move  icon appears.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Discovering FC Ports (BUI)

1. Click the info  icon to view the Discovered Ports dialog where you can troubleshoot link errors. In the Discovered Ports dialog, click a WWN in the list to view associated link errors.




2. In the Discovered Ports dialog, click a WWN in the list to view associated link errors.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Creating FC Initiator Groups (BUI)


1. Create and manage initiator groups on the Initiators screen. Click the add  icon to view unaliased ports. Click a WWN in the list to add a meaningful alias in the Alias field.



The screenshot shows a window titled "Initiators" with "CANCEL" and "OK" buttons at the top right. Inside, there is a "World Wide Name" field containing "21:00:00:1b:32:81:e3:39" and an empty "Alias" field below it. Below these fields is a section titled "Unaliased ports" with a subtitle "2 Total". This section contains a table with two columns: "WWN ▲" and "VENDOR".

WWN ▲	VENDOR
21:00:00:1b:32:81:e3:39	QLogic Corporation
21:01:00:1b:32:a1:e3:39	QLogic Corporation

2. On the Initiators page, drag initiators to the FC Initiator Groups list to create new groups or add to existing groups.




3. Click the Apply button to commit the new Initiator Group. Now you can create a LUN that has exclusive access to the client initiator group.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Associating a LUN with an FC Initiator Group (BUI)

1. To create the LUN, roll-over the initiator group and click the add LUN  icon. The Create LUN dialog appears with the associated initiator group selected.



The image shows a 'Create LUN' dialog box with the following fields and options:

- Project:** default
- Name:** (empty text field)
- Volume size:** 0 GB
- Thin provisioned:** ☐
- Volume block size:** 8k
- Online:** ☒
- Target group:** All targets
- Initiator group(s):** A dropdown menu showing 'All initiators', 'initiators-0', and 'initiators-1'.
- LU number:** ☐ 0 ☒ Auto-assign

2. Set the name and size and click Apply to add the LUN to the storage pool.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Changing FC Port Modes (CLI)

- To change FC port modes, use the following CLI commands:

```
dory:configuration san fc targets> set targets="wwn.2101001B32A11639"
      targets = wwn.2101001B32A11639 (uncommitted)
dory:configuration san fc targets> commit
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Discovering FC Ports (CLI)

- To discover FC Ports, use the following CLI commands:

```
dory:configuration san fc targets> show
Properties:
      targets = wwn.2100001B32811639,wwn.2101001B32A12239

Targets:
NAME      MODE      WWN              PORT              SPEED
target-000 target    wwn.2100001B32811639  PCIe 5: Port 1    4 Gbit/s
target-001 initiator wwn.2101001B32A11639  PCIe 5: Port 2    0 Gbit/s
target-002 initiator wwn.2100001B32812239  PCIe 2: Port 1    0 Gbit/s
target-003 target    wwn.2101001B32A12239  PCIe 2: Port 2    0 Gbit/s
dory:configuration san fc targets> select target-000
dory:configuration san fc targets target-000> show
Properties:
      wwn = wwn.2100001B32811639
      port = PCIe 5: Port 1
      mode = target
      speed = 4 Gbit/s
      discovered_ports = 6
      link_failure_count = 0
      loss_of_sync_count = 0
      loss_of_signal_count = 0
      protocol_error_count = 0
```

```

invalid_tx_word_count = 0
invalid_crc_count = 0

Ports:
PORT      WWN                      ALIAS          MANUFACTURER
port-000   wwn.2100001B3281A339  longjaw-1      QLogic Corporation
port-001   wwn.2101001B32A1A339  longjaw-2      QLogic Corporation
port-002   wwn.2100001B3281AC39  thicktail-1    QLogic Corporation
port-003   wwn.2101001B32A1AC39  thicktail-2    QLogic Corporation
port-004   wwn.2100001B3281E339  <none>         QLogic Corporation
port-005   wwn.2101001B32A1E339  <none>         QLogic Corporation

```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Creating FC Initiator Groups (CLI)

- To create FC initiator groups, use the following CLI commands:

```

dory:configuration san fc initiators> create
dory:configuration san fc initiators (uncommitted)> set name=lefteye
dory:configuration san fc initiators (uncommitted)>
    set initiators=wwn.2101001B32A1AC39,wwn.2100001B3281AC39
dory:configuration san fc initiators (uncommitted)> commit
dory:configuration san fc initiators> list
GROUP      NAME
group-001  lefteye
|
+--> INITIATORS
      wwn.2101001B32A1AC39
      wwn.2100001B3281AC39

```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)

- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Associating a LUN with an FC Initiator Group (CLI)

The following example demonstrates creating a LUN called `lefty` and associating it with the `fera` initiator group.

- To associate a LUN with an FC initiator group, use the following CLI commands:

```
dory:shares default> lun lefty
dory:shares default/lefty (uncommitted)> set volsize=10
                                volsize = 10 (uncommitted)
dory:shares default/lefty (uncommitted)> set initiatorgroup=fera
                                initiatorgroup = default (uncommitted)
dory:shares default/lefty (uncommitted)> commit
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Scripting Aliases for Initiators and Initiator Groups (CLI)

Refer to [CLI Usage](#) and [Simple CLI Scripting and Batching Commands](#) for information about how to modify and use the following example script.

- To script aliases for initiators and initiator groups, use the following CLI commands:

```
script
```

```

/*
 * This script creates both aliases for initiators and initiator
 * groups, as specified by the below data structure. In this
 * particular example, there are five initiator groups, each of
 * which is associated with a single host (thicktail, longjaw, etc.),
 * and each initiator group consists of two initiators, each of which
 * is associated with one of the two ports on the FC HBA. (Note that
 * there is nothing in the code that uses this data structure that
 * assumes the number of initiators per group.)
 */
groups = {
  thicktail: {
    'thicktail-1': 'wwn.2100001b3281ac39',
    'thicktail-2': 'wwn.2101001b32a1ac39'
  },
  longjaw: {
    'longjaw-1': 'wwn.2100001b3281a339',
    'longjaw-2': 'wwn.2101001b32a1a339'
  },
  tecopa: {
    'tecopa-1': 'wwn.2100001b3281e339',
    'tecopa-2': 'wwn.2101001b32a1e339'
  },
  spinedace: {
    'spinedace-1': 'wwn.2100001b3281df39',
    'spinedace-2': 'wwn.2101001b32a1df39'
  },
  fera: {
    'fera-1': 'wwn.2100001b32817939',
    'fera-2': 'wwn.2101001b32a17939'
  }
};

for (group in groups) {
  initiators = [];
  for (initiator in groups[group]) {
    printf('Adding %s for %s ... ',
      groups[group][initiator], initiator);
    try {
      run('select alias=' + initiator);
      printf('(already exists)\n');
      run('cd ..');
    } catch (err) {
      if (err.code != EAKSH_ENTITY_BADSELECT)
        throw err;
      run('create');
      set('alias', initiator);
      set('initiator', groups[group][initiator]);
      run('commit');
    }
  }
}

```



```

        printf('done\n');
    }
    run('select alias=' + initiator);
    initiators.push(get('initiator'));
    run('cd ..');
}
printf('Creating group for %s ... ', group);
run('groups');
try {
    run('select name=' + group);
    printf('(already exists)\n');
    run('cd ..');
} catch (err) {
    if (err.code != EAKSH_ENTITY_BADSELECT)
        throw err;
    run('create');
    set('name', group);
    run('set initiators=' + initiators);
    run('commit');
    printf('done\n');
}
run('cd ..');
}

```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Creating an Analytics Worksheet (BUI)

To create an analytics worksheet for observing operations by initiator, complete the following:

1. **Go to the Analytics screen.**
2. **Click the  add icon for Add Statistic. A menu of all statistics appears.**
3. **Select iSCSI operations > Broken down by initiator under the Protocols section of the menu. A graph of the current operations by initiator appears.**

4. To observe more detailed analytics, select the initiator from the field to the left of the graph and click the  icon. A menu of detailed analytics appears.


Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Configuring SAN iSER Targets

In the BUI, iSER targets are managed as iSCSI targets on the Configuration > SAN screen.

1. To configure ibp(x) interfaces, select the ibp(x) interface (or ipmp) you want, and drag it to the Datalinks list to create the datalink on the Configuration > Network screen.
2. Drag the Datalink to the Interfaces list to create a new interface.




3. To create an iSER target, on the Configuration > SAN page, click the iSCSI Targets link.
4. To add a new iSER target with an alias, click the  add icon.
5. To create a target group, drag the target you just created to the iSCSI Target Group list.



6. To create an initiator, click the Initiator link and then click the iSCSI initiators link.
7. To add a new initiator, click the  add icon.
8. Enter the Initiator IQN and an alias and click OK. Creating an initiator group is optional but if you don't create a group, the LUN associated with the target will be available to all initiators.
9. To create a group, drag the initiator to the iSCSI Initiator Groups list.



10. To create a LUN, on the Shares page, click LUN.
11. Click the  add icon and associate the new LUN with target or initiator groups you created already using the Target Group and Initiator Groups menu.



The image shows a 'Create LUN' dialog box with the following fields and options:

- Project:** A dropdown menu showing 'default'.
- Name:** An empty text input field.
- Volume size:** A text input field showing '0' followed by a 'G' unit selector.
- Thin provisioned:** An unchecked checkbox.
- Volume block size:** A dropdown menu showing '8k'.
- Online:** A checked checkbox.
- Target group:** A dropdown menu showing 'All targets'.
- Initiator group(s):** A dropdown menu showing 'All initiators', 'initiators-0', and 'initiators-1'.
- LU number:** Two radio buttons, '0' and 'Auto-assign', with 'Auto-assign' selected.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Adding an iSCSI Target with an Auto-generated IQN (CLI)

- To add an iSCSI target with an auto-generated IQN, use the following CLI commands:

```

ahi:configuration san iscsi targets> create
ahi:configuration san iscsi targets target (uncommitted)> set alias="Target 0"
ahi:configuration san iscsi targets target (uncommitted)> set auth=none
ahi:configuration san iscsi targets target (uncommitted)> set interfaces=igb1
ahi:configuration san iscsi targets target (uncommitted)> commit
ahi:configuration san iscsi targets> list
TARGET      ALIAS
target-000 Target 0
|
+--> IQN
      iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416

```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)

- To add an iSCSI target with a specific IQN and RADIUS authentication, use the following CLI commands:

```

ahi:configuration san iscsi targets> create
ahi:configuration san iscsi targets target (uncommitted)> set alias="Target 1"
ahi:configuration san iscsi targets target (uncommitted)>
    set iqn=iqn.2001-02.com.acme:12345
ahi:configuration san iscsi targets target (uncommitted)> set auth=radius
ahi:configuration san iscsi targets target (uncommitted)> set interfaces=igb1

```

```
ahi:configuration san iscsi targets target (uncommitted)> commit
ahi:configuration san iscsi targets> list
TARGET      ALIAS
target-000 Target 0
      |
      +--> IQN
              iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416
target-001 Target 1
      |
      +--> IQN
              iqn.2001-02.com.acme:12345
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Adding an iSCSI Initiator with CHAP Authentication (CLI)

- To add an iSCSI initiator with CHAP authentication, use the following CLI commands:

```
ahi:configuration san iscsi initiators> create
ahi:configuration san iscsi initiators initiator (uncommitted)>
    set initiator=iqn.2001-02.com.acme:initiator12345
ahi:configuration san iscsi initiators initiator (uncommitted)> set alias="Init 0"
ahi:configuration san iscsi initiators initiator (uncommitted)>
    set chapuser=thisismychapuser
ahi:configuration san iscsi initiators initiator (uncommitted)>
    set chapsecret=123456789012abc
ahi:configuration san iscsi initiators initiator (uncommitted)> commit
ahi:configuration san iscsi initiators> list
NAME      ALIAS
initiator-000 Init 0
      |
      +--> INITIATOR
              iqn.2001-02.com.acme:initiator12345
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Adding an iSCSI Target Group (CLI)

- To add an iSCSI target group, use the following CLI commands:

```

ahi:configuration san iscsi targets groups> create
ahi:configuration san iscsi targets group (uncommitted)> set name=tg0
ahi:configuration san iscsi targets group (uncommitted)>
    set targets=iqn.2001-02.com.acme:12345,
        iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416
ahi:configuration san iscsi targets group (uncommitted)> commit
ahi:configuration san iscsi targets groups> list
GROUP      NAME
group-000  tg0
          |
          +-> TARGETS
                iqn.2001-02.com.acme:12345
                iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416

```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Adding an iSCSI Initiator Group (CLI)

- To add an iSCSI initiator group, use the following CLI commands:

```
ahi:configuration san iscsi initiators groups> create
ahi:configuration san iscsi initiators group (uncommitted)> set name=ig0
ahi:configuration san iscsi initiators group (uncommitted)>
    set initiators=iqn.2001-02.com.acme:initiator12345
ahi:configuration san iscsi initiators group (uncommitted)> commit
ahi:configuration san iscsi initiators groups> list
GROUP      NAME
group-000  ig0
          |
          +--> INITIATORS
                iqn.2001-02.com.acme:initiator12345
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Configuring SRP Target (BUI)

This procedure describes the steps for configuring SRP targets.

1. **Connect HCA ports to IB interfaces.**
2. **The targets are automatically discovered by the appliance.**
3. **To create the target group, go to the Configuration > SAN screen.**
4. **Click the Target link and then click SRP targets.**
5. **The SRP targets page appears.**
6. **To create the target group, use the  move icon to drag a target to the Target Groups list.**
7. **Click Apply.**

8. (Optional) To create an initiator and initiator group on the Initiator screen, click the  icon, collect GUID from initiator, assign it a name, and drag it to initiator group.
9. To create a LUN and associate it with the SRP target and initiators you created in the previous steps, go to the Shares screen.
10. Click the LUN link and then click the LUN  icon. Use the Target Group and Initiator Group menus on the Create LUN dialog to select the SRP groups to associate with the LUN.

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

▼ Configuring SRP Targets (CLI)

The following example demonstrates how to create an SRP target group named targetSRPgroup using the CLI configuration `san targets srp groups` context:

- To configure SRP targets, use the following CLI commands:

```
swallower:configuration san targets srp groups> create
swallower:configuration san targets srp group (uncommitted)> set name=targetSRPgroup
                           name = targetSRPgroup (uncommitted)
swallower:configuration san targets srp group (uncommitted)>
set targets=eui.0002C903000489A4
                           targets = eui.0002C903000489A4 (uncommitted)
swallower:configuration san targets srp group (uncommitted)> commit
swallower:configuration san targets srp groups> list
GROUP      NAME
group-000  targetSRPgroup
|
+--> TARGETS
      eui.0002C903000489A4
```

Example 1 Creating a LUN associated with the Target SRP Group using the CLI

The following example shows how to create a LUN and associate it with the targetSRPgroup using the CLI shares CLI context:

```
swallower:shares default> lun mylun
swallower:shares default/mylun (uncommitted)> set targetgroup=targetSRPgroup
        targetgroup = targetSRPgroup (uncommitted)
swallower:shares default/mylun (uncommitted)> set volsize=10
        volsize = 10 (uncommitted)
swallower:shares default/mylun (uncommitted)> commit
swallower:shares default> list
Filesystems:
NAME          SIZE    MOUNTPOINT
test          38K     /export/test
LUNs:
NAME          SIZE    GUID
mylun         10G     600144F0E9D19FFB00004B82DF490001
```

Related Topics

- [Understanding SAN](#)
- [SAN Fibre Channel Configuration](#)
- [SAN iSCSI Configuration](#)
- [SAN iSER Target Configuration](#)
- [SAN SRP Configuration](#)
- [SAN Terminology](#)

Understanding SAN

These three components remain the same regardless of which protocol is used on the network. In some cases, the network may even be a cable between the initiator and the target, but in most cases, there is some type of switching involved.

Targets and initiators are configured by protocol. Refer to the documentation on a particular protocol ([“SAN Fibre Channel Configuration” on page 192](#), [iSCSI](#) or [“SRP Configuration” on page 370](#)) for details.

Target and initiator groups define sets of targets and initiators that can be associated with LUNs. A LUN that is associated with a target group can only be seen via the targets in the group. If a LUN is not explicitly associated with a target group, it is in the *default target group* and will be accessible via all targets, regardless of protocol. Similarly, a LUN can only be seen by the


initiators in the group or groups to which it belongs. If a LUN is not explicitly associated with an initiator group, it is in the *default initiator group* and can be accessed by all initiators. While using the default initiator group can be useful for evaluation purposes, its use is discouraged since it may result in exposure of the LUN to unwanted or conflicting initiators.

To avoid possible LUN conflicts when an initiator belongs to multiple groups, configure initiators within all groups before associating groups with LUNs.

To configure targets, go to the Configuration > SAN BUI page, use Fibre Channel, iSCSI, and SRP to navigate, and then configure the Ports, Initiator, and Target Groups controls.

To associate a LUN, go to the Shares > Shares > Protocols page and then configure the Target Group and Initiator Group controls.

FIGURE 18 Associate a LUN



Use the configuration `san` context of the CLI to operate on targets and initiators by protocol type. Then, use the `shares` CLI context to create LUNs and associate them with target and initiator groups.

Related Topics

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)
- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)

SAN Fibre Channel Configuration

Fibre Channel (FC) is a gigabit-speed networking technology used nearly exclusively as a transport for SCSI. FC is one of several block protocols supported by the appliance; to share LUNs via FC, the appliance must be equipped with one or more optional FC cards.

By default, all FC ports are configured to be in target mode. If the appliance is used to connect to a tape SAN for backup, one or more ports must be configured in initiator mode. To configure a port for initiator mode, the appliance must be reset. Multiple ports can be configured for initiator mode simultaneously.

Each FC port is assigned a World Wide Name (WWN), and, as with other block protocols, FC targets may be grouped into SAN target and initiator groups, allowing port bandwidth to be

dedicated to specific LUNs or groups of LUNs. Once an FC port is configured as a target, the remotely discovered ports can be examined and verified.


Refer to the *Implementing Fibre Channel SAN Boot with Oracle ZFS Storage Appliance* white paper at <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/fc-sanboot-081412-pdf-1735984.pdf> for details on FC SAN boot solutions using the appliance.

In a cluster, initiators will have two paths (or sets of paths) to each LUN: one path (or set of paths) will be to the head that has imported the storage associated with the LUN; the other path (or set of paths) will be to that head's clustered peer. The first path (or set of paths) is *active*; the second path (or set of paths) is *standby*. In the event of a takeover, the active paths will become unavailable, and the standby paths will (after a short time) be transitioned to be active, after which I/O will continue. This approach to multipathing is known as asymmetric logical unit access (ALUA) and, when coupled with an ALUA-aware initiator, allows cluster takeover to be transparent to higher-level applications.

Initiators are identified by their WWN. As with other block protocols, aliases can be created for initiators. To aid in creating aliases for FC initiators, a WWN can be selected from the WWNs of discovered ports. Also as with other block protocols, initiators can be collected into groups. When a LUN is associated with a specific initiator group, the LUN will only be visible to initiators in the group. In most FC SANs, LUNs will always be associated with the initiator group that corresponds to the system(s) for which the LUN has been created.

The appliance is an ALUA-compliant array. Properly configuring an FC initiator in an ALUA environment requires an ALUA-aware driver and may require initiator-specific tuning. See "Oracle ZFS Storage Appliance: How to set up Client Multipathing" (Doc ID 1628999.1) for more information.

FC performance can be observed via Analytics, whereby one can breakdown operations or throughput by initiator, target, or LUN:

FIGURE 19 FC Performance

For operations, one can also breakdown by offset, latency, size and SCSI command, allowing one to understand not just the *what* but the *how* and *why* of FC operations.

The appliance has been designed to utilize a global set of resources to service LUNs on each head. It is therefore not generally necessary to restrict queue depths on clients as the FC ports in the appliance can handle a large number of concurrent requests. Even so, there exists the remote possibility that these queues can be overrun, resulting in SCSI transport errors. Such queue overruns are often associated with one or more of the following:

- Overloaded ports on the front end - too many hosts associated with one FC port and/or too many LUNs accessed through one FC port

- Degraded appliance operating modes, such as a cluster takeover in what is designed to be an active-active cluster configuration

While the possibility of queue overruns is remote, it can be eliminated entirely if one is willing to limit queue depth on a per-client basis. To determine a suitable queue depth limit, one should take the number of target ports multiplied by the maximum concurrent commands per port (2048) and divide the product by the number of LUNs provisioned. To accommodate degraded operating modes, one should sum the number of LUNs across cluster peers to determine the number of LUNs, but take as the number of target ports the minimum of the two cluster peers. For example, in an active-active 7420 dual-headed cluster with one head having 2 FC ports and 100 LUNs and the other head having 4 FC ports and 28 LUNs, one should take the pessimal maximum queue depth to be two ports times 2048 commands divided by 100 LUNs plus 28 LUNs -- or 32 commands per LUN.

Tuning the maximum queue depth is initiator specific, but on Oracle Solaris, this is achieved by adjusting the global variable `ssd_max_throttle`.

To troubleshoot link-level issues such as broken optics or a poorly seated cable, look at the error statistics for each FC port. If any number is either significantly non-zero or increasing, that may be an indicator that link-level issues have been encountered, and that link-level diagnostics should be performed.

Related Topics

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)
- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)

SAN iSCSI Configuration

Internet SCSI is one of several block protocols supported by the appliance for sharing SCSI based storage.

When using the iSCSI protocol, the target portal refers to the unique combination of an IP address and TCP port number by which an initiator can contact a target.

When using the iSCSI protocol, a target portal group is a collection of target portals. Target portal groups are managed transparently; each network interface has a corresponding target portal group with that interface's active addresses. Binding a target to an interface advertises that iSCSI target using the portal group associated with that interface.

Note - Multiple connections per session are not supported.

An IQN (iSCSI qualified name) is the unique identifier of a device in an iSCSI network. iSCSI uses the form `iqn.date.authority:uniqueid` for IQNs. For example, the appliance may use the IQN: `iqn.1986-03.com.sun:02:c7824a5b-f3ea-6038-c79d-ca443337d92c` to identify one of its iSCSI targets. This name shows that this is an iSCSI device built by a company registered in March of 1986. The naming authority is just the DNS name of the company reversed, in this case, "com.sun". Everything following is a unique ID that Oracle uses to identify the target.

TABLE 32 iSCSI Target Properties

Target Property	Description
Target IQN	The IQN for this target. The IQN can be manually specified or auto-generated.
Alias	A human-readable nickname for this target.
Authentication mode	One of None, CHAP, or RADIUS.
CHAP name	If CHAP authentication is used, the CHAP username.
CHAP secret	If CHAP authentication is used, the CHAP secret.
Network interfaces	The interfaces whose target portals are used to export this target.

In addition to those properties, the BUI indicates whether a target is online or offline:

TABLE 33 Target Status Icons

icon	description
	Target is online
	Target is offline

On clustered platforms, targets which have at least one active interface on that cluster node will be online. Take care when assigning interfaces to targets; a target may be configured to use portal groups on disjoint head nodes. In that situation, the target will be online on both heads yet will export different LUNs depending on the storage owned by each head node. As network interfaces migrate between cluster heads as part of takeover/failback or ownership changes, iSCSI targets will move online and offline as their respective network interfaces are imported and exported.

Targets which are bound to an IPMP interface will be advertised only via the addresses of that IPMP group. That target will not be reachable via that group's test addresses. Targets bound to interfaces built on top of a LACP aggregation will use the address of that aggregation. If a LACP aggregation is added to an IPMP group, a target can no longer use that aggregation's interface, as that address will become an IPMP test address.

Related Topics

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)
- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)

SAN iSCSI Initiator Configuration

iSCSI initiators have the following configurable properties.

TABLE 34 SAN iSCSI Initiator Properties

Property	Description
Initiator IQN	The IQN for this initiator.
Alias	A human-readable nickname for this initiator.
Use CHAP	Enables or disables CHAP authentication
CHAP name	If CHAP authentication is used, the CHAP username.
CHAP secret	If CHAP authentication is used, the CHAP secret.

When planning your iSCSI client configuration, you'll need the following information:

- What initiators (and their IQNs) will be accessing the SAN?
- If you plan on using CHAP authentication, what CHAP credentials does each initiator use?
- How many iSCSI disks (LUNs) are required, and how big should they be?
- Do the LUNs need to be shared between multiple initiators?

To allow the Appliance to perform CHAP authentication using RADIUS, the following pieces of information must match:

- The Appliance must specify the address of the RADIUS server and a secret to use when communicating with this RADIUS server
- The RADIUS server (e.g. in its clients file) must have an entry giving the address of this Appliance and specifying the same secret as above
- The RADIUS server (e.g. in its users file) must have an entry giving the CHAP name and matching CHAP secret of each initiator
- If the initiator uses its IQN name as its CHAP name (the recommended configuration) then the Appliance does not need a separate Initiator entry for each Initiator box -- the RADIUS server can perform all authentication steps.
- If the initiator uses a separate CHAP name, then the Appliance must have an Initiator entry for that initiator that specifies the mapping from IQN name to CHAP name. This Initiator entry does NOT need to specify the CHAP secret for the initiator.

For tips on troubleshooting common iSCSI misconfiguration, see [iSCSI](#).

iSCSI performance can be observed via Analytics, whereby one can breakdown operations or throughput by initiator, target, or LUN.

Related Topics

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)

- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)

SAN SRP Configuration

SCSI RDMA Protocol, is a protocol supported by the appliance for sharing SCSI based storage over a network that provides RDMA services (i.e. InfiniBand).

SRP ports are shared with other IB port services such as IPoIB and RDMA. The SRP service may only operate in target mode. SRP targets have the following configurable properties.

TABLE 35 SRP Target Properties

Property	Description
Target EUI	The Extended Unique Identifier (EUI) for this target. The EUI is automatically assigned by the system and is equal to the HCA GUID over which the SRP port service is running.
Alias	A human-readable nickname for this target.

In addition to those properties, the BUI indicates whether a target is online or offline:

TABLE 36 SRP Target Status Icons

icon	description
	Target is online
	Target is offline

On clustered platforms, peer targets should be configured into the same target group for highly available (multi-pathed) configurations. SRP multipathed I/O is an initiator-side configuration option.

SRP initiators have the following configurable properties.

TABLE 37 SRP Initiator Properties

Property	Description
Initiator EUI	The EUI for this initiator.
Alias	A human-readable nickname for this initiator.

SRP performance can be observed via Analytics, whereby one can breakdown operations or throughput by initiator or target.

Related Topics

- [Configuring FC Port Modes \(BUI\)](#)
- [Discovering FC Ports \(BUI\)](#)
- [Creating FC Initiator Groups \(BUI\)](#)
- [Associating a LUN with an FC Initiator Group \(BUI\)](#)
- [Changing FC Port Modes \(CLI\)](#)
- [Discovering FC Ports \(CLI\)](#)
- [Creating FC Initiator Groups \(CLI\)](#)
- [Associating a LUN with an FC Initiator Group \(CLI\)](#)
- [Scripting Aliases for Initiators and Initiator Groups \(CLI\)](#)
- [Configuring SAN iSCSI Initiators](#)
- [Creating an Analytics Worksheet \(BUI\)](#)
- [Adding an iSCSI Target with an Auto-generated IQN \(CLI\)](#)
- [Adding an iSCSI Target with a Specific IQN and RADIUS Authentication \(CLI\)](#)
- [Adding an iSCSI Initiator with CHAP Authentication \(CLI\)](#)
- [Adding an iSCSI Target Group \(CLI\)](#)
- [Adding an iSCSI Initiator Group \(CLI\)](#)
- [Configuring SRP Target \(BUI\)](#)
- [Configuring SRP Targets \(CLI\)](#)

SAN Terminology

To configure the appliance to operate on a SAN, you should understand some basic SAN terms:

TABLE 38 SAN Terminology

Term	Description
SCSI Target	A SCSI Target is a storage system end-point that provides a service of processing SCSI commands and I/O requests from an initiator. A SCSI Target is created by the storage system's administrator, and is identified by unique addressing methods. A SCSI Target, once configured, consists of zero or more logical units.
SCSI Initiator	A SCSI Initiator is an application or production system end-point that is capable of initiating a SCSI session, sending SCSI commands and I/O requests. SCSI Initiators are also identified by unique addressing methods (See SCSI Targets).
Logical Unit	A Logical Unit is a term used to describe a component in a storage system. Uniquely numbered, this creates what is referred to as a Logical Unit Number, or LUN. A storage system, being highly configurable, may contain many LUNS. These LUNS, when associated with one or more SCSI Targets, forms a unique SCSI device, a device that can be accessed by one or more SCSI Initiators.
iSCSI	Internet SCSI, a protocol for sharing SCSI based storage over IP networks. The appliance supports the SCSI-3 Persistent Reservations specification.
iSER	iSCSI Extension for RDMA, a protocol that maps the iSCSI protocol over a network that provides RDMA services (i.e. InfiniBand). The iSER protocol is transparently selected by the iSCSI subsystem, based on the presence of correctly configured IB hardware. In the CLI and BUI, all iSER-capable components (targets and initiators) are managed as iSCSI components.
FC	Fibre Channel, a protocol for sharing SCSI based storage over a storage area network (SAN), consisting of fiber-optic cables, FC switches and HBAs. The appliance supports 4GB and 8GB Fibre Channel Arbitrated Loop (FC-AL) topologies.
SRP	SCSI RDMA Protocol, a protocol for sharing SCSI based storage over a network that provides RDMA services (i.e. InfiniBand).
IQN	An iSCSI qualified name, the unique identifier of a device in an iSCSI network. iSCSI uses the form <code>iqn.date.authority:uniqueid</code> for IQNs. For example, the appliance may use the IQN: <code>iqn.1986-03.com.sun:02:c7824a5b-f3ea-6038-c79d-ca443337d92c</code> to identify one of its iSCSI targets. This name shows that this is an iSCSI device built by a company registered in March of 1986. The naming authority is just the DNS name of the company reversed, in this case, "com.sun". Everything following is a unique ID that Sun uses to identify the target.

Term	Description
Target portal	When using the iSCSI protocol, the target portal refers to the unique combination of an IP address and TCP port number by which an initiator can contact a target.
Target portal group	When using the iSCSI protocol, a target portal group is a collection of target portals. Target portal groups are managed transparently; each network interface has a corresponding target portal group with that interface's active addresses. Binding a target to an interface advertises that iSCSI target using the portal group associated with that interface.
CHAP	Challenge-handshake authentication protocol, a security protocol which can authenticate a target to an initiator, an initiator to a target, or both.
RADIUS	A system for using a centralized server to perform CHAP authentication on behalf of storage nodes.
Target group	A set of targets. LUNs are exported over all the targets in one specific target group.
Initiator group	A set of initiators. When an initiator group is associated with a LUN, only initiators from that group may access the LUN.
Target	A storage system end-point that provides a service of processing SCSI commands and I/O requests from an initiator. A target is created by the storage system administrator, and is identified by unique addressing methods. A target, once configured, consists of zero or more logical units.
Initiator	An application or production system end-point that is capable of initiating a SCSI session, sending SCSI commands and I/O requests. Initiators are also identified by unique addressing methods.

Each LUN has several properties which control how the volume is exported. See [Protocols](#) for more information.

Configuring Users

This section describes *users* for the appliance, *roles* to manage authorizations granted to users, and how to add them to the system using the BUI or CLI.

To configure users and roles, use the following sections:

- Adding an Administrator or User [BUI](#), [CLI](#)
- Changing a User Password [BUI](#), [CLI](#)

- Editing Exceptions for a User [BUI](#), [CLI](#)
- Deleting Exceptions for a User [BUI](#), [CLI](#)
- Adding a Role [BUI](#), [CLI](#)
- Editing Authorizations for a Role [BUI](#), [CLI](#)
- Deleting Authorizations from a Role [BUI](#), [CLI](#)
- Adding a User Who can Only View the Dashboard [BUI](#)

To understand users and roles, see the following sections:

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Adding an Administrator or User (BUI)

Use the following procedure to create a user with or without the administrator role. For a description of user types, see [“Understanding Users and Roles” on page 218](#).

1. **Go to Configuration > Users.**
2. **Click the add icon  next to Users.**

3. Choose the appropriate type of user from the drop-down menu.

Properties

This is an appliance administrator managed by a directory service.

Type

☒ Directory
 ☐ Local
 ☐ Data
 ☐ No-login

Username

User ID

Full Name

Password

Confirm

Require session annotation

☐

Kiosk user

☐

Kiosk screen

<https://ar7320-230:215/#status/dashboard>

Roles Exceptions

1 Total

NAME ▲	DESCRIPTION
☒ basic	Basic administration

4. Enter the required properties.

5. (Optional) To assign roles to Local and Directory users, click the checkboxes for the appropriate roles.

Newly created Local and Directory users default to the "basic" role.

6. (Optional) To add exceptions for Local and Directory users:

a. Click Exceptions.

b. Click the checkboxes for the exceptions you want to add.

c. Click ADD in the Exceptions section.

7. Click ADD at the top of the dialog box.

The new user appears in the Users list.

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Adding an Administrator or User (CLI)

Use the following procedure to create a user with or without the administrator roles. For a description of user types, see [“Understanding Users and Roles” on page 218](#).

1. **Go to configuration users.**

```
hostname:> configuration users
```

2. **Type one of the following user types followed by a name.**

directory- for a Directory user (NIS, LDAP).

local- for a Local user.

data- for a Data-only user.

nologin- for a no-login user.

3. **Type get to list the required properties that need to be set.**

```
hostname:configuration users john (uncommitted)> get
      logname = john
        uid = (unset)
      fullname = (unset)
  initial_password = (unset)
    require_annotation = false
```

4. **Type set and the property you want to set, and then type commit.**

```
hostname:configuration users john (uncommitted)> set initial_password=password
      initial_password = ***** (uncommitted)
hostname:configuration users john (uncommitted)> commit
```

At this point you have a created user, but haven't customized all their properties yet.

5. **(Optional) To add roles for Local or Directory users, type select and a username.**

6. (Optional) Type show to see the full list of preferences.

You can now add roles and authorization exceptions for the user.

Example 2 Creating a Local User

```
hostname:configuration users > local john
hostname:configuration users john (uncommitted) > get
    logname = john
    uid = (unset)
    fullname = (unset)
    initial_password = (unset)
    require_annotation = false
hostname:configuration users john (uncommitted) > set initial_password=password
    initial_password = ***** (uncommitted)
hostname:configuration users john (uncommitted) > commit
hostname:configuration users > select john
hostname:configuration users john > show
Properties:
    logname = john
    type = local
    uid =
    fullname =
    initial_password = *****
    require_annotation = false
    roles =
    kiosk_mode = false
    kiosk_screen = status/dashboard

Children:
    exceptions => Configure this user's exceptions
    preferences => Configure user preferences

hostname:configuration users john > set roles=
basic      basic2      test_role1 test_role2
hostname:configuration users john > set roles=basic
    roles = basic (uncommitted)
hostname:configuration users john > commit
hostname:configuration users > select john
hostname:configuration users john > show
Properties:
    logname = john
    type = local
    uid =
    fullname =
    initial_password = *****
    require_annotation = false
```

```
roles = basic
kiosk_mode = false
kiosk_screen = status/dashboard
```

Children:

```
exceptions => Configure this user's exceptions
preferences => Configure user preferences
```

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Changing a User Password (BUI)

Use the following procedure to change a user's password. To change the password for any user other than yourself, you must have Super-User (root) privileges or a role with the user authorization/exception.

1. **Go to Configuration > Users.**
2. **Click the edit icon  next to the User for which you want to change the password.**
3. **In the Edit Local User dialog box, type a new Password and then type it again to confirm it.**
4. **Click APPLY.**

Related Topics

- [Editing Exceptions for a User BUI](#)
- [Editing Authorizations for a Role BUI](#)

▼ Changing a User Password (CLI)

Use the following procedure to change a user's password. To change the password for any user other than yourself, you must have Super-User (root) privileges or a role with the user authorization/exception.

1. **Go to configuration users and then enter show to view a list of users.**

```
hostname:> configuration users
hostname:configuration users > show
Users:
```

NAME	USERNAME	UID	TYPE
Super-User	root	0	Loc

2. **Enter select and the username of the user for which you want to change the password. Then enter show.**

```
hostname:configuration users > select root
hostname:configuration users root > show
Properties:
```

```
    logname = root
    fullname = Super-User
    initial_password = *****
    require_annotation = false
```

Children:

```
    preferences => Configure user preferences
```

3. **Enter set initial_password= and the new password.**

```
hostname:configuration users root > set initial_password=[new password]
    initial_password = ***** (uncommitted)
```

4. **Enter commit.**

```
hostname:configuration users root > commit
```

Related Topics

- Editing Exceptions for a User [CLI](#)
- Editing Authorizations for a Role [CLI](#)

▼ Editing Exceptions for a User (BUI)

Use the following procedure to edit exceptions for a user.

1. **Go to Configuration > Users.**

2. **Hover over the user in the Users list, and click the edit icon .**
3. **Click on Exceptions.**
4. **Select Scope.**
If filters are available for this scope, they are listed below the Scope selector.
5. **Click the checkbox for each exception you want to add.**
6. **Click ADD in the Exceptions section.**

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Editing Exceptions for a User (CLI)

Use the following procedure to edit exceptions for a user.

1. **Go to configuration users.**
2. **Type select followed by the username.**
3. **Type exceptions.**
4. **Type create.**
5. **Type set scope= followed by the scope name. Use tab-completion to see the list.**
6. **Type show to list properties.**
7. **Type set to set the desired properties to true.**
8. **Type commit.**
The exception has now been added.

Example 3 Adding an Exception to Exclude Scope Authorizations

This example adds an exception to exclude svc scope authorizations for the user "brendan":

```

hostname:configuration users brendan > exceptions
hostname:configuration users brendan exceptions > create
hostname:configuration users brendan auth (uncommitted) > show
Properties:
    scope = (unset)
hostname:configuration users brendan auth (uncommitted) > set scope=svc
    scope = svc
hostname:configuration users brendan auth (uncommitted) > show
Properties:
    scope = svc
    service = *
    allow_administer = false
    allow_configure = false
    allow_restart = false
hostname:configuration users brendan auth (uncommitted) > commit
hostname:configuration users brendan exceptions > show
Auths:

NAME          OBJECT          PERMISSIONS
auth-000      svc.*           none

hostname:configuration users brendan exceptions > select auth-000
hostname:configuration users brendan auth-000 > show
Properties:
    scope = svc
    service = *
    allow_administer = false
    allow_configure = false
    allow_restart = false

hostname:configuration users brendan auth-000 >

```

Example 4 Adding an Exception to Include Scope Authorizations

This example adds an exception to include a scope authorization that is not part of the role "webadmin":

```

hostname:configuration users brendan exceptions > create
hostname:configuration users brendan auth (uncommitted) > set scope=appliance
    scope = appliance
hostname:configuration users brendan auth (uncommitted) > show
Properties:
    scope = appliance
    service = *
    allow_audit = false
    allow_factoryReset = false

```

```

allow_powerOff = false
allow_reboot = false
allow_setName = false
allow_shell = false

hostname:configuration users brendan auth (uncommitted) > set allow_audit=true
allow_audit = true (uncommitted)
hostname:configuration users brendan auth (uncommitted) > commit
hostname:configuration users brendan exceptions > show
Auths:

NAME            OBJECT            PERMISSIONS
auth-000        svc.*             none
auth-001        appliance.*       audit

hostname:configuration users brendan exceptions >

```

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Deleting Exceptions for a User (BUI)

Use the following procedure to delete exceptions for a user.

1. **Go to Configuration > Users.**
2. **Hover over the user in the Users list, and click the edit icon .**
3. **Click on Exceptions.**
4. **Hover over the exception in the bottom list, and click the trash icon .**
5. **Click APPLY at the top of the dialog box.**

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Deleting Exceptions for a User (CLI)

Use the following procedure to delete exceptions for a user.

1. **Go to configuration users.**
2. **Type select followed by the username.**
3. **Type exceptions.**
4. **Type show to list the exceptions.**
5. **Type destroy followed by the exception name. The exception has now been destroyed.**

Example 5 Deleting an Exception for a User

```
hostname:configuration users > select john
hostname:configuration users john > ls
Properties:
    logname = john
    type = local
    uid = 2000000001
    fullname = john
    initial_password = *****
    require_annotation = false
    kiosk_mode = false
    kiosk_screen = status/dashboard

Children:
    exceptions => Configure this user's exceptions
    preferences => Configure user preferences

hostname:configuration users john > exceptions
hostname:configuration users john exceptions > show
Auths:

NAME      OBJECT      PERMISSIONS
auth-000  ad.*        domain
                        workgroup

hostname:configuration users john exceptions > destroy auth-000
This will destroy "auth-000". Are you sure? (Y/N)
hostname:configuration users john exceptions > show
hostname:configuration users john exceptions >
```


Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Adding a Role (BUI)

Use the following procedure to add a role.

1. **Go to Configuration > Users.**
2. **Click the add icon  next to Roles.**
3. **Set the name of the role and description.**
4. **(Optional) Under Authorizations, select a scope.**
If filters are available for this scope, they will appear below the Scope selector.
5. **(Optional) Select filters for the scope if appropriate.**
6. **(Optional) Click the checkbox for each authorization to add.**
7. **Click ADD at the top of the dialog box.**
The new role appears in the Roles list.

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Adding a Role (CLI)

Use the following procedure to add a role.

1. **Go to configuration roles.**
2. **Type role followed by the role name you want to create.**

3. **Set the description, then type `commit` to commit the role.**
4. **(Optional) Type `authorizations`.**
5. **(Optional) Type `create` to add an authorization.**
6. **(Optional) Type `set scope=` followed by the scope name. Use tab-completion to see the list.**
7. **(Optional) Type `show` to see both available filters and authorizations.**
8. **(Optional) Type `set` to set the desired authorizations to true, and set the filters (if available). Tab-completion helps show which filter settings are valid.**
9. **Type `commit`.**
The new role has now been added.

Example 6 Creating the Role "webadmin"

```
hostname:> configuration roles
hostname:configuration roles > role webadmin
hostname:configuration roles webadmin (uncommitted) > set
                        description="web server administrator"
                        description = web server administrator (uncommitted)
hostname:configuration roles webadmin (uncommitted) > commit
hostname:configuration roles > show
Roles:
```

NAME	DESCRIPTION
basic	Basic administration
webadmin	web server administrator

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Editing Authorizations for a Role (BUI)

A role is a collection of privileges that can be assigned to a user. Use the following procedure to edit authorizations for a role.

1. **Go to Configuration > Users.**
2. **Hover over the role in the Roles list, and click the edit icon .**
3. **Under Authorizations, select a scope.**
If filters are available for this scope, they will appear below the Scope selector.
4. **Select filters for the scope if appropriate.**
5. **Click the checkbox for each authorization to add.**
6. **Click ADD in the Authorizations section.**
The authorizations are added to the bottom of the dialog box.
7. **Click APPLY at the top of the dialog box.**

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Editing Authorizations for a Role (CLI)

A role is a collection of privileges that can be assigned to a user. Use the following procedure to edit authorizations for a role.

1. **Go to configuration roles.**
2. **Type select followed by the role name.**
3. **Type authorizations.**
4. **Type create to add an authorization.**
5. **Type set scope= followed by the scope name. Use tab-completion to see the list.**
6. **Type show to see both available filters and authorizations.**
7. **Type set to set the desired authorizations to true, and set the filters (if available). Tab-completion helps show which filter settings are valid.**

8. Type commit.

The authorization has now been added.

Example 7 Adding the Authorization to Restart the HTTP Service

This example adds the authorization to restart the HTTP service. This example also shows the output of tab-completion, which lists valid input and is useful when determining the valid scopes and filter options.

```
hostname:configuration roles > select webadmin
hostname:configuration roles webadmin > authorizations
hostname:configuration roles webadmin authorizations > create
hostname:configuration roles webadmin auth (uncommitted) > set scope=tab
ad          cluster    net          schema      update
alert       hardware    replication  stat        user
appliance   nas            role        svc         worksheet
hostname:configuration roles webadmin auth (uncommitted) > set scope=svc
scope = svc
hostname:configuration roles webadmin auth (uncommitted) > show
Properties:
    scope = svc
    service = *
    allow_administer = false
    allow_configure = false
    allow_restart = false

hostname:configuration roles webadmin auth (uncommitted) > set service=tab
*          ftp         ipmp        nis         ssh
ad         http        iscsi       ntp         tags
smb        identity    ldap        routing     vscan
datalink:igb0 idmap      ndmp        scrk
dns        interface:igb0 nfs         snmp
hostname:configuration roles webadmin auth (uncommitted) > set service=http
service = http (uncommitted)
hostname:configuration roles webadmin auth (uncommitted) > set allow_restart=true
allow_restart = true (uncommitted)
hostname:configuration roles webadmin auth (uncommitted) > commit
hostname:configuration roles webadmin authorizations > list
NAME      OBJECT      PERMISSIONS
auth-000  svc.http    restart
```

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)

- [Managing User Properties](#)

▼ Deleting Authorizations from a Role (BUI)

Use the following procedure to delete authorizations from a role.

1. **Go to Configuration > Users.**
2. **Hover over the role in the Roles list, and click the edit icon .**
3. **Hover over the authorization in the bottom list, and click the trash icon .**
4. **Click APPLY at the top of the dialog box.**

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Deleting Authorizations from a Role (CLI)

Use the following procedure to delete authorizations from a role.

1. **Go to configuration roles.**
2. **Type select followed by the role name.**
3. **Type authorizations.**
4. **Type show to list authorizations.**
5. **Type destroy followed by the authorization name.**

Example 8 Deleting an Authorization from a Role

```
hostname:configuration roles > select test_role1
hostname:configuration roles test_role1 > authorizations
```

```
hostname:configuration roles test_role1 authorizations > show
Auths:

NAME      OBJECT      PERMISSIONS
auth-000  ad.*        domain
          workgroup

hostname:configuration roles test_role1 authorizations > destroy auth-000
This will destroy "auth-000". Are you sure? (Y/N)
hostname:configuration roles test_role1 authorizations > show
hostname:configuration roles test_role1 authorizations >
```

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

▼ Adding a User Who Can View the Dashboard

Use the following procedure to add a user who can only view the dashboard.

1. **Add either a Directory or Local user as described in “[Adding an Administrator or User \(BUI\)](#)” on page 203.**
2. **Select the Kiosk user checkbox. Ensure the Kiosk screen is set to status/dashboard.**
3. **Click ADD.**

The user should now be able to login, but only view the dashboard.

Related Topics

- [Understanding Users and Roles](#)
- [User Authorizations](#)
- [Managing User Properties](#)

Understanding Users and Roles

A user can be one of the following types:

Administrator User Type	Non-Administrator User Type
Local - A locally defined appliance administrator; can be granted privileges by assigning custom roles; can optionally have a UID specified. Although local users are supported for data services, local groups are not supported.	Data-only - A data-only user defined locally for data (SMB, NFS, FTP, etc.) with no administrator access; can optionally have a UID specified.
Directory - An appliance administrator managed by a directory service (NIS or LDAP); can be granted privileges by assigning custom roles.	No-login - A username and UID reserved for identity mapping purposes. This user type is not allowed to log in to the appliance and can optionally have a UID specified.

Local and Directory users are administrator types, and can be granted privileges by assigning custom *roles*.

A role is a collection of privileges that can be assigned to an administrator user type. Newly created administrator users default to the "basic" role, which enables logging in to the administrative interface, but does not allow changes. All administrator users can read most system configuration parameters, and any role can be edited to add or delete authorizations.

The use of roles is more secure than giving everyone the *root* password. Roles restrict users to necessary authorizations only, and also attribute their actions to their individual username in the log. For example, you can create *administrator* and *operator* roles, with different authorization levels. Staff members can be assigned any role that is suitable for their needs, without assigning unnecessary privileges.

Related Topics

- Adding an Administrator or User [BUI](#), [CLI](#)
- Changing a User Password [BUI](#), [CLI](#)
- Editing Exceptions for a User [BUI](#), [CLI](#)
- Deleting Exceptions for a User [BUI](#), [CLI](#)
- Adding a Role [BUI](#), [CLI](#)
- Editing Authorizations for a Role [BUI](#), [CLI](#)
- Deleting Authorizations from a Role [BUI](#), [CLI](#)
- Adding a User Who can Only View the Dashboard [BUI](#)

User Authorizations

Authorizations allow users to perform specific tasks, such as creating shares, rebooting the appliance, and updating the system software. Authorizations are grouped into *scopes*, and each scope may have a set of optional filters to narrow the scope of the authorization. For example, rather than an authorization to restart all services, a filter can be used so that this authorization can restart the HTTP service only.

The following table shows the available scopes:

TABLE 39 User Available Scopes, Filters, and Authorizations

Scope BUI	Scope CLI	Filters	Authorizations
Active Directory	ad	Domain or workgroup name	■ Join an Active Directory domain ■ Join a workgroup
Alerts	alert	-	Configure alert filters and thresholds
Analytics	stat	List of drilldowns	■ Configure analytics hostname lookup policy ■ Create a statistic with this drilldown present ■ Read a statistic with this drilldown present
Appliance	appliance	Appliance name	■ Emit an audit log entry ■ Restore the appliance to factory defaults ■ Power down the appliance ■ Reboot the appliance ■ Modify the appliance name ■ Access the underlying Solaris shell ■ Configure system certificates ■ Configure trusted certificates
Clustering	cluster	-	■ Failback resources to a cluster peer ■ Reset a failed cluster I/O device ■ Takeover resources from a cluster peer ■ Transfer resources to a cluster peer
Datasets	dataset	-	Configure dataset retention policies
Hardware	hardware	-	■ Online and offline disks ■ Configure LEDs on disks, appliance, and external enclosures ■ Configure network properties for the service processor ■ Remove a drive as a hot spare ■ Configure a storage pool ■ Unconfigure a storage pool
Keystores	keystore	Keystore name	■ List keys present in a per-user keystore ■ Permit keystore modifications ■ Permit read access to sensitive values in a keystore
Networking	net	-	Configure networking devices, datalinks, and interfaces
Projects and shares	nas	■ Storage pool ■ Project ■ Share	■ Configure who can access a share ■ Change general properties on a share ■ Configure protocol-specific properties ■ Change quota and reservation on a share ■ Change user and group quotas on a share ■ Clear locks held on behalf of an NFS client

Scope BUI	Scope CLI	Filters	Authorizations
			■ Clone a snapshot to a normal filesystem ■ Create a project ■ Create a filesystem or LUN ■ Remove a project or share ■ Manage encryption keys for a project or share ■ Promote a clone ■ Rename a project or share ■ Rollback a filesystem to a previous snapshot ■ Configure data replication to other appliances ■ Manage data replicated from other appliances ■ Configure a recurring schedule of snapshots ■ Check a storage pool for errors ■ Manage shadow migration on a share ■ Take a manual snapshot
Roles	role	Role name	■ Configure authorizations for a role ■ Change a description of a role ■ Create a role ■ Destroy a role
SAN	stmf	-	Configure SAN hosts and targets
Services	svc	Service name	■ Enable or disable service ■ Configure service properties and settings ■ Restart service
Shares property schema	schema	-	Modify property schema
Update	update	-	■ Delete system software ■ Update system software ■ Upload system updates
Users	user	Username	■ Configure authorizations for a user ■ Change a password ■ Configure preferences for a user ■ Configure properties for a user ■ Configure roles for a user ■ Create a user ■ Destroy a user
Workflow	workflow	■ Owner ■ Name	■ Delete workflow ■ Execute workflow
Worksheet	worksheet	■ Owner ■ Name	■ Modify worksheet ■ Read worksheet

Related Topics

- Adding an Administrator or User [BUI](#), [CLI](#)

- Changing a User Password [BUI](#), [CLI](#)
- Editing Exceptions for a User [BUI](#), [CLI](#)
- Deleting Exceptions for a User [BUI](#), [CLI](#)
- Adding a Role [BUI](#), [CLI](#)
- Editing Authorizations for a Role [BUI](#), [CLI](#)
- Deleting Authorizations from a Role [BUI](#), [CLI](#)
- Adding a User Who can Only View the Dashboard [BUI](#)

Managing User Properties

The Configuration > Users page lists both users and groups, along with buttons for administration. Hover over an entry to expose its clone, edit, and destroy buttons. Double-click a user or role, or click its edit icon , to view its edit screen. The icons are as follows:

TABLE 40 Users BUI Page Icons

Icon	Description
	Add a new user or role. A new dialog box is displayed where you enter the required properties.
	Open a search box. Enter a search string and hit enter to search the user or role lists for that text, and only display entries that match. Click the icon again to return to the full listing.
	Clone a user or role. Add a new user or role starting with fields based on the values from this entry.
	Edit a user or role.
	Remove a user, role, or authorization.

Depending on the type of user, all of the following properties can be set when adding a user, and a subset of these when editing a user:

TABLE 41 User Properties

Property	Description
Type	For a description of user types, see “Understanding Users and Roles” on page 218 .
Username	Unique name for user
User ID	Enabled only for Local, Data, and No-login users. You can either choose to have an automatically assigned user

Property	Description
	ID or specify a user ID on your own. The self-assigned ID is not allowed to be less than 100, or be greater than 2147483646, or equal to 60001, 60002, or 65534.
Full Name	User description
Password/Confirm	For Local and Data users, type the initial password in both of these fields
Require session annotation	If enabled, when users log in to the appliance they must provide a text description of the purpose of their login. This annotation may be used to track work performed for requests in a ticketing system, and the ticket ID can be used as the session annotation. The session annotation appears in the log.
Kiosk user	If enabled, the user will only be able to view the screen in the "Kiosk screen" setting. This can be used to restrict a user to only see the dashboard, for example. A kiosk user will not be able to access the appliance via the CLI.
Kiosk screen	The screen that this kiosk user is restricted to, if "Kiosk user" is enabled
Roles	The roles assigned to a Directory or Local user
Exceptions	These authorizations are included or excluded from those normally available for the selected roles. For example, an exception can be added to exclude all or some scope authorizations for a role assigned to a user, or to include scope authorizations for a user without that scope defined in a role.

The following properties can be set when managing roles:

TABLE 42 Role Properties

Property	Description
Name	Name of the role as it will be shown in lists
Description	Verbose description of role if desired
Exceptions	Exceptions for this role

Related Topics

- Adding an Administrator or User [BUI](#), [CLI](#)
- Changing a User Password [BUI](#), [CLI](#)
- Editing Exceptions for a User [BUI](#), [CLI](#)
- Deleting Exceptions for a User [BUI](#), [CLI](#)
- Adding a Role [BUI](#), [CLI](#)

- Editing Authorizations for a Role [BUI](#), [CLI](#)
- Deleting Authorizations from a Role [BUI](#), [CLI](#)
- Adding a User Who can Only View the Dashboard [BUI](#)

Setting Appliance Preferences

This section contains preference settings for your locality, session properties, advanced analytics, and SSH keys.

To configure your preferences, use the following sections:

- Setting Preferences - [BUI](#), [CLI](#)
- Setting SSH Public Keys - [BUI](#), [CLI](#)
- [Preference Properties](#)

▼ Setting Preferences (BUI)

Use the following procedure to set preferences for the current user account. If you log into the BUI with other than your own account, the preferences are saved for that user, such as the root user.

To change preferences for user accounts other than the one currently logged in to, see “[Setting Preferences \(CLI\)](#)” on page 225.

1. **Go to Configuration > Preferences.**
2. **Modify the properties with values described in [Preference Properties](#).**

The screenshot shows the 'General' tab of a configuration interface. At the top right are 'REVERT' and 'APPLY' buttons. The 'Initial login screen' field contains the URL 'https://ar7320-230:215/#status/dashboard'. The 'Locality' dropdown menu is set to 'C'. The 'Session timeout' is set to '15' minutes. There is an 'Edit...' link next to the 'Current session annotation' field. At the bottom, there is a checkbox labeled 'Make available advanced analytics statistics' which is currently unchecked.

3. Click **APPLY**.

▼ Setting Preferences (CLI)

Use the following examples to set preferences for user accounts. If you log into the CLI with other than your own account, the preferences are saved for that user, such as the root user. See Example 2 for how to change preferences for user accounts other than the one currently logged in to.

● Choose the appropriate example.

Example 9 Setting Preferences for the Current User Account

To set preferences for the current user account, use the following CLI commands.

This example shows setting the session annotation property, which can only be set for the currently logged in user.

```
hostname:> configuration preferences
hostname:configuration preferences> show
Properties:
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    session_annotation =
    advanced_analytics = false

Children:
    keys => Manage SSH public keys

hostname:configuration preferences> set session_annotation="Editing my user preferences"
    session_annotation = Editing my user preferences (uncommitted)
hostname:configuration preferences> commit
```

Example 10 Setting Preferences for a Different User Account

To set preferences for a different user account, use the following CLI commands. Note that you cannot set a session annotation for a user other than the currently logged in user.

This example shows enabling advanced analytics for a selected user.

```
hostname:> configuration users
hostname:configuration users> select brendan
```

```
hostname:configuration users brendan> preferences
hostname:configuration users brendan preferences> show
Properties:
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    advanced_analytics = false

Children:
    keys => Manage SSH public keys

hostname:configuration users brendan preferences> set advanced_analytics=true
    advanced_analytics = true (uncommitted)
hostname:configuration users brendan preferences> commit
```

▼ Setting SSH Public Keys (BUI)

SSH public keys can be used to allow SSH connections without the use of passwords. This feature is useful for administrator convenience and for automated execution of scripts.

Use the following procedure to set SSH public keys for the current user. To set keys for other users, see [“Setting SSH Public Keys \(CLI\)” on page 226](#).

1. **Go to Configuration > Preferences.**
2. **Click the add icon  next to SSH Public Keys.**
3. **Select a Type, and then type the SSH public key and a key comment.**
4. **Click ADD.**

▼ Setting SSH Public Keys (CLI)

SSH public keys can be used to allow SSH connections without the use of passwords. This feature is useful for administrator convenience and for automated execution of scripts.

Use the following examples to set SSH public keys for user accounts. If you log into the CLI with other than your own account, the keys are saved for that user, such as the root user. See Example 2 for how to change keys for user accounts other than the one currently logged in to.

- **Choose the appropriate example.**

Example 11 Setting SSH Public Keys for the Current User Account

To set SSH public keys for the current user account, use the following CLI commands.

```
hostname:> configuration preferences
hostname:configuration preferences> show
Properties:
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    advanced_analytics = false

Children:
    keys => Manage SSH public keys

hostname: configuration preferences> keys
hostname:configuration preferences keys> create
hostname:configuration preferences key (uncommitted)> set type=DSA
hostname:configuration preferences key (uncommitted)> set key="...DSA key text..."
    key = ...DSA key text...(uncommitted)
hostname:configuration preferences key (uncommitted)> set comment="fw-log1"
    comment = fw-log1 (uncommitted)
hostname:configuration preferences key (uncommitted)> commit
hostname:configuration preferences keys> show
Keys:
```

NAME	MODIFIED	TYPE	COMMENT
key-000	07/12/2015 10:54:58	DSA	fw-log1

Example 12 Setting SSH Public Keys for a Different User Account

To set SSH public keys for a different user account, use the following CLI commands.

```
hostname:> configuration users
hostname:configuration users> select john
hostname:configuration users john> preferences show
Properties:
    locale = C
    login_screen = status/dashboard
    session_timeout = 15
    advanced_analytics = false

Children:
    keys => Manage SSH public keys
```

```

hostname: configuration users john> preferences keys
hostname: configuration users john preferences keys> create
hostname: configuration users john preferences key (uncommitted)> set type=DSA
hostname: configuration users john preferences key (uncommitted)> set key="...DSA key text..."
      key = ...DSA key text...(uncommitted)
hostname: configuration users john preferences key (uncommitted)> set comment="fw-log2"
      comment = fw-log2 (uncommitted)
hostname: configuration users john preferences key (uncommitted)> commit
hostname: configuration users john preferences keys> show
Keys:

```

NAME	MODIFIED	TYPE	COMMENT
key-001	07/13/2015 10:57:58	DSA	fw-log2

Preference Properties

The following table describes the properties for setting user preferences.

TABLE 43 Preference Properties

Property	Description
Initial login screen	First page the BUI will load after a successful login. By default, this is the Status Dashboard.
Locality	C by default. C and POSIX Localities support only ASCII characters or plain text. ISO 8859-1 supports the following languages: Afrikaans, Basque, Catalan, Danish, Dutch, English, Faeroese, Finnish, French, Galician, German, Icelandic, Irish, Italian, Norwegian, Portuguese, Spanish and Swedish.
Session timeout	Time after navigating away from the BUI that the browser will automatically logout the session
Current session annotation	Annotation text added to audit logs
Advanced analytics statistics	This will make available additional statistics in Analytics
SSH Public Keys	RSA/DSA public keys. Text comments can be associated with the keys to help administrators track why they were added. In the BUI, these keys apply only for the current user; to add keys for other users, use the CLI.

Configuring Alerts

This section describes system Alerts, how they are customized, and where to find alert logs. To monitor statistics from Analytics, create custom threshold alerts. To configure the system to respond to certain types of alerts, use Alert actions.

To configure alerts, use the following sections:

- [Adding an Alert Action \(BUI\)](#)
- [Adding an Alert Action \(CLI\)](#)
- [Sending Email Alerts \(CLI\)](#)
- [Sending an SNMP Trap \(CLI\)](#)

To learn more about alerts, see the following sections:

- [Alert Categories](#)
- [Sending Syslog Messages](#)
- [Executing a Workflow](#)
- [Threshold Alerts](#)
- [Resuming/Suspending Datasets and Worksheets](#)

▼ Adding an Alert Action (BUI)

1. Click the add icon next to "Alert actions".
2. Select the Category, or pick "All events" for everything.
3. Either pick All Events, or a Subset of Events. If the subset is selected, customize the checkbox list to match the desired alerts events.
4. Use the drop down menu in "Alert actions" to select which alert type.
5. Enter details for the Alert action. The "TEST" button can be clicked to create a test alert and execute this alert action (useful for checking if email or SNMP is configured correctly).
6. The add icon next to "Alert actions" can be clicked to add multiple alerts actions.
7. Click "ADD" at the top right.

Related Topics

- [Alert Categories](#)
- [Sending Syslog Messages](#)
- [Executing a Workflow](#)
- [Threshold Alerts](#)
- [Resuming/Suspending Datasets and Worksheets](#)

▼ Adding an Alert Action (CLI)

1. Enter the configuration `alerts actions` context, and enter the `create` command.
2. Go to the "category" property by entering `get category = (unset)`.
3. Enter `set category=thresholds`.
4. Enter `set thresholdid=where` [id is the identifier that was automatically created for the threshold alert].
5. Enter `commit`.
6. Enter `list` to determine the name, including number, of the new alert action. Look for a threshold without an assigned action and handler.
7. Enter `select actions-where` [number is the same number identified in the previous step].
8. Enter `action`, and then enter `get`.
9. By default, the alert type is email. If this is what you want, skip to the next step. If not, enter `set handler=where` [type is either `snmptrap`, `syslog`, `resumedataset`, `suspenddataset`, `resumeworksheet`, `suspendworksheet`, or `executeworkflow`. Then enter `get` to view the needed arguments. Only `snmptrap` and `syslog` do not have arguments.
10. Set each needed argument. For example, to set a subject line for an email alert, enter `set subject=where` [subject is the desired email subject line].
11. Use the `show` command to ensure all arguments have been entered.
12. Enter `commit`, and then enter `list`. If necessary, correct any arguments now.
13. Enter `done`, and then enter `done` again.

Related Topics

- [Alert Categories](#)
- [Sending Syslog Messages](#)
- [Executing a Workflow](#)

- [Threshold Alerts](#)
- [Resuming/Suspending Datasets and Worksheets](#)

Sending Email Alerts (CLI)

An email containing the alert details can be sent. The configuration requires an email address and email subject line. The following example shows an email threshold alert. Details on how the appliance sends mail can be configured on the SMTP service screen.

```
From: aknobody@caji.com Mon Oct 13 15:24:47 2009
Date: Mon, 13 Oct 2009 15:24:21 +0000 (GMT)
From: Appliance on caji <noreply@caji.com>
Subject: High CPU on caji
To: admin@hostname.com

SUNW-MSG-ID: AK-8000-TT, TYPE: Alert, VER: 1, SEVERITY: Minor
EVENT-TIME: Mon Oct 13 15:24:12 2009
PLATFORM: i86pc, CSN: 0809QAU005, HOSTNAME: caji
SOURCE: svc:/appliance/kit/akd:default, REV: 1.0
EVENT-ID: 15a53214-c4e7-eae4-dae6-a652a51ea29b
DESC: cpu.utilization threshold of 90 is violated.
AUTO-RESPONSE: None.
IMPACT: The impact depends on what statistic is being monitored.
REC-ACTION: The suggested action depends on what statistic is being monitored.

SEE: https://192.168.2.80:215/#maintenance/alert=15a53214-c4e7-eae4-dae6-a652a51ea29b
```

Related Topics

- [Alert Categories](#)
- [Sending Syslog Messages](#)
- [Executing a Workflow](#)
- [Threshold Alerts](#)
- [Resuming/Suspending Datasets and Worksheets](#)

Sending an SNMP Trap (CLI)

An SNMP trap containing alert details can be sent, if an SNMP trap destination is configured in the SNMP service, and that service is online. The following example sends an SNMP trap, as seen from the Net-SNMP tool `snmpttrapd -P`.

```
# /usr/sfw/sbin/snmptrapd -P
2009-10-13 15:31:15 NET-SNMP version 5.0.9 Started.
2009-10-13 15:31:34 caji.com [192.168.2.80]:
    iso.3.6.1.2.1.1.3.0 = Timeticks: (2132104431) 246 days, 18:30:44.31
    iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.42.2.225.1.3.0.1
    iso.3.6.1.4.1.42.2.225.1.2.1.2.36.55.99.102.48.97.99.100.52.45.51.48.
99.49.45.52.99.49.57.45.101.57.99.98.45.97.99.50.55.102.55.49.50.54.
98.55.57 = STRING: "7cf0acd4-30c1-4c19-e9cb-ac27f7126b79"
    iso.3.6.1.4.1.42.2.225.1.2.1.3.36.55.99.102.48.97.99.100.52.45.51.48.
99.49.45.52.99.49.57.45.101.57.99.98.45.97.99.50.55.102.55.49.50.54.
98.55.57 = STRING: "alert.ak.xmlrpc.threshold.violated"
    iso.3.6.1.4.1.42.2.225.1.2.1.4.36.55.99.102.48.97.99.100.52.45.51.
48.99.49.45.52.99.49.57.45.101.57.99.98.45.97.99.50.55.102.55.49.50.
54.98.55.57 = STRING: "cpu.utilization threshold of 90 is violated."
```

Related Topics

- [Alert Categories](#)
- [Sending Syslog Messages](#)
- [Executing a Workflow](#)
- [Threshold Alerts](#)
- [Resuming/Suspending Datasets and Worksheets](#)

Alert Categories

Important appliance events trigger alerts, which includes hardware and software faults. These alerts appear in the Logs and may also be configured to execute any of the Alert actions.

Alerts are grouped into the following categories:

TABLE 44 Alert Categories

Category	Description
Cluster	Cluster events, including link failures and peer errors
Custom	Events generated from the custom alert configuration
Hardware Events	Appliance boot and hardware configuration changes
Hardware Faults	Any hardware fault
NDMP operations	NDMP TAR/DUMP backup and restore start and finish events. This group is available as "NDMP: backup only" and "NDMP: restore only"
Network	Network port, datalink, and IP interface events and failures

Category	Description
Phone Home	Support bundle upload events
Remote replication	Send and receive events and failures. This group is available as "Remote replication: source only" and "Remote replication: target only", for just source or target events
Service failures	Software services failure events
Thresholds	Custom alerts based on Analytics statistics
ZFS pool	Storage pool events, including scrub and hot space activation

Sending Syslog Messages

When the Syslog service is enabled, a syslog message containing alert details can be sent to one or more remote systems. For more information about sending syslog messages, see [Syslog Relay service](#).

Executing a Workflow

Workflows can be optionally executed as alert actions. To let workflow be eligible as an alert action, its alert action must be set to true. For more information about executing a workflow, see [“Maintenance Workflows” on page 657](#).

Threshold Alerts

These are alerts based on the statistics from Analytics. The following are properties when creating threshold alerts:

TABLE 45 Threshold Alert Properties

Property	Description
Threshold	The threshold statistic is from Analytics, and is self descriptive (such as "Protocol: NFSv4.0 operations per second")
exceeds/falls below	defines how the threshold value is compared to the current statistic
Timing: for at least	Duration which the current statistic value must exceed/ fall below the threshold
only between/only during	These properties may be set so that the threshold is only sent during certain times of day - such as business hours

Property	Description
Repost alert every ... this condition persists.	If enabled, this will re-execute the alert action (such as sending email) every set interval while the threshold breach exists
Also post alert when this condition clears for at least ...	Send a followup alert if the threshold breach clears for at least the set interval

The "Add Threshold Alert" dialog has been organized so that it can be read as though it is a paragraph describing the alert. The default reads:

Threshold CPU: percent utilization exceeds 95 percent.

Timing for at least 5 minutes only between 0:00 and 0:00 only during weekdays.

Repost alert every 5 minutes while this condition persists.

Also post alert when this condition clears for at least 5 minutes.

Related Topics

- [“Configuring a Threshold Alert \(BUI\)” in Oracle ZFS Storage Appliance Analytics Guide, Release OS8.7.x](#)
- [“Configuring a Threshold Alert \(CLI\)” in Oracle ZFS Storage Appliance Analytics Guide, Release OS8.7.x](#)
- [Adding an Alert Action \(BUI\)](#)
- [Adding an Alert Action \(CLI\)](#)
- [Sending Email Alerts \(CLI\)](#)
- [Sending an SNMP Trap \(CLI\)](#)

Resuming/Suspending Analytics Datasets and Worksheets

Analytics datasets may be resumed or suspended. This is particularly useful when tracking down sporadic performance issues, and when enabling these datasets 24x7 is not desirable.

For example, imagine you noticed a spike in CPU activity once or twice a week, and other analytics showed an associated drop in NFS performance. You enable some additional datasets, but you don't quite have enough information to prove what the problem is. If you could enable the NFS by hostname and filename datasets, you are certain you will understand the cause a lot better. However those particular datasets can be heavy handed - leaving them enabled 24x7 will degrade performance for everyone. This is where the resume/suspend dataset actions may be

of use. A threshold alert could be configured to *resume* paused NFS by hostname and filename datasets, only when the CPU activity spike is detected; a second alert can be configured to then *suspend* those datasets, after a short interval of data is collected. The end result - you collect the data you need only during the issue, and minimize the performance impact of this data collection.

For more information on datasets, see [“About Analytics Datasets” in Oracle ZFS Storage Appliance Analytics Guide, Release OS8.7.x.](#)

These actions are to resume or suspend an entire Analytics worksheet, which may contain numerous datasets. The reasons for doing this are similar to those for resuming and suspending datasets. For more information, see [“Worksheet Graphs and Plots” in Oracle ZFS Storage Appliance Analytics Guide, Release OS8.7.x.](#)

Related Topics

- [“Configuring a Threshold Alert \(BUI\)” in Oracle ZFS Storage Appliance Analytics Guide](#)
- [“Configuring a Threshold Alert \(CLI\)” in Oracle ZFS Storage Appliance Analytics Guide](#)
- [“Adding an Alert Action \(BUI\)” on page 229](#)
- [“Adding an Alert Action \(CLI\)” on page 230](#)
- [“Sending Email Alerts \(CLI\)” on page 231](#)
- [“Sending an SNMP Trap \(CLI\)” on page 231](#)

Configuring Certificates

This section describes the use of public key certificates. Public key certificates and their trust chains provide a mechanism to digitally identify a system without having to manually exchange any secret information.

A public key certificate is a blob of data that encodes a public key value, some information about the generation of the certificate, such as a name and who signed it, a hash or checksum of the certificate, and a digital signature of the hash. Together, these values form the certificate. The digital signature ensures that the certificate has not been modified.

The appliance supports customer-owned certificates. The life cycle of a certificate starts with generating a certificate signing request (CSR). The CSR is then sent to the certificate authority (CA) for signature. After the signed certificate is returned from the CA, it can be installed on the appliance. If a certificate is signed by a non-root CA, you must also obtain certificates from the second- and higher-level CAs.

There are two types of certificates, that you can manage. System certificates identify the current system. Trusted certificates are those that identify remote systems.

To manage system certificates, use the following tasks:

- Creating a New System Certificate - [BUI](#), [CLI](#)
- Uploading CA Certificates from Non-root CAs - [BUI](#), [CLI](#)
- Viewing CSR and System Certificate Details - [BUI](#), [CLI](#)
- Destroying a CSR or System Certificate - [BUI](#), [CLI](#)
- Setting the Appliance or Default System Certificate - [BUI](#), [CLI](#)

To manage trusted certificates, use the following tasks:

- Uploading a Trusted Certificate - [BUI](#), [CLI](#)
- Viewing Trusted Certificate Details - [BUI](#), [CLI](#)
- Destroying a Trusted Certificate - [BUI](#), [CLI](#)
- Assigning a Certificate to a Service - [BUI](#), [CLI](#)

▼ Creating a New Server Certificate (BUI)

To create a new server certificate, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the System tab.**
3. **To create a new CSR, click the add item icon .**
To create a new CSR based on an existing CSR or certificate, hover over an existing entry and click the copy icon .
4. **Complete the CSR form.**
5. **Click CREATE.**
6. **When prompted to open the CSR or save it, select Save File and click OK to save the CSR now, or click Cancel to save the CSR later.**

To save the CSR later, hover over the entry and click the download icon .

7. **Transfer the CSR to your CA in the prescribed manner.**

8. After receiving the signed certificate from the CA, click the upload icon .
9. Browse to the signed certificate and select it.
10. Click **UPLOAD**.

▼ Creating a New Server Certificate (CLI)

To create a new server certificate, use the following steps.

1. To create a new CSR, enter the configuration settings `certificates system` context, and enter the `create` command.

Or to create a new CSR based on existing CSR or certificate, enter the above context and then the command `clone CSR` or `certificate` number. For example:

```
hostname:configuration settings certificates system> clone cert-000
```

2. To complete the CSR form, use the following CLI commands.

```
hostname:configuration settings certificates system (uncommitted)> get
    subject_commonname = hostname.us.example.com
    subject_organizationname = (unset)
    subject_organizationalunitname = (unset)
    subject_localityname = (unset)
    subject_stateorprovincename = (unset)
    subject_countryname = (unset)
    subject_emailaddress = (unset)
    dns = hostname.us.example.com
    ip = 192.0.2.1
    uri = (unset)
    comment = (unset)
hostname:configuration settings certificates system (uncommitted)> set comment="test
certificate"
    comment = test certificate (uncommitted)
hostname:configuration settings certificates system (uncommitted)> commit
```

3. To view the CSR, use following commands.

```
hostname:configuration settings certificates system> show
Properties:
    default = auto
System Certificates:
```

CERT	TYPE	SUBJECT	ISSUER	EXPIRES
cert-000	req	hostname.us.example.com		
cert-001	CA	Joe Test CA	Joe Test CA	2038-1-19
cert-002	cert	hostname.us.example.com	Joe Test CA	2038-1-21

```
hostname:configuration settings certificates system> dump cert-000
-----BEGIN CERTIFICATE REQUEST-----
MIICwzCCAAsCAQIwIjEgMB4GA1UEAxMG9zdG5hbWUudXMuzXhhbXBsZS5jb20w
...
lhwb1MXqR/3xptwym1vy5dYBJsQLKroA8nr/xFb3nhJB8nI+dxSN
-----END CERTIFICATE REQUEST-----
```

4. **Copy the CSR and transfer it to your CA in the prescribed manner.**
5. **After receiving the signed certificate from the CA, enter the configuration settings certificates system context, and enter the import command.**

```
hostname:configuration settings certificates system> import
("." to end)> -----BEGIN CERTIFICATE-----
("." to end)> MIID0DCCArigAwIBAgIBQDANBgkqhkiG9w0BAQUFADCBmDELMAKGA1UEBhMCVVMx
...

("." to end)> 2ai9ZwREdTkcjcgQDxeHNZCpcHk=
("." to end)> -----END CERTIFICATE-----
("." to end)> .
```

6. **To check the imported certificates, view all certificate entries using the command show.**

```
hostname:configuration settings certificates system> show
Properties:
    default = auto
System Certificates:
CERT      TYPE SUBJECT                      ISSUER                      EXPIRES
cert-000  req  hostname.us.example.com
cert-001  CA   Joe Test CA                      Joe Test CA                  2038-1-19
cert-002  cert hostname.us.example.com          Joe Test CA                  2038-1-21
```

▼ Uploading CA Certificates from Non-root CAs (BUI)

If your server certificate is signed by a non-root CA, you need to obtain certificates for the second- and higher-level CAs also. After obtaining these CA certificates, use the following steps to upload them.

1. Go to Configuration > Settings.
2. Click the System tab.
3. Click the upload icon .
4. Browse to the signed certificate and select it.
5. Click UPLOAD.
6. Repeat steps 3 through 5 for each signed certificate.

▼ Uploading CA Certificates from Non-root CAs (CLI)

If your server certificate is signed by a non-root CA, you also need to obtain certificates for the second- and higher-level CAs. After obtaining these CA certificates, use the following steps to upload them.

1. To upload a certificate, enter the configuration settings certificates system context, and enter the `import` command.

```
hostname:configuration settings certificates system> import
("." to end)> -----BEGIN CERTIFICATE-----
("." to end)> MIID0DCCArigAwIBAgIBQDANBgkqhkiG9w0BAQUFADCbmDELMAkGA1UEBhMCVVMx
...
("." to end)> 2ai9ZwREdTkcjcgQDxeHNZCpcHk=
("." to end)> -----END CERTIFICATE-----
("." to end)> .
```

2. Repeat step 1 for each signed certificate.
3. To check the imported certificates, view all certificate entries using the command `show`.

