

DR Series System (Version 4.0.4)

Administrator's Guide



© 2021 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.



Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.
-  **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

DR Series System Administrator's Guide
Updated - April 2021
Version - (Version 4.0.4)

Contents

Introducing the DR Series system documentation	10
Understanding the DR Series system documentation	10
Other information you might need	10
What's new in this release	11
Source code availability	11
Introducing the DR Series system	12
Understanding the DR Series system models	13
Drive and available physical capacities	15
DR Series data storage concepts	15
Data deduplication and compression	15
Encryption at rest	16
Streams and connections	16
Replication	17
Replication seeding	18
Reverse replication	18
Reverse replication: alternate method	19
Rapid Air Gap	20
Secure erase	20
Supported file system protocols	20
NFS	21
CIFS	21
CIFS ACL support	21
Access control list support in containers	22
Unix permissions guidelines	23
Windows permissions guidelines	23
Rapid NFS and Rapid CIFS	24
DR Rapid for the DR Series system	24
RDA with OST for the DR Series system	25
Software components and operational guidelines	26
Supported virtual tape library access protocols	27
NDMP	27
iSCSI	27
Fibre channel	27
Understanding the DR Series system hardware and data operations	28
DR Series expansion shelf	29
Understanding the process for adding a DR Series expansion shelf	30
Understanding DR Series system hardware expansion shelf cabling	30
Supported software and hardware	32

Terminal emulation applications	33
Setting up the DR Series system	34
Interacting with the DR Series system	34
Networking prerequisites for the DR Series system	34
Supported connections for initializing the DR Series system	36
Default IP and subnet mask addresses	36
Local console connection	37
DR4000	37
DR4100/DR6000	38
DR4300e/DR4300/DR6300	39
iDRAC connection	41
Accessing iDRAC6/iDRAC7 by using RACADM	42
Logging on and initializing the DR Series system	42
Logging onto the system GUI for the first time	44
Enabling active scripting in IE	45
Disabling the Compatibility View settings for IE	46
Configuring the DR Series system settings	47
Viewing system configuration settings	47
Using the Configuration Wizard	48
Registering with the Support Portal	49
Configuring Active Directory settings	49
Adding a login group to an ADS domain	50
Removing a login group	51
Configuring system date and time settings	51
Configuring DR Series system enclosures	52
Managing and viewing DR Series system licenses	53
Configuring networking settings	55
Understanding the networking page and Ethernet port values	58
Understanding system operation scheduling	59
Configuring cleaner schedules	60
Viewing cleaner status	61
Viewing cleaner statistics	61
Configuring replication schedules	62
Configuring an SSL certificate for your DR Series system	63
Installing an SSL certificate	63
Resetting the SSL Certificate	64
Installing a CA certificate	64
Generating a CSR	65
Managing storage groups	65
Viewing storage group information	66
Adding a storage group	67

Modifying a storage group	68
Deleting a storage group	70
Managing users	70
Viewing users	70
Adding a user	71
Modifying a user	72
Changing a user password	72
Deleting a user	72
Modifying password reset options	73
Configuring email notification settings	73
Configuring share-level security for CIFS shares	74
Managing containers	76
Viewing containers	76
Creating a container	77
Adding an NFS or CIFS connection type container	78
Moving data into a CIFS type container	80
Creating an OST or RDS connection type container	81
Creating a VTL type container	81
Viewing VTL tape information	84
Editing container settings	85
Deleting a container	86
Viewing container statistics	86
Displaying container statistics by using the CLI	87
Managing replications	88
Guidelines and prerequisites for replication	88
Viewing replication information	89
Viewing replication statistics by using the CLI	90
Adding replication relationships	91
Modifying replication relationships	92
Deleting replication relationships	93
Starting and stopping replication	93
Adding a cascaded replica	93
Monitoring the DR Series system	95
Using the Dashboard page	95
Viewing DR Series system statistics by using the CLI	96
Monitoring system alerts	97
Monitoring clients	97
Monitoring system events	99
Monitoring system health	99

Monitoring system usage	100
Using GlobalView	102
Understanding GlobalView	102
Prerequisites for GlobalView	103
Viewing and using the GlobalView page	103
Adding a DR Series system to GlobalView	107
Removing a DR Series system from GlobalView	108
Configuring and using Rapid NFS and Rapid CIFS	109
Rapid NFS and Rapid CIFS benefits	109
Best practices: Rapid NFS	110
Best practices: Rapid CIFS	111
Setting client-side optimization	112
Installing the Rapid NFS plug-in	112
Installing the Rapid CIFS plug-in	113
Determining if your system is using Rapid NFS or Rapid CIFS	114
Viewing the Rapid NFS and Rapid CIFS logs	114
Viewing Rapid NFS Logs	115
Viewing Rapid CIFS Logs	115
Monitoring performance	115
Uninstalling the Rapid NFS plug-in	116
Uninstalling the Rapid CIFS plug-in	116
Configuring and using Rapid Data Access with NetVault Backup and with vRanger	117
Overview	117
Guidelines for using RDA with NetVault Backup and with vRanger	118
Best Practices: RDA with NetVault Backup and vRanger and the DR Series System	118
Setting client-side optimization	119
Adding RDS devices in NetVault Backup	119
Removing RDS Devices From NetVault Backup	120
Backing Up Data on the RDS container using NetVault Backup	120
Replicating Data to a RDS Container using NetVault Backup	121
Restoring data from a DR Series system by using NetVault Backup	122
Supported DR Series system CLI commands for RDS	123
Configuring and using RDA with OST	125
Understanding RDA with OST	125
Guidelines	126
Terminology	127
Supported RDA with OST software and components	127

Best Practices: RDA with OST and the DR Series System	128
Setting client-side optimization	128
Configuring an LSU	128
Installing the RDA with OST plug-in	129
Understanding the RDA with OST Plug-in (Linux)	129
Understanding the RDA with OST Plug-in (Windows)	130
Installing the RDA with OST plug-in for Backup Exec on Windows	130
Installing the RDA with OST plug-in for NetBackup on Windows	131
Uninstalling the RDA with OST plug-in for Windows	131
Installing the RDA with OST plug-in for NetBackup on Linux	132
Uninstalling the RDA with OST plug-in for Linux	133
Configuring DR Series system information using NetBackup	133
Using NetBackup CLI to add DR Series system name (Linux)	133
Using NetBackup CLI to add DR Series system name (Windows)	134
Configuring NetBackup for the DR Series system	134
Configuring NetBackup for optimized synthetic backups	135
Creating disk pools from LSUs	136
Creating storage units by using the disk pool	136
Backing up data from a DR Series system with NetBackup	137
Restoring data from a DR Series system with NetBackup	137
Duplicating backup images between DR Series systems with NetBackup	138
Using Backup Exec with a DR Series system (Windows)	138
RDA with OST Plug-In and Supported Versions	138
Installation Prerequisites for the RDA with OST Plug-in for Backup Exec	138
Configuring the DR Series system with the Backup Exec GUI	139
Creating backups on the DR Series system using Backup Exec	140
Optimizing duplication between DR Series systems using Backup Exec	140
Restoring data from a DR Series system with Backup Exec	141
Understanding the OST CLI commands	141
Supported DR Series System CLI Commands for RDA with OST	142
Understanding RDA with OST Plug-In Diagnostic Logs	143
Rotating RDA with OST Plug-In Logs for Windows	143
Modifying Log Rotation Size	143
Collecting diagnostics by using a Linux utility	144
Rotating RDA with OST plug-in logs for Linux	144
Guidelines for gathering media server information	145
NetBackup on Linux Media Servers	145
NetBackup on Windows Media Servers	145
Backup Exec on Windows Media Servers	146
Configuring and using VTL	148
Understanding VTL	148

Terminology	148
Supported virtual tape library access protocols	149
NDMP	149
iSCSI	149
Fibre channel	149
VTL and DR Series specifications	150
Guidelines for configuring VTL	151
Plan your Environment	151
Create Containers of Type VTL	152
Authentication/User Management Considerations	152
Verify the Tape Library Creation	153
Configure the Library in the DMA	153
Configuring and Using Encryption at Rest	154
Understanding Encryption at Rest	155
Encryption at Rest Terminology	156
Encryption at Rest and DR Series Considerations	157
Understanding the encryption process	158
Support, maintenance, and troubleshooting	159
Using the DR Series system support options	159
Viewing support information	159
Viewing system diagnostic log files	160
Understanding diagnostics collection	161
Generating a diagnostics log file	162
Downloading diagnostics log files	162
Deleting a Diagnostics Log File	163
Rebooting the DR Series system	163
Shutting down the DR Series system	163
Viewing the software upgrade page	163
Upgrading the DR Series system software	164
Restore Manager (RM)	165
Downloading the Restore Manager	165
Creating the Restore Manager USB key	166
Running the Restore Manager (RM)	166
Resetting the Boot LUN Setting in PERC H700 BIOS After Running RM	167
Hardware removal or replacement	167
Proper shutdown and startup of the system	168
DR Series System NVRAM	168
NVRAM Backup Failure Recovery	169
NVRAM Field Replacement	169
Troubleshooting error conditions	169
DR Series system alert and event messages	170

About the DR Series system maintenance mode 199

Supported Ports in a DR Series system 202

About us 204

 Contacting Quest 204

 Technical support resources 204

Introducing the DR Series system documentation

The DR Series system documentation contains topics that describe how to perform data storage operations and to manage storage and replication containers.

Understanding the DR Series system documentation

The topics in this administrator's guide introduce and describe how to use the DR Series web-based graphical user interface (GUI) to manage your system. It describes how to access the comprehensive system GUI and the associated DR Series system features and capabilities, how to perform a wide variety of data storage and replication operations, how to manage the system, as well as how to manage the related storage and replication containers.

In addition to the DR Series system GUI, you can manage the DR Series system by using a command-line interface (CLI). In some instances, the DR Series system GUI provides additional features and options that are not available in the DR Series system CLI and vice versa. For example, GlobalView is only available in the GUI, while the ability to add and remove clients is only available in the CLI. For more information about the DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

This documentation is written for an administrator.

i **NOTE:** For information about the supported web browsers you can use with the DR Series system, see the *DR Series System Interoperability Guide*.

Other information you might need

Refer to the safety and regulatory information that shipped with your system. Warranty information may be included within this document or as a separate document. Other DR Series system related documentation includes the following documents, which are available at support.quest.com/dr-series by selecting your specific DR model and then navigating to **Technical Documentation**.

- *DR Series System Owner's Manual* — provides information about solution features and describes how to troubleshoot the system as well as how to install or replace hardware versions of the DR Series system components.
- *DR Series System Command Line Reference Guide* — provides information about managing DR Series system data backup and replication operations using the DR Series system command line interface (CLI).
- *DR Series System Getting Started Guide* — provides an overview of how to set up your DR Series system hardware and includes technical specifications.
- *Setting Up Your DR Series System* — provides information about network, initial setup, and user account settings needed to initialize the DR Series system.
- *DR Series System Interoperability Guide* — provides detailed information about the supported hardware and software requirements for the DR Series system.
- *DR2000v Deployment Guide* — provides information about deploying the virtual DR Series system, DR2000v.
- *DR Series System Release Notes* — provides the latest information about new features and known issues with a specific product release.
- Any media that ships with your system that provides documentation and tools for configuring and managing your system, including those pertaining to the operating system, system management software, system updates, and system components that you purchased with your system.

i **NOTE:** Always check for the latest release notes at support.quest.com/dr-series and read the release notes first because they contain the most recently documented information about known issues with a specific product release.

What's new in this release

For a list of the features, enhancements, and changes in the latest release, see the section, “What’s New In This Release” in the *DR Series System Release Notes*. If you are upgrading from a previous software version, please see, “Upgrade Notes,” in the *DR Series System Release Notes*. You can download the latest documentation, including release notes, at support.quest.com/dr-series by selecting your specific DR model and then navigating to Technical Documentation.

Source code availability

A portion of the DR Series system software may contain or consist of open source software, which you can use under the terms and conditions of the specific license under which the open source software is distributed. Under certain open source software licenses, you are also entitled to obtain the corresponding source files. For more information or to find the corresponding source files for respective programs, see the Quest website at opensource.quest.com.

Introducing the DR Series system

The DR Series system is a high-performance, disk-based backup and recovery appliance that is simple to deploy and manage and offers unsurpassed Total Cost of Ownership benefits. Features such as innovative firmware and an all-inclusive licensing model ensure optimal functionality and provide the assurance of no hidden costs for valuable future features.

i | **NOTE:** Unless otherwise noted, later references in this guide to "the system" or "DR Series system" are used interchangeably to represent the DR Series system.

A purpose-built backup-to-disk appliance, the DR Series system provides advanced deduplication and compression technology to store data most efficiently. The DR Series hardware appliances are 2U, rack-based, system backup storage repositories, that include deduplication and compression technology in their operating systems. A virtual machine (VM) version is also available to provide robust, disk-based data backup capability on Virtual Machine host servers, while taking advantage of replicating to a deduplication-enabled appliance.

The Quest Data Protection | DR Series of backup and deduplication appliances support all of the major backup software applications in use today and can lower your backup storage costs to as little as \$.16/GB while reducing your total cost of ownership. (For a complete list of supported backup software, see the *DR Series System Interoperability Guide*.) The purpose-built appliances achieve these results using patented Rapid technology as well as built-in, variable block-based deduplication and compression. The DR Series helps you to:

- Reduce your backup storage footprint
- Speed up recovery
- Reduce or eliminate the need for physical tapes for backup
- Optimize network bandwidth by lowering the amount of data sent to disaster recovery sites

Other benefits include:

- Supports major backup applications for easy deployment.
- Lowers backup storage costs to as little as \$.16/GB using deduplication and compression.
- Speeds data ingest by up to 29TB/hr with built-in protocol accelerators.
- Decreases TCO with all-inclusive licensing that includes replication, encryption, protocol accelerators and all future feature releases.

- Enhances data protection with built-in software safeguards (early write verify and continuous data protection).
- Provides best-in-class hardware features (NVRAM, data integrity scans, RAID6 storage, hot spares).
- Contains built-in AES 256-bit encryption for data in motion or at rest.
- Allows backup to VTL libraries using iSCSI, NDMP, FC protocols.
- Incorporates 13th generation of Dell PowerEdge servers (DR4300e, DR4300 and DR6300).
- Offers in-place capacity expansion (DR4300e).
- Delivers the highest density deduplication target appliances on the market today.

The DR Series system includes the following features:

- Advanced data protection and disaster recover
- Two management interfaces: a command line interface (CLI) or a system graphical user interface (GUI) for the system software to manage storage containers.
- Support for a wide variety of data backup installations and environments.
- A simple installation process that provides full, intuitive remote setup and management capabilities.

The DR Series system is available in a variety of drive capacities and is ideal for SMB, enterprise, and remote office environments. For details about specific drive capacities and models available, see the *DR Series System Interoperability Guide* or the latest *DR Series System Release Notes*.

i NOTE: DR Series system hardware also supports the use of external data storage expansion shelves (also known as expansion enclosures). An added expansion shelf enclosure must be equal to or greater than each DR Series system internal drive slot capacity (0–11). For more information about expansion enclosures, see the topic, “Expansion Unit Limits,” in the *DR Series System Interoperability Guide* and the related Expansion Shelf topics in this guide.

Understanding the DR Series system models

The DR Series system is currently available in the following models:

- DR2000v—a virtual appliance, based on a Virtual Machine (VM) template for ESX and Hyper-V.
- DR4300e core—consists of preinstalled DR Series system software on a modified Dell PowerEdge R730xd appliance platform with no WAM2 card installed.
- DR4300e standard—consists of preinstalled DR Series system software on a modified Dell PowerEdge R730xd appliance platform.
- DR4300—consists of preinstalled DR Series system software on a modified Dell PowerEdge R730xd appliance platform and offers a higher base capacity than the DR4300e.
- DR6300—consists of preinstalled DR Series system software on a modified Dell R730xd appliance platform and offers a higher base capacity than the DR4300.

The DR Series system software is also supported on the following earlier models of the DR Series system.