```
hostname:configuration settings certificates system> show
Properties:
    default = auto
System Certificates:
CERT   TYPE SUBJECT                                ISSUER                                EXPIRES
cert-000 req  hostname.us.example.com                    Joe Test CA                        2038-1-19
cert-001 CA   Joe Test CA
```

```
cert-002 cert hostname.us.example.com Joe Test CA
```

```
2038-1-21
```

▼ Viewing CSR and Certificate Details (BUI)

To view CSR and certificate details, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the Systems tab.**
3. **Hover over an existing entry and click its information icon .**
4. **When finished, click OK to close the Details window.**

▼ Viewing CSR and Certificate Details (CLI)

To view CSR and certificate details, use the following steps.

1. **To view all certificate entries, go to the configuration settings certificates system context, and enter the show command.**

```
hostname:configuration settings certificates system> show
Properties:
```

```
    default = auto
```

```
System Certificates:
```

CERT	TYPE	SUBJECT	ISSUER	EXPIRES
cert-000	req	hostname.us.example.com		
cert-001	CA	Joe Test CA	Joe Test CA	2038-1-19
cert-002	cert	hostname.us.example.com	Joe Test CA	2038-1-21

2. **To view the details of a CSR or certificate, use the following commands.**

```
hostname:configuration settings certificates system> select cert-000
```

```
hostname:configuration settings certificates system cert-000> show
```

```
Properties:
```

```
    uuid = 195071da-66ac-43a6-edfa-bbbd7451f1d5
```

```
    subject_commonname = hostname.us.example.com
```

```
    issuer_commonname = Joe Test CA
```

```
    dns = hostname.us.example.com
```

```
    ip = 192.0.2.1
```

```
    comment = test certificate
```

```

notbefore = 2014-12-4 00:31:33
notafter = 2038-01-19 00:31:33
shalfingerprint = 81:A2:4B:C4:06:A9:14:1E:3E:0B:8A:70:FB:1A:30:45:2D:93:
DD:02
md5fingerprint = B7:B2:F4:3B:BB:04:8E:11:A2:64:3D:69:BF:8A:79:CC
hostname:configuration settings certificates system cert-000> done

```

▼ Destroying a CSR or Certificate (BUI)

To destroy a CSR or certificate, use the following steps.

Note - Destroying a CSR also destroys the associated private key. Therefore, importing a certificate derived from that CSR will not be possible. Destroying a certificate also destroys the associated private key. Therefore, re-importing that certificate will not be possible.

1. **Go to Configuration > Settings.**
2. **Click the System tab.**
3. **Hover over an existing entry and click the trash icon .**
4. **Click DESTROY.**

▼ Destroying a CSR or Certificate (CLI)

To destroy a CSR or certificate, use the following steps.

Note - Destroying a CSR also destroys the associated private key. Therefore, importing a certificate derived from that CSR will not be possible. Destroying a certificate also destroys the associated private key. Therefore, re-importing that certificate will not be possible.

1. **To view all certificate entries, go to the configuration settings certificates system context, and enter the show command.**

```

hostname:configuration settings certificates system> show
Properties:
    default = auto
System Certificates:
CERT      TYPE SUBJECT                                ISSUER                                EXPIRES

```

```
cert-000 req hostname.us.example.com
cert-001 CA Joe Test CA Joe Test CA 2038-1-19
cert-002 cert hostname.us.example.com Joe Test CA 2038-1-21
```

2. **To destroy a CSR or certificate, use the following commands.**

```
hostname:configuration settings certificates system> destroy cert-002
Caution: Destroying a certificate issued by a certificate authority also
destroys the associated private key. Re-importing the certificate will not be possible.
Destroy appliance certificate? (Y/N) Y
```

▼ Setting the Appliance Certificate (BUI)

To set the appliance of default certificate, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the Systems tab.**
3. **From the drop-down menu of system certificates, select the certificate that you want to set as the default.**
4. **Click APPLY.**

▼ Setting the Appliance Certificate (CLI)

To set the appliance or default certificate, use the following steps.

1. **To view all certificate entries, go to the configuration settings certificates system context, and enter the show command.**

```
hostname:configuration settings certificates system> show
Properties:
    default = auto
System Certificates:
CERT    TYPE SUBJECT                ISSUER                EXPIRES
cert-000 req hostname.us.example.com
cert-001 CA Joe Test CA Joe Test CA 2038-1-19
cert-002 cert hostname.us.example.com Joe Test CA 2038-1-21
```

2. **To set a certificate as the default, use the following commands.**

```
hostname:configuration settings certificates system> set default=cert-000
                        default= cert-000 (uncommitted)
hostname:configuration settings certificates system> commit
```

▼ Uploading Trusted Certificates (BUI)

Use the following steps to upload a trusted certificate.

1. **Go to Configuration > Settings.**
2. **Click the Trusted tab.**
3. **Click the upload icon .**
4. **Browse to the signed certificate and select it.**
5. **Click UPLOAD.**
6. **Repeat steps 3 through 5 for each signed certificate.**

▼ Uploading Trusted Certificates (CLI)

Use the following steps to upload a trusted certificate.

1. **To upload a certificate, enter the configuration settings certificates trusted context, and enter the import command.**

```
hostname:configuration settings certificates trusted> import
("." to end)> -----BEGIN CERTIFICATE-----
("." to end)> MIID0CCAriGAWIBAgIBQDANBgkqhkiG9w0BAQUFADCBmDELMAKGA1UEBhMCVVMx
...
("." to end)> 2ai9ZwREdTkcjcgQDxeHNZCpcHk=
("." to end)> -----END CERTIFICATE-----
("." to end)> .
```

2. **Repeat step 1 for each signed certificate.**
3. **To check the imported certificates, view all certificate entries using the command show.**

```
hostname:configuration settings certificates trusted> show
Properties:
    default = auto
Trusted Certificates:
CERT      TYPE SUBJECT                                ISSUER                                EXPIRES
cert-002  cert hostname.us.example.com    Joe Test CA                          2038-1-21
```

▼ Viewing Trusted Certificate Details (BUI)

To view trusted certificate details, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the Trusted tab.**
3. **Hover over an existing entry and click its information icon .**
4. **When finished, click OK to close the Details window.**

▼ Viewing Trusted Certificate Details (CLI)

To view CSR and certificate details, use the following steps.

1. **To view all certificate entries, go to the configuration settings certificates trusted context, and enter the show command.**

```
hostname:configuration settings certificates system> show
Properties:
    default = auto
Trusted Certificates:
CERT      TYPE SUBJECT                                ISSUER                                EXPIRES
cert-002  cert hostname.us.example.com    Joe Test CA                          2038-1-21
```

2. **To view the details of a certificate, use the following commands.**

```
hostname:configuration settings certificates system> select cert-000
hostname:configuration settings certificates system cert-000> show
Properties:
    uuid = 195071da-66ac-43a6-edfa-bbbd7451f1d5
    subject_commonname = hostname.us.example.com
```



```

issuer_commonname = Joe Test CA
      dns = hostname.us.example.com
      ip = 192.0.2.1
      comment = test certificate
      notbefore = 2014-12-4 00:31:33
      notafter = 2038-01-19 00:31:33
      shalfingerprint = 81:A2:4B:C4:06:A9:14:1E:3E:0B:8A:70:FB:1A:30:45:2D:93:
DD:02
      md5fingerprint = B7:B2:F4:3B:BB:04:8E:11:A2:64:3D:69:BF:8A:79:CC
hostname:configuration settings certificates system cert-000> done

```

▼ Destroying a Trusted Certificate (BUI)

To destroy a trusted certificate, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the Trusted tab.**
3. **Hover over an existing entry and click the trash icon .**
4. **Click DESTROY.**

▼ Destroying a Trusted Certificate (CLI)

To destroy a trusted certificate, use the following steps.

1. **To view all certificate entries, go to the configuration settings certificates trusted context, and enter the show command.**

```

hostname:configuration settings certificates trusted> show
Properties:
      default = auto
Trusted Certificates:
CERT    TYPE SUBJECT                ISSUER                EXPIRES
cert-002 cert hostname.us.example.com Joe Test CA          2038-1-21

```

2. **To destroy a certificate, use the following commands.**

```

hostname:configuration settings certificates system> destroy cert-002
Caution: Destroying a certificate issued by a certificate authority also

```

destroys the associated private key. Re-importing the certificate will not be possible.
Destroy appliance certificate? (Y/N) Y

▼ Assigning a Certificate to a Service (BUI)

To assign a certificate to the LDAP service, use the following steps.

1. **Go to Configuration > Settings.**
2. **Click the Trusted tab.**
3. **From the drop-down menu of system certificates, select the certificate that you want to assign.**
4. **Click the edit icon .**
5. **Select the ldap service from the list of services at the bottom of the page.**

▼ Assigning a Certificate to a Service (CLI)

To assign a certificate to the LDAP service, use the following steps.

1. **To view all certificate entries, go to the configuration settings certificates trusted context, and enter the show command.**

```
hostname:configuration settings certificates trusted> show
Properties:
    default = auto
System Certificates:
CERT      TYPE SUBJECT                                ISSUER                                EXPIRES
cert-002  cert hostname.us.example.com Joe Test CA                          2038-1-21
```

2. **Select the certificate that you want to assign the service to.**

```
hostname:configuration settings certificates trusted> select cert-002
hostname:configuration settings certificates trusted cert-002> set services=ldap
services= ldap (uncommitted)
hostname:configuration settings certificates trusted cert-002> commit
```

Appliance Services

Appliance services are easily managed from the BUI Configuration > Services screen or the CLI configuration services context.

Use the following tasks for viewing and managing appliance services:

- [“Viewing a Service in the BUI” on page 248](#)
- [“Selecting a Service in the CLI” on page 249](#)
- Enabling a Service - [BUI](#), [CLI](#)
- Disabling a Service - [BUI](#), [CLI](#)
- [“Viewing Service States in the CLI” on page 251](#)
- [“Viewing Service Help in the CLI” on page 251](#)
- Setting Service Properties - [BUI](#), [CLI](#)
- Viewing Service Logs - [BUI](#), [CLI](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

For information on configuring an individual service, select one of the services from the following table:

Data Services	Directory Services	System Settings	Remote Access
NFS	NIS	DNS	SSH
iSCSI	LDAP	IPMP	RESTful API
SMB	Active Directory	NTP	HTTPS
FTP	Identity Mapping	Phone Home	
HTTP		Dynamic Routing	
NDMP		Service Tags	
Remote Replication		SMTP	
Shadow Migration		SNMP	
SFTP		Syslog	
SRP		System Identity	

Data Services	Directory Services	System Settings	Remote Access
TFTP		Kerberos	
Virus Scan			

Managing Services

For information about managing appliance services, use the following tasks:

- [“Viewing a Service in the BUI” on page 248](#)
- [“Selecting a Service in the CLI” on page 249](#)
- Enabling a Service - [BUI](#), [CLI](#)
- Disabling a Service - [BUI](#), [CLI](#)
- [“Viewing Service States in the CLI” on page 251](#)
- [“Viewing Service Help in the CLI” on page 251](#)
- Setting Service Properties - [BUI](#), [CLI](#)
- Viewing Service Logs - [BUI](#), [CLI](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Viewing a Service in the BUI

1. Go to **Configuration > Services**.
2. To view or edit the properties for a specific service, hover over the service status icon that is to the left of the service name.

The status icon turns into an arrow icon .

3. Click the arrow icon  to display the properties screen for the selected service.
4. In any of the services screens, you can show a side panel of all services by clicking the small arrow icon to the left of the Services title (near the top left of each screen). Click this icon again to hide the list.

Related Topics

- [“Enabling a Service \(BUI\)” on page 249](#)

- [“Setting Service Properties \(BUI\)” on page 252](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Selecting a Service in the CLI

After you select a service, you can view its state, enable it, disable it, and set its properties.

1. **Go to configuration services.**
2. **Select a service by entering its name. For example, enter `nis`:**

```
hostname:configuration services> nis
hostname:configuration services nis>
```

Related Topics

- [“Enabling a Service \(CLI\)” on page 250](#)
- [“Setting Service Properties \(CLI\)” on page 253](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Enabling a Service (BUI)

Use the following procedure to enable a service that is not online.

1. **Go to Configuration > Services.**
2. **Click the power icon  to bring the service online .**

Related Topics

- [“Disabling a Service \(BUI\)” on page 250](#)
- [“Setting Service Properties \(BUI\)” on page 252](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Enabling a Service (CLI)

Use the following procedure to enable a service that is not online.

1. **Go to configuration services.**
2. **Select a service, then enter the `enable` command to enable a service.**

```
hostname:configuration services> nis
hostname:configuration services nis> enable
```

Related Topics

- [“Disabling a Service \(CLI\)” on page 250](#)
- [“Setting Service Properties \(CLI\)” on page 253](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Disabling a Service (BUI)

Use the following procedure to disable a service that is online.

1. **Go to Configuration > Services.**
2. **Click the power icon  to take the service offline .**

Related Topics

- [“Enabling a Service \(BUI\)” on page 249](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Disabling a Service (CLI)

Use the following procedure to disable a service that is online.

1. **Go to configuration services.**

2. **Select the service, then enter `disable` command to disable it.**

```
hostname:configuration services> nis
hostname:configuration services nis> disable
```

Related Topics

- [“Enabling a Service \(CLI\)” on page 250](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Viewing Service States in the CLI

Use the following procedure to view service states.

1. **Go to `configuration services`.**
2. **Enter the `show` command to list the current state of all services.**
3. **To view the state of an individual service, select the service and then enter `show`.**

```
hostname:configuration services> nis
hostname:configuration services nis> show
Properties:
    <status> = online
    domain = fishworks
    broadcast = true
    ypservers =
```

Related Topics

- [“Enabling a Service \(CLI\)” on page 250](#)
- [“Setting Service Properties \(CLI\)” on page 253](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Viewing Service Help in the CLI

Use the following procedure to display available commands for a service.

1. **Go to configuration services.**
2. **Select the service and enter help.**

```
hostname:configuration services> nis
hostname:configuration services nis> help
Subcommands that are valid in this context:
```

help [topic]	=> Get context-sensitive help. If [topic] is specified, it must be one of "builtins", "commands", "general", "help", "script" or "properties".
show	=> Show information pertinent to the current context
commit	=> Commit current state, including any changes
done	=> Finish operating on "nis"
enable	=> Enable the nis service
disable	=> Disable the nis service
get [prop]	=> Get value for property [prop]. ("help properties" for valid properties.) If [prop] is not specified, returns values for all properties.
set [prop]	=> Set property [prop] to [value]. ("help properties" for valid properties.) For properties taking list values, [value] should be a comma-separated list of values.

Related Topics

- [“Setting Service Properties \(CLI\)” on page 253](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Setting Service Properties (BUI)

The Configuration > Services screens allow you to view and modify the services. The following table describes the icons and buttons in the services screens.

Icon	Description
	Go to the service screen to configure properties and view logs. This button appears when you hover over a service.
	The service is enabled and working normally.
	The service is offline or disabled.
	The service has a problem and requires operator attention.
	Enables or disables the service.
	Restarts the service.
	Enable/disable not available for this service.
	Restarts the currently unavailable service. You must enable the service first).

1. **Go to Configuration > Services.**
2. **Double-click a service.**
3. **Change the properties and then click APPLY.**
To reset properties, click REVERT.

Related Topics

- [“Enabling a Service \(BUI\)” on page 249](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Setting Service Properties (CLI)

Use the following procedure to define properties for a service. Property names are similar to their names in the BUI, but CLI names are usually shorter and sometimes abbreviated.

1. **Go to configuration services.**
2. **Select a service and enter `show` to view the list of properties you can set for that service, along with their current values.**

```
hostname:configuration services> nis
hostname:configuration services nis> show
Properties:
    <status> = online
    domain =
    broadcast = true
    ypservers =
```

3. Use the set command to set the properties.

```
hostname:configuration services nis> set domain="mydomain"
domain = mydomain (uncommitted)
```

4. After setting the properties, enter commit to save and activate the new configuration.

```
hostname:configuration services nis> commit
hostname:configuration services nis> show
Properties:
    <status> = online
    domain = mydomain
    broadcast = true
    ypservers =
```

Related Topics

- [“Enabling a Service \(CLI\)” on page 250](#)
- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Viewing Service Logs (BUI)

Some services provide service logs with information to help you diagnose service issues. If a Logs button exists in the top right of a service screen, that service provides a log. Service logs can provide the times when a service changed state, and the error messages from the service. Log content is specific to each individual service and is subject to change.

- 1. Go to Configuration > Services and double-click a service.**
- 2. Click the Logs button at the top right of a service screen.**

The following are common example messages:

Example Log Message	Description
Executing start method	The service is starting up
Method "start" exited with status 0	The service reported a successful start (0 == success)
Method "refresh" exited with status 0	The service successfully refreshed its configuration based on its service settings
Executing stop method	The service is being shut down
Enabled	The service state was checked to see if it should be started (such as during system boot), and it was found to be in the enabled state
Disabled	The service state was checked to see if it should be started (such as during system boot), and it was found to be in the disabled state

The following log example is from the NTP service:

```
[ Oct 15 21:05:31 Enabled. ]
[ Oct 15 21:07:37 Executing start method (...). ]
[ Oct 15 21:13:38 Method "start" exited with status 0. ]
```

The first log event in the example shows that the system was booted at 21:05. The second entry at 21:07:37 records that the service began startup, which completed at 21:13:38. Due to the nature of NTP and system clock adjustment, this service can take minutes to complete startup, as shown by the log.

Related Topics

- [“List of Available Appliance Services” on page 256](#)
- [“Required Service Ports” on page 258](#)

▼ Viewing Service Logs (CLI)

- You cannot view service logs from the CLI. Use the BUI as described in [“Viewing Service Logs \(BUI\)” on page 254](#).

List of Available Appliance Services

This section lists the available appliance services, along with short descriptions and port information. Certain services are always on and cannot be disabled, as described in the following table.

TABLE 46 Data Services

Service	Description	Ports Used
NFS	Filesystem access via the NFSv3, NFSv4.0, and NFS v4.1 protocols	111 and 2049
iSCSI	LUN access via the iSCSI protocol	3260 and 3205
SMB	Filesystem access via the SMB protocol	SMB-over-NetBIOS 139 SMB-over-TCP 445 NetBIOS Datagram 138 NetBIOS Name Service 137
FTP	Filesystem access via the FTP protocol	21
HTTP	Filesystem access via the HTTP protocol	80
NDMP	NDMP host service	10000
Remote Replication	Remote replication	216 and 217
Shadow Migration	Shadow data migration	
SFTP	Filesystem access via the SFTP protocol	218
SRP	Block access via the SRP protocol	
TFTP	Filesystem access via the TFTP protocol	
Virus Scan	Filesystem virus scanning	

Note - UIDs and GIDs from 0 to 99 are reserved by the operating system vendor for use in future applications. Their use by end-system users or vendors of layered products is not supported and may cause security-related issues with future applications.

TABLE 47 Directory Services

Service	Description	Ports Used
NIS	Authenticate users and groups from an NIS service.	

Service	Description	Ports Used
LDAP	Authenticate users and groups from an LDAP directory.	389
Active Directory	Authenticate users with a Microsoft Active Directory Server. Note - Once enabled, this becomes an always-on service, and cannot be disabled.	
Identity Mapping	Map between Windows entities and UNIX IDs. Note - Always-on service, cannot be disabled.	

TABLE 48 Service Settings

Service	Description	Ports Used
DNS	Domain name service client Note - Always-on service, cannot be disabled.	53
Dynamic Routing	RIP and RIPng dynamic routing protocols	
IPMP	IP Multipathing for IP fail-over Note - Always-on service, cannot be disabled.	
Kerberos	Kerberos authentication	88
NTP	Network time protocol client	
Phone Home	Product registration and support configuration	8000 (It depends on the port opened up in the proxy)
Service Tags	Product inventory support	6481
SMTP	Configure outgoing mail server Note - Always-on service, cannot be disabled.	
SNMP	SNMP for sending traps on alerts and serving appliance status information	
Syslog	Syslog Relay for sending syslog messages on alerts and forwarding service syslog messages	
System Identity	System name and location Note - Always-on service, cannot be disabled.	

TABLE 49 Remote Access Services

Service	Description	Ports Used
SSH	SSH for CLI access	22
REST	RESTful API	

Service	Description	Ports Used
Note - Always-on service, cannot be disabled.		

Required Service Ports

To provide security on a network, you can deploy firewalls within your network architecture. Port numbers are used for creating firewall rules and to uniquely identify a transaction over a network by specifying the host and the service.

The following list shows the minimum ports required for creating firewall rules that allow full functionality of the appliance:

Inbound Ports

- icmp/0-65535 (PING)
- tcp/1920 (EM)
- tcp/215 (BUI)
- tcp/22 (SSH)
- udp/161 (SNMP)

Outbound Ports

- tcp/80 (WEB)
- tcp/443 (SSL WEB)

Note - An outbound port of tcp/443 is used for sending Phone Home messages, uploading support bundles, and update notifications. For replication, use Generic Routing Encapsulation (GRE) tunnels when possible. This lets traffic run on the back end interfaces and avoid the firewall where traffic could be slowed. If GRE tunnels are not available on the NFS core, you must run replication over the front end interface. In this case, port 216 and port 217 must also be open.

Configuring Services

For information about configuring a service, select one of the services from the following table:

Data Services	Directory Services	System Settings	Remote Access
NFS	NIS	DNS	SSH
iSCSI	LDAP	IPMP	RESTful API

Data Services	Directory Services	System Settings	Remote Access
SMB	Active Directory	NTP	HTTPS
FTP	Identity Mapping	Phone Home	
HTTP		Dynamic Routing	
NDMP		Service Tags	
Remote Replication		SMTP	
Shadow Migration		SNMP	
SFTP		Syslog	
SRP		System Identity	
TFTP		Kerberos	
Virus Scan			

Related Topics

- [“List of Available Appliance Services” on page 256](#)

Active Directory Configuration

The Active Directory service provides access to a Microsoft Active Directory database, which stores information about users, groups, shares, and other shared objects. This service has two modes: domain and workgroup mode, which dictate how SMB users are authenticated. When operating in domain mode, SMB clients are authenticated through the AD domain controller. In workgroup mode, SMB clients are authenticated locally as local users. See [“Configuring Users” on page 202](#) for more information on local users.

To configure Active Directory, see the following sections:

- [“Joining an AD Domain \(BUI\)” on page 259](#)
- [“Joining an AD Workgroup \(BUI\)” on page 260](#)
- [“Configuring Active Directory \(CLI\)” on page 260](#)
- [“Active Directory Join Domain” on page 262](#)
- [“Active Directory Domains and Workgroups” on page 263](#)
- [“Active Directory Windows Server Support” on page 264](#)

▼ Joining an AD Domain (BUI)

1. (Optional) Configure an Active Directory site in the SMB context.
2. (Optional) Configure a preferred domain controller in the SMB context.

3. **Enable NTP, or ensure that the clocks of the appliance and domain controller are synchronized to within five minutes.**
4. **Ensure that your DNS infrastructure correctly delegates to the Active Directory domain, or add your domain controller's IP address as an additional name server in the DNS context.**
5. **Go to Configuration > Services > Active Directory and click Join Domain.**
6. **Configure the Active Directory domain, administrative user name, and administrative password.**
7. **Click APPLY to commit the configuration.**

Related Topics

- [“Joining an AD Workgroup \(BUI\)” on page 260](#)
- [“Configuring Active Directory \(CLI\)” on page 260](#)
- [“Active Directory Join Domain” on page 262](#)
- [“Active Directory Domains and Workgroups” on page 263](#)
- [“Active Directory Windows Server Support” on page 264](#)

▼ **Joining an AD Workgroup (BUI)**

1. **Go to Configuration > Services > Active Directory and click Join Workgroup.**
2. **Enter the Windows workgroup name.**
3. **Click APPLY to commit the configuration.**

Related Topics

- [“Joining an AD Domain \(BUI\)” on page 259](#)
- [“Configuring Active Directory \(CLI\)” on page 260](#)
- [“Active Directory Join Domain” on page 262](#)
- [“Active Directory Domains and Workgroups” on page 263](#)
- [“Active Directory Windows Server Support” on page 264](#)

▼ **Configuring Active Directory (CLI)**

1. **Go to configuration services ad.**


```
hostname:> configuration services ad
```

2. To view an existing configuration, enter show.

```
hostname:configuration services ad> show
Properties:
    <status> = online
    mode = domain
    domain = eng.test.com
```

```
Children:
    domain => Join an Active Directory domain
    workgroup => Join a Windows workgroup
```

Observe that the appliance is currently operating in the domain `eng.test.com`.

3. To leave the domain mode and join a Windows workgroup, enter the following commands:

```
hostname:configuration services ad> workgroup
hostname:configuration services ad workgroup> set workgroup=WORKGROUP
hostname:configuration services ad workgroup> commit
hostname:configuration services ad workgroup> done
hostname:configuration services ad> show
Properties:
    <status> = disabled
    mode = workgroup
    workgroup = WORKGROUP
```

4. To configure the site and preferred domain controller in preparation for joining another domain, enter the following commands:

```
hostname:configuration services ad> done
hostname:> configuration services smb
hostname:configuration services smb> set ads_site=sf
hostname:configuration services smb> set pdc=192.0.2.21
hostname:configuration services smb> commit
hostname:configuration services smb> show
Properties:
    <status> = online
    lmauth_level = 4
    pdc = 192.168.3.21
    ads_site = sf
hostname:configuration services smb> done
```

5. To join the new domain after the properties are configured, enter the following commands.

When joining an AD domain, you must set the user and password each time you commit the node.

```
hostname:> configuration services ad
hostname:configuration services ad> domain
hostname:configuration services ad domain> set domain=example.com
hostname:configuration services ad domain> set user=Administrator
hostname:configuration services ad domain> set password=*****
hostname:configuration services ad domain> commit
hostname:configuration services ad domain> done
hostname:configuration services ad> show
Properties:
                                <status> = online
                                mode = domain
                                domain = example.com
```

Related Topics

- [“Joining an AD Domain \(BUI\)” on page 259](#)
- [“Joining an AD Workgroup \(BUI\)” on page 260](#)
- [“Active Directory Join Domain” on page 262](#)
- [“Active Directory Domains and Workgroups” on page 263](#)
- [“Active Directory Windows Server Support” on page 264](#)

Active Directory Join Domain

If an account does not already exist in Active Directory by default, a machine trust account for the system is automatically created in the default container for computer accounts (cn=Computers) as part of the domain join operation. The following users are allowed to perform domain join:

- **Domain administrator** - Can join any number of systems to the domain with machine trust accounts placed in any containers.
- **Delegated administrator with authority over one or more Organizational Units** - Can join any number of systems to a domain with machine account location designated in the Organizational Units they are responsible for.
- **Normal user with machine accounts pre-staged by administrator** - Can join a system to the domain as pre-authorized by an administrator.
- **Normal user** - Normally authorized to join a limited number of systems.

The following properties for joining an Active Directory domain are available:

- **Active Directory Domain** - The fully-qualified name or NetBIOS name of an Active Directory domain

- **User** - An AD user who has credentials to create a computer account in Active Directory
- **Password** - The administrative user's password
- **Organizational Unit** - Specifies an alternative organizational unit in which the system's machine trust account will be created. The organizational unit is specified as a comma-separated list of one or more name-value pairs using the domain-relative distinguished name (DN) format, for example, ou=innerOU,ou=outerOU.
- **Use Pre-created Account** - If the system's account exists and the specified Organizational Unit is not the one that the account is in, use the pre-created account.

Related Topics

- [“Joining an AD Domain \(BUI\)” on page 259](#)
- [“Joining an AD Workgroup \(BUI\)” on page 260](#)
- [“Configuring Active Directory \(CLI\)” on page 260](#)
- [“Active Directory Domains and Workgroups” on page 263](#)
- [“Active Directory Windows Server Support” on page 264](#)

Active Directory Domains and Workgroups

The configurable property for joining a workgroup is Windows Workgroup.

Instead of enabling and disabling the service directly, the service is modified by joining a domain or a workgroup. Joining a domain involves creating an account for the appliance in the given Active Directory domain. The account name can be a maximum of 15 characters, and must be unique to other names registered within the Active Directory domain. Otherwise, conflicts may occur with similarly named appliances and cause issues with functionality. After the computer account has been established, the appliance can securely query the database for information about users, groups, and shares.

Joining a workgroup implicitly leaves an Active Directory domain, and SMB clients who are stored in the Active Directory database will be unable to connect to shares.

Active Directory LDAP Signing

There is no configuration option for LDAP signing, as that option is negotiated automatically when communicating with a domain controller. LDAP signing operates on communication between the storage appliance and the domain controller, whereas SMB signing operates on communication between the SMB clients and the storage appliance.

Related Topics

- “Joining an AD Domain (BUI)” on page 259
- “Joining an AD Workgroup (BUI)” on page 260
- “Configuring Active Directory (CLI)” on page 260
- “Active Directory Join Domain” on page 262
- “Active Directory Windows Server Support” on page 264

Active Directory Windows Server Support

Windows Server 2012 is fully supported in software version 2011.1.5 (and later).

TABLE 50 Active Directory Windows Server 2008 Support

Windows Version	Supported Software Versions	Workarounds
Windows Server 2003	All	None
Windows Server 2008 SP1	2009.Q2 3.1 and earlier	Apply hotfix for KB957441 as needed, see Section B.
	2009.Q2 4.0 - 2011.1.1	Must apply hotfix for KB951191 and apply hotfix for KB957441 as needed, see Sections A and B.
	2011.1.2 and later	Must apply hotfix for KB951191, see Section A.
Windows Server 2008 SP2	2009.Q2 4.0 - 2011.1.1	See Section C.
	2011.1.2 and later	None
Windows Server 2008 R2	2009.Q2 4.0 - 2011.1.1	See Section C.
	2011.1.2 and later	None

Active Directory Windows Server 2008 Support Section A: Kerberos issue (KB951191)

- If you upgrade to 2009.Q2.4.0 or later and your Windows 2008 domain controller is running Windows Server 2008 SP2 or R2, no action is required.
- If you upgrade to 2009.Q2.4.0 or later and your Windows 2008 domain controller is running Windows Server 2008 SP1, you must apply the hotfix described in KB951191 or install Windows 2008 SP2.

Active Directory Windows Server 2008 Support Section B: NTLMv2 issue (KB957441)

- The following applies only if your appliance is running a software version prior to 2011.1.2:
- If your Domain Controller is running Windows Server 2008 SP1 you should also apply the hotfix for <http://support.microsoft.com/kb/957441/> (<http://support.microsoft.com/kb/957441/>) which resolves an NTLMv2 issue that prevents the appliance from joining the domain with its default LMCompatibilityLevel setting.

- If the LMCompatibilityLevel on the Windows 2008 SP1 domain controller is set to 5, this hot fix must be installed. After applying the hotfix you must create and set a new registry key as described in KB957441.
- If you upgrade to 2011.1.2 or later, you do not need the hotfix mentioned above.

Active Directory Windows Server 2008 Support Section C: Note on NTLMv2

- The following applies only if your appliance is running a software version prior to 2011.1.2: If your Domain Controller is running Windows Server 2008 SP2 or R2 you do not need to apply the hotfix but you must apply the registry setting as described in KB957441.
- If you upgrade to 2011.1.2 or later, no action is required.

Related Topics

- [“Joining an AD Domain \(BUI\)” on page 259](#)
- [“Joining an AD Workgroup \(BUI\)” on page 260](#)
- [“Configuring Active Directory \(CLI\)” on page 260](#)
- [“Active Directory Join Domain” on page 262](#)
- [“Active Directory Domains and Workgroups” on page 263](#)

DNS Configuration

The DNS (Domain Name Service) client provides the ability to resolve IP addresses to hostnames and vice versa, and can be enabled or disabled. Optionally, secondary hostname resolution via NIS and/or LDAP, if configured and enabled, may be requested for hostnames and addresses that cannot be resolved using DNS. Hostname resolution is used throughout the appliance user interfaces, including in Logs to indicate the location from which a user performed an auditable action and in Analytics to provide statistics on a per-client basis.

To configure and manage DNS, use these tasks:

- Configuring DNS - [BUI](#), [CLI](#)
- [“Testing Hostname Resolution \(CLI\)” on page 267](#)
- Adding a DNS Server - [BUI](#), [CLI](#)
- Viewing DNS Server Status - [BUI](#), [CLI](#)

To understand DNS usage for the appliance, use these topics:

- [“DNS Properties and Logs” on page 270](#)
- [“Active Directory and DNS” on page 271](#)

- [“Non-DNS Resolution” on page 271](#)
- [“DNS-Less Operation” on page 272](#)

▼ Configuring DNS (BUI)

DNS is usually configured during initial configuration, as described in [“Performing Initial Configuration \(BUI\)” in Oracle ZFS Storage Appliance Installation Guide](#). To change your DNS settings after initial configuration, use the following procedure.

1. Go to Configuration > Services > DNS.

2. Under General Settings, set the following properties:

For more information about DNS properties, see [“DNS Properties and Logs” on page 270](#).

- **DNS Domain** - Enter a domain name.
- **DNS Search Domain(s)** - Click the add icon  to add search domain(s). To remove a domain, click the remove icon  beside it.
- **Allow IPv4 non-DNS resolution** - Check this box to enable IPv4 non-DNS resolution. See [“Non-DNS Resolution” on page 271](#).
- **Allow IPv6 non-DNS resolution** - Check this box to enable IPv6 non-DNS resolution. See [“Non-DNS Resolution” on page 271](#).

3. Click APPLY.

Related Topics

- [“Adding a DNS Server \(BUI\)” on page 267](#)

▼ Configuring DNS (CLI)

DNS is usually configured during initial configuration, as described in [“Performing Initial Configuration \(CLI\)” in Oracle ZFS Storage Appliance Installation Guide](#). To change your DNS settings after initial configuration, use the following procedure.

1. Go to configuration services dns and then enter show.

```
hostname:> configuration services dns
hostname:configuration services dns> show
Properties:
```

```

<status> = online
  domain = example.com
  servers = 192.0.2.254
  search =
allow_alternate_v4 = false
allow_alternate_v6 = false

```

2. Set the domain, servers, and search domain, and enable or disable non-DNS resolution.

For more information, see [“DNS Properties and Logs” on page 270](#) and [“Non-DNS Resolution” on page 271](#).

```

hostname:configuration services dns> set domain=example.com
      domain = example.com (uncommitted)
hostname:configuration services dns> set servers=192.0.2.253
      servers = 192.0.2.253 (uncommitted)
hostname:configuration services dns> set search=example.com
      search = example.com (uncommitted)
hostname:configuration services dns> set allow_alternate_v4=true
      allow_alternate_v4 = true (uncommitted)

```

3. Enter commit.

```
hostname:configuration services dns> commit
```

Related Topics

- [“Adding a DNS Server \(CLI\)” on page 268](#)

Testing Hostname Resolution (CLI)

The CLI includes built-ins for `nslookup` and `getent hosts`, which can be used to test that hostname resolution is working:

```

hostname:> nslookup deimos
198.51.100.1  deimos.sf.fishworks.com
hostname:> getent hosts deimos
198.51.100.1  deimos.sf.fishworks.com

```

▼ Adding a DNS Server (BUI)

1. Go to Configuration > Services > DNS.

2. Click the add icon  beside DNS Servers.
3. In the New DNS Server dialog box, enter the server IP address.
4. Click ADD.

A query is sent to the affected DNS servers to validate the changes. If a valid response is not received, a message appears to confirm the settings. You may confirm your changes regardless of whether the server is valid.

Related Topics

- [“Viewing DNS Server Status \(BUI\)” on page 269](#)

▼ Adding a DNS Server (CLI)

1. Go to `configuration services dns` and then enter `create`.

```
hostname:> configuration services dns
hostname:configuration services dns> create
```

2. Enter `show`.

```
hostname:configuration services server (uncommitted)> show
Properties:
    address = (unset)
    status = unavailable
    rtt = unavailable
    err_msg =
```

3. Enter `set address=` and the server address.

```
hostname:configuration services server (uncommitted)> set address=192.0.2.254
address = 192.0.2.254 (uncommitted)
```

4. Enter `show`.

```
hostname:configuration services server (uncommitted)> show
Properties:
    address = 192.0.2.254
    status = online
    rtt = 1.812ms
    err_msg =
```

5. Enter `commit`.

A query is sent to the affected DNS servers to validate the changes. If a valid response is not received, a message appears to confirm the settings. You may confirm your changes regardless of whether the server is valid.

```
hostname:configuration services server (uncommitted)> commit
```

Related Topics

- [“Viewing DNS Server Status \(CLI\)” on page 269](#)

▼ Viewing DNS Server Status (BUI)

Details about the DNS servers are displayed beside each entry in the BUI. A status indicator shows if the server status is online, offline, or unknown. The RTT column indicates the round-trip time, in milliseconds, to receive a valid response.

1. **Go to Configuration > Services > DNS.**
2. **Under DNS Servers, check the status indicator beside each server entry:**
 - **Green icon**  - Online
 - **Amber icon**  - Offline
 - **Gray icon**  - Unknown

▼ Viewing DNS Server Status (CLI)

Select a DNS server to view its properties. The status property indicates if the server status is online, offline, or unknown. The rtt property indicates the round-trip time, in milliseconds, to receive a valid response. If the server status is offline, the err_msg property displays the reason, for example, Connection timed out.

1. **Go to configuration services dns and then enter show.**

```
hostname:> configuration services dns
hostname:configuration services dns> show
SERVER      STATUS    ADDRESS
server-000  online    198.51.100.1
server-001  offline   198.51.100.2
```

2. **Select the server for which you want to view its status.**

```
hostname:configuration services dns> select server-000
```

3. Enter show.

```
hostname:configuration services server-000> show
Properties:
    address = 198.51.100.1
    status = online
    rtt = 1.768ms
    err_msg =
```

DNS Properties and Logs

The configurable properties for the DNS client include a base domain name and a list of servers, specified by IP address. You must supply a domain name and at least one server address; the server must be capable of returning an NS (NameServer) record for the domain you specify, although it need not itself be authoritative for that domain.

TABLE 51 DNS Properties

Property	Description
DNS Domain	Domain name to search first when performing partial hostname lookups.
DNS Server(s)	One or more DNS servers. IP addresses must be used.
DNS Search Domain(s)	List of up to four domains to be searched for after the Active Directory domain, the deprecated Active Directory search domain, and the specified DNS domain.
Allow IPv4 non-DNS resolution	IPv4 addresses may be resolved to hostnames, and hostnames to IPv4 addresses, using NIS and/or LDAP if configured and enabled.
Allow IPv6 non-DNS resolution	IPv4 and IPv6 addresses may be resolved to hostnames, and hostnames to IPv4 and IPv6 addresses, using NIS and/or LDAP if configured and enabled.

Changing services properties is documented in [“Setting Service Properties \(BUI\)” on page 252](#) and [“Setting Service Properties \(CLI\)” on page 253](#). The CLI property names are shorter versions of those listed above.

TABLE 52 DNS Logs

Log	Description
network-dns-client:default	Logs the DNS service events

Related Topics

- [“Active Directory and DNS” on page 271](#)

Active Directory and DNS

If you plan to use Active Directory, the servers must be able to resolve hostname and server records in the Active Directory portion of the domain namespace. For example, if your appliance resides in the domain example.com and the Active Directory portion of the namespace is redmond.example.com, your nameservers must be able to reach an authoritative server for example.com, and they must provide delegation for the domain redmond.example.com to one or more Active Directory servers serving that domain. These are requirements imposed by Active Directory, not the appliance itself. If they are not satisfied, you will be unable to join an Active Directory domain.

Note - With OS8.6.0 (and later), if the primary DNS domain suffix does not match the DNS name of the Active Directory, the configuration results in a disjoint namespace. If you do not want a disjoint namespace, ensure that the DNS domain and the Active Directory domain are the same.

Related Topics

- [“DNS Configuration” on page 265](#)
- [“Active Directory Configuration” on page 259](#)

Non-DNS Resolution

DNS is a standard, enterprise-grade, highly-scalable and reliable mechanism for mapping between hostnames and IP addresses. Use of working DNS servers is a best practice and will generally yield the best results. In some environments, there may be a subset of hosts that can be resolved only in NIS or LDAP maps. If this is the case in your environment, enable non-DNS host resolution and configure the appropriate directory service(s). If LDAP is used for host resolution, the hosts map must be located at the standard DN in your database: ou=Hosts, (Base DN), and must use the standard schema. When this mode is used with NFS sharing by netgroups, it may be necessary for client systems to use the same hostname resolution mechanism configured on the appliance, or NFS sharing exceptions may not work correctly.

When non-DNS host resolution is enabled, DNS will still be used. Only if an address or hostname cannot be resolved using DNS will NIS (if enabled) and then LDAP (if enabled) be used to resolve the name or address. This can have confusing and seemingly inconsistent results. You can validate host resolution results using the `getent` CLI command described above.

Use of these options is strongly discouraged.

DNS-Less Operation

If the appliance is unable to access any DNS servers from its installed location in the network, you may elect to operate without DNS by supplying the server address 127.0.0.1. To operate without DNS:

- **BUI:** Go to Configuration > Services > DNS. In the field for DNS Server(s), enter **127.0.0.1**.
- **CLI:** Go to `configuration services dns` and enter `show`. Enter `set servers=127.0.0.1`, and then enter `commit`.

Use of this mode is strongly discouraged, because several features will not work correctly, including:

- Analytics will be unable to resolve client addresses to host names.
- The Active Directory feature will not function (you will be unable to join a domain).
- Use of SSL-protected LDAP will not work properly with certificates containing host names.
- Alert and threshold actions that involve sending e-mail can only be sent to mail servers on an attached subnet, and all addresses must be specified using the mail server's IP address.
- Some operations may take longer than normal due to hostname resolution timeouts.

These limitations may be partially mitigated by using an alternate host resolution service; see [“Non-DNS Resolution” on page 271](#).

Related Topics

- Enabling a Service [BUI](#), [CLI](#)
- Disabling a Service [BUI](#), [CLI](#)

Dynamic Routing Configuration

The Routing Information Protocol (RIP) is a distance-vector dynamic routing protocol that is used by the appliance to automatically configure optimal routes based on messages received from other RIP-enabled on-link hosts (typically routers). The appliance supports both RIPv1 and RIPv2 for IPv4, and RIPng for IPv6.

Routes that are configured via these protocols are marked as type "dynamic" in the routing table. RIP and RIPng listen on UDP ports 520 and 521 respectively.

TABLE 53 Dynamic Routing

Log	Description
network-routing-route:default	Logs RIP service events
network-routing-ripng:quagga	Logs RIPng service events

FTP Configuration

The FTP (File Transfer Protocol) service allows filesystem access from FTP clients. Anonymous logins are not allowed, users must authenticate with whichever name service is configured in Services.

FTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see [“Kerberos Configuration” on page 293](#).

For added security when configuring FTP, you can specify the SSL/TLS versions and ciphers, as described in [“FTP Properties” on page 273](#).

To configure FTP, use the following sections:

- [“Adding FTP Access to a Share \(BUI\)” on page 273](#)
- [“FTP Properties” on page 273](#)
- [“FTP Logs” on page 275](#)

▼ Adding FTP Access to a Share (BUI)

1. Go to Configuration > Services.
2. Ensure that the FTP service is enabled and online. If not, enable the service.
3. Select or add a share in the Shares screen.
4. Click the Protocols tab, and check that FTP access is enabled.
5. (Optional) Set the Share mode access to Read only or Read/write.

Related Topics

- [“FTP Properties” on page 273](#)
- [“FTP Logs” on page 275](#)

FTP Properties

TABLE 54 FTP General Properties

Property	Description
Port for incoming connections	The port on which FTP listens. The default is 21.
Maximum # of connections ("0" for unlimited)	This is the maximum number of concurrent FTP connections. Set this to cover the anticipated number

Property	Description
	of concurrent users. By default this is 30, since each connection creates a system process and allowing too many (thousands) could constitute a DoS attack.
Turn on delay engine to prevent timing attacks	This inserts small delays during authentication to fool attempts at user name guessing via timing measurements. Turning this on will improve security.
Default login root	The FTP login location. The default is "/" and points to the top of the shares hierarchy. All users will be logged into this location after successfully authenticating with the FTP service.
Logging level	The verbosity of the proftpd log.
Permissions to mask from newly created files and directories	File permissions to remove when files are created. Group and world write are masked by default, to prevent recent uploads from being writeable by everyone.

TABLE 55 FTP Security Properties

Property	Description
Enable SSL/TLS	Allow SSL/TLS encrypted FTP connections. This will ensure that the FTP transaction is encrypted. The default is disabled.
SSL/TLS versions and ciphers	SSL/TLS protocol versions and ciphers for FTP connections. The defaults are TLSv1.1, TLSv1.2 and their associated ciphers. TLSv1.0 is not enabled by default due to security concerns, but it can be enabled for backward compatibility. In the BUI, the list of available ciphers changes based on the selected versions. Some selected SSL/TLS protocol versions and/or ciphers are removed after a software upgrade if they are no longer supported. To avoid service unavailability, keep the default settings unless otherwise needed or as instructed by Oracle Support.
Port for incoming SSL/TLS connections	The port that the SSL/TLS encrypted FTP service listens on. The default is 21.
Permit root login	Allow FTP logins for the root user. This is off by default, since FTP authentication is plain text which poses a security risk from network sniffing attack.
Maximum # of allowable login attempts	The number of failed login attempts before an FTP connection is disconnected, and the user must reconnect to try again. The default is 3.
Permit foreign data connection addresses	Permits foreign FTP connections to enable direct transfer of files between FTP servers. This property is off by default.

Related Topics

- [“Adding FTP Access to a Share \(BUI\)” on page 273](#)
- [“FTP Logs” on page 275](#)

FTP Logs

TABLE 56 FTP Logs

Log	Description
proftpd	Logs FTP events, including successful logins and unsuccessful login attempts.
proftpd_xfer	File transfer log.
proftpd_tls	Logs FTP events related to SSL/TLS encryption.

Related Topics

- [“Adding FTP Access to a Share \(BUI\)” on page 273](#)
- [“FTP Properties” on page 273](#)

HTTP Configuration

The HTTP service provides access to filesystems using the HTTP WebDAV (Web based Distributed Authoring and Versioning) protocol. This service allows clients to access shared filesystems through a web browser, or as a local filesystem if supported by the client software. The URL to access these HTTP shares has the following format:

- `http://hostname/shares/mountpoint/share_name`

HTTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see [“Kerberos Configuration” on page 293](#).

For added security when configuring HTTP, you can specify the SSL/TLS versions and ciphers, as described in [“HTTP Properties and Logs” on page 276](#).

To configure HTTP, see the following sections:

- [“Adding HTTP Access to a Share \(BUI\)” on page 276](#)
- [“HTTP Properties and Logs” on page 276](#)
- [“HTTP Authentication and Access Control” on page 278](#)
- [“Object Store Configuration” on page 278](#)

▼ Adding HTTP Access to a Share (BUI)

1. Go to Configuration > Services.
2. Check that the HTTP service is enabled and online. If not, enable the service.
3. Select or add a share in the Shares screen.
4. Click the Protocols tab, and check that HTTP access is enabled.
5. (Optional) Set the Share mode access to Read only or Read/write.
For HTTP Object Store, set the access to Read/write.

Related Topics

- [“HTTP Properties and Logs” on page 276](#)
- [“HTTP Authentication and Access Control” on page 278](#)
- [“Object Store Configuration” on page 278](#)

HTTP Properties and Logs

TABLE 57 HTTP General Properties

BUI Label	CLI Property	Description
N/A	status	Read-only property showing the status of the HTTP service
Protocols	protocols	Select which access methods to support: HTTP, HTTPS, or both.
HTTP port (for incoming connections)	listen_port	HTTP port. The default is 80.
HTTPS port (for incoming secure connections)	https_port	Secure HTTP port. The default is 443.

TABLE 58 HTTP Security Properties

BUI Label	CLI Property	Description
SSL/TLS versions	tls_version	SSL/TLS protocol versions for HTTP connections. The default TLS versions are TLSv1.1, TLSv1.2 and their associated ciphers. TLSv1.0 is not enabled by default due to security concerns, but it can be enabled for backward compatibility. In the BUI, the list of available ciphers changes based on the selected versions. Some

BUI Label	CLI Property	Description
		selected SSL/TLS protocol versions or ciphers are removed after a software upgrade if they are no longer supported. To avoid service unavailability, keep the default settings unless otherwise needed or as instructed by Oracle Support.
List of ciphers	ciphers	List of ciphers for HTTP connections. In the BUI, the list of available ciphers changes based on the selected versions. Some selected SSL/TLS ciphers are removed after a software upgrade if they are no longer supported. To avoid service unavailability, keep the default settings unless otherwise needed or as instructed by Oracle Support.

TABLE 59 Object Store Properties

BUI Label	CLI Property	Description
Enable Object Store	objectstore_enabled	When selected, enables the HTTP Object Store feature.
Default Path	objectstore_default_path	Sets the location used when a user does not set one.
Use OpenStack Identity Service	os_ids_on	When selected, enables properties required for communication with a keystone server: Authentication URI, Role, Tenant, User, Password.
Authentication URI	os_ids_auth_server	Identity Service URI, for example: <code>http://keystone:5000/v2.0</code> .
Role	os_ids_role	OpenStack Identity Service type.
Tenant	os_ids_tenant	OpenStack Identity Service tenant name.
User	os_ids_user	OpenStack Identity Service user's name.
Password	os_ids_password	OpenStack Identity Service user's password.

TABLE 60 HTTP WebDAV Properties

BUI Label	CLI Property	Description
Enable WebDAV	webdav_enabled	When selected, enables the HTTP WebDAV feature.
Require client login	require_login	Clients must authenticate before share access is allowed, and files they create will have their ownership. If this property is not set, files created will be owned by the HTTP service with user nobody. See “HTTP Authentication and Access Control” on page 278.

HTTP Authentication and Access Control

If the "Require client login" option is enabled, the appliance will deny access to clients that do not supply valid authentication credentials for a local user, a NIS user, or an LDAP user. Active Directory authentication is not supported.

Only basic HTTP authentication is supported. Note that unless HTTPS is being used, this transmits the username and password unencrypted, which may not be appropriate for all environments.

Normally, authenticated users have the same permissions with HTTP that they would have with NFS or FTP. Files and directories created by an authenticated user will be owned by that user, as viewed by other protocols. Privileged users (those having a UID less than 100) will be treated as nobody for the purposes of access control. Files created by privileged users will be owned by nobody.

If the "Require client login" option is disabled, the appliance will not try to authenticate clients (even if they do supply credentials). Newly created files are owned by nobody, and all users are treated as nobody for the purposes of access control.

Regardless of authentication, no permissions are masked from created files and directories. Created files have UNIX permissions 666 (readable and writable by everyone), and created directories have UNIX permissions 777 (readable, writable, and executable by everyone).

Related Topics

- [“Adding HTTP Access to a Share \(BUI\)” on page 276](#)
- [“HTTP Properties and Logs” on page 276](#)
- [“Object Store Configuration” on page 278](#)

Object Store Configuration

The object store service enables an appliance to save data as storage objects into the Oracle ZFS filesystem through the OpenStack Object Storage protocol.

Enabling the Object Store Service

To enable the object store service, go to **Configuration > Services > HTTP** and select the check box for **Enable Object Store**. Before applying the change, configure the object store properties.

After enabling and configuring the object store service, enable the feature on individual filesystems by going to **Shares > Filesystems**. Double-click on a filesystem to view its details,

and then select the **Protocols** tab. In the HTTP section, for the Object store mode option, select Read/write to enable the feature for the filesystem.

Object Store Properties

Once the object store service is enabled, you can configure it with the object store properties, as shown in [Table 59, “Object Store Properties,” on page 277](#).

If the OpenStack Identity Service option was selected when you applied your changes for the HTTP service, a test connection is made to the specified server with the provided properties information. If the connection fails, a dialog box opens, and you can either apply or cancel all changes to the HTTP service.

See Also

- [“Adding HTTP Access to a Share \(BUI\)” on page 276](#)
- [“HTTP Properties and Logs” on page 276](#)
- [“HTTP Authentication and Access Control” on page 278](#)

HTTPS Configuration

The HTTPS service provides the ability to manage the appliance using the HTTPS protocol. This service allows clients to manage the connection to the appliance BUI and the RESTful API service.

For added security when configuring HTTPS, you can specify the SSL/TLS versions and ciphers, as described in [“HTTPS Properties and Logs” on page 279](#).

HTTPS Properties and Logs

TABLE 61 HTTPS Security Properties

Property	Description
SSL/TLS versions and ciphers	SSL/TLS protocol versions and ciphers for HTTPS connections to the appliance BUI and the RESTful API service. The defaults are TLSv1.1, TLSv1.2 and their associated ciphers. TLSv1.0 is not enabled by default due to security concerns, but it can be enabled for backward compatibility. In the BUI, the list of available ciphers changes based on the selected versions. Some selected

Property	Description
	SSL/TLS protocol versions and/or ciphers are removed after software upgrades if they are no longer supported. To avoid service unavailability, keep the default settings unless otherwise needed or as instructed by Oracle Support.

TABLE 62 HTTPS Logs

Log	Description
appliance-kit-http:default	HTTPS service log

Identity Mapping Configuration

Identity mapping allows you to associate Windows and UNIX identities, thereby allowing an SMB client and an NFS client access to the same set of files. The identity mapping service manages Windows and UNIX user identities simultaneously by creating and maintaining a database of mappings between UNIX user identifiers (UIDs) and group identifiers (GIDs), and Windows security identifiers (SIDs).

To manage identity mapping, use these tasks:

- Configuring Identity Mapping - [BUI](#), [CLI](#)
- Creating a Mapping Rule - [BUI](#), [CLI](#)
- “[Viewing a Mapping \(BUI\)](#)” on page 287
- Flushing Mappings from the Cache - [BUI](#), [CLI](#)

To understand identity mapping, use these topics:

- “[Identity Mapping Best Practices](#)” on page 289
- “[Identity Mapping Concepts](#)” on page 289
- “[Cached and Ephemeral Mappings](#)” on page 290
- “[Identity Mapping Case Sensitivity](#)” on page 290
- “[Mapping Rule Directional Symbols](#)” on page 291

▼ Configuring Identity Mapping (BUI)

Use the following procedure to configure identity mapping.

Before You Begin Ensure that you are joined to at least one Active Directory domain. For information about active directories, see “[Active Directory Configuration](#)” on page 259.

1. **Go to Configuration > Services > Identity Mapping > Properties.**
2. **Select one of the following mapping modes.**
 - **Rule-based**
 - **Directory-based** - Set all of the following attributes.
 - **AD Attribute - UNIX User Name** - Name in the Active Directory database of the equivalent UNIX user name
 - **AD Attribute - UNIX Group Name** - Name in the Active Directory database of the equivalent UNIX group name
 - **Native LDAP Attribute - Windows User Name** - Name in the LDAP database of the equivalent Windows identity
 - **IDMU**
3. **To save your settings, click APPLY. To clear your settings, click REVERT.**

Related Topics

- For information on the different mapping modes, see [“Identity Mapping Concepts” on page 289.](#)
- To create an "allow" or "deny" mapping rule, see [“Creating a Mapping Rule \(BUI\)” on page 282.](#)

▼ Configuring Identity Mapping (CLI)

Use the following procedure to configure identity mapping.

Before You Begin Ensure that you are joined to at least one Active Directory domain.

1. **Go to configuration services idmap.**
2. **Enter get to view the identity mapping properties.**

```
hostname:configuration services idmap> get
```

```
<status> = online
ad_unixuser_attr =
ad_unixgroup_attr =
nldap_winname_attr =
directory_based_mapping = none
```

The three *_attr properties correspond to the three fields on C>S>Identity

Mapping>Properties.

3. Set `directory_based_mapping` to one of the following mapping modes.

- To use rule-based mapping, set `directory_based_mapping` to none.

```
hostname:configuration services idmap> set directory_based_mapping=none
hostname:configuration services idmap>
```

- To use directory-based mapping, set `directory_based_mapping` to name and assign each of the following attributes.

- **ad_unixuser_attr** - Name in the Active Directory database of the equivalent UNIX user name
- **ad_unixgroup_attr** - Name in the Active Directory database of the equivalent UNIX group name
- **nldap_winname_attr** - Name in the LDAP database of the equivalent Windows identity

```
hostname:configuration services idmap> set directory_based_mapping=name
hostname:configuration services idmap> set ad_unixuser_attr=demo_unixuser
hostname:configuration services idmap> set ad_unixgroup_attr=demo_group
hostname:configuration services idmap> set nldap_winname_attr=demo_winuser
```

- To use Identity Management for UNIX (IDMU), set `directory_based_mapping` to idmu.

```
hostname:configuration services idmap> set directory_based_mapping=idmu
hostname:configuration services idmap>
```

Related Topics

- For information on the different mapping modes, see [“Identity Mapping Concepts” on page 289](#).
- To create an "allow" or "deny" mapping rule, see [“Creating a Mapping Rule \(CLI\)” on page 284](#).

▼ Creating a Mapping Rule (BUI)

Use the following procedure to grant or deny credentials for specific users through the identity mapping service. An "allow" mapping rule grants Windows identity credentials from a UNIX identity or vice versa. A "deny" mapping rule blocks a Windows identity from receiving the credentials of a UNIX identity or vice versa.

Note - If you create a mapping rule that blocks a particular user, and the user's name then changes, the mapping no longer blocks that user.

Before You Begin Configure rule-based mapping as described in [“Configuring Identity Mapping \(BUI\)” on page 280](#).

1. **Go to Configuration > Services > Identity Mapping > Rules.**
2. **Click the add item icon  next to Rules.**
3. **In the Add Mapping Rule dialog box, choose either Allow or Deny for the mapping type.**
4. **Complete the remaining fields according to the selected mapping type.**
 - **Allow mapping:**
 - **Mapping Direction** - Choose a direction.
 - **Windows Domain** - Type the Active Directory domain of the Windows identity, or select All.
 - **Windows Identity** - Type the name of the Windows identity.
 - **Unix Identity** - Type the name of the UNIX identity.
 - **Unix Identity Type** - Select either User or Group.
 - **Deny mapping:**
 - a. **For Mapping Direction, choose one of the two options.**
 - **Block Windows identity mapping** - Prevents a Windows identity from gaining the credentials of a UNIX identity
 - **Block UNIX identity mapping** - Prevents a UNIX identity from gaining the credentials of a Windows identity
 - b. **Enter the Windows or UNIX identity information.**
 - **If you selected Block Windows identity mapping, type the Windows domain and identity you want to block.**
 - **If you selected Block UNIX identity mapping, type the UNIX identity and identity type you want to block.**

5. **Click ADD.**

The new mapping appears in the Rules list.

Related Topics

- [“Mapping Rule Directional Symbols” on page 291](#)

▼ **Creating a Mapping Rule (CLI)**

Use the following procedure to grant or deny credentials for specific users through the identity mapping service. An "allow" mapping rule grants Windows identity credentials from a UNIX identity or vice versa. A "deny" mapping rule blocks a Windows identity from receiving the credentials of a UNIX identity or vice versa.

Note - if you create a mapping rule that blocks a particular user and the user's name then changes, the mapping no longer blocks that user.

Before You Begin Configure rule-based mapping as described in [“Configuring Identity Mapping \(CLI\)” on page 281](#).

1. **Go to configuration services idmap.**

2. **Enter create.**

```
hostname:configuration services idmap> create
hostname:configuration services idmap (uncommitted)>
```

3. **Set the properties appropriately.**

You can use the `list` command to view the available properties.

```
hostname:configuration services idmap (uncommitted)> list
Properties:
    windomain = (unset)
    winname = (unset)
    direction = (unset)
    unixname = (unset)
    unixtype = (unset)
```

a. **windomain** - Active Directory domain of the Windows identity.

b. **winname** - Set to one of the following options.

- To create an "allow" mapping, set `winname` to the name of the Windows identity.

Enter `*` to indicate all users within the specified domain.

- To create a "deny" mapping that blocks a UNIX identity from receiving the credentials of a Windows identity, set to the name of the Windows identity.
- To create a "deny" mapping that blocks a Windows identity from receiving the credentials of a UNIX identity, do not set `winname`.

c. `direction` - Set to the direction of the mapping:

- `win2unix` - Mapping from Windows to UNIX
- `unix2win` - Mapping from UNIX to Windows
- `bi` - Bidirectional mapping

d. `unixname` - Set to one of the following options:

- To create an "allow" mapping, set to the name of the UNIX identity, or enter `*` to indicate all users of the specified type.
- To create a "deny" mapping that blocks a Windows identity from receiving the credentials of a UNIX identity, set to the name of the UNIX identity.
- To create a "deny" mapping that blocks a UNIX identity from receiving the credentials of a Windows identity, do not set `unixname`.

e. `unixtype` - Set to either `user` or `group` for the UNIX identity type.

```
hostname:configuration services idmap (uncommitted)> set windomain=demo.domain.com
hostname:configuration services idmap (uncommitted)> set winname=*
hostname:configuration services idmap (uncommitted)> set direction=win2unix
hostname:configuration services idmap (uncommitted)> set unixname=
hostname:configuration services idmap (uncommitted)> set unixtype=user
```

4. Enter `commit` to commit the changes and create the mapping rule.

```
hostname:configuration services idmap (uncommitted)> commit
hostname:configuration services idmap>
```

You can use the `list` command to view the new rule in the Rules list.

```
hostname:configuration services idmap> list
```

MAPPING	WINDOWS ENTITY	DIRECTION	UNIX ENTITY
idmap-000	Alice@demo.domain.com	(U) ==	wdp (U)
idmap-001	*@demo.domain.com	(U) =>	"" (U)

Example 13 Creating a Bi-Directional Mapping (CLI)

This example creates a bi-directional name-based mapping between a Windows user and UNIX user.

```
hostname:> configuration services idmap
hostname:configuration services idmap> create
hostname:configuration services idmap (uncommitted)> set
    windomain=eng.fishworks.com
hostname:configuration services idmap (uncommitted)> set winname=Bill
hostname:configuration services idmap (uncommitted)> set direction=bi
hostname:configuration services idmap (uncommitted)> set unixname=wdp
hostname:configuration services idmap (uncommitted)> set unixtype=user
hostname:configuration services idmap (uncommitted)> commit
hostname:configuration services idmap> list
MAPPING      WINDOWS ENTITY      DIRECTION    UNIX ENTITY
idmap-000    Bill@eng.fishworks.com  (U) ==      wdp (U)
```

Example 14 Creating a Deny Mapping (CLI)

This example creates a deny mapping to prevent all Windows users in a domain from obtaining credentials.

```
hostname:configuration services idmap> create
hostname:configuration services idmap (uncommitted)> list
Properties:
    windomain = (unset)
    winname = (unset)
    direction = (unset)
    unixname = (unset)
    unixtype = (unset)

hostname:configuration services idmap (uncommitted)> set
    windomain=guest.fishworks.com
hostname:configuration services idmap (uncommitted)> set winname=*
hostname:configuration services idmap (uncommitted)> set direction=win2unix
hostname:configuration services idmap (uncommitted)> set unixname=
hostname:configuration services idmap (uncommitted)> set unixtype=user
```

```

hostname:configuration services idmap (uncommitted)> commit
hostname:configuration services idmap> list
MAPPING      WINDOWS ENTITY      DIRECTION  UNIX ENTITY
idmap-000    Bill@eng.fishworks.com    (U) ==      wdp (U)
idmap-001    *@guest.fishworks.com    (U) =>      "" (U)

```

▼ Viewing a Mapping (BUI)

Use the following procedure to view an existing mapping.

1. **Go to Configuration > Services > Identity Mapping > Show Mappings.**
2. **Choose either Windows or UNIX for the platform from which the identity is mapped.**
3. **Enter the Windows or UNIX identity information.**

- **If you selected Windows, type the Windows domain and name of the user.**
- **If you selected UNIX, choose either User or Group for the type, and type the entity name.**

4. **Click SHOW MAPPING.**

The identity user or group properties are displayed. The mapping source and backend origin are also displayed:

Source

- **New mapping** - The mapping was newly created and was neither retrieved from the cache nor predefined.
- **Cached mapping** - The mapping was retrieved from the cache, where mappings are stored for 10 minutes after they are requested.
- **Hard coded mapping** - The mapping is predefined and fixed on the appliance. These mappings were created for default UNIX and Windows identities.
- **Algorithmic mapping** - A non-ephemeral UNIX UID or GID could not be mapped by name, so it was mapped to an algorithmically generated SID.

Backend

- **AD Directory** - This is a directory-based mapping that was created using annotations in the Active Directory.
- **Native LDAP Directory** - This is a directory-based mapping that was created using annotations in the LDAP directory.

- **IDMU** - The mapping was created using the Windows feature Identity Management for UNIX.
- **Name rule** - The mapping was created using a name rule.
- **Ephemeral** - Since there was no equivalent identity at the time the mapping was created, the system created a temporary one using an ephemeral UID or GID.
- **Local SID** - A non-ephemeral UNIX UID or GID could not be mapped by name, so it was mapped to an algorithmically generated local SID.
- **Well-known mapping** - The mapping uses a "well-known SID." These Windows SIDs identify generic users or generic groups. Their values remain constant across all operating systems.

▼ Flushing Mappings from the Cache (BUI)

Use the following procedure to flush, or expire, all mappings from the cache.

After a requested mapping has been provided, it is stored in the cache for 10 minutes and then expires. You can immediately expire a mapping by using the flush function, which expires all cached mappings.

1. **Go to Configuration > Services > Identity Mapping > Show Mappings.**
2. **Click FLUSH MAP CACHE.**
All cached mappings are expired.

▼ Flushing Mappings from the Cache (CLI)

Use the following procedure to flush, or expire, all mappings from the cache.

After a requested mapping has been provided, it is stored in the cache for 10 minutes and then expires. You can immediately expire a mapping by using the flush function, which expires all cached mappings.

1. **Go to configuration services idmap.**
2. **Enter flush.**

```
hostname:configuration services idmap> flush  
hostname:configuration services idmap>
```

All cached mappings are expired.

Identity Mapping Best Practices

- Configure user-specific identity mapping rules when you want a user to have access to a common set of files through both NFS and SMB clients. If NFS and SMB clients are accessing disjointed filesystems, there is no need to configure any identity mapping rules.
- Reconfiguring the identity mapping service does not affect active SMB sessions. Connected users remain connected, and their previous name mapping is available for authorizing access to additional shares for up to 10 minutes. To prevent unauthorized access, configure the mappings before exporting shares.
- The security that your identity mappings provide is only as good as their synchronization with your directory services. For example, if you create a name-based mapping that denies access to a particular user, and the user's name changes, the mapping no longer denies access to that user.
- You can only have one bidirectional mapping for each Windows domain that maps all users in the Windows domain to all UNIX identities. If you want to create multiple domain-wide rules, be sure to specify that those rules map *only* from Windows to UNIX.
- Use the IDMU mapping mode instead of directory-based mapping whenever possible.

Identity Mapping Concepts

The SMB service uses the identity mapping service to associate Windows and UNIX identities. When the SMB service authenticates a user, it uses the identity mapping service to map the user's Windows identity to the appropriate UNIX identity. If no UNIX identity exists for a Windows user, the service generates a temporary identity using an ephemeral UID and GID. These mappings allow a share to be exported and accessed concurrently by SMB and NFS clients. By associating Windows and UNIX identities, NFS and SMB clients can share the same identity, thereby allowing access to the same set of files.

In the Windows operating system, an access token contains the security information for a login session and identifies the user, the user's groups, and the user's privileges. Administrators define Windows users and groups in a Workgroup, or in a SAM database, which is managed on an Active Directory domain controller. Each user and group has a SID, which uniquely identifies the user or group, both within a host and a local domain, and across all possible Windows domains.

UNIX creates user credentials based on user authentication and file permissions. Administrators define UNIX users and groups in local password and group files or in a name or directory service, such as NIS or LDAP. Each UNIX user and group has a UID and GID. Typically, the UID or GID uniquely identifies a user or group within a single UNIX domain. However, these values are not unique across domains.

The following options are available when selecting a mapping mode:

- **Rule-based Mapping** - Use for creating various rules that map identities by name, thus establishing equivalences between Windows and UNIX identities. Mapping rules are useful when you want a user to access the same set of files through both SMB and NFS clients.
- **Directory-based Mapping** - Use for annotating an LDAP or Active Directory object with information about how the identity maps to an equivalent identity on the opposite platform.
- **IDMU-based Mapping** - Identity Management for UNIX (IDMU) is a feature that Microsoft offers for Windows Server 2003, and is bundled with Windows Server 2003 R2 and later. IDMU supports Windows as a NIS/NFS server by adding a "UNIX Attributes" panel to the Active Directory Users and Computers user interface. This allows administrators to specify a number of UNIX-related parameters, including UID, GID, login shell, and home directory. These parameters are made available through Active Directory using a schema similar to, but not the same as, RFC 2307, and through the NIS service. When the IDMU mapping mode is selected, the identity mapping service consumes these UNIX attributes to establish mappings between Windows and UNIX identities. This approach is very similar to directory-based mapping, except that the identity mapping service queries the property schema established by the IDMU software instead of allowing a custom schema. When this approach is used, no other directory-based mapping may occur.

Cached and Ephemeral Mappings

When the identity mapping service provides a name mapping, it stores the mapping in the cache for 10 minutes, at which point the mapping expires. Within its 10-minute life, a mapping is persistent across restarts of the identity mapping service. Changes to the mappings or to the name service directories do not affect existing connections within the 10-minute life of a mapping. The service evaluates mappings only when the client tries to connect to a share and there is no unexpired mapping. For example, if the SMB server requests a mapping for the user after the mapping has expired, the service re-evaluates the mapping.

If no name-based mapping rule applies for a particular user, that user will be given temporary credentials through an ephemeral mapping unless the user is blocked by another mapping. When a Windows user with an ephemeral UNIX name creates a file on the system, Windows clients accessing the file using SMB see that the file is owned by that Windows identity. However, NFS clients see that the file is owned by "nobody".

Identity Mapping Case Sensitivity

Windows names are not case sensitive, but UNIX names are case sensitive. The user names JSMITH, JSmith, and jsmith are equivalent names in Windows, but they are three distinct names in UNIX. Case sensitivity affects name mappings differently depending on the direction of the mapping.

- For a Windows-to-UNIX mapping to produce a match, the case of the Windows user name must match the case of the UNIX user name. For example, only Windows user name "jsmith" matches UNIX user name "jsmith". Windows user name "Jsmith" does not match.
- An exception to the case matching requirement for Windows-to-UNIX mappings occurs when the mapping uses the wildcard character "*" to map multiple user names.

If the identity mapping service encounters a mapping that maps Windows user *@some.domain to UNIX user "*", it first searches for a UNIX name that matches the Windows name exactly. If it does not find a match, the service converts the entire Windows name to lower case and searches again for a matching UNIX name. For example, the Windows user name "JSmith@some.domain" maps to UNIX user name "jsmith". If the service does not find a match after using lowercase for the Windows user name, the user does not obtain a mapping.

You can create a rule to match strings that differ only in case. For example, you can create a user-specific mapping to map the Windows user "JSmith@some.domain" to UNIX user "jSmith". Otherwise, the service assigns an ephemeral ID to the Windows user.

- For a UNIX-to-Windows mapping to produce a match, the case does not have to match. For example, UNIX user name "jsmith" matches any Windows user name with the letters "JSMITH" regardless of case.

Mapping Rule Directional Symbols

After creating a name-based mapping, the following symbols indicate the semantics of each rule.

-  - Maps Windows identity to UNIX identity and UNIX identity to Windows identity
-  - Maps Windows identity to UNIX identity
-  - Maps UNIX identity to Windows identity
-  - Prevents Windows identity from obtaining credentials
-  - Prevents UNIX identity from obtaining credentials

If an icon is gray instead of black, the rule matches a UNIX identity that cannot be resolved.

IPMP Configuration

IPMP (Internet Protocol Network Multipathing) allows multiple network interfaces to be grouped as one, for both improved network bandwidth and reliability (interface redundancy). Some properties can be configured in this section. For the configuration of network interfaces in IPMP groups, see [“Network Configuration” on page 89](#).

TABLE 63 IPMP Properties

Property	Description
Failure detection latency	Time for IPMP to declare a network interface has failed, and to fail over its IP addresses
Enable fail-back	Allow the service to resume connections to a repaired interface

Changing services properties is documented in [“Setting Service Properties \(BUI\)” on page 252](#) and [“Setting Service Properties \(CLI\)” on page 253](#). The CLI property names are shorter versions of those listed above.

TABLE 64 IPMP Logs

Log	Description
network-initial:default	Logs the network configuration process

iSCSI Configuration

When you configure a LUN on the appliance you can export that volume over an Internet Small Computer System Interface (iSCSI) target. The iSCSI service allows iSCSI initiators to access targets using the iSCSI protocol.

The service supports discovery, management, and configuration using the iSNS protocol. The iSCSI service supports both unidirectional (target authenticates initiator) and bidirectional (target and initiator authenticate each other) authentication using CHAP. Additionally, the service supports CHAP authentication data management in a RADIUS database.

The system performs authentication first, and authorization second, in two independent steps.

Note - For examples of configuring iSCSI initiators and targets, see [“Configuring Storage Area Network \(SAN\)” on page 170](#).

TABLE 65 iSCSI Service Properties

Property	Description
Use iSNS	Whether iSNS discovery is enabled
iSNS Server	An iSNS server
Use RADIUS	Whether RADIUS is enabled

Property	Description
RADIUS Server	A RADIUS server
RADIUS Server Secret	The RADIUS server's secret

If the local initiator has a CHAP name and a CHAP secret, the system performs authentication. If the local initiator does not have the CHAP properties, the system does not perform any authentication and therefore all initiators are eligible for authorization.

The iSCSI service allows you to specify a global list of initiators that you can use within initiator groups.

If your initiator cannot connect to your target:

- Make sure the IQN of the initiator matches the IQN identified in the initiators list.
- Check that IP address of iSNS server is correct and that the iSNS server is configured.
- Check that the IP address of the target is correct on the initiator side.
- Check that initiator CHAP names and secrets match on both sides.
- Make sure that the target CHAP name and secret do not match those of any of the initiators.
- Check that the IP address and secret of the RADIUS server are correct, and that the RADIUS server is configured.
- Check that the initiator accessing the LUN is a member of that LUN's initiator group.
- Check that the targets exporting that LUN are online.
- Check that the LUN's operational status is online.
- Check the logical unit number for each LUN.

If, during the failover / failbacks, the iSER Reduced Copy I/Os from the Red Hat client are not surviving, modify the `node.session.timeo.replacement_timeout` parameter in the `/etc/iscsi/iscsid.conf` file to 300sec.

Related Topics

- Setting Service Properties [BUI](#), [CLI](#)

Kerberos Configuration

Kerberos is a network protocol that uses secret-key cryptography to authenticate communication between a client and a host machine or service. It uses a Key Distribution Center (KDC) server to issue time-stamped tickets. You can use the appliance to import Kerberos principals and keys created on the KDC, or you can configure principals for the

KDC using the appliance, and their keys are automatically created. Although you can use both methods, importing is the best practice and most commonly used. All keys are encrypted using the Kerberos password and stored within the appliance keytab file.

Both Kerberos and Active Directory can be enabled at the same time because they have distinct realms and keys. When both are active, the Kerberos realm is the default.

The appliance can use Kerberos to authenticate users for administrative login and for access to services, including NFS, HTTP, FTP, SFTP, and SSH. An appliance user must have a Kerberos principal by the same name to use Kerberos authentication for these services. Kerberos can also be used to set security for individual shares that use the NFS protocol, as described in [“Configuring Kerberos Realms for NFS” on page 330](#). Since the Kerberos service uses time stamps, configure the appliance NTP service first.

To configure Kerberos, see the following sections:

- Creating a Kerberos Realm - [BUI](#), [CLI](#)
- Importing Kerberos Keys - [BUI](#), [CLI](#)
- Creating Kerberos Principals and Keys - [BUI](#), [CLI](#)
- Deleting Kerberos Principals and Keys - [BUI](#), [CLI](#)
- Destroying a Kerberos Realm - [BUI](#), [CLI](#)
- “Kerberos Service Properties” on page 309
- “Kerberos Properties and Logs” on page 309
- “Configuring Kerberos Realms for NFS” on page 330

▼ Creating a Kerberos Realm (BUI)

Use the following procedure to create a Kerberos realm, set the KDC(s), and select strong or weak encryption types. Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#).

Before You Begin Ensure that you have configured the [NTP service](#).

1. **Go to Configuration > Services.**
2. **To enable the Kerberos service, click the enable icon  for Kerberos.**
3. **Click Kerberos.**
4. **In the Realm field, type the Kerberos realm.**

For familiarity, the realm name can be the same as your DNS domain name, except that the realm name is in uppercase.

5. **In the KDC(s) field, type the host name of the KDC administrative server.**
If your Kerberos configuration includes DNS support for KDC lookup, leave this field blank.
6. **If you have another KDC, click the add icon  next to KDC(s) and type its host name. Repeat for each additional KDC.**
If your configuration includes DNS support, do not complete this step.
7. **To allow support for weak encryption types, such as DES and Exportable ArcFour with HMAC/md5, select Allow weak encryption types.**
The default does not support weak encryption types.
8. **Click APPLY.**
To reset the properties to their original values, click REVERT instead.

Next Steps

Choose one of the following options:

- [“Importing Kerberos Keys \(BUI\)” on page 297](#)
- [“Creating Kerberos Principals and Keys \(BUI\)” on page 301](#)

▼ Creating a Kerberos Realm (CLI)

Use the following procedure to create a Kerberos realm, set the KDC(s), and select strong or weak encryption types. Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#) and [“Kerberos Properties and Logs” on page 309](#).

Before You Begin Ensure that you have configured the [NTP service](#).

1. **Go to configuration services kerberos and enter show.**

```
hostname:configuration services kerberos> show
Properties:
    <status> = disabled
    allow_weak_crypto = false
```

2. **To enable the Kerberos service, enter enable and then enter commit.**
3. **To allow support for weak encryption types, such as DES and Exportable ArcFour with HMAC/md5, enter set allow_weak_crypto=true and then enter commit.**
The default does not support weak encryption types.

4. **To create a realm, enter create and the realm name, and then enter commit.**
For familiarity, the realm name can be the same as your DNS domain name, except that the realm name is in uppercase.

```
hostname:configuration services kerberos> create TEST.NET
hostname:configuration services kerberos TEST.NET (uncommitted)> commit
```

5. **Enter done.**
6. **To view all realms, enter list.**

```
hostname:configuration services kerberos> list
REALM          KDC
TEST.NET
```

7. **Select the realm.**

```
hostname:configuration services kerberos> select TEST.NET
hostname:configuration services kerberos TEST.NET>
```

8. **To configure the KDC server(s), enter set kdc= and the KDC administrative server host name. If you have additional KDCs, add them to the same line and separate them by commas. Then enter commit.**

If your Kerberos configuration includes DNS support for KDC lookup, do not perform this step.

```
hostname:configuration services kerberos TEST.NET> set kdc=kdc1.us.oracle.com,kdc2.us.
oracle.com
kdc = kdc1.us.oracle.com,kdc2.us.oracle.com (uncommitted)
hostname:configuration services kerberos TEST.NET> commit
```

Next Steps

Choose one of the following options:


- [“Importing Kerberos Keys \(CLI\)” on page 298](#)
- [“Creating Kerberos Principals and Keys \(CLI\)” on page 303](#)

▼ Importing Kerberos Keys (BUI)

Use the following procedure to import Kerberos keys that were created on the KDC. The keys are then stored in the appliance keytab. This task does not require login credentials on the KDC. Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#).


Before You Begin Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in [“Creating a Kerberos Realm \(BUI\)” on page 294](#).

1. **Go to Configuration > Services.**
2. **Click Kerberos.**
3. **Click Keys and click IMPORT KEYS.**



4. **In the Import Keys dialog box, click Browse and select the Kerberos keytab file.**

5. Click **UPLOAD**.



The list of keys is displayed.

Keys 16 Total				1-16	
				CREATE PRINCIPALS AND KEYS	IMPORT KEYS
KVNO	PRINCIPAL	ENCRYPTION TYPE	TYPE #		
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23		
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16		
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17		
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18		
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23		
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16		
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17		
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18		
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23		
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16		
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17		
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18		
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23		
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16		
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17		
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18		

▼ Importing Kerberos Keys (CLI)

Use the following procedure to import Kerberos keys that were created on the KDC. The keys are then stored in the appliance keytab. This task does not require login credentials on the KDC.

Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#) and [“Kerberos Properties and Logs” on page 309](#).

Before You Begin Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in [“Creating a Kerberos Realm \(CLI\)” on page 296](#).

1. **Go to configuration services kerberos importkeytab and enter show to view the properties.**

```
hostname:configuration services kerberos importkeytab (uncommitted)> show
Properties:
    url = (unset)
    user = (unset)
    password = (unset)
```

2. **Enter set url= and the URL of the Kerberos keytab file.**

```
hostname:configuration services kerberos importkeytab (uncommitted)> set url=http://
akbuild1/shares/export/123456/demo.keytab
url = http://akbuild1/shares/export/123456/demo.keytab
```

3. **Enter set user= and the user name for URL access.**

```
hostname:configuration services kerberos importkeytab (uncommitted)> set user=myusername
user = myusername
```

4. **Enter set password= and the password for URL access, and then enter commit.**

```
hostname:configuration services kerberos importkeytab (uncommitted)> set
password=letmein
password = *****
hostname:configuration services kerberos importkeytab (uncommitted)> commit
Transferred 718 of 718 (100%) . . . done
Imported 8 keys.
```

5. **Enter show to view the realms and KDCs.**

```
hostname:configuration services kerberos> show
Properties:
    <status> = online
    allow_weak_crypto = true
Realms:
REALM      KDC
TEST.NET   kdc1.us.oracle.com
```

6. **To view the principals for a realm, select a realm and enter show.**

```
hostname:configuration services kerberos> select TEST.NET
hostname:configuration services kerberos TEST.NET> show
Properties:
    kdc = kdc1.us.oracle.com
Keytab entries:
NAME      KEYS  PRINCIPAL
principal-000  4    host/hostname.us.oracle.com@TEST.NET
principal-001  4    nfs/hostname.us.oracle.com@TEST.NET
```

7. To view the keys for a principal, select a principal and enter show.

```
hostname:configuration services kerberos TEST.NET> select principal-001
hostname:configuration services kerberos principal-001> show
Properties:
    name = nfs/hostname.us.oracle.com@TEST.NET
Keys:
KEY      KVNO  ENCTYPENO  ENCTYPE
key-000  28    18         AES-256 CTS mode with 96-bit SHA-1 HMAC
key-001  28    17         AES-128 CTS mode with 96-bit SHA-1 HMAC
key-002  28    16         Triple DES cbc mode with HMAC/sha1
key-003  28    23         ArcFour with HMAC/md5
key-004  28    24         Exportable ArcFour with HMAC/md5
key-005  28    3          DES cbc mode with RSA-MD5
key-006  28    1          DES cbc mode with CRC-32
```

Legend for column headings:

- KEY = Key name
- KVNO = Key version number
- ENCTYPENO = Encryption type number
- ENCTYPE = Encryption type

8. To view the properties of a key, select a key and enter show.


```
hostname:configuration services kerberos principal-001> select key-003
hostname:configuration services kerberos principal-001 key-003> show
Properties:
    principal = nfs/hostname.us.oracle.com@TEST.NET
    kvno = 28
    enctype = ArcFour with HMAC/md5
    enctypeno = 23
```


▼ Creating Kerberos Principals and Keys (BUI)

Use the following procedure to create Kerberos principals on the KDC administrative server using the appliance. Keys are generated for each principal and stored in the appliance keytab. Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#).


- Before You Begin**
- Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in [“Creating a Kerberos Realm \(BUI\)” on page 294](#).
 - Ensure that you have login credentials on the KDC.

1. **Go to Configuration > Services.**
2. **Click Kerberos.**
3. **Click Keys and click CREATE PRINCIPALS AND KEYS.**



4. **In the KDC Admin Login dialog box, complete the following fields:**
 - **Realm** - This field is auto-populated and cannot be modified.
 - **Admin server** - KDC administrative server host name. This field is auto-populated, but can be modified.
 - **Admin principal** - KDC administrator name for the realm.

- **Password** - Password for the KDC administrator.



A screenshot of a 'KDC Admin login' dialog box. The dialog has a title bar with 'KDC Admin login' and two buttons: 'CANCEL' and 'OK'. Inside the dialog, there are four labeled text input fields arranged vertically. The first field is labeled 'Realm' and contains the text 'TEST.NET'. The second field is labeled 'Admin server' and contains 'kdc1.us.oracle.com'. The third field is labeled 'Admin principal' and contains 'kdc/admin@TEST.NET'. The fourth field is labeled 'Password' and contains seven dots, indicating a masked password.

Realm	TEST.NET
Admin server	kdc1.us.oracle.com
Admin principal	kdc/admin@TEST.NET
Password	•••••••

5. Click OK.
6. In the confirmation box, click OK.

Keys 16 Total		CREATE PRINCIPALS AND KEYS		IMPORT KEYS
KVNO	PRINCIPAL	ENCRYPTION TYPE	TYPE #	
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23	
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16	
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17	
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18	
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23	
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16	
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17	
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18	
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23	
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16	
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17	
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18	
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23	
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16	
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17	
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18	

▼ Creating Kerberos Principals and Keys (CLI)

Use the following procedure to create Kerberos principals on the KDC administrative server using the appliance. Keys are generated for each principal and stored in the appliance keytab. Descriptions of each property are located in [“Kerberos Service Properties” on page 309](#) and [“Kerberos Properties and Logs” on page 309](#).

Before You Begin

- Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in [“Creating a Kerberos Realm \(CLI\)”](#) on page 296.
- Ensure that you have login credentials on the KDC.

1. Go to configuration services kerberos and enter list.

```
hostname:configuration services kerberos> list
REALM          KDC
TEST.NET
```

- 2. Select the realm.**

```
hostname:configuration services kerberos> select TEST.NET
hostname:configuration services kerberos TEST.NET>
```

3. **To create the principals, enter principals and then enter show to view the properties.**

```
hostname:configuration services kerberos TEST.NET> principals
hostname:configuration services kerberos TEST.NET principals (uncommitted)> show
Properties:
    realm = TEST.NET
    server = kdc1.us.oracle.com
    admin = (unset)
    password = (unset)
```

4. **(Optional) To change the KDC server, enter set kdcs= and the KDC server host name. Then enter commit.**

```
hostname:configuration services kerberos TEST.NET> set kdcs=kdc2.us.oracle.com
    kdcs = kdc2.us.oracle.com (uncommitted)
hostname:configuration services kerberos TEST.NET> commit
```

5. **Enter set admin= and the KDC administrator name for the realm.**

```
hostname:configuration services kerberos TEST.NET principals (uncommitted)> set
    admin=kdc/admin
```

6. **Enter set password= and the KDC administrator password, and then enter commit.**

```
hostname:configuration services kerberos TEST.NET principals (uncommitted)> set
    password=test123
    password = *****
hostname:configuration services kerberos TEST.NET principals (uncommitted)> commit
```

7. **Enter show to view the principals for the KDC.**

```
hostname:configuration services kerberos TEST.NET> show
Properties:
    kdcs = kdc1.us.oracle.com
Keytab entries:
NAME      KEYS  PRINCIPAL
principal-000  4    host/hostname.us.oracle.com@TEST.NET
principal-001  4    nfs/hostname.us.oracle.com@TEST.NET
```

8. **To view the keys for a principal, select a principal and enter show.**

```
hostname:configuration services kerberos TEST.NET> select principal-001
```

```
hostname:configuration services kerberos principal-001> show
Properties:
    name = nfs/hostname.us.oracle.com@TEST.NET
Keys:
KEY      KVNO  ENCTYPENO  ENCTYPE
key-000  28    18         AES-256 CTS mode with 96-bit SHA-1 HMAC
key-001  28    17         AES-128 CTS mode with 96-bit SHA-1 HMAC
key-002  28    16         Triple DES cbc mode with HMAC/sha1
key-003  28    23         ArcFour with HMAC/md5
key-004  28    24         Exportable ArcFour with HMAC/md5
key-005  28    3          DES cbc mode with RSA-MD5
key-006  28    1          DES cbc mode with CRC-32
```

Legend for column headings:

- KEY = Key name
- KVNO = Key version number
- ENCTYPENO = Encryption type number
- ENCTYPE = Encryption type

9. To view the properties of a key, select a key and enter `show`.

```
hostname:configuration services kerberos principal-001> select key-003
hostname:configuration services kerberos principal-001 key-003> show
Properties:
    principal = nfs/hostname.us.oracle.com@TEST.NET
    kvno = 28
    enctype = ArcFour with HMAC/md5
    enctypeno = 23
```

▼ Deleting Kerberos Principals and Keys (BUI)

Use the following procedure to delete individual keys.

- 1. Go to Configuration > Services.**
- 2. Click Kerberos.**
- 3. Click Keys.**

The list of principals and keys is displayed.

Keys 16 Total			
CREATE PRINCIPALS AND KEYS IMPORT KEYS			
KVNO	PRINCIPAL	ENCRYPTION TYPE	TYPE #
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18

- (Optional) To sort by a column, such as **PRINCIPAL**, click the column heading.
- To delete an individual key, hover over the appropriate row, click its trash icon , and confirm your action.

If you delete all keys for a principal, you effectively delete the principal from the appliance.

▼ Deleting Kerberos Principals and Keys (CLI)

Use the following procedure to delete individual keys, or to delete all keys for a principal.

- Go to configuration services kerberos and enter list.**

```
hostname:configuration services kerberos> list
REALM          KDC
TEST.NET
```

2. Select the realm.

```
hostname:configuration services kerberos> select TEST.NET
hostname:configuration services kerberos TEST.NET>
```

3. Enter **show** to view the principals for the KDC.

```
hostname:configuration services kerberos TEST.NET> show
Properties:
    kdc = kdc1.us.oracle.com
Keytab entries:
NAME      KEYS  PRINCIPAL
principal-000  4    host/hostname.us.oracle.com@TEST.NET
principal-001  4    nfs/hostname.us.oracle.com@TEST.NET
```

4. To delete all of the keys for a principal, enter **destroy** and the principal name, and confirm your action.

To delete an individual key, see the next step.

```
hostname:configuration services kerberos TEST.NET> destroy principal-000
This will delete all keys for "principal-000". Are you sure? (Y/N) Y
```

5. To delete an individual key for a principal, first select a principal and enter **show** to view the list of keys.

```
hostname:configuration services kerberos TEST.NET> select principal-001
hostname:configuration services kerberos principal-001> show
Properties:
    name = nfs/hostname.us.oracle.com@TEST.NET
Keys:
KEY      KVNO  ENCTYPENO  ENCTYPE
key-000  28    18         AES-256 CTS mode with 96-bit SHA-1 HMAC
key-001  28    17         AES-128 CTS mode with 96-bit SHA-1 HMAC
key-002  28    16         Triple DES cbc mode with HMAC/sha1
key-003  28    23         ArcFour with HMAC/md5
key-004  28    24         Exportable ArcFour with HMAC/md5
key-005  28    3          DES cbc mode with RSA-MD5
key-006  28    1          DES cbc mode with CRC-32
```

Legend for column headings:

- KEY = Key name
- KVNO = Key version number
- ENCTYPENO = Encryption type number
- ENCTYPE = Encryption type

6. To view the properties of a key, select a key and enter **show**.

```
hostname:configuration services kerberos principal-001> select key-003
hostname:configuration services kerberos principal-001 key-003> show
Properties:
    principal = nfs/hostname.us.oracle.com@TEST.NET
    kvno = 28
    enctype = ArcFour with HMAC/md5
    enctypeno = 23
```

7. To delete a key or view a different key, enter **done** to return to the principal context.

```
hostname:configuration services kerberos principal-001 key-003> done
hostname:configuration services kerberos principal-001>
```

8. To delete the key, enter **destroy** and the key name, and confirm your action.

```
hostname:configuration services kerberos principal-001> destroy key-003
This will delete key "key-003". Are you sure? (Y/N) Y
```

▼ Destroying a Kerberos Realm (BUI)

Destroying a realm also destroys its corresponding keys.

1. Go to **Configuration > Services**.
2. Click **Kerberos**.
3. Clear the **Realm** field, click **APPLY**, and confirm your action.

▼ Destroying a Kerberos Realm (CLI)

Destroying a realm also destroys its corresponding keys.

1. Go to **configuration services kerberos**.
2. Enter **destroy** and the realm name, and then confirm your action.