- **DR4000**—consists of pre-installed DR Series system software on a Dell PowerEdge R510 appliance platform.
- **DR4100**—consists of preinstalled DR Series system software on a Dell PowerEdge R720xd appliance platform.
- **DR6000**—consists of preinstalled DR Series system software on a Dell PowerEdge R720xd appliance platform and offers a higher level of base system hardware.

The DR Series system consists of the following components:

- **Software** — System software and firmware is pre-installed, which supports record linkage and context-based lossless data compression methods.

i **NOTE:** The DR2000v system offers a VM template in various capacities for ESX and HyperV that can be deployed on your existing VM infrastructure.

- **Hardware** — The hardware supporting the DR Series systems is listed below:
 - **DR4000 system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and two cabled 2.5-inch SAS drives for the operating system. The operating system is installed on two 2.5-inch internal drives that are in a RAID 1 configuration in the DR4000 system.
 - **DR4100 system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.
 - **DR6000 system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.
 - **DR4300e core system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.
 - **DR4300e standard system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.
 - **DR4300 system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.
 - **DR6300 system:** Includes twelve 3.5-inch SAS or Nearline SAS chassis drives that are hot-swappable, two power supplies for power redundancy, and includes two 2.5-inch drives that are hot-pluggable in the rear.

i **NOTE:** For the DR4000, DR4100 and DR6000 systems, there is a global hot spare for both OS and data drives. For the DR4300e core and standard, DR4300 and DR6300 systems, there is a dedicated hot spare for data drives only (and not for OS drives).

i **NOTE:** For slot locations for the twelve 3.5-inch drives in the hardware appliance-based DR Series system types, see the topic, “DR Series System and Data Operations.”

- **Expansion shelf**—The hardware system appliance supports the addition of external Dell PowerVault MD1200 (for DR4000, DR4100, DR6000 systems) and MD1400 (for DR4300e core and standard, DR4300, DR6300 systems) data storage expansion shelf enclosures. Adding an expansion shelf provides additional data storage for the DR Series system and also requires a license. Each added expansion shelf enclosure must be equal to or greater than each DR Series system internal drive slot capacity (0–11). For more information, see the topic, “Expansion Unit Limits,” in the *DR Series System Interoperability Guide* and the related Expansion Shelf topics in this guide.

Drive and available physical capacities

The internal system drive capacity and available physical capacities of the DR Series system vary, depending on the system type and the drives installed. For detailed information, see the *DR Series System Interoperability Guide*, which describes the internal system drive capacity and available physical capacity (in decimal and binary values). This guide also includes the available capacities per virtual machine operating system (OS) for the DR2000v.

DR Series data storage concepts

Data deduplication and compression

The DR Series system design uses various data-reduction technologies, including advanced deduplication algorithms, in addition to the generic and custom compression solutions that prove effective across many differing file types. Data deduplication and compression are addressed in the following areas:

- **DR Series System** — The DR Series system backup and recovery appliances provide both efficient and high-performance disk-based data protection to leverage the advanced deduplication and compression capabilities in the DR Series system software. The DR Series systems provide a key component that performs backup, recovery, and data protection operations.
- **Deduplication** — This technology eliminates redundant copies of data and in the process it decreases disk capacity requirements and reduces the bandwidth needed for data transfer. Deduplication can be a major asset for companies that are dealing with increasing data volumes and require a means for optimizing their data protection.
- **Compression** — This technology reduces the size of data that is stored, protected, and transmitted. Compression helps companies improve their backup and recovery times while helping reduce infrastructure and network resource constraints.

In general, DR Series systems are disk-based data protection appliances that offer advanced deduplication and compression capabilities to reduce the time and cost associated with backing up and restoring data. Based on deduplication and compression technology, the DR Series systems eliminate the need to maintain multiple copies of the same data. This lets customers keep more data online longer and reduce the need for tape backup dependency.

Using its deduplication and compression technology, DR Series systems can help achieve an expected data reduction ratio of 15:1. Achieving this reduction in data means that you need fewer incremental storage operations to run and it provides you with a smaller backup footprint. By removing redundant data, DR Series systems deliver fast reliable backup and restore functionality, reduce media usage and power and cooling requirements, and improve your overall data protection and retention costs.

You can extend the benefits of data deduplication across the enterprise as well by using the DR Series system deduplication replication function—to provide a complete backup solution for multi-site environments. With 64:1 deduplicated replication (32:1 for DR4000 and DR4100, 8:1 for DR2000v), up to 64 nodes can be replicated simultaneously to separate, individual containers on one node. The DR Series systems use compression with replication to shrink the data that is needed to be moved across the wire to a container.

Replication can be scheduled based on your settings to occur during non-peak periods. The replication schedule you create can be set and prioritized to ingest data over replication data to ensure the most optimal back up windows based on your needs.

Unlike NFS and CIFS containers, OST and RDS container replication is handled by the Data Management Applications (DMAs) media servers.

The DR Series system supports the 64:1 replication of data (32:1 if on the DR4000 or DR 4100 and 8:1 for the DR2000v), whereby up to 64 source DR Series systems can write data to different individual containers on a single, target DR Series system. This supports, for example, the use case where branch or regional offices can each write their own data to a separate, distinct container on a main corporate DR Series system.

NOTE: Be aware that the storage capacity of the target DR Series system is directly affected by the number of source systems writing to its containers and by the amount being written by each of the source systems.

If the source and target systems reside in different Active Directory (AD) domains, then the data that resides on the target DR Series system may not be accessible. When AD is used for authentication for DR Series systems, the AD information is saved with the file. This can serve to restrict user access to the data based on the type of AD permissions that are in place.

NOTE: This same authentication information is replicated to the target DR Series system when you have replication configured. To prevent domain access issues, ensure that both the target and source systems reside in the same Active Directory domain.

For a complete list of supported management application, refer to the *DR Series System Interoperability Guide*.

Encryption at rest

Data that resides on the DR Series system can be encrypted. When encryption is enabled, the DR Series system uses the Industry standard FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) encryption algorithm for encrypting and decrypting user data. The content encryption key is managed by the key manager, which operates in either a Static mode or an Internal mode. In Static mode, a global, fixed key is used to encrypt all data. In internal mode, key lifecycle management is performed in which the keys are periodically rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days. A user-defined passphrase is used to generate a pass phrase key, which is used to encrypt the content encryption keys. It is mandatory to define a passphrase to enable encryption. The system supports up to a limit of 1023 different content encryption keys.

NOTE: Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

Streams and connections

This topic describes the differences between data streams and application connections.

Streams refer to the number of files written at the same time to a DR Series system. The DR Series system tracks the number of files being written and assembles the data into 4-MB chunks before processing that section of the

data. If the stream count is exceeded, the data is processed out of order and overall deduplication savings can be affected. For details on maximum stream count, see the *DR Series System Interoperability Guide*.

Connections are created by applications; and, within a single connection, there can be multiple streams depending on the application and the number backup jobs running in parallel over that single connection. Replication can use up to 16 streams over a single port using one connection.

For example, suppose you are running backups using Backup Exec and using DR4100 and the CIFS protocol. If you have:

- One Backup Exec server connected to the DR4100 over CIFS and one backup running, you have **one connection** and **one stream**.
- One Backup Exec server connected to the DR4100 over CIFS with 10 concurrent backups running, you have **one connection** and **ten streams**. This means that Backup Exec is writing ten different files to the DR4100.

Replication

Replication is the process by which key data is saved from storage locations, with the goal of maintaining consistency between redundant resources in data storage environments. Data replication improves the level of fault-tolerance, which improves the reliability of maintaining saved data and permits accessibility to the same stored data.

The DR Series system uses an active form of replication that lets you configure a primary-backup scheme. During replication, the system processes data storage requests from a specified source to a specified replica target, which acts as a replica of the original source data. This replica can then be cascaded optionally to a third location called a Cascaded replica for an additional copy.

i NOTE: The DR Series system software includes version checking that limits replication only between other DR Series systems that run the same system software release version. If versions are incompatible, the administrator is notified by an event.

Replicas/Cascaded replicas are read-only and are updated with new or unique data during scheduled or manual replications. The DR Series system can be considered to act as a form of a storage replication process in which the backup and deduplication data is replicated in real-time or via a scheduled window in a network environment. In a replication relationship between two or three DR Series systems, this means that a relationship exists between a number of systems. One system acts as the source and the other as a replica, with an optional third cascaded replica if you have chosen to keep two instances of replicated data in your backup workflow.

Replication is done at the container level and is one directional from source to replica to optional cascaded replica; however, since replication is done at the container level you can set up various containers to meet your specific replication requirements for your specific workflow. This form of replication is supported for the CIFS, NFS, Rapid CIFS, and Rapid NFS protocols and is fully handled by the DR Series system.

i NOTE: Refer to the *DR Series Interoperability Guide* for information about the maximum number of files replicated per container at a time per DR Series system.

For VTL type containers, replication is done at the cartridge level, and the system replicates the media/tape cartridges from a source DR Series system to a target DR Series system. The media is replicated to the target with the same markers. This condition restores successfully for a multi-domain environment; however, if the source and target are in the same domain, the media must be re-serialized on the target side once replication completes to successfully restore. This requires you to activate the VTL container, and change the bar codes of the tapes on the target DMA.

Unlike VTL, NFS, CIFS, Rapid NFS or Rapid CIFS containers, RDA with OST, RDA with NetVault Backup, and RDA with vRanger container replication is handled by the media servers of the respective Data Management Applications (DMAs).

The DR Series system supports the 64:1 replication of data (32:1 if on DR4X00 and 8:1 on DR2000v), whereby up to 64 source DR Series systems can write data to different individual containers on a single, target DR Series system. This supports the use case where branch or regional offices can each write their own data to a separate, distinct container on a main corporate DR Series system.

i NOTE: The storage capacity of the target DR Series system is directly affected by the number of source systems writing to its containers, and by the amount being written by each of the source systems.

If the source and target systems (replica or cascaded replica) are in different Active Directory (AD) domains, then the data that resides on the target system may not be accessible. When AD is used to perform authentication for DR Series systems, the AD information is saved with the file. This can act to restrict user access to the data based on the type of AD permissions that are in place.

i NOTE: This same authentication information is replicated to the target DR Series system when you have replication configured. To prevent domain access issues, ensure that both the target and source systems reside in the same Active Directory domain.

Replication seeding

The DR Series system supports replication seeding, which provides the ability to create a local seed and place it in a remote system. The seed backup is a process on the source DR Series system, which collects all of the unique data chunks from the containers and stores them on the target device. This is helpful if you have a new replication target DR to set up, the amount of data to be replicated is very large, and the network bandwidth is low. You can seed the target replica with the source data saved on a third party device, for example, a CIFS—mounted share, attach it to the target DR and then get the data into the target DR. Once the seeding is complete, replication is enabled between source and target and replication re-synchronization is done to complete any pending data transfers. Thereby, continuous replication can be done, which reduces network traffic significantly, and data can be replicated and synced with the target in a short amount of time.

i NOTE: The following scenarios are not supported for seeding:

- Import AND export from one share/device cannot occur at the same time.
- Import from one share/device cannot be completed from multiple locations at the same time.
- Export to a mount point can be completed only from one seed job. Multiple seed export jobs cannot send data to a single mount point.

You can initiate seeding using the DR Series system command line interface (CLI), and the data to be seeded is gathered in an organized manner and stored in the target devices. Refer to the *DR Series System Command Line Reference Guide* for more information about replication seeding support.

Reverse replication

The concept of reverse replication is not a supported operation on the DR Series system. This is because replica containers are always in a R-O (read-only) mode on the DR Series system, thus making write operations a non-supported operation.

Alternate ways to retrieve data

Under very specific conditions, it could be possible for replica containers to support a type of write operation whose sole function is to restore data from an archival target. For example, data could be replicated back to the remote site where a data management application (DMA), or backup software, is connected to allow this data to be restored directly.

This specific type of case applies only to configurations where data is backed up from a remote location to a local container, and then replicated over a WAN to a replica container that is backed up to tape. The data needs to be restored from the tape backup to the original location; first back to a DR Series system replica container, and then back to the original source location of the data on the other side of the WAN link.

i NOTE: If you choose to use this alternate workaround method, you must set up a new data storage unit in your DMA, and import the images before a restore to the original location can occur.

To leverage this type of deduplication across the WAN, complete the following:

1. Make sure that the replication operation has completed (between source and target).
2. Delete the current replication relationship, and re-create a replication relationship (reversing the source and target roles).
3. Restore data to the original source container (now the target).
4. Make sure that the replication operation has completed.
5. Delete the replication relationship and re-create a replication relationship (restoring original source and target destinations).

Under this scenario, a fraction of the data to be recovered is sent across the WAN link. This could speed up a remote restore significantly. However, there are some downsides to this type of scenario:

- If step 1 is not followed correctly, any changes not fully replicated are lost.
- During steps 2 and 3, any data that is written to the original DR Series system source container may be lost.
- During step 4, if the data is not fully replicated back before the switch is made, it may be lost.

Alternatively, you could still support this type of effort by completing the following:

1. Create a new container on the target DR Series system.
2. Set up replication from this container back to the source DR Series system container.
3. Set up a new disk storage unit in the DMA and make sure that the DMA is aware of any new images.
4. Import the old images back into the DMA from the target DR Series system (the original source location).
5. Use a new disk storage unit in the DMA, and then restore the data back to the original client.

Reverse replication: alternate method

For an alternate method of reverse replication, complete the following steps:

1. Create a new container on the target DR Series system.
2. Set up replication from this container back to the source DR Series system container.
3. Set up a new disk storage unit in your Data Management Application (DMA) and make sure that the DMA is aware of any new images.

4. Import the old images back into the DMA from the target DR Series system (the original source location).
5. Use a new disk storage unit in the DMA, and then restore the data back to the original client.

Rapid Air Gap

Rapid Air Gap is a replication feature in which a secondary target DR Series system in a three-way replication scenario exists in an isolated private network. The Secondary target is available for replication for only a specific period of time (when the Air gap closes). Only during this period of time will replication between the Primary and Secondary target DR systems occur. Only a subset of services are available on the secondary target so that it is secured from the external environment. The primary target DR Series system acts as the target as well as the source for the Secondary target. The primary target will be in two different networks to cater both Primary DR and Secondary target networks. The Secondary target also has a (user-editable) retention period that retains the data on it even after the respective data is deleted on the Secondary source. Commands for using this feature are available in the DR Series system command line interface (CLI). For information about setting up and using this feature, see the *DR Rapid Air Gap Setup Guide* and the *DR Series System Command Line Reference Guide*.

Secure erase

Secure Erase is a process for securely deleting data that follows standards developed by the Defense Security Service (DSS). These standards were developed to solve the problem of secure and permanent removal of data, and this capability is now used by many commercial enterprises. These standards require multiple passes to erase data. In the process of erasing, the system overwrites data with zeros, a random pattern of data, or ones (1s) to make the original contents unreadable.

In the DR Series system, two modes of secure erase have been adopted: 3-pass and 7-pass standards.

- 3 passes (US DoD 5220.22-M (C))
 - Pass 1: Writing zeros.
 - Pass 2: Writing pseudo random pattern.
 - Pass 3: Writing ones.
- 7 passes (US DoD 5220.22-M (ECE))
 - Pass 1-3: Writing pattern according to the 3-pass mode.
 - Pass 4: Writing a pseudo random pattern.
 - Pass 5-7: Writing pattern according to the 3-pass mode.

The Secure Erase process can be completed at the system level by using the DR Series system CLI. This feature securely erases all data with a 3-pass or 7-pass mode and runs during the system initialization. This method erases all of the stored data while keeping the system in tact. See the *DR Series System Command Line Reference Guide* for more information.

i | **NOTE:** This capability is not supported on the DR2000v.

Supported file system protocols

The DR Series system supports the following file system protocols. The Rapid Data Access (RDA) protocols below provide a logical disk interface that can be used with network storage devices to store data and support

data storage operation.