```
hostname:configuration services kerberos> destroy TEST.NET
This will destroy "TEST.NET". Are you sure? (Y/N) Y
```


Kerberos Service Properties

The following properties are available for the Kerberos service:

- **Realm** - A string, the name of the realm.
- **KDC(s)** - A list of zero or more host names, the Key Distribution Center(s) for the realm. The first Key Distribution Center (KDC) listed is assumed to be the Admin Server, which is relevant if creating principals on the KDC from the appliance, but not when importing keys. The list may be empty if at least one KDC is published for the realm in DNS.
- **Allow weak encryption types** - A Boolean value. This enables/disables support for deprecated weak encryption types (des-cbc-crc, des-cbc-md5, and arcfour-hmac-exp). This property is disabled by default.
- **Admin** - A string, the name of the Kerberos admin principal (administrator). By convention, a principal name is divided into three components: the primary, the instance, and the realm. You can specify a principal as joe, joe/admin, or joe/admin@ENG.EXAMPLE.COM. This property is used only if creating service principals, and is not retained.
- **Password** - Kerberos admin password - A string, the password for the administrator. This property is used only if creating service principals, and is not retained.

Changing services properties is documented in [“Setting Service Properties \(BUI\)” on page 252](#) and [“Setting Service Properties \(CLI\)” on page 253](#).

Kerberos Properties and Logs

The following table describes the mapping between Kerberos CLI properties and their BUI property descriptions.

Note - Older Kerberos properties associated with the NFS service have been deprecated and will continue to function in scripts and workflows.

TABLE 66 Kerberos Properties

CLI Property	BUI Property
realm	Realm
kdc	KDC(s)
allow_weak_crypto	Allow weak encryption types. Permits weak encryption types in Kerberos (arcfour-hmac-md5-exp, des-cbc-md5, and des-cbc-crc).
principals	Kerberos administrator principal
principals - server	Admin server - KDC host name
principals - admin	Admin principal - Administrator login name on KDC

CLI Property	BUI Property
principals - password	Password - Administrator login password on KDC
importkeytab	Import Keys - Imports keys in a keytab file from the KDC
importkeytab - url	URL of keytab file
importkeytab - user	User name for URL access
importkeytab - password	Password for URL access

The following log is available for the Kerberos service.

TABLE 67 Logs Available for Kerberos Service

Log	Description
appliance-kit-kerberos:default	Log of appliance Kerberos service events

LDAP Configuration

Lightweight Directory Access Protocol (LDAP) is a directory service for centralizing management of users, groups, hostnames, and other resources (called objects). This service on the appliance acts as an LDAP client so that:

- LDAP users can log in to the FTP and HTTP services.
- LDAP user names (instead of numerical ids) can be used to configure root directory ACLs on a share.
- LDAP users can be granted privileges for appliance administration. The appliance supplements LDAP information with its own privilege settings.
- The LDAP server's certificate can be self-signed.
- You cannot supply a list of trusted CA certificates; each certificate must be individually accepted by the appliance administrator.
- When an LDAP server's certificate expires, you must delete the server from the list and then add it again to accept its new certificate.

Note - UIDs from 0-99 inclusive are reserved by the operating system vendor for use in future applications. Their use by end system users or vendors of layered products is not supported and can cause security issues with other applications.

To configure LDAP and monitor LDAP servers, see the following sections:

- [“Adding an Appliance Administrator \(BUI\)” on page 311](#)
- [“Setting Properties with Multiple Attribute Value Pairs \(CLI\)” on page 311](#)
- Configuring LDAP Security Settings - [BUI](#), [CLI](#)

- Monitoring LDAP Server Status - [BUI](#), [CLI](#)
- “LDAP Properties” on page 316
- “LDAP Custom Mappings” on page 317

▼ Adding an Appliance Administrator (BUI)

To let an existing LDAP user log in using LDAP credentials and administer the appliance, use the following procedure.

Note - If both NIS and LDAP are configured on the appliance and the services return different information for a particular item, the appliance will use the data provided by NIS.

1. **Go to Configuration > Services > LDAP, and enter the properties that you want to use.**
For information about the available properties, see “LDAP Properties” on page 316.
2. **To apply the properties you selected, click Apply or click Revert to start over.**
3. **To add LDAP servers, in the Servers section click the add item icon .**
For information about servers, see the Servers section in “LDAP Properties” on page 316.
4. **To configure the LDAP server, in the New LDAP Server box, enter the LDAP server Address and select the LDAP Certificate source that you want to use.**
For the Certificate source, select Server to search the current server and retrieve the certificate (in an insecure manner), and use it to validate the certificate presented later.
5. **Go to Configuration > Users, and add users as needed using LDAP usernames.**
For information about adding users, see “Configuring Users” on page 202.

Related Topics

- “Setting Properties with Multiple Attribute Value Pairs (CLI)” on page 311
- “LDAP Properties” on page 316
- “LDAP Custom Mappings” on page 317

▼ Setting Properties with Multiple Attribute Value Pairs (CLI)

To set LDAP property values that have multiple attribute value pairs with equal signs (=), surround each attribute value pair with double quotation marks.

1. **Go to configuration services ldap.**
2. **To set multiple attribute value pairs, use the following commands:**

```
hostname:configuration services ldap> set user_mapattr="uid=uid",  
"uidnumber=uidNumber","gidnumber=gidNumber",  
"gecos=displayName","description=distinguishedName",  
"homedirectory=unixHomeDirectory"
```

Related Topics

- [“Adding an Appliance Administrator \(BUI\)” on page 311](#)
- [“LDAP Properties” on page 316](#)
- [“LDAP Custom Mappings” on page 317](#)

▼ Configuring LDAP Security Settings (BUI)

To configure security settings for the LDAP service, use the following procedure.

1. **Go to Configuration > Services.**
2. **Under Directory Services, select LDAP.**
3. **For Security Settings, select one of the following authentication options:**
 - **Anonymous** - Gives the appliance access only to data that is available to everyone. To optionally enable the SSL and TLS protocols, select the check box for **Enable SSL/TLS**.
 - **Self** - Authenticates the appliance using the user's identity and credentials. Self authentication uses Kerberos encryption and the SASL/GSSAPI authentication method.
 - **Proxy** - Specifies authentication through a proxy for a specific user account. Set the following options:
 - (Optional) **Enable SSL/TLS** - Select this check box to enable the SSL and TLS protocols, which is highly recommended when using the simple authentication method so the user's distinguished name and password are not sent as plain text.
 - **Authentication Method** - Select either **Simple** (RFC 4513) or **SASL/DIGEST-MD5**.
 - **Proxy DN** - Distinguished name of account used for proxy authentication.
 - **Proxy Password** - Password for account used for proxy authentication.
4. **Click APPLY.**

Note - After clicking APPLY, the LDAP server configuration is validated. If the Proxy DN or Proxy Password fails or times out, a warning is displayed.

Related Topics

- [“LDAP Properties” on page 316](#)

▼ Configuring LDAP Security Settings (CLI)

To configure security settings for the LDAP service, use the following procedure. For valid property setting combinations, see the table at the end of this task.

1. **Go to configuration services ldap and enter show to view the properties.**

```
hostname:configuration services ldap> show
Properties:
    <status> = enabled
    default_servers =
    proxy_dn =
    proxy_password =
    base_dn =
    search_scope = one
    cred_level = anonymous
    auth_method = none
    use_tls = false
    user_search =
    user_mapattr =
    user_mapobjclass =
    group_search =
    group_mapattr =
    group_mapobjclass =
    netgroup_search =
    netgroup_mapattr =
    netgroup_mapobjclass =
```

2. **To set the credential level, enter set cred_level= and one of the following options:**

- **anonymous** - Allows anonymous authentication for access to data available to everyone.
- **self** - Provides self-authentication for users based on their identity and credentials. Self-authentication uses Kerberos encryption and the SASL/GSSAPI authentication method.
- **proxy** - Specifies authentication through a proxy for a specific user account.

```
hostname:configuration services ldap> set cred_level=proxy
cred_level = proxy (uncommitted)
```

3. **To set the authorization method, enter set auth_method= and one of the following options:**

- **none** - None (use with anonymous)
- **sasl/GSSAPI** - SASL/GSSAPI (use with self)
- **simple** - Simple, RFC 4513 (use with proxy)
- **sasl/DIGEST-MD5** - SASL/DIGEST-MD5 (use with proxy)

```
hostname:configuration services ldap> set auth_method=simple
auth_method = simple (uncommitted)
```

4. To enable or disable SSL/TLS, enter set use_tls= and either true or false.

Enabling SSL/TLS is highly recommended when using the simple authentication method so the user's distinguished name and password are not sent in plain text.

```
hostname:configuration services ldap> set use_tls=true
use_tls = true (uncommitted)
```

5. If the credential level is set to proxy, enter set proxy_dn= and the distinguished name of the account used for proxy authentication. Then enter set proxy_password= and the password for the account.

```
hostname:configuration services ldap> set proxy_dn=ProxyName
proxy_dn = ProxyName (uncommitted)
hostname:configuration services ldap> set proxy_password=MyPassword5
proxy_password = ***** (uncommitted)
```

6. Enter commit.

```
hostname:configuration services ldap> commit
```

Note - Changes to the LDAP server configuration will be validated when committed. If the proxy_dn or proxy_password validation fails or times out, a warning message is displayed.

Refer to the following table for valid security property setting combinations:

cred_level	auth_method	use_tls
anonymous	none	true
	none	false
self	sasl/GSSAPI	false
proxy	simple	true
	simple	false

cred_level	auth_method	use_tls
		<i>Permitted, but not recommended because the user's distinguished name (DN) and password will be sent in plain text.</i>
	sasl/DIGEST-MD5	true
	sasl/DIGEST-MD5	false

Related Topics

- [“LDAP Properties” on page 316](#)

▼ Monitoring LDAP Server Status (BUI)

To monitor LDAP server status, use the following procedure.

1. **Go to Configuration > Services > LDAP > Properties.**
2. **To check the status of LDAP servers, look under the LDAP Servers section. For each server, the following status is displayed:**
 -  - Indicator light representing the status of the server. The indicator light is either online  or unavailable .
 - Last Seen - The total time since the last response was received from each LDAP server.
 - RTT - The round-trip time to get the response from the server.
3. **To view LDAP server logs:**
 - a. **Click the Logs tab at the top of the LDAP page.**
 - b. **Select an LDAP server from the drop-down menu.**
The log entries contain a time and description for specific LDAP server alerts.

Related Topics

- [“LDAP Properties” on page 316](#)

▼ Monitoring LDAP Server Status (CLI)

To monitor LDAP server status, use the following procedure.

1. **Go to configuration services ldap.**

```
hostname:> configuration services ldap
```

2. **Enter list to show the status of LDAP servers.**

```
hostname:configuration services ldap> list
SERVER      STATUS      LDAP SERVER
server-000   unavailable ldap-server1.us.example.com:636
server-001   online      ldap-server2.us.example.com:484
```

Related Topics

- [“LDAP Properties” on page 316](#)

LDAP Properties

For the appropriate settings for your environment, consult your LDAP server administrator.

Schema

- **Base search DN** - Supplies the distinguished name of the base object which is the starting point for directory searches.
- **Search scope** - Defines which objects in the LDAP directory are searched, relative to the base object. Search results can be limited only to objects directly beneath the base search object (one-level) or they can include any object beneath the base search object (subtree). The default is one-level.
- **Schema definition** - Schema used by the appliance. This property lets administrators override the default search descriptor, attribute mappings, and object class mappings for users, groups, and netgroups. For more information, see [“LDAP Custom Mappings” on page 317](#).

Security Settings

- **Authenticate As** - Credentials used to authenticate the appliance to the LDAP server.
- **Enable SSL/TLS** - Toggles TLS (Transport Layer Security, the descendant of SSL) to establish secure connections to the LDAP server. If authenticating as Self, this option is unavailable because Self uses Kerberos encryption.
- **Authentication method** - Method used to authenticate the appliance to the LDAP server. You can only configure this setting if authenticating as Proxy.

LDAP Servers

- **Servers** - List of LDAP servers to use. If only one server is specified, the appliance uses only that server and LDAP services are unavailable if that server fails. If multiple servers are specified, any functioning server can be used at any time without preference. If any server fails, another server in the list is used. LDAP services remain available unless all specified servers fail. To monitor the configured LDAP servers and their status, see Monitoring LDAP Server Status - [BUI](#), [CLI](#).

Related Topics

- Configuring LDAP Security Settings - [BUI](#), [CLI](#)
- Monitoring LDAP Server Status - [BUI](#), [CLI](#)

LDAP Custom Mappings

To look up users and groups in the LDAP directory, the appliance uses a search descriptor and must know which object classes correspond to users and groups and which attributes correspond to the properties needed. By default, the appliance uses object classes specified by RFC 2307 (*posixAccount* and *posixGroup*) and the default search descriptors shown in the following list, but this can be customized for different environments. The base search DN used in the examples below is `dc=example,dc=com`:

TABLE 68 LDAP Custom Mappings

Search descriptor	Default value	Example
users	<code>ou=people,base search DN</code>	<code>ou=people,dc=example,dc=com</code>
groups	<code>ou=group,base search DN</code>	<code>ou=group,dc=example,dc=com</code>
netgroups	<code>ou=netgroup,base search DN</code>	<code>ou=netgroup,dc=example,dc=com</code>

The search descriptor, object classes, and attributes used can be customized using the Schema definition property. To override the default search descriptor, enter the entire DN you wish to use. The appliance will use this value unmodified, and will ignore the values of the Base search DN and Search scope properties. To override user, group, and netgroup attributes and objects, choose the appropriate tab ("Users", "Groups", or "Netgroups") and specify mappings using the *default = new* syntax, where *default* is the default value and *new* is the value you want to use. For examples:

- To use *unixaccount* instead of *posixAccount* as the user object class, enter `posixAccount = unixaccount` in Object class mappings on the Users tab.
- To use *employeenumber* instead of *uid* as the attribute for user objects, enter `uid = employeenumber` in Attribute mappings on the Users tab.

- To use *unixgroup* instead of *posixGroup* as the group object class, type `posixGroup = unixgroup` in Object class mappings on the Groups tab.
- To use *groupaccount* instead of *cn* as the attribute for group objects, enter `cn = groupaccount` in Attribute mappings on the Groups tab.

The following is a list of object classes and attributes that you might want to map:

Classes

- `posixAccount`
- `posixGroup`
- `shadowAccount`

Attributes - Users

- `uid`
- `uidNumber`
- `gidNumber`
- `gecos`
- `homeDirectory`
- `loginShell`
- `userPassword`

Attributes - Groups

- `uid`
- `memberUid`
- `cn`
- `userPassword`
- `gidNumber`
- `member`
- `uniqueMember`
- `memberOf`
- `isMemberOf`

Related Topics

- [“Adding an Appliance Administrator \(BUI\)” on page 311](#)
- [“Setting Properties with Multiple Attribute Value Pairs \(CLI\)” on page 311](#)
- [“LDAP Properties” on page 316](#)

NDMP Configuration

The Network Data Management Protocol (NDMP) service enables the system to participate in NDMP-based backup and restore operations controlled by a remote NDMP client called a Data Management Application (DMA). Using NDMP, data stored in administrator-created shares on the appliance can be backed up and restored to both locally attached tape devices and remote systems. Locally-attached tape devices can also be exposed to the DMA for backing up and restoring remote systems.

NDMP cannot be used to back up and restore system configuration data. Instead, use the Configuration Backup and Restore feature (see [“Backing Up the Configuration” in Oracle ZFS Storage Appliance Customer Service Manual](#)).

To configure NDMP, see the following sections:

- [“NDMP Local vs. Remote Configurations” on page 319](#)
- [“NDMP Backup Formats and Types” on page 320](#)
- [“NDMP Backup with Types dump and tar” on page 320](#)
- [“NDMP Backup with Type zfs” on page 321](#)
- [“NDMP Incremental Backups” on page 323](#)
- [“Replica Backups” on page 324](#)
- [“NDMP Properties and Logs” on page 326](#)

NDMP Local vs. Remote Configurations

The appliance supports backup and restore using both a *local* configuration, in which tape drives are physically attached to the appliance, and a *remote* configuration, in which data is streamed to another system on the same network. In both cases, the backup must be managed by a supported DMA.

In local configurations, supported tape devices, including both drives and changers (robots), are physically connected to the system using a supported SCSI or Fibre Channel (FC) card configured in Initiator mode. These devices can be viewed on the NDMP Status screen. The NDMP service presents these devices to a DMA when the DMA scans for devices. Once configured in the DMA, these devices are available for backup and restore of the appliance or other systems on the same network. After adding tape drives or changers to the system or removing such devices from the system, a reboot may be required before the changes will be recognized by the NDMP service. After that, the DMA may need to be reconfigured because tape device names may have changed.

In remote configurations, the tape devices are not physically connected to the system being backed up and restored (the data server) but rather to the system running the DMA or a separate system (the tape server). These are commonly called "3-way configurations" because the DMA

controls two other systems. In these configurations, the data stream is transmitted between the data server and the tape server over an IP network.

NDMP Backup Formats and Types

The NDMP protocol does not specify a backup data format. The appliance supports three backup types corresponding to different implementations and on-tape formats. DMAs can select a backup type using the following values for the NDMP environment variable TYPE:

TABLE 69 NDMP Backup Formats and Types

Backup type	Details
dump	File-based for filesystems only. Supports file history and direct access recovery (DAR).
tar	File-based for filesystems only. Supports file history and direct access recovery (DAR).
zfs	Share-based for both filesystems and volumes. Does not support file history or direct access recovery (DAR), but may be faster for some datasets. Only supported with NDMPv4.

There is no standard NDMP data stream format, so backup streams generated on the appliance can only be restored on ZFS storage appliances running compatible software. Future versions of appliance software can generally restore streams backed up from older versions of the software, but the reverse is not necessarily true. For example, the "zfs" backup type is new in 2010.Q3 and systems running 2010.Q1 (or earlier) cannot restore backup streams created using type "zfs" under 2010.Q3.

NDMP Backup with Types dump and tar

When backing up with "dump" and "tar" backup types, administrators specify the data to backup by a filesystem path, called the *backup path*. For example, if the administrator configures a backup of */export/home*, then the share mounted at that path will be backed up. Similarly, if a backup stream is restored to */export/code*, then that's the path where files will be restored, even if they were backed up from another path.

Only paths that are mountpoints of existing shares, or contained within existing shares, may be specified for backup. If the backup path matches a share's mountpoint, only that share is backed up. Otherwise the path must be contained within a share, in which case only the portion of that share under that path is backed up. In both cases, other shares mounted inside the specified share under the backup path will not be backed up; these shares must be specified separately for backup.

Snapshots - If the backup path specifies a live filesystem (such as `/export/code`) or a path contained within a live filesystem (such as `/export/code/src`), the appliance immediately takes a new snapshot and backs up the given path from that snapshot. When the backup completes, the snapshot is destroyed. If the backup path specifies a snapshot (e.g., `/export/code/zfs/snapshot/mysnap`), no new snapshot is created and the system backs up from the specified snapshot.

Share metadata - To simplify backup and restore of complex share configurations, "dump" and "tar" backups include share metadata for projects and shares associated with the backup path. This metadata describes the share configuration on the appliance, including protocol sharing properties, quota properties, and other properties configured on the Shares screen. This is not to be confused with filesystem metadata like directory structure and file permissions, which is also backed up and restored with NDMP.

For example, if you back up `/export/proj`, the share metadata for all shares whose mountpoints start with `/export/proj` will be backed up, as well as the share metadata for their parent projects. Similarly, if you back up `/export/someshare/somedir`, and a share is mounted at `/export/someshare`, that share and its project's share metadata will be backed up.


When restoring, if the destination of the restore path is not contained inside an existing share, projects and shares in the backup stream will be recreated as needed with their original properties as stored in the backup. For example, if you back up `/export/foo`, which contains project `proj1` and shares `share1` and `share2`, and then destroy the project and restore from the backup, then these two shares and the project will be recreated with their backed-up properties as part of the restore operation.

During a restore, if a project exists that would have been automatically recreated, the existing project is used and no new project is automatically created. If a share exists that would have been automatically recreated, and if its mountpoint matches what the appliance expects based on the original backup path and the destination of the restore, then the existing share is used and no new share is automatically created. Otherwise, a new share is automatically created from the metadata in the backup. If a share with the same name already exists (but has a different mountpoint), then the newly created share will be given a unique name starting with `ndmp-` and with the correct mountpoint.

It is recommended that you either restore a stream whose datasets no longer exist on the appliance, allowing the appliance to recreate datasets as specified in the backup stream, or precreate a destination share for restores. Either of these practices avoids surprising results related to the automatic share creation described earlier.

NDMP Backup with Type `zfs`

When backing up with type "zfs", administrators specify the data to backup by its canonical name on the appliance. This can be found underneath the name of the share in the BUI:

FIGURE 20 NDMP Share Name

or in the CLI as the value of the `canonical_name` property. Canonical names do not begin with a leading '/', but when configuring the backup path the canonical name must be prefixed with '/'.

Both projects and shares can be specified for backup using type "zfs". If the canonical name is specified as-is, then a new snapshot is created and used for the backup. A specific snapshot can be specified for backup using the `@snapshot` suffix, in which case no new snapshot is created and the specified snapshot is backed up. For example:

TABLE 70 Canonical Names and Shares Backed Up

Canonical name	Shares backed up
<code>pool-0/local/default</code>	New snapshot of the local project called <code>default</code> and all of its shares.
<code>pool-0/local/default@yesterday</code>	Named snapshot <code>yesterday</code> of local project <code>default</code> , and all of its shares having snapshot <code>yesterday</code> .
<code>pool-0/local/default/code</code>	New snapshot of share <code>code</code> in local project <code>default</code> . <code>code</code> could be a filesystem or volume.
<code>pool-0/local/default/code@yesterday</code>	Named snapshot <code>yesterday</code> of share <code>code</code> in local project <code>default</code> . <code>code</code> could be a filesystem or volume.

Because level-based incremental backups using the "zfs" backup type require a base snapshot from the previous incremental, the default behavior for level backups for which a new snapshot is created is to keep the new snapshot so that it can be used for subsequent incremental backups. If the DMA indicates that the backup will not be used for subsequent incremental backups by setting `UPDATE=n`, the newly created snapshot is destroyed after the backup. Existing user snapshots are never destroyed after a backup. For details, see [“NDMP Incremental Backups” on page 323](#).

Share metadata - Share metadata (i.e., share configuration) is always included in "zfs" backups. When restoring a full backup with type "zfs", the destination project or share must not already exist. It will be recreated from the metadata in the backup stream. When restoring an incremental backup with type "zfs", the destination project or share must already exist. Its properties will be updated from the metadata in the backup stream. For details, see [“NDMP Incremental Backups” on page 323](#).

NDMP Incremental Backups

The appliance supports level-based incremental backups for all of the above backup types. To specify a level backup, DMAs typically specify the following three environment variables:

Variable	Details
LEVEL	Integer from 0 to 9 identifying the backup level.
DMP_NAME	Specifies a particular incremental backup set. Multiple sets of level incremental backups can be used concurrently by specifying different values for DMP_NAME.
UPDATE	Indicates whether this backup can be used as the base for subsequent incremental backups

By definition, a level-N backup includes all files changed since the previous backup of the same backup set (specified by "DMP_NAME") of the same share using LEVEL less than N. Level-0 backups always include all files. If UPDATE has value "y" (the default), then the current backup is recorded so that future backups of level greater than N will use this backup as a base. These variables are typically managed by the DMA and need not be configured directly by administrators.

Below is a sample incremental backup schedule:

TABLE 71 Sample Incremental Backup Schedule

Day	Details
First of month	Level-0 backup. Backup contains all files in the share.
Every 7th, 14th, 21st of month	Level-1 backup. Backup contains all files changed since the last full (monthly) backup
Every day	Level-2 backup. Backup contains all files changed since the last level-1 backup

To recover the filesystem's state as it was on the 24th of the month, an administrator typically restores the Level-0 backup from the 1st of the month to a new share, then restores the Level-1

backup from the 21st of the month, and then restores the Level-2 backup from the 24th of the month.

To implement level-based incremental backups the appliance must keep track of the level backup history for each share. For "tar" and "dump" backups, the level backup history is maintained in the share metadata. Incremental backups traverse the filesystem and include files modified since the time of the previous level backup. At restore time, the system simply restores all the files in the backup stream. In the above example, it would therefore be possible to restore the Level-2 backup from the 24th onto any filesystem and the files contained in that backup stream will be restored even though the target filesystem may not match the filesystem where the files were backed up. However, best practice suggests using a procedure like the above which starts from an empty tree restores the previous level backups in order to recover the original filesystem state.

To implement efficient level-based incremental backups for type "zfs", the system uses a different approach. Backups that are part of an incremental set do not destroy the snapshot used for the backup but rather leave it on the system. Subsequent incremental backups use this snapshot as a base to quickly identify the changed filesystem blocks and generate the backup stream. As a consequence, the snapshots left by the NDMP service after a backup must not be destroyed if you want to create subsequent incremental backups.

Another important consequence of this behavior is that in order to restore an incremental stream, the filesystem state must exactly match its state at the base snapshot of the incremental stream. In other words, in order to restore a level-2 backup, the filesystem must look exactly as it did when the previous level-1 backup completed. Note that the above commonly-used procedure guarantees this because when restoring the Level-2 backup stream from the 24th, the system is exactly as it was when the Level-1 backup from the 21st completed because that backup has just been restored.

The NDMP service will report an error if you attempt to restore an incremental "zfs" backup stream to a filesystem whose most recent snapshot doesn't match the base snapshot for the incremental stream, or if the filesystem has been changed since that snapshot. You can configure the NDMP service to rollback to the base snapshot just before the restore begins by specifying the NDMP environment variable "ZFS_FORCE" with value "y" or by configuring the "Rollback datasets" property of the NDMP service (see [“NDMP Properties and Logs” on page 326](#)).

Replica Backups

The Oracle ZFS Storage Appliance product supports direct backup of replicas and replica snapshots with the "zfs" backup type. It is not necessary to first clone a replica dataset (project or share) in order to back it up.

Note - Because the backup is of a replica, the source dataset properties are backed up rather than those of the target.

Enabling Replica Backups

- To enable replica backups, apply the corresponding deferred update. For more information, see [“Deferred Updates” in Oracle ZFS Storage Appliance Customer Service Manual](#).
- Replica backups require software version 2011.1.0 (or later on the source).
- If the replica backup will be restored to the source with the original replicated dataset, the source must run software version 2013.1.4 (or later).

Replica Backup Syntax

To back up a replicated project or share, input the ZFS dataset name without a snapshot extension into the DMA. ndmpd uses the appliance software to determine the latest complete replica snapshot to back up. To specify a replica dataset for backup, use copy and paste to avoid mistyping long replica dataset names which may include a UUID.

If a user-generated snapshot extension is included, ndmpd backs up the indicated user snapshot. If a system-generated extension is included (begins with .rr) the backup fails and generates a message that is logged to the DMA console.

Replica Backup Persistent Holds

Persistent holds are taken on backed-up snapshots when the backups complete. This is necessary for future incremental backups, which use current snapshots as a base, otherwise the replication subsystem may delete replica snapshots it no longer needs. Holds are released by ndmpd when the snapshots are no longer needed.

Persistent holds can be cleared manually. When deleting a replica snapshot with a hold on it, a confirmation is displayed warning of the potential impact to ongoing or future NDMP backups. Snapshots required by the replication subsystem cannot be deleted.

If incremental backups are not needed, prevent persistent holds by setting the DMA UPDATE parameter to no (UPDATE=n). UPDATE=y is the default mode. For more information about the UPDATE NDMP environment variable, see the whitepaper [NDMP Implementation Guide for the Sun ZFS Storage Appliance \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html).

Incremental Replica Backups

Continuing (incrementing) a backup series across a replication reversal or sever is not supported; instead, start a new backup series. Use a full (Level-0) backup for the first backup after a replication state change has occurred, such as on a new source after a reversal or sever has taken place.

Static snapshot extensions that do not change per level are NOT supported for user-generated replica snapshots (snapshots not starting with `.rr`). This prevents name collisions, which generate an error and can cause replication to fail.

Some DMAs do not support zfs-type replica incremental backup and restore operations for snapshot extension name changes per level. To conserve appliance space and ensure that such snapshots are not preserved for future incremental backups, set `UPDATE=n` at the time of backup of replica. User-generated snapshots can be removed manually.

Even if no user data has changed in a restored dataset, changed metadata can cause incremental replica restores to fail. To avoid this, always roll back to the base snapshot before incremental replica restores by setting the `ZFS rollback before restore` parameter to `Always`.

For non-incremental replica backups, such as for one-off backups, set `UPDATE=n` so future snapshots are not saved and consume space. Some older replica snapshots preserved for future incremental backups, such as those created by setting `UPDATE=y`, may no longer be needed and waste space. These snapshots are safe to manually destroy. Snapshots needed by the replication subsystem cannot be deleted. Unneeded snapshots can be deleted after confirming the warning message about possible impacts to ongoing or future NDMP backups if the snapshot is deleted.

NDMP Properties and Logs

The NDMP service configuration consists of the following properties and logs:

TABLE 72 NDMP Properties

Property	Description
Version	The version of NDMP that your DMA supports.
TCP port (v4 only)	The NDMP default connection port is 10000. NDMPv3 always uses this port. NDMPv4 allows a different port if needed.
Ignore metadata-only changes	Directs the system to backup only files in which content has changed, ignoring files for which only metadata, such as permissions or ownership, has changed. This option only applies to incremental "tar" and "dump" backups and is disabled by default.
Target restore pool(s)	When you perform a full restore using "tar" or "dump", the system re-creates datasets if there is no share

Property	Description
	mounted at the target. Because the NDMP protocol specifies only the mount point, the system chooses a pool in which to recreate projects and shares. On a system with multiple pools, this property lets you specify one or more pools. Multiple pools only need to be specified in a cluster with active pools on each head. You must ensure that this list is kept in sync with any storage configuration changes. If none of the pools exist or are online, the system will select a default pool at random.
Allow token-based backup	Enables or disables token-based method for ZFS backup. This property is off by default.
ZFS rollback before restore (v4 only)	Only applies to backups with type "zfs". Determines whether when restoring an incremental backup the system rolls back the target project and share to the snapshot used as the base for the incremental restore. If the project and shares are rolled back, then any changes made since that snapshot will be lost. This setting is normally controlled by the DMA via the "ZFS_FORCE" environment variable, but this property can be used to override the DMA setting to always rollback these data sets or never roll them back. For details, see “NDMP Incremental Backups” on page 323 . Not rolling them back will cause the restore to fail unless they have already been manually rolled back. This property is intended for use with DMAs that do not allow administrators to configure custom environment variables like ZFS_FORCE.
Allow direct access recovery	Enables the system to locate files by position rather than by sequential search during restore operations. Enabling this option reduces the time it takes to recover a small number of files from many tapes. You must specify this option at backup time in order to be able to recover individual files later.
Restore absolute paths (v3 only)	Specifies that when a file is restored, the complete absolute path to that file is also restored (instead of just the file itself). This option is disabled by default.
Share creation on restore	Configures the restore operation to create a new share based on the backup type: ■ All - Allows all backup types to create a new share on restore ■ Tar_Dump - Allows backup types "tar" and "dump" to create a new share on restore ■ ZFS - Allows backup type "zfs" to create a new share on restore ■ None - No backup type can create a new share on restore
DMA tape mode (for locally attached drives)	Specifies whether the DMA expects System V or BSD semantics. The default is System V, which is recommended for most DMAs. This option is only

Property	Description
	applicable for locally attached tape drives exported via NDMP. Consult your DMA documentation for which mode your DMA expects. Changing this option only changes which devices are exported when the DMA scans for devices, so you will need to reconfigure the tape devices in your DMA after changing this setting.
DMA username and password	Used to authenticate the DMA. The system uses MD5 for user authentication

TABLE 73 NDMP Logs

Log	Description
system-ndmpd:default	NDMP service log

NFS Configuration

Network File System (NFS) is an industry standard protocol to share files over a network. The Oracle ZFS Storage Appliance supports NFS versions 2, 3, 4.0 and 4.1. For more information on how the filesystem namespace is constructed, see [“Working with Filesystem Namespace” on page 439](#). For information about NFS with local users, see [“Configuring Users” on page 202](#).



Caution - To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

To configure NFS, see the following sections:

- [“NFS Service Properties” on page 328](#)
- [“Configuring Kerberos Realms for NFS” on page 330](#)
- [“NFS Logs and Analytics” on page 331](#)
- [“NFS Properties” on page 331](#)
- [“NFS Naming Service Dependencies” on page 332](#)
- [“Sharing a Filesystem Over NFS” on page 333](#)

NFS Service Properties

The following NFS Service properties are available in Configuration > Services. Note that NFSv4 is also known as NFSv4.0.

- **Minimum supported version** - Use this drop-down list to control which versions of NFS the appliance supports.

- **Maximum supported version** - Use this drop-down list to control which versions of NFS the appliance supports.

Note - Setting the NFS minimum and maximum versions to the same value causes the appliance to only communicate with clients using that version. This may be useful if you find an issue with one NFS version or the other (such as the performance characteristics of an NFS version with your workload), and you want to force clients to only use the version that works best.

- **Maximum # of server threads** - Define the maximum number of concurrent NFS requests (from 20 to 3000). This should at least cover the number of concurrent NFS clients that you anticipate. The default value is 1500.
- **Grace period** - Define the number of seconds that all clients have to recover locking state after an appliance reboot (from 15 to 600 seconds) from an unplanned outage. This property affects only NFSv4.0 and NFSv4.1 clients (NFSv3 is stateless so there is no state to reclaim). During this period, the NFS service only processes reclaims of the old locking state. No other requests for service are processed until the grace period is over. The default grace period is 90 seconds. Reducing the grace period lets NFS clients resume operation more quickly after a server reboot, but increases the probability that a client cannot recover all of its locking state. The Oracle ZFS Storage Appliance provides grace-less recovery of the locking state for NFSv4.0 and NFSv4.1 clients during planned outages. Planned outages occur during events such as updates and appliance reboot using the CLI maintenance system reboot command or the BUI power icon . For planned outages, the NFS service processes all requests for service without incurring the grace period delay.
- **Custom NFSv4 identity domain** - Use this property to define the domain for mapping NFSv4.0 and NFSv4.1 users and group identities. If you do not set this property, the appliances uses DNS to obtain the identity domain, first by checking for a _nfsv4idmapdomain DNS resource record, and then by falling back to the DNS domain itself.
- **Use NFSv4 numeric id strings** - Use this property to allow NFSv4.0 and NFSv4.1 clients to use numeric strings for user and group IDs. If you do not set this property, user and group IDs are exchanged in the form of user@domain, the default. This property applies only when the authentication type is AUTH_SYS. The CLI property is use_numeric_ids.
- **Enable NFSv4 delegation** - Select this property to allow clients to cache files locally and make modifications without contacting the server. This option is enabled by default and typically results in better performance; but in rare circumstances it can cause problems. You should only disable this setting after careful performance measurements of your particular workload and after validating that the setting has a measurable performance benefit. This option only affects NFSv4.0 and NFSv4.1 mounts.
- **Mount visibility** - This property lets you limit the availability of information about share access lists and remote mounts from NFS clients. Full allows full access. Restricted restricts

access such that a client can see only the shares which it is allowed to access. A client cannot see access lists for shares defined at the server or remote mounts from the server done by other clients. The property is set to Full by default.

- **Oracle Intelligent Storage Protocol** - The NFSv4.0 and NFSv4.1 services include support for the Oracle Intelligent Storage Protocol, which lets Oracle Database NFSv4.0 and NFSv4.1 clients pass optimization information to the ZFS Storage Appliance NFSv4.0 and NFSv4.1 server. For more information, see [“Oracle Intelligent Storage Protocol” on page 688](#).

Related Topics

- [“NFS Properties” on page 331](#)
- Setting Service Properties [BUI](#), [CLI](#).

Configuring Kerberos Realms for NFS

Configuring a Kerberos realm creates certain service principals and adds the necessary keys to the system's local keytab. The NTP service must be configured before configuring Kerberized NFS. The following service principals are created and updated to support Kerberized NFS:

```
host/node1.example.com@EXAMPLE.COM
nfs/node1.example.com@EXAMPLE.COM
```

If you clustered your appliances, principals and keys are generated for each cluster node:

```
host/node1.example.com@EXAMPLE.COM
nfs/node1.example.com@EXAMPLE.COM
host/node2.example.com@EXAMPLE.COM
nfs/node2.example.com@EXAMPLE.COM
```

If these principals have already been created, configuring the realm resets the password for each of those principals.

For information on setting up KDCs and Kerberized clients, see [Oracle Solaris 11.1 Administration: Security Services](#) (http://docs.oracle.com/cd/E26502_01/html/E29015/index.html). For information about the appliance Kerberos service, see [“Kerberos Configuration” on page 293](#). After configuring Kerberos, change the Security mode on the Shares->Filesystem->Protocols screen to a mode using Kerberos.

Note - Kerberized NFS clients must access the appliance using an IP address that resolves to an FQDN for those principals. For example, if an appliance is configured with multiple IP addresses, only the IP address that resolves to the appliance's FQDN can be used by its Kerberized NFS clients.

NFS Logs and Analytics

These logs are available for the NFS service:

TABLE 74 Logs Available for NFS

Log	Description
network-nfs-server:default	Master NFS server log
appliance-kit-nfsconf:default	Log of appliance NFS configuration events
network-nfs-cbd:default	Log for the NFSv4.0 and NFSv4.1 callback daemon
network-nfs-mapid:default	Log for the NFSv4.0 and NFSv4.1 mapid daemon - which maps NFSv4.0 and NFSv4.1 user and group credentials
network-nfs-status:default	Log for the NFS statd daemon - which assists crash and recovery functions for NFS locks
network-nfs-nlockmgr:default	Log for the NFS lockd daemon - which supports record locking operations for files

You can monitor NFS activity in the Analytics section. This includes:

- NFS operations per second
- ... by type of operation (read/write/...)
- ... by share name
- ... by client hostname
- ... by accessed filename
- ... by access latency

NFS Properties

The following table describes the mapping between CLI properties and the BUI property descriptions above.



Caution - To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

TABLE 75 NFS Properties

CLI Property	BUI Property
version_min	Minimum supported version
version_max	Maximum supported version
nfsd_servers	Maximum # of server threads

CLI Property	BUI Property
grace_period	Grace period
mapid_domain	Custom NFSv4 and NFSv4.1 identity domain
use_numeric_ids	Use NFSv4 and NFSv4.1 numeric string ids
enable_delegation	Enable NFSv4 and NFSv4.1 delegation
mount_visibility	Client share information restriction level

NFS Naming Service Dependencies

Naming services, such as DNS, NIS, and LDAP, are used by the appliance to resolve host names and corresponding IP addresses, user identities, and Analytics statistics. This topic describes NFS naming service dependencies and resulting problems if naming services are not configured for the appliance.

NFS depends on information from each of the following naming services:

TABLE 76 NFS Naming Service Dependencies

Description	Service
IP address and corresponding host name of NFS clients and servers	DNS
User identity number and corresponding user name	NIS/LDAP
Group identity number and corresponding group name	NIS/LDAP
Clients belonging to netgroups	NIS/LDAP

If the appliance is unable to access any DNS network servers or DNS mappings are unpopulated, the following problems can occur:

- Filesystem mount failure
- Client is denied access to NFS shares after the filesystem has been successfully mounted
- Client receives "weak authentication" errors
- NFS server unresponsive
- User and group lookup failures, using either NFSv3 or NFSv4, as listed in the following table.

NFS Version	Problem	Setting
NFSv3 or NFSv4	Cannot access particular files or directories when the user is a member of 16 or more groups.	
NFSv4	Cannot retrieve ownership information (users and groups might be shown as nobody).	Set the "Use NFSv4 numeric id strings" option.

NFS Version	Problem	Setting
NFSv4	Cannot change ownership of files.	On the client, set the equivalent of the "Use NFSv4 numeric id strings" option.
NFSv4	Cannot retrieve ACL information.	Set the "Use NFSv4 numeric id strings" option.
NFSv4	Cannot change ACLs, including inability to change entries unrelated to the affected entries.	Set the "Use NFSv4 numeric id strings". On the client, set the equivalent of the "Use NFSv4 numeric id strings" option.

Related Topics

- [“DNS Configuration” on page 265](#)
- [“DNS-Less Operation” on page 272](#)

▼ Sharing a Filesystem Over NFS

1. **Go to Configuration > Services.**
2. **Check that the NFS service is enabled and online. If not, enable the service.**



Caution - To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

3. **Go to the Shares screen and edit an existing share or create a new share.**
4. **Click the Protocols tab of the share you are editing and check that NFS sharing is enabled.**

You can also configure the NFS share mode (read/read+write) in this screen.

NIS Configuration

Network Information Service (NIS) is a name service for centralized management. The appliance can act as an NIS client for users and groups, so that:

- NIS users can log in to the FTP and HTTP services.
- NIS users can be granted privileges for appliance administration. The appliance supplements NIS information with its own privilege settings.

Note - UIDs and GIDs from 0-99 inclusive are reserved by the operating system vendor for use in future applications. Their use by end system users or vendors of layered products is not supported and may cause security related issues with future applications.

To configure NIS, see the following sections:

- [“Adding an Appliance Administrator from NIS \(BUI\)” on page 334](#)
- [“NIS Properties and Logs” on page 334](#)

▼ Adding an Appliance Administrator from NIS (BUI)

If you have an existing user in NIS who would like to log in using their NIS credentials and administer the appliance:

Note - If both NIS and LDAP are configured on the appliance and the services return different information for a particular item, the appliance will use the data provided by NIS.

1. **Go to Configuration > Services > NIS.**
2. **Set the NIS domain and server properties.**
3. **Click APPLY to commit the configuration.**
4. **Go to Configuration > Users.**
5. **Add a user with type "directory".**
6. **Set the username to their NIS username.**
7. **Continue with the instructions in [“Configuring Users” on page 202](#) for adding authorizations to this user.**

Related Topics

- [“NIS Properties and Logs” on page 334](#)

NIS Properties and Logs

TABLE 77 NIS Properties

Property	Description
Domain	The NIS domain to use.

Property	Description
Server(s): Search using broadcast	The appliance sends a NIS broadcast to locate NIS servers for that domain.
Server(s): Use listed servers	NIS server hostnames or IP addresses.

The appliance will connect to the first NIS server listed, or a server found using broadcast, and switch to the next if it stops responding.

TABLE 78 NIS Logs

Log	Description
network-nis-client:default	NIS client service log.
appliance-kit-nsswitch:default	Log of the appliance name service, through which NIS queries are made.
system-identity:domain	Log of the appliance domain name configurator.

Related Topics

- [“Adding an Appliance Administrator from NIS \(BUI\)” on page 334](#)

NTP Configuration

The Network Time Protocol (NTP) service can be used to keep the appliance clock accurate. This is important for recording accurate timestamps in the filesystem, and for protocol authentication. The appliance records times using the UTC timezone. The times that are displayed in the BUI use the timezone offset of your browser.

To the right of the BUI screen are times from both the appliance (Server Time) and your browser (Client Time). If the NTP service is not online, the SYNC button can be clicked to set the appliance time to match your client browser time.

If you are sharing filesystems using SMB, the client clocks must be synchronized to within five minutes of the appliance clock to avoid user authentication errors. One way to ensure clock synchronization is to configure the appliance and the SMB clients to use the same NTP server.

TABLE 79 NTP Clock Synchronization

Log	Description
network-ntp:default	Log for the NTP service

To configure NTP, see the following sections:

- [“Setting Clock Synchronization \(BUI\)” on page 336](#)
- [“Configuring NTP \(CLI\)” on page 336](#)
- [“NTP Properties” on page 337](#)

▼ Setting Clock Synchronization (BUI)

This will set the appliance time to match the time of your browser.

1. **Go to Configuration > Services > NTP.**
2. **Disable the NTP service.**
3. **Click SYNC.**

▼ Configuring NTP (CLI)

1. **Under configuration services ntp, edit authorizations with the authkey command:**

```
hostname:configuration services ntp> authkey
hostname:configuration services ntp authkey>
```

2. **From this context, new keys can be added with the create command:**

```
hostname:configuration services ntp authkey> create
hostname:configuration services ntp authkey-000 (uncommitted)> get
                                keyno = (unset)
                                type = (unset)
                                key = (unset)
hostname:configuration services ntp authkey-000 (uncommitted)> set keyno=1
                                keyno = 1 (uncommitted)
hostname:configuration services ntp authkey-000 (uncommitted)> set type=A
                                type = A (uncommitted)
hostname:configuration services ntp authkey-000 (uncommitted)> set key=coconuts
                                key = ***** (uncommitted)
hostname:configuration services ntp authkey-000 (uncommitted)> commit
hostname:configuration services ntp authkey>
```

3. **To associate authentication keys with servers via the CLI, the serverkeys property should be set to a list of values in which each value is a key to be associated with the corresponding server in the servers property.**

If a server does not use authentication, the corresponding server key should be set to 0. For example, to use the key created above to authenticate the servers "gefilte" and "carp":

```
hostname:configuration services ntp> set servers=gefilte,carp
      servers = gefilte,carp (uncommitted)
hostname:configuration services ntp> set serverkeys=1,1
      serverkeys = 1,1 (uncommitted)
hostname:configuration services ntp> commit
hostname:configuration services ntp>
```

4. **To associate authentication keys with servers, set the serverkeys property to a list of values in which each value is a key to be associated with the corresponding server in the servers property.**

If a server does not use authentication, the corresponding server key should be set to 0. For example, to use the key created above to authenticate the servers "gefilte" and "carp":

```
hostname:configuration services ntp> set servers=gefilte,carp
      servers = gefilte,carp (uncommitted)
hostname:configuration services ntp> set serverkeys=1,1
      serverkeys = 1,1 (uncommitted)
hostname:configuration services ntp> commit
hostname:configuration services ntp>
```

5. **To authenticate the server "gefilte" with key 1, "carp" with key 2 and "dory" with key 3:**

```
hostname:configuration services ntp> set servers=gefilte,carp,dory
      servers = gefilte,carp,dory (uncommitted)
hostname:configuration services ntp> set serverkeys=1,2,3
      serverkeys = 1,2,3 (uncommitted)
hostname:configuration services ntp> commit
hostname:configuration services ntp>
```

6. **To authenticate the servers "gefilte" and "carp" with key 1, and to additionally have an unauthenticated NTP server "dory":**

```
hostname:configuration services ntp> set servers=gefilte,carp,dory
      servers = gefilte,carp,dory (uncommitted)
hostname:configuration services ntp> set serverkeys=1,1,0
      serverkeys = 1,1,0 (uncommitted)
hostname:configuration services ntp> commit
hostname:configuration services ntp>
```

NTP Properties

The following NTP properties are available at Configuration > Services > NTP:

TABLE 80 NTP Properties

Property	Description	Examples
Discover NTP server via multicast address	Enter a multicast address here for an NTP server to be located automatically	224.0.1.1
Manually specify NTP server(s)	Enter one or more NTP servers (and their corresponding authentication keys, if any) for the appliance to contact directly	0.pool.ntp.org
NTP Authentication Keys	Enter one or more NTP authentication keys for the appliance to use when authenticating the validity of NTP servers. See Table 81, “NTP Private Keys and Integers,” on page 338.	Auth key: 10, Type: ASCII, Private Key: SUN7000

Validation - If an invalid configuration is entered, a warning message is displayed and the configuration is not committed. This will happen if:

- A multicast address is used but no NTP response is found.
- An NTP server address is used, but that server does not respond properly to NTP.

Authentication - To prevent against NTP spoofing attacks from rogue servers, NTP has a private key encryption scheme whereby NTP servers are associated with a private key that is used by the client to verify their identity. These keys are not used to encrypt traffic, and they are not used to authenticate the client -- they are only used by the NTP client (that is, the appliance) to authenticate the NTP server. To associate a private key with an NTP server, the private key must first be specified. Each private key has a unique integer associated with it, along with a type and key. The type must be one of the following:

TABLE 81 NTP Private Keys and Integers

Type	Description	Example
DES	A 64-bit hexadecimal number in DES format	0101010101010101
NTP	A 64-bit hexadecimal number in NTP format	8080808080808080
ASCII	A 1-to-8 character ASCII string	topsecret
MD5	A 1-to-8 character ASCII string, using the MD5 authentication scheme.	md5secret

After the keys have been specified, an NTP server can be associated with a particular private key. For a given key, all of the key number, key type, and private key values must match between client and server for an NTP server to be authenticated.

Phone Home Configuration

The Phone Home service screen is used to manage the appliance registration as well as the Phone Home remote support service.

Registration connects your appliance with [Oracle Auto Service Request \(http://www.oracle.com/us/support/auto-service-request/index.html\)](http://www.oracle.com/us/support/auto-service-request/index.html). Oracle ASR automatically opens Service Requests (SR) for specific problems reported by your appliance. Registration also connects your appliance with My Oracle Support (MOS) to detect update notifications.

The Phone Home service communicates with Oracle support to provide:

- **Fault reporting** - The system reports active problems to Oracle for automated service response. Depending on the nature of the fault, a support case may be opened. Details of these events can be viewed in the Active Problem Display. For more information, see [“Working with Problems” in Oracle ZFS Storage Appliance Customer Service Manual](#).
- **Heartbeats** - Daily heartbeat messages are sent to Oracle to indicate that the system is up and running. Oracle support may notify the technical contact for an account when one of the activated systems fails to send a heartbeat for too long.
- **System configuration** - Periodic messages are sent to Oracle describing current software and hardware versions and configuration as well as storage configuration. No user data or metadata is transmitted in these messages.
- **Support bundles** - The Phone Home service must be enabled before support bundles can be uploaded to Oracle Support. See [“Working with Support Bundles” in Oracle ZFS Storage Appliance Customer Service Manual](#) for more information.
- **Update Notifications** - Creates an Alert when new software updates are available on My Oracle Support (MOS). See [“Working with Software Notifications and Updates” in Oracle ZFS Storage Appliance Customer Service Manual](#) for more information.

You must register to use the Phone Home service.

You need a valid Oracle Single Sign-On account user name and password to use the fault reporting and heartbeat features of the Phone Home service. Go to [My Oracle Support \(http://support.oracle.com\)](http://support.oracle.com) and click Register to create your account.

To configure Phone Home, see the following sections:

- Registering the Appliance [BUI, CLI](#)
- [“Changing Account Information \(BUI\)” on page 341](#)
- [“Phone Home Properties” on page 341](#)

▼ Registering the Appliance (BUI)

1. **Go to Configuration > Services > Phone Home.**
2. **Enter your Oracle Single Sign-On Account user name and password.**
Click Privacy Statement for information about privacy policy. It can be viewed at any time in both the BUI and CLI.
3. **Click APPLY to commit your changes.**
4. **Use My Oracle Support (<http://support.oracle.com/>) to complete Auto Service Request (ASR) (<http://www.oracle.com/us/support/auto-service-request/index.html>) activation.**

Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

▼ Registering the Appliance (CLI)

1. **Go to configuration services scrk.**
2. **Set soa_id and soa_password to the user name and password for your Oracle Single Sign-On Account, respectively.**
3. **Commit your changes.**
4. **Use My Oracle Support (<http://support.oracle.com/>) to complete Auto Service Request (ASR) (<http://www.oracle.com/us/support/auto-service-request/index.html>) activation.**

Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

Example 15 CLI Registration

```
hostname:> configuration services scrk
hostname:configuration services scrk>set soa_id=myuser
soa_id = myuser(uncommitted)
hostname:configuration services scrk> set soa_password=mypass
soa_password = ***** (uncommitted)
hostname:configuration services scrk> commit
```


▼ Changing Account Information (BUI)

1. **Go to Configuration > Services > Phone Home.**
2. **Click 'Change account...' to change the Oracle Single Sign-On Account used by the appliance.**
3. **Commit your changes.**
4. **Use My Oracle Support to complete Auto Service Request (ASR) activation.**
Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

Phone Home Properties

If the appliance is not directly connected to the Internet, you may need to configure an HTTP proxy through which the Phone Home service can communicate with Oracle. These proxy settings will also be used to upload support bundles. See [“Working with Support Bundles” in Oracle ZFS Storage Appliance Customer Service Manual](#) for more details on support bundles.

TABLE 82 Phone Home Web Proxy Settings

Property	Description
Use web proxy	Connect via a web proxy
Host : port	Web proxy hostname or IP address, and port
Username	Web proxy username
Password	Web proxy password

TABLE 83 Phone Home Status

Property	Description
Last heartbeat sent at	Time last heartbeat was sent to Oracle support

If the Phone Home service is enabled before a valid Oracle Single Sign-On account has been entered, it will appear in the maintenance state. You must enter a valid Oracle Single Sign-On account to use the Phone Home service.

There is a log of Phone Home events in Maintenance > Logs > Phone Home.

RESTful API Configuration

The Oracle ZFS Storage Appliance RESTful API lets you manage the appliance using simple requests such as GET, PUT, POST, and DELETE HTTP against managed resource URL paths.

The appliance RESTful based architecture is defined as a layered client-server model. Advantages of this model mean that services can be transparently redirected through standard hubs, routers, and other network systems without client configuration. This architecture supports caching of information and is useful when many clients request the same static resources.

For complete Oracle ZFS Storage Appliance RESTful API documentation, see [Oracle ZFS Storage Appliance RESTful API Guide](#).

Service Tags Configuration

Service Tags are used to facilitate product inventory and support, by allowing the appliance to be queried for data such as:

- System serial number
- System type
- Software version numbers

You can register the service tags with Oracle support, allowing you to easily keep track of your Oracle equipment and also expedite service calls. The service tags are enabled by default.

TABLE 84 UDP/TCP Port Properties

Property	Description
Discovery Port	UDP port used for service tag discovery. Default is 6481
Listener Port	TCP port used to query service tag data. Default is 6481

SFTP Configuration

The SFTP (SSH File Transfer Protocol) service allows filesystem access from SFTP clients. Anonymous logins are not allowed, users must authenticate with whichever name service is configured in Services.

SFTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see [“Kerberos Configuration” on page 293](#).

For added security when configuring SFTP, you can specify the ciphers and MACs, as described in [“SFTP Properties, Ports, and Logs” on page 345](#).

To configure SFTP, see the following sections:

- [“Adding SFTP Access to a Share \(BUI\)” on page 343](#)
- [“Configuring SFTP for Remote Access \(CLI\)” on page 343](#)
- [“SFTP Properties, Ports, and Logs” on page 345](#)

▼ Adding SFTP Access to a Share (BUI)

1. Go to **Configuration > Services**.
2. Check that the **SFTP service** is enabled and online. If not, enable the service.
3. Go to **Shares > Shares** and select or add a share.
4. Go to the **Protocols** tab, and check that **SFTP access** is enabled.
5. (Optional) Set the Share mode access to **Read only** or **Read/Write**.

Related Topics

- [“Configuring SFTP for Remote Access \(CLI\)” on page 343](#)
- [“SFTP Properties, Ports, and Logs” on page 345](#)

▼ Configuring SFTP for Remote Access (CLI)

1. Create a local user or network user (LDAP or NIS) with an appropriate administrator role. (See [“Configuring Users” on page 202](#)).
2. Generate an SSH authentication key by entering the command `ssh-keygen -t dsa` on the Oracle Solaris host/client.
3. Enter a file name in which to store the key.
4. Enter a passphrase if required, or leave this field blank to log on directly to the SFTP share.

The location is displayed for the key. The key looks similar to the following:

```
ssh-dss AAAAB3NzaC1kc3MAAACBAPMMs5h8UWk1NPf/  
VJDDEo0OAwt+s6iZxkCmmrgAmLfTX9izWk+ bsvNldOlXN/6EgkusLjo/
```

```
+UaEt5+704vMHClRaq3AlVHLS5tVjeX3iCs+fDo0qwXZg3Brh8QBAaWk3
ywr2osuII1tHh4v/HwEAHZq5mVWXav0pO3bgmxl0/
+VAAAAFQDIJxnm52DfyEdQQMTY+jRVvzGwMQA AAlAhTP6Ey
+2gGFICKkvUofsc04d8pbqH8duE9P6Y88s0+opuj52GkAdRUt2fRrdM9Cf3h4lIOc8Bw9
bZIBzrCKBNWBUDZG56tsfLdilW6vS6gxKrmL2v7fSp9WYPsxZGhOLfU29zW4n2WVcVHbGyFEoVe
+taq aq+AYJaWoHnjZL1/
LpQAAAIaOLc8+uc3hDOcK3pAkYdg8b2rYIGOAZU4py0rq24DGPeVHd5h5jbe4p
WDM70uYqGCOPYiOKEoMnJpczRX5qjI+BfoUY4sH24WWwsKkT8XX9PUAa0WT
+7axEqg2N6YelaTJ95J vMaj6E7HkAIra2Sj2H/LSDktL42UL+j1Wx5A== username sunray
```

5. **Go to Configuration > Services > SFTP. Under Keys, click the plus (+) sign.**
6. **In the New Key window, select DSA.**
7. **Copy only the key portion (beginning with AAAA and ending with Wx5A== in the example above) and paste into the Key field.**

Note - The key should not contain any white spaces.

8. **Enter the user name and add a comment as a reminder.**
9. **Go to Shares > Shares and click the add item icon  to create a filesystem.**
10. **In the Create Filesystem window, enter the filesystem name (for example, sftp), change the permissions to Read/ Write for the share, and click APPLY.**
11. **Click the edit icon  to set up the share properties. (See [“Filesystem Properties” on page 423.](#))**
12. **To access the share, use the `sftp` command as shown in these examples:**

```
sftp -o "port=218" <username> 10.x.x.151:/export/sftp
Connecting to 10.x.xx.151...
Changing to: /export/sftp
sftp>
```

Example with -v option:

```
sftp -v -o "IdentityFile=/home/<username>/.ssh/id_dsa" -o "port=218"
root 10.x.xx.151:/export/sftp
```

Related Topics

- [“Adding SFTP Access to a Share \(BUI\)” on page 343](#)

- [“SFTP Properties, Ports, and Logs” on page 345](#)

SFTP Properties, Ports, and Logs

SFTP Properties

TABLE 85 SFTP Properties

Property	Description
Port (for incoming connections)	The port SFTP listens on. The default is 218.
Permit root login	Allows SFTP logins for the root user. This property is off by default.
Logging level	The verbosity of SFTP log messages
SFTP Keys	RSA/DSA public keys for SFTP authentication. Text comments can be associated with the keys to help administrators track why they were added. As of the 2011.1 software release, key management for SFTP has changed to increase security. When creating an SFTP key, it is required to include the user property with a valid user assignment. SFTP keys are grouped by user and are authenticated via SFTP with the user's name. It is recommended to recreate any existing SFTP keys that do not include the user property, even though they will still authenticate.

TABLE 86 SFTP Security Properties

Property	Description
Ciphers	Ciphers for SFTP connections.
MACs	Message authentication codes (MACs) for SFTP connections.

SFTP Ports

The SFTP service uses a non-standard port number for connections to the appliance. This is to avoid conflicts with administrative SSH connections to port 22. By default, the SFTP port is 218 and must be specified on the SFTP client prior to connecting. For example, an Oracle Solaris client using SFTP, would connect with the following command:

```
manta# sftp -o "Port 218" root@guppy
```

SFTP Logs

network-sftp:default - Logs SFTP service events

Related Topics

- [“Adding SFTP Access to a Share \(BUI\)” on page 343](#)
- [“Configuring SFTP for Remote Access \(CLI\)” on page 343](#)

Shadow Migration Configuration

The shadow migration service allows for automatic migration of data from external or internal sources. This functionality is described in detail in [“Shadow Migration” on page 469](#). The service itself only controls automatic background migration. Regardless of whether the service is enabled or not, data will be migrated synchronously for in-band requests.

The service should only be disabled for testing purposes, or if the load on the system due to shadow migration is too great. When disabled, no filesystems will ever finish migrating. The primary purpose of the service is to allow tuning of the number of threads dedicated to background migration.

Number of Threads Property - Number of threads to devote to background migration of data. These threads are global to the entire machine, and increasing the number can increase concurrency and the overall speed of migration at the expense of increased resource consumption (network, I/O, and CPU).

SMB Configuration

The SMB service provides access to filesystems using the SMB protocol. The supported SMB versions are: SMB 1, SMB 2.0, SMB 2.1, and SMB 3.0. To share filesystems over SMB, configure the filesystem as described in [“Filesystem Properties” on page 423](#). The following tables show the supported and unsupported features for SMB 3.0 and SMB 2.1.

It is strongly advised to upgrade clients from SMB 1 to at least SMB 2.0 because SMB 1 has known security and performance issues that are resolved in later SMB versions.

TABLE 87 SMB 3.0 Supported and Unsupported Features

SUPPORTED FEATURES	UNSUPPORTED FEATURES
Transparent failover (Continuously Available shares)	Encryption
Multichannel	SMB over Remote Direct Memory Access (RDMA)
	Volume Shadow Copy Service (VSS) for SMB filesystems
	Directory leasing

TABLE 88 SMB 2.1 Supported and Unsupported Features

SUPPORTED FEATURES	UNSUPPORTED FEATURES
Lease	Branch cache
Multi-protocol negotiate request	Resilient handles
Individual write-through operations	
Multi-credit operations	

Local accounts and user IDs are mapped to Windows user IDs. Note that the *guest* account is a special, read-only account and cannot be configured for read/write in the appliance.

To configure SMB, see the following sections:

- [“SMB Service Properties” on page 347](#)
- [“Setting Properties to Export Shares over SMB” on page 349](#)
- [“NFS/SMB Interoperability” on page 350](#)
- [“SMB DFS Namespaces” on page 350](#)
- [“SMB Microsoft Stand-alone DFS Namespace Management Tools Support Matrix” on page 351](#)
- [“Adding DFS Namespaces to a Local SMB Group” on page 353](#)
- [“SMB Autohome” on page 353](#)
- [“Adding SMB Autohome Rules \(CLI\)” on page 354](#)
- [“Adding a User to an SMB Local Group” on page 355](#)
- [“SMB MMC Integration” on page 356](#)
- [“SMB Share Management” on page 357](#)
- [“SMB Users, Groups, and Connections” on page 359](#)
- [“Listing SMB Services” on page 360](#)
- [“Configuring SMB \(BUI\)” on page 362](#)
- [“Configuring SMB Active Directory \(BUI\)” on page 364](#)
- [“Configuring SMB Project and Share \(BUI\)” on page 364](#)
- [“Configuring SMB Data Service \(BUI\)” on page 365](#)

SMB Service Properties

Changing service properties is documented in [“Setting Service Properties \(BUI\)” on page 252](#) and [“Setting Service Properties \(CLI\)” on page 253](#).

- **Minimum supported version** - Choose the minimum version of SMB that the appliance supports.

- **Maximum supported version** - Choose the maximum version of SMB that the appliance supports.
- **System comment** - Meaningful text string.
- **Idle Session timeout** - Timeout setting for session inactivity.
- **Preferred domain controller** - The preferred domain controller to use when joining an Active Directory domain. If this controller is not available, Active Directory will rely on DNS SRV records and the Active Directory site to locate an appropriate domain controller. For more information, see [“Active Directory Configuration” on page 259](#).
- **Active Directory site** - The site to use when joining an Active Directory domain. A site is a logical collection of machines which are all connected with high bandwidth, low latency network links. When this property is configured and the preferred domain controller is not specified, joining an Active Directory domain will prefer domain controllers located in this site over external domain controllers.
- **Maximum # of server threads** - The maximum number of simultaneous server threads (workers). Default is 1024.
- **Enable Dynamic DNS** - Choose whether the appliance will use Dynamic DNS to update DNS records in the Active Directory domain. Default is off.
- **Enable oplocks** - Choose whether the appliance will grant opportunistic locks to SMB clients. This will improve performance for most clients. Default is on. The SMB server grants an oplock to a client process so that the client can cache data while the lock is in place. When the server revokes the oplock, the client flushes its cached data to the server.
- **Restrict anonymous access to share list** - If this option is enabled, clients must authenticate to the SMB service before receiving a list of shares. If disabled, anonymous clients may access the list of shares.
- **Primary WINS server** - Primary WINS address configured in the TCP/IP setup.
- **Secondary WINS server** - Secondary WINS address configured in the TCP/IP setup.
- **Excluded IP addresses from WINS** - IP addresses excluded from registration with WINS.
- **LAN Manager compatibility level** - Authentication modes supported (LM, NTLM, LMv2, NTLMv2). For more information on the supported authentication modes within each compatibility level, consult the Oracle Solaris Information Library for *smb*. NTLMv2 is the recommended minimum security level to avoid publicly known security vulnerabilities.
- **SMB signing enabled** - Enables interoperability with SMB clients using the SMB signing feature. If a packet has been signed, the signature will be verified. If a packet has not been signed it will be accepted without signature verification (if SMB signing is not required - see below).
- **SMB signing required** - When SMB signing is required, all SMB packets must be signed or they will be rejected, and clients that do not support signing will be unable to connect to the server.
- **Ignore zero VC** - When an SMB client establishes a new connection, it may request that the appliance clean up all previous connections and file locks from this client by specifying a Virtual Circuit (VC) number of zero. This protocol artifact however, does not respect

network address translation (NAT) for clients or multiple DNS entries assigned to the same host. In combination, zero VC requests between masked or redundant network locations may result in unrelated active connections being reset. By default, zero VC requests are honored to prevent stale file locking, however if SMB sessions are being disconnected in error, ignoring zero VC requests may resolve the issue.

- **Share visibility** - Use this property to set the access-based enumeration (ABE) policy for displaying available shares to clients. Valid values are "Full" and "Restricted." While "Full" allows full access, "Restricted" limits access to only shares that the client is allowed to see. Access to shares is determined by the SMB exceptions and the share's ACL. This property is set to "Full" by default.

Setting Properties to Export Shares over SMB

Several share properties must be set in certain ways when exporting a share over SMB.

TABLE 89 SMB Share Properties

Property	Description
Case Sensitivity	SMB clients expect case-insensitive behavior, so this property must be "mixed" or "insensitive". See “Static Properties” on page 414 .
Reject non UTF-8	If non-UTF-8 filenames are allowed in a filesystem, SMB clients may function incorrectly. See “Static Properties” on page 414 .
Non-Blocking Mandatory Locking	This property must be enabled to allow byte range locking to function correctly. See “Static Properties” on page 414 .
Resource name	The name by which clients refer to the share. For information about how this name is inherited from a project, see “Share and Project Protocols” on page 442 .
Share-level ACL	An ACL which adds another layer of access control beyond the ACLs stored in the filesystem. For more information on this property, see “Access Control Lists for Filesystems” on page 456 .

The case sensitivity and reject non UTF-8 properties can only be set when creating a share.

No two SMB shares on the same system may share the same resource name. Resource names inherited from projects have special behavior. For details, see [“Shares and Projects” on page 383](#). Resource names must be less than 80 characters, and can contain any alphanumeric characters besides the following characters:

" / \ [] : | < > + ; , ? * =

When access-based enumeration is enabled, clients may see directory entries for files which they cannot open. Directory entries are filtered only when the client has no access to that file. For example, if a client attempts to open a file for read/write access but the ACL grants only read access, that open request will fail but that file will still be included in the list of entries.

NFS/SMB Interoperability

The appliance supports NFS and SMB clients accessing the same shares concurrently. To correctly configure the appliance for NFS/SMB interoperability, you must configure the following components:

- Configure the Active Directory service. See [“Active Directory Configuration” on page 259](#).
- Establish an Identity Mapping strategy and configure the service. See [“Identity Mapping Configuration” on page 280](#).
- Configure SMB. See [“SMB Configuration” on page 346](#).
- Configure access control, ACL entries, and ACL inheritance on shares.

SMB and NFSv3 do not use the same access control model. For best results, configure the ACL on the root directory from a SMB client as the SMB access control model is a more verbose model. For information on inheritable trivial ACL entries, see [“Access Control Lists for Filesystems” on page 456](#).

SMB DFS Namespaces

The Distributed File System (DFS) is a virtualization technology delivered over the SMB and MSRPC protocols. DFS allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. An administrator can select which shared folders to present in the namespace, design the hierarchy in which those folders appear and determine the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside in a single, high-capacity file system. Users can navigate the folders in the namespace without needing to know the server names or shared folders hosting the data.

Only one share per system may be provisioned as a standalone DFS namespace. Domain-based DFS namespaces are not supported. Note that one DFS namespace may be provisioned per cluster, even if each cluster node has a separate storage pool. To provision a SMB share as a DFS namespace, use the DFS Management MMC Snap-in to create a standalone namespace.

When the appliance is not joined to an Active Directory domain, additional configuration is necessary to allow Workgroup users to modify DFS namespaces. To enable an SMB local user

to create or delete a DFS namespace, that user must have a separate local account created on the server. For information about steps to let the SMB local user `dfsadmin` manipulate DFS namespaces, see [“Adding DFS Namespaces to a Local SMB Group” on page 353](#).

SMB Microsoft Stand-alone DFS Namespace Management Tools Support Matrix

The following table lists operations (subcommands/options) of the Microsoft DFS tools on various Windows operating system versions. It identifies which of these are supported by the DFS service on the appliance for managing a standalone DFS namespace on the appliance.

- **y** - supported
- **n** - not supported
- **NA** - not applicable

Microsoft Windows systems	XP	2003	2003	Vista	2008	2008	Win7
	SP3	SP2	SP2	SP2	SP2	SP1	SP1
dfscmd CLI:							
/map [comment] [/restore]	y	y	y	y	y	y	y
/unmap	y	y	y	y	y	y	y
/add [/restore]	y	y	y	y	y	y	y
/remove	y	y	y	y	y	y	y
/view [/partial /full]	y	y	y	y	y	y	y
dfsutil CLI (old format):							
/addstdroot [/comment]	y	y	y	n	n	y	y
/remstdroot	y	y	y	n	n	y	y
/root:<DfsName> /view	n	n	n	y	y	y	y
/addlink [/comment]	NA	NA	NA	y	y	y	y
/removelink	NA	NA	NA	y	y	y	y
/state /display	NA	NA	NA	y	y	y	y
/state /enable	NA	NA	NA	y	y	y	y
/state /disable	NA	NA	NA	y	y	y	y
/ttl /display	NA	NA	NA	y	y	y	y
/ttl /set	NA	NA	NA	y	y	y	y
/server:<MachineName> /view	y	y	y	y	y	y	y
dfsutil CLI (new format):							

root addstd [comment]	NA	NA	NA	n	n	y	y
root remove	NA	NA	NA	n	n	y	y
root (view namespace)	NA	NA	NA	y	y	y	y
link add [comment]	NA	NA	NA	y	y	y	y
link remove	NA	NA	NA	y	y	y	y
link (view)	NA	NA	NA	y	y	y	y
target add	NA	NA	NA	y	y	y	y
target remove	NA	NA	NA	y	y	y	y
target (view)	NA	NA	NA	y	y	y	y
property comment (view)	NA	NA	NA	y	y	y	y
property comment set	NA	NA	NA	y	y	y	y
property ttl (view)	NA	NA	NA	y	y	y	y
property ttl set	NA	NA	NA	y	y	y	y
property state (view)	NA	NA	NA	y	y	y	y
property state offline	NA	NA	NA	y	y	y	y
property state online	NA	NA	NA	y	y	y	y
DFS GUI:							
add standalone root	y	y	y	n	n	n	n
remove standalone root	y	y	y	n	n	n	n
change root comment	y	y	y	n	n	n	n
change root timeout	y	y	y	n	n	n	n
add link	y	y	y	n	n	n	n
remove link	y	y	y	n	n	n	n
change link comment	y	y	y	n	n	n	n
change link timeout	y	y	y	n	n	n	n
add link's target	y	y	y	n	n	n	n
remove link's target	y	y	y	n	n	n	n
enable link's referral (target)	y	y	y	n	n	n	n
disable link's referral (target)	y	y	y	n	n	n	n
hide root	y	y	y	y	y	y	y
show root	y	y	y	y	y	y	y
display links	y	y	y	n	n	n	n
display targets	y	y	y	n	n	n	n
	XP 2003 2003 Vista 2008 2008 Win7						
			R2			R2	
	SP3	SP2	SP2	SP2	SP2	SP1	SP1

Note that:

- Oracle Solaris does not verify the DFS link target.
- CLI commands for modifying and viewing comment and timeout (TTL) are applicable to both root and link.
- CLI commands for viewing state are applicable to root, root's target, link, and link's target.
- CLI commands for modifying state are only applicable for link and link's target.

▼ Adding DFS Namespaces to a Local SMB Group

1. Create a local user account on the server for user `dfsadmin`. Be sure to use the same password as when the local user was first created on the Windows machine.
2. Add `dfsadmin` to the local SMB group Administrators.
3. Log in as `dfsadmin` on the Windows machine from which the DFS namespace will be modified.

SMB Autohome

For Windows file sharing, Autohome provides access to filesystems using the SMB protocol. Autohome defines and maintains home directory shares for users that access the system through SMB. Autohome rules map SMB clients to home directories.

FIGURE 21 Setting Autohome Rules



- **Use Name Service Switch** - Toggles Name Service Switch (NSS) on or off. You cannot create an NSS rule and an rule for all users at the same time.

- **AD Container** - Sets the Active Directory container, for example: dc=com,dc=fishworks,ou=Engineering,CN=myhome.
- **User** - Sets the Autohome rule for all All users or for the user you specify. When you specify a user, the wildcards "&" and "?" refer to a user's login and its corresponding first character.
- **Directory** - Sets the directory for the rule, for example: /export/wdp.

▼ Adding SMB Autohome Rules (CLI)

1. **Go to configuration services smb.**
2. **Use the create command to add autohome rules, and the list command to list existing rules.**

This example adds a rule for the user "Bill" then lists the rules:

```
hostname:> configuration services smb
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=false
hostname:configuration services rule (uncommitted)> set user=Bill
hostname:configuration services rule (uncommitted)> set directory=/export/wdp
hostname:configuration services rule (uncommitted)> set container="dc=com,dc=fishworks,
ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE      NSS      USER      DIRECTORY      CONTAINER
rule-000  false    Bill      /export/wdp    dc=com,dc=fishworks,
ou=Engineering,CN=myhome
```

3. **Create Autohome rules using wildcard characters.**

The & character matches the users' username, and the ? character matches the first letter of the users' username. The following uses wildcards to match all users:

```
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=false
hostname:configuration services rule (uncommitted)> set user=*
hostname:configuration services rule (uncommitted)> set directory=/export/?/&
hostname:configuration services rule (uncommitted)> set container="dc=com,dc=fishworks,
ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE      NSS      USER      DIRECTORY      CONTAINER
rule-000  false    Bill      /export/wdp    dc=com,dc=fishworks,
ou=Engineering,CN=myhome
```

4. The name service switch can also be used to create autohome rules:

```
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=true
hostname:configuration services rule (uncommitted)> set container="dc=com,dc=fishworks,
ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE      NSS      USER      DIRECTORY      CONTAINER
rule-000  true                      dc=com,dc=fishworks,
ou=Engineering,CN=myhome
```

▼ Adding a User to an SMB Local Group

Local groups are groups of domain and/or local users that grant additional privileges to those users.

SMB Local Groups:

- **Administrators** - Administrators can bypass file permissions to change the ownership on files.
- **Backup Operators** - Backup Operators can bypass file access controls to backup and restore files.

1. Go to configuration services smb groups.

```
hostname:configuration services smb> groups
```

2. Enter create.

```
hostname:configuration services smb groups> create
```

3. Specify the user you want to add to the group:

```
hostname:configuration services smb member (uncommitted)> set user=Bill
```

4. Enter the group name, and then commit the change:

```
hostname:configuration services smb member (uncommitted)> set group="Backup Operators"
hostname:configuration services smb member (uncommitted)> commit
```

5. Enter list to confirm the user was added to the specified group:

```
hostname:configuration services smb groups> list
MEMBER      USER      GROUP
```

member-000 WINDOMAIN\Bill

Backup Operators

SMB MMC Integration


The Microsoft Management Console (MMC) is an extensible framework of registered components, known as snap-ins, that provide comprehensive management features for both the local system and remote systems on the network. Computer Management is a collection of Microsoft Management Console tools, that may be used to configure, monitor and manage local and remote services and resources.

In order to use the MMC functionality on the appliance in workgroup mode, be sure to add the Windows administrator who will use the management console to the Administrators local group on the appliance. Otherwise you may receive an `Access is denied` or similar error on the administration client when attempting to connect to the appliance using the MMC.

The appliance supports the following Computer Management facilities:

The Event Viewer MMC snap-in displays the Application log, Security log, and System log. These logs show the contents of the alert, audit, and system logs of the appliance.

The following screen shows an example of the Application log and the properties dialog for an error event.

FIGURE 22 SMB Event Viewer

SMB Share Management

Support for share management includes the following:


- Listing shares
- Setting ACLs on shares
- Changing share permissions
- Setting the description of a share

Features not currently supported via MMC include the following:

- Adding or Deleting a share
- Setting client side caching property
- Setting maximum allowed or number of users property

The following screen shows an example of permission properties for a share.

FIGURE 23 SMB Share Permission Properties




SMB Users, Groups, and Connections

The following features are supported:

- Viewing local SMB users and groups
- Listing user connections, including listing the number of open files per connection
- Closing user connections
- Listing open files, including listing the number of locks on the file and file open mode
- Closing open files


The following screen shows an example of open files per connection.

FIGURE 24 Open Files per Connection



The following screen shows an example of open sessions.


FIGURE 25 Open Sessions



Listing SMB Services

Listing of appliance services is supported using the MMC application. However, you cannot enable or disable services using the MMC application. Support includes listing of appliance services. Services cannot be enabled or disabled using the Computer Management MMC application.

The following screen shows an example of General properties for the vscan Service.

FIGURE 26 vscan Properties

To ensure that only the appropriate users have access to administrative operations, there are some access restrictions on the operations performed remotely using MMC.

TABLE 90 Users and Allowed Operations

User	Allowed Operations
Regular users	List shares.
Members of the Administrators or Power Users groups	Manage shares, list user connections.
Members of the Administrators group	List open files and close files, disconnect user connections, view services and event log.

▼ Configuring SMB (BUI)

Initial configuration of the appliance may be completed using the BUI or the CLI and should take less than 20 minutes. Initial Setup may also be performed again later using the Maintenance > System contexts of the BUI or CLI. Initial configuration takes you through the following steps.

1. **Configure Network Devices, Datalinks, and Interfaces.**
 - a. Create interfaces using the Datalink add or Interface  icons or by using drag-and-drop of devices to the datalink or interface lists.
 - b. Set the desired properties and click the Apply button to add them to the list.
 - c. Set each interface to active or standby as appropriate.
 - d. Click APPLY at the top of the page to commit your changes.
2. **Configure DNS.**
 - a. Provide the base domain name.
 - b. Provide the IP address of at least one server that is able to resolve hostname and server records in the Active Directory portion of the domain namespace.
3. **Configure NTP authentication keys to ensure clock synchronization.**
 - a. Click the  icon to add a new key.
 - b. Specify the number, type, and private value for the new key and apply the changes.

The default team for each account is the same as the account user name, prefixed with a '\$'.

- d. **Commit your initial configuration changes.**

▼ **Configuring SMB Active Directory (BUI)**

1. **Create an account for the appliance in the Active Directory domain.**
For detailed instructions, refer to [“Active Directory Configuration” on page 259](#).
2. **On the Configuration > Services > Active Directory screen, click the Join Domain button.**
3. **Specify the Active Directory domain, administrative user, administrative password.**
4. **Click APPLY to commit the changes.**

▼ **Configuring SMB Project and Share (BUI)**

1. **Go to Shares > Shares.**
2. **Create a project.**
 - a. **On the Shares screen, click the panel open icon  to expand the Projects panel.**
 - b. **Click the add icon  to add a new project.**
 - c. **Specify the project name and click APPLY.**
3. **Select the new project from the Projects panel.**
4. **Click the add item icon  to add a filesystem.**
5. **Click the edit icon  for the filesystem.**
6. **Click the General tab and deselect the Inherit from project checkbox.**
7. **Choose a mountpoint under /export, even though SMB shares are accessed by resource name, and click APPLY.**

8. Click the Protocols tab for the project and set the SMB resource name to on.
9. Enable sharesmb and share-level ACL for the Project.
10. Click APPLY to activate the configuration.

▼ Configuring SMB Data Service (BUI)

1. Go to Configuration > Services > SMB and click the power icon  to enable the service.
2. Set SMB properties and click APPLY to activate the configuration. See [“SMB Service Properties” on page 347](#).
3. Click the Autohome tab on the Configuration > Services > SMB screen to set autohome rules to map SMB clients to home directories as described in [“SMB Autohome” on page 353](#). Click APPLY to activate the configuration.
4. Click the Local Groups tab on the Configuration > Services > SMB screen and use the add item icon  to add administrators or backup operator users to local groups as described in [“Adding a User to an SMB Local Group” on page 355](#). Click APPLY to activate the configuration.

SMTP Configuration

The SMTP service sends all mail generated by the appliance, typically in response to alerts as configured on the Alerts screen. The SMTP service does not accept external mail; it only sends mail generated automatically by the appliance itself.

By default, the SMTP service uses DNS (MX records) to determine where to send mail. If DNS is not configured for the appliance's domain, or the destination domain for outgoing mail does not have DNS MX records setup properly, the appliance can be configured to forward all mail through an outgoing mail server, commonly called a smarthost.

TABLE 91 SMTP Properties

Property	Description
Send mail through smarthost	If enabled, all mail is sent through the specified outgoing mail server. Otherwise, DNS is used to determine where to send mail for a particular domain.

Property	Description
Smarthost hostname	Outgoing mail server hostname.
Allow customized from address	If enabled, the From address for email is set to the Custom from address property. It may be desirable to customize this if the default From address is being identified as spam, for example.
Custom from address	The From address to use for outbound email.

When changing properties, you can use Alerts to send a test email to verify that the properties are correct. A common reason for undelivered email is misconfigured DNS, which prevents the appliance from determining which mail server to deliver the mail to; as described earlier, a smarthost could be used if DNS cannot be configured.

TABLE 92 SMTP Logs

Log	Description
network-smtp:sendmail	Logs the SMTP service events
mail	Log of SMTP activity (including mails sent)

SNMP Configuration

The Simple Network Management Protocol (SNMP) service provides two different functions on the appliance:

- Appliance status information can be served by SNMP.
- Alerts can be configured to send SNMP traps. See [“Configuring Alerts” on page 228](#).

SNMP versions v1, v2c, and v3 are available when this service is enabled. The appliance supports a maximum of 128 physical and logical network interfaces. More than 128 network interfaces could cause time outs for such commands as `snmpwalk` and `snmpget`. If you need more than 128 network interfaces, contact Oracle Support.

To configure SNMP, see the following sections:

- [“Configuring SNMP to Serve Appliance Status \(BUI\)” on page 367](#)
- [“Configuring SNMP to Send Traps \(BUI\)” on page 367](#)
- [“SNMP Properties” on page 367](#)
- [“SNMP MIBs” on page 368](#)
- [“Sun FM MIB” on page 368](#)
- [“Sun AK MIB” on page 369](#)

▼ Configuring SNMP to Serve Appliance Status (BUI)

1. Go to Configuration > Services > SNMP.
2. Set the community name, authorized network and contact string.
3. (Optional) Set the trap destination to a remote SNMP host, else set this to 127.0.0.1.
4. Click APPLY to commit the configuration.

▼ Configuring SNMP to Send Traps (BUI)

1. Go to Configuration > Services > SNMP.
2. Set the community name, contact string, and trap destination(s).
3. (Optional) Set the authorized network to allow SNMP clients, else set this to 127.0.0.1/8.
4. Click APPLY to commit the configuration.
5. You must configure alerts to send the traps you want to receive.
For more information about alerts, see [“Configuring Alerts” on page 228](#).

Related Topics

- [“SNMP Properties” on page 367](#)

SNMP Properties

TABLE 93 SNMP Properties

Property	Description
Version	Toggles between v1/2c and v3.
Community name	Toggles between public and user-input. If you select user-input, you must also enter a community name. If you select v3, this property is not available.
Authorized network/subnet	Enter an appropriate IPv4 address and subnet (integers from 0-32). If you select v3, this property is not available.
Appliance contact	Enter an appropriate appliance contact.

Property	Description
Username/password	Enter a valid username (max 501 characters) and password (8-501 characters). If you select v1/2c, this property is not available.
Authentication	Toggles between MD5 and SHA authentication algorithms. If you select v1/2c, this property is not available.
Privacy	Toggles between None and DES encryption algorithm. If you select v1/2c, this property is not available.
Engine ID	The EngineID value hashed by snmpd. If SNMP was not previously enabled, the label shows "0x000".
Trap destinations	Lets you add IPv4 addresses. Use the "+" and "-" buttons to add or remove addresses.

The SNMP service also provides the MIB-II location string. This property is sourced from the [System Identity](#) configuration.

SNMP MIBs

If the SNMP services is online, authorized networks will have access to the following MIBs (Management Information Bases):

TABLE 94 SNMP MIBs

MIB	Purpose
.1.3.6.1.2.1.1	MIB-II system - generic system information, including hostname, contact and location
.1.3.6.1.2.1.2	MIB-II interfaces - network interface statistics
.1.3.6.1.2.1.4	MIB-II IP - Internet Protocol information, including IP addresses and route table
.1.3.6.1.4.1.42	Sun Enterprise MIB (SUN-MIB.mib.txt)
.1.3.6.1.4.1.42.2.195	Sun FM - fault management statistics (MIB file linked below)
.1.3.6.1.4.1.42.2.225	Sun AK - appliance information and statistics (MIB file linked below)

Note - Sun MIB files are available at <https://your IP address or host name:215/help/docs/snmp/>.

Sun FM MIB

The Sun FM MIB (SUN-FM-MIB.mib) provides access to SUN Fault Manager information such as:

- Active problems on the system
- Fault Manager events
- Fault Manager configuration information

There are four main tables to read:

TABLE 95 Sun FM MIBs

OID	Contents
.1.3.6.1.4.1.42.2.195.1.1	Fault Management problems
.1.3.6.1.4.1.42.2.195.1.2	Fault Management fault events
.1.3.6.1.4.1.42.2.195.1.3	Fault Management module configuration
.1.3.6.1.4.1.42.2.195.1.5	Fault Management faulty resources

See the MIB file for the full descriptions.

Note - Sun MIB files are available at <https://your IP address or host name:215/help/docs/snmp/>.

Sun AK MIB

The Sun AK MIB (SUN-AK-MIB.mib) provides the following information:

- Product description string and part number
- Appliance software version
- Appliance and chassis serial numbers
- Install, update and boot times
- Cluster state
- Share status (share name, size, used and available bytes)

There are three main tables to read:

TABLE 96 Sun AK MIBs

OID	Contents
.1.3.6.1.4.1.42.2.225.1.4	General appliance info
.1.3.6.1.4.1.42.2.225.1.5	Cluster status
.1.3.6.1.4.1.42.2.225.1.6	Share status

See the MIB file for the full descriptions.

Note - Sun MIB files are available at <https://your IP address or host name:215/help/docs/snmp/>.

SRP Configuration

When you configure a LUN on the appliance you can export that volume over a SCSI Remote Protocol (SRP) target. The SRP service allows initiators to access targets using the SRP protocol.

For information on SRP targets and initiators, see [“Configuring Storage Area Network \(SAN\)” on page 170](#).

SSH Configuration

The SSH (Secure Shell) service allows users to log in to the appliance CLI and perform most of the same administrative actions that can be performed in the BUI. The SSH service can also be used as a way to execute automated scripts from a remote host, such as for retrieving daily logs or Analytics statistics.

SSH keys can be configured for individual accounts using the preferences function described in [“Setting Appliance Preferences” on page 224](#).

SSH can also be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see [“Kerberos Configuration” on page 293](#).

For added security when configuring SSH, you can specify the ciphers and MACs, as described in [“SSH Properties and Logs” on page 371](#).

To configure SSH, see the following sections:

- [“Disabling root SSH Access \(CLI\)” on page 370](#)
- [“SSH Properties and Logs” on page 371](#)

▼ Disabling root SSH Access (CLI)

1. **Go to configuration services ssh.**
2. **Set permit root login to false.**
3. **Commit the configuration.**

SSH Properties and Logs

TABLE 97 SSH Properties

Property	Description	Examples
Server key length	The number of bits in the ephemeral key.	768
Key regeneration interval	Ephemeral key regeneration interval, in seconds.	3600
Login grace period	The SSH connection will be disconnected after this many seconds if the client has failed to authenticate.	120
Permit root login	Allows the root user to log in using SSH.	yes
Port (for incoming connections)	The designated port for incoming connections.	22

TABLE 98 SSH Security Properties

Property	Description
Ciphers	Ciphers for SSH connections.
MACs	Message authentication codes (MACs) for SSH connections.

TABLE 99 SSH Logs

Log	Description
network-ssh:default	Log of the SSH service events and errors

Syslog Configuration

The Syslog Relay service provides two different functions on the appliance:

- Alerts can be configured to send Syslog messages to one or more remote systems. See [“Configuring Alerts” on page 228](#).
- Services on the appliance that are syslog capable will have their syslog messages forwarded to remote systems.

A *syslog message* is a small event message transmitted from the appliance to one or more remote systems (or as we like to call it: intercontinental printf). The message contains the following elements:

- A facility describing the type of system component that emitted the message.
- A severity describing the severity of the condition associated with the message.
- A timestamp describing the time of the associated event in UTC.
- A hostname describing the canonical name of the appliance
- A tag describing the name of the system component that emitted the message. See [“SYSLOG Alert Message Format” on page 374](#) for details of the message format.
- A message describing the event itself. See [“SYSLOG Alert Message Format” on page 374](#) for details of the message format.

Syslog receivers are provided with most operating systems, including Oracle Solaris and Linux. A number of third-party and open-source management software packages also support Syslog. Syslog receivers allow administrators to aggregate messages from a number of systems on to a single management system and incorporated into a single set of log files.

The Syslog Relay can be configured to use the "classic" output format described by RFC 3164, or the newer, versioned output format described by RFC 5424. Syslog messages are transmitted as UDP datagrams. Therefore they are subject to being dropped by the network, or may not be sent at all if the sending system is low on memory or the network is sufficiently congested. Administrators should therefore assume that in complex failure scenarios in a network some messages may be missing and were dropped.

Syslog Properties

- **Protocol Version** - The version of the Syslog protocol to use, either Classic Syslog (RFC 3164) or Updated Syslog (RFC 5424).
- **Destinations** - The list of destination IPv4, IPv6, and FQDN addresses to which messages are relayed.

To configure syslog, see the following sections:

- [“Classic Syslog: RFC 3164” on page 372](#)
- [“Updated Syslog: RFC 5424” on page 373](#)
- [“SYSLOG Message Format” on page 373](#)
- [“SYSLOG Alert Message Format” on page 374](#)
- [“Example Configuring an Oracle Solaris Receiver \(CLI\)” on page 375](#)
- [“Example Configuring a Linux Receiver \(CLI\)” on page 376](#)

Classic Syslog: RFC 3164

The Classic Syslog protocol includes the facility and level values encoded as a single integer priority, the timestamp, a hostname, a tag, and the message body.

The tag will be one of the tags described in [“SYSLOG Message Format” on page 373](#).

The hostname will be the canonical name of the appliance as defined by the System Identity configuration. For more information, see [“System Identity Configuration” on page 377](#).

Updated Syslog: RFC 5424

The Classic Syslog protocol includes the facility and level values encoded as a single integer priority, a version field (1), the timestamp, a hostname, a app-name, and the message body. Syslog messages relayed by the Sun Storage systems will set the RFC 5424 procid, msgid, and structured-data fields to the nil value (-) to indicate that these fields do not contain any data.

The app-name will be one of the tags described in [“SYSLOG Message Format” on page 373](#).

The hostname will be the canonical name of the appliance as defined by the System Identity configuration. For more information, see [“System Identity Configuration” on page 377](#).

SYSLOG Message Format

The Syslog protocol itself does not define the format of the message payload, leaving it up to the sender to include any kind of structured data or unstructured human-readable string that is appropriate. Sun Storage appliances use the syslog subsystem tag `ak` to indicate a structured, parseable message payload, described next. Other subsystem tags indicate arbitrary human-readable text, but administrators should consider these string forms *unstable* and subject to change without notice or removal in future releases of the Sun Storage software.

TABLE 100 SYSLOG Message Formats

Facility	Tag Name	Description
daemon	ak	Generic tag for appliance subsystems. All alerts will be tagged <code>ak</code> , indicating a SUNW-MSG-ID follows.
daemon	idmap	Identity Mapping service for POSIX and Windows identity conversion. See “Identity Mapping Configuration” on page 280 .
daemon	smbd	SMA Data Protocol for accessing shares. See “SMB Configuration” on page 346 .

SYSLOG Alert Message Format

If an alert is configured with the Send Syslog Message action, it will produce a syslog message payload containing localized text consisting of the following standard fields. Each field will be prefixed with the field name in CAPITAL letters followed by a colon and whitespace character.

TABLE 101 SYSLOG Alert Message Formats

Field Name	Description
SUNW-MSG-ID	The stable Sun Fault Message Identifier associated with the alert. Each system condition and fault diagnosis that produces an administrator alert is assigned a persistent, unique identifier in Sun's Fault Message catalog. These identifiers can be easily read over the phone or scribbled down in your notebook, and link to a corresponding knowledge article found at My Oracle Support (https://support.oracle.com/) "Predictive Self-Healing" (Doc ID 1154428.1).
TYPE	The type of condition. This will be one of the labels: Fault, indicating a hardware component or connector failure; Defect indicating a software defect or misconfiguration; Alert, indicating a condition not associated with a fault or defect, such as the completion of a backup activity or remote replication.
VER	The version of this encoding format itself. This description corresponds to version "1" of the SUNW-MSG-ID format. If a "1" is present in the VER field, parsing code may assume that all of the subsequent fields will be present. Parsing code should be written to handle or ignore additional fields if a decimal integer greater than one is specified.
SEVERITY	The severity of the condition associated with the problem that triggered the alert. The list of severities is shown below.
EVENT-TIME	The time corresponding to this event. The time will be in the form "Day Mon DD HH:MM:SS YYYY" in UTC. For example: Fri Aug 14 21:34:22 2009.
PLATFORM	The platform identifier for the appliance. This field is for Oracle Service use only.
CSN	The chassis serial number of the appliance.
HOSTNAME	The canonical name of the appliance as defined by the System Identity configuration. See System Identity .
SOURCE	The subsystem within the appliance software that emitted the event. This field is for Oracle Service use only.
REV	The internal revision of the subsystem. This field is for Oracle Service use only.

Field Name	Description
EVENT-ID	The Universally Unique Identifier (UUID) associated with this event. Oracle's Fault Management system associates a UUID with each alert and fault diagnosis such that administrators can gather and correlated multiple messages associated with a single condition, and detect duplicate messages. Oracle Service personnel can use the EVENT-ID to retrieve additional postmortem information associated with the problem that may help Oracle respond to the issue.
DESC	Description of the condition associated with the event.
AUTO-RESPONSE	The automated response to the problem, if any, by the Fault Management software included in the system. Automated responses include capabilities such as proactively offlining faulty disks, DRAM memory chips, and processor cores.
REC-ACTION	The recommended service action. This will include a brief summary of the recommended action, but administrators should consult the knowledge article and this documentation for information on the complete repair procedure.

The SEVERITY field will be set to one of the following values:

TABLE 102 SYSLOG Severity Fields

Severity	Syslog Level	Description
Minor	LOG_WARNING	A condition occurred that does not currently impair service, but the condition needs to be corrected before it becomes more severe.
Major	LOG_ERR	A condition occurred that does impair service but not seriously.
Critical	LOG_CRIT	A condition occurred that seriously impairs service and requires immediate correction.

Example Configuring an Oracle Solaris Receiver (CLI)

Most operating systems include a syslog receiver, but some configuration steps may be required to turn it on. Consult the documentation for your operating system or management software for specific details of syslog receiver configuration.

Oracle Solaris includes a bundled syslogd that can act as a syslog receiver, but the remote receive capability is disabled by default. To enable Oracle Solaris to receive syslog traffic, use svccfg and svcadm to modify the syslog settings as follows:

```
# svccfg -s system/system-log setprop config/log_from_remote = true
# svcadm restart system/system-log
```

Oracle Solaris syslogd only understands the classic Syslog protocol. Refer to the Oracle Solaris syslog.conf(4) man page for information on how to configure filtering and logging of the received messages.

By default, Oracle Solaris syslogd records messages to /var/adm/messages and a test alert would be recorded as follows:

```
Aug 14 21:34:22 poptart.example.us.com poptart ak: SUNW-MSG-ID: AK-8000-LM, \
TYPE: alert, VER: 1, SEVERITY: Minor\nEVENT-TIME: Fri Aug 14 21:34:22 2009\n\
PLATFORM: i86pc, CSN: 12345678, HOSTNAME: poptart\n\
SOURCE: jsui.359, REV: 1.0\n\
EVENT-ID: 92dfcb39-6e15-e2d5-a7d9-dc3e221becea\n\
DESC: A test alert has been posted.\n\
AUTO-RESPONSE: None.\nIMPACT: None.\nREC-ACTION: None.
```

Example Configuring a Linux Receiver (CLI)

Most operating systems include a syslog receiver, but some configuration steps may be required to turn it on. Consult the documentation for your operating system or management software for specific details of syslog receiver configuration.

Most Linux distributions include a bundled syslogd(8) daemon that can act as a syslog receiver, but the remote receive capability is disabled by default. To enable Linux to receive syslog traffic, edit the /etc/sysconfig/syslog configuration file such that the -r option is included (enables remote logging):

```
SYSLOGD_OPTIONS="-r -m 0"
```

and then restart the logging service:

```
# /etc/init.d/syslog stop
# /etc/init.d/syslog start
```

Some Linux distributions have an ipfilter packet filter that will reject syslog UDP packets by default, and the filter must be modified to permit them. On these distributions, use a command similar to the following to add an INPUT rule to accept syslog UDP packets:

```
# iptables -I INPUT 1 -p udp --sport 514 --dport 514 -j ACCEPT
```

By default, Linux syslogd records messages to /var/log/messages and a test alert would be recorded as follows:

```
Aug 12 22:03:15 192.168.1.105 poptart ak: SUNW-MSG-ID: AK-8000-LM, \
```

```

TYPE: alert, VER: 1, SEVERITY: Minor EVENT-TIME: Wed Aug 12 22:03:14 2009 \
PLATFORM: i86pc, CSN: 12345678, HOSTNAME: poptart SOURCE: jsui.3775, REV: 1.0 \
EVENT-ID: 9d40db07-8078-4b21-e64e-86e5cac90912 \
DESC: A test alert has been posted. AUTO-RESPONSE: None. IMPACT: None. \
REC-ACTION: None.

```

System Identity Configuration

This service provides configuration for the system name and location. You might need to change these if the appliance is moved to a different network location, or repurposed. You can change this data in the BUI by going to Configuration > Service > System Identity. To access the same data in the CLI, go to the configuration `service identity` context.

System Identity Properties and Logs

The System Identity properties are described in the following table.

TABLE 103 System Identity Properties

BUI Text	CLI Property	Description
System Name	nodename	A single canonical identifying name for the appliance that is shown in the user interface. This name is separate from any DNS names that are used to connect to the system (which would be configured on remote DNS servers). This name can be changed at any time.
System Location	syslocation	A text string to describe where the appliance is physically located. If SNMP is enabled, this will be exported as the <i>syslocation</i> string in MIB-II.

Changing services properties is documented in [“Setting Service Properties \(BUI\)” on page 252](#) and [“Setting Service Properties \(CLI\)” on page 253](#). The CLI property names are shorter versions of those listed above.

TABLE 104 System Identity Logs

Log	Description
system-identity:node	Logs the System Identity service events and errors

To view service logs, refer to [“Using Logs” in Oracle ZFS Storage Appliance Customer Service Manual](#).

TFTP Configuration

Trivial File Transfer Protocol (TFTP) is a simple protocol to transfer files. TFTP is designed to be small and easy to implement, therefore, lacks most of the features of a regular FTP. TFTP only reads and writes files (or mail) from/to a remote server. It cannot list directories, and currently has no provisions for user authentication..

TABLE 105 TFTP Properties

Property	Description
Default Root Directory	The TFTP login location. The default is /export and points to the top of the shares hierarchy. All users will be logged into this location after successfully authenticating with the TFTP service

To use TFTP with a share, see [“Adding TFTP Access to a Share \(BUI\)” on page 378](#).

▼ Adding TFTP Access to a Share (BUI)

1. **Go to Configuration > Services.**
2. **Check that the TFTP service is enabled and online. If not, enable the service.**
3. **Go to Shares > Shares and select or add a share.**
4. **Go to the Protocols tab, and check that TFTP access is enabled.**
5. **(Optional) Set the Share mode access to Read only or Read/Write.**

Virus Scan Configuration

The Virus Scan service will scan for viruses at the filesystem level. When a file is accessed from any protocol, the Virus Scan service will first scan the file, and both deny access and quarantine the file if a virus is found. Once a file has been scanned with the latest virus definitions, it is not rescanned until it is next modified. Files accessed by NFS clients that have cached file data, or been delegated read privileges by the NFSv4.0 or NFSv4.1 servers, may not be immediately quarantined.

To configure Virus Scan, see the following sections:

- [“Configuring Virus Scanning for a Share \(BUI\)” on page 379](#)
- [“Virus Scan Properties and Logs” on page 379](#)
- [“Virus Scan File Extensions” on page 380](#)
- [“Scanning Engines” on page 381](#)

▼ Configuring Virus Scanning for a Share (BUI)

1. Go to Configuration > Services > Virus Scan.
2. Enable the service.
3. Set the appropriate properties.
4. Click APPLY to commit the configuration.
5. Go to Shares.
6. Edit a filesystem or a project.
7. Select the General tab.
8. Enable the Virus scan option.

Related Topics

- [“Virus Scan Properties and Logs” on page 379](#)
- [“Virus Scan File Extensions” on page 380](#)
- [“Scanning Engines” on page 381](#)

Virus Scan Properties and Logs

TABLE 106 Virus Scan Properties

Property	Description
Maximum file size to scan	Files larger than this size will not be scanned, to avoid significant performance penalties. These large files are unlikely to be executable themselves (such as database files), and so are less likely to pose a risk to vulnerable clients. The default value is 1GB.

Property	Description
Allow access to files that exceed maximum file size	Enabled by default, this allows access to files larger than the maximum scan size (which are therefore unscanned prior to being returned to clients). Administrators at a site with more stringent security requirements may elect to disable this option and increase the maximum file size, so that all accessible files are known to be scanned for viruses.

TABLE 107 Virus Scan Logs

Log	Description
vscan	Log of the Virus Scan service.

Related Topics

- [“Configuring Virus Scanning for a Share \(BUI\)” on page 379](#)
- [“Virus Scan File Extensions” on page 380](#)
- [“Scanning Engines” on page 381](#)

Virus Scan File Extensions

This section describes how to control which files are scanned. The default value, " * ", causes all files to be scanned. Scanning all files may impact performance so you can designate a subset of files to scan.

For example, to scan only high-risk files, including zip files, but not files with names that match the pattern "data-archive*.zip", you could configure the following settings:

TABLE 108 Virus Scan File Extensions

Action	Pattern
Scan	exe
Scan	com
Scan	bat
Scan	doc
Scan	zip
Don't Scan	data-archive*.zip
Don't Scan	*

Note - You must use "Don't Scan *" to exclude all other file types not explicitly included in the scan list. A file named `file.name.exe.bat.jpg123` would *not* be scanned, as only the "jpg123" portion of the name, the extension, would be compared against the rules.

Do *not* use exclude settings before include settings. For example, do not use a "Don't Scan *" setting before include settings since that would exclude all file types that come after it.

TABLE 109 Virus Scan Actions

Action	Pattern
Don't Scan	*
Scan	exe
Scan	com
Scan	bat
Scan	doc
Scan	zip
Don't Scan	data-archive*.zip

Related Topics

- [“Configuring Virus Scanning for a Share \(BUI\)” on page 379](#)
- [“Virus Scan Properties and Logs” on page 379](#)
- [“Scanning Engines” on page 381](#)

Scanning Engines

In this section, specify which scanning engines to use. A scanning engine is an external third-party virus scanning server which the appliance contacts using ICAP (Internet Content Adaptation Protocol, RFC 3507) to have files scanned.

TABLE 110 Scanning Engines Properties

Property	Description
Enable	Use this scan engine.
Host	Hostname or IP address of the scan engine server.
Maximum Connections	Maximum number of concurrent connections. Some scan engines operate better with connections limited to 8.
Port	Port for the scan engine.

Related Topics

- [“Configuring Virus Scanning for a Share \(BUI\)” on page 379](#)
- [“Virus Scan Properties and Logs” on page 379](#)
- [“Virus Scan File Extensions” on page 380](#)

Shares and Projects

The Oracle ZFS Storage Appliance product uses storage pools, projects, and shares to organize data. Shares are filesystems and LUNs that are exported over supported data protocols to clients of the appliance. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. For more information about how the appliance organizes data, see [“About Storage Pools, Projects, and Shares” on page 402](#).

To create and modify projects, use these tasks:

- Creating a Project - [BUI](#), [CLI](#)
- Editing a Project - [BUI](#), [CLI](#)
- Renaming a Project - [BUI](#), [CLI](#)
- Deleting a Project - [BUI](#), [CLI](#)

To create and modify filesystems and LUNs, use these tasks:

- Creating a Filesystem or LUN in a Project - [BUI](#), [CLI](#)
- Editing a Filesystem or LUN - [BUI](#), [CLI](#)
- Renaming a Filesystem or LUN - [BUI](#), [CLI](#)
- Moving a Filesystem or LUN to a Different Project - [BUI](#), [CLI](#)
- Deleting a Filesystem or LUN - [BUI](#), [CLI](#)
- Setting User or Group Quotas - [BUI](#), [CLI](#)

To understand more about how the appliance organizes storage, see these topics:

- [“About Storage Pools, Projects, and Shares” on page 402](#)
- [“Space Management for Shares” on page 436](#)
- [“Project and Share Properties” on page 403](#)
- [“Working with Filesystem Namespace” on page 439](#)
- [“Share Usage Statistics” on page 441](#)
- [“Share and Project Protocols” on page 442](#)
- [“Access Control Lists for Filesystems” on page 456](#)
- [“Working with Schemas” on page 464](#)

- [“Snapshots and Clones” on page 481](#)
- [“Remote Replication” on page 511](#)

▼ Creating a Project (BUI)

Use this task to create an unencrypted project. To create an encrypted project, see [“Creating an Encrypted Project \(BUI\)” on page 636](#).

1. **Go to Shares > Projects.**
2. **Click the add icon  next to Projects or in the expanded Projects panel. To expand the Projects panel, click the arrow icon.**
3. **In the Create Project window, enter a name for the new project.**
A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters _ - . :
4. **Click APPLY.**
The new project is added to the Projects list.

Related Topics

- [“Project Properties” on page 417](#)
- [“Creating an Encrypted Project \(BUI\)” on page 636](#)

▼ Creating a Project (CLI)

Use this task to create an unencrypted project. To create an encrypted project, see [“Creating an Encrypted Project \(CLI\)” on page 637](#).

1. **Go to shares.**

`hostname:> shares`
2. **Enter project and a project name.**
A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters _ - . :

```
hostname:shares> project home
```

3. To list the project properties, use the get command.

```
hostname:shares home(uncommitted)> get
      mountpoint = /export (default)
      quota = 0 (default)
      reservation = 0 (default)
      sharesmb = off (default)
      sharenfs = on (default)
      encryption = off (default)
      sharedav = off (default)
      shareftp = off (default)
      sharesftp = off (default)
      sharetftp = off (default)
      default_group = other (default)
      default_permissions = 700 (default)
      default_sparse = true (default)
      default_user = nobody (default)
      default_volblocksize = 8K (default)
      default_volsize = 0 (default)
      aclinherit = (default)
      aclmode = (default)
      atime = (default)
      checksum = (default)
      compression = (default)
      dedup = (default)
      copies = (default)
      logbias = (default)
      readonly = (default)
      recordsize = (default)
      rstchown = (default)
      secondarycache = (default)
      nbmand = (default)
      snapdir = (default)
      vscan = (default)
      defaultuserquota = (default)
      defaultgroupquota = (default)
      snaplabel = (default)
      canonical_name = (default)
      keyname = (default)
      keystore = (default)
      exported = (default)
      nodeestroy = (default)
hostname:shares home (uncommitted)>
```

4. To modify the project properties, use the set command. Project properties are described in [“Project Properties” on page 417](#).

5. **Enter `commit`.**

```
hostname:shares home> commit
```

Related Topics

- [“Project Properties” on page 417](#)
- [“Creating an Encrypted Project \(CLI\)” on page 637](#)

▼ **Editing a Project (BUI)**

To modify project properties, use these steps.

1. **Go to Shares > Projects.**
2. **Select a project in one of the following ways:**
 - Hover over the project and click the edit icon .
 - Double-click the project name
 - Click the arrow icon next to Projects to expand the panel, then click on the project name.

The project is selected, and tabs are displayed for editing the properties.

3. **Click one of the tabs to edit the project properties.**
4. **Modify the project properties for the project as needed. See [“Project Properties” on page 417](#).**

Related Topics

- [“Snapshots and Clones” on page 481](#)
- [“Remote Replication” on page 511](#)

▼ **Editing a Project (CLI)**

To modify project properties, use these steps.

1. **Go to shares.**

```
hostname:> shares
```

2. Enter select and a project name.

```
hostname:shares> select home
```

3. To list the project properties, use the get command.

```
hostname:shares home> get
    aclinherit = restricted
    aclmode = discard
    atime = true
    checksum = fletcher4
    compression = off
    compressratio = 100
    copies = 1
    creation = Thu Oct 23 2009 17:30:55 GMT+0000 (UTC)
    mountpoint = /export
    quota = 0
    readonly = false
    recordsize = 128K
    reservation = 0
    rstchown = true
    secondarycache = all
    nbmand = false
    sharesmb = off
    sharenfs = on
    snapdir = hidden
    snaplabel = project1:share1
    vscan = false
    defaultuserquota = 0
    defaultgroupquota = 0
    encryption = off
    snaplabel =
    sharedav = off
    shareftp = off
    sharesftp = off
    sharetftp = off
    pool = Pool1
    canonical_name = Pool1/local/default
    default_group = other
    default_permissions = 700
    default_sparse = false
    default_user = nobody
    default_volblocksize = 8K
    default_volsize = 0
    space_data = 43.9K
    space_unused_res = 0
    space_unused_res_shares = 0
    space_snapshots = 0
```

```
space_available = 12.0T
space_total = 43.9K
origin =
hostname:shares home>
```

4. **To modify the project properties, use the `set` command. Project properties and values are described in [“Project Properties” on page 417](#).**

For example, to enable `vsan` for this project, enter the following command:

```
hostname:shares home >set vsan=true
```

5. **Enter `commit`.**

```
hostname:shares home> commit
```

▼ Renaming a Project (BUI)



Caution - Changing a project name will disrupt active client I/O operations.

1. **Disconnect any active clients connected to the project.**
2. **Go to Shares > Projects.**
3. **Click on the project name in the Projects list.**
4. **Enter the new name for the project.**
A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters `_ - . :`
5. **Press Return.**
6. **Click OK to confirm.**

▼ Renaming a Project (CLI)



Caution - Changing a project name will disrupt active client I/O operations.

1. **Disconnect any active clients connected to the project.**
2. **Go to shares.**


```
hostname:> shares
```

3. To view the projects, use the `list` command.

```
hostname:shares> list
default
home
```

4. Enter `rename`, the existing project name, and the new project name.

A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters `_ - . :`

```
hostname:shares> rename home project1
```

5. To confirm the project was renamed, use the `list` command.

```
hostname:shares> list
default
project1
```

Related Topics

- [“Project Properties” on page 417](#)

▼ Deleting a Project (BUI)



Caution - Deleting a project destroys all data in the project by deleting its filesystems and LUNs.

1. Go to **Shares > Projects**.
2. Hover over the project you want to delete and click the destroy icon .
3. Click **OK**.

▼ Deleting a Project (CLI)



Caution - Deleting a project destroys all data in the project by deleting its filesystems and LUNs.

1. Go to shares.

```
hostname:> shares
```

2. Enter destroy and a project name.

```
hostname:shares> destroy home
```

3. Enter Y.

```
This will destroy all data in "home"! Are you sure? (Y/N)
hostname:shares> Y
```

▼ Creating a Filesystem or LUN in a Project (BUI)

A filesystem or LUN created within a project inherits the properties of the project. For a list of standard properties that can be inherited, see [“Inherited Properties” on page 404](#). If the project is encrypted, a filesystem or LUN created within it is also encrypted.

If you are adding a filesystem or LUN to a non-default project, the project must already exist. To create a new project, see [“Creating a Project \(BUI\)” on page 384](#).

1. Go to Shares > Shares.

2. Select Filesystems or LUNs.

3. Click the add icon .

4. Complete the fields in the Create Filesystem or Create LUN dialog box.

- For a filesystem, select a project and enter a name.
- For a LUN, select a project, enter a name and specify the volume size.

A name must consist of 1 to 64 characters, but not include spaces or begin with a period.

Allowable characters are: alphanumeric and special characters _ - . :

5. Click Apply.

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

- [“Inherited Properties” on page 404](#)
- [“Creating an Encrypted Filesystem or LUN \(BUI\)” on page 640](#)

▼ Creating a Filesystem or LUN in a Project (CLI)

A filesystem or LUN created within a project inherits the properties of the project. For a list of standard properties that can be inherited, see [“Inherited Properties” on page 404](#). If the project is encrypted, a filesystem or LUN created within it is also encrypted.

If you are adding a filesystem or LUN to a non-default project, the project must already exist. To create a new project, see [“Creating a Project \(CLI\)” on page 384](#).

1. Go to shares.

```
hostname:> shares
```

2. Enter `select` and the project name. In this example, the default project is selected.

```
hostname:shares > select default
```

3. Enter `filesystem` and a filesystem name, or `lun` and a LUN name.

A name must consist of 1 to 64 characters, but not include spaces or begin with a period.

Allowable characters are: alphanumeric and special characters `_ - . :`

The following example creates a filesystem named `fs-1` in the default project.

```
hostname:shares default> filesystem fs-1
hostname:shares default/fs-1 (uncommitted)>
```

4. If creating a LUN, enter `set volsize=` and the volume size.

```
hostname:shares default/lun1 (uncommitted)> set volsize=2G
volsize = 2G (uncommitted)
```

5. Enter `commit`.

```
hostname:shares default/fs-1 (uncommitted)> commit
```

6. Enter `select` and a filesystem or LUN name.

```
hostname:shares default> select fs-1
```

7. List the properties of the share using the `get` command:

```
hostname:shares default/fs-1> get
    aclinherit = restricted (inherited)
    aclmode = discard (inherited)
    atime = true (inherited)
casesensitivity = mixed
    checksum = fletcher4 (inherited)
    compression = off (inherited)
    dedup = false (inherited)
    compressratio = 100
    copies = 1 (inherited)
    creation = Wed Apr 29 2015 17:57:18 GMT+0000(UTC)
    logbias = latency (inherited)
    mountpoint = /export/fs-1 (inherited)
    normalization = none
    quota = 0
    quota_snap = true
    readonly = false (inherited)
    recordsize = 128K (inherited)
    reservation = 0
reservation_snap = true
    rstchown = true(inherited)
    secondarycache = all (inherited)
    shadow = none
    nbmand = false (inherited)
    sharesmb = off (inherited)
    sharenfs = on (inherited)
    snapdir = hidden (inherited)
    utf8only = false
    vscan = false (inherited)
    encryption = off (inherited)
    snaplabel =
    sharedav = off (inherited)
    shareftp = off (inherited)
    sharesftp = off (inherited)
    sharetftp = off (inherited)
    pool = pool_demo
canonical_name = pool_demo/local/default/fs-1
    exported = true (inherited)
    nodeestroy = false
    maxblocksize = 1M (inherited)
    lz4supported = (inherited)
    space_data = 31K
space_unused_res = 0
space_snapshots = 0
space_available = 29.4T
space_total = 31K
```

```

root_acl =
owner@:rwxpDaARWcCo:allow,group@:aRc:allow,everyone@:aRc:allow
root_group = other
root_permissions = 700
root_user = nobody
origin =
smbshareacl =

```

8. **To modify the filesystem or LUN properties, use the `set` command. Properties are described in [“Project and Share Properties” on page 403](#).**

For example, to disable the NFS protocol for the filesystem named `fs-1`, enter:

```

hostname:shares default/fs-1> set sharenfs=off
sharenfs = off (uncommitted)

```

9. **Enter `commit`.**

```

hostname:shares default/fs-1> commit
hostname:shares default/fs-1>

```

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)
- [“Inherited Properties” on page 404](#)
- [“Creating an Encrypted Filesystem or LUN \(CLI\)” on page 641](#)

▼ Editing a Filesystem or LUN (BUI)

To modify properties for an individual filesystem or LUN, use these steps.

1. **Go to Shares > Shares.**
2. **Select Filesystems or LUNs.**
3. **Hover over the filesystem or LUN and click the edit icon , or double-click the filesystem or LUN you want to edit.**
The general properties are displayed for the filesystem or LUN.
4. **Click the Protocols, Access, Snapshots, or Replication tab.**

5. **Modify the filesystem or LUN properties described in “[Filesystem Properties](#)” on page 423 and “[LUN Properties](#)” on page 431.**

Related Topics

- “[Project and Share Properties](#)” on page 403

▼ Editing a Filesystem or LUN (CLI)

To modify properties for an individual filesystem or LUN, use these steps.

1. **Go to shares.**

```
hostname:> shares
```

2. **Enter `select` and the project name that contains the filesystem or LUN you want to edit.**

```
hostname:shares> select default
```

3. **Enter `select` and a filesystem name or LUN name.**

```
hostname:shares default> select fs-1
```

4. **List the properties of the share using the `get` command:**

```
hostname:shares default/fs-1> get
aclinherit = restricted (inherited)
aclmode = discard (inherited)
atime = true (inherited)
casesensitivity = mixed
checksum = fletcher4 (inherited)
compression = off (inherited)
dedup = false (inherited)
compressratio = 100
copies = 1 (inherited)
creation = Wed Apr 29 2015 17:57:18 GMT+0000(UTC)
logbias = latency (inherited)
mountpoint = /export/fs-1 (inherited)
normalization = none
quota = 0
quota_snap = true
readonly = false (inherited)
recordsize = 128K (inherited)
reservation = 0
reservation_snap = true
```

```

        rstchown = true(inherited)
        secondarycache = all (inherited)
        shadow = none
        nbmand = false (inherited)
        sharesmb = off (inherited)
        sharenfs = on (inherited)
        snapdir = hidden (inherited)
        utf8only = false
        vscan = false (inherited)
        encryption = off (inherited)
        snaplabel =
        sharedav = off (inherited)
        shareftp = off (inherited)
        sharesftp = off (inherited)
        sharetftp = off (inherited)
        pool = pool_demo
        canonical_name = pool_demo/local/default/fs-1
        exported = true (inherited)
        nodeestroy = false
        maxblocksize = 1M (inherited)
        space_data = 31K
        space_unused_res = 0
        space_snapshots = 0
        space_available = 29.4T
        space_total = 31K
        root_acl =
owner@:rwxpDaARWcCo:allow,group@:aRc:allow,everyone@:aRc:allow
        root_group = other
        root_permissions = 700
        root_user = nobody
        origin =
        smbshareacl =

```

5. **Use the `set` command to modify the filesystem or LUN properties described in [“Filesystem Properties” on page 423](#) and [“LUN Properties” on page 431](#).**

For example, to disable the NFS protocol for the filesystem named `fs-1`, enter:

```

hostname:shares default/fs-1> set sharenfs=off
        sharenfs = off (uncommitted)

```

6. **Enter `commit`.**

```

hostname:shares default/fs-1> commit

```

Related Topics

- [“Project and Share Properties” on page 403](#)

▼ Renaming a Filesystem or LUN (BUI)



Caution - Changing a share name will disrupt active client I/O operations.

1. **Disconnect all active clients connected to the filesystem or LUN you want to rename.**
2. **Go to Shares > Shares.**
3. **Select Filesystems or LUNs.**
4. **Click on the filesystem or LUN name in the list.**
5. **Enter the new name for the filesystem or LUN.**
A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters _ - . :
6. **Press Return.**
7. **Click OK to confirm.**

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Renaming a Filesystem or LUN (CLI)



Caution - Changing a share name will disrupt active client I/O operations.

1. **Disconnect all active clients connected to the filesystem or LUN.**
2. **Go to shares.**
3. **To view the projects, use the list command.**

```
hostname:> shares
```

```
hostname:shares>list
default
```


home

4. **Enter `select` and the project name that contains the filesystem or LUN you want to rename.**

```
hostname:shares>select default
```

5. **Enter `rename`, the existing filesystem or LUN name, and the new filesystem or LUN name.**

A name must consist of 1 to 64 characters, but not include spaces or begin with a period.

Allowable characters are: alphanumeric and special characters `_ - . :`

```
hostname:shares default> rename fs-1 fs-2
```

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Moving a Filesystem or LUN to a Different Project (BUI)

Filesystems and LUNs within a project inherit the properties of the project.

1. **Go to Shares > Shares.**
2. **Select Filesystems or LUNs.**
3. **Hover over the filesystem or LUN and click the move icon .**
4. **Drag the filesystem or LUN to the different project under Projects.**

If the project panel is not expanded, the panel will automatically expand until the share is dropped onto a project.

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Moving a Filesystem or LUN to a Different Project (CLI)

Filesystems and LUNs within a project inherit the properties of the project.

1. **Go to shares and select the project that contains the filesystem or LUN to be moved.**

In this example, the default project contains the filesystem or LUN to be moved.

```
hostname> shares
hostname:shares> select default
```

2. **Enter move, the name of the filesystem or LUN to be moved, and the name of the project to move it to.**

```
hostname:shares default> move foo home
```

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Deleting a Filesystem or LUN (BUI)



Caution - Deleting a filesystem or LUN destroys all data in the share and cannot be undone.

1. **Go to Shares > Shares.**
2. **Select Filesystems or LUNs.**
3. **Hover over the filesystem or LUN you want to delete and click the destroy icon .**
4. **Click OK.**

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Deleting a Filesystem or LUN (CLI)



Caution - Deleting a filesystem or LUN destroys all data in the share and cannot be undone.

1. **Go to shares.**

```
hostname> shares
```

2. **Enter select and the project name that contains the filesystem or LUN.**

```
hostname:shares> select default
```

3. **Enter select and the filesystem or LUN name.**

```
hostname:shares default>select fs-1
```

4. **Enter destroy.**

```
hostname:shares default/fs-1> destroy
This will destroy all data in "fs-1"! Are you sure? (Y/N)
```

5. **Enter Y.**

```
hostname:shares default> Y
```

Related Topics

- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

▼ Setting User or Group Quotas (BUI)

Quotas can be set for a user or group at the project or filesystem level.

1. **Go to Shares > Shares and select a project or share.**
2. **Click the General tab.**
3. **In the Space Usage - Users & Groups section, select User, Group, or User or Group from the drop down menu.**

Note - Any user that is not consuming any space on the filesystem, and does not have any quota set, does not appear in the list of active users.

4. **To set a quota at the project level, select one of three options:**

- None - No quota is set for this filesystem.
- Default - Sets the quota to the default quota at the project level; if no default was set, no quota is set for this filesystem.
- Click the radio button, enter a quota in the size field, and select a measurement.

5. Click APPLY.

The user and group quota properties are validated separately from the other properties. However, you may only see one validation error if an invalid user/group as well as another invalid property is entered. Correcting one error and applying the changes will show any remaining error messages.

If you see an error message that an invalid property has been entered, it may be an invalid user/group, another invalid property on the page, or both. Fixing one invalid property and then applying the changes will reveal any remaining error messages.

Related Topics

- [“Setting User or Group Quotas” on page 438](#)

▼ Setting User or Group Quotas (CLI)

Quotas can be set for a user or group at the project or filesystem level.

1. Go to shares, select a project, then select a share, as shown in this example:

```
hostname:> shares select default select eschrock
```

2. Enter users, then list to see the current users.

```
hostname:shares default/eschrock> users
hostname:shares default/eschrock users> list
USER      NAME      USAGE  QUOTA  SOURCE
user-000  root      321K   -      -
user-001  ahl       9.94K  -      -
user-002  eschrock  20.0G  -      -
```

Note - Any user that is not consuming any space on the filesystem, and does not have any quota set, does not appear in the list of active users.

3. Enter select and the name= of the user.

```

hostname:shares default/eschrock users> select name=eschrock
hostname:shares default/eschrock user-002> get
      name = eschrock
  unixname = eschrock
    unixid = 132651
    winname = (unset)
    winid   = (unset)
    usage   = 20.0G
    quota   = (unset)
    source   = (unset)

```

4. Enter quota= and a value. Enter commit and done.

Note - To clear a quota, set the value to '0'.

```

hostname:shares default/eschrock user-002> set quota=100G
      quota = 100G (uncommitted)
hostname:shares default/eschrock user-002> commit
hostname:shares default/eschrock user-002> done

```

5. To set a quota for such a user or group, use the quota command, after which the name and quota can be set.

The Source column displays "local" if the quota was set at the filesystem level, "default" if set at the project level, or "-" if no quota was set. In the following example, the default user quota set at the project level is 50 GB.

If a default user or group quota was set at the project level, this procedure overrides that value.

```

hostname:shares default/eschrock users> quota
hostname:shares default/eschrock users quota (uncommitted)> set name=bmc
      name = bmc (uncommitted)
hostname:shares default/eschrock users quota (uncommitted)> set quota=default
      quota = default (uncommitted)
hostname:shares default/eschrock users quota (uncommitted)> commit
hostname:shares default/eschrock users> list

```

USER	NAME	USAGE	QUOTA	SOURCE
user-000	root	321K	-	-
user-001	ahl	9.94K	-	-
user-002	eschrock	20.0G	100G	local
user-003	bmc	-	50G	default

Related Topics

- [“Setting User or Group Quotas” on page 438](#)

About Storage Pools, Projects, and Shares

The Oracle ZFS Storage Appliance product manages physical storage using a pooled storage model where all filesystems and LUNs share common space. This topic describes how storage is organized using storage pools, projects, and shares.

Storage Pools

The appliance is based on the ZFS filesystem, which groups underlying storage devices into pools. Filesystems and LUNs, collectively referred to as shares, allocate from this storage pool as needed. Before creating filesystems or LUNs, you must first configure storage on the appliance. Once a storage pool is configured, there is no need to statically size filesystems, though this behavior can be achieved by using quotas and reservations.

While multiple storage pools are supported, this type of configuration is generally discouraged because it provides significant drawbacks as described in the [“Configuring Storage” on page 122](#) section. Multiple pools should only be used where the performance or reliability characteristics of two different profiles are drastically different, such as a mirrored pool for databases and a RAID-Z pool for streaming workloads.

When multiple pools are active on a single host, the BUI displays a drop-down list in the menu bar that can be used to switch between pools. In the CLI, the name of the current pool will be displayed in parenthesis, and can be changed by setting the 'pool' property. If there is only a single pool configured, then these controls will be hidden. When multiple pools are selected, the default pool chosen by the UI is arbitrary, so any scripted operation should be sure to set the pool name explicitly before manipulating any shares.

Projects

All filesystems and LUNs are grouped into projects. A project can be considered a consistency group, that defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level, as well as the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space) can be accessed from a single point.

By default, the appliance creates a single default project when a storage pool is first configured. It is possible to create all shares within this default project, although for reasonably sized environments creating additional projects is strongly recommended, if only for organizational purposes.

Shares

Shares are filesystems and LUNs that are exported over supported data protocols to clients of the appliance. Exported filesystems can be accessed over SMB, NFS, HTTP/WebDav, and FTP. LUNs export block-based volumes and can be accessed over iSCSI or Fibre Channel.

The project/share is a unique identifier for a share within a pool. Projects within a pool cannot contain shares with the same name. If you attempt to name or rename a share using a name that is already in use, a mount point error occurs.

In addition to the default properties, you can configure shares and projects with any number of additional properties. These properties are given basic types for validation purposes, and are inherited like most other standard properties. The values are never consumed by the software in any way, and exist solely for end-user consumption. The property schema is global to the system, across all pools, and is synchronized between cluster peers.

Related Topics

- [“Space Management for Shares” on page 436](#)
- [“Project and Share Properties” on page 403](#)
- [“Snapshots and Clones” on page 481](#)

Project and Share Properties

All projects and shares have a number of associated properties which can be set using the BUI or CLI. For a list of property names and descriptions, click on one of these links:

- [“Project Properties” on page 417](#)
- [“Filesystem Properties” on page 423](#)
- [“LUN Properties” on page 431](#)

Project and share properties can be one of the following types:

TABLE 111 Project and Share Property Types

Property Type	Description
Inherited	Inherited properties, the most common property type, represent most of the configurable project and share properties. Shares that are part of a project can either have local settings for properties, or they can inherit their settings from the parent project. By default, shares inherit all properties from the project. If a property is changed on a project, all shares that inherit that property are updated to reflect the new value. When inherited, all properties have the same value as the parent project, with the exception of the mountpoint and SMB properties. When inherited, these properties concatenate the project setting with their own share name.
Read only	Read-only properties represent statistics about the project and share and cannot be changed. The most common properties of this type are space usage statistics.

Property Type	Description
Space management	Space management properties (quota and reservation) apply to both shares and projects, but are not inherited. A project with a quota of 100G will be enforced across all shares, but each individual share will have no quota unless explicitly set.
Static (Create time)	Static properties are specified at filesystem or LUN creation time, but cannot be changed once the share has been created. These properties control the on-disk data structures, and include internationalization settings, case sensitivity, and volume block size.
Project default	Project default properties are set on a project, but do not affect the project itself. They are used to populate the initial settings when creating a filesystem or LUN, and can be useful when shares have a common set of non-inheritable properties. Changing these properties do not affect existing shares, and the properties can be changed before or after creating the share.
Filesystem local	Filesystem local properties apply only to filesystems, and are convenience properties for managing the root directory of the filesystem. They cannot be set on projects. These access control properties can also be set by in-band protocol operations.
LUN local	LUN local properties apply only to LUNs and are not inherited. They cannot be set on projects.
Custom	Custom properties are user-defined properties.

Inherited Properties

Inherited properties are standard properties that can either be inherited from the project or explicitly set on the share. The BUI only allows the properties to be inherited all at once, while the CLI allows for individual properties to be inherited.

Shares that are part of a project can either have local settings for properties, or they can inherit their settings from the parent project. By default, shares inherit all properties from the project. If a property is changed on a project, all shares that inherit that property are updated to reflect the new value. When inherited, all properties have the same value as the parent project, with the exception of the mountpoint and SMB properties. When inherited, these properties concatenate the project setting with their own share name.

Mountpoint

The mountpoint property is the location where the filesystem is mounted. This property is only valid for filesystems.

The following restrictions apply to the mountpoint property:

- Must be under /export
- Cannot conflict with another share
- Cannot conflict with another share on cluster peer to allow for proper failover

When inheriting the mountpoint property, the current dataset name is appended to the project's mountpoint setting, joined with a slash ('/'). For example, if the "home" project has the mountpoint setting /export/home, then "home/bob" would inherit the mountpoint /export/home/bob.

SMB shares are exported via their resource name, and the mountpoint is not visible over the protocol. However, even SMB-only shares must have a valid unique mountpoint on the appliance.

Mountpoints can be nested underneath other shares, though this has some limitations. For more information, see [“Working with Filesystem Namespace” on page 439](#).

Read only

The read-only property controls whether the filesystem contents are read only. This property is only valid for filesystems.

The contents of a read-only filesystem cannot be modified, regardless of any protocol settings. This setting does not affect the ability to rename, destroy, or change properties of the filesystem. In addition, when a filesystem is read only, access control properties cannot be altered, because they require modifying the attributes of the root directory of the filesystem.

Update access time on read

The update access time on read property controls whether the access time for files is updated on read. This property is only valid for filesystems.

POSIX standards require that the access time for a file properly reflect the last time it was read. This requires issuing writes to the underlying filesystem even for a mostly read-only workload. For working sets consisting primarily of reads over a large number of files, turning off this property may yield performance improvements at the expense of standards conformance. These updates happen asynchronously and are grouped together, so its effect should not be visible except under heavy load.

Non-blocking mandatory locking

The non-blocking mandatory locking property controls whether SMB locking semantics are enforced over POSIX semantics. This property is only valid for filesystems.

By default, filesystems implement file behavior according to POSIX standards. These standards are fundamentally incompatible with the behavior required by the SMB protocol. For shares where the primary protocol is SMB, this option should always be enabled. Changing this property requires all clients to be disconnected and reconnect.

Data Deduplication

The data deduplication property controls whether duplicate copies of data are eliminated. Deduplication is synchronous, pool-wide, block-based, and can be enabled on a per project or share basis.

Before deduplication can be enabled on a project or share, configure the storage pool with meta devices. Meta devices are designated cache devices used to store specific types of metadata to optimize use cases like deduplication.

Deduplication is also only available on datasets with a record size 128K or above.

To enable deduplication, select the Data Deduplication checkbox on the general properties screen for projects or shares. The deduplication ratio will appear in the usage area of the Status Dashboard. Data written with deduplication enabled is entered into the deduplication table indexed by the data checksum. Deduplication forces the use of the cryptographically strong SHA-256 checksum. Subsequent writes will identify duplicate data and retain only the existing copy on disk. Deduplication can only happen between blocks of the same size, data written with the same record size. For best results, set the record size to that of the application using the data; for streaming workloads, use a large record size.

Note - Starting with the Data Deduplication deferred update in OS8.7, if replication is configured for a deduplicated project or share, a compatibility test will run to determine if the replication target has the required software and meta device for receiving deduplicated updates. If the target is running OS8.6 or earlier, replication updates with deduplication enabled will fail, and an alert will post indicating that data deduplication needs to be disabled at the source. If the target is running OS8.7, but does not have the required meta device, the target will ignore the incoming data deduplication property, and an alert will post to show that the target system will deliberately ignore the deduplication settings during replication receive.

If a OS8.6 or earlier source has been replicating to a target with deduplication enabled, the compatibility test will check the target for a meta device or deduplicated share in the package. If the compatibility test finds either of these, the target will preserve the deduplication settings when receiving replication updates.

If your data does not contain any duplicates, enabling data deduplication will add overhead (a more CPU-intensive checksum and on-disk deduplication table entries) without providing any benefit. If your data does contain duplicates, enabling data deduplication will both save space by storing only one copy of a given block regardless of how many times it occurs. Deduplication could impact performance in that the checksum is more expensive to compute and the metadata of the deduplication table must be accessed and maintained.

Note that deduplication has no effect on the calculated size of a share, but does affect the amount of space used for the pool. For example, if two shares contain the same 1GB file, each will appear to be 1GB in size, but the total for the pool will be just 1GB and the deduplication ratio will be reported as 2x.

There are 3 sets of analytics used to monitor performance of deduplication:

- **ZFS DMU operations (by DMU object type)** - This analytic will show you how many operations are being performed against the Data Deduplication Table compared to other ZFS operations.
- **Meta device bytes used (by pool)** - Amount of space used on the metadata devices.
This statistic will remain blank until at least 1% of the meta device capacity is used.
- **Meta device percent used (by pool)** - Percent of space used on the metadata devices.
This statistic will remain blank until at least 1% of the meta device capacity is used.

To use deduplication with encryption, keep in mind that only AES with the CCM mode encryption is compatible with deduplication. For more information, see [“Managing Encryption Keys” on page 652](#).

Data Compression

The data compression property controls whether data is compressed before being written to disk. Shares can optionally compress data before writing to the storage pool. This allows for much greater storage utilization at the expense of increased CPU utilization. By default, no compression is done. If the compression does not yield a minimum space savings, it is not committed to disk to avoid unnecessary decompression when reading back the data. Before choosing a compression algorithm, it is recommended that you perform any necessary performance tests and measure the achieved compression ratio.

BUI value	CLI value	Description
Off	off	No compression is done.
LZ4	lz4	An algorithm that typically consumes less CPU than GZIP-2, but compresses better than LZJB, depending on the data that is compressed.
LZJB (Fastest)	lzjb	A simple run-length encoding that only works for sufficiently simple inputs, but doesn't consume much CPU.
GZIP-2 (Fast)	gzip-2	A lightweight version of the gzip compression algorithm.
GZIP (Default)	gzip	The standard gzip compression algorithm.
GZIP-9 (Best Compression)	gzip-9	Highest achievable compression using gzip. This consumes a significant amount of CPU and can often yield only marginal gains.

Checksum

The checksum property controls the checksum used for data blocks. On the appliance, all data is checksummed on disk, and in such a way to avoid traditional pitfalls (phantom reads and write in particular). This allows the system to detect invalid data returned from the devices. The default checksum (fletcher4) is sufficient for normal operation, but users can increase the checksum strength at the expense of additional CPU load. Metadata is always checksummed using the same algorithm, so this only affects user data (files or LUN blocks).

BUI value	CLI value	Description
Fletcher 2 (Legacy)	fletcher2	16-bit fletcher checksum

BUI value	CLI value	Description
Fletcher 4 (Standard)	fletcher4	32-bit fletcher checksum
SHA-256 (Extra Strong)	sha256	SHA-256 checksum
SHA-256-MAC	sha256mac	

Cache device usage

The cache device usage property controls whether cache devices are used for the share. By default, all datasets make use of any cache devices on the system. Cache devices are configured as part of the storage pool and provide an extra layer of caching for faster tiered access. For more information on cache devices, see [“Configuring Storage” on page 122](#). This property is independent of whether there are any cache devices currently configured in the storage pool. For example, it is possible to have this property set to "all" even if there are no cache devices present. If any such devices are added in the future, the share will automatically take advantage of the additional performance. This property does not affect use of the primary (DRAM) cache.

BUI Value	CLI Value	Description
All data and metadata	all	All normal file or LUN data is cached, as well as any metadata.
Metadata only	metadata	Only metadata is kept on cache devices. This allows for rapid traversal of directory structures, but retrieving file contents may require reading from the data devices.
Do not use cache devices	none	No data in this share is cached on the cache device. Data is only cached in the primary cache or stored on data devices.

Synchronous write bias

The synchronous write bias property controls the behavior when servicing synchronous writes. By default, the system optimizes synchronous writes for latency, which leverages the log devices to provide fast response times. In a system with multiple disjointed filesystems, this can cause contention on the log devices that can increase latency across all consumers. Even with multiple filesystems requesting synchronous semantics, it may be the case that some filesystems are more latency-sensitive than others.

A common case is a database that has a separate log. The log is extremely latency sensitive, and while the database itself also requires synchronous semantics, it is heavier bandwidth and not latency sensitive. In this environment, setting this property to 'throughput' on the main database

while leaving the log filesystem as 'latency' can result in significant performance improvements. This setting will change behavior even when no log devices are present, though the effects may be less dramatic.

The synchronous write bias setting can be bypassed by the Oracle Intelligent Storage Protocol. Instead of using the write bias defined in the file system, the Oracle Intelligent Storage Protocol can use the write bias value provided by the Oracle Database NFSv4.0 or NFSv4.1 client. The write bias value sent by the Oracle Database NFSv4.0 or NFSv4.1 client is used only for that write request.

BUI Value	CLI Value	Description
Latency	latency	Synchronous writes are optimized for latency, leveraging the dedicated log device(s), if any.
Throughput	throughput	Synchronous writes are optimized for throughput. Data is written to the primary data disks instead of the log device(s), and the writes are performed in a way that optimizes for total bandwidth of the system. Log devices will be used for small amounts of metadata associated with the data writes.

Database record size

The database record size property specifies a suggested block size for files in the file system. This property is only valid for filesystems and is designed for use with database workloads that access files in fixed-size records. The system automatically tunes block sizes according to internal algorithms optimized for typical access patterns. For databases that create very large files but access them in small random chunks, these algorithms may be suboptimal. Specifying a record size greater than or equal to the record size of the database can result in significant performance gains. Use of this property for general purpose file systems is strongly discouraged, and may adversely affect performance.

The default record size is 128 KB. The size specified must be a power of two greater than or equal to 512 and less than or equal to 1 MB. Changing the file system's record size affects only files created afterward; existing files and received data are unaffected. If block sizes greater than 128K are used for projects or shares, replication of those projects or shares to systems that do not support large block sizes will fail.

The database record size setting can be bypassed by the Oracle Intelligent Storage Protocol. Instead of using the record size defined in the file system the Oracle Intelligent Storage Protocol can use the block size value provided by the Oracle Database NFSv4.0 or NFSv4.1 client. The

block size provided by the Oracle Database NFSv4.0 or NFSv4.1 client can only be applied when creating a new database files or table. Block sizes of existing files and tables will not be changed. For more information, see [“Oracle Intelligent Storage Protocol” on page 688](#).

Additional replication

The additional replication property controls the number of copies stored of each block, above and beyond any redundancy of the storage pool. Metadata is always stored with multiple copies, but this property allows the same behavior to be applied to data blocks. The storage pool attempts to store these extra blocks on different devices, but it is not guaranteed. In addition, a storage pool cannot be imported if a complete logical device (RAID stripe, mirrored pair, etc) is lost. This property is not a replacement for proper replication in the storage pool, but can be reassuring for paranoid administrators.

Virus scan

The virus scan property controls whether the filesystem is scanned for viruses. This property is only valid for filesystems. This property setting is independent of the state of the virus scan service. Even if the Virus Scan service is enabled, filesystem scanning must be explicitly enabled using this property. Similarly, virus scanning can be enabled for a particular share even if the service itself is off. For more information about configuration virus scanning, see [Virus Scan](#).

Prevent destruction

When set, the share or project cannot be destroyed. This includes destroying a share through dependent clones, destroying a share within a project, or destroying a replication package. However, it does not affect shares destroyed through replication updates. If a share is destroyed on an appliance that is the source for replication, the corresponding share on the target will be destroyed, even if this property is set. To destroy the share, the property must first be explicitly turned off as a separate step. This property is off by default.

Restrict ownership change

By default, ownership of files cannot be changed except by a root user (on a suitable client with a root-enabled export). This property can be turned off on a per-filesystem or per-project basis by turning off this property. When off, file ownership can be changed by the owner of the file or directory, effectively allowing users to "give away" their own files. When ownership

is changed, any setuid or setgid bits are stripped, preventing users from escalating privileges through this operation.

LUN Local Properties

These properties apply only to LUNs and are not inherited. They cannot be set on projects.

Volume size

The volume size property is the logical size of the LUN as exported over iSCSI. This property controls the size of the LUN. By default, LUNs reserve enough space to completely fill the volume. Changing the size of a LUN while actively exported to clients may yield undefined results. It may require clients to reconnect and/or cause data corruption on the filesystem on top of the LUN. Check best practices for your particular iSCSI client before attempting this operation.

Thin provisioned

The thin provisioned property controls whether space is reserved for the volume. By default, a LUN reserves exactly enough space to completely fill the volume. This ensures that clients will not get out-of-space errors at inopportune times. This property allows the volume size to exceed the amount of available space. When set, the LUN will consume only the space that has been written to the LUN. While this allows for thin provisioning of LUNs, most filesystems do not expect to get "out of space" from underlying devices, and if the share runs out of space, it may cause instability and/or data corruption on clients.

When not set, the volume size behaves like a reservation excluding snapshots. It therefore has the same pathologies, including failure to take snapshots if the snapshot could theoretically diverge to the point of exceeding the amount of available space. For more information, see the Reservation property in [“Managing Filesystem and Project Space” on page 437](#).

Volume block size

The volume block size property sets the native block size for LUNs. This can be any power of 2 from 512 bytes to 1M, and the default is 8K. This property is static; it is set when the LUN is created and cannot be changed.

Note - LUNs with a volume block size smaller than 4K may cause performance degradation.

Other Properties

The following properties are available: Project Default, Filesystem Local, Space Management, Read-only, and Custom.

Project default

The project default properties are set on a project, but do not affect the project itself. They are used to populate the initial settings when creating a filesystem or LUN, and can be useful when shares have a common set of non-inheritable properties. Changing these properties do not affect existing shares, and the properties can be changed before or after creating the share.

Filesystem local

The filesystem local properties apply only to filesystems, and are convenience properties for managing the root directory of the filesystem. They are not inherited and cannot be set on projects. These access control properties can also be set by in-band protocol operations.

Space management

The space management properties (quota and reservation) apply to both shares and projects, but are not inherited. A project with a quota of 100G will be enforced across all shares, but each individual share will have no quota unless explicitly set.

Read only

The read-only properties represent statistics about the project and share and cannot be changed. The most common properties of this type are space usage statistics.

Custom

Custom properties are user-defined using a schema. For more information, see [“Working with Schemas” on page 464](#).

Static Properties

Static (create time) properties are specified at filesystem or LUN creation time, but cannot be changed once the share has been created. These properties control the on-disk data structures, and include internationalization settings, case sensitivity, and volume block size.

In the BUI, static properties can be viewed on the left side of the interface when editing a filesystem or LUN.

TABLE 112 Filesystem and LUN Static Properties

BUI Name	CLI Name	Description
Creation date	creation	Indicates the date of creation.
Compression ratio	compressratio	Current compression ratio for the filesystem or LUN, which is a product of the compression algorithm. For more information, see “Compression ratio” on page 415 .
Case sensitivity	casesensitivity	The case sensitivity property controls whether directory lookups are case-sensitive or case-insensitive. For more information, see “Case sensitivity” on page 415 .
Reject non UTF-8	utf8only	This property enforces UTF-8 encoding for all files and directories. For more information, see “Reject non UTF-8” on page 416 .
Normalization	normalization	The normalization property controls what unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encodings. For more information, see “Normalization” on page 416 .
Volume block size (LUNs only)	volblocksize	The volume block size property sets the native block size for LUNs. For more information, see “Volume block size” on page 417 .
Origin	origin	Shows the name of the snapshot from which it was cloned. For more information, see “Origin” on page 417 .
Data migration source (Filesystems only)	shadow	Location of the source if the filesystem is actively shadowing an existing filesystem, either locally or over NFS. For more information, see “Data Migration Source” on page 417 .

Compression ratio

If compression is enabled, this property shows the compression ratio currently achieved for the share. This is expressed as a multiplier. For example, a compression of 2x means that the data is consuming half as much space as the uncompressed contents. For more information about selecting a compression algorithm, see "Data Compression" described in ["Inherited Properties" on page 404](#).

Case sensitivity

The case sensitivity property controls whether directory lookups are case-sensitive or case-insensitive. It supports the following options:

BUI Value	CLI Value	Description
Mixed	mixed	Case sensitivity depends on the protocol being used. For NFS, FTP, and HTTP, lookups are case-sensitive. For SMB, lookups are case-insensitive. This is default, and prioritizes conformance of the various protocols over cross-protocol consistency. When using this mode, it's possible to create files that are distinct over case-sensitive protocols, but clash when accessed over SMB. In this situation, the SMB server will create a "mangled" version of the conflicts that uniquely identify the filename.
Insensitive	insensitive	All lookups are case-insensitive, even over protocols (such as NFS) that are traditionally case-sensitive. This can cause confusion for clients of these protocols, but prevents clients from creating name conflicts that would cause mangled names to be used over SMB. This setting should only be used where SMB is the primary protocol and alternative protocols are considered second-class, where conformance to expected standards is not an issue.
Sensitive	sensitive	All lookups are case-sensitive, even over SMB where lookups are traditionally case-insensitive. In general, this setting should not be used because the SMB server can

BUI Value	CLI Value	Description
		deal with name conflicts via mangled names, and may cause Windows applications to behave strangely.

Reject non UTF-8

This property enforces UTF-8 encoding for all files and directories. When set, attempts to create a file or directory with an invalid UTF-8 encoding will fail. This only affects NFSv3, where the encoding is not defined by the standard. NFSv4.0 and NFSv4.1 always use UTF-8, and SMB negotiates the appropriate encoding. This setting should normally be "on", or else SMB (which must know the encoding in order to do case sensitive comparisons, among other things) will be unable to decode filenames that are created with an invalid UTF-8 encoding. This setting should only be set to "off" in pre-existing NFSv3 deployments where clients are configured to use different encodings. Enabling SMB, NFSv4.0 or NFSv4.1 when this property is set to "off" can yield undefined results if a NFSv3 client creates a file or directory that is not a valid UTF-8 encoding. This property must be set to "on" if the normalization property is set to anything other than "none".

Normalization

The normalization property controls what unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encodings. Without normalization, the on-disk name stored will be different, and lookups using one of the alternative forms will fail depending on how the file was created and how it is accessed. If this property is set to anything other than "none" (the default), the "Reject non UTF-8" property must also be set to "on".

BUI Value	CLI Value	Description
None	none	No normalization is done.
Form C	formC	<i>Normalization Form Canonical Composition (NFC)</i> - Characters are decomposed and then recomposed by canonical equivalence.
Form D	formD	<i>Normalization Form Canonical Decomposition (NFD)</i> - Characters are decomposed by canonical equivalence.
Form KC	formKC	<i>Normalization Form Compatibility Composition (NFKC)</i> - Characters are decomposed by compatibility

BUI Value	CLI Value	Description
		equivalence, then recomposed by canonical equivalence.
Form KD	formKD	<i>Normalization Form Compatibility Decomposition (NFKD)</i> - Characters are decomposed by compatibility equivalence.

Volume block size

The volume block size property sets the native block size for LUNs. This can be any power of 2 from 512 bytes to 1M, and the default is 8K.

Note - LUNs with a volume block size smaller than 4K may cause performance degradation.

Origin

If this is a clone, this is the name of the snapshot from which it was cloned.

Data Migration Source

If set, then this filesystem is actively shadowing an existing filesystem, either locally or over NFS. For more information about data migration, see [“Shadow Migration” on page 469](#).

Project Properties

Note - In the CLI, use the `get` command to see a list of all properties.

Use the `list` command to list all children.

TABLE 113 Projects Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
Create Project	Name	project	Static	Defines the name of the project.
	Encryption	encryption	Inherited	Defines the encryption type.

Project Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
	Key	key	Inherited	Sets a specific LOCAL or OKM key.
	Keyname	keyname	Static	Identifies the key.
General - Space Usage - Data	Quota	quota	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.
	Reservation	reservation	Space management	Represents a guarantee of space that can be consumed by any particular entity.
General - Space Usage - Users & Groups	Default user quota	defaultuserquota	Space management	Sets a limit on the amount of space that can be consumed by the user.
	Default group quota	defaultgroupquota	Space management	Sets a limit on the amount of space that can be consumed by the group.
	User and group	users / groups	--	Specifies users and/or groups.
	Usage	--	Space management	Shows the amount of data used by users and/or groups.
General - Inherited Properties	Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see “Mountpoint” on page 404.
	Read only	readonly	Inherited	Controls whether the filesystem contents are read only. For more information, see “Read only” on page 405.
	Update access time on read	atime	Inherited	Controls whether the access time for files is updated on read. For more information, see “Update access time on read” on page 405.
	Non-blocking mandatory locking	nbmand	inherited	Controls whether SMB locking semantics are enforced over POSIX semantics. For more information, see “Non-blocking mandatory locking” on page 406.

BUI Location	BUI Name	CLI Name	Property Type	Description
	Data deduplication	dedup	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see “Data Deduplication” on page 406 .
	Data compression	compression	Inherited	Controls whether data is compressed before being written to disk. For more information, see “Data Compression” on page 408 .
	Checksum	checksum	Inherited	Controls the checksum used for data blocks. For more information, see “Checksum” on page 408 .
	Cache device usage	secondarycache	Inherited	Controls whether cache devices are used for the share. For more information, see “Cache device usage” on page 409 .
	Synchronous write bias	logbias	Inherited	Controls the behavior when servicing synchronous writes. For more information, see “Synchronous write bias” on page 409 .
	Database record size	recordsize	Inherited	Specifies a suggested block size for files in the filesystem. For more information, see “Database record size” on page 410 .
	Additional Replication	copies	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see “Additional replication” on page 411 .
	Virus scan	vscan	Inherited	Controls whether a filesystem is scanned for viruses. For more information, see “Virus scan” on page 411 .
	Prevent destruction	nodestroy	Inherited	Prevents shares or projects from being destroyed when set. For more

Project Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
				information, see “Prevent destruction” on page 411.
	Restrict ownership change	rstchown	Inherited	Controls the ownership and can be turned off on a per-filesystem or per-project basis. For more information, see “Restrict ownership change” on page 411.
General - Custom Properties	Schema	custom	--	Can be added as needed to attach user-defined tags to projects and shares. For more information, see “Schema Properties” on page 466.
General - Default Settings - Filesystems	User	default_user	Creation default	Specifies a user ID or user name.
	Group	default_group	Creation default	Specifies a group ID or group name.
	Permissions	default_permissions	Creation default	Sets the default permissions for filesystem.
General - Default Settings - LUNs	Volume size	default_volsize	LUN only, creation default	Shows the maximum volume size and unit of measurement. For more information, see “Volume size” on page 412.
	Thin provisioned	default_sparse	LUN only, creation default	Indicates only the amount of space physically consumed by data is used when selected. For more information, see “Thin provisioned” on page 412.
	Volume block size	default_volblocksize	Creation default	Shows the native block size for LUNs and can be set from 512 bytes to 1M (the default is 8K). For more information, see “Volume block size” on page 412.
Protocols - NFS	NFS	share nfs	Inherited	NFS Protocol property settings and values are described in “NFS Protocol Properties” on page 443.
Protocols - NFS Exceptions	Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see “NFS Share Mode Exceptions” on page 444			

BUI Location	BUI Name	CLI Name	Property Type	Description
Protocols - SMB	SMB	sharesmb	Inherited	SMB Protocol property settings and values are described in “SMB Protocol Properties” on page 449.
Protocols - SMB Exceptions	Exceptions to the overall sharing modes may be defined for clients or collections of clients. See “SMB Protocol Share Mode Exceptions” on page 452.			
Protocols - HTTP (Inherit from project)	Share mode	sharedav	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - FTP (Inherit from project)	Share mode	shareftp	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - SFTP (Inherit from project)	Share mode	sharesftp	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - TFTP (Inherit from project)	Share mode	sharetftp	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Access	ACL behavior on mode change	aclmode	Inherited	Controls how a mode change request interacts with the existing ACL.
	ACL inheritance behavior	aclinherit	Inherited	Controls how a new file or directory inherits existing ACL settings from the parent directory.
Snapshots - Properties	.zfs/snapshot visibility	snapdir	Inherited	Controls whether filesystem snapshots can be accessed over data protocols at .zfs/snapshot in the root of the filesystem.
	Scheduled snapshot label	snaplabel	Inherited	Appends a user-defined label to each scheduled snapshot and is blank by default.

Project Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
Snapshots - Snapshots	Name	snapshot name	--	Specifies the name of the snapshot.
	Creation	creation	--	Specifies the date and time when the snapshot is created.
	Unique	space_unique	--	Indicates the amount of unique space used by the snapshot.
	Total	space_data	--	Indicates the total amount of space referenced by the snapshot.
Snapshots - Schedules	Frequency	frequency	Create time	Indicates how often the snapshot is taken.
	Keep at most	keep	Create time	Controls the retention policy for snapshots.
Replication (Inherit from project)/Create New Actions	Target	target	Inherited	Identifies the replication target system.
	Pool	pool	Inherited	Specifies the storage pool on the target where the project will be replicated.
	Export data path	export_path	Inherited	Indicates the export data path.
	Limit bandwidth	max_bandwidth	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
	Enable SSL-encryption	use_ssl	Inherited	Controls whether to encrypt data on the wire using SSL.
	Disable compression	--	Inherited	Controls whether the compression is enabled or disabled.
	Include snapshot	include_snaps	Inherited	Controls whether replication updates include non-replication snapshots.
	Retain user snapshots on target	retain_user_snaps_on_target	Inherited	When set, keeps user-generated snapshots on the replication target. Continues to retain snapshots on the target until disabled.
	Include clone origin as data	include_clone_origin_as_data	Inherited	Controls the replication of each share that was

BUI Location	BUI Name	CLI Name	Property Type	Description
				cloned from a share that is external to the replication package on the target.
	Recovery point objective	recovery_point_objective	Inherited	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage.
	Replica lag warning alert	replica_lag_warning_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a minor alert is generated.
	Replica lag error alert	replica_lag_error_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a major alert is generated.
	Update frequency	continuous	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.

Filesystem Properties

Note - In the CLI, use the get command to see a list of all properties.

TABLE 114 Filesystems Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
Create Filesystem	Project	select project_name	--	Defines which project the filesystem uses to inherit parameter settings. You can also select the default project.
	Name	filesystem	--	Defines the name of the filesystem.
	Data migration source	shadow	Create time	Shows the location of the source if you are migrating data.
	User	root-user	Filesystem local	Specifies the owner of the root directory.

Filesystem Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
	Group	root_group	Filesystem local	Specifies the group of the root directory.
	Permissions or Use Windows default permissions	root_permission	Filesystem local	Specifies standard UNIX permissions for the root directory, or Windows default permissions.
	Inherit mountpoint	--	--	Indicates the mountpoint is inherited if selected.
	Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see “Mountpoint” on page 404 .
	Reject non UTF-8	utf8only	Create time	Enforces UTF-8 encoding for all filesystems and directories. For more information, see “Reject non UTF-8” on page 416 .
	Case sensitivity	casesensitivity	Create time	Controls whether directory lookups are case-sensitive, case-insensitive, or mixed. For more information, see “Case sensitivity” on page 415 .
	Normalization	normalization	Create time	Controls what unicode normalization, if any, is performed on filesystems and directories. For more information, see “Normalization” on page 416 .
	Encryption	encryption	Inherited	Defines the encryption type. For more information see, “Managing Encryption Keys” on page 652 .
	Inherit Key	--	--	Indicates the encryption key is inherited from the parent project if selected.
	Key	key	Inherited	Sets a specific LOCAL or OKM key and is used when the key is not inherited from the parent project.
	Keyname	keyname	Static	Identifies the key.

BUI Location	BUI Name	CLI Name	Property Type	Description
General - Space Usage - Data	Quota	quota	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.
	Quota Include snapshots	quota_snap	Space management	Sets a limit on the amount of space that can be consumed by any particular entity including the snapshots.
	Reservation	reservation	Space management	Represents a guarantee of space that can be consumed by any particular entity.
	Reservation Include snapshots	reservation_snap	Space management	Represents a guarantee of space that can be consumed by any particular entity including the snapshots.
General - Space Usage - Users & Groups	Users & Groups	--	--	Specifies the users and/or groups.
	Usage	--	--	Shows the amount of data used by the users and/or groups.
	Quota	quota	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.
General - Properties (Inherit from project)	Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see “Mountpoint” on page 404 .
	Read only	readonly	Inherited	Controls whether the filesystem contents are read only. For more information, see “Read only” on page 405 .
	Update access time on read	atime	Inherited	Controls whether the access time for files is updated on read. For more information, see “Update access time on read” on page 405 .
	Non-blocking mandatory locking	nbmand	inherited	Controls whether SMB locking semantics are enforced over POSIX semantics. For more information, see “Non-

BUI Location	BUI Name	CLI Name	Property Type	Description
				blocking mandatory locking on page 406.
	Data deduplication (warning)	dedup	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see “Data Deduplication” on page 406.
	Data compression	compression	Inherited	Controls whether data is compressed before being written to disk. For more information, see “Data Compression” on page 408.
	Checksum	checksum	Inherited	Controls the checksum used for data blocks. For more information, see “Checksum” on page 408.
	Cache device usage	secondarycache	Inherited	Controls whether cache devices are used for the share. For more information, see “Cache device usage” on page 409.
	Synchronous write bias	logbias	Inherited	Controls the behavior when servicing synchronous writes. For more information, see “Synchronous write bias” on page 409.
	Database record size	recordsize	Inherited	Specifies a suggested block size for files in the filesystem. For more information, see “Database record size” on page 410.
	Additional Replication	copies	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see “Additional replication” on page 411.
	Virus scan	vscan	Inherited	Controls whether a filesystem is scanned for viruses. For more information, see “Virus scan” on page 411.
	Prevent destruction	nodestroy	Inherited	Prevents shares or projects from

BUI Location	BUI Name	CLI Name	Property Type	Description
				being destroyed when set. For more information, see “Prevent destruction” on page 411.
	Restrict ownership change	rstchown	Inherited	Controls the ownership and can be turned off on a per-filesystem or per-project basis. For more information, see “Restrict ownership change” on page 411.
General - Custom Properties (Inherit from Project)	--	custom	--	Custom properties can be added as needed to attach user-defined tags to projects and shares.
Protocols - NFS	NFS	sharenfs	Inherited	NFS Protocol property settings and values are described in “NFS Protocol Properties” on page 443.
Protocols - NFS Exceptions	Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see “NFS Share Mode Exceptions” on page 444			
Protocols - SMB	SMB	sharesmb	Inherited	SMB Protocol property settings and values are described in “SMB Protocol Properties” on page 449.
Protocols - SMB Exceptions	Exceptions to the overall sharing modes may be defined for clients or collections of clients. See “SMB Protocol Share Mode Exceptions” on page 452.			
Protocols - Share Level ACL	Type	--	--	Indicates the type of the ACL.
	Target	--	--	Indicates the target fo the ACL.
	Access	--	--	Indicates whether the ACL access is allowed or denied.
	Permissions: Inheritance	--	--	Specifies standard UNIX permissions for the ACL.
Protocols - HTTP (Inherit from project)	Share mode	sharedav	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - FTP (Inherit from project)	Share mode	shareftp	Inherited	Determines whether the share is available for reading only, for reading

Filesystem Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
				and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - SFTP (Inherit from project)	Share mode	sharesftp	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Protocols - TFTP (Inherit from project)	Share mode	sharetftp	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, "on" is an alias for "rw."
Access - Root Directory Access	User	root_user	Filesystem local	Specifies the owner of the root directory.
	Group	root_group	Filesystem local	Specifies the group of the root directory.
	Permissions	root_permissions	Filesystem local	Specifies standard UNIX permissions for the root directory.
Access - ACL Behavior (Inherit from project)	ACL behavior on mode change	aclmode	Inherited	Controls how a mode change request interacts with the existing ACL.
	ACL inheritance behavior	aclinherit	Inherited	Controls how a new file or directory inherits existing ACL settings from the parent directory.
Access - Root Directory ACL	Type	--	--	Indicates the type of the ACL.
	Target	--	--	Indicates the target of the ACL.
	Access	--	--	Indicates whether the ACL access is allowed or denied.
	Permissions:Inheritance	--	--	Specifies standard UNIX permissions for the ACL.
Snapshots - Properties (Inherit from project)	.zfs/snapshot visibility	snapdir	Inherited	Controls whether filesystem snapshots can be accessed over data protocols at .zfs/snapshot in the root of the filesystem.
	Scheduled snapshot label	snaplabel	Inherited	Appends a user-defined label to each scheduled

BUI Location	BUI Name	CLI Name	Property Type	Description
				snapshot and is blank by default.
Snapshots - Snapshots	Name	--	--	Specifies the name of the snapshot.
	Creation	--	--	Specifies the date and time when the snapshot is created.
	Unique	--	--	Indicates the amount of unique space used by the snapshot.
	Total	--	--	Indicates the total amount of space referenced by the snapshot. This represents the size of the filesystem at the time the snapshot was taken, and any snapshot can theoretically take up an amount of space equal to the total size as data blocks are rewritten.
	Clones	--	--	Shows the number of clones of the snapshot.
Snapshots - Schedule	Frequency	frequency	Create time	Indicates how often the snapshot is taken.
	Keep at most	keep	Create time	Controls the retention policy for snapshots.
Replication (Inherit from project)/Create New Actions	Target	target	Inherited	Identifies the replication target system.
	Pool	pool	Inherited	Specifies the storage pool on the target where the project will be replicated.
	Export data path	export_path	Inherited	Indicates the export data path.
	Limit bandwidth	max_bandwidth	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
	Enable SSL-encryption	use_ssl	Inherited	Controls whether to encrypt data on the wire using SSL.
	Disable compression	compression	Inherited	Controls whether the compression is enabled or disabled.

Filesystem Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
	Include snapshot	include_snaps	Inherited	Controls whether replication updates include non-replication snapshots.
	Retain user snapshots on target	retain_user_snaps_on_target	Inherited	When set, keeps user-generated snapshots on the replication target. Continues to retain snapshots on the target until disabled.
	Include clone origin as data	include_clone_origin_as_data	Inherited	Controls the replication of each share that was cloned from a share that is external to the replication package on the target.
	Recovery point objective	recovery_point_objective	Inherited	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage.
	Replica lag warning alert	replica_lag_warning_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a minor alert is generated.
	Replica lag error alert	replica_lag_error_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a major alert is generated.
	Update frequency	continuous	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.
Usage	Referenced data	space_data	Read-only	Shows the total amount of space referenced by the active share, independent of any snapshots.
	Unused Reservation	space_unused_res	Read-only	Shows the amount of remaining space that is reserved for the filesystem.
	Snapshot data	space_snapshots	Read-only	Shows the total amount of data currently held by all snapshots of the share.
	Available data	space_available	Read-only	Shows any quotas on the share or project, or the

BUI Location	BUI Name	CLI Name	Property Type	Description
				absolute capacity of the pool.
	Total space	space_total	Read-only	Shows the sum of referenced data, snapshot data, and unused reservation.

LUN Properties

Note - In the CLI, use the `get` command to see a list of all properties.

TABLE 115 LUNs Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
Create LUN	Project	--	--	Defines which project the LUN uses to inherit parameter settings.
	Name	--	--	Defines the name of the LUN.
	Volume size	volsize	LUN local	Defines the maximum volume size and unit of measurement. For more information, see “Volume size” on page 412 .
	Thin provisioned	sparse	LUN local	Indicates only the amount of space physically consumed by data is used when selected. For more information, see “Thin provisioned” on page 412 .
	Volume block size	volblocksize	Create time	Native block size for the LUN; any power of 2 from 512 bytes to 1M, and the default is 8K.
	Online	status	LUN local	Indicates whether it is online or not.
	Target group	targetgroup	LUN local	Shows groups of targets used when exporting a LUN.
	Initiator group(s)	initiatorgroup	LUN local	Shows groups of initiators that can access the LUN.

LUN Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
	Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see “Mountpoint” on page 404 .
	LU number	lunumber	LUN local	Sets the logical unit number to zero or automatically assigns the number.
	Encryption	encryption	Inherited	Defines the encryption type.
	Inherit key	--	--	Indicates the encryption key is inherited from the parent project if selected. For more information see, “Managing Encryption Keys” on page 652 .
	Key	key	Inherited	Sets a specific LOCAL or OKM key and is used when the key is not inherited from the parent project.
	Keyname	keyname	Static	Identifies the key.
	GUID	lunguid	Read-only, LUN local	A globally unique, read-only identifier that identifies the SCSI device.
General - Space Usage - Data	Volume size	volsize	LUN local	Defines the maximum volume size and unit of measurement. For more information, see “Volume size” on page 412 .
	Thin provisioned	sparse	LUN local	Indicates only the amount of space physically consumed by data is used when selected. For more information, see “Thin provisioned” on page 412 .
General - Properties (Inherit from project)	Data deduplication (warning)	dedup	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see “Data Deduplication” on page 406 .
	Data compression	compression	Inherited	Controls whether data is compressed before being written to disk. For more

BUI Location	BUI Name	CLI Name	Property Type	Description
				information, see “Data Compression” on page 408.
	Checksum	checksum	Inherited	Controls the checksum used for data blocks. For more information, see “Checksum” on page 408.
	Additional replication	copies	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see “Additional replication” on page 411.
	Cache device usage	secondarycache	Inherited	Controls whether cache devices are used for the share. For more information, see “Cache device usage” on page 409.
	Synchronous write bias	logbias	Inherited	Controls the behavior when servicing synchronous writes. For more information, see “Synchronous write bias” on page 409.
	Prevent destruction	nodestroy	Inherited	Prevents shares or projects from being destroyed when set. For more information, see “Prevent destruction” on page 411.
Bandwidth Properties	Read limit	readlimit	--	Sets the maximum bytes per second that can be read from a LUN. M indicates megabytes and G indicates gigabytes. The default setting is unlimited, which provides no I/O throttling.
	Write limit	writelimit	--	Sets the maximum bytes per second that can be written to a LUN. M indicates megabytes and G indicates gigabytes. The default setting is unlimited, which

LUN Properties

BUI Location	BUI Name	CLI Name	Property Type	Description
				provides no I/O throttling.
	Effective read limit	effectivewritelimit	--	Read-only property that reports the lowest read limit for a LUN.
	Effective write limit	effectivewritelimit	--	Read-only property that reports the lowest write limit for a LUN.
Custom Properties (Inherit from Project)	Schema	schema	--	Can be added as needed to attach user-defined tags to projects and shares. For more information, see “Working with Schemas” on page 464 .
Protocols - Sharing Options	Online	status	LUN local	Indicates whether it is online or not.
	Target group	targetgroup	LUN local	Shows groups of targets used when exporting a LUN.
	Initiator group(s): LU number	initiatorgroup	LUN local	Shows groups of initiators that can access the LUN.
Protocols - Write Cache Behavior	Write cache enabled	writocache	LUN local	Controls whether the LUN caches writes.
Snapshots - Properties (Inherit from project)	Scheduled snapshot label	snaplabel	Inherited	Appends a user-defined label to each scheduled snapshot and is blank by default.
Snapshots - Snapshots	Name	--	--	Specifies the name of the snapshot.
	Creation	--	--	Specifies the date and time when the snapshot is created.
	Unique	--	--	Indicates the amount of unique space used by the snapshot.
	Total	--	--	Indicates the total amount of space referenced by the snapshot.
	Clones	--	--	Shows the number of clones of the snapshot.
Snapshots - Schedules	Frequency	frequency	Create time	Indicates how often the snapshot is taken.
	Keep at most	keep	Create time	Controls the retention policy for snapshots.

BUI Location	BUI Name	CLI Name	Property Type	Description
Replication (Inherit from project)/Create New Actions	Target	target	Inherited	Identifies the replication target system.
	Pool	pool	Inherited	Specifies the storage pool on the target where the project will be replicated.
	Export data path	export_path	Inherited	Indicates the export data path.
	Limit bandwidth	max_bandwidth	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
	Enable SSL-encryption	use_ssl	Inherited	Controls whether to encrypt data on the wire using SSL.
	Disable compression	compression	Inherited	Controls whether the compression is enabled or disabled.
	Include snapshot	include_snaps	Inherited	Controls whether replication updates include non-replication snapshots.
	Retain user snapshots on target	retain_user_snaps_on_target	Inherited	When set, keeps user-generated snapshots on the replication target.
	Update frequency	continuous	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.
Usage	Referenced data	space_data	Read-only	Shows the total amount of space referenced by the active share, independent of any snapshots.
	Snapshot data	space_snapshots	Read-only	Shows the total amount of data currently held by all snapshots of the share.
	Available data	space_available	Read-only	Shows any quotas on the share or project, or the absolute capacity of the pool.
	Total space	space_total	Read-only	Shows the sum of referenced data, snapshot data, and unused reservation.

Space Management for Shares

The appliance manages physical storage using a pooled storage model where all filesystems and LUNs share common space. Filesystems never have an explicit size assigned to them, and only take up as much space as they need. LUNs reserve enough physical space to write the entire contents of the device, unless they are thinly provisioned, in which case they behave like filesystems and use only the amount of space physically consumed by data.

This system provides maximum flexibility and simplicity of management in an environment when users are generally trusted to do the right thing. A stricter environment, where user's data usage is monitored and/or restricted, requires more careful management.

These topics define terminology and how to manage space usage, on a per-share or per-user basis, using quotas and reservations.

- [“Shares Terminology” on page 436](#)
- [“Managing Filesystem and Project Space” on page 437](#)
- [“Setting User or Group Quotas” on page 438](#)
- [“Working with Identity Management” on page 439](#)
- [“Share Usage Statistics” on page 441](#)

Shares Terminology

Before getting into details, it is important to understand some basic terms used when talking about space usage on the appliance.

- **Physical Data** - Size of data as stored physically on disk. Typically, this is equivalent to the logical size of the corresponding data, but can be different in the phase of compression or other factors. This includes the space of the active share as well as all snapshots. Space accounting is generally enforced and managed based on physical space.
- **Logical Data** - The amount of space logically consumed by a filesystem. This does not factor into compression, and can be viewed as the theoretical upper bound on the amount of space consumed by the filesystem. Copying the filesystem to another appliance using a different compression algorithm will not consume more than this amount. This statistic is not explicitly exported and can generally only be computed by taking the amount of physical space consumed and multiplying by the current compression ratio.
- **Referenced Data** - This represents the total amount of space referenced by the active share, independent of any snapshots. This is the amount of space that the share would consume should all snapshots be destroyed. This is also the amount of data that is directly manageable by the user over the data protocols.

- **Snapshot Data** - This represents the total amount of data currently held by all snapshots of the share. This is the amount of space that would be free should all snapshots be destroyed.
- **Quota** - A quota represents a limit on the amount of space that can be consumed by any particular entity. This can be based on filesystem, project, user, or group, and is independent of any current space usage.
- **Reservation** - A reservation represents a guarantee of space for a particular project or filesystem. This takes available space away from the rest of the pool without increasing the actual space consumed by the filesystem. This setting cannot be applied to users and groups. The traditional notion of a statically sized filesystem can be created by setting a quota and reservation to the same value.

Managing Filesystem and Project Space

The simplest way of enforcing quotas and reservations is on a per-project or per-filesystem basis. Quotas and reservations do not apply to LUNs, though their usage is accounted for in the total project quota or reservations.

Data Quotas - A data quota enforces a limit on the amount of space a filesystem or project can use. By default, it will include the data in the filesystem and all snapshots. Clients attempting to write new data will get an error when the filesystem is full, either because of a quota or because the storage pool is out of space. As described in [“Snapshot Space Management” on page 482](#), this behavior may not be intuitive in all situations, particularly when snapshots are present. Removing a file may cause the filesystem to write new data if the data blocks are referenced by a snapshot, so it may be the case that the only way to decrease space usage is to destroy existing snapshots.

If the 'include snapshots' property is unset, then the quota applies only to the immediate data referenced by the filesystem, not any snapshots. The space used by snapshots is enforced by the project-level quota but is otherwise not enforced. In this situation, removing a file referenced by a snapshot will cause the filesystem's referenced data to decrease, even though the system as a whole is using more space. If the storage pool is full (as opposed to the filesystem reaching a preset quota), then the only way to free up space may be to destroy snapshots.

Data quotas are strictly enforced, which means that as space usage nears the limit, the amount of data that can be written must be throttled as the precise amount of data to be written is not known until after writes have been acknowledged. This can affect performance when operating at or near the quota. Because of this, it is generally advisable to remain below the quota during normal operating procedures.

Quotas are managed through the BUI under Shares > General > Space Usage > Data. They are managed in the CLI as the `quota` and `quota_snap` properties.

Data Reservations - A data reservation is used to make sure that a filesystem or project has at least a certain amount of available space, even if other shares in the system try to use more space. This unused reservation is considered part of the filesystem, so if the rest of the pool (or project) reaches capacity, the filesystem can still write new data even though other shares may be out of space.

By default, a reservation includes all snapshots of a filesystem. If the 'include snapshots' property is unset, then the reservation only applies to the immediate data of the filesystem. The behavior when taking snapshots may not always be intuitive. If a reservation on filesystem data (but not snapshots) is in effect, then whenever a snapshot is taken, the system must reserve enough space for that snapshot to diverge completely, even if that never occurs. For example, if a 50G filesystem has a 100G reservation without snapshots, then taking the first snapshot will reserve an additional 50G of space, and the filesystem will end up reserving 150G of space total. If there is insufficient space to guarantee complete divergence of data, then taking the snapshot will fail.

Reservations are managed through the BUI under Shares > General > Space Usage > Data. They are managed in the CLI as the `reservation` and `reservation_snap` properties.

Space Management for Replicating LUNs - When you create a LUN the full physical space you configure for the LUN is reserved and cannot be used by other file systems (unless it is thinly provisioned). For replication, when you take a snapshot of a LUN of any given size, up to twice the size of the LUN is also reserved, depending on how much of the LUN space has been used.

The following list shows the maximum overhead space required when replicating a LUN:

- Up to 100% on the source between updates
- Up to 200% on the source during an update
- Up to 200% on the target

Setting User or Group Quotas

Quotas can be set on a user or group at the filesystem level, as well as the project level. These enforce physical data usage based on the POSIX or Windows identity of the owner or group of the file or directory. There are some significant differences between user and group quotas and filesystem and project data quotas:

- User and group quotas can be applied to filesystems and projects.
- Default quotas can be set at the project level and inherited by the project's filesystems.
- Default quotas set at the project level can be changed at the filesystem level.
- Default quotas can be retrieved or modified over the SMB protocol.

- User and group quotas are implemented using *delayed enforcement*. This means that users will be able to exceed their quota for a short period of time before data is written to disk. Once the data has been pushed to disk, the user will receive an error on new writes, just as with the filesystem-level quota case.
- User and group quotas are always enforced against referenced data. This means that snapshots do not affect any quotas, and a clone of a snapshot will consume the same amount of effective quota, even though the underlying blocks are shared.
- User and group reservations are not supported.
- User and group quotas, unlike data quotas, are stored with the regular filesystem data. This means that if the filesystem is out of space, you will not be able to make changes to user and group quotas. You must first make additional space available before modifying user and group quotas.
- User and group quotas are sent as part of any remote replication. It is up to the administrator to ensure that the name service environments are identical on the source and destination.
- NDMP backup and restore of an entire share will include any user or group quotas. Restores into an existing share will not affect any current quotas.

Working with Identity Management

User and group quotas leverage the identity mapping service on the appliance. This allows users and groups to be specified as either UNIX or Windows identities, depending on the environment. Like file ownership, these identities are tracked in the following ways:

- If there is no UNIX mapping, a reference to the windows ID is stored.
- If there is a UNIX mapping, then the UNIX ID is stored.

This means that the canonical form of the identity is the UNIX ID. If the mapping is changed later, the new mapping will be enforced based on the new UNIX ID. If a file is created by a Windows user when no mapping exists, and a mapping is later created, new files will be treated as a different owner for the purposes of access control and usage format. This also implies that if a user ID is reused (i.e. a new user name association created), then any existing files or quotas will appear to be owned by the new user name.

It is recommended that any identity mapping rules be established before attempting to actively use filesystems. Otherwise, any change in mapping can sometimes have surprising results.

Working with Filesystem Namespace

Every filesystem on the appliance must be given a unique mountpoint which serves as the access point for the filesystem data. Projects can be given mountpoints, but these serve only as

a tool to manage the namespace using inherited properties. Projects are never mounted, and do not export data over any protocol.

All shares must be mounted under `/export`. While it is possible to create a filesystem mounted at `/export`, it is not required. If such a share doesn't exist, any directories will be created dynamically as necessary underneath this portion of the hierarchy. Each mountpoint must be unique within a cluster.

- **Namespace Nested Mountpoints** - It is possible to create filesystems with mountpoints beneath that of other filesystems. In this scenario, the parent filesystems are mounted before children and vice versa. The following cases should be considered when using nested mountpoints:
 - If the mountpoint doesn't exist, one will be created, owned by root and mode 0755. This mountpoint may or may not be torn down when the filesystem is renamed, destroyed, or moved, depending on circumstances. To be safe, mountpoints should be created within the parent share before creating the child filesystem.
 - If the parent directory is read-only, and the mountpoint doesn't exist, the filesystem mount will fail. This can happen synchronously when creating a filesystem, but can also happen asynchronously when making a large-scale change, such as renaming filesystems with inherited mountpoints.
 - When renaming a filesystem or changing its mountpoint, all children beneath the current mountpoint as well as the new mountpoint (if different) will be unmounted and remounted after applying the change. This will interrupt any data services currently accessing the share.
 - Support for automatically traversing nested mountpoints depends on protocol, as outlined below.
- **Namespace NFSv2 / NFSv3 / NFSv4.0 / NFSv4.1** - Under NFS, each filesystem is a unique export made visible via the MOUNT protocol. NFSv2 and NFSv3 have no way to traverse nested filesystems, and each filesystem must be accessed by its full path. While nested mountpoints are still functional, attempts to cross a nested mountpoint will result in an empty directory on the client. While this can be mitigated through the use of automount mounts, transparent support of nested mountpoints in a dynamic environment requires NFSv4.0 or NFSv4.1.

NFSv4.0 and NFSv4.1 have several improvements over NFSv3 when dealing with mountpoints. First is that parent directories can be mounted, even if there is no share available at that point in the hierarchy. For example, if `/export/home` was shared, it is possible to mount `/export` on the client and traverse into the actual exports transparently. More significantly, some NFSv4.0 and NFSv4.1 clients (including Linux) support automatic client-side mounts, sometimes referred to as "mirror mounts". With such a client, when a user traverses a mountpoint, the child filesystem is automatically mounted at the appropriate local mountpoint, and torn down when the filesystem is unmounted on the client. From the

server's perspective, these are separate mount requests, but they are stitched together onto the client to form a seamless filesystem namespace.

- **Namespace SMB** - The SMB protocol does not use mountpoints, as each share is made available by resource name. However, each filesystem must still have a unique mountpoint. Nested mountpoints (multiple filesystems within one resource) are not currently supported, and any attempt to traverse a mountpoint will result in an empty directory.
- **Namespace FTP / FTPS / SFTP** - Filesystems are exported using their standard mountpoint. Nested mountpoints are fully supported and are transparent to the user. However, it is not possible to not share a nested filesystem when its parent is shared. If a parent mountpoint is shared, then all children will be shared as well.
- **Namespace HTTP / HTTPS** - Filesystems are exported under the /shares directory, so a filesystem at /export/home will appear at /shares/export/home over HTTP/HTTPS. Nested mountpoints are fully supported and are transparent to the user. The same behavior regarding conflicting share options described in the FTP protocol section also applies to HTTP.

Share Usage Statistics

On the left side of the view (beneath the project panel when collapsed) is a table explaining the current space usage statistics. These statistics are either for a particular share (when editing a share) or for the pool as a whole (when looking at the list of shares). If any properties are zero, then they are excluded from the table.

Some of the usage statistics are also displayed in the CLI context `shares show`.

The following table describes the BUI and CLI usage properties.

BUI Name	CLI Name	Description
Available space	--	This statistic is implicitly shown as the capacity in terms of capacity percentage in the title. The available space reflects any quotas on the share or project, or the absolute capacity of the pool. The number shown here is the sum of the total space used and the amount of available space.
Referenced data	usage_data	The amount of data referenced by the data. This includes all filesystem data or LUN blocks, in addition to requisite metadata. With compression, this value may be much less than the logical size of the data contained within the share. If the share is a clone of a snapshot, this value may be less than the physical storage it could theoretically include, and may be zero.
Snapshot data	usage_snapshots	The amount of space used by all snapshots of the share, including any project snapshots. This size is not equal to

BUI Name	CLI Name	Description
		the sum of unique space consumed by all snapshots. Blocks that are referenced by multiple snapshots are not included in the per-snapshot usage statistics, but will show up in the share's snapshot data total.
Unused reservation	--	If a filesystem has a reservation set, this value indicates the amount of remaining space that is reserved for the filesystem. This value is not set for LUNs. The appliance prevents other shares from consuming this space, guaranteeing the filesystem enough space. If the reservation does not include snapshots, then there must be enough space when taking a snapshot for the entire snapshot to be overwritten. For more information on reservations, see “Filesystem Properties” on page 423 .
Total space	usage_total	The sum of referenced data, snapshot data, and unused reservation.

Share and Project Protocols

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project.

For iSCSI, initiators can discover the target through one of the mechanisms described in [“Configuring Storage Area Network \(SAN\)” on page 170](#).

For information about supported protocol properties, see the following sections:

- [“NFS Protocol” on page 442](#)
- [“SMB Protocol” on page 449](#)
- [“HTTP Protocol” on page 455](#)
- [“FTP Protocol” on page 456](#)
- [“SFTP Protocol” on page 456](#)
- [“TFTP Protocol” on page 456](#)

Related Topics

- [“NFS Configuration” on page 328](#)
- [“SMB Configuration” on page 346](#)

NFS Protocol

This section contains the following topics:

- “NFS Protocol Properties” on page 443
- “NFS Share Mode Exceptions” on page 444
- “NFS Protocol Character Set Encodings” on page 447
- “NFS Protocol Security Modes” on page 448

For more information about the NFS protocol, use these topics:

- “NFS Configuration” on page 328
- “Filesystem Properties” on page 423
- “Project Properties” on page 417
- NFSv2 and NFSv3 Security (RFC 2623) (<http://www.ietf.org/rfc/rfc2623.txt>)
- NFSv4 Protocol (RFC 7530) (<http://www.ietf.org/rfc/rfc7530.txt>)
- NFSv4.1 Protocol (RFC 5661) (<https://tools.ietf.org/html/rfc5661>)

For information about other supported protocols, see the following sections:

- “SMB Protocol” on page 449
- “HTTP Protocol” on page 455
- “FTP Protocol” on page 456
- “SFTP Protocol” on page 456
- “TFTP Protocol” on page 456

NFS Protocol Properties

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. The following table shows NFS protocol properties and possible values.

TABLE 116 NFS Protocol Properties

Property	CLI Value(s)	Property Type	Description
Share mode	off/rw/ro	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. See “Share and Project Protocols” on page 442.
Disable setuid/setgid file creation	nosuid	Inherited	If selected, clients will not be able to create files with the setuid (S_ISUID) and setgid (S_ISGID) bits set, nor enable these bits on existing files via the chmod(2) system call.
Prevent clients from mounting subdirectories	nosub	Inherited	If selected, clients will be prevented from directly mounting subdirectories. They will be forced to mount the root of the share. Note: This

Property	CLI Value(s)	Property Type	Description
			only applies to the NFSv2 and NFSv3 protocols, not to NFSv4.0 or NFSv4.1.
Anonymous user mapping	anon	Inherited	Unless the "root" option is in effect for a particular client, the root user on that client is treated as an unknown user, and all attempts by that user to access the sharer's files will be treated as attempts by a user with this uid. This file's access bits and ACLs will then be evaluated normally.
Character set	See Character Set Encodings for possible values.	Inherited	Sets the character set default for all clients.
Security mode	sec= See Security Modes for list of possible values.	Inherited	Sets the security mode for all clients.
Enforce reserved ports for system authentication	resvport	Inherited	When set on a share or project in conjunction with the system authentication security mode, requires NFS clients to use low-numbered ("reserved") TCP ports. Some NFS clients, such as Oracle Solaris and Linux, use low-numbered TCP ports by default. Other clients, such as Windows, may require configuration.

NFS Share Mode Exceptions

Exceptions to the global sharing mode may be defined for clients or collections of clients by setting client-specific share modes or *exceptions*. To restrict access to certain clients, set the global sharing mode to none and increasingly grant access to smaller and smaller groups. For example, you could create a share with the global sharing mode set to none, which denies access to all clients, and then grant read-only access to a subset of the clients. Further, you could grant read/write access to an even smaller subset of the clients and, finally, only trusted hosts might have read/write and root-enabled access.

Client-specific share modes take precedence over the global share mode. A client is granted access according to the client-specific share mode that is specified in an exception. In the absence of exceptions, the client is granted access according to the global share mode.

TABLE 117 Client Types

Type	CLI Prefix	Description	Example
Host(FQDN) or Netgroup	none	A single client whose IP address resolves to the specified fully qualified name, or a	caji.sf. example. com

Type	CLI Prefix	Description	Example
		netgroup containing fully qualified names to which a client's IP address resolves.	
DNS Domain	.	All clients whose IP addresses resolve to a fully qualified name ending in this suffix.	sf.example.com
IPv4 Subnet	@	All clients whose IP addresses are within the specified IPv4 subnet, expressed in CIDR notation.	192.0.2.254/22
IPv6 Subnet	@	All clients whose IP addresses are within the specified IPv6 subnet, expressed in CIDR notation.	2001:db8:410:d43::/64

For each client or collection of clients, you specify whether the client has read-only or read-write access to the share. If you are setting an NFS exception, you also specify whether the client has root user privileges or is treated as a user without root access.

Managing Netgroups

Netgroups can be used to control access for NFS exports. However, managing netgroups can be complex. Consider using IP subnet rules or DNS domain rules instead.

If netgroups are used, they will be resolved from NIS or LDAP, depending on which service is enabled. If LDAP is used, each netgroup must be located at the default location, `ou=Netgroup`, (Base DN), and must use the standard schema.

The username component of a netgroup entry typically has no effect on NFS; only the hostname is significant. Hostnames contained in netgroups must be canonical and, if resolved using DNS, fully qualified. That is, the NFS subsystem will attempt to verify that the IP address of the requesting client resolves to a canonical hostname that matches either the specified FQDN, or one of the members of one of the specified netgroups. This match must be exact, including any domain components; otherwise, the exception will not match and the next exception will be tried. For more information on hostname resolution, see [DNS](#).

As of the 2013.1.0 software release, UNIX client users may belong to a maximum of 1024 groups without any performance degradation. Prior releases supported up to 16 groups per UNIX client user.

NFS Share Modes and Exception Options

In the CLI, all NFS share modes and exceptions are specified using a single options string for the `share nfs` property. This string is a comma-separated list of values. It should begin with one of `ro`, `rw`, `on`, or `off`, as an analogue to the global share modes described for the BUI.

TABLE 118 NFS Share Mode Values (BUI and CLI)

BUI Share Mode Value	CLI Share Mode Value	Description	Example
None	off	Share mode is disabled.	sharenfs=off
	on	The share name is the dataset name and is available for reading and writing or reading only if the rw or ro NFS exceptions are defined. For all other clients, share mode is disabled.	sharenfs="on, ro=sf.example.com"
	<resource name>	The share name is the resource name and is available for reading and writing or reading only if the rw or ro NFS exceptions are defined. For all other clients, share mode is disabled.	sharenfs="myshare, ro=sf.example.com"
Read/write	on	The share name is the dataset name and is available for reading and writing for all clients if there are no NFS exceptions.	sharenfs=on
	rw	The share name is the dataset name and is available for reading and writing for all clients except those for which the ro exception is defined.	sharenfs=rw or sharenfs="rw, ro=sf.example.com"
	<resource name>	The share name is the resource name and is available for reading and writing for all clients if there are no NFS exceptions.	sharenfs=myshare
	<resource name>,rw	The share name is the resource name, is available for reading and writing for all clients except those for which the ro exception is defined. NFS exceptions may or may not be defined.	sharenfs="myshare, rw" or sharenfs="myshare, rw, ro=sf.example.com"
Read only	ro	The share name is the dataset name and is available for reading only for all hosts except those for which the rw exception is defined.	sharenfs="ro, rw=sf.example.com"
	<resource name>,ro	The share name is the resource name, is available for reading only for all clients except those for which the rw exception is defined. NFS exceptions may or may not be defined.	sharenfs="myshare, ro" or sharenfs="myshare, ro, rw=sf.example.com"

The following example sets the share mode for all clients to read-only. The root users on all clients will access the files on the share as if they were the generic "nobody" user.