- Network File System (NFS)
- Common Internet File System (CIFS)
- DR Rapid
 - Rapid NFS
 - Rapid CIFS
 - RDA with OpenStorage Technology (OST)
 - RDA with NetVault Backup
 - RDA with vRanger

NFS

The Network File System (NFS) is a file system protocol that is designated to be a file server standard, and its protocol uses the Remote Procedure Call (RPC) method of communication between computers. Clients can access files via the network similar to the way that local storage is accessed.

NFS is a client-server application in which a client can view, store, and update files on a remote system just like they are working on a local system. System or Network Administrators can mount all or a portion of a file system, and the file system (or portion) that is mounted can be accessed using the privileges assigned to each file.

i **NOTE:** If you want to do a mount on AIX, you must set the `nfs_use_reserved_ports` and `portcheck` parameters first. The parameters cannot be set to 0. For example: `root@aixhost1 / # nfso -po portcheck=1 root@aixhost1 / # nfso -po nfs_use_reserved_ports=1`

CIFS

The Common Internet File System (CIFS) remote file access protocol is supported by the DR Series system, and is also known as a Server Message Block (SMB). SMB occurs more commonly than the Network File System (NFS) protocol on systems that run the Microsoft Windows operating system. CIFS allows programs to request files or services on remote computers.

CIFS also uses the client-server programming model, whereby the client requests access to a file or passes a message to a program running on the server. Servers review all requested actions and return a response. CIFS is a public (or open) variation of the SMB that was originally developed and used by Microsoft.

i **NOTE:** The DR Series system currently supports version 2.0 of the Server Message Block (SMB).

i **NOTE:** For complete details on CIFS feature restrictions, see the *DR Series System Interoperability Guide*, at support.quest.com/dr-series.

CIFS ACL support

The DR Series system software supports the use of access control lists (ACLs) for CIFS and share-level permissions. By definition, an ACL is simply a list of permissions that can be associated with any network resource.

Each ACL can contain access control entries (ACEs) that define or describe the permissions for an individual user or a group of users. An ACL can consist of zero (meaning that all users have access) or a number of ACEs that define specific permissions on a per-user or per-group basis.

i **NOTE:** If an ACE list is empty (meaning that it contains zero entries), this means that all access requests will be granted.

An ACL describes the entities that are allowed to access a specific resource. ACLs are a built-in access control mechanism in the Windows operating systems.

i **NOTE:** The DR Series system supports setting up share-level permissions for a CIFS share using a Microsoft Windows administrative tool. Share-level permissions let you control access to shares. For more information, see [Configuring Share-Level Security](#).

Access control list support in containers

All new containers apply a default Access Control List (ACL) at the root of the container. This default ACL is the same as that which would be created by a Microsoft Windows 2003 Server. Therefore, these new containers with the default ACL support the following permission types:

i **NOTE:** Any user that is part of BUILTIN\Administrators can edit ACLs on CIFS shares. The local DR Series system administrator is included in the BUILTIN\Administrators group. To add additional domain groups to the BUILTIN\Administrators group, you can use the Computer Manager tool on a Windows client to connect to the DR Series system as Domain administrator and add any groups you want. This capability allows users other than the Domain administrator to modify an ACL as needed.

- **BUILTIN\Administrators:**
 - **Allows:** Full access, object inherit, and container inherit.
 - **Applies to:** This folder, subfolders, and files.
- **CREATOR OWNER:**
 - **Allows:** Full access, inherit only, object inherit, and container inherit.
 - **Applies to:** Subfolders and files only.
- **EVERYONE:**
 - **Allows:** Traverse folders, execute files, list folders, read data, read attributes, and read extended attributes.
 - **Applies to:** This folder only.
- **NT AUTHORITY\SYSTEM:**
 - **Allows:** Full access, object inherit, and container inherit.
 - **Applies to:** This folder, subfolders, and files.
- **BUILTIN\Users:**
 - **Allows:** Create folders and append data, inherit-only, and container inherit.
 - **Applies to:** This folder, subfolders, and files.
- **BUILTIN\Users:**
 - **Allows:** Read and execute, and container inherit.
 - **Applies to:** This folder, subfolders, and files.

- **BUILTINUsers:**

- **Allows:** Create files and write data, object inherit, and container inherit.
- **Applies to:** Subfolders only.

i **NOTE:** If these permissions are unsuitable for your needs, you can modify the default ACL to suit your own requirement using the Windows ACL Editor (for example, using Properties → Security from Windows Explorer).

i **NOTE:** The system does not understand the Owner Rights permission and sets the owner of new files/folders created by the Domain Administrators as **DOMAdministrator** rather than as **BUILTINAdministrators**.

Unix permissions guidelines

For a user to create, delete, or rename a file or a directory requires Write access to the parent directory that contains these files. Only the owner of a file (or the root user) can change permissions.

Permissions are based on the user IDs (UIDs) for the file Owner and group IDs (GIDs) for the primary group. Files have owner IDs and group owner IDs. To enable Unix access, the DR Series system supports three levels of users:

- Owner (of the file)
- Group (group in which the owner belongs)
- Other (other users with an account on the system)

Each of these three user types support the following access permissions:

- Read (read access that allows user to read files)
- Write (write access that allows user to create or write to a file)
- Execute (access that allows user to execute files or traverse directories in the filesystem)

i **NOTE:** A root user has all levels of permission access, and a user can be a member of a single group or of multiple groups (up to 32 groups are allowed in Unix).

Windows permissions guidelines

To enable Windows access, the DR Series system supports access control lists (ACLs) that contain zero or more access control entries (ACEs), and an empty ACE list grants all access requests. The Windows New Technology File System (NTFS) uses ACLs as part of the security descriptor (SD) process, which requires permissions to access such filesystem objects as files and directories. ACLs support two levels of users:

- Owners
- Groups

Both Owners and Groups have Security IDs (SIDs) that define and identify an object owner or the group owning an object. ACEs in an ACL consist of a SID, a specific permission that either allows or denies access and also defines which of the following inheritance settings apply:

- IO—inherit-only: not used for access checking.
- OI—object inherit: new files get this ACE added.

- CI—container inherit: new directories get this ACE added.

Windows NTFS ACLs include the following read, write, append, execute, and delete permissions that allow users to:

- Synchronize access
- Read data or list the directory
- Write data or add a file
- Append data or add a folder
- Read Extended Attributes (EAs)
- Write EAs
- Execute file or traverse folders
- Delete child or delete folders
- Delete a file

The Owner user type has two default permissions:

- Write discretionary ACL
- Read control

Rapid NFS and Rapid CIFS

Rapid NFS and Rapid CIFS enable write operation acceleration on clients that use DR replication and NFS or CIFS file system protocols. Similar to OST and RDS, these accelerators allow for better coordination and integration between DR Series system backup, restore, and optimized deduplication operations with Data Management Applications (DMAs) such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of qualified DMAs, see the *DR Series System Interoperability Guide*.

Rapid NFS is a new client file system type that ensures that only unique data is written to the DR Series system. It uses user space components and file system in user space (FUSE) to accomplish this. Metadata operations such as file creates and permission changes go through the standard NFS protocol, whereas write operations go through RDNFS.

Rapid CIFS is a Windows-certified filter driver that also ensures that only unique data is written to the DR Series system.

All chunking and hash computations are done at the media or client server level.

Rapid NFS and Rapid CIFS require you to install a plug-in on the client or media server, depending on your DMA and configuration. For details, see the Configuring and Using Rapid NFS and Rapid CIFS chapter.

DR Rapid for the DR Series system

DR Rapid is developed by Quest and provides a logical disk interface for use with network storage devices. DR Rapid allows for better coordination and integration between DR Series system backup, restore, and optimized duplication operations with backup applications, such as NetVault Backup.

The DR Series system and backup application integration is done by using DR Rapid plugins developed by Quest. Plugins allow backup application control over backup image creation, deletion, and duplication. They also allow deduplication and compression operations to happen on the client-side so that network traffic can be reduced.

DR Rapid allows the supported backup applications to communicate directly with the DR Series system and determine whether a specific chunk of data already exists on the system. If the data already exists, only the pointers need to be updated on the DR Series system, and the duplicate chunk of data does not need to be transferred to the system. This process provides two benefits: it improves the overall backup speed, and also reduces the network load.