```
set sharenfs=ro
```

Either or both of the nosuid and anon options can also be appended. Therefore, to define the mapping of all unknown users to the uid 153762, you might specify the following:

```
set sharenfs="ro,anon=153762"
```

Note - CLI property values that contain the "=" character must be quoted.

Additional NFS exceptions can be specified by appending text of the form "option=collection", where "option" is one of `ro`, `rw`, or `root`, defining the type of access to be granted to the client collection. The collection is specified by the prefix character from Client Types table and either a DNS hostname/domain name or CIDR network number. For example, to grant read-write access to all hosts in the `sf.example.com` domain and root access to those in the `192.168.44.0/24` network, you might use:

```
set sharenfs="ro,anon=153762,rw=.sf.example.com,root=@192.168.44.0/24"
```

Note - This example only applies to NFS exceptions.

Netgroup names can be used anywhere an individual fully qualified hostname can be used. For example, you can permit read-write access to the "engineering" netgroup as follows:

```
set sharenfs="ro,rw=engineering"
```

NFS Protocol Character Set Encodings

Normally, the character set encoding used for filename is unspecified. The NFSv3 and NFSv2 protocols do not specify the character set. NFSv4.0 and NFSv4.1 are supposed to use UTF-8, but not all clients do and this restriction is not enforced by the server. If the UTF-8 only option is disabled for a share, these filenames are written verbatim to the filesystem without any knowledge of their encoding. This means that they can only be interpreted by clients using the same encoding. SMB, however, requires filenames to be stored as UTF-8 so that they can be interpreted on the server side. This makes it impossible to support arbitrary client encodings while still permitting access over SMB.

In order to support such configurations, the character set encoding can be set share-wide or on a per-client basis. The following character set encodings are supported:

cp932	euc-tw	iso8859-7	koi8-r
euc-cn	iso8859-1	iso8859-8	shift_jis
euc-jp	iso8859-2	iso8859-9	
euc-jpms	iso8859-5	iso8859-13	
euc-kr	iso8859-6	iso8859-15	

The default behavior is to leave the character set encoding unspecified (pass-through). The BUI allows the character set to be chosen through the standard exception list mechanism. In the CLI, each character set itself becomes an option with one or more hosts, with '*' indicating the share-wide setting. For example, the following:

```
hostname:shares default> set sharenfs="rw,euc-kr=*
```

Will share the filesystem with 'euc-kr' as the default encoding. The following:

```
hostname:shares default> set sharenfs="rw,euc-kr=host1.domain.com,euc-jp=host2.domain.com"
```

Use the default encoding for all clients except 'host1' and 'host2', which will use 'euc-kr' and 'euc-jp', respectively. The format of the host lists follows that of other CLI NFS options.

Note that some NFS clients do not correctly support alternate locales; consult your NFS client documentation for details.

NFS Protocol Security Modes

Security modes are set on a per-share basis. The following list describes Kerberos security settings:

- **krb** - End-user authentication through Kerberos V5
- **krb5i** - krb5 plus integrity protection (data packets are tamper proof)
- **krb5p** - krb5i plus privacy protection (data packets are tamper proof and encrypted)

Security modes are specified by appending text in the form "*option=mode*" where *option* is *sec* and *mode* is the security setting. For example:

```
hostname: shares default> set sharenfs="sec=krb5"
```

Note - CLI property values that contain the "=" character must be quoted.

Combinations of Kerberos types can be specified in the security mode setting. The combination security modes let clients mount with any Kerberos type listed, as shown in the following table.

TABLE 119 Combinations of Kerberos types

Setting	Description
sys	System Authentication
krb5	Kerberos v5 only - Clients must mount using this flavor.
krb5:krb5i	Kerberos v5, with integrity - Clients may mount using any flavor listed.
krb5i	Kerberos v5 integrity only - Clients must mount using this flavor.
krb5:krb5i:krb5p	Kerberos v5, with integrity or privacy - Clients may mount using any flavor listed.

Setting	Description
krb5p	Kerberos v5 privacy only - Clients may mount using this flavor.

Reserved Ports

To set reserved ports for system authentication, use `resvport` as shown in this example:

```
set sharenfs="sec=sys,rw,resvport"
```

Note that `resvport` can only be used with the system authentication security mode `sec=sys`.

SMB Protocol

This section contains the following topics:

- [“SMB Protocol Properties” on page 449](#)
- [“Client-side Caching Property” on page 450](#)
- [“Opportunistic Locks Property” on page 451](#)
- [“SMB Protocol Share Mode Exceptions” on page 452](#)
- [“Share-Level ACLs” on page 455](#)

For more information about the SMB protocol, use these topics:

- [“SMB Configuration” on page 346](#)
- [“Filesystem Properties” on page 423](#)
- [“Project Properties” on page 417](#)

For information about other supported protocols, see the following sections:

- [“NFS Protocol” on page 442](#)
- [“HTTP Protocol” on page 455](#)
- [“FTP Protocol” on page 456](#)
- [“SFTP Protocol” on page 456](#)
- [“TFTP Protocol” on page 456](#)

SMB Protocol Properties

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. The following table shows SMB protocol properties and possible values.

TABLE 120 SMB Protocol Properties

Property	CLI Value(s)	Property Type	Description
Share mode	off/on/rw/ro	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. See “SMB Protocol Share Mode Exceptions” on page 452 .
Resource name	resource_name/off/on	Inherited	Shows the name by which SMB clients refer to this share. The resource name off indicates no SMB client may access the share, and the resource name on indicates the share will be exported with the filesystem's name.
Enable access-based enumeration	abe	Inherited	Performs access-based enumeration when enabled.
Enable guest access	guestok	Inherited	Grants guest access when enabled. This property is disabled by default.
Is a DFS namespace	dfsroot	Inherited	Indicates whether this share is provisioned as a standalone DFS namespace.
Client-side caching policy	csc	Inherited	Indicates per-share configuration options provided to support client-side caching. For more information, see “Client-side Caching Property” on page 450 .
Opportunistic locks policy	oplocks	Inherited	Enables or disables opportunistic locks at the share level. For more information, see “Opportunistic Locks Property” on page 451 .
Enable continuous availability	cont_avail	Inherited	When enabled, SMB3 clients can request persistent file handles for the share. This allows the appliance to store the state associated with a persistent file handle in stable, persistent storage. The state can be transparently restored in the event of a controller failure, such as a takeover and failback operation on clustered controllers. Continuously available SMB shares are not allowed to be shared over NFS or used on workloads such as Home Directory that have a very high number of opens/closes. Continuously available SMB shares are only recommended for enterprise applications that have limited number of opens/closes.

Client-side Caching Property

The client-side caching property (csc) controls whether files and programs from the share are cached on the local client for offline use when disconnected from the appliance.

BUI value	CLI value	Description
No caching	none	Disables client-side caching for the share. No files or programs from the share are available offline. This option blocks Offline Files on the client computers from making copies of the files and programs on the shared folder.
Manual caching	manual	Only specified files and programs are cached on the local client and available offline. This is the default option when you set up a shared folder. By using this option, no files or programs are available offline by default. You can control which files and programs to access when you are not connected to the network.
Automatic document caching	documents	All files accessed from the share are cached on the local client and available offline. Files are automatically reintegrated when the local client is online again. Programs accessed from the share are not available offline unless previously cached locally.
Automatic program caching	programs	All programs accessed from the share are cached on the local client and available offline. When online, the programs are run from the local client. Additionally, all files accessed from the share are cached on the local client and available offline. Files are automatically reintegrated when the local client is online again.

Opportunistic Locks Property

Opportunistic locks are a client-caching mechanism that facilitates local caching to reduce network traffic and improve performance. The property (oplocks) controls whether the server grants or denies opportunistic locks at the share level, and applies to both lease (SMB 2.1 and above) and legacy (SMB 2.0 and below) opportunistic locks.

The client requests an opportunistic lock on a file within a share, and that request is either granted or denied depending on the server configuration and the current state of the file. If the client attempts to access a file in a manner inconsistent with the opportunistic locks that have already been granted for that file, a conflict occurs. In such cases, the server initiates a process to break the existing opportunistic locks before proceeding with the conflicting operation.

Enabling opportunistic locks improves performance when files within a share are accessed by a single client. In some scenarios, however, such as when the same file is accessed simultaneously by multiple clients, it can introduce unnecessary overhead. Opportunistic locks can thus be enabled or disabled per share, instead of globally controlled, based on the expected pattern of workloads.

If an opportunistic locks property is not defined at the share level, the default is the global opportunistic locks property set at the service level. For more information, see "Enable oplocks" in section ["SMB Service Properties" on page 347](#).

BUI Value	CLI Value	Description	Example
Enabled	enabled	Enables opportunistic locks for a share	set sharesmb="myshare, oplocks=enabled, abe=off, dfsroot=false"
Disabled	disabled	Disables opportunistic locks for a share	set sharesmb="myshare, oplocks=disabled, abe=off, dfsroot=false"
<empty>	--	The opportunistic locks property is neither enabled or disabled. Uses the global opportunistic locks property when the share-level property is not set.	set sharesmb="myshare, abe=off, dfsroot=false"

SMB Protocol Share Mode Exceptions

Exceptions to the global sharing mode may be defined for clients or collections of clients by setting client-specific share modes or *exceptions*. To restrict access to certain clients, set the global sharing mode to none and increasingly grant access to smaller and smaller groups. For example, you could create a share with the global sharing mode set to none, which denies access to all clients, and then grant read-only access to a subset of the clients. Further, you could grant read/write access to an even smaller subset of the clients and, finally, only trusted hosts might have read/write access.

TABLE 121 Client Types

Type	CLI Prefix	Description	Example
Host(FQDN) or Netgroup	none	A single client whose IP address resolves to the specified fully qualified name, or a netgroup containing fully qualified names to which a client's IP address resolves.	caji.sf.example.com

Type	CLI Prefix	Description	Example
DNS Domain	.	All clients whose IP addresses resolve to a fully qualified name ending in this suffix.	sf.example.com
IPv4 Subnet	@	All clients whose IP addresses are within the specified IPv4 subnet, expressed in CIDR notation.	192.0.2.254/22
IPv6 Subnet	@	All clients whose IP addresses are within the specified IPv6 subnet, expressed in CIDR notation.	2001:db8:410:d43::/64

For each client or collection of clients, you specify whether the client has read-only or read-write access to the share.

Managing netgroups - Netgroups can be used to control access for SMB exports. However, managing netgroups can be complex. Consider using IP subnet rules or DNS domain rules instead.

If netgroups are used, they will be resolved from NIS or LDAP, depending on which service is enabled. If LDAP is used, each netgroup must be located at the default location, `ou=Netgroup`, (Base DN), and must use the standard schema.

The username component of a netgroup entry typically has no effect on SMB; only the hostname is significant. Hostnames contained in netgroups must be canonical and, if resolved using DNS, fully qualified. That is, the SMB subsystem will attempt to verify that the IP address of the requesting client resolves to a canonical hostname that matches either the specified FQDN, or one of the members of one of the specified netgroups. This match must be exact, including any domain components; otherwise, the exception will not match and the next exception will be tried. For more information on hostname resolution, see [DNS](#).

As of the 2013.1.0 software release, UNIX client users may belong to a maximum of 1024 groups without any performance degradation. Prior releases supported up to 16 groups per UNIX client user.

SMB Share Modes and Exception Options

In the CLI, all SMB share modes and exceptions are specified using a single options string for the `sharesmb` property. This string is a comma-separated list of values. It should begin with one of `ro`, `rw`, `on`, or `off`, as an analogue to the global share modes described for the BUI.

TABLE 122 SMB Share Mode Values (BUI and CLI)

BUI Share Mode Value	CLI Share Mode Value	Description	Example
None	off	Share mode is disabled.	sharesmb=off
	on	The share name is the dataset name and is available for reading and writing or reading only if the rw or ro SMB exceptions are defined. For all other clients, share mode is disabled.	sharesmb="on, ro=sf.example.com"
	<resource name>	The share name is the resource name and is available for reading and writing or reading only if the rw or ro SMB exceptions are defined. For all other clients, share mode is disabled.	sharesmb="myshare, ro=sf.example.com"
Read/write	on	The share name is the dataset name and is available for reading and writing for all clients if there are no SMB exceptions.	sharesmb=on
	rw	The share name is the dataset name and is available for reading and writing for all clients except those for which the ro exception is defined.	sharesmb=rw or sharesmb="rw, ro=sf.example.com"
	<resource name>	The share name is the resource name and is available for reading and writing for all clients if there are no SMB exceptions.	sharesmb=myshare
	<resource name>,rw	The share name is the resource name, is available for reading and writing for all clients except those for which the ro exception is defined. SMB exceptions may or may not be defined.	sharesmb="myshare, rw" or sharesmb="myshare, rw, ro=sf.example.com"
Read only	ro	The share name is the dataset name and is available for reading only for all hosts except those for which the rw exception is defined.	sharesmb="ro, rw=sf.example.com"
	<resource name>,ro	The share name is the resource name, is available for reading only for all clients except those for which the rw exception is defined. SMB exceptions may or may not be defined.	sharesmb="myshare, ro" or sharesmb="myshare, ro, rw=sf.example.com"

The following example sets the share mode for all clients to read-only.

```
set sharesmb=ro
```

Additional SMB exceptions can be specified by appending text of the form "option=collection", where "option" is either ro or rw. You cannot grant root access with SMB exceptions. The collection is specified by the prefix character from table 114, and either a DNS hostname/domain name or CIDR network number.

For example, to grant read-write access to all hosts in the sf.example.com domain:

```
set sharesmb="ro, rw=.sf.example.com"
```

This example grants read-only access to clients with the IP addresses 2001:db8:410:d43::/64 and 192.0.2.254/22:

```
set sharesmb="on,ro=@[2001:db8:410:d43::/64]:@192.0.2.254/22"
```

Netgroup names can be used anywhere an individual fully qualified hostname can be used. For example, you can permit read-write access to the "engineering" netgroup as follows:

```
set sharesmb="ro,rw=engineering"
```

Share-Level ACLs

A share-level access control list (ACL), when combined with the ACL of a file or directory in the share, determines the effective permissions for that file. By default, this ACL grants everyone full control. This ACL provides another layer of access control above the ACLs on files and allows for more sophisticated access control configurations. This property may only be set once the filesystem has been exported by configuring the SMB resource name. If the filesystem is not exported over the SMB protocol, setting the share-level ACL has no effect.

When access-based enumeration is enabled, clients may see directory entries for files which they cannot open. Directory entries are filtered only when the client has no access to that file. For example, if a client attempts to open a file for read/write access but the ACL grants only read access, that open request will fail but that file will still be included in the list of entries.

For more information about ACLs, see [“Access Control Lists for Filesystems” on page 456](#).

HTTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the HTTP protocol (shar dav) and for an Object Store, users can set the share mode to determine if the filesystem is available for read only (ro), read and write (rw or on), or neither (off).

Related Topics

- [“Project Properties” on page 417](#)
- [“Filesystem Properties” on page 423](#)

FTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the FTP protocol (`shareftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- [“Project Properties” on page 417](#)
- [“Filesystem Properties” on page 423](#)

SFTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the SFTP protocol (`sharesftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- [“Project Properties” on page 417](#)
- [“Filesystem Properties” on page 423](#)

TFTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the TFTP protocol (`sharetftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- [“Project Properties” on page 417](#)
- [“Filesystem Properties” on page 423](#)

Access Control Lists for Filesystems

You can set options to control ACL behavior as well as control access to the root directory of the filesystem.

Note - ACLs are available only for filesystems.

For more information about ACLs, see the following topics:

- [“Root Directory Access” on page 457](#)
- [“ACL Behavior on Mode Change” on page 458](#)
- [“ACL Inheritance Behavior” on page 459](#)
- [“Root Directory ACL” on page 460](#)

Root Directory Access

To set basic access control for the root directory of the filesystem, go to `Shares > Shares > filesystem > Access`. These settings can be managed in-band via whatever protocols are being used, but they can also be specified here for convenience. These properties cannot be changed on a read-only filesystem, as they require changing metadata for the root directory of the filesystem.

- **User** - The owner of the root directory. This can be specified as a user ID or user name. For more information on mapping UNIX and Windows users, see [Identity Mapping](#). For UNIX-based NFS access, this can be changed from the client using the `chown` command.
- **Group** - The group of the root directory. This can be specified as a group ID or group name. For more information on mapping UNIX and Windows groups, see [Identity Mapping](#). For UNIX-based NFS access, this can be changed from the client using the `chgrp` command.
- **Permissions** - Standard UNIX permissions for the root directory. For UNIX-based NFS access, this can be changed from the client using the `chmod` command. The permissions are divided into three types.

Access Type	Description
User	User that is the current owner of the directory.
Group	Group that is the current group of the directory.
Other	All other accesses.

For each access type, the following permissions can be granted.

Type	Description
Read R	Permission to list the contents of the directory.

Type		Description
Write	W	Permission to create files in the directory.*
Execute	X	Permission to look up entries in the directory. If users have execute permissions but not read permissions, they can access files explicitly by name but not list the contents of the directory.

Related Topics

- [“ACL Behavior on Mode Change” on page 458](#)
- [“ACL Inheritance Behavior” on page 459](#)
- [“Root Directory ACL” on page 460](#)

ACL Behavior on Mode Change

When an ACL is modified via `chmod(2)` using the standard UNIX user/group/other permissions, the simplified mode change request will interact with the existing ACL in different ways depending on the setting of this property. To edit the ACL behavior on mode change, see [Editing a Project BUI](#), [CLI](#).

TABLE 123 Mode Change Values

BUI Value	CLI Value	Description
Discard ACL	discard	All ACL entries that do not represent the mode of the directory or file are discarded. This is the default behavior.
Mask ACL with mode	mask	The permissions are reduced, such that they are no greater than the group permission bits, unless it is a user entry that has the same UID as the owner of the file or directory. In this case, the ACL permissions are reduced so that they are no greater than owner permission bits. The mask value also preserves the ACL across mode changes, provided an explicit ACL set operation has not been performed.
Do not change ACL	passthrough	No changes are made to the ACL other than generating the necessary ACL entries to represent the new mode of the file or directory.

Related Topics

- [“Root Directory ACL” on page 460](#)
- [“ACL Inheritance Behavior” on page 459](#)
- [“Root Directory ACL” on page 460](#)

ACL Inheritance Behavior

When a new file or directory is created, it is possible to inherit existing ACL settings from the parent directory. This property controls how this inheritance works. These property settings usually only affect ACL entries that are flagged as inheritable - other entries are not propagated regardless of this property setting. However, all trivial ACL entries are inheritable when used with SMB. A trivial ACL represents the traditional UNIX owner/group/other entries. To edit the ACL inheritance behavior, see [Editing a Project BUI](#), [CLI](#).

TABLE 124 ACL Inheritance Behavior Values

BUI Value	CLI Value	Description
Do not inherit entries	discard	No ACL entries are inherited. The file or directory is created according to the client and protocol being used.
Only inherit deny entries	noallow	Only inheritable ACL entries specifying "deny" permissions are inherited.
Inherit all but "write ACL" and "change owner"	restricted	Removes the "write_acl" and "write_owner" permissions when the ACL entry is inherited, but otherwise leaves inheritable ACL entries untouched. This is the default.
Inherit all entries	passthrough	All inheritable ACL entries are inherited. The "passthrough" mode is typically used to cause all "data" files to be created with an identical mode in a directory tree. An administrator sets up ACL inheritance so that all files are created with a mode, such as 0664 or 0666.
Inherit all but "execute" when not specified	passthrough-x	Same as 'passthrough', except that the owner, group, and everyone ACL entries inherit the execute permission only if the file creation mode also requests the execute bit. The "passthrough" setting works as expected for data files, but you might want to optionally include the execute bit from the file creation

BUI Value	CLI Value	Description
		mode into the inherited ACL. One example is an output file that is generated from tools, such as "cc" or "gcc". If the inherited ACL doesn't include the execute bit, then the output executable from the compiler won't be executable until you use <code>chmod(1)</code> to change the file's permissions.
Inherit all, but preserve mode from client	<code>passthrough-mode-preserve</code>	Inheritable ACL entries are inherited, while preserving the creation mode specified by the application. This preserves the inheritance bits so SMB creates ACLs that interoperate well with shares accessed over NFS and SMB simultaneously. This property setting is only available after applying the deferred update for ACL Passthrough with Mode Preservation. For more information, see "Deferred Updates" in Oracle ZFS Storage Appliance Customer Service Manual .

When using SMB to create a file in a directory with a trivial ACL, all ACL entries are inherited. As a result, the following behavior occurs:

- Inheritance bits display differently when viewed in SMB or NFS. When viewing the ACL directory in SMB, inheritance bits are displayed. In NFS, inheritance bits are not displayed.
- When a file is created in a directory using SMB, its ACL entries are shown as inherited; however, when viewed through NFS, the directory has no inheritable ACL entries.
- If the ACL is changed so that it is no longer trivial, e.g., by adding an access control entry (ACE), this behavior does not occur.
- If the ACL is modified using SMB, the resulting ACL will have the previously synthetic inheritance bits turned into real inheritance bits.

Related Topics

- ["Project Properties" on page 417](#)

Root Directory ACL

Fine-grained access on files and directories is managed via Access Control Lists. An ACL describes what permissions are granted, if any, to specific users or groups. The appliance supports NFSv4.0 and NFSv4.1-style ACLs, also accessible over SMB. POSIX draft ACLs

(used by NFSv3) are not supported. Some trivial ACLs can be represented over NFSv3, but making complicated ACL changes may result in undefined behavior when accessed over NFSv3.

Like root directory access, this property only affects the root directory of the filesystem. ACLs can be controlled through in-band protocol management; BUI and CLI provide a way to set the ACL just for the root directory of the filesystem. You can use in-band management tools if the BUI is not an option. Changing this ACL does not affect existing files and directories in the filesystem. Depending on the ACL inheritance behavior, these settings may or may not be inherited by newly created files and directories. However, all ACL entries are inherited when SMB is used to create a file in a directory with a trivial ACL.

An ACL is composed of any number of ACEs (access control entries). Each ACE describes a type/target, a set of permissions, inheritance flags and a mode. ACEs are applied in order, starting at the beginning of the ACL, to determine whether a given action should be permitted. For information on in-band configuration ACLs through data protocols, consult the appropriate client documentation. The BUI interface for managing ACLs and the effect on the root directory are described here.

TABLE 125 Share - ACL Types

Type	Description
Owner	Current owner of the directory. If the owner is changed, this ACE will apply to the new owner.
Group	Current group of the directory. If the group is changed, this ACE will apply to the new group.
Everyone	Any user.
Named User	User named by the 'target' field. The user can be specified as a user ID or a name resolvable by the current name service configuration.
Named Group	Group named by the 'target' field. The group can be specified as a group ID or a name resolvable by the current name service configuration.

TABLE 126 Share - ACL Modes

Mode	Description
 Allow	The permissions are explicitly granted to the ACE target.
 Deny	The permissions are explicitly denied to the ACE target.

TABLE 127 Share - ACL Permissions

Permission	Description
Read	

	Permission	Description
(r)	Read Data/List Directory	Permission to list the contents of a directory. When inherited by a file, permission to read the data of the file.
(x)	Execute File/Traverse Directory	Permission to traverse (lookup) entries in a directory. When inherited by a file, permission to execute the file.
(a)	Read Attributes	Permission to read basic attributes (non-ACLs) of a file. Basic attributes are considered to be the stat level attributes, and allowing this permission means that the user can execute <code>ls</code> and <code>stat</code> equivalents.
(R)	Read Extended Attributes	Permission to read the extended attributes of a file or do a lookup in the extended attributes directory.
	Write	
(w)	Write Data/Add File	Permission to add a new file to a directory. When inherited by a file, permission to modify a file's data anywhere in the file's offset range. This include the ability to grow the file or write to any arbitrary offset.
(p)	Append Data/Add Subdirectory	Permission to create a subdirectory within a directory. When inherited by a file, permission to modify the file's data, but only starting at the end of the file. This permission (when applied to files) is not currently supported.
(d)	Delete	Permission to delete a file.
(D)	Delete Child	Permission to delete a file within a directory. As of the 2011.1 software release, if the sticky bit is set, a child file can only be deleted by the file owner.
(A)	Write Attributes	Permission to change the times associated with a file or directory.
(W)	Write Extended Attributes	Permission to create extended attributes or write to the extended attributes directory.
	Admin	
(c)	Read ACL/Permissions	Permission to read the ACL.
(C)	Write ACL/Permissions	Permission to write the ACL or change the basic access modes.

	Permission	Description
(o)	Change Owner	Permission to change the owner.
	Inheritance	
(f)	Apply to Files	Inherit to all newly created files in a directory.
(d)	Apply to Directories	Inherit to all newly created directories in a directory.
(i)	Do not apply to self	The current ACE is not applied to the current directory, but does apply to children. This flag requires one of "Apply to Files" or "Apply to Directories" to be set.
(n)	Do not apply past children	The current ACE should only be inherited one level of the tree, to immediate children. This flag requires one of "Apply to Files" or "Apply to Directories" to be set.

When the option to use Windows default permissions is used at share creation time, an ACL with the following three entries is created for the share's root directory:

TABLE 128 Share Root Directory Entities

Type	Action	Access
Owner	Allow	Full Control
Group	Allow	Read and Execute
Everyone	Allow	Read and Execute

In the CLI, set the root directory ACL properties after navigating to the shares context and selecting a project and filesystem. Use colons to separate the ACE properties, and separate multiple ACE entries with commas. The target and inheritance fields are optional. To set the properties, enter `set root_acl=ace1,ace2,ace3,...`, where `acen` is:

`type:<target:>permissions:<inheritance:>mode`

Examples:

```
set root_acl=owner@:r:allow
```

```
set root_acl=everyone@:rwx:fd:allow
```

```
set root_acl=user:root:r:allow
```

Working with Schemas

In addition to the standard built in properties, you can configure any number of additional properties that are available on all shares and projects. These properties are given basic types for validation purposes, and are inherited like most other standard properties. The values are never consumed by the software in any way, and exist solely for end-user consumption. The property schema is global to the system, across all pools, and is synchronized between cluster peers.

To work with schemas, see the following sections:

- [“Creating a Schema \(BUI\)” on page 464](#)
- [“Creating a Schema \(CLI\)” on page 464](#)
- [“Schema Properties” on page 466](#)

▼ Creating a Schema (BUI)

1. Go to **Shares > Schema**.
2. Click the '+' icon to add a new property to the schema property list.
3. Enter the name of the property ("contact").
4. Enter a description of the property ("Owner Contact").
5. Choose a type for the new property ("Email Address").
6. Click **Apply**.
7. Navigate to an existing share or project.
8. Change the "Owner Contact" property under the "Custom Properties" section.

▼ Creating a Schema (CLI)

1. Go to the schema context (`shares schema`).
2. Create a new property named "contact" (`create contact`).
3. Set the description for the property (`set description="Owner Contact"`).

4. **Set the type of the property (set type=EmailAddress).**
5. **Commit the changes (commit).**
6. **Go to an existing share or project.**
7. **Set the "custom:contact" property.**

Example 16 Example Schema

The schema context can be found at Shares > Schema".

```
carp:> shares schema
carp:shares schema> show
Properties:
```

NAME	TYPE	DESCRIPTION
owner	EmailAddress	Owner Contact

Each property is a child of the schema context, using the name of the property as the token. To create a property, use the create command:

```
carp:shares schema> create department
carp:shares schema department (uncommitted)> get
    type = String
    description = department
carp:shares schema department (uncommitted)> set description="Department Code"
    description = Department Code (uncommitted)
carp:shares schema department (uncommitted)> commit
carp:shares schema>
```

Within the context of a particular property, fields can be set using the standard CLI commands:

```
carp:shares schema> select owner
carp:shares schema owner> get
    type = EmailAddress
    description = Owner Contact
carp:shares schema owner> set description="Owner Contact Email"
    description = Owner Contact Email (uncommitted)
carp:shares schema owner> commit
```

Once custom properties have been defined, they can be accessed like any other property under the name "custom:<property>":

```

carp:shares default> get
...
                custom:department = 123-45-6789
                custom:owner =
...
carp:shares default> set custom:owner=bob@corp
                custom:owner = bob@corp (uncommitted)
carp:shares default> commit

```

Schema Properties

To define custom properties, access the Shares > Schema navigation item. The current schema is displayed as a list, and entries can be added or removed as needed. Each property has the following fields:

TABLE 129 Schema Property Fields

Field	Description
NAME	The CLI name for this property. This must contain only alphanumeric characters or the characters ". : _ \".
DESCRIPTION	The BUI name for this property. This can contain arbitrary characters and is used in the help section of the CLI
TYPE	The property type, for validation purposes. This must be one of the types described below.

The valid types for properties are:

TABLE 130 Valid Types for Properties

BUI Type	CLI Type	Description
String	String	Arbitrary string data. This is the equivalent of no validation.
Integer	Integer	A positive or negative integer
Positive Integer	PositiveInteger	A positive integer
Boolean	Boolean	A true/false value. In the BUI this is presented as a checkbox, while in the CLI it must be one of the values "true" or "false".
Email Address	EmailAddress	An email address. Only minimal syntactic validation is done.
Hostname or IP	Host	A valid DNS hostname or IP (v4 or v6) address.

Once defined, the properties are available under the general properties tab, using the description provided in the property table. Properties are identified by their CLI name, so renaming a property will have the effect of removing all existing settings on the system. A property that is removed and later renamed back to the original name will still refer to the previously set values. Changing the types of properties, while supported, may have undefined results on existing properties on the system. Existing properties will retain their current settings, even if they would be invalid given the new property type.

Shadow Migration

A common task for administrators is to move data from one location to another. In the most abstract sense, this problem encompasses a large number of use cases, from replicating data between servers to keeping user data on laptops in sync with servers. There are many external tools available to do this, but the appliance has two integrated solutions for migrating data that addresses the most common use cases. The first, replication, is intended for replicating data between one or more appliances, and is covered separately; see [“Remote Replication” on page 511](#). The second, shadow migration, is described here.

Shadow migration is a process for migrating data from external NAS sources with the intent of replacing or decommissioning the original once the migration is complete. This is most often used when introducing a new appliance into an existing environment in order to take over file sharing duties of another server, but a number of other novel uses are possible, outlined below.


To use shadow migration, see the following sections:

- [“Understanding Shadow Migration” on page 470](#)
- [“Creating a Shadow Filesystem” on page 472](#)
- [“Managing Background Migration” on page 473](#)
- [“Handling Migration Errors” on page 473](#)
- [“Monitoring Migration Progress” on page 474](#)
- [“Canceling Migration” on page 477](#)
- [“Snapshotting Shadow File Systems” on page 477](#)
- [“Backing Up Shadow File Systems” on page 478](#)
- [“Replicating Shadow File Systems” on page 478](#)
- [“Migrating Local File Systems” on page 478](#)
- [“Using Shadow Migration Analytics” on page 479](#)
- [“Testing Potential Shadow Migration using the CLI” on page 479](#)
- [“Migrating Data from an Active NFS Server using the CLI” on page 480](#)

Understanding Shadow Migration

Shadow migration uses interposition, but it is integrated into the appliance and does not require a separate physical machine. When filesystems are created, they can optionally "shadow" an existing directory, either locally or over NFS. In this scenario, downtime is scheduled once, where the source appliance X is placed into read-only mode, a share is created with the shadow property set, and clients are updated to point to the new share on the appliance. Clients can then access the appliance in read-write mode.

FIGURE 27 Shadow Migration



Once the shadow property is set, data is transparently migrated in the background from the source appliance locally. If a request comes from a client for a file that has not yet been migrated, the appliance will automatically migrate this file to the local server before responding

to the request. This may incur some initial latency for some client requests, but once a file has been migrated, all accesses are local to the appliance and have native performance. It is often the case that the current working set for a filesystem is much smaller than the total size, so once this working set has been migrated, regardless of the total native size on the source, there will be no perceived impact on performance.

The downside to shadow migration is that it requires a commitment before the data has finished migrating, though this is the case with any interposition method. During the migration, portions of the data exist in two locations, which means that backups are more complicated, and snapshots may be incomplete and/or exist only on one host. It is therefore extremely important that any migration between two hosts first be tested thoroughly to make sure that identity management and access controls are setup correctly. This need not test the entire data migration, but it should be verified that files or directories that are not world readable are migrated correctly, ACLs (if any) are preserved, and identities are properly represented on the new system.

Shadow migration is implemented using on-disk data within the filesystem, so there is no external database and no data stored locally outside the storage pool. If a pool is failed over in a cluster, or both system disks fail and a new head node is required, all data necessary to continue shadow migration without interruption will be kept with the storage pool.

The following lists the restrictions on the shadow source:

- In order to properly migrate data, the source filesystem or directory *must* be read-only*. Changes made to files source may or may not be propagated based on timing, and changes to the directory structure can result in unrecoverable errors on the appliance.
- Shadow migration supports migration only from NFS sources. NFSv4.0 and NFSv4.1 filesystems will yield the best results. NFSv2 and NFSv3 migration are possible, but ACLs will be lost in the process and files that are too large for NFSv2 cannot be migrated using that protocol. Migration from SMB sources is not supported.
- Shadow migration of LUNs is not supported.

During migration, if the client accesses a file or directory that has not yet been migrated, there is an observable effect on behavior. The following lists the shadow file system semantics:

- For directories, client requests are blocked until meta-data infrastructure is created on the migration target for any intervening directories for which infrastructure is not yet established. For files, only the portion of the file being requested is migrated, and multiple clients can migrate different portions of a file at the same time.
- Files and directories can be arbitrarily renamed, removed, or overwritten on the shadow filesystem without any effect on the migration process.
- For files that are hard links, the hard link count may not match the source until the migration is complete.
- The majority of file attributes are migrated when the directory is created, but the on-disk size (`st_nblocks` in the UNIX `stat` structure) is not available until a read or write operation is

done on the file. The logical size will be correct, but a `du(1)` or other command will report a zero size until the file contents are actually migrated.

- If the appliance is rebooted, the migration will pick up where it left off originally. While it will not have to re-migrate data, it may have to traverse some already-migrated portions of the local filesystem, so there may be some impact to the total migration time due to the interruption.
- Data migration makes use of private extended attributes on files. These are generally not observable except on the root directory of the filesystem or through snapshots. Adding, modifying, or removing any extended attribute that begins with `SUNWshadow` will have undefined effects on the migration process and will result in incomplete or corrupt state. In addition, filesystem-wide state is stored in the `.SUNWshadow` directory at the root of the filesystem. Any modification to this content will have a similar effect.
- Once a filesystem has completed migration, an alert will be posted, and the shadow attribute will be removed, along with any applicable metadata. After this point, the filesystem will be indistinguishable from a normal filesystem.
- Data can be migrated across multiple filesystems into a single filesystem, through the use of NFSv4.0 or NFSv4.1 automatic client mounts (sometimes called "mirror mounts") or nested local mounts.

Use the following rules to migrate identity information for files, including ACLs:

- The migration source and target appliance must have the same name service configuration.
- The migration source and target appliance must have the same NFSv4.0 or NFSv4.1 mapid domain
- The migration source must support NFSv4.0 and NFSv4.1. Use of NFSv3 is possible, but some loss of information will result. Basic identity information (owner and group) and POSIX permissions will be preserved, but any ACLs will be lost.
- The migration source must be exported with root permissions to the appliance.

If you see files or directories owned by "nobody", it likely means that the appliance does not have name services setup correctly, or that the NFSv4.0 or NFSv4.1 mapid domain is different. If you get 'permission denied' errors while traversing filesystems that the client should otherwise have access to, the most likely problem is failure to export the migration source with root permissions.

Creating a Shadow Filesystem

The shadow migration source can only be set when a filesystem is created. In the BUI, this is available in the filesystem creation dialog. In the CLI, it is available as the `shadow` property. The property takes one of the following forms:

- **Local** - file:///<path>
- **NFS** - nfs://<host>/<path>

The BUI also allows the alternate form <host>:/<path> for NFS mounts, which matches the syntax used in UNIX systems. The BUI also sets the protocol portion of the setting (file:// or nfs://) via the use of a pull down menu. When creating a filesystem, the server will verify that the path exists and can be mounted.

Managing Background Migration

When a share is created, it will automatically begin migrating in the background, in addition to servicing inline requests. This migration is controlled by the shadow migration service. There is a single global tunable, which is the number of threads dedicated to this task. Increasing the number of threads will result in greater parallelism at the expense of additional resources.

The shadow migration service can be disabled, but this should only be used for testing purposes, or when the active of shadow migration is overwhelming the system to the point where it needs to be temporarily stopped. When the shadow migration service is disabled, synchronous requests are still migrated as needed, but no background migration occurs. With the service disabled, no shadow migration will ever complete, even if all the contents of the filesystem are read manually. It is highly recommended to always leave the service enabled.

Handling Migration Errors

Because shadow migration requires committing new writes to the server prior to migration being complete, it is very important to test migration and monitor for any errors. Errors encountered during background migration are kept and displayed in the BUI as part of shadow migration status. Errors encountered during other synchronous migration are not tracked, but will be accounted for once the background process accesses the affected file. For each file, the remote filename as well as the specific error are kept. Clicking on the information icon next to the error count will bring up this detailed list. The error list is not updated as errors are fixed, but simply cleared by virtue of the migration completing successfully.

Shadow migration will not complete until all files are migrated successfully. If there are errors, the background migration will continually retry the migration until it succeeds. This allows the administrator to fix any errors (such as permission problems), let the migration complete, and be assured of success. If the migration cannot complete due to persistent errors, the migration can be canceled, leaving the local filesystem with whatever data was able to be migrated. This should only be used as a last resort; once migration has been canceled, it cannot be resumed.

Monitoring Migration Progress

To monitor shadow migration progress, the appliance provides such statistics as:

- Size of data transferred so far
- Estimate of remaining size to be migrated
- Migration time so far
- Migration time remaining
- Migration errors

At the beginning of migration, the appliance obtains the source filesystem statistics and calculates its size. It uses these values to provide a reasonably accurate visual representation of migration progress and an estimation of the remaining data to be migrated. Of note, the remaining bytes is an estimate based on the assumption that an entire filesystem is being migrated. If only part of the source file system is migrated, the remaining bytes estimate is inaccurate. If the source filesystem has nested filesystems, the total filesystem size is recalculated when the nested mount point is discovered during migration, and the remaining bytes are re-estimated based on the newly calculated total. Estimation of remaining bytes may be inaccurate if the source filesystem uses compression. These values are available in the BUI and CLI through both the standard filesystem properties as well as properties of the shadow migration node (or UI panel).


Note - When a sparse file (a file with empty blocks) is migrated, the target file will be smaller than the source file size. Shadow migration does not write the empty blocks to the target file, resulting in less space usage.

The following tasks describe how to monitor shadow migration progress and view any resulting errors. To view shadow migration errors using the RESTful API, see [“Filesystem Operations” in Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.7.x.](#)

▼ Monitoring Migration Progress and Errors (BUI)

1. **Go to Shares > Shares, and select a filesystem with shadow migration source.**

2. Under Shadow Migration, examine the progress bar and shadow migration status.



3. To view shadow migration errors, click the edit icon .

▼ Monitoring Migration Progress and Errors (CLI)

1. Go to a filesystem with shadow migration source and enter `shadow`, and then enter `list`.

```
hostname:shares default/file_sys1> shadow
hostname:shares default/file_sys1 shadow> list
Properties:
    source = nfs://zfs0000-15/sm/errors
    transferred = (unset)
    remaining = 1.37G
    elapsed = 0h3m
    errors = 23
    complete = true
    time = (unset)

Children:
    errors => Shadow Migration Errors
hostname:shares default/file_sys1 shadow>
```

Note - If there is no shadow migration source defined `shadow=none`, then the `shadow` command is invalid for the filesystem:

```
hostname:shares default/xyz_1> shadow
error: invalid command "shadow"
```

2. To view shadow migration errors, enter child node `errors` and enter `help` to view a list of subcommands.

```
hostname:shares default/file_sys1 shadow> errors
```

```
hostname:shares default/file_sys1 shadow errors> help
```

Subcommands that are valid in this context:

```

help [topic]          => Get context-sensitive help. If [topic] is specified,
                        it must be one of "builtins", "commands", "general",
                        "help" or "script".

show                  => Show information pertinent to the current context

done                  => Finish operating on "errors"

select [entry]        => Select the specified entry to get its properties,
                        set its properties, or run a subcommand

list                  => Lists up to the first 100 errors. The "-a" option may be
                        used to list all the errors if there are more than 100
                        errors. The "-number" option may be used to list the first
                        (number) of errors. Format: list -a or list -xx

```

3. Enter show to view individual migration errors in the current context.

```
hostname:shares default/file_sys1 shadow errors> show
Errors:
```

PATH	REASON
ak-2013-dev-on-clone.tar.gz	Permission denied
test_dir/CREDITS.html	Permission denied
test_dir/CREDITS_ja.html	Permission denied
test_dir/CREDITS_pt_BR.html	Permission denied
test_dir/CREDITS_zh_CN.html	Permission denied
test_dir/DISTRIBUTION.txt	Permission denied
test_dir/LEGALNOTICE.txt	Permission denied
test_dir/LICENSE.txt	Permission denied
test_dir/README.html	Permission denied
test_dir/README_ja.html	Permission denied
test_dir/README_pt_BR.html	Permission denied
test_dir/README_zh_CN.html	Permission denied
test_dir/THIRDPARTYLICENSE.txt	Permission denied
test_dir/bin	Permission denied
test_dir/cnd2	Permission denied
test_dir/etc	Permission denied
test_dir/gsf1	Permission denied
test_dir/ide10	Permission denied
test_dir/modCluster.properties	Permission denied


```
test_dir/nb6.5           Permission denied
test_dir/forms.css       Permission denied
test_dir/platform9       Permission denied
test_dir/websvccommon1   Permission denied
```

4. **To view an individual error, enter the `select` command and an individual error name. Then enter `show`.**

To view individual error properties, use the `get` command.

```
hostname:shares default/file_sys1 shadow errors> select test_dir/nb6.5

hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> show
Properties:
    path = test_dir/nb6.5
    reason = Permission denied

hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> get path
    path = test_dir/nb6.5
hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> get reason
    reason = Permission denied
hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> done
hostname:shares default/file_sys1 shadow errors>
```

Canceling Migration

Migration can be canceled, but this should only be done in extreme circumstances when the source is no longer available. Once migration has been canceled, it cannot be resumed. The primary purpose is to allow migration to complete when there are uncorrectable errors on the source. If the complete filesystem has finished migration except for a few files or directories, and there is no way to correct these errors (i.e. the source is permanently broken), then canceling the migration will allow the local filesystem to resume status as a 'normal' filesystem.

To cancel migration in the BUI, click the close icon next to the progress bar in the left column of the share in question. In the CLI, migrate to the shadow node beneath the filesystem and run the `cancel` command.

Snapshotting Shadow File Systems

Shadow filesystems can be snapshotted; however, the state of what is included in the snapshot is arbitrary. Files that have not yet been migrated will not be present, and implementation

details (such as SUNWshadow extended attributes) may be visible in the snapshot. This snapshot can be used to restore individual files that have been migrated or modified since the original migration began. Because of this, it is recommended that any snapshots be kept on the source until the migration is completed, so that unmigrated files can still be retrieved from the source if necessary. Depending on the retention policy, it may be necessary to extend retention on the source in order to meet service requirements. While snapshots can be taken, these snapshots cannot be rolled back to, nor can they be the source of a clone.

Backing Up Shadow File Systems

Filesystems that are actively migrating shadow data can be backed using NDMP as with any other filesystem. The shadow setting is preserved with the backup stream, but will be restored only if a complete restore of the filesystem is done and the share doesn't already exist. Restoring individual files from such a backup stream or restoring into existing filesystems may result in inconsistent state or data corruption. During the full filesystem restore, the filesystem will be in an inconsistent state (beyond the normal inconsistency of a partial restore) and shadow migration will not be active. Only when the restore is completed is the shadow setting restored. If the shadow source is no longer present or has moved, the administrator can observe any errors and correct them as necessary.

Replicating Shadow File Systems

Filesystems that are actively migrating shadow data can be replicated using the normal mechanism, but only the migrated data is sent in the data stream. As such, the remote side contains only partial data that may represent an inconsistent state. The shadow setting is sent along with the replication stream, so when the remote target is failed over, it will keep the same shadow setting. As with restoring an NDMP backup stream, this setting may be incorrect in the context of the remote target. After failing over the target, the administrator can observe any errors and correct the shadow setting as necessary for the new environment.

Migrating Local File Systems

In addition to its primary purpose of migrating data from remote sources, the same mechanism can also be used to migrate data from local filesystem to another on the appliance. This can be used to change settings that otherwise cannot be modified, such as creating a compressed version of a filesystem or changing the recordsize for a filesystem after the fact. In this model, the old share (or subdirectory within a share) is made read-only or moved aside, and a new

share is created with the shadow property set using the file protocol. Clients access this new share, and data is written using the settings of the new share.

Using Shadow Migration Analytics

In addition to standard monitoring on a per-share basis, it is also possible to monitor shadow migration system-wide through Analytics. The shadow migration analytics are available under the "Data Movement" category. There are three basic statistics available:

- **Shadow Migration Requests** - This statistic tracks requests for files or directories that are not cached and known to be local to the filesystem. It does account for both migrated and unmigrated files and directories, and it can be used to track the latency incurred as part of shadow migration, as well as track the progress of background migration. It can be broken down by file, share, project, or latency. It currently encompasses both synchronous and asynchronous (background) migration, so it is not possible to view only latency visible to clients.
- **Shadow Migration Bytes** - This statistic tracks bytes transferred as part of migrating file or directory contents. This does not apply to metadata (extended attributes, ACLs, etc). It gives a rough approximation of the data transferred, but source datasets with a large amount of metadata will show a disproportionately small bandwidth. The complete bandwidth can be observed by looking at network analytics. This statistic can be broken down by local filename, share, or project.
- **Shadow Migration Operations** - This statistic tracks operations that require going to the source filesystem. This can be used to track the latency of requests from the shadow migration source. It can be broken down by file, share, project, or latency.

▼ Testing Potential Shadow Migration using the CLI

Before attempting a complete migration, it is important to test the migration to make sure that the appliance has appropriate permissions and security attributes are translated correctly. Once you are confident that the basic setup is functional, the filesystems can be setup for the final migration.

Note - As part of capacity planning, remember to take into account default/user group quotas because the quotas could be exceeded if the source is larger than the destination. Also, shadow migration will fail if the target runs out of disk space.

1. **Configure the source so that the appliance has root access to the share.** This typically involves adding an NFS host-based exception or setting

the anonymous user mapping (the latter having more significant security implications).

2. Create a share on the local filesystem with the shadow attribute set to 'nfs://<host>/<snapshotpath>' in the CLI or just '<host>/<snapshotpath>' in the BUI (with the protocol selected as 'NFS'). The snapshot should be read-only copy of the source. If no snapshots are available, a read-write source can be used, but may result in undefined errors.
3. Validate that file contents and identity mapping are correctly preserved by traversing the file structure.
4. If the data source is read-only (as with a snapshot), let the migration complete and verify that there were no errors in the transfer.

▼ Migrating Data from an Active NFS Server using the CLI

Use the following procedure to migrate data from an active NFS server using the CLI. Note that shadow migration fails if it encounters files under procfs or the following special file types: doors, sockets, and event ports.

1. Schedule downtime during which clients can be quiesced and reconfigured to point to a new server.
2. Configure the source so that the appliance has root access to the share. This typically involves adding an NFS host-based exception or setting the anonymous user mapping (the latter having more significant security implications).
3. Configure the source to be read-only. This step is technically optional, but it is much easier to guarantee compliance if it is impossible for misconfigured clients to write to the source while migration is in progress.
4. Create a share on the local filesystem with the shadow attribute set to 'nfs://<host>/<path>' in the CLI or just '<host>/<path>' in the BUI (with the protocol selected as 'NFS').
5. Reconfigure clients to point at the local share on the appliance.

At this point, shadow migration should be running in the background, and client requests should be serviced as necessary. You can observe the progress as described above. Multiple filesystems can be created during a single scheduled downtime through scripting the CLI.

Snapshots and Clones

Note - Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Using snapshots and clones, you can make point-in-time copies of a share or project. These copies can be useful as backups or as different working versions.

A snapshot is a read-only copy of a filesystem, LUN, or project. Taking a project snapshot is equivalent to snapshotting all of the shares in the project. A snapshot takes up no additional space when it is first created, but as the active share changes, the snapshot takes up additional space, with a maximum equivalent to the size of the share at the time the snapshot was taken.

A clone is a writable copy of a filesystem or LUN snapshot and can be treated as an independent share. Clones of projects are not supported. Like a snapshot, a clone consumes no additional space when it is first created, but as new data is written to the clone, the space required for the changes are associated with the clone.

You can take snapshots manually, or you can set a schedule so that snapshots are taken automatically every half-hour, hour, day, week, or month. Some snapshots are taken by the appliance automatically during replication updates; these will appear on the snapshot page with `.ndmp` and `.rr` in their names.

For information about snapshot space management, see the following:

[“Snapshot Space Management” on page 482](#)

To take manual snapshots or schedule automatic snapshots of projects or shares, use the following tasks:

- Taking a Snapshot [BUI](#), [CLI](#)
- Scheduling Snapshots [BUI](#), [CLI](#)
- Setting a Scheduled Snapshot Label [BUI](#), [CLI](#)

You can make clones of a snapshot, which can be useful to create numerous working versions of one share. To make clones, use the following tasks:

- Cloning a Snapshot [BUI](#), [CLI](#)
- “Cloning a Clone” on page 508
- Cloning a Replication Package [BUI](#), [CLI](#)

To determine the relationships between existing snapshots and clones, use the following tasks:

- Viewing Clones of a Snapshot [BUI](#), [CLI](#)
- Viewing a Clone Origin [BUI](#), [CLI](#)

To view and edit existing snapshots, snapshot schedules, and snapshot retention policies, use the following tasks:

- Viewing Snapshots and Schedules [BUI](#), [CLI](#)
- Renaming a Snapshot [BUI](#), [CLI](#)
- Editing a Snapshot Retention Policy [BUI](#), [CLI](#)

You can look at the contents of filesystem snapshots through the `.zfs/snapshot` filesystem directory. LUN snapshots cannot be accessed directly, though they can be used as a rollback target or as the source of a clone. To manage and access the `.zfs/snapshot` directory, use the following tasks:

- Making a Filesystem Snapshot Directory Visible [BUI](#), [CLI](#)
- “Accessing a Hidden Filesystem Snapshot Directory (CLI)” on page 498
- “Accessing a Visible Filesystem Snapshot Directory (CLI)” on page 498

You can use an existing snapshot to restore a filesystem or LUN to the exact state it was in when the snapshot was taken. To roll back a filesystem, LUN, or project to an existing snapshot, use the following tasks:

- Rolling Back to a Snapshot [BUI](#), [CLI](#)

To destroy snapshots, use the following tasks:

- Destroying a Snapshot [BUI](#), [CLI](#)

Snapshot Space Management

Snapshots present an interesting dilemma for space management. They represent the set of physical blocks referenced by a share at a given point in time. Initially, this snapshot consumes no additional space. But as new data is overwritten in the new share, the blocks in the active share will only contain the new data, and older blocks will be “held” by the most recent (and

possibly older) snapshots. Gradually, snapshots can consume additional space as the content diverges in the active share. If you take a snapshot of a filesystem of any given size, and re-write 100% of the data within the filesystem, you must maintain references to twice the data that was originally in the filesystem.

Each snapshot has two associated space statistics: unique space and referenced space. The amount of referenced space is the total space consumed by the filesystem at the time the snapshot was taken. It represents the theoretical maximum size of the snapshot should it remain the sole reference to all data blocks. The unique space indicates the amount of physical space referenced only by the current snapshot. When a snapshot is destroyed, the unique space is made available to the rest of the pool.

Note that the amount of space consumed by all snapshots is not equivalent to the sum of unique space across all snapshots. With a share and a single snapshot, all blocks must be referenced by one or both of the snapshot or the share. With multiple snapshots, however, it's possible for a block to be referenced by some subset of snapshots, and not any particular snapshot. For example, if a file is created, two snapshots X and Y are taken, the file is deleted, and another snapshot Z is taken, the blocks within the file are held by X and Y, but not by Z. In this case, destroying Z will not free up the space, but destroying both X and Y will. Because of this, destroying any snapshot can affect the unique space referenced by neighboring snapshots, though the total amount of space consumed by snapshots will always decrease.

The total size of a project or share always accounts for space consumed by all snapshots, though the usage breakdown is also available. Quotas and reservations can be set at the project level to enforce physical constraints across this total space. In addition, quotas and reservations can be set at the filesystem level, and these settings can apply to only referenced data or total data.

Whether or not quotas and reservations should be applied to referenced data or total physical data depends on the administrative environment. If users are not in control of their snapshots (i.e. an automatic snapshot schedule is set for them), then quotas should typically not include snapshots in the calculation. Otherwise, the user may run out of space but be confused when files cannot be deleted. Without an understanding of snapshots or means to manage those snapshots, it is possible for such a situation to be unrecoverable without administrator intervention. In this scenario, the snapshots represent an overhead cost that is factored into operation of the system in order to provide backup capabilities. On the other hand, there are environments where users are billed according to their physical space requirements, and snapshots represent a choice by the user to provide some level of backup that meets their requirements given the churn rate of their dataset. In these environments, it makes more sense to enforce quotas based on total physical data, including snapshots. The users understand the cost of snapshots, and can be provided a means to actively manage them (as through dedicated roles on the appliance).

Related Topics

- [“Space Management for Shares” on page 436](#)


- [“Managing User-Generated Snapshots” on page 557](#)

▼ Taking a Snapshot (BUI)

Use the following procedure to take a manual snapshot of a filesystem, LUN, or project.

To schedule automatic snapshots at regular intervals, see [“Scheduling Snapshots \(BUI\)” on page 485](#).

1. Go to the share or project you want to snapshot.
 - a. To take a snapshot of a filesystem, go to Shares > Shares.
 - b. To take a snapshot of a LUN, go to Shares > Shares and click LUNs.
 - c. To take a snapshot of a project, go to Shares > Projects.
2. Hover over the share or project and click the Edit icon .
3. Click the Snapshots tab.
4. Click the Add icon  next to Snapshots.
5. Type a name for the snapshot.



6. Click **APPLY**.

▼ Taking a Snapshot (CLI)

Use the following procedure to take a manual snapshot of a filesystem, LUN, or project.

To schedule automatic snapshots at regular intervals, see [“Scheduling Snapshots \(CLI\)” on page 487](#).

1. **Go to the share or project you want to snapshot.**

- a. **To take a snapshot of a project, go to shares and select the project.**

```
hostname:shares> select myproject
hostname:shares myproject>
```

- b. **To take a snapshot of a filesystem or LUN, go to shares and select the project containing the share, then select the share.**

```
hostname:shares> select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. **Enter snapshots.**

```
hostname:shares myproject/demo_share> snapshots
```

3. **Use the snapshot command followed by the name you want to give the new snapshot.**

```
hostname:shares myproject/demo_share snapshots> snapshot demo_snap
```

▼ Scheduling Snapshots (BUI)

Use the following procedure to configure automatic snapshots of a filesystem, LUN, or project and set a retention policy for those snapshots.

Automatic snapshots can be taken half-hourly, hourly, daily, weekly, or monthly and are named `.auto[-<snaplabel>]-<timestamp>`. In the Creation column of the Snapshots list, times are displayed in the local (client browser) time zone. However, times are stored and executed in UTC format, without regard to such conventions as daylight saving time. For example, a snapshot scheduled for 10:00 a.m. PST (UTC-8) is stored and executed at 18:00 UTC, and this is the time that will appear as the timestamp in the snapshot name.

Automatic snapshots can be set on a project or a share, but not both. Otherwise, overlapping schedules and retention policies would make it impossible to guarantee both schedules.

Removing an interval, or changing its retention policy, will immediately destroy any automatic snapshots not covered by the new schedule. Automatic snapshots with clones are ignored.

Note - Previous versions of the software allowed for automatic snapshots at the frequency of a minute. To help users avoid placing undue stress on the system, this feature was removed with the 2010.Q3 release. If the software is rolled back, existing minutes will be preserved. Previous instances will expire according to the existing schedule, but no new snapshots will be taken. An alert will be posted if a share or project with this frequency is found.

1. **Go to the share or project.**
 - a. **To schedule snapshots of a filesystem, go to Shares > Shares.**
 - b. **To schedule snapshots of a LUN, go to Shares > Shares and click LUNs.**
 - c. **To schedule snapshots of a project, go to Shares > Projects.**
2. **Hover over the share or project and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Click Schedules.**
5. **Click the Add icon  next to Schedules.**
6. **Set each field appropriately.**
 - a. **Set the frequency to half-hourly, hourly, daily, weekly, or monthly to indicate how often the snapshot is automatically taken.**
 - b. **Set the precise time the snapshot is automatically taken.**

For half-hourly or hourly snapshots, you can choose how many minutes after the half-hour or hour the snapshot is taken. For daily snapshots, you can choose the hour and minute the snapshot is taken, and for weekly or monthly snapshots, you can specify the day, hour, and minute.
 - c. **Set the "Keep at most" property to specify how many snapshots should be retained, or uncheck the checkbox to set no retention policy.**

Automatic snapshots can be kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively), or they can be limited to a certain number. When the number of snapshots exceeds the number you have specified here, the oldest snapshots will be deleted first.

7. Click **APPLY**.

▼ Scheduling Snapshots (CLI)

Use the following procedure to configure automatic snapshots of a share and set a retention policy for those snapshots.

Automatic snapshots can be taken half-hourly, hourly, daily, weekly, or monthly and are named `.auto[-<snaplabel>]-<timestamp>`. Snapshot creation times are stored and executed in UTC format, without regard to such conventions as daylight saving time. For example, a snapshot scheduled for 10:00 a.m. PST (UTC-8) is stored and executed at 18:00 UTC, and this is the time that will appear as the timestamp in the snapshot name.

Automatic snapshots can be set on a project or a share, but not both. Otherwise, overlapping schedules and retention policies would make it impossible to guarantee both schedules. Removing an interval, or changing its retention policy, will immediately destroy any automatic snapshots not covered by the new schedule. Automatic snapshots with clones are ignored.

Note - Previous versions of the software allowed for automatic snapshots at the frequency of a minute. To help users avoid placing undue stress on the system, this feature was removed with the 2010.Q3 release. If the software is rolled back, existing minutes will be preserved. Previous instances will expire according to the existing schedule, but no new snapshots will be taken. An alert will be posted if a share or project with this frequency is found.

1. Go to shares and select the project or share you want to snapshot.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. Enter snapshots automatic.

```
hostname:shares myproject/demo_share> snapshots automatic
hostname:shares myproject/demo_share snapshots automatic>
```

3. Use the create command to enter an uncommitted schedule context.

```
hostname:shares myproject/demo_share snapshots automatic> create
hostname:shares myproject/demo_share snapshots automatic (uncommitted)>
```

4. Use the set command to set each field appropriately.

- a. Set the frequency to `halfhour`, `hour`, `day`, `week`, or `month` to indicate how often the snapshot is automatically taken.

b. Set the day, hour, and minute to specify the precise time the snapshot is automatically taken.

For half-hourly or hourly snapshots, you can choose how many minutes after the half-hour or hour the snapshot is taken. For daily snapshots, you can choose the hour and minute the snapshot is taken, and for weekly or monthly snapshots, you can specify the day, hour, and minute.

c. Set the `keep` property to the number of snapshots you want to retain for this schedule.

Automatic snapshots can be kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively), or they can be limited to a certain number. When the number of snapshots exceeds the number you have specified here, the oldest snapshots will be deleted first.

```
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set
frequency=day
frequency = day (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set hour=14
hour = 14 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set minute=30
minute = 30 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set keep=7
keep = 7 (uncommitted)
```

You can use the `get` command to view the current uncommitted settings.

```
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> get
frequency = day (uncommitted)
day = (unset)
hour = 14 (uncommitted)
minute = 30 (uncommitted)
keep = 7 (uncommitted)
```

5. Enter `commit` to commit the changes and create the automatic snapshot schedule.

```
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> commit
```

You can use the `list` command to view your new schedule.

```
hostname:shares myproject/demo_share snapshots automatic> list
NAME          FREQUENCY    DAY          HH:MM KEEP
automatic-000 day          -            14:30    7
```

6. Enter `done` to finish.

```
hostname:shares myproject/demo_share snapshots automatic> done
hostname:shares myproject/demo_share snapshots>
```

▼ Setting a Scheduled Snapshot Label (BUI)

Use the following procedure to set a label for scheduled snapshots of a filesystem, LUN, or project.

This optional property appends a user-defined label to each scheduled snapshot and is blank by default. The label can either be set for an individual share, or it can be set for a project and inherited by its shares, but not both.

Snapshot labels can help identify the project or share for which a snapshot was taken. For example, "project1:share1" could indicate a scheduled snapshot taken on share1 within project1. Labels can be up to 35 alphanumeric characters and can include the special characters _ - . : .

1. **Go to the share or project for which you want to set the scheduled snapshot label.**
 - a. To set a label for a filesystem, go to **Shares > Shares**.
 - b. To set a label for a LUN, go to **Shares > Shares** and click **LUNs**.
 - c. To set a label for a project, go to **Shares > Projects**.
2. **Hover over the appropriate share or project and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Under Properties, type the label you want to set into the Scheduled Snapshot Label field.**
5. **Click APPLY to save the change.**

This label will be included in the name of each scheduled snapshot taken from now on. The label will appear before the timestamp, so that the snapshot name is .auto-<snaplabel>-<timestamp>.

▼ Setting a Scheduled Snapshot Label (CLI)

Use the following procedure to set a label for scheduled snapshots of a filesystem, LUN, or project.

This optional property appends a user-defined label to each scheduled snapshot and is blank by default. The label can either be set for an individual share, or it can be set for a project and inherited by its shares, but not both.

Snapshot labels can help identify the project or share for which a snapshot was taken. For example, "project1:share1" could indicate a scheduled snapshot taken on share1 within project1. Labels can be up to 35 alphanumeric characters and can include the special characters _ - . : .

1. **Go to shares and select the filesystem, LUN, or project for which you want to set the label.**

```
hostname:shares myproject> select demo_share  
hostname:shares myproject/demo_share>
```

2. **Use the set snapLabel command to create a scheduled snapshot label.**

```
hostname:shares myproject/demo_share> set snapLabel=myproject:demo_share
```

▼ Viewing Snapshots and Schedules (BUI)

Use the following procedure to view the snapshots and automatic snapshot schedules of a particular filesystem, LUN, or project.

1. **Go to the share or project.**
 - a. **To view snapshots and snapshot schedules of a filesystem, go to Shares > Shares.**
 - b. **To view snapshots and snapshot schedules of a LUN, go to Shares > Shares and click LUNs.**
 - c. **To view snapshots and snapshot schedules of a project, go to Shares > Projects.**
2. **Hover over the share or project and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **View the snapshots or snapshot schedules.**
 - a. **View the snapshots of that share under Snapshots and optionally select MANUAL or SCHEDULED to view only manual or only scheduled snapshots.**

For each snapshot, you can see the name (NAME), the date and time of creation (CREATION), the amount of unique space used by the snapshot (UNIQUE), the total amount of space referenced by the snapshot (TOTAL), and the number of clones of the snapshot (CLONES).

b. Click Schedules to view the automatic snapshot schedules for that share.

For each schedule, you can see the frequency that a snapshot is taken, the precise day and time at which it is taken, and how many snapshots are kept.

▼ Viewing Snapshots and Schedules (CLI)

Use the following procedure to view the snapshots and automatic snapshot schedules of a particular filesystem, LUN, or project.

1. Go to shares and select the project or share.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. Enter snapshots.

```
hostname:shares myproject/demo_share> snapshots
hostname:shares myproject/demo_share snapshots>
```

3. View the snapshots or snapshot schedules using the appropriate commands.

a. Enter list to view a list of the snapshots of this share or project.

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
hostname:shares myproject/demo_share snapshots>
```

You can select a snapshot and use the `list` command to see the following properties:

- `creation` - the date and time of creation of the snapshot in UTC format
- `numclones` - the number of clones of the snapshot
- `isauto` - whether the snapshot was created manually (`false`) or with an automatic snapshot schedule (`true`)
- `pool` - what storage pool the snapshot is in
- `canonical_name` - the location of the snapshot

- shadowsnap - whether the snapshot was taken during shadow migration (true) or not (false)
- space_unique - the amount of unique space the snapshot uses
- space_data - the total amount of space the snapshot references

```
hostname:shares myproject/demo_share snapshots> select demo_snap1
hostname:shares myproject/demo_share snapshots demo_snap1> list
Properties:
    creation = Thu Jan 22 2015 20:19:49 GMT+0000(UTC)
    numclones = 1
    isauto = false
    pool = pool1
    canonical_name = pool1/local/myproject/demo_share@demo_snap1
    shadowsnap = false
    space_unique = 0
    space_data = 31K
```

- b. **Enter `automatic` and use the `list` command to view a list of the automatic snapshot schedules of this share or project.**

```
hostname:shares myproject/demo_share snapshots> automatic
hostname:shares myproject/demo_share snapshots automatic> list
Properties:
    convert = false

Automatics:

NAME          FREQUENCY  DAY  HH:MM  KEEP
automatic-000 day        -    00:00   4
automatic-001 month      01    00:00  12
```

▼ Editing a Snapshot Retention Policy (BUI)

Use the following procedure to edit a snapshot retention policy for a filesystem, LUN, or project. Snapshot retention policies are included in automatic snapshot schedules.

1. **Go to the appropriate project or share.**
 - a. **If the schedule applies to a project, go to Shares > Projects.**
 - b. **If the schedule applies to a filesystem, go to Shares > Shares.**
 - c. **If the schedule applies to a LUN, go to Shares > Shares and click LUNs.**

2. **Hover over the appropriate project or share and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Click Schedules.**
5. **Set the "Keep at most" property for a schedule to specify how many automatic snapshots should be retained from that schedule, or uncheck the checkbox to set no retention policy.**

Automatic snapshots can be kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively), or they can be limited to a certain number. When the number of snapshots exceeds the number you have specified here, the oldest snapshots will be deleted first.

▼ Editing a Snapshot Retention Policy (CLI)

Use the following procedure to edit a snapshot retention policy for a filesystem, LUN, or project. Snapshot retention policies are included in automatic snapshot schedules.

1. **Go to shares and select the project or share.**

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. **Enter `snapshots automatic`.**

You can use the `list` command to view a list of the automatic snapshot schedules for this share or project.

```
hostname:shares myproject/demo_share> snapshots automatic
hostname:shares myproject/demo_share snapshots automatic> list
Properties:
    convert = false
```

Automatics:

NAME	FREQUENCY	DAY	HH:MM	KEEP
automatic-000	hour	-	-:02	4

3. **Select the schedule you want to edit.**

```
hostname:shares myproject/demo_share snapshots automatic> select automatic-000
```

4. **Set the `keep` property to the number of snapshots you want to retain for this schedule.**

Set the **keep** property to the number of snapshots you want to retain for this schedule (for half-hourly and hourly snapshots, this value is capped at 48 and 24, respectively). When the number of snapshots exceeds the number you specify, the oldest snapshots will be deleted first.

You can set **keep=0** to set no retention policy. In this case, automatic snapshots from that schedule are kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively).

```
hostname:shares myproject/demo_share snapshots automatic-000> set keep=4
keep=4 (uncommitted)
```

5. **Enter `commit` to save the change.**

```
hostname:shares myproject/demo_share snapshots automatic-000> commit
```

▼ Removing a Snapshot Schedule (BUI)

Use the following procedure to delete an automatic snapshot schedule for a filesystem, LUN, or project.

1. **Go to the appropriate project or share.**
 - a. If the schedule applies to a project, go to **Shares > Projects**.
 - b. If the schedule applies to a filesystem, go to **Shares > Shares**.
 - c. If the schedule applies to a LUN, go to **Shares > Shares** and click **LUNs**.
2. **Hover over the appropriate project or share and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Click Schedules.**
5. **Hover over the schedule you want to remove and click the remove icon .**
A window appears warning you that existing automatic snapshots might be destroyed.
6. **If you want to keep existing automatic snapshots, click **CONVERT** to convert them to manual snapshots. Otherwise, click **DISCARD** to destroy them.**

▼ Removing a Snapshot Schedule (CLI)

Use the following procedure to delete an automatic snapshot schedule for a filesystem, LUN, or project.

1. Go to and select the project or share.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
```

2. Enter snapshots automatic.

```
hostname:shares myproject/demo_share> snapshots automatic
hostname:shares myproject/demo_share snapshots automatic>
```

You can use the `list` command to view a list of the automatic snapshot schedules for this share or project.

```
hostname:shares myproject/demo_share snapshots automatic> list
Properties:
    convert = false
```

Automatics:

NAME	FREQUENCY	DAY	HH:MM	KEEP
automatic-000	day	-	00:00	4

3. If you want to keep existing automatic snapshots taken with this schedule, set the `convert` property to `true`.

If you keep this property as `convert = false`, automatic snapshots taken with this schedule will be discarded when you destroy the schedule.

```
hostname:shares myproject/demo_share snapshots automatic> set convert=true
    convert = true
hostname:shares myproject/demo_share snapshots automatic> commit
```

You can use the `list` command to view the change.

```
hostname:shares myproject/demo_share snapshots automatic> list
Properties:
    convert = true
```

Automatics:

NAME	FREQUENCY	DAY	HH:MM	KEEP
automatic-000	day	-	00:00	4

4. **Use the `destroy` command followed by the name of the automatic snapshot schedule you want to destroy.**

You are asked to confirm.

```
hostname:shares myproject/demo_share snapshots automatic> destroy automatic-000  
This will destroy "automatic-000". Are you sure? (Y/N)
```

5. **Type `y` to confirm.**

```
This will destroy "automatic-000". Are you sure? (Y/N) y
```

▼ Making a Filesystem Snapshot Directory Visible (BUI)

Use the following procedure to set the `.zfs/snapshot` directory, which is hidden by default, to appear like any other directory in a filesystem.

The `.zfs/snapshot` directory contains a list of all snapshots on the filesystem. The snapshots can be accessed just like normal filesystem data, but are read-only. By default, the `.zfs` directory is not visible when listing directory contents. This setting prevents backup software from inadvertently backing up snapshots in addition to new data.

Note - Setting the `.zfs/snapshot` directory to "visible" may cause backup software to back up snapshots in addition to live data.

1. **Go to Shares > Shares.**
2. **Hover over the filesystem and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Uncheck the "Inherit from project" box next to Properties, or click the lock icon next to ".zfs/snapshot visibility."**
5. **Select "Visible" from the drop-down menu next to ".zfs/snapshot visibility."**
6. **Click APPLY to save the changes.**

To make the directory hidden again, return to this page and select "Hidden" from the drop-down menu, then click APPLY.

▼ Making a Filesystem Snapshot Directory Visible (CLI)

Use the following procedure to set the `.zfs/snapshot` directory, which is hidden by default, to appear like any other directory in a filesystem.

The `.zfs/snapshot` directory contains a list of all snapshots on the filesystem. The snapshots can be accessed just like normal filesystem data, but are read-only. By default, the `.zfs` directory is not visible when listing directory contents. This setting prevents backup software from inadvertently backing up snapshots in addition to new data.

Note - Setting the `.zfs/snapshot` directory to "visible" may cause backup software to back up snapshots in addition to live data.

1. Go to and select the filesystem share.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. Use the `set snapdir` command to set the filesystem snapshot directory to `visible`.

```
hostname:shares myproject/demo_share> set snapdir=visible
snapdir=visible(uncommitted)
```

3. Type `commit` to save the change.

```
hostname:shares myproject/demo_share> commit
```

4. To make the directory hidden again, return to this context and use the `set snapdir` command to set the directory to hidden, then type `commit` to save the change.

```
hostname:shares myproject/demo_share> set snapdir=hidden
snapdir=hidden(uncommitted)
hostname:shares myproject/demo_share> commit
```

Related Topics

- [“Accessing a Hidden Filesystem Snapshot Directory \(CLI\)” on page 498](#)
- [“Making a Filesystem Snapshot Directory Visible \(BUI\)” on page 496](#)
- [“Accessing a Visible Filesystem Snapshot Directory \(CLI\)” on page 498](#)

▼ Accessing a Hidden Filesystem Snapshot Directory (CLI)

Use the following procedure to access filesystem snapshots over data protocols at `.zfs/snapshot` in the root of your filesystem.

The `.zfs/snapshot` directory contains a list of all snapshots on the filesystem. The snapshots can be accessed just like normal filesystem data, but are read-only. By default, the `.zfs` directory is not visible when listing directory contents, but it can be accessed by explicitly looking it up. This prevents backup software from inadvertently backing up snapshots in addition to new data.

1. **In a terminal window, go to the directory where you mounted the share.**

2. **Look up `.zfs/snapshot`.**

From here, you can list snapshots of this filesystem and look at the contents of each snapshot.

Example 17 Accessing `.zfs/snapshot`

In this example, there are two snapshots of a filesystem. The first snapshot contains three files.

```
$ ls -l /mnt/demo
$ ls -l /mnt/demo/.zfs/snapshot
demo_snap1
demo_snap2
$ ls -l /mnt/demo/.zfs/snapshot/demo_snap1
file1
file2
file3
```

Related Topics

- [“Making a Filesystem Snapshot Directory Visible \(BUI\)” on page 496](#)
- [“Making a Filesystem Snapshot Directory Visible \(CLI\)” on page 497](#)
- [“Accessing a Visible Filesystem Snapshot Directory \(CLI\)” on page 498](#)

▼ Accessing a Visible Filesystem Snapshot Directory (CLI)

Use the following procedure to access filesystem snapshots in the `.zfs/snapshot` directory after making it visible.

Before You Begin Set the `.zfs/snapshot` directory to Visible as described in [“Making a Filesystem Snapshot Directory Visible \(CLI\)” on page 497](#).

1. **Go to the directory where you mounted the share.**
2. **Go to `.zfs/snapshot` within that directory.**

From here, you can list snapshots of this filesystem and look at the contents of each snapshot.

▼ Renaming a Snapshot (BUI)

Use the following procedure to rename an existing manual snapshot. An automatic snapshot, which has `.auto`, `.rr`, or `.ndmp` in its name, cannot be renamed.

If a share snapshot that is part of a larger project snapshot is renamed, it will no longer be considered part of the same snapshot, and if any snapshot is renamed to have the same name as a snapshot in the parent project, it will be treated as part of the project snapshot.

- Before You Begin**
- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
 - `renameSnap` - Allows renaming snapshots.
 - `rename` - Allows renaming projects and shares, including snapshot names.
 - To add authorizations to a role, see [“Editing Authorizations for a Role \(BUI\)” on page 214](#).

1. **Go to the share or project that contains the snapshot you want to rename.**
 - To rename a filesystem snapshot, go to **Shares > Shares**.
 - To rename a LUN snapshot, go to **Shares > Shares** and click **LUNs**.
 - To rename a project snapshot, go to **Shares > Projects**.
2. **Hover over the share or project that contains the snapshot you want to rename and click the edit icon .**
3. **Click the Snapshots tab.**
4. **Under Snapshots, click the name of the snapshot you want to rename.**

The snapshot name changes to a text input box.
5. **Type the new name for the snapshot.**

A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters `_ - . :`

6. **Press Enter to commit the change.**

Related Topics

- [“Understanding Users and Roles” on page 218](#)
- [“User Authorizations” on page 219](#)

▼ Renaming a Snapshot (CLI)

Use the following procedure to rename an existing manual snapshot. Automatic snapshots cannot be renamed.

If a share snapshot that is part of a larger project snapshot is renamed, it will no longer be considered part of the same snapshot, and if any snapshot is renamed to have the same name as a snapshot in the parent project, it will be treated as part of the project snapshot.

- Before You Begin**
- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
 - `renameSnap` - Allows renaming snapshots.
 - `rename` - Allows renaming projects and shares, including snapshot names.
 - To add authorizations to a role, see [“Editing Authorizations for a Role \(CLI\)” on page 215](#).

1. **Go to shares and select the project, or select the project and then a share.**

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
```

2. **Enter snapshots.**

```
hostname:shares myproject/demo_share> snapshots
```

3. **Enter `list` to view the list of snapshots for the project or share.**

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
```

4. **To rename the snapshot, enter `rename` followed by the current snapshot name, a space, and then the new snapshot name.**

A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters `_ - . :`

```
hostname:shares myproject/demo_share snapshots> rename demo_snap1 new_name
```


Related Topics

- [“Understanding Users and Roles” on page 218](#)
- [“User Authorizations” on page 219](#)

▼ Rolling Back to a Snapshot (BUI)

Use the following procedure to roll back, or restore, a filesystem or LUN to an existing snapshot.

When a rollback occurs, any newer snapshots (and clones of newer snapshots) are destroyed, and the active data are reverted to the state when the snapshot was taken. Snapshots only include data, not properties, so any property settings changed since the snapshot was taken will remain. Changes to filesystem root directory access are lost during rollback.



Caution - This procedure cannot be undone.

1. **Go to the share or project that contains the snapshot you want to rename.**
 - **To restore a filesystem snapshot, go to Shares > Shares.**
 - **To restore a LUN snapshot, go to Shares > Shares and click LUNs.**
2. **Hover over the share that contains the snapshot you want to restore and click the edit icon .**
3. **Click the Snapshots tab.**
4. **Hover over the snapshot you want to restore, click its rollback icon , and confirm your action.**

▼ Rolling Back to a Snapshot (CLI)

Use the following procedure to roll back, or restore, a filesystem or LUN to an existing snapshot.

Restoring a snapshot requires destroying any newer snapshots and their clones, and it reverts the share contents to what they were at the time the snapshot was taken. Property settings on the share are not affected, but changes to filesystem root directory access are lost during rollback.



Caution - This procedure cannot be undone.

1. **Go to and select the share that contains the snapshot you want to restore.**

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
```

2. **Enter snapshots.**

```
hostname:shares myproject/demo_share> snapshots
```

3. **Enter list to view the list of snapshots for the project or share.**

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
```

4. **Select the snapshot you want to restore, then enter the rollback command.**

```
hostname:shares myproject/demo_share snapshots> select demo_snap1
hostname:shares myproject/demo_share@demo_snap1> rollback
```

5. **Type y to confirm.**

```
hostname:shares myproject/demo_share@demo_snap1> rollback
Rolling back will revert data to snapshot, destroying newer data. Active initiators will
be disconnected.
```

```
Continue? (Y/N)
hostname: shares myproject/demo_share@demo_snap1> Y
```

▼ Destroying a Snapshot (BUI)

Use the following procedure to destroy a snapshot.

- Before You Begin**
- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
 - `destroySnap` - Allows destroying snapshots
 - `destroy` - Allows destroying projects and shares, including snapshot names
 - To add authorizations to a role, see [“Editing Authorizations for a Role \(BUI\)” on page 214](#).

1. **Go to the snapshot.**

- For a snapshot of a filesystem or LUN, go to Shares > Shares and click either Filesystems or LUNs, depending on whether the snapshot you want to destroy is of a filesystem or LUN.
 - For a snapshot of a project, go to Shares > Projects.
2. Hover over the appropriate share and click the edit icon .
 3. Click the Snapshots tab.
 4. Hover over the snapshot you want to destroy and click the destroy icon .

A confirmation dialog box appears.

If clones have been made of this snapshot, you are prompted with a list of the clones that will be affected. Destroying a snapshot also destroys any clones of that snapshot and descendents of those clones.
 5. Click OK to confirm.

Related Topics

- [“Understanding Users and Roles” on page 218](#)
- [“User Authorizations” on page 219](#)

▼ Destroying a Snapshot (CLI)

Use the following procedure to destroy a snapshot.

- Before You Begin**
- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the Projects and shares scope:
 - `destroySnap` - Allows users to only destroy snapshots.
 - `destroy` - Grants privileges to remove projects and shares, including snapshots.
 - To add authorizations to a role, see [“Editing Authorizations for a Role \(CLI\)” on page 215](#).
1. Go to shares and select the project, or select the project and then a share.


```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
```
 2. Enter snapshots.

```
hostname:shares myproject/demo_share> snapshots
```

3. **Enter `list` to view the list of snapshots for the project or share.**

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
```

4. **Use the `destroy` command to delete an individual snapshot using one of two methods:**

- Select the snapshot you want to delete, then type `destroy`.

```
hostname:shares myproject/demo_share snapshots> select demo_snap1
hostname:shares myproject/demo_share@demo_snap1> destroy
```

- Type **`destroy`** followed by the snapshot name.

```
hostname:shares myproject/demo_share snapshots> destroy demo_snap1
```

5. **Type `y` to confirm your action.**

```
This will destroy all data in "demo_snap1"! Are you sure? (Y/N) Y
```

Related Topics

- [“Understanding Users and Roles” on page 218](#)
- [“User Authorizations” on page 219](#)

▼ Cloning a Snapshot (BUI)

Note - Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing snapshot of a filesystem or LUN.

Note - Clones of projects are not supported.

1. **Go to the share you want to clone.**
 - a. **To clone a filesystem, go to `Shares > Shares`.**

- Snapshots and Clones 505

If the box is unchecked, the keystore and keyname of the clone will be that of the parent share. Alternatively, select a different keystore and keyname from the drop-down menu.

- f. **(Optional) Check the Retain Other Local Settings checkbox to cause any inherited properties to be preserved as local settings in the new clone.**

This field determines whether inherited properties will come from the parent dataset or the destination project. By default, the box is unchecked, meaning that all inherited properties will come from the destination project for the new clone. If you check the box, all currently inherited properties will be preserved as local settings in the new clone.

6. **Click APPLY to confirm the settings and create the clone.**

The clone appears in the list of shares for the destination project you set. You can work with a clone just like any other share.

Related Topics

- To perform share operations on a clone, see [“Shares and Projects” on page 383](#).
- To make a clone of a clone, see [“Cloning a Clone” on page 508](#).
- To view all the clones of a particular snapshot, see [“Viewing Clones of a Snapshot \(BUI\)” on page 508](#).
- To determine the snapshot from which a clone was made, see [“Viewing a Clone Origin \(BUI\)” on page 509](#).
- To make a clone of a snapshot in a replication package, see [“Cloning a Snapshot in a Replication Package \(BUI\)” on page 567](#)

▼ Cloning a Snapshot (CLI)

Note - Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing snapshot of a filesystem or LUN.

Note - Clones of projects are not supported.

1. **Go to the appropriate filesystem, LUN, or project and type snapshots.**

```
hostname:shares myproject/demo_share> snapshots
hostname:shares myproject/demo_share snapshots>
```

2. Select the snapshot you want to clone.

```
hostname:shares myproject/demo_share snapshots> select snap1
```

3. Use the `clone` command, optionally followed by the name of the project in which you want to create the clone.

By default, the clone is created in the same project as the snapshot being cloned.

```
hostname:shares myproject/demo_share@snap1> clone project1
```

You are placed into an uncommitted share context. From here, you can adjust properties as needed before committing the changes to create the clone.

4. Use the `get` command to view properties.

```
hostname:shares myproject/demo_clone (uncommitted clone)> get
aclinherit = restricted (inherited)
aclmode = discard (inherited)
atime = true (inherited)
checksum = fletcher4 (inherited)
compression = off (inherited)
copies = 1 (inherited)
mountpoint = /export/testbed (inherited)
quota = 0 (default)
readonly = false (inherited)
recordsize = 128K (inherited)
reservation = 0 (default)
secondarycache = all (inherited)
nbmand = false (inherited)
sharesmb = off (inherited)
sharenfs = on (inherited)
snapdir = hidden (inherited)
vscan = false (inherited)
sharedav = off (inherited)
shareftp = off (inherited)
root_group = other (default)
root_permissions = 777 (default)
root_user = nobody (default)
quota_snap = true (default)
reservation_snap = true (default)
```

5. Use the `set` command to adjust properties.

```
hostname:shares myproject/demo_clone (uncommitted clone)> set quota=10G
quota = 10G (uncommitted)
```

6. Use the `commit` command to commit the changes and create the clone.

```
hostname:shares myproject/demo_clone (uncommitted clone)> commit  
hostname:shares myproject/demo_share@demo_clone>
```

Related Topics

- To perform share operations on a clone, see [“Shares and Projects” on page 383](#).
- To make a clone of a clone, see [“Cloning a Clone” on page 508](#).
- To view all the clones of a particular snapshot, see [“Viewing Clones of a Snapshot \(CLI\)” on page 509](#).
- To determine the snapshot from which a clone was made, see [“Viewing a Clone Origin \(CLI\)” on page 510](#).
- To make a clone of a snapshot in a replication package, see [“Cloning a Snapshot in a Replication Package \(CLI\)” on page 570](#)

▼ Cloning a Clone

Note - Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing clone.

Once you have created a clone from a snapshot of a filesystem or LUN, you can work with that clone as with any other share. You can take a snapshot of the clone, and you can make a clone of that snapshot. You can continue this process to make clones of clones indefinitely.

1. **Take a snapshot of the clone using one of these procedures:**
 - [“Taking a Snapshot \(BUI\)” on page 484](#)
 - [“Taking a Snapshot \(CLI\)” on page 485](#)
2. **Clone the snapshot using one of these procedures:**
 - [“Cloning a Snapshot \(BUI\)” on page 504](#)
 - [“Cloning a Snapshot \(CLI\)” on page 506](#)

▼ Viewing Clones of a Snapshot (BUI)

Use the following procedure to view a list of all clones created from a particular snapshot. These are also called the "dependent clones" of the snapshot.

1. **Go to Shares > Shares and click either Filesystems or LUNs, depending on whether you want to view clones of a filesystem or a LUN.**
2. **Hover over the appropriate share and click the Edit icon .**
3. **Click the Snapshots tab.**
4. **Hover over the appropriate snapshot and click "Show..." under Clones.**
A window appears with a list of the snapshot's dependent clones and the projects in which they are located.
If the "Show..." link does not appear, the snapshot has no clones.
5. **Click OK to close the window.**

▼ Viewing Clones of a Snapshot (CLI)

Use the following procedure to view a list of all clones created from a particular snapshot.

1. **Go to and select the snapshot.**

```
hostname:shares myproject/demo_share> snapshots
hostname:shares myproject/demo_share snapshots> select snap1
hostname:shares myproject/demo_share@snap1>
```

2. **Use the `list clones` command.**

```
hostname:shares myproject/demo_share@snap1> list clones

Clones: 2 total

PROJECT      SHARE
myproject    demo_clone1
myproject    demo_clone2
hostname:shares myproject/demo_share@snap1>
```

The result shows how many clones exist, the project in which each resides, and the name of each clone.

▼ Viewing a Clone Origin (BUI)

Use the following procedure to determine the snapshot from which a clone was made.

1. **Go to the clone.**
 - a. **Go to Shares > Shares.**
 - b. **Hover over the clone and click the Edit icon .**
2. **Under Static Properties on the left, click "Show" next to "Clone origin".**

A window appears giving the name of the snapshot from which the clone was made.

▼ Viewing a Clone Origin (CLI)

Use the following procedure to determine the snapshot from which a clone was made.

1. **Go to shares and select the project that contains the clone, then select the clone.**

```
hostname:> shares select myproject
hostname:shares myproject> select demo_clone
hostname:shares myproject/demo_clone>
```

2. **Use the `get origin` command.**

The command returns the location and name of the snapshot from which the clone was made.

```
hostname:shares myproject/demo_clone> get origin
origin = myproject/demo_share@demo_snapshot
```

Remote Replication

Note - Replication and Cloning are licensed features for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Oracle ZFS Storage Appliance supports snapshot-based replication of projects and shares from a source appliance to a replication target, to a different pool on the same appliance, or to an NFS server for offline replication. You can configure replication to be executed manually, on a schedule, or continuously. Use cases for remote replication include disaster recovery, data distribution, disk-to-disk backup, and data migration between appliances when upgrading hardware or rebalancing storage.

To configure, monitor, and manage remote replication, use the following tasks:

- [“Remote Replication Workflow” on page 511](#)
- [“Configuring Remote Replication” on page 512](#)
- [“Monitoring Remote Replication” on page 548](#)
- [“Managing Replication Packages” on page 556](#)
- [“Disaster Recovery with Remote Replication” on page 580](#)

For details about remote replication, see:

- [“Remote Replication Concepts” on page 588](#)

▼ Remote Replication Workflow

The following steps summarize the basic steps for using remote replication. For information about remote replication concepts, see [“Remote Replication Concepts” on page 588](#).

- 1. Check software compatibility on source and target appliances.**

For information about software compatibility, see [“Checking Source and Target Compatibility” on page 513](#).

2. Set up network interfaces and routing.

For information about setting up network routing, see [Setting Up Network Interfaces and Static Routing - BUI, CLI](#).

3. Create a replication target.

For information about creating a replication target, see [Creating a Replication Target - BUI, CLI](#).

4. Create a replication action.

For information about creating a replication action, see [Creating a Replication Action - BUI, CLI](#).

5. Send a replication update as specified by the replication action.

Replication updates occur at a specified frequency, on a continuous basis, or manually. See [“Replication Update Frequency” on page 592](#).

6. Optionally, configure offline replication.

For information about offline replication, see [Configuring Offline Replication - BUI, CLI](#).

Configuring Remote Replication

Use the following tasks to configure remote replication:

- [“Checking Source and Target Compatibility” on page 513](#)
- [Setting Up Network Interfaces and Static Routing - BUI, CLI](#)
- [Creating a Replication Target - BUI, CLI](#)
- [Creating a Replication Action - BUI, CLI](#)
- [Configuring Auto Snapshot Management on Target - BUI, CLI](#)
- [Manually Sending a Replication Update - BUI, CLI](#)
- [“Configuring Replication for a Clustered Configuration” on page 526](#)
- [Configuring Offline Replication - BUI, CLI](#)
- [Disabling Replication Compression - BUI, CLI](#)
- [Editing a Replication Target - BUI, CLI](#)
- [Editing a Replication Action - BUI, CLI](#)

▼ Checking Source and Target Compatibility

Remote Replication is compatible between most Oracle ZFS Storage Appliance software versions. Compatibility failures are caused if a replication update uses a feature that is not supported on the replication target. Features are delivered with software updates or as deferred updates.

For details about compatibility and deferred update features for each software version, see the Oracle ZFS Storage Appliance Remote Replication Compatibility document (Doc ID 1958039.1) on [My Oracle Support \(http://support.oracle.com/\)](http://support.oracle.com/).

1. **Check the current software version on the source and target appliances.**
If you are using the BUI, go to Maintenance > System. If you are using the CLI, navigate to `maintenance system updates` and enter `show`.
2. **Ensure the replication target provides support for any deferred update feature used by the source project or share.**
For example, if the source share uses large blocks, ensure that the replication target provides support for this feature.
3. **Update software and apply deferred updates on the replication target, as needed.**
For more information, see “Deferred Updates” in *Oracle ZFS Storage Appliance Customer Service Manual*.


▼ Setting Up Network Interfaces and Static Routing (BUI)

To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure IPv4 static /32 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done. The two source cluster nodes can replicate to the same target node only if the target node provides two IPv4 addresses, one for use by each node in the source cluster. Replicating to the same target IPv4 address from both nodes of a source cluster is not supported.

1. **Go to the Configuration > Network > Routing page.**

2. Click the add icon .



Insert Static Route [CANCEL] [ADD]

Family ☒ IPv4
☐ IPv6

Kind ☐ Default
☒ Network

Destination

Gateway

Interface ▼

3. In the Insert Static Route screen, specify the following:

- **Family:** select IPv4.

Note - Remote replication does not support IPv6.

- **Kind:** select Network.
- **Destination:** enter the IPv4 address and netmask /32 of the target appliance.
- **Gateway:** enter the gateway address of the target appliance.
- **Interface:** enter the interface name.

4. Click Add.
5. After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.
6. To verify that traffic is routed through the correct source and target interfaces, use the `traceroute` command.

For information about using `traceroute`, see [“Configuring Network Routing” on page 119](#).

Note - When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- [“Example: Replication Configuration for Clustered Appliances” on page 604](#)
- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Setting Up Network Interfaces and Static Routing (CLI)

To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure IPv4 static /32 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done. The two source cluster nodes can replicate to the same target node only if the target node provides two IPv4 addresses, one for use by each node in the source cluster. Replicating to the same target IPv4 address from both nodes of a source cluster is not supported.

1. Navigate to configuration services routing on the source appliance.

Use a static /32 (host-specific) route to the target system IPv4 address via the dedicated network interface. In the following example, mask=32 means this is a host-specific route.

Note - Remote replication does not support IPv6.

```
host_source:configuration services routing> create

host_source:configuration services route (uncommitted)> get
    family = (unset)
    destination = (unset)
    mask = (unset)
    gateway = (unset)
    interface = (unset)
host_source:configuration services route (uncommitted)> set family=IPv4
host_source:configuration services route (uncommitted)> set destination=203.34.56.78
host_source:configuration services route (uncommitted)> set mask=32
host_source:configuration services route (uncommitted)> set gateway=203.34.56.254
```

```

host_source:configuration services route (uncommitted)> set interface=nge3
host_source:configuration services route (uncommitted)> commit
host_source:configuration services routing> show
route-000  0.0.0.0/0                203.24.30.254  nge0          static
route-001  203.24.30.0/32            203.24.30.28   nge0          dynamic
route-002  203.24.150.0/32                  203.24.150.10  ibd0          dynamic
route-003  203.24.101.65/32                 203.24.30.254  nge1          inactive
route-005  203.34.56.78/32                  203.34.56.254  nge3          static

```

2. **After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.**
3. **To ensure traffic is routed through the correct source and target interfaces, use the `traceroute` command.**

For information about using `traceroute`, see [“Configuring Network Routing” on page 119](#).

Note - When an interface is deleted, all routes associated with the interface are also removed.

Related Topics


- [“Example: Replication Configuration for Clustered Appliances” on page 604](#)
- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Creating a Replication Target (BUI)

A replication target establishes a secure communication connection between source and target appliances. To create a replication target:

1. **From the BUI of the source appliance, go to Configuration > Services > Remote Replication.**

2. Next to Targets, click the add icon .



3. Enter the following:

- **Name** - a name of the target to display in the BUI and CLI of the source appliance.
- **Hostname** - an IPv4 address, or host name, of the target appliance.

Note - Use an IPv4 address configured with a static route to force traffic over a specific network interface.

- **Root password** - the root password of the target appliance, to authorize the connection.

4. Click Add.

Related Topics

- [“Replication Targets” on page 590](#)
- [“Remote Replication Workflow” on page 511](#)
- [“Replicating an Encrypted Share” on page 656](#)

▼ Creating a Replication Target (CLI)

A replication target establishes a secure communication connection between replication source and target. To create a replication target:

1. **From the source appliance, navigate to the configuration services replication targets node, and enter target to configure a remote appliance as a replication target.**

```
host_source:> configuration services replication targets> target
host_source:configuration services replication target (uncommitted)>
```

2. **Set the target hostname, root_password, and label.**

- hostname - Target appliance host name or IPv4 address
- password - Target appliance password
- label - Target appliance name to display in the source appliance's BUI and CLI

Note - To force traffic over a specific network interface, use an IPv4 address configured with a static route.

```
host_source:configuration services replication target (uncommitted)> set hostname=203.123.225.201
host_source:configuration services replication target (uncommitted)> set root_password=password
host_source:configuration services replication target (uncommitted)> set label=repl_1
```

3. **Commit the changes.**

```
host_source:configuration services replication target (uncommitted)> commit
```

Related Topics

- [“Replication Targets” on page 590](#)
- [“Remote Replication Workflow” on page 511](#)
- [“Replicating an Encrypted Share” on page 656](#)

▼ Creating a Replication Action (BUI)

A replication action describes the project or share to be replicated, where to send the replication, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link.

- If you are setting up remote replication for the first time, it may be useful to replicate a minimal amount of data, by either replicating an empty project or by choosing not to replicate the snapshots in the project/shares.

- If you are replicating a large data set and bandwidth is limited due to distance between source and target appliances, you can export the replication to offline media, as described in [“Configuring Offline Replication \(BUI\)” on page 527](#).

1. From the BUI of the source appliance, go to **Shares > Projects**.
2. Select the project or share, and click the **Replication** tab.
3. Next to **Actions**, click the add icon .
4. Select a target and a pool.

Add Replication Action
CANCEL
ADD

Properties

Target ar7320-230

Pool Storage_Pool_230

Export data path ☐ nfs://

Limit bandwidth ☐ 0 M/s

Enable SSL-encryption ☒

Disable compression ☐

Enable deduplication ☐

Include snapshots ☒

Retain user snapshots on target ☐

Include clone origin as data ☐

Recovery point objective ☐ hours

Replica lag warning alert ☐ % of Recovery Point Objective

Replica lag error alert ☐ % of Recovery Point Objective

Update frequency ☒ Scheduled ☐ Continuous

 **Schedule** ... **Snapshots**

No schedule entries are configured for this action.

5. Select properties for this action.
See [“Replication Action Properties” on page 594](#) for a description of all properties.
6. Select **Scheduled**, and set a frequency for the replication update, or select **Continuous** to send replication updates continuously.

7. Click Add.

The replication action is added to the Actions list.

+ Actions		
TARGET ▲	UPDATES	STATUS
 sashimi Manual	Never Synced Never Attempted	 Sync now
 sushi Scheduled	2015-7-1 19:52:03 Synced 2015-7-1 19:52:03 Attempted	 2015-7-1 20:05:00 Next

Related Topics

- [“Replication Action Properties” on page 594](#)
- [“Replication Actions and Packages” on page 591](#)
- [“Manually Sending a Replication Update \(BUI\)” on page 525](#)

▼ Creating a Replication Action (CLI)

A replication action describes the project or share to be replicated, the replication target, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link.

- If you are setting up remote replication for the first time, it may be useful to replicate a minimal amount of data to ensure the sync completes successfully. You can either replicate an empty project or choose not to replicate the snapshots in the project/shares.
- If you are replicating a large data set and bandwidth is limited due to distance between source and target appliances, you can export the replication as described in [“Configuring Offline Replication \(CLI\)” on page 532](#).

1. Navigate to the project or share, and enter action:

```
host_source:shares PROJECT1/SHARE1 replication> action
```

2. Display the properties.

```
host_source:shares PROJECT1/SHARE1 action (uncommitted)> get
Properties:
    target = (unset)
    pool = (unset)
    enabled = true
```

```

        continuous = false
        include_snaps = true
        max_bandwidth = unlimited
        use_ssl = true

```

3. Set the properties for this action.

See [“Replication Action Properties” on page 594](#) for a description of CLI properties.

```

host_source:shares PROJECT1/SHARE1 action (uncommitted)> set target=repl_sys
                        target = repl_sys (uncommitted)
host_source:shares PROJECT1/SHARE1 action (uncommitted)> set pool=pool-0
                        pool = pool-0 (uncommitted)
host_source:shares PROJECT1/SHARE1 action (uncommitted)> set include_snaps=false
                        include_snaps = false (uncommitted)
host_source:shares PROJECT1/SHARE1 action (uncommitted)> set use_ssl=false
                        use_ssl = false (uncommitted)
host_source:shares PROJECT1/SHARE1 action-000> schedule
host_source:shares PROJECT1/SHARE1 action-000 schedule (uncommitted)> set frequency=day
                        frequency = day (uncommitted)
host_source:shares PROJECT1/SHARE1 action-000 schedule (uncommitted)> set hour=23
                        hour = 23 (uncommitted)
host_source:shares PROJECT1/SHARE1 action-000 schedule (uncommitted)> set minute=05
                        minute = 05 (uncommitted)

```

4. Commit the new replication action.

```

host_source:shares PROJECT1/SHARE1 action (uncommitted)> commit

```

5. To view the properties of the newly created action, enter ls:

```

host_source:shares PROJECT1/SHARE1 replication> ls
Properties:
    inherited = false
Actions:
    TARGET      STATUS      NEXT
    action-000  repl_sys   idle       manual

host_source:shares PROJECT1/SHARE1 action-000> ls
Properties:
    id = a751dc0f-abcd-1234-6789-f5e8315eaffa
    target = repl_sys
    enabled = true
    continuous = false
    include_snaps = false
    max_bandwidth = unlimited
    use_ssl = false
    compression = on
    export_path =

```

```
state = idle
state_description = Idle (no update pending)
export_pending = false
offline = false
next_update = Wed Sep 01 2016 23:05:00 GMT+0000 (UTC)
last_sync = Wed Sep 01 2016 10:24:05 GMT+0000 (UTC)
last_try = Wed Sep 01 2016 10:24:05 GMT+0000 (UTC)
last_result = success
```

6. **To view the ID of the newly created action, use the `last` command, which navigates to the node with the new action, combined with `get id`, which retrieves the action ID.**

The ID is used later to select the correct replication action node.

```
host_source:shares PROJECT1/SHARE1 replication>last get id
id =
fb1bb3fd-3361-42e1-e4a1-b06c426172fb
```

Related Topics

- [“Replication Action Properties” on page 594](#)
- [“Replication Actions and Packages” on page 591](#)

▼ Configuring Automatic Snapshot Retention on a Target (BUI)

Use this procedure to set a different number of retained automatic snapshots on the replication target than what was set on the source appliance. Before performing this procedure, you must set a replication action on a project or share, and schedule automatic snapshots.

When performing this task, modify the replication action and the snapshot schedule from the appropriate project or share. If the replication action and snapshot schedule are set at different levels, edit the replication action at the same level where the schedule is configured, as shown in the following table:

Action	Snapshot Schedule	Modify automatic snapshot retention on:
Project level	Project level	Action configured at the project level.
Share level	Share level	Action configured at the share level.
Project level	Share level	Action visible at the share level. (Project-level action inherited by share.)

Action	Snapshot Schedule	Modify automatic snapshot retention on:
Share level	Project level	Action configured at the share level.

See [“Replication Automatic Snapshot Management” on page 614](#) for more information.

Before You Begin Before configuring automatic snapshot retention on a target, you must first do the following:

- Create a replication action on a project or share: [“Creating a Replication Action \(BUI\)” on page 518](#).
- Create an automatic snapshot schedule for the project or share: [“Scheduling Snapshots \(BUI\)” on page 485](#).

1. **On the source appliance, go to Shares > Projects.**
2. **Select the project or share with the replication action and snapshot schedule, and click the Replication tab.**
3. **Click the edit icon  of the action you want to modify.**
4. **At the bottom of the Edit Replication Action window, click the Snapshots tab.**
The automatic snapshot schedules appear.
5. **In the field Keep At Most, change the value to specify the number of automatic snapshots to be retained on the target.**
6. **Click Apply.**

Note - Automatic snapshot retention for replication has special processing during reverse replication. For more information, see [“Replication Snapshot Management” on page 612](#).

▼ Configuring Automatic Snapshot Retention on a Target (CLI)

Use this procedure to set a different number of retained automatic snapshots on the replication target than what was set on the source appliance. Before performing this procedure, you must set a replication action on a project or share, and schedule automatic snapshots.

When performing this task, modify the replication action and the snapshot schedule from the appropriate project or share. If the replication action and snapshot schedule are set at different

levels, edit the replication action at the same level where the schedule is configured, as shown in the following table:

Action	Snapshot Schedule	Modify automatic snapshot retention on:
Project level	Project level	Action configured at the project level.
Share level	Share level	Action configured at the share level.
Project level	Share level	Action visible at the share level. (Project-level action inherited by share.)
Share level	Project level	Action configured at the share level.

See [“Replication Automatic Snapshot Management” on page 614](#) for more information.

Before You Begin Before configuring automatic snapshot retention on a target, you must first do the following:

- Create a replication action on a project or share: [“Creating a Replication Action \(BUI\)” on page 518](#).
- Create an automatic snapshot schedule for the project or share: [“Scheduling Snapshots \(BUI\)” on page 485](#).

1. **On the source appliance, go to Shares and select the appropriate project or share with the replication action and snapshot schedule.**

```
hostname:> shares select MyProject
hostname:shares MyProject> select MyShare
hostname:shares MyProject/MyShare>
```

2. **Enter replication, then enter show to display existing actions.**

```
hostname:shares MyProject/MyShare> replication
hostname:shares MyProject/MyShare replication> show
Properties:
            inherited = false

Actions:

            TARGET      STATUS      NEXT
action-000  local      idle       Thu Jan 05 2017 12:04:00 GMT+0000 (UTC)
```

3. **Select the action for which you want to modify retention settings.**

```
hostname:shares MyProject/MyShare replication> select action-000
hostname:shares MyProject/MyShare action-000>
```

4. **Enter autosnaps, then enter show to display snapshot schedules.**


```
hostname:shares MyProject/MyShare action-000> autosnaps
hostname:shares MyProject/MyShare action-000 autosnaps> show
Properties:
  autosnaps_retention_policies = independent

Automatics:

NAME          FREQUENCY    DAY          HH:MM KEEP
automatic-000 hour        -            -:04      3
```

5. **Select the automatic snapshot schedule that you want to modify. Then use `set keep=` to set the number of automatic snapshots to be retained on the target.**

```
hostname:shares MyProject/MyShare action-000 autosnaps> select automatic-000
hostname:shares MyProject/MyShare action-000 automatic-000> set keep=10
      keep = 10 (uncommitted)
```

6. **Enter `commit` to save the changes, and enter `show` to verify that the `keep` property has changed.**

```
hostname:shares MyProject/MyShare action-000 automatic-000> commit
hostname:shares MyProject/MyShare action-000 automatic-000> show
Properties:
      frequency = hour
      day =
      hour =
      minute = 04
      keep = 10
```

Note - Automatic snapshot retention for replication has special processing during reverse replication. For more information, see [“Replication Snapshot Management” on page 612](#).

▼ Manually Sending a Replication Update (BUI)

If continuous or scheduled replication is already configured, replication updates are performed automatically. You can also perform a manual update using the BUI.

1. **From the source appliance, go to `Shares > Projects`.**
2. **Open a project, and click the `Replication` tab.**
3. **Click the `Sync now` icon .**

Note - This action is not available (or will not work) if an update is actively being sent. Ensure there is enough disk space on the target to replicate the entire project before sending an update.

The BUI displays a progress bar and indicates when the update completes.

4. **If the replication update does not complete successfully, remove any old actions or snapshots and start it again.**

Related Topics

- [“Creating a Replication Action \(BUI\)” on page 518](#)
- [“Canceling a Replication Update \(BUI\)” on page 558](#)

▼ Manually Sending a Replication Update (CLI)

If continuous or scheduled replication is already configured, replication updates are performed automatically. You can also perform a manual update using the CLI.

1. **Navigate to the share and enter the `sendupdate` command:**

```
host_source:shares PROJECT1/SHARE1 action-000> sendupdate
```

When the update is currently active, the CLI shows a state of sending.

2. **If the replication does not complete successfully, remove any old actions or snapshots and start it again.**

Related Topics

- [“Creating a Replication Action \(CLI\)” on page 520](#)
- [“Canceling a Replication Update \(CLI\)” on page 559](#)

▼ Configuring Replication for a Clustered Configuration

This task describes how to configure replication in a clustered environment. Follow these steps to configure replication properly to ensure that projects continue to replicate after a cluster takeover, cluster failback, or after performing reverse replication on a target appliance.

Before You Begin If you are configuring replication for clustered appliances for the first time, review [“Example: Replication Configuration for Clustered Appliances”](#) on page 604.

1. **On the replication source and target appliances, select network interfaces and IP addresses to be used for replication traffic, using these guidelines:**
 - a. **Always select a singleton network interface to ensure it is taken over by the peer node after a cluster takeover or failback operation.**
 - b. **On the source system, ensure that the selected network interface and the storage pool, from which the data will be replicated, are both assigned to the same node. This is always the case when the source cluster is in the CLUSTERED state.**
 - c. **On the target system, assign the selected network interface on the target appliance and the storage pool, into which the replicated data will be put, to the same node. This maintains the association when the replication configuration is performed while the target cluster is in the CLUSTERED state.**
 - d. **Ensure that the source and the target systems can communicate using the selected network interfaces and IP addresses.**
2. **On the source and target appliances, create static /32 (host-based) network routes using the selected network interfaces and IP addresses.**
3. **On the source appliance, configure the replication target object using the selected IP address of the target.**

Configuring Offline Replication (BUI)

Use the following steps to configure offline replication:

- [“Setting Up an NFS Server for Offline Replication”](#) on page 528
- [“Setting Up an Export Path to the NFS Server \(BUI\)”](#) on page 528
- [“Exporting a Replication Update \(BUI\)”](#) on page 529
- [“Verifying the Replication Stream On the NFS Server”](#) on page 529
- [“Importing the Replication Stream from the NFS Server \(BUI\)”](#) on page 530
- [“Performing a Manual Network Update \(BUI\)”](#) on page 530
- [“Reversing an Offline Replication \(BUI\)”](#) on page 531

▼ Setting Up an NFS Server for Offline Replication

The steps for setting up an NFS server will vary depending on the NFS server type you use. Refer to your NFS server documentation for specific instructions.

1. **Identify a server that is network ready and has NFS Services enabled.**
2. **As root of the NFS server, create a filesystem or share.**
3. **Set the file permissions to expose the NFS share only to the IP address of the source and target appliances.**
4. **To encrypt the replication stream, enable on-disk encryption for the NFS share on the NFS server.**

Note - An exported replication stream is never encrypted by the appliance.

5. **Export the share for access by the NFS client.**
6. **Verify that the filesystem is shared.**

Next Steps

- [“Setting Up an Export Path to the NFS Server \(BUI\)” on page 528](#)

▼ Setting Up an Export Path to the NFS Server (BUI)

Before You Begin Identify or create a target, see [“Creating a Replication Target \(BUI\)” on page 516](#).

1. **From the BUI of the source appliance, go to Shares > Projects.**
2. **Open the project, and click the Replication tab.**
3. **Click the add icon  next to Actions.**
4. **In the Add Replication Action screen, select Export data path and enter the path of the NFS share in the form: `nfs://server/path`.**
5. **Select additional properties for this action, and then click Add.**

Note - If you configure a schedule or select continuous replication mode, the update will occur automatically after the export and the import operations have completed.

Next Steps

- [“Exporting a Replication Update \(BUI\)” on page 529](#)

▼ Exporting a Replication Update (BUI)

1. From the source appliance, go to Shares > Projects.
2. Open the project, and click the Replication tab.
3. Click the Export replication data icon .
4. Check the replication status, and wait until the replication completes.

Next Steps

- [“Verifying the Replication Stream On the NFS Server” on page 529](#)

▼ Verifying the Replication Stream On the NFS Server

1. Navigate to the NFS directory, check the MD5, and view the metadata.

```
bigfish25# pwd
/export/init_repl/rr_updates/96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
bigfish25# ls -l
total 67
-rw-r--r--  1 nobody  nobody      633 Nov 17 21:46 metadata.xml
-rw-----  1 nobody  nobody    31016 Nov 17 21:46 stream
-rw-----  1 nobody  nobody      33 Nov 17 21:46 stream.md5
bigfish25# md5sum stream
25b4671c9aaf34455a63e203bcecff49  stream
bigfish25# cat stream.md5
25b4671c9aaf34455a63e203bcecff49
bigfish25# cat metadata.xml
<?xml version="1.0"?>
<!DOCTYPE nvlist SYSTEM "/usr/share/lib/xml/dtd/nvlist.dtd.1">
<nvlist>
  <nvpair name='offline_rr_version'><string value='1.1' /></nvpair>
  <nvpair name='source_asn'><string value='2ea4670f-bc17-cf8f-a420-9211d6edda04' /></
nvpair>
```

```
<nvpair name='project'><string value='default' /></nvpair>
<nvpair name='pkgid'><string value='96366bf2-0b3c-4eec-e85b-e36e1b5bc18c' /></nvpair>
<nvpair name='basesnap'><string value='/'></nvpair>
<nvpair name='newsnap'><string value='.rr-96366bf2-0b3c-4eec-e85b-e36e1b5bc18c-1' /></nvpair>
nvpair>
  <nvpair name='compression'><string value='on' /></nvpair>
</nvlist>
bigfish25#
```

2. **Physically move the NFS server to the target appliance site, or copy the rr_updates folder to external media and prepare for shipping.**

Next Steps

- [“Importing the Replication Stream from the NFS Server \(BUI\)” on page 530](#)

▼ Importing the Replication Stream from the NFS Server (BUI)

1. **Go to Shares > Projects > Replicas.**
2. **Select the replica that shows source: awaiting import.**
3. **Click the Replication tab.**
4. **In the Import Data Path field, enter the path of the replica.**
5. **Click the Import update from external media icon  to start the import.**

Next Steps

After the replication stream is imported to the target appliance, proceed with one of the following:

- [“Performing a Manual Network Update \(BUI\)” on page 530](#)
- [“Reversing an Offline Replication \(BUI\)” on page 531](#)

▼ Performing a Manual Network Update (BUI)

After importing the offline replication stream to the target appliance, confirm future network updates will work correctly. If continuous or scheduled replication is already configured, the update will be performed automatically. Otherwise, perform a manual update.

1. Go to the source appliance.
2. See [“Manually Sending a Replication Update \(BUI\)” on page 525](#).

▼ Reversing an Offline Replication (BUI)

Follow this procedure to move an offline replication package to a new local project, configured to replicate back to a source appliance.

1. Import the offline replication package from an NFS server to the target appliance, as described in [“Importing the Replication Stream from the NFS Server \(BUI\)” on page 530](#).
2. From the target appliance, go to **Shares > Projects > Replica** and locate the replicated package.

The project is named *target_appliance: new_project/share*.

3. Select the project and click its reverse replication direction icon .
4. In the Reverse Replication window, enter a name for the new local project.
This action moves the contents of this package to a new local project configured to replicate back to the source. Any data or metadata changes made on the source since the last successful update will be lost when the new project is replicated back to the source. If replication actions on the source are not disabled, future updates to this package will fail.
5. Go to **Shares > Projects**.
6. Open the project and click the **Replication** tab.
7. Click the **Export replication data** icon .
8. Check the replication status, and wait until the replication completes.
9. After the replication update is complete, navigate to the newly reversed package on the new target.

The state description should be *Idle (awaiting import)*.

10. Import the update from the NFS server.

Related Topics

- [“Configuring Offline Replication \(BUI\)” on page 527](#)

Configuring Offline Replication (CLI)

Use the following steps to configure offline replication:

- [“Setting Up an NFS Server for Offline Replication” on page 532](#)
- [“Setting Up an Export Path to the NFS Server \(CLI\)” on page 533](#)
- [“Exporting a Replication Update \(CLI\)” on page 534](#)
- [“Verifying a Replication Stream On the NFS Server” on page 535](#)
- [“Importing a Replication Stream from the NFS Server \(CLI\)” on page 536](#)
- [“Performing a Manual Network Update \(CLI\)” on page 537](#)
- [“Reversing an Offline Replication \(CLI\)” on page 539](#)

▼ Setting Up an NFS Server for Offline Replication

The steps for setting up an NFS server will vary depending on the NFS server type you use. Refer to your NFS server documentation for specific instructions.

1. **Identify a server that is network ready and has NFS Services enabled.**
2. **As root of the NFS server, create a filesystem or share.**
3. **Set the file permissions to expose the NFS share only to the IP address of the source and target appliances.**
4. **To encrypt the replication stream, enable on-disk encryption for the NFS share on the NFS server.**

Note - An exported replication stream is never encrypted by the appliance.

5. **Export the share for access by the NFS client.**
6. **Verify that the filesystem is shared.**

Next Steps

- [“Setting Up an Export Path to the NFS Server \(CLI\)” on page 533](#)

▼ Setting Up an Export Path to the NFS Server (CLI)

1. Identify or create a replication target.
2. Create a replication action, set the `export_path`, and commit the new action.

```
source:shares default replication> action
source:shares default action (uncommitted)> set target=target_a
      target = target_a (uncommitted)
source:shares default action (uncommitted)> set pool=pool2
      pool = pool2 (uncommitted)
source:shares default action (uncommitted)> set export_path=nfs://nfs_server/export/
init_repl
      export_path = nfs://nfs_server/export/init_repl (uncommitted)
source:shares default action (uncommitted)>commit
```

Note - Optionally, you can set a scheduled or continuous replication mode, which will start the update after the export and the import operations have completed.

3. Navigate back to the replication action that you just created, and view the current status.

```
source:shares default replication> ls
Actions:
      TARGET      STATUS      NEXT
action-000  target_a      idle      Export replication data
```

```
source:shares default replication> last
source:shares default action-000> ls
Properties:
      id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
      target = target_a
      enabled = true
      continuous = false
      include_snaps = true
      max_bandwidth = unlimited
      bytes_sent = 0
      estimated_size = 0
      estimated_time_left = 00:00:00
      average_throughput = 0B/s
      use_ssl = true
      compression = on
      export_path = nfs://nfs_server/export/init_repl
      state = idle
      state_description = Idle (export pending)
      export_pending = true
```

```
offline = false
next_update = Export replication data
last_sync = <unknown>
last_try = <unknown>
last_result = <unknown>
```

Next Steps

- [“Exporting a Replication Update \(CLI\)” on page 534](#)

▼ Exporting a Replication Update (CLI)

1. To export the replication update to the NFS server, use the `sendupdate` command.

```
source:shares default action-000>sendupdate
```

2. Enter `ls` to view the status, as shown in this example:

```
source:shares default action-000> ls
Properties:
      id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
      target = target_a
      enabled = true
      continuous = false
      include_snaps = true
      max_bandwidth = unlimited
      bytes_sent = 0
      estimated_size = 0
      estimated_time_left = 00:00:00
      average_throughput = 0B/s
      use_ssl = true
      compression = on
      export_path = nfs://nfs_server/export/init_repl
      state = sending
      state_description = Exporting update
      export_pending = true
      offline = false
      next_update = Export replication data
      last_sync = <unknown>
      last_try = <unknown>
      last_result = <unknown>
```

3. To determine when the export has completed, enter `ls` to view the status.

Look for `last_result=success`, as shown in this example:

```
source:shares default action-000> ls
```

Properties:

```

        id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
        target = target_a
        enabled = true
        continuous = false
        include_snaps = true
        max_bandwidth = unlimited
        bytes_sent = 0
        estimated_size = 0
        estimated_time_left = 00:00:00
        average_throughput = 0B/s
        use_ssl = true
        compression = on
        export_path =
        state = idle
        state_description = Idle (no update in progress)
        export_pending = false
        offline = true
        next_update = Sync now
        last_sync = <unknown>
        last_try = Tue Nov 18 2014 04:40:40 GMT+0000 (UTC)
        last_result = success
source:shares default action-000>

```

Next Steps

- [“Verifying a Replication Stream On the NFS Server” on page 535](#)

▼ Verifying a Replication Stream On the NFS Server

1. **Navigate to the NFS directory, check the MD5, and view the metadata.**

```

nfs_server# pwd
/export/init_repl/rr_updates/96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
nfs_server# ls -l
total 67
-rw-r--r--  1 nobody  nobody    633 Nov 17 21:46 metadata.xml
-rw-----  1 nobody  nobody   31016 Nov 17 21:46 stream
-rw-----  1 nobody  nobody    33 Nov 17 21:46 stream.md5
nfs_server# md5sum stream
25b4671c9aaf34455a63e203bcecff49  stream
nfs_server# cat stream.md5
25b4671c9aaf34455a63e203bcecff49
nfs_server# cat metadata.xml
<?xml version="1.0"?>
<!DOCTYPE nvlist SYSTEM "/usr/share/lib/xml/dtd/nvlist.dtd.1">
<nvlist>

```

```

    <nvpair name='offline_rr_version'><string value='1.1'/></nvpair>
    <nvpair name='source_asn'><string value='2ea4670f-bc17-cf8f-a420-9211d6edda04'/></
nvpair>
    <nvpair name='project'><string value='default'/></nvpair>
    <nvpair name='pkgid'><string value='96366bf2-0b3c-4eec-e85b-e36e1b5bc18c'/></nvpair>
    <nvpair name='basesnap'><string value='/></nvpair>
    <nvpair name='newsnap'><string value='.rr-96366bf2-0b3c-4eec-e85b-e36e1b5bc18c-1'/></
nvpair>
    <nvpair name='compression'><string value='on'/></nvpair>
</nvlist>
nfs_server#

```

2. **Physically move the NFS server to the target appliance site, or copy the `rr_updates` folder to external media and prepare for shipping.**

Next Steps

- [“Importing a Replication Stream from the NFS Server \(CLI\)” on page 536](#)

▼ Importing a Replication Stream from the NFS Server (CLI)

1. **To import the replication stream from the NFS server, navigate to `shares replication packages` on the target, and then enter `ls` to list the packages.**

```

target_a:> shares replication packages
target_a:shares replication packages> ls
Packages:

```

ID	STATE	DATA_TIMESTAMP	SOURCE	DATASET
package-000	idle	unknown	sourceA	<unknown>

2. **Select the package you want to import.**

To view the properties, enter `ls`.

```

target_a:shares replication packages> select package-000
target_a:shares replication package-000> ls
Properties:

```

```

    id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
    source_name = sourceA
    source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
    source_ip = 10.000.000.000:216
    target_pool = poolA
    replica_of = <unknown>
    enabled = true
    state = idle

```

```

state_description = Idle (no update in progress)
  offline = false
  import_path =
data_timestamp = unknown
  last_sync = unknown
  last_try = unknown
  last_result = unknown

```

3. Set the import path of the replicated data, and then enter **commit**.

```

target_a:shares replication package-000> set import_path=
nfs://nfs_server/export/init_repl
      import_path = nfs://nfs_server/export/init_repl (uncommitted)
target_a:shares replication package-000> commit
target_a:shares replication package-000> ls
Properties:

```

```

      id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
source_name = sourceA
source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
source_ip = 10.000.000.000:216
target_pool = poolA
replica_of = <unknown>
  enabled = true
  state = receiving
state_description = Importing update
  offline = true
  import_path = nfs://nfs_server/export/init_repl
data_timestamp = unknown
  last_sync = unknown
  last_try = unknown
  last_result = unknown

```

Next Steps

After the replication stream is imported to the target appliance, proceed with one of the following:

- [“Performing a Manual Network Update \(CLI\)” on page 537](#)
- [“Reversing an Offline Replication \(CLI\)” on page 539](#)

▼ Performing a Manual Network Update (CLI)

After importing the offline replication stream to the target appliance, confirm future network updates will work correctly. If continuous or scheduled replication is already configured, the

update will be performed automatically. Otherwise, perform a manual update as shown in the following example.

1. Go to the source appliance, and navigate to the share.

```
source:shares default action-000> ls
Properties:
    id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
    target = target_a
    enabled = true
    continuous = false
    include_snaps = true
    max_bandwidth = unlimited
    bytes_sent = 0
    estimated_size = 0
    estimated_time_left = 00:00:00
    average_throughput = 0B/s
    use_ssl = true
    compression = on
    export_path =
    state = idle
    state_description = Idle (no update in progress)
    export_pending = false
    offline = true
    next_update = Sync now
    last_sync = <unknown>
    last_try = Tue Nov 18 2014 04:40:40 GMT+0000 (UTC)
    last_result = success
```

2. Start the update using `sendupdate`, and then view the status using the `ls` command.

```
source:shares default action-000> sendupdate
source:shares default action-000> ls
Properties:
    id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
    target = target1
    enabled = true
    continuous = false
    include_snaps = true
    max_bandwidth = unlimited
    bytes_sent = 0
    estimated_size = 0
    estimated_time_left = 00:00:00
    average_throughput = 0B/s
    use_ssl = true
    compression = on
    export_path =
```

```

        state = sending
state_description = Ready (awaiting available resources to send update)
export_pending = false
    offline = true
next_update = Sync now
    last_sync = <unknown>
    last_try = Tue Nov 18 2014 04:40:40 GMT+0000 (UTC)
    last_result = success

source:shares default action-000> ls
Properties:
        id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
        target = target1
        enabled = true
        continuous = false
        include_snaps = true
        max_bandwidth = unlimited
        bytes_sent = 0
        estimated_size = 0
estimated_time_left = 00:00:00
average_throughput = 0B/s
        use_ssl = true
        compression = on
        export_path =
            state = idle
state_description = Idle (no update in progress)
export_pending = false
        offline = false
        next_update = Sync now
        last_sync = Tue Nov 18 2014 04:40:40 GMT+0000 (UTC)
        last_try = Tue Nov 18 2014 04:40:40 GMT+0000 (UTC)
        last_result = success

```

▼ Reversing an Offline Replication (CLI)

Follow this procedure to move an offline replication package to a new local project, configured to replicate back to a source appliance.

Before You Begin Import the offline replication stream from an NFS server to the replication target, as described in [“Importing a Replication Stream from the NFS Server \(CLI\)” on page 536](#).

1. **From the replication target, navigate to the replicated package and locate the project:**

```

target:> shares replication packages
target: shares replication packages> select package-000

```

```
target:shares replication package-000> ls
Properties:
```

```

        id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
source_name = bigfish78
source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
source_ip = 10.000.000.000:216
target_pool = poolA
replica_of = proj1
enabled = true
state = idle
state_description = Idle (no update in progress)
offline = false
import_path =
data_timestamp = Thu Feb 16 2017 19:10:59 GMT+0000 (UTC)
last_sync = Fri Feb 17 2017 03:10:11 GMT+0000 (UTC)
last_try = Fri Feb 17 2017 03:10:11 GMT+0000 (UTC)
last_result = success
```

```
Projects:
```

```
proj1
```

2. Enter pkgreverse.

```
target:shares replication package-000> pkgreverse
```

3. (Optional) Set a new project name and enable the action using the following commands:

```
target:shares replication package-000 pkgreverse> set new_project_name=new-kmm3
new_project_name = new-kmm3
target:shares replication package-000 pkgreverse> set enable_action_upon_reversal=true
enable_action_upon_reversal = true
```

Note - The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the appliance at the production site, the reverse operation will fail.

4. Enter show to confirm the properties, and then enter commit:

```
target:shares replication package-000 pkgreverse> show
Properties:
        new_project_name = new-kmm3
        enable_action_upon_reversal = true

host-prod:shares replication package-000 pkgreverse> commit
```


This action will move the contents of this package to a new local project configured to replicate back to the source. Any data or metadata changes made on the source since the last successful update will be lost when the new project is replicated back to the source. If replication actions on the source are not disabled, future updates to this package will fail.

5. Navigate to shares replication actions.

```
target:shares replication packages> cd /
target:> shares replication actions
```

6. Select the newly created action using the package ID from the previous steps.

Use the package ID as `origin_pkg_id` to select the action.

```
target:shares replication actions> select origin_pkg_id=
1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
target:shares replication action-000> ls
Properties:
        id = 6a10ce61-cc87-4850-89dd-8673f7734d03
        origin_pkg_id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
        target = new_target
        target_pool = p
        source_pool = p
        replication_of = dataset1
        enabled = true
        continuous = false
        include_snaps = false
        retain_user_snaps_on_target = false
        dedup = false
        include_clone_origin_as_data = false
        max_bandwidth = unlimited
        bytes_sent = 0
        estimated_size = 0
        estimated_time_left = 00:00:00
        average_throughput = 0B/s
        use_ssl = false
        compression = on
        export_path =
        state = idle
        state_description = Idle (no update in progress)
        export_pending = false
        offline = false
        next_update = Sync now
        replica_data_timestamp = Thu Feb 09 2017 16:17:25 GMT+0000 (UTC)
        last_sync = <unknown>
        last_try = <unknown>
        last_result = <unknown>
        replica_lag = 461:55:40
```

```

recovery_point_objective =
replica_lag_warning_alert =
replica_lag_error_alert =
replica_lag_over_warning_limit = false
replica_lag_over_error_limit = false

```

7. **To export the first replication update after reversal to an NFS server, enter `export_path` and the pathname of the NFS server. Enter `commit` and then enter `sendupdate`:**

```

target:shares replication action-000> set export_path=nfs://nfs_server/export/init_repl
      export_path = nfs://nfs_server/export/init_repl (uncommitted)
target:shares replication action-000> commit
target:shares replication action-000> sendupdate
target:shares replication action-000> ls
Properties:
      id = 6a10ce61-cc87-4850-89dd-8673f7734d03
      origin_pkg_id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
      target = new_target
      target_pool = p
      source_pool = p
      replication_of = dataset1
      enabled = true
      continuous = false
      include_snaps = false
      retain_user_snaps_on_target = false
      dedup = false
      include_clone_origin_as_data = false
      max_bandwidth = unlimited
      bytes_sent = 0
      estimated_size = 0
      estimated_time_left = 00:00:00
      average_throughput = 0B/s
      use_ssl = false
      compression = on
      export_path = nfs://nfs_server/export/init_repl
      state = idle
      state_description = Idle (no update in progress)
      export_pending = false
      offline = false
      next_update = Sync now
      replica_data_timestamp = Thu Feb 09 2017 16:17:25 GMT+0000 (UTC)
      last_sync = <unknown>
      last_try = <unknown>
      last_result = <unknown>
      replica_lag = 461:55:40
      recovery_point_objective =
      replica_lag_warning_alert =

```

```

        replica_lag_error_alert =
replica_lag_over_warning_limit = false
        replica_lag_over_error_limit = false

```

8. **After the replication update is complete, navigate to the newly reversed package on the new target. The state description should be Idle, as shown in the example:**

```

new_target:shares replication packages> ls
Packages:

```

ID	STATE	DATA_TIMESTAMP	SOURCE	DATASET
package-000	idle	unknown	target	<unknown>

9. **Select the package and enter `ls` to list its properties.**

```

new_target:shares replication packages> select package-000
new_target:shares replication package-000> ls

```

Properties:

```

        id = 6a10ce61-cc87-4850-89dd-8673f7734d03
        source_name = target
        source_asn = ddbd5d4e-daff-4f52-9417-cd6e893c694a
        source_ip = 00.000.00.000:216
        source_pool = poolA
        target_pool = poolA
        replica_of = <unknown>
        enabled = true
        state = idle
state_description = Idle (no update in progress)
        offline = true
        import_path =
data_timestamp = unknown
        last_sync = unknown
        last_try = unknown
        last_result = unknown

```

10. **Import the update from the NFS server.**

```

new_target:shares replication package-000> set import_path=nfs://nfs_server/export/
init_repl
        import_path = nfs://nfs_server/export/init_repl (uncommitted)

```

11. **Enter `commit` and then list the package properties to confirm the update has completed.**

The property `last_result` displays success.

```

new_target:shares replication package-000> commit

```

```
new_target:shares replication package-000> ls
Properties:
      id = 6a10ce61-cc87-4850-89dd-8673f7734d03
    source_name = target
    source_asn = ddbd5d4e-daff-4f52-9417-cd6e893c694a
    source_ip = 00.000.00.000:216
    source_pool = poolA
    target_pool = poolA
    replica_of = <unknown>
      enabled = true
      state = idle
state_description = Idle (no update in progress)
      offline = false
    import_path =
    data_timestamp = unknown
      last_sync = Fri Jul 31 2015 22:11:32 GMT+0000 (UTC)
      last_try = Fri Jul 31 2015 22:11:32 GMT+0000 (UTC)
    last_result = success
```

Related Topics

- [“Importing a Replication Stream from the NFS Server \(CLI\)” on page 536](#)
- [“Configuring Offline Replication \(CLI\)” on page 532](#)

▼ Disabling Replication Compression (BUI)

You can disable compression when you create or edit a replication action. By default, all replication streams are compressed before being sent over the network.

1. **From the source appliance, go to Shares > Projects > and double-click the project you want to edit.**
2. **Click the Replication tab.**
3. **Click the edit icon .**
4. **Click Disable compression, and click Apply.**

Related Topics

- [“Compressed Replication” on page 620](#)
- [“Remote Replication Workflow” on page 511](#)

▼ Disabling Replication Compression (CLI)

You can disable compression when you create or edit a replication action. By default, all replication streams are compressed before being sent over the network. For more information, see [“Compressed Replication” on page 620](#).

1. **To disable compression, navigate to the project or share and set the compression property, as shown in the following example:**

```
eel:shares proj1 action-000> set compression=off
```

2. **Enter commit and then show to confirm the compression property is set to off.**

```
eel:shares proj1 action-000> commit
eel:shares proj1 action-000> show
Properties:
    id = 67f0d3d6-10af-6f30-9d4c-a60d19eb1200
    target = goby-10g
    enabled = true
    continuous = false
    include_snaps = false
    max_bandwidth = unlimited
    bytes_sent = 0
    estimated_size = 0
    estimated_time_left = 00:14:35
    average_throughput = 0B/s
    use_ssl = false
    compression = off
    export_path =
    state = idle
    state_description = idle (no update in progress)
    export_pending = false
    offline = false
    next_update = Sync now
    last_sync = <unknown>
    last_try =
    last_result =
```

Related Topics

- [“Compressed Replication” on page 620](#)
- [“Remote Replication Workflow” on page 511](#)

▼ Editing a Replication Target (BUI)

1. **Go to Configuration > Services > Remote Replication > Targets.**
2. **For the target you want to edit, move the cursor over the target name, and click the edit icon .**
3. **Change the Name and/or Hostname.**
The hostname or IP address must resolve to the same appliance as before (checked by the serial number of the target).

Note - If you want to point to a different appliance than previously configured, you must create a new target to authenticate against the new appliance.

4. **Click Apply to save the changes.**

Related Topics

- [“Remote Replication Concepts” on page 588](#)
- [“Remote Replication Workflow” on page 511](#)

▼ Editing a Replication Target (CLI)

1. **Navigate to shares replication targets to set or unset the target hostname, root_password, and label.**

```
hostname:> shares replication targets
```

2. **From this context, you can:**
 - Add new targets.
 - View the actions configured with the existing target.
 - Edit the unique identifier (label) and/or hostname for the target.
 - Destroy a target, if no actions are using it.

Note - A target should not be destroyed while actions are using it. Such actions will be permanently broken. The system makes a best effort to enforce this but cannot guarantee that no actions exist in exported storage pools that are using a given target.

3. If the share you are replicating is encrypted, be sure the target also supports data encryption.

Related Topics

- [“Replicating an Encrypted Share” on page 656](#)
- [“Replication Targets” on page 590](#)

▼ Editing a Replication Action (BUI)

1. Navigate to the project or share, and click the Replication tab.
2. Select the project or share you want to edit.
3. Click the edit icon .
4. From the Edit Replication Action screen, modify the properties, and click Apply.
For a description of replication actions, see [“Replication Action Properties” on page 594](#).

Related Topics

- [“Replication Actions and Packages” on page 591](#)
- [“Remote Replication Workflow” on page 511](#)

▼ Editing a Replication Action (CLI)

1. Navigate to `shares replication actions`, then enter `ls` to list available actions.

```
hostname:> shares replication actions
hostname:shares replication actions> ls
Actions:
```

ID	STATE	REPLICA_DATA_TSTAMP	TARGET	DATASET
action-007	idle	2017-02-17 23:01:19	targetA	berries
action-008	idle	2017-02-17 23:01:40	targetA	cherries
action-003	idle	2017-02-15 23:48:15	targetA	ocean
action-002	disbl	<unknown>	targetA	oceanR
action-004	idle	<unknown>	targetA	berries

2. Select the action you want to edit

```
hostname:shares replication actions> select action-007
hostname:shares replication action-007>
```

3. **Modify the properties as necessary, using the set command.**

For a list of replication action commands for the CLI, see [“Replication Action Properties” on page 594](#).

Related Topics

- [“Creating a Replication Action \(CLI\)” on page 520](#)
- [“Remote Replication Workflow” on page 511](#)

Monitoring Remote Replication

Use the following tasks to monitor replication progress, alerts and events, and replication delays. To investigate replication performance in detail, use replication analytics statistics.

- [Monitoring Replication Progress - BUI, CLI](#)
- [“Setting Replication Alerts” on page 551](#)
- [“Replication Audit Actions” on page 552](#)
- [Monitoring Replication Delays and RPO - BUI, CLI](#)
- [“Using Replication Analytics” on page 556](#)

▼ Monitoring Replication Progress (BUI)

1. **To monitor the progress of a replication update, go to Shares > Projects and select the replicated project, or select a project and then select the replicated share.**

2. **Click the Replication tab.**

The replication initial stages are displayed below the progress bar.

TARGET ▲	UPDATES	STATUS
cleo Manual	2016-5-2 10:59:32 Synced 2016-5-2 10:59:32 Attempted	Connecting to replication target

The different stages are:

- Connecting to replication target
 - Receiving checkpoint from replication target
 - Estimating size of update
 - Building deduplication tables
3. After the replication action is sending data, you can view the percentage of bytes sent, estimated size, average throughput, and estimated remaining time.

TARGET ▲	UPDATES	STATUS
 brmzs3-2-250 Scheduled	2014-12-24 02:49:52 Synced 2014-12-24 02:49:52 Attempted	  78% of 2.5T @50MB/s (-03:13:10)

Related Topics

- [“Replication Audit Actions” on page 552](#)
- [“Using Replication Analytics” on page 556](#)
- [“Deduplicated Replication” on page 600](#)

▼ Monitoring Replication Progress (CLI)

1. To monitor the progress of a replication update, navigate to the project or share, and enter the replication node.

```
hostname:shares> select TestProj
hostname:shares TestProj> replication
hostname:shares TestProj replication>
```

2. Select the replication action, then enter get:

```
hostname:shares TestProj replication> select action-000
hostname:shares TestProj action-000> get
Properties:
      id = aed46331-160b-48ec-8727-dcd563adbd78
    target_id = 4fd3483e-b1f5-4bdc-9be3-b3a4becd0c42
      target = target1
    enabled = true
  continuous = false
    include_snaps = true
retain_user_snaps_on_target = false
      dedup = true
```

```
include_clone_origin_as_data = false
max_bandwidth = unlimited
bytes_sent = 0
estimated_size = 0
estimated_time_left = 00:00:00
average_throughput = 0B/s
use_ssl = true
compression = on
export_path =
state = sending
state_description = Connecting to replication target
export_pending = false
offline = false
next_update = Sync now
replica_data_timestamp = Thu Apr 28 2016 22:18:11 GMT+0000 (UTC)
last_sync = <unknown>
last_try = <unknown>
last_result = <unknown>
replica_lag = 00:00:09
recovery_point_objective =
replica_lag_warning_alert =
replica_lag_error_alert =
replica_lag_over_warning_limit = false
replica_lag_over_error_limit = false
```

3. Review property `state_description` for information on replication progress.

The different states are:

- Connecting to replication target
- Receiving checkpoint from target
- Estimating size of update
- Building deduplication tables
- Sending update

State Building deduplication tables is displayed only if the project or share has deduplication enabled.

4. If `state_description` is Sending update, determine the replication progress by reviewing the following properties:

- `bytes_sent`
- `estimated_size`
- `estimated_time_left`
- `average_throughput`

Related Topics

- [“Replication Audit Actions” on page 552](#)
- [“Using Replication Analytics” on page 556](#)
- [“Deduplicated Replication” on page 600](#)

▼ Setting Replication Alerts

Use this task to configure how the system responds to replication alert events. For more information on Replication Alerts, see [“Replication Alerts” on page 616](#).

The appliance posts alerts when any of the following replication events occur:

- Manual or scheduled replication update starts or finishes successfully (both source and target).
- Any replication update fails, including as a result of explicit cancellation by an administrator (both source and target).
- A scheduled replication update is skipped because another update for the same action is already in progress.
- When a continuous replication starts for the first time, fails, or resumes after a failure.
- A replica time lag exceeds its specified threshold.

To configure how the system responds to alert events:

1. **Go to Configuration > Alerts.**
2. **Click the add icon  next to Alert actions.**
3. **Select one of the replication categories:**
 - Remote replication
 - Remote replication: source only

- Remote replication: target only

Add alert action [CANCEL] [ADD]

Events

Category: Active Directory

Events: Active Directory, All events, Analytics, Appliance Software, Cluster, Custom, Hardware events, Hardware faults, NDMP operations, NDMP: backup only, NDMP: restore only, Network, Phone home, **Remote replication**, Remote replication: source only, Remote replication: target only, Service Failures, Shadow migration events, Thresholds, ZFS pool

Alert actions

Send email [v] Send to: [] Subject: []

[TEST]

4. **Select all events or a subset of events.**
5. **Set one of the following alert actions:**
email, SNMP trap, syslog message, resume dataset, suspend dataset, resume worksheet, suspend worksheet, or execute workflow.
6. **(Optional) Click TEST to test the selected alert action.**
7. **Click Add.**

Related Topics

- Monitoring Replication Delays and RPO [BUI](#), [CLI](#)

Replication Audit Actions

The following replication configuration actions are tracked and written to the audit log. To view audit log entries in the BUI, go to Maintenance > Logs > Audit.

- Creating, modifying, or destroying replication actions
- Adding or removing shares from a replication group
- Creating, modifying, cloning, reversing, severing or destroying replication packages on the target
- Creating, modifying, or destroying replication targets


▼ Monitoring Replication Delays and RPO (BUI)

With asynchronous replication, a time delay occurs when writing data from the source to its replica on the target. A warning and error alert can be set to notify the administrator when a replication delay is approaching or exceeds the replication point objective (RPO). These alerts prompt the administrator to check for networking problems, application problems, and to examine the health of the storage appliance using analytics.

A replication delay alert can be set when creating or editing a replication action.

To set replication delay alerts:

1. **Go to Shares > Projects.**
2. **Select a project or share, and click the Replication tab.**
3. **Next to Actions, click the add icon .**
4. **Select the properties for this action. See [“Replication Action Properties” on page 594](#).**
5. **Select "Recovery point objective" and enter a value. Then specify days, hours, minutes, or seconds.**



Recovery point objective hours

Replica lag warning alert % of Recovery Point Objective

Replica lag error alert % of Recovery Point Objective

Update frequency ☒ Scheduled ☐ Continuous

6. **Select "Replica lag warning alert" and "Replica lag error alert" and specify a percentage of the RPO for each property.**

Setting these properties will generate warnings at different times. For example, enter 50 to generate a minor alert when the replication delay exceeds 50% of the RPO. Enter 180 to generate a major alert when the replication delay exceeds 180% of the RPO.

When the replica lag falls below the set values, a minor alert reports the replica lag is within the warning or error limit.

7. **Set an alert action as described in [“Setting Replication Alerts” on page 551](#).**

Related Topics

- [“Creating a Replication Action \(BUI\)” on page 518](#)
- [“Replication Action Properties” on page 594](#)
- [“Setting Replication Alerts” on page 551](#)

▼ Monitoring Replication Delays and RPO (CLI)

With asynchronous replication, a time delay occurs when writing data from the source to its replica on the target. A warning and error alert can be set to notify the administrator when a replication delay is approaching or exceeds the replication point objective (RPO). These alerts prompt the administrator to check for networking problems, application problems, and to examine the health of the storage appliance using analytics.

A replication delay alert can be set when creating or editing a replication action.

To set replication delay alerts:

1. **Navigate to the project or share, and enter action:**

```
host_source:shares New_Project replication> action
```

2. **Display the properties by entering get.**

```
host_source:shares New_Project action (uncommitted)> get
      origin_pkg_id =
        target = replication-target
        pool = demo_pool
        enabled = true
        continuous = false
        include_snaps = true
      retain_user_snaps_on_target = false
        dedup = false
      include_clone_origin_as_data = false
```

```

        max_bandwidth = unlimited
        bytes_sent = 0
        estimated_size = 0
        estimated_time_left = 00:00:00
        average_throughput = 0B/s
        use_ssl = true
        compression = on
        export_path =
        state = idle
        state_description = Idle (no update in progress)
        export_pending = false
        offline = false
        next_update = Sync now
        replica_lag = P1H30M
        replica_data_timestamp = Wed Feb 15 2016 12:12:05 GMT+0000 (UTC)
        last_sync = Wed Feb 15 2016 22:32:59 GMT+0000 (UTC)
        last_try = Wed Feb 15 2016 22:32:59 GMT+0000 (UTC)
        last_result = success
        recovery_point_objective =
        replica_lag_warning_alert =
        replica_lag_error_alert =

```

3. **Set the RPO and replica lag properties for this action as shown in the following example:**

```

host_source:shares New_Project action (uncommitted)> set recovery_point_objective=50min
        recovery_point_objective = 50 minutes (uncommitted)
host_source:shares New_Project action (uncommitted)> set replica_lag_warning_alert=50
        replica_lag_warning = 50% (uncommitted)
host_source:shares New_Project action (uncommitted)> set replica_lag_error_alert=180
        replica_lag_error = 180% (uncommitted)

```

For a description of all properties, see [“Replication Action Properties” on page 594](#).

4. **Commit the changes for this action.**

```

host_source:shares New_Project action (uncommitted)> commit

```

5. **To view the current properties for the action, enter `ls`. The RPO and replica lag portion of the output is shown in the following example:**

```

host_source:shares New_Project action (uncommitted)> ls
.
.
        replica_lag = P1H30M
        replica_data_timestamp = Wed Feb 15 2016 12:12:05 GMT+0000 (UTC)
        last_sync = Wed Feb 15 2016 22:32:59 GMT+0000 (UTC)
        last_try = Wed Feb 15 2016 22:32:59 GMT+0000 (UTC)

```

```
last_result = success
recovery_point_objective = 50 minutes
replica_lag_warning_alert = 50%
replica_lag_error_alert = 180%
```

Related Topics

- [“Creating a Replication Action \(CLI\)” on page 520](#)
- [“Replication Action Properties” on page 594](#)
- [“Setting Replication Alerts” on page 551](#)

Using Replication Analytics

The following analytics statistics are available for monitoring replication progress:

- Data Movement: Replication bytes
- Data Movement: Replication operations
- Data Movement: Replication latencies (advanced analytics)
- Data Movement: Replication send/receive bytes (advanced analytics)

Each statistic can be broken down by direction, type of operation, peer, pool name, project, data set, or captured as a raw statistic. For more information about analytics and statistics, see [“Working with Analytics” in *Oracle ZFS Storage Appliance Analytics Guide*](#).

Managing Replication Packages

When projects and shares are replicated to a target, the data is stored in a replication package. Replication packages can be used for disk-to-disk backup, failover, testing purposes, or mounted from a client for read-only access.

Use any of the following methods to access data within a replication package:

- **Export a replication package** - When you export a selected replication package, the exported shares contain the data from the most recently completed replication update. After exporting the shares, you can access any specific snapshot data for any of the shares by navigating to the `.zfs/snapshot` directory in the share's root directory. The `.zfs` directory is normally invisible, but can be accessed by specifying a subdirectory or file name explicitly, as described in [Accessing a Hidden Filesystem Snapshot Directory](#). For information about changing the snapshot visibility property, see Making a Filesystem Snapshot Directory Visible [BUI](#), [CLI](#).

- **Clone a replication package** - Creating a clone of a replication package uses the most recent data received from the source appliance. Cloning converts a replication package into a new project allowing read-write access to all data in the project. The cloned project must have a unique name, mountpoint, and SMB resource name that does not conflict with any existing project. For more information, see [Cloning a Replication Package BUI, CLI](#).
- **Clone shares from a replication package** - Individual share snapshots, created prior to the most recent replication, can be cloned from a replication package. Cloning an individual snapshot from a replication package provides read-write access to data as it existed on the source appliance at the time a snapshot was created. For more information, see [Cloning a Snapshot BUI, CLI](#).
- **Sever a replication package** - The sever operation converts a replication package to a new project, allowing read-write access to data within the project. The replication connection between the source appliance and target appliance is severed after this operation. Note that the new project must have a unique name, mountpoint, and SMB resource name that does not conflict with any existing projects. For more information, see [Severing a Replication Package BUI, CLI](#).
- **Reverse a replication package** - Reverse replication converts a replication package into a new project, allowing read-write access to data within the project. The replication connection is preserved and a new replication action is created that allows replication back to the original source appliance. For more information, see [“Reverse the Direction of Replication” on page 625](#).

Additional tasks related to replication packages include:

- [Managing User-Generated Snapshots](#)
- [Canceling a Replication Update - BUI, CLI](#)
- [Cloning a Replication Package - BUI, CLI](#)
- [Severing a Replication Package - BUI, CLI](#)
- [Editing a Replication Package - BUI, CLI](#)
- [Disabling a Replication Package - BUI, CLI](#)

Managing User-Generated Snapshots

Disk-to-disk backup can also be achieved by setting the property "Retain user-generated snapshots on target". User-generated snapshots, created on a source appliance, are replicated to a target which serves as an incremental-forever backup repository.

Setting this property allows you to manage user-generated snapshots independently on a source appliance and replication target. Normally, when you destroy user-generated snapshots on the source appliance, the snapshots are immediately destroyed at the replication target after a replication update. To keep user-generated snapshots on the target, set this property when creating or editing a replication action.

When user-generated snapshots are no longer needed, manually destroy them on the replication target. To destroy snapshots, see [Destroying a Snapshot BUI, CLI](#).

Related Topics


- [“Replication Action Properties” on page 594](#)
- [Creating a Replication Action BUI, CLI](#)
- [Editing a Replication Action BUI, CLI](#)
- [“Snapshot Space Management” on page 482](#)

▼ Canceling a Replication Update (BUI)

Replication packages are displayed as projects under the Replica filter.

1. **From the target appliance, go to Shares > Projects, and click Replica.**
2. **Click the Replication tab.**

If an update is in progress, you will see a barber pole progress bar with a cancel icon  next to it.



3. **Click the cancel icon .**

It might take several seconds for the cancellation to complete.

After canceling an update, the next scheduled update sends the remainder of the interrupted data stream, followed by an incremental update that is based on the interrupted data stream.

Note - A manual update cannot be initiated from the target appliance. You must log into the source appliance to initiate a manual update.

Related Topics

- [“Replication Packages” on page 621](#)
- [“Resumable Replication” on page 615](#)
- [“Manually Sending a Replication Update \(BUI\)” on page 525](#)

▼ Canceling a Replication Update (CLI)

You can cancel in-progress replication updates from the replication target.

1. **From the replication target, navigate to `shares replication packages` and then enter `ls` to list the packages.**

```
hostname:> shares replication packages
hostname:shares replication packages> ls
```

ID	STATE	DATA	TIMESTAMP	SOURCE	DATASET
package-002	idle	2015-10-02	19:26:37	hostsource	berries
package-001	idle	2015-10-02	19:26:10	hostsource	berries
package-004	idle	2015-10-02	20:53:51	hostsource	berries/blackberry
package-003	recv	2015-10-02	20:59:52	hostsource	cherries/maraschino

Entries are sorted by source, dataset, and data timestamp respectively. The most recent replica snapshot is indicated by data timestamp.

2. **Select a package.**

```
hostname:shares replication packages> select package-001
```

3. **Enter `cancelupdate`.**

```
hostname:shares replication package-001> cancelupdate
```

It might take several seconds for the cancellation to complete.

After canceling an update, the next scheduled update sends the remainder of the interrupted data stream, followed by an incremental update that is based on the interrupted data stream.

Note - A manual update cannot be initiated from the replication target. You must log into the source appliance to initiate a manual update.

Related Topics

- [“Replication Packages” on page 621](#)

- [“Manually Sending a Replication Update \(CLI\)” on page 526](#)

▼ Cloning a Replication Package (BUI)

A clone of a replication package is based on the most recently received replication snapshot.

When creating a cloned project, avoid naming conflicts by following these guidelines:

- The cloned project must have a unique name that does not conflict with any of the existing projects within the same pool.
- The mountpoint and SMB resource name for any of the shares of the cloned project must not conflict with any existing mountpoint or SMB resource names.
- For shares that inherit project properties, resolve conflicts by overriding the project-level mountpoint and/or SMB resource name.
- For shares that do not inherit properties from the project, set a suffix that will be appended to the mountpoint and/or SMB resource names to resolve conflicts, or override the mountpoint and/or SMB resource name individually for each share.

Use the following procedure to clone a replication package.

1. **From the target appliance, navigate to the replication package you want to clone.**
2. **Click the Replication tab.**
3. **Click the clone icon .**
4. **In the Clone Replication Project dialog box, complete the following fields:**
 - a. **New project - Enter a unique name for the new project (clone).**

A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters _ - . :
 - b. **(Optional) Mountpoint - Enter a unique project-level mountpoint for the clone.**

This setting applies to shares that inherit the mountpoint from the project. Entering a unique mountpoint helps avoid conflicts.
 - c. **(Optional) Disable SMB Sharing - Check to disable SMB.**


This setting applies to shares that inherit SMB sharing from the project. Shares that do not inherit SMB sharing from the project are not affected. Disabling SMB sharing on the

project level will not disable SMB sharing, for shares that do not inherit SMB sharing from project.

d. (Optional) SMB Resource Name Prefix - Enter an SMB resource name.

This setting applies to shares that inherit the SMB resource name from the project. Entering a unique resource name helps avoid conflicts.

When SMB is enabled, you can share the clone over SMB. The SMB Resource Name Prefix used to share inherited shares of the new cloned project will be constructed using the prefix you add plus the name of the corresponding share.



Clone Replication Project

CANCEL APPLY

New project

Mountpoint

Disable SMB Sharing ☐

SMB Resource Name Prefix

Tip: Enter unique Mountpoint and/or Resource Name in order to avoid conflicts.

5. Click APPLY.

If there are no mountpoint or resource name conflicts, the clone operation is initiated.

If the project name is already in use, an alert appears and a new project name must be entered.

6. (Optional) If conflicts are detected, use the additional dialog boxes to resolve them.

a. Resolve conflicts that inherit a mountpoint and/or SMB resource names from the project.




The dialog box is titled "Clone Replication Project (Resolve Conflicts)" and has "CANCEL" and "APPLY" buttons. It contains the following fields and options:

- "New project: clone-project" (text label)
- A message: "Conflicts detected for mountpoints and/or SMB resource names that inherit from the project. Use project level setting to resolve the conflicts."
- "Mountpoint" label with a text input field containing "/export"
- "Disable SMB Sharing" label with an unchecked checkbox
- "SMB Resource Name Prefix" label with an empty text input field
- A section titled "List of Conflicting Shares" with a "Total: 1" indicator and a table with one row of dashes.

- i. Enter a unique mountpoint.
 - ii. Disable SMB Sharing or enter a unique SMB resource name prefix.
 - iii. Click APPLY.
- b. Resolve conflicts for shares that do not inherit a mountpoint and/or SMB resource names from the project.

This dialog box appears only after you have resolved all conflicts for inheriting shares.



- i. Enter a unique suffix to append to the mountpoint and/or SMB resource names of these shares.
 - ii. Click APPLY.
7. If conflicts still exist, repeat step 6 to resolve the appropriate conflicts.

Related Topics

- [“Cloning a Replication Package or Share” on page 621](#)
- [“Replication Packages” on page 621](#)

▼ Cloning a Replication Package (CLI)

A clone of a replication package is based on the most recently received replication snapshot.

When creating a cloned project, avoid naming conflicts by following these guidelines:

- The cloned project must have a unique name that does not conflict with any of the existing projects within the same pool.
- The mountpoint and SMB resource name for any of the shares of the cloned project must not conflict with any existing mountpoint or SMB resource names.
- For shares that inherit project properties, resolve conflicts by overriding the project-level mountpoint and/or SMB resource name.
- For shares that do not inherit properties from the project, set a suffix that will be appended to the mountpoint and/or SMB resource names to resolve conflicts, or override the mountpoint and/or SMB resource name individually for each share.

Use the following procedure to clone a replication package.

1. From the replication target, navigate to shares replication packages and list the packages.