RDA with OST for the DR Series system

OpenStorage Technology (OST), by Veritas, provides a logical disk interface for use with network storage devices. The DR Series system appliance can use OST via DR Rapid plug-in software to integrate its data storage operations with NetBackup and Backup Exec. OST is part of DR Rapid at Quest.

RDA with OST allows for better coordination and tighter integration between DR Series system backup, restore, and optimized duplication operations and data management applications. For a list of the supported applications, see the *DR Series System Interoperability Guide*.

Integration is done via a RDA with OST plug-in developed for the DR Series system, through which data management applications can control when the backup images are created, duplicated, and deleted. The major benefit of RDA with OST is that it allows the deduplication operations to happen on the client side so that network traffic can be reduced.

The RDA with OST plug-in allows data management applications to take full advantage of such DR Series system features as data deduplication, replication, and energy efficiency. DR Series systems can access the OpenStorage API code through the plug-in, which can be installed on the media server platform choice you make (Windows or Linux). The OST protocol allows the supported backup applications to communicate directly with the DR Series system and determine whether a specific chunk of data already exists on the system. This process means that if the data already exists, only the pointers need to be updated on the DR Series system, and the duplicate chunk of data does not need to be transferred to the system. This process provides two benefits: it improves the overall backup speed, and also reduces the network load.

When RDA with OST is used with the DR Series system, it offers the following benefits:

- OST protocol provides faster and improved data transfers:
 - Focused on backups with minimal overhead
 - Accommodates larger data transfer sizes
 - Provides throughput that is significantly better than CIFS or NFS
- RDA with OST and DMA integration:
 - OpenStorage API enables the DMA-to-media server software communications
 - DR Series system storage capabilities can be used without extensive changes to DMAs
 - Backup and replication operations are simplified by using built-in DMA policies
- DR Series system and RDA with OST:
 - Control channel uses TCP port 10011
 - Data channel uses TCP port 11000
 - Optimized write operations enable client-side deduplication

- Replication operations between DR Series systems:
 - No configuration required on source or target DR Series systems
 - Replication is file-based, not container-based
 - Triggered by DMA optimized duplication operation
 - DR Series system transfers the data file (not the media server)
 - After duplication, DR Series system notifies DMA to update its catalog (acknowledging the second backup)
 - Supports different retention policies between source and replica

Software components and operational guidelines

To better coordinate and integrate OpenStorage Technology (OST) with the DR Series system data storage operations, the following guidelines list the required components and supported operations. For details on the supported operating systems and data management application (DMA) versions, see the *DR Series System Interoperability Guide*.

The DR Series system licensing is all-inclusive, so that no additional licensing is required to use OST or the optimized duplication capability. The OST plug-in that gets installed on a supported Linux or Windows media server platform is a free download. However, Veritas NetBackup requires that you purchase an OpenStorage Disk Option license. Similarly, Veritas Backup Exec requires that you purchase the Deduplication Option to enable the OST feature.

- OST Media Server Component:
 - An OST server component resides on the DR Series system
 - For Linux media server installations, use the Linux OST plug-in and the Red Hat Package Manager (RPM) installer
 - For Windows media server installations, use the Windows OST plug-in and the Microsoft (MSI) installer
- Windows-based OST plug-in
- Linux-based 64-bit OST plug-in
- Supported OpenStorage (OST) protocol:
 - Version 9
 - Version 10
- Supported Veritas DMAs
 - NetBackup
 - Backup Exec
- Supported OST operations
 - Backup (Passthrough writes and Optimized writes)
 - Restore
 - Replication

Supported virtual tape library access protocols

The DR Series system supports the following virtual tape library (VTL) tape access protocols.

- Network Data Management Protocol (NDMP)
- Internet Small Computer System Interface (iSCSI)
- Fibre Channel (FC)

NDMP

The Network Data Management protocol (NDMP) is used to control data backup and recovery between primary and secondary storage in a network environment. For example, a NAS server (Filer) can talk to a tape drive for the purposes of a backup.

You can use the protocol with a centralized data management application (DMA) to back up data on file servers running on different platforms to tape drives or tape libraries located elsewhere within the network. The protocol separates the data path from the control path and minimizes demands on network resources. With NDMP, a network file server can communicate directly to a network-attached tape drive or virtual tape library (VTL) for backup or recovery.

The DR Series system VTL container type is designed to work seamlessly with the NDMP protocol.

iSCSI

iSCSI or **Internet Small Computer System Interface** is an Internet Protocol (IP)-based storage networking standard for storage subsystems. It is a carrier protocol for SCSI. SCSI commands are sent over IP networks by using iSCSI. It also facilitates data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over LANs or WANs.

In iSCSI, clients are called *initiators* and SCSI storage devices are *targets*. The protocol allows an *initiator* to send SCSI commands (*CDBs*) to the *targets* on remote servers. It is a storage area network (SAN) protocol, allowing organizations to consolidate storage into data center storage arrays while providing hosts (such as database and web servers) with the illusion of locally attached disks. Unlike traditional Fibre Channel, which requires different cabling, iSCSI can be run over long distances using existing network infrastructure.

iSCSI is a low-cost alternative to Fibre Channel, which requires dedicated infrastructure except in FCoE (Fibre Channel over Ethernet). Note that the performance of an iSCSI SAN deployment can be degraded if not operated on a dedicated network or subnet.

The VTL container type is designed to work seamlessly with the iSCSI protocol. For details, see the topic, [Creating Storage Containers](#).

Fibre channel

Fibre Channel (FC) is a high-speed network technology primarily used to connect computer data storage to servers in storage area networks (SAN) in enterprise storage. Fibre Channel networks are known as a Fabric because they operate in unison as one big switch. Fibre Channel mainly runs on optical fiber cables within and between data centers. Virtual tape libraries (VTLs) can ingest data over a Fibre Channel interface, which enables seamless integration with many existing backup infrastructures and processes.

The DR Series system VTL container type is designed to work seamlessly with the FC interface. With FC, the DR Series system can direct attach to NAS filers or Fibre Channel switches and supports SAN devices.

A FC VTL container on a DR Series system supports multiple initiators, making it possible for the VTL to be shared across multiple clients of a Data Management Application (DMA).

i | **NOTE:** VTL access via FC is only available on the DR4300 and DR6300.

Understanding the DR Series system hardware and data operations

Data is stored and resides on the DR Series system hardware appliances (two-rack unit (RU) appliances), which have DR Series system software pre-installed.

The DR Series system hardware consists of a total of 14 drives. Two of these drives are 2.5-inch drives that are configured as a Redundant Array of Independent Disks (RAID) 1 on the RAID Controller, and this is considered to be volume 1. On the DR4000 system, these drives are internal; while in the DR4100, DR6000, DR4300e core and standard, DR4300, and DR6300 systems, these drives are accessible from the rear of the appliance. The data that is being backed up is stored on the 12 virtual disks that reside on the DR Series system. The DR Series system also supports additional storage in the form of external expansion shelf enclosures (see the *DR Series Expansion Shelf* section in this topic). The hot-swappable data drives that are attached to the RAID controller are configured as:

- 11 drives that operate as RAID 6, which act as virtual-disks for data storage (drives 1–11).
- The remaining drive (drive 0) acts as the global hot-spare drive for RAID 6 for the system for the DR4000, DR4100 and DR6000 and a dedicated hot spare for the DR4300e core and standard, DR4300, and DR6300.

The DR Series system supports RAID 6, which allows the appliance to continue read and write requests to the RAID array virtual disks even in the event of up to two concurrent disk failures, providing protection to your mission-critical data. In this way, the system design supports double-data drive failure survivability.

If the system detects that one of the 11 virtual drives has failed, then the dedicated hot spare (drive slot 0) becomes an active member of the RAID group. Data is then automatically copied to the hot spare as it acts as the replacement for the failed drive. The dedicated hot spare remains inactive until it is called upon to replace a failed drive. This scenario is usually encountered when a faulty data drive is replaced. The hot spare can act as replacement for both internal mirrored drives and the RAID 6 drive arrays.