```
target:> shares replication packages
target:shares replication packages> list
```

ID	STATE	DATA	TIMESTAMP	SOURCE	DATASET
package-002	idle	2015-10-02	19:26:37	hostsource	berries
package-001	idle	2015-10-02	19:26:10	hostsource	berries
package-004	idle	2015-10-02	20:53:51	hostsource	berries/blackberry
package-003	recv	2015-10-02	20:59:52	hostsource	cherries/maraschino

Entries are sorted by source, dataset, and data timestamp respectively. The most recent replica snapshot is indicated by data timestamp.

2. Select the replication package you want to clone.

```
target:shares replication packages> select package-001
```

3. Enter clone to create a new clone project.

```
target:shares replication package-001> clone
target:shares replication package-001 clone>
```

4. Set target_project to the project name.

The project name must be unique, or the clone operation will fail.

A project name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters _ - . :

```
target:shares replication package-001 clone> set target_project=clone
```



```
target_project = clone
```

5. Enter conflicts to check for conflicts.

If conflicts exist, a message similar to the following appears.

```
target:shares replication package-001 clone> conflicts
```

Cloning cannot proceed because the following shares have mountpoints or SMB resource names that are invalid or conflict with those of other shares (either on the system or also being failed over). Please specify valid mountpoints or SMB resource names for these shares:

SHARE	MOUNTPOINT	SHARESMB
share1	/export/share1	share1
clothes	/export/clothes (inherited)	clothes (inherited)
electronics	/export/electronics	electronics
furniture	/export/furniture (inherited)	furniture (inherited)
groceries	/export/groceries (inherited)	groceries (inherited)
health	/export/health (inherited)	health (inherited)
toys	/export/toys	toys

```
target:shares replication packages package-001 clone>
```

Note - The conflicts command can be used at any point in this procedure to check for mountpoint or naming conflicts.

6. (Optional) Set project-level properties to resolve conflicts for shares that inherit properties from a project.

Use the get command to view the properties of the clone.

```
target:shares replication package-001 clone> get
      target_project = clone2
      rename_suffix =
      original_mountpoint = /export
      mountpoint = /export/clone
      original_smb_resource_name = off
      smb_resource_name = off
```

The property mountpoint shows the current mountpoint. The property smb_resource_name shows the current resource name.

a. Enter a unique project-level mountpoint for the clone.

This setting applies to shares that inherit a mountpoint from a project. Use set mountpoint to specify a unique mountpoint for the clone.

```
target:shares replication package-001 clone> set mountpoint=/export/clone
mountpoint = /export/clone
```

b. Enter a unique project-level SMB resource name.

This setting applies to shares that inherit the SMB resource name from a project. Set **smb_resource_name** to a unique SMB resource name.

```
target:shares replication package-001 clone> set smb_resource_name=clone
smb_resource_name = clone
```

c. Set rename_suffix to resolve remaining share conflicts.

This property creates a suffix that is appended to a mountpoint and/or SMB resource names, if a conflict occurs.

```
target:shares replication package-001 clone> set rename_suffix=-clone
rename_suffix = -clone
```

Note - This operation overrides inheritance. For example, if a share originally inherits its mountpoint from the project, but the mountpoint is renamed with a suffix during the clone operation, the share in the new cloned project no longer inherits its mountpoint, but instead uses the unique renamed mountpoint.

7. (Optional) To set properties for an individual share:

a. Select a share.

```
target:shares replication package-001 clone> select share1
```

b. Override its mountpoint and/or SMB resource name.

The following example overrides the mountpoint for the share.

```
target:shares replication package-001 clone share1> set mountpoint=/export/
appliances-clone
target:shares replication package-001 clone share1> set sharesmb=appliances-clone
```

8. Enter confirm commit to initiate the clone operation.

```
target:shares replication package-001 clone> confirm commit
```

If name conflicts are detected, a message similar to the following appears:

Cloning cannot proceed because the following shares have
mountpoints or SMB resource names that are invalid or conflict
with those of other shares (either on the system or also being

failed over). Please specify valid mountpoints or SMB resource names for these shares:

SHARE	MOUNTPOINT	SHARESMB
share1	/export/share1	share1
clothes	/export/clothes (inherited)	clothes (inherited)
electronics	/export/electronics	electronics
furniture	/export/furniture (inherited)	furniture (inherited)

target:shares replication package-001 clone>

9. **(Optional) Resolve any remaining name conflicts and confirm the cloning.**
Repeat steps 6 and 7, as appropriate, until no conflicts remain, and then enter commit.

Related Topics

- [“Cloning a Replication Package or Share” on page 621](#)
- [“Remote Replication Workflow” on page 511](#)

▼ Cloning a Snapshot in a Replication Package (BUI)

Note - Cloning is a licensed feature. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the appropriate Licensing Information User Manual for your product.

Note - Snapshots within a replication package can be temporary. It's possible for a replication snapshot to be destroyed by replication updates, to accommodate new snapshots from the source. Therefore, this procedure recommends disabling replication updates for the package before cloning the replication snapshot.


Use the following procedure to clone a replication snapshot within a replication package.

1. **Go to the package that contains the share you want to clone.**
 - a. **Select Shares > Projects.**
 - b. **Click Replica above the list of projects.**
 - c. **Hover over the package and click the edit icon .**
2. **Suspend replication updates for the replication package.**


This will disable replication updates for a package entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.

a. In the replication package details page, select **Replication**.

b. Click the power icon  to disable replication updates.




The status icon on the left indicates the status of the package.



The package remains disabled until you re-enable it using the same power icon .

3. Navigate to the share you want to clone.

a. In the replication package details page, select the **Shares** tab.



b. Hover over a filesystem or LUN and click the edit icon .

c. Click the **Snapshots** tab.

4. **Hover over the snapshot you want to clone and click the clone icon .**
5. **In the Create Clone dialog box, set the following fields.**
 - a. **From the Project drop-down menu, select the destination project.**
The clone is created within the current project, by default, but you can specify a different project.
 - b. **Type a name for the clone.**
 - c. **(Optional) Click the lock icon  next to Mountpoint and set a mountpoint for the clone.**
If you leave this field locked, the mountpoint for the clone remains as /export/<sharename>.
 - d. **(Optional) Click the lock icon  next to Resource name and enter one of the following values:**
 - **off** - To disable SMB.
 - **on** - To enable SMB, so you can share the clone over SMB. The name of the clone in SMB matches the name of the clone in the appliance.
 - **<resource_name>** - SMB is enabled, so you can share the clone over SMB. The name of the clone in SMB is the name you specify here instead of the name of the clone in the appliance.

If you leave the Resource field locked, the Resource name property is inherited from the snapshot you are cloning.
 - e. **(Optional) Check the Inherit key checkbox or uncheck the checkbox and select the keystore (Local or OKM) and name of the encryption key you want the clone to inherit.**
If the box is checked, the keystore and keyname of the clone will be that of the destination project.
If the box is unchecked, the keystore and keyname of the clone will be that of the parent share. Alternatively, select a different keystore and keyname from the drop-down menu.
 - f. **(Optional) Check the Retain Other Local Settings checkbox to cause any inherited properties to be preserved as local settings in the new clone.**
This field determines whether inherited properties will come from the parent dataset or the destination project. By default, the box is unchecked, meaning that all inherited

properties will come from the destination project for the new clone. If you check the box, all currently inherited properties will be preserved as local settings in the new clone.

6. Click APPLY to confirm the settings and create the clone.

The clone appears in the list of shares for the destination project you set. You can work with a clone just like any other share.

Note - Once a replication package snapshot is cloned, it will can no longer be destroyed by replication updates from the source.

7. Re-enable replication updates for the replication package.

a. **Navigate to the parent project of the share you just cloned.**

b. **Click the Replication tab.**

c. **Click the power icon .**

Ensure the status icon on the left is green, indicating that replication updates have been enabled.

Related Topics

- [“Cloning a Replication Package or Share” on page 621](#)

▼ Cloning a Snapshot in a Replication Package (CLI)

Note - Cloning is a licensed feature. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the appropriate Licensing Information User Manual for your product.

Note - Snapshots within a replication package can be temporary. It's possible for a replication snapshot to be destroyed by replication updates, to accommodate new snapshots from the source. Therefore, this procedure recommends disabling replication updates for the package before cloning the replication snapshot.

Use the following procedure to clone a replication snapshot within a replication package.

1. **Go to shares replication packages and enter list to display the available replication packages.**

```
hostname:> shares replication packages
hostname:shares replication packages> list
Packages:
```

ID	STATE	DATA_TIMESTAMP	SOURCE	DATASET
package-005	idle	2017-04-28 22:28:08	src1	data1
package-004	idle	2017-04-28 15:44:38	src1	data1
package-003	disbl	2017-04-27 23:46:20	src1	data1
package-002	idle	2017-04-27 23:14:10	src1	data1
package-001	idle	2017-04-17 17:27:05	src2	data2

2. **Select the package that contains the share you want to clone.**

```
hostname:shares replication packages> select package-005
```

3. **Suspend replication updates for the replication package.**

This action disables replication updates for a package entirely, which cancels any ongoing updates and causes new updates from the source appliance to fail.

- a. **Enter set enabled=false.**

```
hostname:shares replication package-005> set enabled=false
enabled = false (uncommitted)
```

- b. **Enter commit.**

```
hostname:shares replication package-005> commit
```

4. **Select the project that contains the share you want to clone.**

- a. **Enter show to display the project name.**

```
hostname:shares replication package-005> show
Properties:
    id = 7e184188-2738-432b-f304-123412341234de
    ...
    ...
    ...

Projects:
    proj1
```

- b. **Select the project.**

```
hostname:shares replication package-005> select proj1
```

5. Select the share you want to clone.

a. Enter show to display the available shares.

```
hostname:shares replication package-005 proj1> show
Properties:
```

```
    aclinherit = restricted
    ...
    ...
    ...
```

```
Shares:
```

```
Filesystems:
```

NAME	SIZE	ENCRYPTED	MOUNTPPOINT
share1	36K	off	/export/share1
share2	36K	off	/export/share2

b. Select the share.

```
hostname:shares replication package-005 proj1> select share1
```

6. Select the snapshot you want to use to clone the share.

a. Enter snapshots.

```
hostname:shares replication package-005 proj1/share1> snapshots
```

b. Enter list to display the available snapshots.

```
hostname:shares replication package-005 proj1/share1 snapshots> list
Snapshots:
```

```
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-ec
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f2
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f3
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f4
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f5
```

```
Children:
```

```
    automatic => Configure automatic snapshots
```

c. Select the snapshot you want to clone.


```
hostname:shares replication package-005 proj1/share1 snapshots> select .rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb
```

7. Clone the snapshot.

- a. Use the `clone` command, followed by the name of the project in which you want to create the clone, and the name for the clone.

```
hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb> clone proj_name clone1
```

- b. Use the `get` command to view properties.

```
hostname:shares proj_name/clone1 (uncommitted clone)> get
Properties:
    aclinherit = restricted (inherited)
    aclmode = discard (inherited)
    atime = true (inherited)
    checksum = fletcher4 (inherited)
    compression = off (inherited)
    copies = 1 (inherited)
    logbias = latency (inherited)
    mountpoint = /export/clone1 (inherited)
    quota = 0 (default)
    readonly = false (inherited)
    ...
    ...
    ...
```

- c. Use the `set` command to adjust properties.

```
hostname:shares proj_name/clone1 (uncommitted clone)> set mountpoint=/export/clone_mountpoint_name
mountpoint = /export/clone_mountpoint_name (uncommitted)
```

- d. Use the `commit` command to commit the changes and create the clone.

```
hostname:shares proj_name/clone1 (uncommitted clone)> commit
hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb>
```

8. Re-enable replication updates for the replication package.

- a. Enter `cd /` to return to the top level.

```
hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-  
fa116972d26f-cb> cd /
```

- b. Enter `shares replication`, then use `select` and the package name.**

```
hostname:> shares replication  
hostname:shares replication> select package-005
```

- c. Use `set enabled=true` to re-enable replication updates to the package. Then enter `commit` to save the changes.**

```
hostname:shares replication package-005> set enabled=true  
enabled = true (uncommitted)  
hostname:shares replication package-005> commit
```

- 9. Check your specified project destination to see the clone.**

- a. Enter `cd /` to return to the top level.**

```
hostname:shares replication package-005> cd /
```

- b. Use `shares select` with the project you used for the clone destination.**

```
hostname:> shares select proj_name
```

- c. Enter `show` to list the shares, and look for the cloned share.**

```
hostname:shares proj_name> show  
Properties:  
    aclinherit = restricted  
    aclmode = discard  
    ...  
    ...  
    ...  
  
Shares:  
  
Filesystems:  
  
NAME          SIZE  ENCRYPTED  MOUNTPOINT  
clone1        1K    off       /export/clone_mountpoint_name
```

Note - Once a replication package snapshot is cloned, it will can no longer be destroyed by replication updates from the source.

Related Topics

- [“Cloning a Replication Package or Share” on page 621](#)

▼ Severing a Replication Package (BUI)

Before You Begin Check for mount point or SMB share name conflicts between replicated filesystems and other filesystems on the system. To resolve mount point (and/or SMB resource names) conflicts, change the project or share mount points (or SMB resource names) in the replication package before severing a replication package. For more information, see [“Severing Replication” on page 624](#).

1. **Navigate to the replication package.**
2. **Click the Replication tab.**
3. **Click the sever icon .**
4. **Enter a name for the new local project.**

Note - If a replication update is performed during or after a sever operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.

Related Topics

- [“Severing Replication” on page 624](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Severing a Replication Package (CLI)

Before You Begin Check for mount point or SMB share name conflicts between replicated filesystems and other filesystems on the system. To resolve mount point (and/or SMB resource names) conflicts, change the project or share mount points (or SMB resource names) in the replication package before severing a replication package. For more information, see [“Severing Replication” on page 624](#).

1. **From the replication target, navigate to the replication package.**

```
host-target:shares default replication source-001 package-001>
```

2. Enter the sever and a name for the new local project.

```
host-target:shares default replication source-001 package-001> sever new_project
```

Note - If a replication update is performed during or after a sever operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.


Related Topics

- [“Severing Replication” on page 624](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Editing a Replication Package (BUI)

1. From the target appliance, go to Shares > Projects, and click Replica.

The name, size, and creation date of each replication package is displayed.




The screenshot shows a web interface titled 'Projects' with a sub-header '6 Total'. Below the title are tabs for 'SHOW ALL', 'LOCAL', and 'REPLICA', with 'REPLICA' being the active tab. A table displays two rows of replication packages. The first row has a name 'loader: default', a size of '50.1M', and a creation date of '2010-2-18 09:45:49'. The second row has a name 'loader: default', a size of '100M', and a creation date of '2010-2-18 13:42:25'. A mouse cursor is pointing at the 'REPLICA' tab.

NAME ▲	SIZE	CREATION
loader: default	50.1M	2010-2-18 09:45:49
loader: default	100M	2010-2-18 13:42:25

Note - Packages are displayed in the BUI only after the first replication update has begun. They may not appear in the list until some time after the first update has completed.

2. Select a replication package for editing.

The Shares view for the package's project is displayed.



3. To modify package properties, click the **Replication** tab.

For a list of properties that you can modify, see [“Replication Packages” on page 621](#).

Related Topics

- [“Replication Actions and Packages” on page 591](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Editing a Replication Package (CLI)

Replication packages are organized in the CLI in a flat view under `shares replication packages`, which displays all replication packages in the system.

1. From the replication target, go to `shares replication packages` and enter `list` to list all replication packages in the system.

```
hostname:> shares replication packages
hostname:shares replication packages> list
```

ID	STATE	DATA	TIMESTAMP	SOURCE	DATASET
package-002	idle	2015-10-02	19:26:37	hostsource	berries
package-001	idle	2015-10-02	19:26:10	hostsource	berries
package-004	idle	2015-10-02	20:53:51	hostsource	berries/blackberry
package-003	recv	2015-10-02	20:59:52	hostsource	cherries/maraschino

The packages are sorted by SOURCE, DATASET, and DATA_TIMESTAMP (in descending order).

2. Select a package.

```
hostname:shares replication packages> select package-001
```

```
hostname:shares replication packages package-001> show
```

Properties:

```

    id = d6137c89-7056-4788-a4d1-b5892fe315e0
    source_name = hostsource
    source_asn = a751dc0f-abcd-1234-6789-f5e8315eaffa
    source_ip = 00.000.00.00:000
    source_pool = poolS
    target_pool = poolT
    replica_of = berries
    enabled = true
    state = idle
    state_description = Idle (no update in progress)
    offline = false
    import_path =
    data_timestamp = Fri Oct 02 2015 19:26:10 GMT+0000 (UTC)
    last_sync = Fri Oct 02 2015 19:26:10 GMT+0000 (UTC)
    last_try = Fri Oct 02 2015 19:26:10 GMT+0000 (UTC)
    last_result = success

```

Projects:

```
berries
```

A replication package can be selected directly by specifying its ID, as shown in the following example:

```

hostname:shares replication packages> select d6137c89-7056-4788-a4d1-b5892fe315e0
hostname:shares replication packages package-001>

```

3. To edit project properties and shares, select a project.

```

hostname:shares replication packages package-001> select berries
hostname:shares replication packages package-001 berries> get mountpoint
    mountpoint = /export
hostname:shares replication packages package-001 berries> get sharenfs
    sharenfs = on

```

For a list of package properties that you can modify, see [“Replication Package Properties” on page 597](#).

Related Topics

- [“Replication Packages” on page 621](#)
- [“Remote Replication Concepts” on page 588](#)

▼ Disabling a Replication Package (BUI)

Replication updates for a package can be disabled entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.

1. **From the target appliance, navigate to the package.**
2. **Click the Replication tab.**
3. **Click the power icon .**

The status icon on the left indicates the status of the package (enabled, disabled, or failed). The package remains disabled until you re-enable it using the using the same power icon .

Related Topics

- [“Replication Packages” on page 621](#)
- [“Remote Replication” on page 511](#)

▼ Disabling a Replication Package (CLI)

Replication updates for a package can be disabled entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.

1. **From the target appliance, navigate to the package.**
2. **Modify the `enabled` property.**
3. **Commit your changes.**

The package remains disabled until you set the `enabled` property to on.

Related Topics

- [“Replication Packages” on page 621](#)
- [“Remote Replication Concepts” on page 588](#)

Disaster Recovery with Remote Replication

A two-system disaster recovery site consists of a source appliance at a production site and a replication target located at a recovery site in a geographically different location. In the event of a catastrophic production site failure, the administrator redirects client operations to the recovery site by reversing replication on the replication target, thus ensuring continuous operation. After the production site is restored to normal operation, the administrator updates the production site by reversing replication at the recovery site. To restore the original source-target relationship, replication is then reversed again.

To set up remote replication for disaster recovery, use these tasks:

- Setting Up a Target Appliance at a Recovery Site - [BUI](#), [CLI](#)
- Switching Operations to the Recovery Site - [BUI](#), [CLI](#)
- Updating the Production Site - [BUI](#), [CLI](#)
- Reversing Replication Back to the Production Site - [BUI](#), [CLI](#)

▼ Setting Up a Replication Target at a Recovery Site (BUI)

To create a replication target for disaster recovery:

1. **Identify a replication target at a recovery site.**
The replication target requires a software version compatible with the source appliance. For details, see [MOS DOC ID 1958039.1 \(https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1).
2. **From the source appliance, create a target as described in “Creating a Replication Target (BUI)” on page 516.**
3. **Create a replication action and schedule a continuous replication. See “Creating a Replication Action (BUI)” on page 518.**

Note - Continuous replication minimizes data loss in the event of a disaster at the production site.

Next Steps

- “Switching Operations to the Recovery Site (BUI)” on page 581

▼ Switching Operations to the Recovery Site (BUI)

After a failure occurs at the production site, perform a reverse replication at the recovery site, and then redirect client operations to the recovery site.


1. **From the replication target, go to Shares > Projects > Replica and look for the replication package from the source appliance.**

The replica is named *source_appliance:project/share*.

2. **Double-click the replication package, or click its edit icon **
3. **Click the Replication tab.**
4. **Click the reverse replication direction icon .**

5. **Enter a name for the new local project and enable the action.**

The project name and location of the original action are preserved.



Note - The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the target at the recovery site, the reverse operation will fail.

6. **Click OK.**
7. **Transfer client activity to the IP address of the appliance at the recovery site.**
Depending on the protocol used, map (SMB clients) or remount (NFS clients) the shares using the IP address or name of the appliance at the recovery site.

Next Steps

- [“Updating the Production Site \(BUI\)” on page 582](#)

▼ Updating the Production Site (BUI)

After the production site is restored and back online, replicate the changes written to the recovery site during the outage back to the production site.

1. **From the appliance at the recovery site, go to Shares > Projects > Local and select the new local project.**

The new project is listed with status *Never synced*.

2. **Click Sync Now to start the replication.**

3. **Wait for the replication to complete.**

At the top of the window, *Finished replicating to the new project on the source appliance* is displayed.

Next Steps

- [“Reversing Replication Back to the Production Site \(BUI\)” on page 582](#)

▼ Reversing Replication Back to the Production Site (BUI)

After all changes have been replicated from the recovery site to the production site, reverse replication again to restore the original replication relationship between source and replication targets.

1. **From the production appliance, go to Shares > Projects > Replica and look for the new project name.**

The project is named *target_appliance: new_project/share*.

2. **Select the new project and click its edit icon .**
3. **Click the Replication tab.**

4. Click the reverse replication direction icon .
5. In the Reverse Replication window, enter a name for the new local project and then enable the action.

Note - The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the source at the production site, the reverse operation will fail.

6. Depending on the protocol used, remap (SMB clients) or remount (NFS clients) shares to the appliance at the recovery site.
7. Delete the original project on the source appliance.
 - a. Go to Shares > Projects > Local and look for the original project, which should be empty.
 - b. Select the empty project and click its destroy icon .
 - c. Click OK.

▼ Setting Up a Replication Target at a Recovery Site (CLI)

To create a replication target for disaster recovery:

1. **Identify a replication target at the recovery site.**
The replication target requires a software version compatible with the source appliance. For details, see [MOS DOC ID 1958039.1 \(https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1).
2. **From the source appliance, create a target as described in “Creating a Replication Target (CLI)” on page 517.**
3. **Create a replication action and schedule a continuous replication. See “Creating a Replication Action (CLI)” on page 520.**

Note - Continuous replication minimizes data loss in the event of a disaster at the production site.

Next Steps

- [“Switching Operations to the Recovery Site \(CLI\)” on page 584](#)

▼ Switching Operations to the Recovery Site (CLI)

After a failure occurs at the production site, perform a reverse replication at the recovery site, and then redirect client operations to the recovery site.

1. **From the replication target, enter `shares replication`.**

```
host-offsite:> shares replication
```

2. **Enter `sources` to list the source appliances that are associated with this target.**

```
host-offsite:shares default replication> sources
```

3. **Look for the package replicated by the source appliance.**

In this example, `source-001` is the source appliance number, and `host-prod` is the source appliance name. `kmm2` is the local project name in replication package number `package-001`.

```
source-001 host-prod
PROJECT                STATE          LAST UPDATE
package-000 <unknown>  idle          unknown
package-001 kmm2       idle          Wed May 01 2015 20:06:27 GMT+0000(UTC)
```

4. **Enter `select` and the source appliance number.**

```
host-offsite:shares default replication sources> select source-001
```

5. **Enter `select` and the replication package number.**

```
host-offsite:shares default replication source-001> select package-001
```

6. **Enter `pkgreverse`.**

```
host-prod:shares replication source-005 package-000> pkgreverse
```

The `pkgreverse` command preserves the properties of the original replication action, including schedules.

7. **(Optional) Set a new project name and enable the action using the following commands:**

```

host-prod:shares replication source-000 package-000 pkgreverse> set
  new_project_name=new-kmm2
    new_project_name = new-kmm2
host-prod:shares replication source-000 package-000 pkgreverse> set
  enable_action_upon_reversal=true
    enable_action_upon_reversal = true

```

Note - The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the target at the recovery site, the reverse operation will fail.

8. Enter `show` to confirm the properties, and then enter `commit`:

```

host-prod:shares replication source-000 package-000 pkgreverse> show
Properties:
    new_project_name = new-kmm2
    enable_action_upon_reversal = true

host-prod:shares replication source-000 package-000 pkgreverse> commit
This action will move the contents of this package to a new local project
configured to replicate back to the source. Any data or metadata changes made
on the source since the last successful update will be lost when the new
project is replicated back to the source. If replication actions on the source
are not disabled, future updates to this package will fail.

```

9. Transfer client activity to the IP address of the appliance at the recovery site.

Depending on the protocol used, map (SMB clients) or remount (NFS clients) the shares using the IP address or name of the appliance at the recovery site.

Next Steps

- [“Updating the Production Site \(CLI\)” on page 585](#)

▼ Updating the Production Site (CLI)

After the production site is restored and back online, replicate the changes written to the recovery site during the outage to the production site.

1. From the appliance at the recovery site, go to shares and select the new project.

```

host-offsite:> shares
host-offsite:shares pool> select new-kmm2

```

2. Enter `list` to find the name of the share.

```
host-offsite:shares pool new-kmm2> list
Filesystems:
NAME          SIZE  ENCRYPTED  MOUNTPPOINT
karen2        31K   off       /export/karen2
host:shares pool new-kmm2> replication
host:shares pool new-kmm2 replication> show
Actions:
TARGET        STATUS      NEXT
action-000
host2         idle       Sync now
```

3. Select the action number and then enter `sendupdate` to start replication to the production appliance.

```
host-offsite:shares pool new-kmm2 replication> select action-000
host-offsite:shares pool new-kmm2 action-000> sendupdate
```

4. Wait for the replication to complete.

The state changes to `idle` when the replication has completed.

Next Steps

- [“Reversing Replication Back to the Production Site \(CLI\)” on page 586](#)

▼ Reversing Replication Back to the Production Site (CLI)

After all changes have been replicated from the recovery site to the production site, reverse replication again to restore the original replication relationship between the source appliance and replication target.

1. From the replication target, navigate to `shares replication packages` and list the packages.

```
loader:> shares replication packages
loader:shares replication packages> list

ID          STATE DATA  TIMESTAMP      SOURCE      DATASET
package-002 idle   2015-10-02 19:26:37 hostsource    berries
package-001 idle   2015-10-02 19:26:10 hostsource    berries
package-004 idle   2015-10-02 20:53:51 hostsource    berries/blackberry
```

```
package-003 recv 2015-10-02 20:59:52 hostsource cherries/maraschino
```

The package with the most recent data timestamp for this dataset is the one with the most recent, up-to-date data for the corresponding source dataset.

2. Select the replication package you want to reverse.

```
loader:shares replication packages> select package-002
```

3. Enter pkgreverse.

```
host-prod:shares replication package-002> pkgreverse
```

4. (Optional) Set a new project name and enable the action using the following commands:

```
host-prod:shares replication package-002 pkgreverse> set new_project_name=new-kmm3
new_project_name = new-kmm3
host-prod:shares replication package-002 pkgreverse> set
enable_action_upon_reversal=true
enable_action_upon_reversal = true
```

Note - The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the appliance at the production site, the reverse operation will fail.

5. Enter show to confirm the properties, and then enter commit:

```
host-prod:shares replication package-002 pkgreverse> show
Properties:
```

```
new_project_name = new-kmm3
enable_action_upon_reversal = true
```

```
host-prod:shares replication package-002 pkgreverse> commit
```

This action will move the contents of this package to a new local project configured to replicate back to the source. Any data or metadata changes made on the source since the last successful update will be lost when the new project is replicated back to the source. If replication actions on the source are not disabled, future updates to this package will fail.

6. Depending on the protocol used, remap (SMB clients) or remount (NFS clients) shares to the appliance at the production site.

Related Topics

- [“Reverse the Direction of Replication” on page 625](#)

- [“Remote Replication Concepts” on page 588](#)

Remote Replication Concepts

Oracle ZFS Storage Appliance Replication is a licensed feature for certain models that provides snapshot-based replication of projects and shares from a source appliance to one or more replication targets. Replication performs a full update of an entire project and/or share contents, followed by incremental updates containing only the changes since the previous update.

For license details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

This topic describes key remote replication (or simply "replication") concepts and replication terminology.

Remote replication has the following important characteristics:

- **Snapshot-based asynchronous replication** - The replication subsystem takes a snapshot as part of each update operation. In the case of an initial update, the entire contents of a project and/or share are sent. In the case of an incremental update, only the changes since the last replication snapshot for the same action are sent. Because replication takes snapshots and then sends them, data is already committed to stable storage before replication even begins sending it. Continuous replication effectively sends continuous streams of filesystem changes, but it is still asynchronous with respect to NAS and SAN clients.
- **Block-level** - Each update operation traverses the filesystem at the block level and sends the appropriate filesystem data and metadata to the target.
- **Includes all metadata** - The underlying replication stream serializes both user data and metadata, including most properties configured on the Shares screen. These properties can be modified on the target after the first replication update completes, though not all take effect until the replication connection is severed. For example, this allows sharing over NFS to a different set of hosts than on the source. For more information, see [“Replication Packages” on page 621](#).
- **Secure** - The replication control protocol used among Oracle ZFS Storage Appliance products is secured with Secure Sockets Layer (SSL). Data can optionally be protected with SSL as well. The appliance can only replicate to or from another appliance after an initial manual authentication process. For more information, see [“Replication Targets” on page 590](#).
- **Encrypted projects and shares** - When enabled, transparent data encryption protects individual shares (filesystems and LUNs) and projects. For more information, see [“Data Encryption” on page 629](#).

- **Protocol independent** - The appliance supports both file and block-based storage volumes. The replication mechanism is protocol independent.
- **Compressed replication** - Support for compressed replication streams increases replication performance and improves throughput utilization between multiple sites that have limited bandwidth connections. For more information, see [“Compressed Replication” on page 620](#).

Replication has the following known limitations:

- Actions cannot move between pools
- Network throughput is limited to a maximum of 200 MB/s per project level replication. With compressed replication, the effective data rate can exceed the actual physical network data rate.

For information about remote replication concepts, see:

- [“Replication Terminology” on page 590](#)
- [“Replication Targets” on page 590](#)
- [“Replication Actions and Packages” on page 591](#)
- [“Replication Action Properties” on page 594](#)
- [“Replication Storage Pools” on page 598](#)
- [“Project vs. Share Replication” on page 598](#)
- [“Replication Authorizations” on page 600](#)
- [“Deduplicated Replication” on page 600](#)
- [“Replication Configuration for Clustered Appliances” on page 603](#)
- [“Example: Replication Configuration for Clustered Appliances” on page 604](#)
- [“Replication Snapshots and Data Consistency” on page 611](#)
- [“Replication Snapshot Management” on page 612](#)
- [“iSCSI Configurations and Replication” on page 615](#)
- [“Resumable Replication” on page 615](#)
- [“Replication Failures” on page 617](#)
- [“Compressed Replication” on page 620](#)
- [“Replication Packages” on page 621](#)
- [“Cloning a Replication Package or Share” on page 621](#)
- [“Exporting Replicated Filesystems” on page 623](#)
- [“Severing Replication” on page 624](#)
- [“Reverse the Direction of Replication” on page 625](#)
- [“Destroying a Replication Package” on page 628](#)
- [“Target Replica Backups” on page 628](#)

Replication Terminology

The following is a list of the common replication terms.

- **clone** - A replicated package can be cloned into a mutable project. A clone can be managed like any other project on the system.
- **full sync or full update** - A replication operation that sends the entire contents of a project and some of its shares. The initial sync of a project and/or share is be a full sync.
- **incremental update** - A replication operation that sends only the differences in a project and its shares since the previous update (whether that one was full or incremental).
- **recovery point objective (RPO)** - The maximum tolerable amount of data loss, expressed in unit of time, in the event of a disaster or major outage. The RPO is defined as part of a disaster recovery plan and represents the last consistent set of data that is available for recovery. The lower an RPO value, the less data loss.
- **replica** - Replicated data contained in the replication package on the replication target.
- **replication action** - Describes the data to be replicated, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link. See [“Replication Action Properties” on page 594](#).
- **replication group** - The set of datasets (exactly one project and some number of shares) that are replicated as a unit. See [“Project vs. Share Replication” on page 598](#).
- **replication package** - Exists on the replication target and is associated with a replication action. It is a special object that contains the replica. A replication package can be exported, cloned, severed, or reversed, which allows for write access to the data within the project.
- **replication source** - An appliance that sends replication updates to one or more target appliances, periodically, continuously, or on demand. Individual appliances can act as both a source and a *target*, but are only one of these in the context of a particular replication *action*.
- **replication target** - An appliance that receives and stores data replicated from one or more *source* appliances. This term also refers to a configuration object on the appliance that enables it to replicate to another appliance.
- **reverse replication** - A replication relationship that exchanges the source and target roles. Following a disaster recovery, roles can be reversed again.

Replication Targets

When creating a replication target on the source appliance, a connection is established that enables secure communications between a source and replication target. This operation requires the following details:

- Name of the replication target - used only to identify the target in the BUI and CLI of the source appliance.

- Network IP address - the data interface of the replication target. The address must be an IPv4 address or host name. Remote replication does not support IPv6 addresses.
- Target root password - to authorize the administrator to set up the connection on the replication target.

The appliances exchange keys used to securely identify each other in subsequent communications. These keys are stored persistently as part of the appliance configuration and persist across reboots and upgrades. They will be lost if the appliance is factory reset or reinstalled. The root password is never stored persistently, so changing the root password on either appliance does not require any changes to the replication configuration. The password is never transmitted in the clear because this initial identity exchange (like all replication control operations) is protected with SSL.

Note - If a replication source uses NIS or LDAP and directly maps these service's users or groups in the share configuration, the equivalent setup must present on the replication target. Otherwise, replication sever and reverse operations could fail.

By default, the replication target connection is not bidirectional. For example, if an administrator configures replication from *source A* to *target B*, *target B* cannot automatically use *source A* as a target. However, the system supports reversing the direction of replication, which automatically creates a target for *source A* on *target B* (if it does not already exist) so that *target B* can replicate back to *source A*. For more information, see [“Reverse the Direction of Replication” on page 625](#).

To configure replication targets, see [Creating a Replication Target BUI, CLI](#).

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication Concepts” on page 588](#)

Replication Actions and Packages

A replication action specifies where and how to replicate a project or share. Replication actions are created on the source appliance specifying the following:

- A replication group that includes either a project or individual shares
- Name of the replication target
- Name of a storage pool on the replication target (used only during the initial setup)
- Frequency (scheduled or continuous) of the update
- Number of Auto-Snapshots (Scheduled Snapshots) that are retained on the target

- Additional options such as encryption of the data stream or disabling compression

A replication group is specified implicitly by the project or share on which the action is configured (see [“Replication Storage Pools” on page 598](#)). The replication target and storage pool cannot be changed after the action is created, but the other options can be modified at any time. If a replication update is in progress when an option is changed, then the new value takes effect when the next update begins (With the exception of the `max_bandwidth` parameter, which takes effect immediately after the modification).

When a replication action is created on the source appliance, a package on the target in the specified storage pool is created. The package on the replication target contains an exact copy of the source project and shares on which the action is configured as of the start time of the last replication update. Actions are the primary unit of replication configuration on the appliance.


Replication Update Frequency

Replications can be executed manually or configured in the replication action to be sent continuously or at scheduled times. The three replication modes are:

- **Manual** - Replication is started manually, at any time, by the administrator. A manual replication update can be useful for testing purposes and for applications that require data to be in a certain state before replication can occur. See [Manually Sending a Replication Update BUI](#), [CLI](#).
- **Scheduled** - Replication is automatically executed according to a selected schedule. The scheduled frequency can be set to replicate to the target every 5, 10, 15, 20 or 30 minutes, every 1, 2, 4, 8 or 12 hours, every day, every week, or every month. More granular update frequencies can be set by defining multiple schedules for a single replication action.

The Auto selection, which is available when creating the first schedule for a replication action, is a start time generated by the appliance. When multiple replication actions are configured on an appliance, the auto-generated start time can minimize overlapping replication updates and improve load balancing.

The scheduled frequency can also be set to replicate to the target based on the automatic snapshot schedules configured in the project or share. When this option is selected, replication updates are performed when the scheduled automatic snapshots are created.



Update frequency ☒ Scheduled ☐ Continuous

Schedule :: Snapshots

FREQUENCY

auto-snapshots ▼

- **Continuous** - Replication is executed continuously. As soon as one replication update is complete, a subsequent update is started. Changes are transmitted as frequently as possible, resulting in sending a constant stream of all filesystem changes to the target system. For filesystems with a lot of churn (many files created and destroyed in short intervals), this can result in replicating much more data than is actually necessary. However, as long as replication can keep up with the data changes, this results in the minimum amount of data lost in the event of a data-loss disaster on the source system.

Replication Action and Package Relationship

A replication action and package are bound to each other. If the package is somehow corrupted or destroyed, the action cannot send replication updates, even if the target still has the data and snapshots associated with the action. Similarly, if the action is destroyed, the package will be unable to receive new replication updates (even if the source still has the same data and snapshots). A warning will occur, in both the BUI and CLI, if you attempt to perform an operation that would destroy the action-package connection. If an error or explicit administrative operation breaks the action-package connection such that an incremental update is no longer possible, you must sever or destroy the package and action, then create a new action on the source.

Note - The appliance avoids destroying data on the target unless explicitly requested to do so by the administrator. As a result, if the initial replication update for an action fails after replicating some of the data and leaving incomplete data inside the package, subsequent replication updates using the same action will fail because the appliance cannot overwrite the previously received data. To resolve this, administrators should destroy the existing action and package, create a new action, then restart replication.

Related Topics

- [“Replication Action Properties” on page 594](#)
- [“Replication Packages” on page 621](#)

Replication Action Properties

The replication action properties in the BUI and CLI differ slightly, as described in the following table.

Note - When you change a replication action property, the new setting takes effect with the next replication update unless specified otherwise.

TABLE 131 Replication Action Properties (BUI and CLI)

BUI Property	CLI Property	Description
Disable compression	compression	The replication stream is compressed by default. Disable if compression is provided by another means, such as a WAN accelerator. For more information, see “Compressed Replication” on page 620 .
Update frequency: Scheduled or Continuous	continuous (True or False)	Select Scheduled or Continuous. For more information, see Replication Update Frequency .
Enable deduplication	dedup	When set, enables deduplication on replication streams. For more information, see “Deduplicated Replication” on page 600 .
Power icon 	enabled	When enabled (true in the CLI), replication updates can be sent. When disabled (false in the CLI), the power icon  is not highlighted and replication updates cannot be sent.
Export data path	export_path	Specifies the path to an NFS share for this action, using the format: nfs://server/path. This property exports the replication stream to a file on an NFS server, which can be physically moved to the remote target site, and then imported to a replication target. For procedures, see Creating an Offline Replication BUI, CLI .
Include clone origin as data	include_clone_origin	Controls the replication of each share that was cloned from a share that is external to the replication package on the target. Select this option to insert a complete copy of the clone origin snapshot's data into the replica of the clone. If you deselect this option, a clone created from an external origin snapshot will share storage with the replica of the clone origin snapshot that resides in replication target's pool. Sharing storage saves space, but if the replication target does not contain the external clone origin snapshot, the replication of the clone will fail. For details, see “Cloning a Replication Package or Share” on page 621 .
Include Snapshots	include_snaps	Determines whether replication updates include non-replication snapshots. For details, see “Replication Snapshot Management” on page 612 .
Limit bandwidth	max_bandwidth	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second). Use this property to limit bandwidth, especially during an initial replication update. If you change this property during a replication update, the new setting takes effect immediately.

BUI Property	CLI Property	Description
Pool	pool	Storage pool on the target where this project will be replicated. This property is specified when an action is initially configured and not shown thereafter.
Recovery point objective [] <i>unit of time</i>	recovery_point_objective	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage. The recovery point objective (RPO) can be specified as days, hours, minutes, or seconds. Set this property when creating or editing a replication action. This property is only used in conjunction with the replica lag warning and error alerts.
Replica lag error alert [] % of Recovery Point Objective	replica_lag_error_alert	Specifies a limit represented as a percentage of the RPO. The source appliance generates a major alert when the replica lag exceeds the specified limit. When the replica lag falls below this value, a minor alert indicates the replica lag is within the error limit.
Replica lag warning alert [] % of Recovery Point Objective	replica_lag_warning_alert	Specifies a limit represented as a percentage of the RPO. The source appliance generates a minor alert when the replica lag exceeds the specified limit. When the replica lag falls below this value, a minor alert indicates the replica lag is within the warning limit.
Retain user snapshots on target	retain_user_snapshots	Retains user-generated snapshots on the replication target for the associated action, even after the snapshots are destroyed on the source appliance. The snapshots retained on the target can then be used as part of a disk-to-disk or data backup solution. For details, see “Managing User-Generated Snapshots” on page 557 .
Target	target	Unique identifier for the replication target system. This property is specified when an action is initially configured and cannot be edited.
Target Pool	target_pool	Storage pool on the target where this project will be replicated. This property is displayed when editing an existing replication action.
Enable SSL-encryption	use_ssl	When set, encrypts data on the wire using SSL. Using this feature may have an impact on per-action replication performance.

The following table describes the CLI child nodes of replication actions.

TABLE 132 Replication Action Child Nodes (CLI only)

Replication Action Child Node	Description
autosnaps	Automatically scheduled snapshots, with sub-node automatic. To change the number of snapshots retained for a replication package, select an individual automatic node and modify the keep property. For information about automatic snapshot management, see “Replication Automatic Snapshot Management” on page 614. For more information about configuring auto snapshots, see “Configuring Automatic Snapshot Retention on a Target (CLI)” on page 523.
stats	Statistics for the most recent replication update, and the accumulated statistics over the lifetime of this replication action. Statistics are updated after replication update completion.

Replication Action Child Node	Description
	For more information on the stats node properties, see “Deduplicated Replication” on page 600 .

The following table describes the CLI read-only replication action properties.

TABLE 133 Replication Action Properties (CLI Read-only)

CLI Property	Description
average_throughput	Describes the average replication throughput.
bytes_sent	Number of bytes sent to the replication target.
estimated_size	Estimated size of the data to be replicated.
estimated_time_left	Estimated time remaining until completion of the replication update.
export_pending	Indicates whether an export is pending. Value is true or false.
id	The replication action's unique identifier. The identifier can be used to select an action or its associated replication package using the syntax: <code>select id=<uniqueid></code> .
last_result	The result of the last update. Value is success or failed.
last_sync	The last time an update was successfully sent. This value may be unknown if the system has not sent a successful update since boot.
last_try	The last time an update was attempted. This value may be unknown if the system has not attempted to send an update since boot.
next_update	Date and time when the next attempt will be made. This value could be a date (for a scheduled update), <code>Sync now</code> , or <code>continuous</code> .
offline	Indicates whether the replication update is offline. Value is true or false.
origin_pkg_id	The unique identifier of the replication package from which this action was created when the package was reversed. The origin package identifier is displayed after the first successful replication update completes.
replica_data_timestamp	Creation time of the snapshot used in the last successful update.
replica_lag	Current replica lag with a format of <code>hh:mm:ss</code> .
replica_lag_over_error_limit	true when the replica lag has exceeded the error limit specified by the combination of the <code>recovery_point_objective</code> and the <code>replica_lag_error_alert</code> threshold.
replica_lag_over_warning_limit	true when the replica lag has exceeded the warning limit specified by the combination of the <code>recovery_point_objective</code> and the <code>replica_lag_warning_alert</code> threshold.
replication_of	Project or share name (under project) where the action is configured.
source_pool	Pool name of the source project/share.
state	Describes if replication is in progress or not. Values are <code>sending</code> or <code>idle</code> .
state_description	Specifies details on replication progress. The different states are: ■ Connecting to replication target

CLI Property	Description
	■ Receiving checkpoint from target ■ Estimating size of update ■ Building deduplication tables ■ Sending update State Building deduplication tables is displayed only if the project or share has deduplication enabled.
target_id	The unique identifier of the replication target object that describes the target of this replication action.

Related Topics

- [“Creating a Replication Action \(BUI\)” on page 518](#)
- [“Creating a Replication Action \(CLI\)” on page 520](#)

Replication Package Properties

The replication package properties in the BUI and CLI differ slightly, as described in the following table.

TABLE 134 Replication Package Properties (BUI and CLI)

BUI Property	CLI Property	Description
Source host	source_name	Name of this package's source.
Source pool	source_pool	Storage pool on the source where this project is replicated from. This property is specified when an action is initially configured.
	enabled	When enabled (true in the CLI), replication updates can be received. When disabled (false in the CLI), the  power icon is not highlighted and replication updates cannot be received.
Data timestamp	data_timestamp	Creation time of the snapshot used in the last successful update.
Last sync	last_sync	Completion time of last successful update.
Last attempt	last_try	Completion time of last update attempt.
Import data path	import_path	External media URI for pending import.
Status	state	Current state of replication updates.
Last result	last_result	Result of last update attempt. Value is success or failed.

The following table describes the CLI read-only replication package properties.

TABLE 135 Replication Package Properties (CLI Read-only)

CLI Property	Description
id	Unique identifier of the replication package. The identifier can be used to select a package, by entering <code>select id= unique-id</code> .
source_asn	ID of the replication action associated with this package.
source_ip	IP address of this package's source.
target_pool	Target pool for this package.
replica_of	Replicated dataset in this package.
state_description	Current state of replication updates.
offline	Indicates that the package is waiting for an offline update.

Replication Storage Pools

When a replication action is initially configured, the administrator is given a choice of which storage pool on the target should contain the replicated data. The storage pool containing an action cannot be changed once the action has been created. Creating the action creates the empty package on the target in the specified storage pool. After this operation the source has no knowledge of the storage configuration on the target. It does not keep track of which pool the action is being replicated to, nor is it updated with storage configuration changes on the target.

When the target is a clustered system, the chosen storage pool must be one owned by same controller which owns the IP address used by the source for replication because only those pools are always guaranteed to be accessible when the source contacts the target using that IP address. This is exactly analogous to the configuration of NAS clients, for example, NFS or SMB, where the IP address and path requested in a mount operation must follow the same constraint. When performing operations that change the ownership of storage pools and IP addresses in a cluster, administrators must consider the impact to sources replicating to the cluster. There is currently no way to move packages from one storage pool to another.

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication” on page 511](#)

Project vs. Share Replication

Although it is possible to configure remote replication on both the project level and share level, project-level replication is recommended for these reasons. The appliance allows administrators to configure remote replication on both the project and share level.

- Replication snapshots are always taken at the project level. Multiple share level replications in a single project can create a substantial amount of overhead and consume space on the pool.
- When reversing share-level replication, the share is placed in its own project. This means that replication reversals will end up splitting the share away from the other shares in the project, unless they are all replicated together.

Like other properties configurable on the Shares screen, each share can either inherit or override the configuration of its parent project. Inheriting the configuration means not only that the share is replicated on the same schedule to the same target with the same options as its parent project is, but also that the share will be replicated in the same stream using the same project-level snapshots as other shares inheriting the project's configuration. This may be important for applications which require consistency between data stored on multiple shares. Overriding the configuration means that the share will not be replicated with any project-level actions, though it may be replicated with its own share-level actions that will include the project. It is not possible to override part of the project's replication configuration and inherit the rest.

More precisely, the replication configuration of a project and its shares define some number of replication *groups*, each of which is replicated with a single stream using snapshots taken simultaneously. All groups contain the project itself (which essentially just includes its properties). One project-level group includes all shares inheriting the replication configuration of the parent project. Any share that overrides the project's configuration forms a new group consisting of only the project and the share itself.

For example, suppose you have:

- A project home and shares bill, cindi, and dave.
- home has replication configured with some number of actions.
- home/bill and home/cindi inherit the project's replication configuration.
- home/dave overrides the project's replication configuration, using its own configuration with some number of actions.

This configuration defines the following replication groups, each of which is replicated as a single stream per action using snapshots taken simultaneously on the project and shares:

- One project-level group including home, home/bill, and home/cindi.
- One share-level group including home and home/dave.

Note - Due to current limitations, do not mix project- and share-level replications within the same project. This avoids unpredictable results when reversing the replication direction or when replicating clones. For more details, see [“Replication Packages” on page 621](#) and [“Cloning a Replication Package or Share” on page 621](#).

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication” on page 511](#)

Replication Authorizations

The replication subsystem provides two user authorizations under the "Projects and Shares" scope:

Authorization	Details
rrsource	Allows administrators to create, edit, and destroy replication targets and actions. Additionally, it allows an administrator to send and cancel updates for replication actions.
rrtarget	Allows administrators to manage replicated packages, including disabling replication at the package level, cloning a package or its members, modifying properties of received datasets, and severing or reversing replication. Other authorizations may be required for some of these operations (like setting properties or cloning individual shares). See the available authorizations in the Projects and Shares scope for details.

The `rrsource` authorization is required to configure replication targets on an appliance, even though this is configured under the Remote Replication service screen. For help with authorizations, see [“Configuring Users” on page 202](#).

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication” on page 511](#)

Deduplicated Replication

Deduplicated Replication provides the ability to reduce the amount of data sent over the network by replication jobs. This feature is useful for reducing the on-the-wire data bandwidth requirements of replication, especially when using a high-latency, low-bandwidth, high-cost network.

Note - This feature imposes a cost in the form of pre-processing and increased memory overhead. The effectiveness of deduplication is highly data dependent, so it is strongly recommended to verify the deduplication savings with representative datasets prior to using this feature in a production environment. Deduplicated Replication is more efficient when there is more duplicate data.

Deduplicated Replication is disabled by default. It can be enabled for individual replication actions, as shown in the following BUI figure.

The screenshot shows the 'Add Replication Action' dialog box. The 'Properties' section is expanded, showing various configuration options. The 'Enable deduplication' checkbox is highlighted with a red rectangle. Other options include Target (cleo), Pool (p0), Export data path (nfs://), Limit bandwidth (0 M/s), Enable SSL-encryption (checked), Disable compression (unchecked), Include snapshots (checked), Retain user snapshots on target (unchecked), Include clone origin as data (unchecked), Recovery point objective (hours), Replica lag warning alert (% of Recovery Point Objective), Replica lag error alert (% of Recovery Point Objective), and Update frequency (Scheduled).

Deduplicated Replication Statistics

Each replication action has a `stats` node, which records information about the most recent replication update, as well as the accumulated statistics over the lifetime of the replication action.

These `stats` fields quantify:

- On-disk compression benefits
- Deduplication benefits
- Replication data stream compression benefits
- Replication update duration
- Deduplication tables construction time (before sending data)
- Deduplication tables maximum memory consumption

The `stats` node of a deduplicated replication stream has the following read-only properties:

TABLE 136 Replication Action: stats Node Properties

Property Name	Description
logical_bytes	Number of bytes that the replication update data stream would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.
phys_bytes	Number of bytes in the internal replication data stream prior to replication deduplication or replication data stream compression.
after_dedup	Number of bytes in the internal replication data stream after any deduplication of the replication data stream.
to_network	Number of bytes that the replication data stream compression pipeline delivered to the network. This shows the consequence of replication data stream compression, if enabled.
duration	Total time required to perform the replication update.
dd_table_build	Time spent building the deduplication tables prior to the actual transmission of the replication update.
dd_table_mem	Maximum amount of memory that was consumed by the deduplication tables.

To list the stats node fields, first navigate to the replication action, enter the stats node and then enter get

```
hostname:shares testproj action-001> stats
hostname:shares testproj action-001 stats>
hostname:shares testproj action-001 stats> get
Properties:
    replica_data_timestamp = Thu Apr 21 2016 06:14:58 GMT+0000 (UTC)
        last_sync = Thu Apr 21 2016 17:50:18 GMT+0000 (UTC)
            last_try = Thu Apr 21 2016 17:50:18 GMT+0000 (UTC)
                last_result = success
                    last_logical_bytes = 5.80401479T
                        last_phys_bytes = 3.57996902T
                            last_after_dedup = 953.489698G
                                last_to_network = 943.954802G
                                    last_duration = 11:35:26
                                        last_dd_table_build = 02:57:10
                                            last_dd_table_mem = 3.5273976G
                                                total_updates = 40
                                                    total_logical_bytes = 232.16591T
                                                        total_phys_bytes = 143.198761T
                                                            total_after_dedup = 90.2222261T
                                                                total_to_network = 90.0359976T
                                                                    total_duration = 404:34:00
                                                                        dd_total_updates = 20
                                                                            dd_total_logical_bytes = 116.080296T
                                                                                dd_total_phys_bytes = 71.5993804T
                                                                                    dd_total_after_dedup = 18.6228456T
                                                                                        dd_total_to_network = 18.4366172T
                                                                                            dd_total_duration = 231:48:40
```

```
dd_total_table_build = 59:03:20
dd_total_table_mem = 70.547952G
```

Recent replication statistics are also recorded as send alerts, which can be seen and accessed through the BUI and CLI. For more information, see [“Replication Alerts” on page 616](#).

Measuring Deduplicated Replication Statistics

When deduplication is enabled for a replication stream, the data is transformed through several layers of deduplication and compression. Data rates are measured and recorded as the data is transformed. These statistics are recorded in the stats node of a replication action.

To determine if deduplication was sufficiently effective for the replication action, examine the replication statistics.

Single Deduplicated Replication Update Benefits Comparison

- In the BUI, use the replication finish alerts to compare the `phys_bytes` and `after_dedup` statistics to gauge the benefit of deduplicated replication. For information on replication alerts, see [“Replication Alerts” on page 616](#).
- In the CLI, use the replication action stats node to compare `last_phys_bytes` and `last_after_dedup` statistics to gauge the benefit of deduplicated replication. For information on the stats node, see [“Deduplicated Replication Statistics” on page 601](#).

Averaged Deduplicated Replications Updates Benefits Comparison

- To gauge the average benefit of all deduplicated replication updates performed by this replication action, use the replication action stats node to compare statistics `dd_total_phys_bytes` and `dd_total_after_dedup`. For information on the stats node, see [“Deduplicated Replication Statistics” on page 601](#).

Replication Configuration for Clustered Appliances

Replication can be configured from any source appliance to any replication target regardless of whether each is part of a cluster and whether the appliance's cluster peer has replication configured in either direction.

The following rules govern the behavior of replication updates for clustered appliances:

- Replication updates for projects and shares are sent from whichever cluster peer has imported the containing storage pool.
- Storage pools that are owned by each controller can replicate to the same replication target, only if OS8.6.0 (or later) is installed on both controllers.
- Replication updates are received by whichever peer has imported the IP address configured in the replication action on the source. Administrators must ensure that the controller using this IP address will always have the storage pool containing the replica imported. This is ensured by assigning the pool and IP address resources to the same controller during cluster configuration.
- Replication updates (both to and from an appliance) that are in progress when an appliance exports the corresponding storage pool or IP address (as part of a takeover or failback) will fail. Replication updates using storage pools and IP addresses unaffected by a takeover or failback operation will be unaffected by the operation.

Related Topics

- [“Example: Replication Configuration for Clustered Appliances” on page 604](#)

Example: Replication Configuration for Clustered Appliances

The goal of this example is to configure replication properly to ensure that projects continue to replicate after a cluster takeover, cluster failback, or after performing reverse replication on a target appliance.

- [Configuration Guidelines](#)
- [Example: Configuring Replication for Clustered Appliances](#)
- [Replication Data Path Illustrated Examples](#)

Configuration Guidelines

When configuring replication for clustered appliances, follow these guidelines:

- Ensure that both replication source and target appliances are in the CLUSTERED state. For details, see [Table 13, “Cluster States,” on page 74](#).
- Select network interfaces and IP addresses to be used for replication traffic on the replication source and target appliances.
 - Select a singleton network interface. Unlike a private network interface, a singleton network interface will be taken over by the surviving controller following the loss

of one of the controllers in the cluster. Using a singleton interface ensures successful replication following a cluster takeover or failback transition. For more information about singleton interfaces, see [Table 12, “Cluster Resource Management,” on page 73](#).

- Ensure that the selected network interface on the source appliance and the pool from which the data will be replicated are both assigned to the same controller. This is always the case when the source cluster is in the CLUSTERED state.
- Similarly for the target cluster, the selected network interface on the target appliance and the pool into which the data will be replicated must both be assigned to the same controller. This association is guaranteed when the replication configuration is performed while the target cluster is in CLUSTERED state.
- The source and the target appliances must be able to successfully communicate using the selected network interfaces and IP addresses.
- Create static /32 host-based routing between target and source appliances to ensure that following replication reversal, the selected interface is used for outbound replication traffic when reversal has transformed the current target into a replication source.
- After the static route has been created, configure the replication target object on the source appliance using the selected IP address of the target.
- When the target appliance is in the OWNER state, all shared resources including network interfaces and storage pools are taken over and owned by the one surviving controller, the controller that is now in the OWNER state. On the controller in the OWNER state, it is possible to select a network interface that is assigned to one controller and use it to deliver replication traffic to a pool that is assigned to a different controller. When the controllers are returned to the CLUSTERED state, the network interfaces and storage pools are returned to their assigned controllers. Therefore, replication updates might not be possible because the source appliance will use the network interface on the target controller that no longer owns the pool. This configuration error cannot arise when replication configuration is performed while the target appliance is in the CLUSTERED state.

Example: Configuring Replication for Clustered Appliances

The example procedure uses the following source and target network interfaces and IP addresses:

The source appliance cluster consists of source controllers S1 and S2. Storage pool sp1 is assigned to S1 and pool sp2 is assigned to S2. The cluster network interfaces consist of:

- Private interface ixgbe0 on S1 with IP address 198.51.100.81/24
- Private interface ixgbe0 on S2 with IP address 198.51.100.82/24
- Singleton interface ixgbe1 with IP address 192.0.2.101/25 assigned to S1
- Singleton interface ixgbe2 with IP address 192.0.2.102/25 assigned to S2
- Singleton interface ixgbe3 with IP address 192.0.2.201/25 assigned to S1

- Singleton interface `ixgbe4` with IP address `192.0.2.202/25` assigned to S2

The appliance is Initially in the CLUSTERED state where:

- S1 owns `sp1`, `ixgbe1`, and `ixgbe3`
- S2 owns `sp2`, `ixgbe2` and `ixgbe4`

The target appliance cluster consists of controllers T1 and T2. Storage pool `tp1` is assigned to T1 and pool `tp2` is assigned to T2. The cluster network interfaces consist of:

- Private interface `ixgbe0` on T1 with IP address `198.51.100.83/24`
- Private interface `ixgbe0` on T2 with IP address `198.51.100.84/24`
- Singleton interface `ixgbe1` with IP address `192.0.2.103/25` assigned to T1
- Singleton interface `ixgbe2` with IP address `192.0.2.104/25` assigned to T2
- Singleton interface `ixgbe3` with IP address `192.0.2.203/25` assigned to T1
- Singleton interface `ixgbe4` with IP address `192.0.2.204/25` assigned to T2

The appliance is initially in the CLUSTERED state where:

- T1 owns `tp1`, `ixgbe1`, `ixgbe3`
- T2 owns `tp2`, `ixgbe2` and `ixgbe4`

The following steps describe how to configure replication using the CLI for projects Red, Blue, and Green.

1. Select network interfaces and IP addresses.
 - Start by selecting network interfaces and IP addresses for replication of project Red. Because the source S is in the CLUSTERED state, it is sufficient to ensure that the selected network interfaces and IP addresses are not private. Thus, on S1 use either `ixgbe1` or `ixgbe3`.
 - The same applies to target T, therefore, use either `ixgbe1` or `ixgbe3` on appliance T1. Because `ixgbe1` and `ixgbe3` on both S1 and T1 belong to the same subnet, select either to perform replication of project Red. For this example, select interface `ixgbe1` on S1 and on T1.
2. Set up a static route on S1.

The following example sets up the static route for replication of project Red on source controller S1:

```
S1:configuration net routing> create
S1:configuration net route (uncommitted)> set family=IPv4
                                     family = IPv4 (uncommitted)
S1:configuration net route (uncommitted)> set destination=192.0.2.103
                                     destination = 192.0.2.103 (uncommitted)
```

```

S1:configuration net route (uncommitted)> set mask=32
      mask = 32 (uncommitted)
S1:configuration net route (uncommitted)> set interface=ixgbe1
      interface = ixgbe1 (uncommitted)
S1:configuration net route (uncommitted)> set gateway=192.0.2.1
      gateway = 192.0.2.1 (uncommitted)
S1:configuration net route (uncommitted)> commit
S1:configuration net routing> list
ROUTE      DESTINATION      GATEWAY      INTERFACE      TYPE      STATUS
...
route-003  192.0.2.103/32    192.0.2.1    ixgbe1         static    active

```

3. Set up a static route on T1.

The following example sets the static route for replicating project Red on target controller T1:

```

T1:configuration net routing> create
T1:configuration net route (uncommitted)> set family=IPv4
      family = IPv4 (uncommitted)
T1:configuration net route (uncommitted)> set destination=192.0.2.101
      destination = 192.0.2.101 (uncommitted)
T1:configuration net route (uncommitted)> set mask=32
      mask = 32 (uncommitted)
T1:configuration net route (uncommitted)> set interface=ixgbe1
      interface = ixgbe1 (uncommitted)
T1:configuration net route (uncommitted)> set gateway=192.0.2.1
      gateway = 192.0.2.1 (uncommitted)
T1:configuration net route (uncommitted)> commit
T1:configuration net routing> list
ROUTE      DESTINATION      GATEWAY      INTERFACE      TYPE      STATUS
...
route-003  192.0.2.101/32    192.0.2.1    ixgbe1         static    active

```

4. Create a replication target on S1.

The following example creates the replication target object on S1 to be used to replicate project Red from sp1 to tp1:

```

S1:shares replication targets>target

S1:shares replication target (uncommitted)> set hostname=192.0.2.103
      hostname = 192.0.2.103 (uncommitted)
S1:shares replication target (uncommitted)> set label=t1-1
      label = t1-1 (uncommitted)
S1:shares replication target (uncommitted)> set root_password=*****

```

```
root_password = (set) (uncommitted)
S1:shares replication target (uncommitted)> commit
```

5. Create a replication action for each project.

- Replicate project Red from pool sp1 to tp1
- Replicate project Blue from pool sp1 to pool tp2
- Replicate project Green from pool sp2 to tp2

The following example creates the replication action for project Red:

```
S1:> shares select Red replication action
S1:shares Red action (uncommitted)> set target=t1-1

target=t1-1 (uncommitted)
S1:shares Red action (uncommitted)> set pool=tp1
pool=tp1 (uncommitted)
S1:shares Red action (uncommitted)> commit
```

6. Set up to replicate project Blue from pool sp1 to tp2.

Start with interface and address selection, and select interfaces S1/ixgbe3 and T2/ixgbe4, knowing that both S and T are in the CLUSTERED state and that the interface addresses are on the same subnet, 192.0.2.128/25. Next, define static routes on both appliances similar to the above examples. Then create replication target object t2-2 on S1, and create the replication action on S1 for project Blue using target object t2-2.

7. Set up to replicate project Green from pool sp2 to tp2.

Start with interface selection and select interfaces S2/ixgbe2 and T2/ixgbe2. Create static routes on S2 and T2 using the selected interfaces and their addresses, define replication target object t2-1 using address of T2/ixgbe2, and finally create the replication action for project Green using target object t2-1.

8. Initiate replication for all three actions.


- a. Start with project Red:

```
S1:> shares select Red replication select action-000
S1:shares Red action-000> sendupdate
```

- b. Initiate replication for the actions for projects Blue and Green by following the previous example.

Replication Data Path Illustrated Examples

The following figures illustrate the replication data paths during replication updates for the replication actions for projects Red, Blue, and Green:

FIGURE 28 Normal Replication Data Path

Assume that controller T2 has been taken down for a maintenance. T1 performed a takeover and now owns all of the resources. If replication updates for projects Blue and Green are in progress during the takeover, they will be canceled. After T1 takes over, these replication updates can be resumed manually, or they will be resumed automatically if schedules are configured for the corresponding replication actions.

After controller T1 has completed the takeover, it owns interfaces ixgbe2 and ixgbe4 which are necessary to continue replication updates for projects Blue and Green. The following figure shows the replication data path after T1 completed the takeover.





FIGURE 30 Data Path After Failback on T1 and Takeover on S1**Related Topics**

- [“Configuring Replication for a Clustered Configuration” on page 526](#)
- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication” on page 511](#)

Replication Snapshots and Data Consistency

The source appliance replicates snapshots atomically to the target, meaning the contents of the replica always exactly matches the source's share at the time the snapshot was taken. Because the snapshots for all shares sent in a particular group are taken at the same time, the entire package contents after the completion of a successful replication update exactly matches the group's content when the snapshot was created on the source.

However, each share's snapshots are replicated separately, so it is possible for some shares within a package to have been updated with a snapshot that is more recent than those of other shares in the same package. This is true during a replication update and after a failed replication update.

To summarize:

- Each share is always point-in-time consistent on the target.
- When no replication update is in progress and the previous replication update succeeded, each package's shares are also point-in-time consistent with each other.
- When a replication update is in progress or the previous update failed, package shares may be inconsistent with each other, but each one will still be self-consistent. If package consistency is important for an application, one must clone the replication package, which always clones the most recent successfully received snapshot of each share.

Related Topics

- [“Replication Snapshot Management” on page 612](#)

Replication Snapshot Management

Snapshots are the basis for replication. The source and target must always share a common snapshot in order to continue replicating incrementally, and the source must know which is the most recent snapshot that the target has. To facilitate this, the replication subsystem creates and manages its own snapshots. Administrators generally do not need to be concerned with them, but the details are described here since snapshots can have significant effects on storage utilization.

Each replication update for a particular action consists of the following steps:

- Determine whether this is an incremental or full update based on whether:
 - An attempt was made to replicate this action before and
 - The target already has the necessary snapshot for an incremental update
- Take a new project-level snapshot.
- Send the update. For a full update, send the entire group's contents up to the new snapshot. For an incremental update, send the difference between from the previous (base) snapshot and the new snapshot.
- Record the new snapshot as the base snapshot for the next update and destroy the previous base snapshot (for incremental updates). The base snapshot remains on the target until the next update is received at which point it is the first thing that is destroyed.

This has several consequences for snapshot management:

- During the first replication update and after the initial update when replication is not active, there is exactly one project-level snapshot for each action configured on the project or any share in the group. A replication action may create snapshots on shares that are in the same project as the share(s) in the group being replicated by the action, but that are not being sent as part of the update for the group.
- During subsequent replication updates of a particular action, there may be two project-level snapshots associated with the action. Both snapshots may remain after the update completes in the event of failure where the source was unable to determine whether the target successfully received the new snapshot (as in the case of a network outage during the update that causes a failure).
- None of the snapshots associated with a replication action can be destroyed by the administrator without breaking incremental replication. The system will not allow administrators to destroy snapshots on either the source or target that are necessary for incremental replication. To destroy such snapshots on the source, one must destroy the action (which destroys the snapshots associated with the action). To destroy such snapshots on the target, one must first sever the package (which destroys the ability to receive incremental updates to that package).
- Administrators must not rollback to snapshots created prior to any replication snapshots. Doing so will destroy the later replication snapshots and break incremental replication for any actions using those snapshots.
- Replication's usage of snapshots requires that administrators using replication understand space management on the appliance, particularly as it applies to snapshots.

Intermediate Replication Snapshots

A replication action can be set to include non-replication snapshots. When the property "Include Snapshots" is set, replication updates include the non-replication snapshots created after the previous replication update (or since the share's creation, in the case of the first full update). This includes automatic snapshots and administrator-created snapshots. This property can be disabled to skip these snapshots and send only the changes between replication snapshots with each update.

The action property `include_snaps` should be enabled in order to replicate any intermediate snapshots, including auto snapshots.

Related Topics

- [“Space Management for Shares” on page 436](#)
- [“Managing User-Generated Snapshots” on page 557](#)

Replication Automatic Snapshot Management

The automatic scheduled snapshots feature allows for automatically creating and destroying snapshots for projects and/or shares based on administrator-provided schedules. The schedule specifies when to create an automatic snapshot, and how many snapshots to retain. Several schedules can be created for a project or a share.

With Remote Replication, snapshots, including automatic snapshots, can be included in replication updates and will be available on the replication target as part of the corresponding replication package.

By default, the number of automatic snapshots retained on the target corresponds to the retention setting (Keep At Most) in the project's or share's snapshot schedule.

The screenshot shows the 'Snapshots' tab for a project named 'Proj'. Under the 'Schedules' section, the 'FREQUENCY' is set to 'every hour' and the 'scheduled time' is '00 minutes past the hour'. The 'KEEP AT MOST' value is set to '5'.

Replication actions can be configured to retain a separate, specific number of automatic snapshots on the target throughout replication updates.

The screenshot shows the 'Schedule : Snapshots' configuration page. Under 'Auto snapshot retention policies', the 'independent' radio button is selected. In the 'Configure auto snapshot retention on the target' section, the 'FREQUENCY' is 'every hour' and the 'KEEP AT MOST' value is '10'.

Reverse Replication and Automatic Snapshot Management

When reversing replication, automatic snapshot retention settings are preserved: The source and target will continue to maintain their retention settings.

Example:

- Source A has configured automatic snapshots and retains **5** snapshots on Source A.
- Through a replication action on Source A, Target B has been configured to retain **10** automatic snapshots.

After reverse replication, the source and target have switched to **Source B** and **Target A**.

- Now Source B has the automatic snapshot schedule, still retaining **10** snapshots.
- Target A is still configured to retain **5** snapshots. This retention setting is now configurable through the replication action on Source B.

Performing another reverse replication will revert the source and target to their original configurations.

For more information on configuring automatic snapshot retention on a target, see:

- [“Configuring Automatic Snapshot Retention on a Target \(BUI\)” on page 522](#)
- [“Configuring Automatic Snapshot Retention on a Target \(CLI\)” on page 523](#)

iSCSI Configurations and Replication

Replication updates include most of the configuration specified on the Shares screen for a project and its shares. This includes any target groups and initiator groups associated with replicated LUNs.

When using non-default target groups and initiator groups, administrators must ensure that the target groups and initiator groups used by LUNs within the project also exist on the replication target. If the target group or initiator group does not exist on the target system, a clone, sever, or reverse replication will fail. An error message reports that the initiator or target group name was either deleted or renamed on the target system.

The SCSI GUID associated with a LUN is replicated with the LUN. As a result, the LUN on the target appliance will have the same SCSI GUID as the LUN on the source appliance. Clones of replicated LUNs, however, will have different GUIDs (just as clones of local LUNs have different GUIDs than their origins).

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Remote Replication” on page 511](#)

Resumable Replication

When a scheduled or continuous replication update is interrupted due to a network failure, system outage, or operator action, data transfer *automatically* resumes from the point of interruption. This feature is available with OS8.7.0 or later, and must be installed on the source and replication target. For example, if a scheduled replication transfers 100 bytes of data before

a failure occurs, the next replication update resumes data transfer at byte 101. The estimated data size, shown in the replication progress monitor, includes data to be sent as part of both actions.

If a failure occurs during a manual replication update, the update is not automatically retried, however the data transfer will resume from the point of interruption with the next replication update.

Replication Alerts

Alerts are posted when any of the following replication events occur:

- A manual or scheduled replication update starts or finishes successfully (both source and target).
- Any replication update fails, including explicit cancellation by an administrator (both source and target).
- A scheduled replication update is skipped because another update for the same action is already in progress.
- A continuous replication starts for the first time, fails, or resumes after a failure.
- A replica time lag exceeds its specified threshold.


To view alerts in the BUI, go to Maintenance > Logs > Alerts.

To view alerts in the CLI, enter `maintenance logs`, and then enter `select alert`. Enter `show` to list the alerts.

```
hostname:> maintenance logs
hostname:maintenance logs> select alert
hostname:maintenance logs alert> show
```

Send Finish Alerts

When the system finishes replicating to a target, an alert appears at the top of the BUI window, providing statistics about the most recent replication update:



In the CLI, the completed replication update is reflected in the update statistics, as shown in the following table:

TABLE 137 Replication Update Statistics (CLI Read-Only)

Property	Description
logical_bytes	Number of bytes that the replication update data stream would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.
phys_bytes	Number of bytes in the internal replication data stream prior to replication deduplication or replication data stream compression.
to_network	Number of bytes that the replication data stream compression pipeline delivered to the network. This shows the consequence of replication data stream compression, if enabled.
duration	Total time required to perform the replication update.

These replication send finish alerts are also recorded in the system's Alert Log.

Deduplicated Replication Finish Alerts

Alerts for deduplicated replication streams provide additional deduplication statistics.

TABLE 138 Deduplicated Replication Update Statistics

Property	Description
after_dedup	Number of bytes in the internal replication data stream after any deduplication of the replication data stream.
dd_table_build	Time spent building the deduplication tables prior to the actual transmission of the replication update.
dd_table_mem	Maximum amount of memory that was consumed by the deduplication tables.

For more information on replication statistics, see [“Deduplicated Replication” on page 600](#).

Replication Failures

Individual replication updates can fail for a number of reasons. The appliance reports the reason for the failure in alerts posted on the source appliance or replication target, or on the Replication screen for the action that failed. You may be able to get details on the failure by clicking the orange alert icon representing the action's status.

The following are some common replication failures:

Failure	Details
Cancelled	The replication update was cancelled by an administrator. Replication can be cancelled on the source or target.

Failure	Details
Network connectivity failure	The appliance was unable to connect to the replication target due to a network problem. Check for a misconfiguration on the source, target, or the network.
Peer verification failed	The appliance failed to verify the identity of the target. This occurs most commonly when the target has been reinstalled or factory reset. A new replication target must be configured on the source appliance for a target which has been reinstalled or factory reset in order to generate a new set of authentication keys. See “ Replication Targets ” on page 590.
Peer RPC failed	A remote procedure call failed on the target system. This occurs most commonly when the replication target is running incompatible software. See Oracle ZFS Storage Appliance: Remote Replication Compatibility [Doc ID 1958039.1] https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1 .
Name collision	Replication of <i><project/share></i> from <i><source></i> failed due to a name collision with <i>@<snapname></i> being held on the target for NDMP. To recover, rename (or remove) the snapshot on the replication source that has the same name as the snapshot held by NDMP on the target (the one named in the alert), unless it starts with <i>.rr</i> . Then either perform a manual sync or allow the replication source to automatically retry the replication update.
No package	Replication failed because no package exists on the target to contain the replicated data. Since the package is created when configuring the action, this error typically happens after an administrator has destroyed the package on the target. This error could also occur if the storage pool containing the package is not imported on the target system, which may occur if the pool is faulted or if storage or networking has been reconfigured on the replication target.
Disabled	Replication failed because it is disabled on the target. Either the replication service is disabled on the target or replication has been disabled for the specific package being replicated.
Target busy	Replication failed because the target system has reached the maximum number of concurrent replication updates. The system limits the maximum number of ongoing replication operations to avoid resource exhaustion. When this limit is reached, subsequent attempts to receive updates will fail with this error, while subsequent attempts to send updates will queue up until resources are available.
Target is missing	The most recent replication update failed because the target is missing. If the target is no longer configured on the source, the action will be permanently disabled.

Failure	Details
	If this error occurs, destroy the replication action and reconfigure the replication target and action.
Out of space	Replication failed because the source system had insufficient space to create a new snapshot. This may be because there is no physical space available in the storage pool or because the project or one of its shares would be over quota because of reservations that do not include snapshots.
Key Unavailability	Replication failed because the encryption key used by the share is not available either on the source or target system. Review the alerts on both the source and replication target to ensure the key is available on both systems. See “Replicating an Encrypted Share” on page 656 for information about replicating encrypted shares and projects.
Incompatible target	Replication failed because the target system is unable to receive the source system's data stream format. This can happen as a result of upgrading a source system and applying deferred updates without having upgraded and applied the same updates on the target. For deferred updates that have remote replication implications, see Oracle ZFS Storage Appliance: Remote Replication Compatibility [Doc ID 1958039.1] https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1 .
iSCSI initiator/target missing	A replication clone, sever, or reverse operation failed because the initiator group or target group LUNs do not exist for the LUNs included in the replication package. The initiator or target group name was either deleted or renamed on the replication target.
Misc	Replication failed, but no additional information is available on the source. Check the alert log on the target system and if necessary contact support for assistance. Some failure modes that currently fall into this category include insufficient disk space on the target to receive the update and attempting to replicate a clone whose origin snapshot does not exist on the target system.

A replication update fails if any part of the update fails. The shares inside a project are replicated serially and changes are not rolled back from a failed update. As a result, when an update fails, some shares on the target may be up to date while others are not. For more information, see [“Replication Snapshots and Data Consistency” on page 611](#).

When a scheduled or continuous replication fails, the system waits several minutes and tries again. The system will continue retrying failed scheduled or continuous replications indefinitely. At any point during the retry procedure, initiating a manual update will immediately begin a retry, circumventing the usual delay between successive retries. If the

manual update completes successfully, it terminates the retry sequence and the replication action reverts to its normal scheduled or continuous updates.

For more information about failed or interrupted replication updates, see [“Resumable Replication” on page 615](#).

When a replication update is in progress and another update is scheduled, the scheduled replication is deferred until the previous update completes, and an alert is posted.

Related Topics

- *How to Troubleshoot Replication Issues* (Doc ID 1397959.1) on [My Oracle Support \(http://support.oracle.com/\)](http://support.oracle.com/)

Compressed Replication

The compressed replication feature improves performance when replicating compressible data between source and target sites that have limited bandwidth. Before a replication stream is sent to the target, it is automatically compressed at a rate based on current CPU utilization and network I/O throughput. The replication stream is then decompressed when received by the replication target. If any part of the data is not compressible, that portion will be sent as if compression were disabled.

All replication streams will be compressed, unless you explicitly disable compression. If your WAN equipment provides compression, for example through a WAN accelerator, disable the compression feature by following the procedure [Disabling Replication Compression BUI](#), [CLI](#).

The source appliance and replication target require software version 2013.1.4.0 (or later) to support replication compression. If the target has an earlier version, a warning icon  is displayed next to the target name. You will need to update the replication target to the minimum version.

You can view replication performance statistics on the source appliance, under the progress bar for the replication.

Actions		
TARGET ▲	UPDATES	STATUS
 hog-10g Manual	Never Synced Never Attempted	  43% of 10.0G @548MB/s (-00:00:11)

Replication Packages

Packages are containers for replicated projects and shares. Each replication action on a source appliance corresponds to one package on a replication target.


You can browse replicated projects, shares, snapshots, and properties much like local projects and shares, using the BUI or CLI. However, because replicated shares must exactly match their counterparts on the source appliance, many management operations are not allowed inside replication packages.

You can modify the following properties of replicated projects and shares:

- **Reservation, Compression, Copies, Deduplication, and Caching** - These properties can be changed on the replication target to effect different cost, flexibility, performance, or reliability policies on the replication target from the source.
- **Mountpoint and Sharing Properties** (e.g., sharenfs, SMB resource name) - These properties control how shares are exported to NAS clients and can be changed to effect different security or protection policies on the replication target from the source.

Such property modifications persist across replication updates.

FIGURE 31 Managing Replication Package Properties



Related Topics

- [“Project and Share Properties” on page 403](#)
- [“Severing Replication” on page 624](#)

Cloning a Replication Package or Share

A *clone* of a replicated package is a local, mutable project that can be managed like any other project on the system. When the clone project is created, the most recently received snapshot

of the replicated shares is used to create the shares within the clone project. These clones share storage with their origin snapshots the same way that clones of share snapshots do (see Cloning a Snapshot [BUI](#), [CLI](#)). This mechanism can be used to failover in the case of a catastrophic problem at the replication source, or simply to provide a local version of the data that can be modified.

As long as a clone exists, its origin snapshot cannot be destroyed. When destroying the snapshot (possibly as a result of destroying the share, project, or replication package of which the snapshot is a member), the system warns administrators of any dependent clones that will be destroyed by the operation. Note that snapshots can also be destroyed on the source at any time and such snapshots are destroyed on the target as part of the subsequent replication update. If such a snapshot has clones, the snapshot will not be destroyed until the last clone has been destroyed.

Replicating Clones

When replicating clones, it is important to understand the relationship between a clone replica and its origin snapshot. By default, the replica of a clone maintains its relationship with its origin snapshot, mandating that a replica of the origin snapshot also exist on the target. A replica of a clone origin snapshot must reside in the same pool as the clone, but does not have to be in the same project.

To maintain the relationship between a replicated clone and its origin snapshot, the origin snapshot must be:

- Replicated to the target before the initial replication of the clone or
- Replicated as part of the same update.

This restriction is not enforced by the appliance software, but must be followed to ensure a successful replication update.

There are several ways to ensure successful replication of a clone so it maintains its relationship with its origin snapshot:

- If the clone's origin snapshot is in the same project, use project-level replication.
- If the share containing the clone origin snapshot is not in the same project or if the clone or its origin share have been omitted from project-level replication, replicate the origin share first and then replicate the clone using project-level or share-level replication.
- On the target system, do not destroy the origin of the clone unless you also intend to destroy the clone itself.

To ensure that the origin snapshot is sent to the target, always set the property "Include snapshots" for the origin's replication action.

Just as a clone and its origin snapshot conserve space on the source appliance, a replicated clone and its replicated origin snapshot conserve space on the replication target. If space conservation

on the replication target is less important, the administrator may set the property `Include clone origin as data`. When this property is set, and the origin snapshot of a clone is *not* replicated in the same update as the clone, the source appliance inserts a copy of the clone origin's data content into the replica clone. Thus, there is no need to replicate the clone origin share first, but the copy of the clone origin data consumes additional storage space on the target.

When `Include snapshots` and `Include clone origin as data` are both set, the replica clone contains only the snapshots that are present in the clone on the source. The source appliance inserts the clone origin data content, not the clone origin snapshots, into the replica clone. This ensures that the snapshots present in the replica clone match the snapshots present in the clone on the source.

The property `Include clone origin as data` does not affect the replication of a clone and its origin snapshot when they are both replicated in the same update. When replicated together by the same replication action, the relationship between the clone and its origin snapshot is preserved and the space sharing benefit is retained on the target.

Related Topics

- [“Project vs. Share Replication” on page 598](#)
- Cloning a Replication Package [BUI](#), [CLI](#)
- [“Project and Share Properties” on page 403](#)


Exporting Replicated Filesystems

Replicated filesystems can be exported read-only to NAS clients. This can be used to verify the replicated data or to perform backups or other intensive operations on the replicated data (offloading such work from the source appliance).

The filesystem's contents always match the most recently received replication snapshot for that filesystem. This may be newer than the most recently received snapshot for the entire package, and it may not match the most recent snapshot for other shares in the same package. For details, see [“Replication Snapshots and Data Consistency” on page 611](#).

Replication updates are applied atomically at the filesystem level. Clients looking at replicated files will see replication updates as an instantaneous change in the underlying filesystem. Clients working with files deleted in the most recent update will see errors. Clients working with files changed in the most recent update will immediately see the updated contents.

Replicated filesystems are not exported by default. They are exported by modifying the exported property of the project or share using the BUI or CLI:

FIGURE 32 Inherited Properties

Inherited Properties

Export ☒

Mountpoint

Read only ☐

Update access time on read ☒

This property is inherited like other share properties. This property is not shown for local projects and shares because they are always exported. Additionally, severing replication (which converts the package into a local project) causes the package's shares to become exported.

Replicated LUNs currently cannot be exported. They must be first cloned or the replication package severed in order to export their contents.

Related Topics

- [“Remote Replication Workflow” on page 511](#)
- [“Inherited Properties” on page 404](#)

Severing Replication

A replication package can be converted into a local, writable project that behaves just like other local projects (that is, without the management restrictions applied to replication packages) by severing the replication connection. Severing a replication package can be used to migrate data between appliances or in other scenarios that do not involve replicating the received data back to the source appliance.

If a replication update is performed during or after a sever operation, the update will fail with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.

A new replication action and a full update of the same project is required to send replication updates to a new replication package.

To avoid mount point or SMB name conflicts, resolve the conflicts before severing the replication package by reconfiguring the project, share mount points, or SMB resource names.

Because all local shares are always exported, and might be shared over SMB, the sever operation will fail if any mount points or SMB resource names conflict between replicated filesystems and other filesystems on the system.

Related Topics

- Severing a Replication Package [BUI](#), [CLI](#)
- [Disaster Recovery with Remote Replication](#)
- [Managing User-Generated Snapshots](#)


Reverse the Direction of Replication

The direction of the replication can be reversed to support two-system disaster recovery plans and disk-to-disk backups.

Reversing Replication for Disaster Recovery

The reverse replication operation converts the replication package into a local project. This operation also configures a replication action on the new local project for incremental replication back to the source appliance. The first update attempt will convert the original project on the source system into a replication package and roll back any changes made since the last successful replication update from that system.

The following figure describes a typical reverse replication sequence of events.

FIGURE 33 Using Remote Replication for Disaster Recovery

Legend	Description
1	The production system is the source appliance serving the client workload and replicating to the replication target located at a recovery site. A complete failure of the source appliance occurs at the production site.

Legend	Description
	From the recovery site, the administrator reverses the direction of replication. This operation converts the replication package to a local, writable project. The administrator redirects client workloads and failover IP addresses to the recovery site.
2	After the production site is restored and back to normal operations, the administrator initiates a replication update from the recovery site to the production site. This operation converts the production copy into a replication package, and rolls back any changes written to the recovery site while the production site was down.
3	Once the production site is updated, the administrator reverses the direction of replication again, which makes the copy at the production site writable. The administrator then redirects client workloads and failover IP address back to the production site. The original relationship between the source appliance at the production site and the replication target at the recovery site is restored.

Share-level and Project-level Reversal

When the original source project is converted into a replication package on the original source appliance (which is now acting as the target), the shares that were replicated as part of the action/package currently being reversed are moved into a new replication package and unexported. The original project remains in the local collection, but may end up empty if the action/package included all of its shares. When share-level replication is reversed, any other shares in the original project remain unchanged.

Before reversing the direction of replication for a package, stop replication updates of that project from the source appliance. If a replication update is in progress when an administrator reverses the direction of replication for a project, administrators cannot know which consistent replication snapshot was used to create the resulting project on the former replication target (now source appliance).

If a replication update is performed during or after a reversal operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target. A new replication action and a full update are required to send updates from the original project to a new replication package.

Because all local shares are exported, all shares in a package are exported when the package is reversed, whether or not they were previously exported. If there are mount point conflicts between replicated filesystems and other filesystems on the system, the reverse operation will fail. These conflicts must be resolved before severing by reconfiguring the mount points of the relevant shares. Because this operation is typically part of the critical path of restoring

production service, it is strongly recommended to resolve these mount point conflicts when the systems are first set up rather than at the time of disaster recovery failover.

Related Topics

- [Disaster Recovery with Remote Replication](#)
- [Managing Replication Packages](#)

Destroying a Replication Package

The project and shares within a package cannot be destroyed without destroying the entire package. The entire package can be destroyed from the BUI by destroying the corresponding project. A package can be destroyed from the CLI using the `destroy` command at the `shares replication packages` node.

When a package is destroyed, subsequent replication updates from the corresponding action will fail. To resume replication, the action will need to be recreated on the source to create a new package on the target into which to receive a new copy of the data.

Target Replica Backups

You can back up target replica datasets (projects or shares) using the NDMP zfs backup type. Replica backup is enabled on the appliance by applying the deferred update [Support for NDMP zfs-type Replica Backup](#). The replica backup feature chooses the most recent system-generated snapshot to be backed up, unless you specify a user-generated (non `.rr` extension) snapshot. For more information, see [“Replica Backups” on page 324](#).

Some older replication snapshots, those originally preserved for future incremental backups, might not be needed and can be deleted. If the snapshot is held by NDMP, a confirmation is displayed warning of the potential impact to ongoing or future NDMP backups.

The following sequence of events causes a replication failure and generates an alert. For information about recovering from this error, see "Name collision" in [“Replication Failures” on page 617](#).

1. A replica snapshot is held by NDMP on the replication target (for an ongoing backup or a future incremental backup).
2. The corresponding snapshot on the source appliance is deleted or renamed.
3. A new snapshot is created on the source appliance with the same name as the replica snapshot held on the replication target.
4. A replication update is attempted.

Data Encryption

Note - Encryption is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Oracle ZFS Storage Appliance offers transparent data encryption for individual shares (filesystems and LUNs), and shares created inside of projects.

To configure and manage encryption, use these tasks:

- [“Data Encryption Workflow” on page 630](#)
- Configuring LOCAL Keystore Encryption - [BUI](#), [CLI](#)
- Configuring OKM Keystore Encryption - [BUI](#), [CLI](#)
- Creating an Encrypted Project - [BUI](#), [CLI](#)
- Changing a Project Encryption Key - [BUI](#), [CLI](#)
- Creating an Encrypted Filesystem or LUN - [BUI](#), [CLI](#)
- Changing a Share Encryption Key - [BUI](#), [CLI](#)
- Backing up a LOCAL Key - [BUI](#), [CLI](#)
- Deleting an Encryption Key - [BUI](#), [CLI](#)
- Restoring a LOCAL Key - [BUI](#), [CLI](#)
- Cloning a Snapshot - [BUI](#), [CLI](#)

To understand data encryption, use these topics:

- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Performance Impact of Encryption” on page 654](#)
- [“Encryption Key Life Cycle” on page 655](#)
- [“Backing up and Restoring Encrypted Data” on page 655](#)
- [“Replicating an Encrypted Share” on page 656](#)

▼ Data Encryption Workflow

The following steps show the general procedure for configuring and using data encryption. For information about encryption properties, see [“Encryption Properties” on page 651](#).

1. **Configure LOCAL keystore or Oracle Key Manager (OKM) keystore encryption.**

For information about creating LOCAL or OKM keystores, see the following topics:

- [“Configuring LOCAL Keystore Encryption \(BUI\)” on page 630](#) or [“Configuring LOCAL Keystore Encryption \(CLI\)” on page 633](#)
- [“Configuring OKM Keystore Encryption \(BUI\)” on page 634](#) or [“Configuring OKM Keystore Encryption \(CLI\)” on page 635](#)

2. **Create LOCAL or OKM encryption keys.**

For information about creating LOCAL or OKM keys, see the following topics:

- [“Configuring LOCAL Keystore Encryption \(BUI\)” on page 630](#) or [“Configuring LOCAL Keystore Encryption \(CLI\)” on page 633](#)
- [“Configuring OKM Keystore Encryption \(BUI\)” on page 634](#) or [“Configuring OKM Keystore Encryption \(CLI\)” on page 635](#)

3. **(Optional) Create a project using one of the LOCAL or OKM encryption keys.**

For information about creating a project, see [“Creating a Project \(BUI\)” on page 384](#).

4. **Create a share in a project that uses an encryption key or create a share using one of the LOCAL or OKM encryption keys.**

For information about creating a share, see [“Shares and Projects” on page 383](#) or [“Creating an Encrypted Project \(CLI\)” on page 637](#).

Related Topics

- [“Managing Encryption Keys” on page 652](#)
- [“Performance Impact of Encryption” on page 654](#)
- [“Encryption Key Life Cycle” on page 655](#)
- [“Backing up and Restoring Encrypted Data” on page 655](#)


▼ Configuring LOCAL Keystore Encryption (BUI)

To configure encryption using the LOCAL keystore, first set up the master passphrase and then create keys for assigning to encrypted shares. For information about encryption properties, see [“Encryption Properties” on page 651](#).

To create a key, you provide the name to be used for assigning the key to projects or shares. You can choose to let the system generate the key value or you can supply a hex-encoded raw 256-bit key. Keys are stored in an encrypted form.


1. To configure the LOCAL keystore, go to the **Shares > Encryption BUI** page.
2. Click **Local**.

The LOCAL key store information is displayed.



3. To configure the master passphrase, type the passphrase supplied by your administrator and then retype it in the next box.
4. To save the master passphrase, click **Apply** or to discard your changes click **Revert**.
5. To create a key, click the **Add item icon** .

The New Key dialog box is displayed.



6. **Type a name for the key.**
7. **To use a hex-encoded raw 256-bit key, uncheck "Generate key automatically" and type the key.**
8. **To save the key, click Add, or to discard the changes click Cancel.**

When you click Add, the new key appears in the list of keys with the creation date.

Related Topics

- [“Configuring LOCAL Keystore Encryption \(CLI\)” on page 633](#)
- [“Configuring OKM Keystore Encryption \(BUI\)” on page 634](#)
- [“Creating a Filesystem or LUN in a Project \(BUI\)” on page 390](#)

▼ Configuring LOCAL Keystore Encryption (CLI)

This procedure assumes that encryption was not previously set up on the appliance. For information about encryption properties, see [“Encryption Properties” on page 651](#).

1. To set up the master passphrase, use the following CLI commands:

```
hostname:> shares encryption
hostname:shares encryption> show
Children:
    okm => Manage encryption keys
    local => Manage encryption keys

hostname:shares encryption> local
hostname:shares encryption local> show
Properties:
    master_passphrase =

Children:
    keys => Manage this Keystore's Keys

hostname:shares encryption local> set master_passphrase
Enter new master_passphrase:
Re-enter new master_passphrase:
    master_passphrase = *****
```

2. To create the first key, use the following CLI commands and type a keyname.

This is the name used in the CLI and BUI when assigning a key to a project or share. You can either leave the key property blank and the system will generate a random key value, or you can enter a hex-encoded raw 256-bit key value.

Note - The keys are stored in an encrypted form using the master passphrase supplied. In this example, the system generates the key value.

```
hostname:shares encryption local> keys create
hostname:shares encryption local key (uncommitted)> show
Properties:
    cipher = AES
    key =
    keyname = (unset)
hostname:shares encryption local key (uncommitted)> set keyname=MyFirstKey
    keyname = MyFirstKey (uncommitted)
hostname:shares encryption local key (uncommitted)> commit
```

Related Topics

- [“Configuring LOCAL Keystore Encryption \(BUI\)” on page 630](#)
- [“Configuring OKM Keystore Encryption \(CLI\)” on page 635](#)
- [“Creating an Encrypted Project \(CLI\)” on page 637](#)

▼ Configuring OKM Keystore Encryption (BUI)

To configure encryption using the Oracle Key Manager (OKM), first set up the Key Manager Server information, and then create keys for assigning to encrypted shares. For information about encryption properties, see [“Encryption Properties” on page 651](#).

Note - If the appliance is clustered, do not use the "one time passphrase" setting when creating the OKM server agent otherwise registration on the other cluster node will fail and keys will not be available on failover.

To create a key, you provide the name to be used for assigning the key to projects or shares.

1. **To configure the OKM keystore, go to the Shares > Encryption BUI page.**
2. **Click OKM.**

The OKM keystore information is displayed.

3. **To configure the server information, type the following information:**

- Key Manager Server
 - User Agent ID
 - Registration PIN
4. **To save the server information, click Apply, or to discard the changes, click Cancel.**
 5. **To create a key, click the Add item icon .**
The New Key dialog box is displayed.
 6. **Type a name for the key.**
 7. **To save the key, click Add, or to discard the changes, click Cancel.**
When you click Add, the new key appears in the list of keys with the creation date.

Related Topics

- [“Configuring LOCAL Keystore Encryption \(BUI\)” on page 630](#)
- [“Configuring OKM Keystore Encryption \(CLI\)” on page 635](#)
- [“Creating a Filesystem or LUN in a Project \(BUI\)” on page 390](#)

▼ Configuring OKM Keystore Encryption (CLI)

To use the Oracle Key Manager (OKM) keystore, configure the following parameters:

- agent_id
- registration_pin (supplied by your OKM security officer)
- server_addr

For information about encryption properties, see [“Encryption Properties” on page 651](#).

Note - If the appliance is clustered, do not use the "one time passphrase" setting when creating the OKM server agent otherwise registration on the other cluster node will fail and keys will not be available on failover.

1. **To configure OKM keystore encryption, use the following CLI commands:**

```
hostname:> shares encryption
hostname:shares encryption> show
```

```

Children:
    okm => Manage encryption keys
    local => Manage encryption keys

hostname:shares encryption> okm
hostname:shares encryption okm> show
Properties:
    agent_id = ExternalClient041
    registration_pin = *****
    server_addr = 10.80.180.109

Children:
    keys => Manage this Keystore's Keys

```

2. To create an OKM key, use the following CLI commands:

```

hostname:shares (pool-290-A) encryption okm keys>
hostname:shares (pool-290-A) encryption okm keys> create
hostname:shares (pool-290-A) encryption okm key-372 (uncommitted)> ls
Properties:
    cipher = AES
    keyname = (unset)
hostname:shares (pool-290-A) encryption okm key-372 (uncommitted)> set
keyname=anykey
    keyname = anykey (uncommitted)
hostname:shares (pool-290-A) encryption okm key-372 (uncommitted)> commit

```

▼ Creating an Encrypted Project (BUI)

Shares (filesystems and LUNs) can be encrypted individually or they can be encrypted at the project level because shares inherit project properties. The following example shows how to encrypt all shares within a project by encrypting the project itself.

Before You Begin To use encryption, you must configure it first; see [“Data Encryption” on page 629](#).

1. **Navigate to Shares > Projects.**
2. **Click the Add icon .**
3. **Name the project.**
4. **Set an encryption key length.**
5. **Choose LOCAL or OKM for the keystore.**

6. **Select a keyname.**
7. **Save the project.**

Related Topics

- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Changing a Project Encryption Key \(BUI\)” on page 638](#)

▼ Creating an Encrypted Project (CLI)

Shares (filesystems and LUNs) can be encrypted individually or they can be encrypted at the project level because shares inherit project properties. The following example shows how to encrypt all shares within a project by encrypting the project itself.

Before You Begin To use encryption, you must configure it first; see [“Data Encryption” on page 629](#).

1. **To create an encrypted project, use the following CLI commands:**

```
hostname:shares> project myproject
hostname:shares myproject (uncommitted)> set encryption=aes-128-ccm
      encryption = aes-128-ccm (uncommitted)
hostname:shares myproject (uncommitted)> set keystore=LOCAL
      keystore = LOCAL (uncommitted)
hostname:shares myproject (uncommitted)> set keyname=MyFirstKey
      keyname = MyFirstKey (uncommitted)
hostname:shares myproject (uncommitted)> commit
hostname:shares>
```

All shares created under this project are automatically encrypted with AES-128 CCM using the key named "MyFirstKey" from the LOCAL keystore.

2. **To create a filesystem in the new project and show that it inherited the encryption properties, use the following CLI commands:**

```
hostname:shares> select myproject
hostname:shares myproject> filesystem f1
hostname:shares myproject/f1 (uncommitted)> commit
hostname:shares myproject> select f1
hostname:shares myproject/f1> get encryption keystore keyname keystorestatus
      encryption = aes-128-ccm (inherited)
      keystore = LOCAL (inherited)
      keyname = MyFirstKey (inherited)
```

```
keystatus = available
hostname:shares myproject/fl> done
```

Related Topics


- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Changing a Share Encryption Key \(CLI\)” on page 644](#)

▼ Changing a Project Encryption Key (BUI)

Changing a project encryption key changes the key for shares that inherit the key from the project. You can change the encryption key associated share at any time, even while it is in use by client systems. To change the key for a project, use the following procedure:

1. **To display the project you want to change, go to Shares > Projects.**
2. **To find the share you want, click Show All, Local, or Replica.**
3. **Move your cursor over the project you want to change, and click the Edit icon**
.
4. **Click General.**

The project parameters are displayed.



5. To change the project encryption key, select Local or OKM and select the key you want to use.
6. To save the changes, click Apply or to discard your changes click Revert. When you click Apply, your changes are saved and the new key appears in the Encryption key area.

Related Topics

- [“Changing a Project Encryption Key \(CLI\)” on page 640](#)
- [“Deleting an Encryption Key \(BUI\)” on page 645](#)
- [“Encryption Properties” on page 651](#)

▼ Changing a Project Encryption Key (CLI)

Changing a project encryption key changes the key for shares that inherit the key from the project. You can change the encryption key associated share at any time, even while it is in use by client systems. To change the key for a project, use the following procedure:

1. **To see a project's current key, navigate to that project and enter `get keyname`.**

```
hostname:shares> select default
hostname:shares default> get keyname
                        keyname = MyFirstKey
```

2. **To change the key used for the project and all associated shares, enter `set keyname=` followed by the new keyname, and then enter `commit`.**

```
hostname:shares default> set keyname=MySecondKey
                        keyname = MySecondKey (uncommitted)
hostname: shares default> commit.
```

3. **Verify the new project key by entering `get keyname`.**

```
shares default> get keyname
                        keyname = MySecondKey
```

Related Topics

- [“Changing a Project Encryption Key \(BUI\)” on page 638](#)
- [“Changing a Share Encryption Key \(CLI\)” on page 644](#)
- [“Deleting an Encryption Key \(CLI\)” on page 648](#)
- [“Encryption Properties” on page 651](#)

▼ Creating an Encrypted Filesystem or LUN (BUI)

Shares (filesystems and LUNs) can be encrypted individually or they can be encrypted at the project level because shares inherit project properties. If the project is encrypted, a filesystem or LUN created within it is also encrypted. To create an individual encrypted filesystem or LUN within an unencrypted project, use the following procedure.

Before You Begin To use encryption, you must first configure a keystore and keys; see [“Data Encryption” on page 629](#).

1. **Go to Shares > Shares.**
2. **Select Filesystems or LUNs.**

3. Click the add icon .
4. Complete the fields in the Create Filesystem or Create LUN dialog box.
 - For a filesystem, select a project and enter a name.
 - For a LUN, select a project, enter a name and specify the volume size.
 - For Encryption, select On.

A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
Allowable characters are: alphanumeric and special characters `_ - . :`

5. Click **APPLY**.

Related Topics

- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Changing a Share Encryption Key \(BUI\)” on page 642](#)

▼ Creating an Encrypted Filesystem or LUN (CLI)

Shares (filesystems and LUNs) can be encrypted individually or they can be encrypted at the project level because shares inherit project properties. If the project is encrypted, a filesystem or LUN created within it is also encrypted. To create an individual encrypted filesystem or LUN that is in an unencrypted project, use the following procedure.

Before You Begin To use encryption, you must configure it first; see [“Data Encryption” on page 629](#).

1. **Go to shares.**

```
hostname:> shares
```
2. **Enter select and the project name.**
 In this example, the default project is selected.

```
hostname:shares > select default
```
3. **Enter filesystem and a filesystem name, or lun and a LUN name.**
 A name must consist of 1 to 64 characters, but not include spaces or begin with a period.
 Allowable characters are: alphanumeric and special characters `_ - . :`
 The following example creates a filesystem named `fs-1` in the default project.

```
hostname:shares default> filesystem fs-1
hostname:shares default/fs-1 (uncommitted)>
```

4. **If creating a LUN, enter `set volsize=` and the volume size.**

```
hostname:shares default/lun1 (uncommitted)> set volsize=2G
volsize = 2G (uncommitted)
```

5. **To enable encryption, enter `set encryption=` and the encryption type.**

```
hostname:shares default/fs-1 (uncommitted)> set encryption=aes-128-ccm
encryption = aes-128-ccm (uncommitted)
```

6. **Configure encryption using either the LOCAL keystore or the Oracle Key Manager (OKM) keystore. Enter `set keystore=` and either LOCAL or OKM.**

```
hostname:shares default/fs-1 (uncommitted)> set keystore=LOCAL
keystore = LOCAL (uncommitted)
```

7. **To set the encryption key, enter `set keyname=` and the key name.**

```
hostname:shares default/fs-1 (uncommitted)> set keyname=MyFirstKey
keyname = MyFirstKey (uncommitted)
```

8. **Enter `commit`.**

```
hostname:shares default/fs-1 (uncommitted)> commit
```

Related Topics


- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Changing a Share Encryption Key \(CLI\)” on page 644](#)

▼ Changing a Share Encryption Key (BUI)

You can change the encryption key associated share at any time, even while it is in use by client systems. To change a key for a share without changing the parent project, use the following procedure:

1. **To display the properties for the share you want to change, go to **Shares > Shares**.**

2. Select Filesystems or LUNs.
3. To find the share you want, click Show All, Local, or Replica.
4. Move your cursor over the share you want to change, and click the Edit icon . The share properties are displayed.



5. If necessary, uncheck Inherit from project.
6. To change the encryption key, select Local or OKM and select the key you want to use.
7. To save the changes, click Apply or to discard your changes click Revert. When you click Apply, your changes are saved and the new key appears in the Encryption key area.

Related Topics

- [“Changing a Project Encryption Key \(BUI\)” on page 638](#)
- [“Deleting an Encryption Key \(BUI\)” on page 645](#)
- [“Encryption Properties” on page 651](#)

▼ Changing a Share Encryption Key (CLI)

You can change the encryption key associated share at any time, even while it is in use by client systems. To change a key for a share without changing the parent project, use the following procedure:

1. **To see a share's current key, navigate to that filesystem or LUN and enter `get keyname`.**

```
hostname:shares default> select fs-1
hostname:shares default/fs-1> get keyname
keyname = MyFirstKey
```

2. **To change the key used for the share, enter `set keyname=` followed by the new keyname, and then enter `commit`.**

```
hostname:shares default/fs-1> set keyname=MySecondKey
keyname = MySecondKey (uncommitted)
hostname: shares default/fs-1> commit.
```

3. **Verify the new share key by entering `get keyname`.**

```
shares default/fs-1> get keyname
keyname = MySecondKey
```

Related Topics


- [“Changing a Project Encryption Key \(CLI\)” on page 640](#)
- [“Deleting an Encryption Key \(CLI\)” on page 648](#)
- [“Encryption Properties” on page 651](#)

▼ Backing Up a LOCAL Key (BUI)

Use the following procedure to retrieve the information for a single LOCAL key in order to back it up.

1. **Navigate to Shares > Encryption > Local.**

2. **Click on the key you want to back up.**
A dialog box appears with the keyname and key value.



3. **Using any method, record this information in a backup location of your choosing and then click OK.**

▼ Backing Up a LOCAL Key (CLI)

Use the following procedure to retrieve the information for a single LOCAL key in order to back it up.

1. **Select the key:**

```
hostname:shares encryption local keys> select keyname=Mykey
```
2. **Get the key value:**

```
hostname:shares encryption local key-005> get key
key = d6a5b801ffb93fcb19ef70a11d662d8092f243c5d4ccd0cd34264b15dd0b7739
```
3. **Using any method, record this information in a backup location of your choosing.**

▼ Deleting an Encryption Key (BUI)

Deleting an encryption key is a fast and effective way to make large amounts of data inaccessible. Keys can be deleted even if they are in use. If the key is in use, a warning is given


and confirmation is required. All shares or projects using that key are unshared and can no longer be accessed by clients.

If you might use a LOCAL key again to access its associated shares, back up the keyname and value before deleting the key. Then you can later perform a restore procedure as described in [“Restoring a LOCAL Key \(BUI\)” on page 649](#).

Use the following procedure to delete a LOCAL or OKM encryption key.

1. **Navigate to Shares > Encryption.**
2. **Select Local or OKM.**
3. **Move your cursor over the key that you want to delete and click the Delete icon .**

The following alert is displayed:



4. To delete the key, click OK, or to keep the key, click Cancel.

When a key is deleted, all of the data in all of the shares that use the key becomes inaccessible. This is equivalent to secure data destruction and is permanent and irrevocable, unless you have prepared for key restoration by backing up the key. For more information about key backup and restoration, see [“Backing Up a LOCAL Key \(BUI\)” on page 644](#) and [“Restoring a LOCAL Key \(BUI\)” on page 649](#).

Related Topics

- [“Changing a Share Encryption Key \(BUI\)” on page 642](#)
- [“Deleting an Encryption Key \(CLI\)” on page 648](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Encryption Key Life Cycle” on page 655](#)

▼ Deleting an Encryption Key (CLI)

Deleting an encryption key is a fast and effective way to make large amounts of data inaccessible. Keys can be deleted even if they are in use. If the key is in use, a warning is given and confirmation is required. All shares or projects using that key are unshared and can no longer be accessed by clients.

If you might use a LOCAL key again to access its associated shares, back up the keyname and value before deleting the key. Then you can later perform a restore procedure as described in [“Restoring a LOCAL Key \(CLI\)” on page 650](#).

Use the following procedure to delete a LOCAL or OKM encryption key.

1. To delete a key, use the following CLI command:

```
hostname:shares encryption local local_keys> destroy keyname=AKTEST_K1
```

This key has the following dependent shares:

```
Pool2/local/BG1
Pool2/local/BG1/BG3
Pool2/local/BG1/fast1
Pool2/local/default/BG2
```

Destroying this key will render the data inaccessible. Are you sure? (Y/N)

2. To verify that a share is no longer accessible using that key, navigate to the share and use the following CLI commands:

```
hostname:> shares select test_project select test_share1
hostname:shares test_project/test_share1> get encryption keystore keyname keystore
```

```
encryption = aes-128-ccm (inherited)
keystore = LOCAL (inherited)
keyname = AKTEST_K1 (inherited)
keystore = unavailable
```

Errors:

```
key_unavailable
```

3. To list dependents, use the following CLI commands:

```
hostname:shares (pool-010) encryption local keys> select keyname=1 hostname:shares
(pool-010) encryption local key-002> list
```

Properties:

```
cipher = AES
keyname = 1
```

```
hostname:shares (pool-010) encryption local key-002> list dependents DEPENDENTS
pool-010/local/default/a hostname:shares (pool-010) encryption local key-002>
```

Related Topics

- [“Changing a Share Encryption Key \(CLI\)” on page 644](#)
- [“Backing Up a LOCAL Key \(CLI\)” on page 645](#)
- [“Restoring a LOCAL Key \(CLI\)” on page 650](#)

▼ Restoring a LOCAL Key (BUI)

To restore a LOCAL key that was deleted, create a new LOCAL key with the same keyname and value as the deleted key. You must have first recorded, or backed up, this information before the key was deleted. The backup procedure is described in [“Backing Up a LOCAL Key \(BUI\)” on page 644](#). Although deleting a LOCAL key renders shares inaccessible, the shares can be made accessible again by recreating the LOCAL key.

For information about restoring keys stored in the OKM keystore, refer to the Oracle Key Manager documentation on the [Oracle Technology Network \(http://www.oracle.com/technetwork/indexes/documentation/index.html\)](http://www.oracle.com/technetwork/indexes/documentation/index.html).

Use the following procedure to restore a backed up LOCAL key.

Note - If the keyname is in use with a different key value for existing shares, change the key used for those shares before restoring the original LOCAL key. For more information, see [“Changing a Share Encryption Key \(BUI\)” on page 642](#).

1. **Retrieve the keyname and value for the LOCAL key from your backup location.**
2. **Navigate to Shares > Encryption > Local and click the Add icon.**
3. **Enter the same keyname as in the backup.**

4. **Uncheck "Generate key automatically" and set the key value based on the backup.**
5. **Save the restored key by clicking ADD.**

If the keyname is used with existing shares, a dialog box appears. To overwrite the key value in the existing shares, click OK. Click Cancel to not add the new key. You can then change the key used for those shares before repeating this procedure and restoring the original key. For more information, see ["Changing a Share Encryption Key \(BUI\)" on page 642](#).

Related Topics

- ["Changing a Share Encryption Key \(BUI\)" on page 642](#)
- ["Backing Up a LOCAL Key \(BUI\)" on page 644](#)
- ["Deleting an Encryption Key \(BUI\)" on page 645](#)

▼ Restoring a LOCAL Key (CLI)

To restore a LOCAL key that was deleted, create a new LOCAL key with the same keyname and value as the deleted key. You must have first recorded, or backed up, this information before the key was deleted. The backup procedure is described in ["Backing Up a LOCAL Key \(CLI\)" on page 645](#). Although deleting a LOCAL key renders shares inaccessible, the shares can be made accessible again by recreating the LOCAL key.

For information about restoring keys stored in the OKM keystore, refer to the Oracle Key Manager documentation on the [Oracle Technology Network \(http://www.oracle.com/technetwork/indexes/documentation/index.html\)](http://www.oracle.com/technetwork/indexes/documentation/index.html).

Use the following procedure to restore a backed up LOCAL key.

Note - If the keyname is in use with a different key value for existing shares, change the key used for those shares before restoring the original LOCAL key. For more information, see ["Changing a Share Encryption Key \(CLI\)" on page 644](#).

1. **Retrieve the keyname and value for the LOCAL key from your backup location.**
2. **Create a key in the LOCAL keystore:**
3. **Name the key based on the backup:**

```
hostname:shares encryption local keys> create
```

```
hostname:shares encryption local key-005 (uncommitted)> set keyname=Mykey
keyname = Mykey (uncommitted)
```

4. Set the key value based on the backup:

```
hostname:shares encryption local key-005 (uncommitted)> set
key=d6a5b801ffb93fcb19ef70a11d662d8092f243c5d4ccd0cd34264b15dd0b7739
  key = d6a5b801ffb93fcb19ef70a11d662d8092f243c5d4ccd0cd34264b15dd0b7739
(uncommitted)
```

5. Save the key:

```
hostname:shares encryption local key-005 (uncommitted)> commit
```

If the keyname is used with existing shares, you will be alerted:

```
Existing shares reference the key Mykey from the LOCAL keystore. Are you sure? (Y/N)
```

To overwrite the key value in the existing shares, type Y. Type N to not add the new key. You can then change the key used for those shares before repeating this procedure and restoring the original key. For more information, see [“Changing a Share Encryption Key \(CLI\)” on page 644](#).

Related Topics

- [“Changing a Share Encryption Key \(CLI\)” on page 644](#)
- [“Backing Up a LOCAL Key \(CLI\)” on page 645](#)
- [“Deleting an Encryption Key \(CLI\)” on page 648](#)

Encryption Properties

The following list shows the encryption properties available for managing keys, creating keys, and creating encrypted projects and shares.

- **LOCAL Key Management Properties**
 - **Master Passphrase** - The master passphrase is used to generate an AES key for encrypting the keys stored in the LOCAL keystore. The PKCS#5 PBKDF algorithm is used to generate the key and the key is randomly generated and managed by the system.
- **LOCAL Key Creation Properties**
 - **Keyname** - Name to identify the key.
 - **Generate Key Automatically** - Automatically generate the key.
 - **Key** - Hex-encoded raw 256-bit key, stored in an encrypted form, if automatic key generation is not selected.
- **OKM Key Management Properties** (supplied by your OKM administrator)
 - **Key Manager Server** - IP address of your Oracle Key Manager (OKM) server.

- **User Agent ID** - Agent ID.
- **Registration PIN** - Registration PIN.
- **OKM Key Creation Properties**
 - **Keyname** - Name to identify the key.
- **Shares Encryption Properties**
 - **Encryption** - AES encryption type and key length (for more information, see [“Understanding Encryption Key Values” on page 653](#)).
 - **Inherit key** - Inherit the encryption key from the parent project.
 - **Key** - Sets a specific LOCAL or OKM key and is used when the key is not inherited from the parent project.
- **Project Encryption Properties**
 - **Name** - Name to identify the project.
 - **Encryption** - AES encryption type and key length (for more information, see [“Understanding Encryption Key Values” on page 653](#)).
 - **Key** - Specific LOCAL or OKM key.

Related Topics

- [“Data Encryption Workflow” on page 630](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Performance Impact of Encryption” on page 654](#)
- [“Encryption Key Life Cycle” on page 655](#)

Managing Encryption Keys

The appliance includes a built-in LOCAL keystore and the ability to connect to the Oracle Key Manager (OKM) system. Each encrypted project or share requires a wrapping key from either the LOCAL or OKM keystores. The data encryption keys are managed by the storage appliance and are stored persistently encrypted by the wrapping key from the LOCAL or OKM keystore.

OKM is a comprehensive key management system (KMS) that addresses the rapidly growing enterprise need for storage-based data encryption. Developed to comply with open standards, this feature provides the capacity, scalability, and interoperability to manage encryption keys centrally over widely distributed and heterogeneous storage infrastructures.

OKM meets the unique challenges of storage key management, including:

- **Long-term key retention** - OKM ensures that archive data is always available, and it securely retains encryption keys for the full data life cycle.

- **Interoperability** - OKM provides the interoperability needed to support a diverse range of storage devices attached to mainframe or open systems under a single storage key management service.
- **High availability** - With active N-node clustering, dynamic load balancing, and automated failover, OKM provides high availability, whether the appliances are sited together or distributed around the world.
- **High capacity** - OKM manages large numbers of storage devices and even more storage keys. A single clustered appliance can provide key management services for thousands of storage devices and millions of storage keys.
- **Flexible Key Configuration** - Per OKM cluster, keys can be generated automatically or created individually for a LOCAL or OKM keystore. Security administrators are responsible for providing the key names which, when combined with the keystore, associate a given wrapping key with a project or share.

Note - If the appliance is clustered, do not use the "one time passphrase" setting when creating the OKM server agent otherwise registration on the other cluster node will fail and keys will not be available on failover.

Maintaining Keys

Shares and projects that use OKM keys that are in a deactivated state remain accessible. To prevent an OKM key from being used, the OKM administrator must explicitly delete the key.

To ensure encrypted shares and projects are accessible, back up your appliance configurations and LOCAL keystore key values. If a key(s) becomes unavailable, any shares or projects that use that key become inaccessible. If a project key is unavailable, new shares cannot be created in that project.

Keys can become unavailable in the following ways:

- Keys are deleted
- Rollback to a release that does not support encryption
- Rollback to a release where the keys are not configured
- Factory reset
- OKM server is not available

Understanding Encryption Key Values

The following table shows the BUI and CLI encryption key values and descriptions. It also indicates if the encryption type works with deduplication.

TABLE 139 Encryption Key Values

BUI Value	CLI Value	Description
Off	off	Share/Project is not encrypted
AES-128-CCM	aes-128-ccm	Lowest CPU impact encryption. Dedupable
AES-192-CCM	aes-192-ccm	Dedupable
AES-256-CCM	aes-256-ccm	Dedupable
AES-128-GCM	aes-128-gcm	NIST SP800-38D recommended, Not-Dedupable
AES-192-GCM	aes-192-gcm	NIST SP800-38D recommended, Not-Dedupable
AES-256-GCM	aes-256-gcm	Highest CPU impact encryption, NIST SP800-38D recommended, Not-Dedupable

Performance Impact of Encryption

Using encryption with shares can have CPU performance impacts, as follows:

- The AES-128-CCM mode has the lowest CPU performance impact and is recommended for all workloads where there are no LOCAL security requirements.
- When encrypted data is read, it is stored decrypted and decompressed in DRAM. For read-dominate workloads that can be serviced read-dominant from the DRAM cache, the impact of decrypting the data is minimal.
- When SSD cache devices are used, data blocks evicted out of DRAM to the cache are compressed and encrypted and must be decrypted and decompressed when retrieved back into DRAM.
- For workloads that are write-dominant and use larger block sizes, especially 128 kilobytes and 1 megabyte, there can be a significant CPU impact resulting in lower throughput. This is particularly likely if the filesystem record size or LUN volume block size is larger than the application block size.

Related Topics

- [“Data Encryption Workflow” on page 630](#)
- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)

- [“Encryption Key Life Cycle” on page 655](#)

Encryption Key Life Cycle

The encryption key life cycle is flexible because you can change keys at any time without taking data services offline.

When a key is deleted from the keystore, all the shares that use it are unmounted and their data becomes inaccessible. Backing up keys in the OKM keystore should be performed using the OKM backup services. Backup of keys in the LOCAL keystore is included as part of the System Configuration Backup. For the LOCAL keystore, it is also possible to supply the key by value at creation time to allow it to be escrowed in an external system, which provides an alternative per-key backup/restore capability.

Related Topics

- [“Data Encryption Workflow” on page 630](#)
- [“Encryption Properties” on page 651](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Performance Impact of Encryption” on page 654](#)

Backing up and Restoring Encrypted Data

When a share is restored using the ZFS restore function, the restored share inherits the encryption properties of the target project if the original share inherited its encryption properties from the source project.

To ensure encryption properties of an original share are maintained in a restored share, configure encryption on the original share instead of inheriting it from its project.

If you want to set encryption differently for an individual share within a project, manually configure encryption for the individual source share, instead of letting the share inherit its properties from the project. This ensures that all shares are backed up and restored with the desired encryption settings.

For more information about NDMP backup, see [“NDMP Configuration” on page 319](#). For information about replication, see [“Remote Replication” on page 511](#).

Related Topics

- [“Data Encryption Workflow” on page 630](#)

- [“Managing Encryption Keys” on page 652](#)
- [“Encryption Key Life Cycle” on page 655](#)
- [“Replicating an Encrypted Share” on page 656](#)

Replicating an Encrypted Share

To replicate an encrypted share, both the source and target must support encryption and meet these requirements:

- Software release 2013.1.3.0 (or later)
- Encryption wrapping keys used by the share
- OKM key name must be identical in the keystore on both replication source and replication targets.
- OKM Agent ID must be unique on the replication source and target replication appliances. Replication peer appliances cannot use the same agent.
- OKM agents for the replication peers should be configured on the OKM server to see the same key groups. In this case, it is not necessary to manually create the OKM key on the replication target system.

The replication will fail if you attempt to replicate an encrypted share and the target does not support encryption. If the wrapping key is not available on the source or target system, or the target software is earlier than 2013.1.3.0, an alert is raised. Review the alerts on both the source and target to determine the reason for the replication failure.

For more information on configuring replication, see [“Remote Replication” on page 511](#).

Related Topics

- [“Data Encryption Workflow” on page 630](#)
- [“Managing Encryption Keys” on page 652](#)
- [“Encryption Key Life Cycle” on page 655](#)
- [“Backing up and Restoring Encrypted Data” on page 655](#)
- Oracle ZFS Storage Appliance: Remote Replication Compatibility [Doc ID 1958039.1]
<https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1>.

Maintenance Workflows

A workflow is a CLI script that is uploaded to and managed by the appliance by itself. Workflows can be parameterized and executed in a first-class fashion from either the browser interface or the command line interface. Workflows may also be optionally executed as alert or at a designated time. As such, workflows allow for the appliance to be *extended* in ways that capture specific policies and procedures, and can be used (for example) to formally encode best practices for a particular organization or application.

To work with workflows, use the following sections:

- [Understanding Workflows](#)
- [Understanding Workflow Parameters](#)
- [Constrained Workflow Parameters](#)
- [Optional Workflow Parameters](#)
- [Workflow Error Handling](#)
- [Workflow Input Validation](#)
- [Workflow Execution Auditing and Reporting](#)
- [Understanding Workflow Versioning](#)
- [Using Workflows for Alert Actions](#)
- [Using Scheduled Workflows](#)
- [Using a Scheduled Workflow](#)
- [Coding Workflow Schedules](#)
- [Creating a Worksheet Based on a Specified Drive Type](#)
- [Uploading and Executing Workflows using the BUI](#)
- [Downloading Workflows using the CLI](#)
- [Listing Workflows using the CLI](#)
- [Executing Workflows using the CLI](#)

Understanding Workflows

A workflow is embodied in a valid ECMAScript file, containing a single global variable, `workflow`. This is an Object that must contain at least three members:

TABLE 140 Required Object Members

Required member	Type	Description
name	String	Name of the workflow
description	String	Description of workflow
execute	Function	Function that executes the workflow

EXAMPLE 18 Hello World Workflow

The following is an example of a basic workflow:

```
var workflow = {
  name: 'Hello world',
  description: 'Bids a greeting to the world',
  execute: function () { return ('hello world!') }
};
```

Uploading this workflow will result in a new workflow named "Hello world"; executing the workflow will result in the output "hello world!"

EXAMPLE 19 Using the Workflow Run Function to Return CPU Utilization

Workflows execute asynchronously in the appliance shell, running (by default) as the user executing the workflow. As such, workflows have at their disposal the appliance scripting facility (see [“Working with CLI Scripting” on page 45](#)), and may interact with the appliance just as any other instance of the appliance shell. That is, workflows may execute commands, parse output, modify state, and so on. Here is a more complicated example that uses the `run` function to return the current CPU utilization:

```
var workflow = {
  name: 'CPU utilization',
  description: 'Displays the current CPU utilization',
  execute: function () {
    run('analytics datasets select name=cpu.utilization');
    cpu = run('csv 1').split('\n')[1].split(',');
    return ('At ' + cpu[0] + ', utilization is ' + cpu[1] + '%');
  }
};
```

Understanding Workflow Parameters

Workflows that do not operate on input have limited scope; many workflows need to be parameterized to be useful. This is done by adding a `parameters` member to the global workflow object. The `parameters` member is in turn an object that is expected to have a member for each parameter. Each `parameters` member must have the following members:

TABLE 141 Required Workflow Parameter Members

Required Member	Type	Description
<code>label</code>	String	Label to adorn input of workflow parameter
<code>type</code>	String	Type of workflow parameter

The `type` member must be set to one of these types:

TABLE 142 Workflow Member Type Names

Type name	Description
<code>Boolean</code>	A boolean value
<code>ChooseOne</code>	One of a number of specified values
<code>EmailAddress</code>	An e-mail address
<code>File</code>	A file to be transferred to the appliance
<code>Host</code>	A valid host, as either a name or dotted decimal
<code>HostName</code>	A valid hostname
<code>HostPort</code>	A valid, available port
<code>Integer</code>	An integer
<code>NetAddress</code>	A network address
<code>NodeName</code>	A name of a network node
<code>NonNegativeInteger</code>	An integer that is greater than or equal to zero
<code>Number</code>	Any number -- including floating point
<code>Password</code>	A password
<code>Permissions</code>	POSIX permissions
<code>Port</code>	A port number
<code>Size</code>	A size
<code>String</code>	A string
<code>StringList</code>	A list of strings

EXAMPLE 20 Workflow Using Two Parameters

Based on the specified types, an appropriate input form will be generated upon execution of the workflow. For example, here is a workflow that has two parameters, the name of a business unit (to be used as a project) and the name of a share (to be used as the share name):

```
var workflow = {
  name: 'New share',
  description: 'Creates a new share in a business unit',
  parameters: {
    name: {
      label: 'Name of new share',
      type: 'String'
    },
    unit: {
      label: 'Business unit',
      type: 'String'
    }
  },
  execute: function (params) {
    run('shares select ' + params.unit);
    run('filesystem ' + params.name);
    run('commit');
    return ('Created new share "' + params.name + '"');
  }
};
```

If you upload this workflow and execute it, you will be prompted with a dialog box to fill in the name of the share and the business unit. When the share has been created, a message will be generated indicating as much.

Constrained Workflow Parameters

For some parameters, one does not wish to allow an arbitrary string, but wishes to rather limit input to one of a small number of alternatives. These parameters should be specified to be of type `ChooseOne`, and the object containing the parameter must have two additional members:

TABLE 143 Constrained Parameters Required Members

Required Member	Type	Description
options	Array	An array of strings that specifies the valid options
optionlabels	Array	An array of strings that specifies the labels associated with the options specified in <code>options</code>

EXAMPLE 21 Using the Workflow ChooseOne Parameter

Using the ChooseOne parameter type, we can enhance the previous example to limit the business unit to be one of a small number of predefined values:

```
var workflow = {
  name: 'Create share',
  description: 'Creates a new share in a business unit',
  parameters: {
    name: {
      label: 'Name of new share',
      type: 'String'
    },
    unit: {
      label: 'Business unit',
      type: 'ChooseOne',
      options: [ 'development', 'finance', 'qa', 'sales' ],
      optionlabels: [ 'Development', 'Finance',
        'Quality Assurance', 'Sales/Administrative' ],
    }
  },
  execute: function (params) {
    run('shares select ' + params.unit);
    run('filesystem ' + params.name);
    run('commit');
    return ('Created new share "' + params.name + '"');
  }
};
```

When this workflow is executed, the unit parameter will not be entered by hand -- it will be selected from the specified list of possible options.

Optional Workflow Parameters

Some parameters may be considered *optional* in that the UI should not mandate that these parameters are set to any value to allow execution of the workflow. Such a parameter is denoted via the optional field of the parameters member:

TABLE 144 Required Members for Optional Parameters

Optional Member	Type	Description
optional	Boolean	If set to true, denotes that the parameter need not be set; the UI may allow the workflow to be

Optional Member	Type	Description
		executed without a value being specified for the parameter.

If a parameter is optional and is unset, its member in the parameters object passed to the execute function will be set to undefined.

Workflow Error Handling

If, in the course of executing a workflow, an error is encountered, an exception will be thrown. If the exception is not caught by the workflow itself (or if the workflow throws an exception that is not otherwise caught), the workflow will fail, and the information regarding the exception will be displayed to the user. To properly handle errors, exceptions should be caught and processed. For example, in the previous example, an attempt to create a share in a non-existent project results in an uncaught exception.

EXAMPLE 22 Workflow Error Handling

This example could be modified to catch the offending error, and create the project in the case that it doesn't exist:

```
var workflow = {
  name: 'Create share',
  description: 'Creates a new share in a business unit',
  parameters: {
    name: {
      label: 'Name of new share',
      type: 'String'
    },
    unit: {
      label: 'Business unit',
      type: 'ChooseOne',
      options: [ 'development', 'finance', 'qa', 'sales' ],
      optionlabels: [ 'Development', 'Finance',
        'Quality Assurance', 'Sales/Administrative' ],
    }
  },
  execute: function (params) {
    try {
      run('shares select ' + params.unit);
    } catch (err) {
      if (err.code != EAKSH_ENTITY_BADSELECT)
        throw (err);
    }
  }
}
```

```

/*
 * We haven't yet created a project that corresponds to
 * this business unit; create it now.
 */
run('shares project ' + params.unit);
run('commit');
run('shares select ' + params.unit);
}

run('filesystem ' + params.name);
run('commit');
return ('Created new share "' + params.name + '"');
}
};

```

Workflow Input Validation

Workflows may optionally validate their input by adding a `validate` member that takes as a parameter an object that contains the workflow parameters as members. The `validate` function should return an object where each member is named with the parameter that failed validation, and each member's value is the validation failure message to be displayed to the user.

EXAMPLE 23 Workflow Input Validation

To extend our example to give a crisp error if the user attempts to create an extant share:

```

var workflow = {
  name: 'Create share',
  description: 'Creates a new share in a business unit',
  parameters: {
    name: {
      label: 'Name of new share',
      type: 'String'
    },
  },
  unit: {
    label: 'Business unit',
    type: 'ChooseOne',
    options: [ 'development', 'finance', 'qa', 'sales' ],
    optionlabels: [ 'Development', 'Finance',
      'Quality Assurance', 'Sales/Administrative' ],
  }
},
validate: function (params) {

```

```
    try {
      run('shares select ' + params.unit);
      run('select ' + params.name);
    } catch (err) {
      if (err.code == EAKSH_ENTITY_BADSELECT)
        return;
    }

    return ({ name: 'share already exists' });
  },
  execute: function (params) {
    try {
      run('shares select ' + params.unit);
    } catch (err) {
      if (err.code != EAKSH_ENTITY_BADSELECT)
        throw (err);

      /*
       * We haven't yet created a project that corresponds to
       * this business unit; create it now.
       */
      run('shares project ' + params.unit);
      set('mountpoint', '/export/' + params.unit);
      run('commit');
      run('shares select ' + params.unit);
    }

    run('filesystem ' + params.name);
    run('commit');
    return ('Created new share "' + params.name + '"');
  }
};
```

Workflow Execution Auditing and Reporting

Workflows may emit audit records by calling the `audit` function. The `audit` function's only argument is a string that is to be placed into the audit log.

For complicated workflows that may require some time to execute, it can be useful to provide clear progress to the user executing the workflow. To allow the execution of a workflow to be reported in this way, the `execute` member should return an array of *steps*. Each array element must contain the following members:

TABLE 145 Required Members for Execution Reporting

Required Member	Type	Description
step	String	String that denotes the name of the execution step
execute	Function	Function that executes the step of the workflow

As with the execute function on the workflow as a whole, the execute member of each step takes as its argument an object that contains the parameters to the workflow.

EXAMPLE 24 Workflow Execution Reporting

As an example, the following is a workflow that creates a new project, share, and audit record over three steps:

```
var steps = [ {
  step: 'Checking for associated project',
  execute: function (params) {
    try {
      run('shares select ' + params.unit);
    } catch (err) {
      if (err.code !== EAKSH_ENTITY_BADSELECT)
        throw (err);

      /*
       * We haven't yet created a project that corresponds to
       * this business unit; create it now.
       */
      run('shares project ' + params.unit);
      set('mountpoint', '/export/' + params.unit);
      run('commit');
      run('shares select ' + params.unit);
    }
  }
}, {
  step: 'Creating share',
  execute: function (params) {
    run('filesystem ' + params.name);
    run('commit');
  }
}, {
  step: 'Creating audit record',
  execute: function (params) {
    audit('created "' + params.name + '" in "' + params.unit);
  }
} ];
```

```
var workflow = {
  name: 'Create share',
  description: 'Creates a new share in a business unit',
  parameters: {
    name: {
      label: 'Name of new share',
      type: 'String'
    },
    unit: {
      label: 'Business unit',
      type: 'ChooseOne',
      options: [ 'development', 'finance', 'qa', 'sales' ],
      optionlabels: [ 'Development', 'Finance',
        'Quality Assurance', 'Sales/Administrative' ],
    }
  },
  validate: function (params) {
    try {
      run('shares select ' + params.unit);
      run('select ' + params.name);
    } catch (err) {
      if (err.code == EAKSH_ENTITY_BADSELECT)
        return;
    }

    return ({ name: 'share already exists' });
  },
  execute: function (params) { return (steps); }
};
```

Using the mail function, workflows can deliver certain outputs of the workflow via email. The mail function must contain the following arguments: an object with to and subject, and a messageBody string.

EXAMPLE 25 Workflow Execution with a Mailer

```
var workflow = {
  name: 'email controller state',
  description: 'email controller state',
  execute: function () {

    // verify state of the controller
    var faulted = run('maintenance hardware "chassis-000" get faulted');

    var messageBody = faulted;
```

```

        emailAddress = 'first.last@xyz.com';
        subjectLine = 'Controller State';
        mail({To: emailAddress, Subject: subjectLine}, messageBody);
    }
};

```

Understanding Workflow Versioning

There are two aspects of versioning with respect to workflows: the first is the expression of the version of the appliance software that the workflow depends on, and the second is the expression of the version of the workflow itself. Versioning is expressed through two optional members to the workflow:

TABLE 146 Optional Members for Versioning

Optional Member	Type	Description
required	String	The minimum version of the appliance software required to run this workflow, including the minimum year, month, day, build and branch.
version	String	Version of this workflow, in dotted decimal (major.minor.micro) form.

Appliance Versioning - To express a minimally required version of the appliance software, add the optional required field to your workflow. The appliance is versioned in terms of the year, month and day on which the software was built, followed by a build number and then a branch number, expressed as "year.month.day.build-branch". For example "2009.04.10,12-0" would be the twelfth build of the software originally build on April 10th, 2009. To get the version of the current appliance kit software, run the "configuration version get version" CLI command, or look at the "Version" field in the "System" screen in the BUI. Here's an example of using the required field:

EXAMPLE 26 Using the Workflow Required Field

Here's an example of using the required field:

```

var workflow = {
    name: 'Configure FC',
    description: 'Configures fibre channel target groups',
    required: '2009.12.25,1-0',
    ...
}

```

If a workflow requires a version of software that is newer than the version loaded on the appliance, the attempt to upload the workflow will fail with a message explaining the mismatch.

Workflow Versioning - In addition to specifying the required version of the appliance software, workflows themselves may be versioned with the `version` field. This string denotes the major, minor and micro numbers of the workflow version, and allows multiple versions of the same workflow to exist on the machine. When uploading a workflow, any *compatible*, *older* versions of the same workflow are deleted. A workflow is deemed to be *compatible* if it has the same major number, and a workflow is considered to be *older* if it has a lower version number. Therefore, uploading a workflow with a version of "2.1" will remove the same workflow with version "2.0" (or version "2.0.1") but not "1.2" or "0.1".

Using Workflows for Alert Actions

Workflows may be optionally executed as an alert. To allow a workflow to be eligible as an alert action, its `alert` action must be set to `true`.

When executed as alert actions, workflows assume the identity of the user that created them. For this reason, any workflow that is to be eligible as an alert action must set `setid` to `true`. Alert actions have a single object parameter that has the following members:

TABLE 147 Required Members for Alert Execution Context

Required Member	Type	Description
<code>class</code>	String	The class of the alert.
<code>code</code>	String	The code of the alert.
<code>items</code>	Object	An object describing the alert.
<code>timestamp</code>	Date	Time of alert.

The `items` member of the parameters object has the following members:

TABLE 148 Required Members for the Items Member

Required Member	Type	Description
<code>url</code>	String	The URL of the web page describing the alert
<code>action</code>	String	The action that should be taken by the user in response to the alert.
<code>impact</code>	String	The impact of the event that precipitated the alert.

Required Member	Type	Description
description	String	A human-readable string describing the alert.
severity	String	The severity of the event that precipitated the alert.

Workflows executing as alert actions may use the `audit` function to generate audit log entries. It is recommended that any relevant debugging information be generated to the audit log via the `audit` function. For example, here is a workflow that executes failover if in the clustered state -- but audits any failure to reboot:

EXAMPLE 27 Workflow Auditing Failure to Reboot

For example, here is a workflow that executes failover if in the clustered state -- but audits any failure to reboot:

```
var workflow = {
  name: 'Failover',
  description: 'Fail the node over to its clustered peer',
  alert: true,
  setid: true,
  execute: function (params) {
    /*
     * To failover, we first confirm that clustering is configured
     * and that we are in the clustered state. We then reboot,
     * which will force our peer to takeover. Note that we're
     * being very conservative by only rebooting if in the
     * AKCS_CLUSTERED state: there are other states in which it
     * may well be valid to failback (e.g., we are in AKCS_OWNER,
     * and our peer is AKCS_STRIPPED), but those states may also
     * indicate aberrant operation, and we therefore refuse to
     * failback. (Even in an active/passive clustered config, a
     * FAILBACK should always be performed to transition the
     * cluster peers from OWNER/STRIPPED to CLUSTERED/CLUSTERED.)
     */
    var uuid = params.uuid;
    var clustered = 'AKCS_CLUSTERED';

    audit('attempting failover in response to alert ' + uuid);

    try {
      run('configuration cluster');
    } catch (err) {
      audit('could not get clustered state; aborting');
      return;
    }
  }
}
```

```
        if ((state = get('state')) != clustered) {
            audit('state is ' + state + '; aborting');
            return;
        }

        if ((state = get('peer_state')) != clustered) {
            audit('peer state is ' + state + '; aborting');
            return;
        }

        run('cd /');
        run('confirm maintenance system reboot');
    }
};
```

Using Scheduled Workflows

Workflows can be started via a timer event by setting up a schedule for them. The property `scheduled` has to be added to the Workflow Object and needs to be set to `true`. Schedules can either be created via the CLI once a workflow is loaded into the appliance or an array type property named `schedule` can be added to the Object Workflow.

Each schedule entry consists of the following properties:

TABLE 149 Workflow Schedule Properties

Property	Type	Description
NAME	String	Name of the schedule, system generated
frequency	String	minute, halfhour, hour, day, week, month
day	String	Specifies specific day and can be set to: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday or Sunday. Can be set when frequency is set to week or month
hour	String	00-23, Specifies the hour part of the schedule and can be specified when the frequency is set to a day, week or month.
minute	String	00-59, Specifies the minute part of the schedule.

Using a Scheduled Workflow

Once a workflow has been loaded into the appliance a schedule can be defined for it via the CLI interface as follows:

EXAMPLE 28 Scheduled Workflow in the CLI

```
dory:> maintenance workflows
dory:maintenance workflows> "select workflow-002''"
dory:maintenance workflow-002> schedules
dory:maintenance workflow-002 schedules> create
dory:maintenance workflow-002 schedule (uncommitted)> set frequency=day
frequency = day (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> set hour=10
hour = 10 (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> set minute=05
minute = 05 (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> commit
dory:maintenance workflow-002 schedules> list
```

NAME	FREQUENCY	DAY	HH:MM
schedule-001	day	-	10:05

```
dory:maintenance workflow-002 schedules> create
dory:maintenance workflow-002 schedule (uncommitted)> set frequency=week
frequency = week (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> set day=Monday
day = Monday (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> set hour=13
hour = 13 (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> set minute=15
minute = 15 (uncommitted)
dory:maintenance workflow-002 schedule (uncommitted)> commit
dory:maintenance workflow-002 schedules> list
```

NAME	FREQUENCY	DAY	HH:MM
schedule-001	day	-	10:05
schedule-002	week	Monday	13:15

```
dory:maintenance workflow-002 schedules>
```

Coding Workflow Schedules

Schedules can also be specified in the workflow code as a property in the Object workflow. The property syntax used here differs from the CLI schedule creation. Here three properties are used,

TABLE 150 Workflow Schedule Properties

Property	Type	Description
offset	Number	Determines the starting point in the defined period
period	Number	Defines the frequency of the Schedule
unit	String	Specifies if either seconds or month are used as unit in the offset and period definition

EXAMPLE 29 Illustrating the Use of Workflow Properties

The following code example illustrates the use of the properties. Note that inline arithmetic helps to make the offset and period declarations more readable.

```
// Example of using Schedule definitions within a workflow
var MyTextObject = {
  MyVersion: '1.0',
  MyName: 'Example 9',
  MyDescription: 'Example of use of Timer',
  Origin: 'Oracle'
};
var MySchedules = [
  // half hr interval
  { offset: 0, period: 1800, units: "seconds" },
  // offset 2 days, 4hr, 30min , week interval
  {offset: 2*24*60*60+4*60*60+30*60, period: 604800,units: "seconds" }
];
var workflow = {
  name: MyTextObject.MyName,
  description: MyTextObject.MyDescription,
  version: MyTextObject.MyVersion,
  alert: false,
  setid: true,
  schedules: MySchedules,
  scheduled: true,
  origin: MyTextObject.Origin,
  execute: function () {
    audit('workflow started for timer; ');
  }
};
```

The property units in the Object MySchedules specifies the type of units used for the properties offset and period. They can be set to either seconds or month. The property period specifies the frequency of the event and the offset specifies the units within the period. In the above example

the period in the second schedule is set for a week, starting at the second day, at 4:30. Multiple schedules can be defined in the property schedules.

The Object MySchedules in the example uses the following three properties:

- offset - This is the starting offset from January 1, 1970 for the schedule. The offset is given in the units defined by the property "units".
- period - This is the period between recurrences of the schedule which is also given in the units defined by the property "units".
- units - This can be defined in seconds or months.

The starting point for weekly schedules is Thursday. This is due to the fact that the epoch is defined as starting on 1 Jan 1970 which was a Thursday.

EXAMPLE 30 Workflow Schedule Shown in the CLI

In the above example the period in the second schedule uses a starting offset of 2 days + 4 hours + 30 minutes. This results in the starting date being January 3, 1970 at 4:30 am. The schedule recurs weekly indefinitely every Saturday at 4:30 am. Below you can see the display of the schedule in the CLI.

```
<small>dory:> maintenance workflows
dory:maintenance workflows> list
WORKFLOW    NAME                                OWNER SETID ORIGIN                VERSION
workflow-000 Configure for Oracle Solaris Cluster NFS root  false Oracle Corporation
1.0.0
workflow-001 Unconfigure Oracle Solaris Cluster NFS root  false Oracle Corporation
1.0.0
workflow-002 Configure for Oracle Enterprise Manager Monitoring root  false Sun
Microsystems, Inc. 1.1
workflow-003 Unconfigure Oracle Enterprise Manager Monitoring root  false Sun
Microsystems, Inc. 1.0</small>
```

```
dory:maintenance workflow-002 schedules>
```

NAME	FREQUENCY	DAY	HH:MM
schedule-000	halfhour	-	--:00
schedule-001	week	Saturday	04:30

Creating a Worksheet Based on a Specified Drive Type

Here is an example workflow that creates a worksheet based on a specified drive type:

EXAMPLE 31 Workflow Device Type Selection

```
var steps = [ {
  step: 'Checking for existing worksheet',
  execute: function (params) {
    /*
     * In this step, we're going to see if the worksheet that
     * we're going to create already exists. If the worksheet
     * already exists, we blow it away if the user has indicated
     * that they desire this behavior. Note that we store our
     * derived worksheet name with the parameters, even though
     * it is not a parameter per se; this is explicitly allowed,
     * and it allows us to build state in one step that is
     * processed in another without requiring additional global
     * variables.
     */
    params.worksheet = 'Drilling down on ' + params.type + ' disks';

    try {
      run('analytics worksheets select name="' +
        params.worksheet + '"');

      if (params.override) {
        run('confirm destroy');
        return;
      }

      throw ('Worksheet called "' + params.worksheet +
        '" already exists!');
    } catch (err) {
      if (err.code !== EAKSH_ENTITY_BADSELECT)
        throw (err);
    }
  }, {
  step: 'Finding disks of specified type',
  execute: function (params) {
    /*
     * In this step, we will iterate over all chassis, and for
     * each chassis iterates over all disks in the chassis,
     * looking for disks that match the specified type.
     */
    var chassis, name, disks;
    var i, j;

    run('cd /');
    run('maintenance hardware');
```

```

chassis = list();
params.disks = [];

for (i = 0; i < chassis.length; i++) {
  run('select ' + chassis[i]);

  name = get('name');
  run('select disk');
  disks = list();

  for (j = 0; j < disks.length; j++) {
    run('select ' + disks[j]);

    if (get('use') == params.type) {
      params.disks.push(name + '/' +
        get('label'));
    }

    run('cd ..');
  }

  run('cd ../../');
}

if (params.disks.length === 0)
  throw ('No ' + params.type + ' disks found');
run('cd /');
}, {
step: 'Creating worksheet',
execute: function (params) {
  /*
   * In this step, we're ready to actually create the worksheet
   * itself: we have the disks of the specified type and
   * we know that we can create the worksheet. Note that we
   * create several datasets: first, I/O bytes broken down
   * by disk, with each disk of the specified type highlighted
   * as a drilldown. Then, we create a separate dataset for
   * each disk of the specified type. Finally, note that we
   * aren't saving the datasets -- we'll let the user do that
   * from the created worksheet if they so desire. (It would
   * be straightforward to add a boolean parameter to this
   * workflow that allows that last behavior to be optionally
   * changed.)
   */
  var disks = [], i;

  run('analytics worksheets');
}

```

```

run('create '' + params.worksheet + '');
run('select name='' + params.worksheet + '');
run('dataset');
run('set name=io.bytes[disk]');

for (i = 0; i < params.disks.length; i++)
    disks.push('' + params.disks[i] + '');

run('set drilldowns=' + disks.join(',')');
run('commit');

for (i = 0; i < params.disks.length; i++) {
    run('dataset');
    run('set name="io.bytes[disk=' +
        params.disks[i] + ']"');
    run('commit');
}
}
} ];

var workflow = {
    name: 'Disk drilldown',
    description: 'Creates a worksheet that drills down on system, ' +
        'cache, or log devices',
    parameters: {
        type: {
            label: 'Create a new worksheet drilling down on',
            type: 'ChooseOne',
            options: [ 'cache', 'log', 'system' ],
            optionlabels: [ 'Cache', 'Log', 'System' ]
        },
        overwrite: {
            label: 'Overwrite the worksheet if it exists',
            type: 'Boolean'
        }
    },
    execute: function (params) { return (steps); }
};

```

Uploading and Executing Workflows Using the BUI

You can upload a workflow to the appliance by clicking on the plus icon . You can execute a workflow by clicking on the row specifying the workflow or by hovering over the workflow

row and clicking the execute icon . You can also edit or delete a workflow by hovering over the workflow row and clicking the appropriate icon.

FIGURE 34 Workflows Seen in the BUI

Workflows Total: 5		
NAME ▲	DESCRIPTION	VERSION
Clear locks	Clear locks held on behalf of an NFS client	1.0.0
Configure for Oracle Enterprise Manager Monitoring	Sets up environment to be monitored by Oracle Enterprise Manager	1.1
Configure for Oracle Solaris Cluster NFS	Sets up environment for Oracle Solaris Cluster NFS	1.0.0   
Unconfigure Oracle Enterprise Manager Monitoring	Removes the artifacts from the appliance used by Oracle Enterprise Manager	1.0
Unconfigure Oracle Solaris Cluster NFS	Removes the artifacts from the appliance used by Oracle Solaris Cluster NFS	1.0.0

To view the list of expanded workflows, hold Shift and click the plus icon . To hide the expanded list, hold Shift and click the plus icon  again.

▼ Downloading Workflows using the CLI

1. **Workflows are downloaded to the appliance via the `download` command, which is similar to the mechanism used for software updates:**

```
dory:maintenance workflows> download
dory:maintenance workflows download (uncommitted)> get
      url = (unset)
      user = (unset)
      password = (unset)
```

2. **You must set the "url" property to be a valid URL for the workflow. This may be either local to your network or over the internet. The URL can be either HTTP (beginning with "http://") or FTP (beginning with "ftp://"). If user authentication is required, it may be a part of the URL (e.g. "ftp://myusername: mypasswd@myserver/export/foo"), or you may leave the username and password out of the URL and instead set the user and password properties.**

```
dory:maintenance workflows download (uncommitted)> set url=
      ftp://foo/example1.akwf
```

```
url = ftp://foo/example1.akwf
dory:maintenance workflows download (uncommitted)> set user=bmc
user = bmc
dory:maintenance workflows download (uncommitted)> set password
Enter password:
password = *****
dory:maintenance workflows download (uncommitted)> commit
Transferred 138 of 138 (100%) ... done
```

▼ Listing Workflows using the CLI

1. **To list workflows, use the `list` command from the maintenance workflows context:**

```
<small>dory:maintenance workflows> list
WORKFLOW    NAME                                OWNER SETID ORIGIN                VERSION
workflow-000 Configure for Oracle Solaris Cluster NFS root  false Oracle Corporation
1.0.0
workflow-001 Unconfigure Oracle Solaris Cluster NFS root  false Oracle Corporation
1.0.0
workflow-002 Configure for Oracle Enterprise Manager Monitoring root  false Sun
Microsystems, Inc. 1.1
workflow-003 Unconfigure Oracle Enterprise Manager Monitoring root  false Sun
Microsystems, Inc. 1.0</small>
```

2. **To view workflows, use the `show` command from the maintenance workflows context:**

```
dory:maintenance workflows> select workflow-001
dory:maintenance workflow-001> show
Properties:
    name = Configure for Oracle Solaris Cluster NFS
    description = Sets up environment for Oracle Solaris Cluster NFS
    owner = root
    origin = Oracle Corporation
    setid = false
    alert = false
    version = 1.0.0
    scheduled = false
```

3. **To select a workflow, use the `select` command:**

```
dory:maintenance workflows> select workflow-000
dory:maintenance workflow-000>
```

4. **To get a workflow's properties, use the `get` command from within the context of the selected workflow:**

```
dory:maintenance workflow-000> get
      name = Hello world
      description = Bids a greeting to the world
      owner = root
      origin = <local>
      setid = false
      alert = false
      scheduled = false
```

▼ Executing Workflows using the CLI

1. **To execute a workflow, use the `run` command from within the context of the selected workflow.**

```
dory:maintenance workflow-000> run
```

2. **The context will become a captive context in which parameters must be specified for a workflow which has parameters:**

```
dory:maintenance workflow-000> run
dory:maintenance workflow-000 run(uncommitted)> get
      type = (unset)
      overwrite = (unset)
```

For a workflow which has no parameters, you can commit directly after entering the captive context:

```
dory:maintenance workflow-000> run
dory:maintenance workflow-000 run(uncommitted)> commit
hello world!
```

3. **Any attempt to commit the execution of the workflow without first setting the requisite parameters will result in an explicit failure:**

```
dory:maintenance workflow-000 run(uncommitted)> commit
error: cannot execute workflow without setting property "type"
```

4. **To execute the workflow, set the specified parameters, and then use the `commit` command:**

```
dory:maintenance workflow-000 run (uncommitted)> set type=system
      type = system
dory:maintenance workflow-000 run (uncommitted)> set overwrite=true
      overwrite = true
dory:maintenance workflow-000 run (uncommitted)> commit
```

5. **If the workflow has specified steps, those steps will be displayed via the CLI, e.g.:**

```
dory:maintenance workflow-000 run (uncommitted)> commit
Checking for existing worksheet ... done
Finding disks of specified type ... done
Creating worksheet ... done
```

▼ Auditing Workflows using the CLI

All workflows have a checksum property computed by the system. This checksum is the SHA-256 digest of the workflow content. To determine if a workflow has changed, compare the workflow checksum against its previous checksum.

- **To obtain the checksum property of a workflow, use the `get checksum` command from the `maintenance workflows` context:**

```
hostname:maintenance workflows> select workflow-001
hostname:maintenance workflow-001> get checksum
checksum = 15f4188643d7add37b5ad8bda6d9b4e7210f1cd890a73df176382e800aec
```

Integration

The Oracle ZFS Storage Appliance product is engineered to seamlessly integrate with other Oracle products as well as third-party and virtualized environments. The appliance provides a full suite of data protocols to communicate with a wide variety of application hosts. To improve application performance, provide effective backup support, or more tightly integrate with your application environment and manage and monitor storage appliances, see these resources for a complete list of available downloadable plug-ins, and for white papers and documentation detailing configuration and deployment best practices and recommendations for maximizing performance:

- [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html \)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html)
- [NAS Storage Documentation \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html)

For information about plug-ins that provide functionality with Oracle and non-Oracle products, see:

- [“Plug-ins for Oracle Products” on page 682](#)
- [“Plug-ins for Non-Oracle Products” on page 684](#)

Please check the Oracle ZFS Storage Appliance Plug-ins Download page for a complete and up-to-date list of available plug-ins and their latest versions and documentation.

Oracle ZFS Storage Appliance can be a backup target for Oracle Database files residing on Oracle Exadata, Oracle SuperCluster, and Oracle Database Appliance. For the most up-to-date configuration best practices and recommendations, check:

- [NAS Storage Documentation, White Papers and Solutions Briefs \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html)
- [My Oracle Support \(MOS\)](#) notes in the Knowledge Center, including:
 - "Guidelines When Using ZFS Storage in an Exadata Environment" (Doc ID 2087231.1)
 - "RMAN Backup From Oracle SuperCluster to Oracle ZFS Storage Appliance" (Doc ID 1517107.1)

Additionally, database administrators can ensure data protection, availability, and rapid testing by efficiently provisioning, cloning, backing up, and restoring Oracle Database files on Oracle ZFS Storage Appliance using the Oracle Snap Management Utility for Oracle Database. This robust software product is available outside of the plug-in download offerings. For more information, see:

- [Oracle Snap Management Utility for Oracle Database \(https://www.oracle.com/storage/nas/snap/index.html\)](https://www.oracle.com/storage/nas/snap/index.html)

Configuring Oracle ZFS Storage Appliance for Oracle Database Clients

The appliance offers a number of unique features designed to integrate with Oracle Database clients, including Hybrid Columnar Compression (HCC) and Oracle Intelligent Storage Protocol (OISP).

To enable these features, the SNMP service on the Oracle ZFS Storage Appliance must be configured to allow SNMP queries by database clients. The client uses this mechanism to identify a storage device as a Oracle ZFS Storage Appliance. For specifying the database client hostname or IP address as the trap destination, see [“Configuring SNMP to Serve Appliance Status \(BUI\)” on page 367](#).

To verify the appliance SNMP service is configured properly, run the `snmpget (1)` command from the client system, replacing `<host>` with the name or IP address of the appliance.

```
-bash-4.1$ snmpget -v1 -c public <host> .1.3.6.1.4.1.42.2.225.1.4.2.0
SNMPv2-SMI::enterprises.42.2.225.1.4.2.0 = STRING: "Oracle ZFS Storage Appliance"
```

Plug-ins for Oracle Products

The following plug-ins provide functionality with Oracle products:

- [“Oracle Enterprise Manager Plug-in for Oracle ZFS Storage Appliance” on page 683](#)
- [“Oracle VM Storage Connect Plug-in for Oracle ZFS Storage Appliance” on page 683](#)
- [“Oracle ZFS Storage Appliance Network File System Plug-in for Oracle Solaris Cluster” on page 684](#)
- [“Oracle ZFS Storage Appliance Plug-in for Oracle Solaris Cluster Geographic Edition” on page 684](#)

Oracle Enterprise Manager Plug-in for Oracle ZFS Storage Appliance

The Oracle Enterprise Manager Plug-in for Oracle ZFS Storage Appliance extends Oracle Enterprise Manager Cloud Control to support monitoring and limited management of Oracle ZFS Storage Appliance. The plug-in has the following features:

- Gathers and presents storage system, configuration, and performance information for Oracle ZFS Storage Appliance in both single target and group presentations
- Raises alerts for pre-selected configuration and monitoring data
- Provides Business Intelligence (BI) Publisher reports that complement Oracle ZFS Storage Analytics reports
- Ties together Oracle Databases deployed on NFS shares from Oracle ZFS Storage Appliance
- Provides configurable job utilities in the Enterprise Manager Job Library that can provision users on target appliances, manage dataset retention policies and manage collections of datasets
- Provides target-based storage provisioning on Oracle ZFS Storage Appliance

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle VM Storage Connect Plug-in for Oracle ZFS Storage Appliance

The Oracle Storage Connect Plug-in for Oracle ZFS Storage Appliance is a component of the Oracle VM software suite that enables Oracle VM to provision and manage Oracle ZFS Storage Appliance for virtualization.

For further information about the Oracle VM software suite, see <http://www.oracle.com/technetwork/server-storage/vm/overview/index.html>.

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Network File System Plug-in for Oracle Solaris Cluster

Oracle ZFS Storage Appliance Network File System Plug-in for Oracle Solaris Cluster provides integration for the appliance with NFS protocol operations directed by Oracle Solaris Cluster software, which provides high availability, including efficient, fast disaster recovery, for applications and virtualized workloads. The plug-in is integrated into the Oracle Solaris Cluster software.

For information on the plug-in, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Plug-in for Oracle Solaris Cluster Geographic Edition

Oracle ZFS Storage Appliance Plug-in for Oracle Solaris Cluster Geographic Edition provides integration for Oracle ZFS Storage Appliance with Oracle Solaris Cluster software, particularly its ability to provide high availability and fast, efficient disaster recovery for geographically distanced systems through the remote replication feature of the appliance.

This plug-in is an application-programming interface to the Oracle Solaris Cluster Geographic Edition software, so it is an integrated component of that download.

For information on the plug-in, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Plug-ins for Non-Oracle Products

The following plug-ins provide functionality with non-Oracle products:

- [“Oracle ZFS Storage Appliance Virtual Storage Manager Plug-ins for VMware vSphere and VMware vSphere Web Client”](#) on page 685
- [“Oracle ZFS Storage Appliance Storage Replication Adapter for VMware Site Recovery Manager”](#) on page 686
- [“Oracle ZFS Storage Appliance Plug-in for VMware vSphere Storage APIs for Array Integration – NAS”](#) on page 686

- “Oracle ZFS Storage Appliance Provider for VMware vSphere APIs for Storage Awareness” on page 686
- “Oracle ZFS Storage Appliance Provider for Volume Shadow Copy Service Software” on page 687
- “Oracle ZFS Storage Appliance Plug-in for Veritas NetBackup OpenStorage” on page 687
- “Oracle ZFS Storage Appliance Plug-in for CommVault Simpana IntelliSnap” on page 688

Oracle ZFS Storage Appliance Virtual Storage Manager Plug-ins for VMware vSphere and VMware vSphere Web Client

Beginning with VMware vSphere version 5.1, VMware provides a web-based interface, vSphere Web Client, to monitor and manage virtual data centers. The Oracle ZFS Storage Appliance Virtual Storage Manager Plug-in for VMware vSphere Web Client is a management interface to Oracle ZFS Storage Appliance that operates within the vSphere Web Client to provide users with a single 'pane-of-glass' view when provisioning and monitoring storage for use in their vSphere virtual data centers.

The Oracle ZFS Storage Appliance Virtual Storage Manager Plug-in for VMware vSphere provides monitoring for all Oracle ZFS Storage Appliance models in the VMware vSphere environment. The plug-in provides the following primary features:

- Gathers and presents share and project information for one or more appliances
- Provides basic share and project provisioning on an appliance
- Provides access to the Analytics facility on an appliance

Designed to support VMware vSphere versions 5.5 and earlier, (VMware vSphere 5.5.x and later employs a Web Client), the plug-in employs a primary interface referred to as the Oracle ZFS Virtual Storage Manager.

For the latest versions of these plug-ins, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Storage Replication Adapter for VMware Site Recovery Manager

Oracle ZFS Storage Appliance Storage Replication Adapter for VMware Site Recovery Manager integrates the appliance into VMware deployments that span multiple sites and require fast recovery in the event of a protected site service disruption. The Storage Replication Adapter plugs into existing VMware vCenter Site Recovery Manager environments and allows the appliance to be managed through VMware vCenter Site Recovery Manager discovery, test, and failover sequences as the recovery plan is tested and run. Usage of the Storage Replication Adapter occurs entirely within the VMware vCenter Site Recovery Manager application.

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Plug-in for VMware vSphere Storage APIs for Array Integration – NAS

Oracle ZFS Storage Appliance Plug-in for VMware vSphere Storage APIs for Array Integration for network attached storage (NAS) offloads common virtual machine operations to the storage hardware in order to free up resources on VMware ESXi servers as well as network bandwidth. This plug-in product has been designed to accelerate VM disk file copy operations on appliances in vSphere environments. The plug-in is packaged as software that is distributed as a VMware installation bundle (VIB).

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Provider for VMware vSphere APIs for Storage Awareness

Oracle ZFS Storage Appliance Provider for VMware vSphere APIs for Storage Awareness is a software component that supports the VMware vStorage APIs for Storage Awareness framework. The provider collects data from appliances and delivers information about storage

topology, LUNs, and shares to the vCenter Server, enabling it to monitor the storage systems. The provider uses the REST API over HTTPS to collect the storage information. This product communicates with the clients (vCenter servers) through SOAP over HTTPS. The provider is capable of serving multiple storage appliances and multiple vCenter servers.

The provider is delivered as a virtual machine template image where the provider software is pre-installed and configured. The provider software runs from the virtual machine once the virtual template is deployed in the vSphere environment.

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Provider for Volume Shadow Copy Service Software

Oracle ZFS Storage Appliance Provider for Volume Shadow Copy Service Software is a hardware provider that interfaces with Microsoft Volume Shadow Copy Service to provide a backup infrastructure that coordinates applications with file system activities. This infrastructure can then create point in time, coalesced copies known as "shadow copies."

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Plug-in for Veritas NetBackup OpenStorage

Oracle ZFS Storage Appliance Plug-in for Veritas NetBackup OpenStorage performs as a remote interface to the appliance that allows NetBackup to back up, duplicate, and delete data from the appliance using the appliance's advanced features.

For the latest version of this plug-in, including full documentation with installation and administration information, see the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle ZFS Storage Appliance Plug-in for CommVault Simpana IntelliSnap

Oracle ZFS Storage Appliance Plug-in for CommVault Simpana IntelliSnap allows management connectivity for Simpana's IntelliSnap functions which provide point-in-time snapshots of data files. The plug-in includes support for Fibre Channel protocol and Microsoft Hyper-V virtual environments.

This plug-in is integrated in the Simpana software. It does not require a separate package installation as long as you have the right minimum Service Pack of Simpana installed. Download and install CommVault Simpana 10 SP10 or higher for this plug-in. The latest release for CommVault Simpana can be downloaded from: <https://ma.commvault.com>.

The online documentation for the plug-in is located at: http://documentation.commvault.com/commvault/v10/article?p=features/snap_backup/oracle_zfs/overview.htm. Additionally, release notes for the plug-in are located on the Oracle Technology Network (OTN) page: [Oracle ZFS Storage Appliance Plug-in Downloads \(http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html\)](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html).

Oracle Intelligent Storage Protocol

Oracle Intelligent Storage Protocol (OISP) enables the Oracle direct NFS (dNFS) client to encode and pass attributes associated with I/O requests to the appliance. These attributes contain such information as the type of database file that the I/O request is targeting, the record size of the file, whether to cache the I/O data, and the identity of the database issuing the I/O request.

The appliance decodes these attributes, using them to simplify database configuration, increase database performance, and provide observability into the source of I/O workloads being generated by database clients.

Database Record Size

The Oracle dNFS client can pass the optimal record size based on the type of file for each I/O request. If a record size is passed, it overrides the "Database record size" property setting on the share or project. The record size can only be set for newly created files. If a file already exists, the record size is not changed.

Synchronous Write Bias Hint

The Oracle dNFS client can pass a write bias "hint" associated with write I/O requests that prompts the appliance to treat I/O requests as either latency sensitive or throughput oriented. If the hint is passed, it overrides the "Synchronous write bias" property setting on the share or project.

Analytics Breakdown by Database Name

The Oracle Database 12c dNFS client can pass the identification of the database (SID) or container database and pluggable database (SID:SID) responsible for issuing I/O requests. Oracle ZFS Storage Appliance analytics can display I/O statistics broken down by the SID name(s) of the database by selecting breakdown or drill by "Application ID."

With OS8.7 and later firmware on Oracle ZFS Storage Appliance, additional OISP database analytics may be displayed. OISP operations by client, file name, database name, database file type, database function, share, project, size file offset and latency are all available.

Caching Hints

The Oracle Database 12.2 or later dNFS client includes caching hints on I/O requests. Negative caching hints are included on I/O requests that do not expect to soon reference the data read or written again, such as datafile blocks read, and backup pieces written as part of Oracle Recovery Manager (Oracle RMAN) backup. This assists the appliance in making the best use of available memory in caching filesystem data. The main negatively cached operations are: Oracle RMAN reads and writes, Oracle Database datafile and redo log file creation, and Oracle Database Archiver reads and writes.

OISP-Capable Protocols and Clients

Protocols: NFSv4.0 and NFSv4.1

Clients: Oracle Database NFS (dNFS) client

Fibre Channel and iSCSI Support with Veritas Dynamic Multi-Pathing and Storage Foundation/InfoScale Foundation

The Oracle ZFS Storage Appliance product integrates with Veritas' business continuity (high availability) and storage management products, (previously Symantec) Storage Foundation and Veritas InfoScale. Refer to the [Veritas Hardware Compatibility List \(HCL\) for Storage Foundation and InfoScale](https://sort.veritas.com/hclcentral/diskarray) (<https://sort.veritas.com/hclcentral/diskarray>) for support information updates.

Supported components/platforms include:

- Veritas Storage Foundation (SF) 5.1, 6.0, 6.1, 6.2
- Storage Foundation High Availability (SFHA) 5.1, 6.0, 6.1, 6.2
- Storage Foundation Cluster File System (High Availability)/Storage Foundation for Oracle RAC (SFCFS(HA)/SF Oracle RAC) 5.1, 6.0, 6.1, 6.2
- InfoScale Foundation 7.0, 7.1, 7.2
- InfoScale Foundation + Availability 7.0, 7.1, 7.2
- InfoScale Storage/Storage+ Availability/Enterprise 7.0, 7.1, 7.2
- Oracle Solaris 10 SPARC and x86
- Oracle Solaris 11 SPARC and x86
- Linux RedHat5
- Oracle Linux

Be sure to check the Note column of the [Veritas Hardware Compatibility List \(HCL\) for Storage Foundation and InfoScale](https://sort.veritas.com/hclcentral/diskarray) (<https://sort.veritas.com/hclcentral/diskarray>) for any restrictions or special requirements.

Refer to Veritas' Host and Storage Configuration Guide at https://www.veritas.com/support/en_US/article.TECH47728.

Storage Foundation 6.1 supports both FC and iSCSI connections to the appliance for the following Windows versions:

- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

InfoScale Foundation 7.2 supports both FC and iSCSI connections to the appliance for the following Windows versions:

- Windows Server 2012
- Windows Server 2012 R2

- Windows Server 2016

For an up-to-date listing of Veritas' supported configurations for Windows with older versions of Storage Foundation refer to https://www.veritas.com/support/en_US/article.TECH77195.

Refer to the [Veritas Hardware Compatibility List \(HCL\) for Storage Foundation and InfoScale](https://sort.veritas.com/hclcentral/diskarray) (<https://sort.veritas.com/hclcentral/diskarray>) for support information updates.