Figure 1: DR Series System Drive Slot Locations



Drive 0 (top)	Drive 3 (top)	Drive 6 (top)	Drive 9 (top)
Drive 1 (middle)	Drive 4 (middle)	Drive 7 (middle)	Drive 10 (middle)
Drive 2 (bottom)	Drive 5 (bottom)	Drive 8 (bottom)	Drive 11 (bottom)

DR Series expansion shelf

The DR Series hardware system appliance supports the installation and connection of Dell PowerVault MD1200 (for DR4000, DR4100, and DR6000) and Dell PowerVault MD1400 (for DR4300e core and standard, DR4300, and DR6300 systems) data storage expansion shelf enclosures. Each expansion shelf contains 12 physical disks in an enclosure, which provides additional data storage capacity for the basic DR Series system. The supported data storage expansion shelves can be added in a variety of capacities based on your DR Series system version; for details, see the *DR Series System Interoperability Guide*.

The physical disks in each expansion shelf are required to be Dell-certified Serial Attached SCSI (SAS) drives, and the physical drives in the expansion shelf uses slots 1–11 configured as RAID 6, with slot 0 being a global hot spare (GHS). When being configured, the first expansion shelf is identified as Enclosure 1 (in the case where two enclosures are added, these would be Enclosure 1 and Enclosure 2). Adding an expansion shelf to support the DR Series system requires a license.

i **NOTE:** The 300 Gigabyte (GB) drive capacity (2.7 TB) version of the DR Series system does not support the addition of expansion shelf enclosures.

i **NOTE:** If you are running a DR Series system with an installed release of system software prior to 2.1, and you intend to upgrade to release 3.x or 4.x system software and add an external expansion shelf (or shelves), Quest recommends that you observe the following best practice sequence of operations to avoid any issues:

- Upgrade the DR Series system with the release 3.x system software
- Power off the DR Series system.
- Connect the external expansion shelf (or shelves) with cabling to the DR Series system.
- Power on the external expansion shelf (or shelves).
- Power on the DR Series system.

i **NOTE:** If you install an expansion shelf enclosure to support a DR Series system, each shelf must use physical disks that have a capacity equal to or greater than each DR Series system internal drive slot capacity (0–11) that they are supporting.

Figure 2: DR Series System Expansion Shelf (MD1200) Drive Slot Locations



Drive 0 (top)	Drive 3 (top)	Drive 6 (top)	Drive 9 (top)
Drive 1 (middle)	Drive 4 (middle)	Drive 7 (middle)	Drive 10 (middle)
Drive 2 (bottom)	Drive 5 (bottom)	Drive 8 (bottom)	Drive 11 (bottom)

Understanding the process for adding a DR Series expansion shelf

The process for adding an expansion shelf requires the following:

- Physically adding or installing the expansion shelf.
- Cabling the expansion shelf to the DR Series system (for more information, see the topic, “Understanding DR Series system hardware expansion shelf cabling”).
- Installing the license for the expansion shelf.
- Using the DR Series system GUI (System Configuration page) to add or detect the expansion shelf (for more information, see the topic, “Viewing and adding a DR Series system expansion shelf”).

Understanding DR Series system hardware expansion shelf cabling

The DR Series system hardware appliance is capable of supporting additional storage capacity by connecting Dell PowerVault MD1200 (DR4000, DR4100, DR6000) or Dell PowerVault MD1400 (DR4300e core and standard, DR4300, DR6300) data storage expansion shelf enclosures. The expansion shelf enclosure contains 12 physical disks that provide additional data storage capacity for a basic DR Series system. For the expansion unit limits and supported capacities, see the *DR Series System Interoperability Guide*.

As an example, this section and the following figures display the recommended method for cabling between the DR Series system’s PERC controller card to the appropriate connectors on the rear of a Dell PowerVault MD1200 expansion shelf enclosure.

Make sure that the Dell PowerVault MD1200 front panel selector switch is set to its Unified mode (with the switch set to its “up” position, indicated by a single Volume icon). The first figure below shows the SAS In ports on the Enclosure Management Module (EMM) on the rear of the Dell MD1200. The next figure shows the recommended redundant path cabling configuration, which includes cable connections from both PERC H800 connectors on the DR4000 system (or the PERC H810 on a DR4100/DR6000 system) to the two SAS In ports on the EMM rear chassis of the Dell PowerVault MD1200.

If you plan on installing multiple expansion shelf enclosures, then the two SAS In ports on the rear chassis of the EMM on the additional enclosure are daisy-chained to the two SAS Out ports on the EMM rear chassis on the first enclosure. This is considered a redundant mode connection via the SAS In/Out connectors on the enclosures with the DR Series system appliance.

If you install multiple enclosures and cable them as described here, make sure to set the enclosure mode switch on the MD1200 front chassis to the top (unified mode) position. For more information, see the *Dell PowerVault MD1200 and MD1220 Storage Enclosures Hardware Owner's Manual* or the *Dell Storage MD1400 Enclosures Hardware Owner's Manual* at dell.com/support/manuals.

Figure 3: Dell PowerVault MD1200 Rear Chassis




Figure 4: Unified Mode Daisy-Chained Redundant Path Dell PowerVault MD1200 Enclosures





Figure 5: SAS Port and Cable Connections (Dell PowerVault MD1200 EMM)



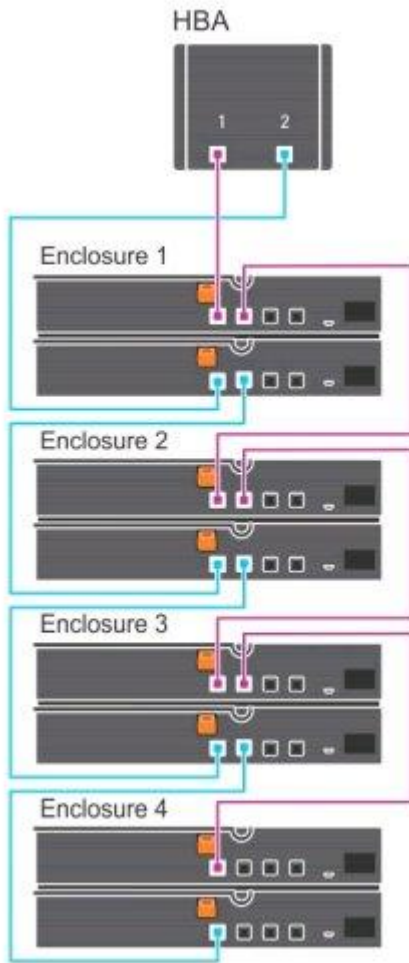
1. SAS cable
2. pull-tab

For the DR4300e core and standard, DR4300, and DR6300, the Dell Storage MD1400 expansion shelf enclosure is used for data storage expansion. The MD1400 has 4 ports in each controller or EMM; and, it is recommended to use ports 1 and 2 on the MD1400.

Figure 6: Dell Storage MD1400 Rear Chassis



Figure 7: Daisy-Chained Redundant Path Dell Storage MD1400 Enclosures



Supported software and hardware

For a complete list of the latest supported software and hardware for the DR Series system, refer to the *DR Series System Interoperability Guide*. You can download this guide by visiting support.quest.com/dr-series, selecting your specific DR model and then navigating to Technical Documentation.

The *DR Series System Interoperability Guide* includes the following supported hardware and software categories:

- Hardware
 - BIOS
 - RAID controllers
 - Hard drives (internal)
 - Hard drives (external)
 - Expansion unit limits
 - Fibre Channel controllers
 - USB flash drives
 - Network interface controllers
 - iDRAC Enterprise
 - Marvell WAM controller
- Software
 - Operating systems
 - Supported virtual environments for the DR2000v system
 - Supported backup software
 - Network file protocols and backup client operating systems
 - Supported iSCSI initiators
 - Supported NAS filers for NDMP and for FC
 - Supported web browsers
 - Supported system limits
 - Supported OST software and components
 - Supported RDS software and components
 - Supported Rapid NFS and Rapid CIFS software and components

Terminal emulation applications

To access the DR Series system command line interface (CLI), the following terminal emulation applications can be used:

- FoxTerm
- Win32 console
- PuTTY
- Tera Term Pro

i **NOTE:** The listed terminal emulation applications are not the only ones that work with the DR Series system. This list is only intended to provide examples of terminal emulation applications that can be used.

Setting up the DR Series system

You can interact with the DR Series system hardware using one of two supported methods: the system graphical user interface (GUI) accessed in a web browser or the command line interface (CLI) accessed by using a terminal emulator application (for example, PuTTY). Before you can interact with your system, you must first, however, ensure that the DR Series system is properly set up.

i **NOTE:** The topics in this section apply to DR Series hardware systems. For information about setting up the virtual DR Series system, DR2000v, see the *DR2000v Deployment Guide* and the *DR Series System Interoperability Guide*. For more information on the DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Interacting with the DR Series system

You can interact with the DR Series system by using the system graphical user interface (GUI) in your web browser.

i **NOTE:** A second method for interacting with the DR Series system is by using the command-line interface (CLI) via a terminal emulator application (for example, PuTTY).

In the system GUI, you can configure your system as well as create and manage containers, which store your backup and deduplicated data. A data container is a shared file system that is imported using a client, and is accessible via file system or tape access protocols. For details, see [Supported File System Protocols](#). The system GUI also provides real-time summary information for monitoring the status of the data capacity, storage savings, and the throughput of your data containers.

i **NOTE:** Before you can start to use the system GUI, you must initialize the system by using the CLI. See the topic, “Initializing the DR Series System,” for more information.

Networking prerequisites for the DR Series system

Before you can start using the DR Series system, ensure that you have satisfied the following networking prerequisites:

- **Network:** An active network is available using Ethernet cables and connections.

i **NOTE:** If your DR Series system is equipped with a 1-GbE NIC, Quest recommends using CAT6 (or CAT6a) copper cabling. If your DR Series system is equipped with a 10-GbE NIC, Quest recommends using CAT6a copper cabling.

i **NOTE:** If your DR Series system is equipped with a 10-GbE enhanced small form-factor pluggable (SFP+) NIC, you must use Quest-supported SFP+ LC fiber-optic transceivers or twin-axial cabling.

- **IP Addresses:** You need to ensure you have IP addresses that you will use for the DR Series system. The DR Series system ships with a default IP address and subnet mask address, which should only be used for an initial system configuration.

i **NOTE:** You need to have an IP address available to replace the default IP address if you choose the static mode of IP addressing, or select to use the DHCP mode of IP addressing.

For the initial configuration, you need the following addresses:

- An IP address for the system
- A subnet mask address
- A default gateway address
- A DNS suffix address
- A primary DNS server IP address
- (Optional) A secondary DNS server IP address
- **NIC Connections:** By default, the DR Series system will configure NIC interfaces together as a bonded team (and only one IP address is needed because the bonded NICs assume the primary interface address). NIC connection bonding can use either of these two configurations:
 - Adaptive load balancing (ALB), which is the default setting, does not require any special network switch support. Ensure that the data source system resides on the same subnet as the DR Series system. For more information, see [Configuring Networking Settings](#).
 - 802.3ad or dynamic link aggregation (using the IEEE 802.3ad standard). 802.3ad requires special switch configuration before using the system (contact your network administrator for an 802.3ad configuration).
- i** **NOTE:** To configure a 10-GbE NIC or 10-GbE SFP+ bonded configuration, connect only the 10-GbE/10-GbE SFP+ NICs. You can use the Advanced Networking feature in the command line interface to modify the default factory configuration.
- **DNS:** you need a DNS domain available, and you need to know the primary DNS server IP address (and a secondary DNS server IP address, if you choose to configure one).
- **Replication ports:** the replication service in the DR Series system requires that enabled fixed ports be configured to support replication operations that are to be performed across firewalls (TCP ports 9904, 9911, 9915, and 9916).

For more information about replication ports, see [Managing Replication Operations](#), and for more information about system ports, see [Supported Ports in a DR Series System](#).

i **NOTE:** For the latest details about supported hardware and software for the DR Series system, see the *DR Series System Interoperability Guide* at support.quest.com/dr-series.

Supported connections for initializing the DR Series system

There are two supported methods for connecting to the DR Series system for logging on and initial system configuration via the DR Series system CLI:

- **Local console connection:** a local access connection made between a local workstation and the DR Series system (with one connection made to a USB keyboard port on the DR Series front/rear chassis, and a second connection made to the VGA monitor port on the DR Series system rear chassis. (See Figure 3 for locations in the DR Series System Rear Chassis Port Locations in the [Local Console Connection](#).)
- **iDRAC connection:** a remote access connection made between an integrated Dell Remote Access Controller (iDRAC) and the dedicated management port on the DR Series system rear chassis. (See Figure 3 for locations in the DR Series System Rear Chassis Port Locations in the [Local Console Connection](#).)

Default IP and subnet mask addresses

This topic lists the following default address values that can be used for initialization of a DR Series system:

- IP address—10.77.88.99
- Subnet mask address—255.0.0.0

There are two key factors related to default address values and initializing a DR Series system:

- Using the local console
- Reserving MAC addresses using DHCP

If the network where the system will reside does not have or does not support DHCP, then the DR Series system can use the default IP (10.77.88.99) and subnet mask (255.0.0.0) addresses provided for initialization. If the network where the system will reside does not have or support reserving an IP address for the MAC address of the NICs in the DHCP server, then DHCP assigns an arbitrary IP address that is unknown (and which is unusable by you) during initialization.

As a result, if your network does not support DHCP or if you cannot reserve an IP address for the specific MAC addresses of the DHCP network interface cards (NICs), then Quest recommends that you use the local console connection method and the Initial System Configuration Wizard.

i NOTE: After successfully initializing and configuring your system, you can modify the IP address to use either a static IP address or use dynamic IP addressing (DHCP), and modify the subnet mask address to be one that is supported by your network.

i NOTE: If you have not run the Initial System Configuration Wizard on one (or more) DR Series system(s) being installed into the same network, there is a potential that the system (or systems) may come up with the same default IP address (10.77.88.99). The default IP address is not user-configurable and it can potentially result in becoming a duplicate IP address in the case of multiple systems.

Initialization issues could include when a network has had a network power outage, the DHCP server in the network is misconfigured, or if the Initial System Configuration Wizard has never been run.

If your network does not accept the default subnet mask address (255.0.0.0), you can establish a connection between the DR Series system and a laptop workstation. In this case, make sure that you connect using SSH, and use the default IP address to run the Initial System Configuration Wizard.

If you are using a known static IP address, you can skip the Initial System Configuration Wizard, and directly configure the DR Series system using the system graphical user interface (GUI). To configure the DR Series system, select System Configuration > Networking, and configure the network settings as desired. For more information, see [Configuring Networking Settings](#).

i | **NOTE:** For details about logging on and using the Initial System Configuration Wizard, see [Configuring Networking Settings](#).

Local console connection

To configure a local console connection, you must make the following two rear chassis cables connections:


- VGA port and your video monitor
- USB port and your keyboard





DR4000

To make local console cable connections for the DR Series system appliance, complete the following steps.

1. Locate the VGA monitor port and the USB ports on the back of your system. See Figure 3 for the VGA and USB port locations and complete steps 1 to 4. For the DR4100/DR6000 system, skip to step 5.
2. Connect the video monitor to the VGA port on the back of your system (see item 1 in the DR4000 System Rear Chassis Port Locations table).
3. Connect the USB keyboard to one of the two USB ports on the back of your system (see item 3 in DR4000 System Rear Chassis Port Locations table).
4. You are now ready to perform initialization using the DR Series system CLI login process. For more information, see [Logging in and Initializing the DR Series System](#).

Figure 8: DR4000 System Rear Chassis Port Locations



Item	Indicator, Button, or Connector	Icon	Description
1	Video connector		Connects a VGA display to the system.
2	iDRAC6 Enterprise port		Dedicated management port for the iDRAC6 Enterprise card.
3	USB connectors (2)		Connects USB devices to the system. The ports are USB 2.0-compliant.
4	Ethernet connectors (2)		Embedded 10/100/1000 NIC connectors.
5	Ethernet Connectors (2) on expansion card		1-GbE/10-GbE/10-GbE SFP+ Ethernet Port

DR4100/DR6000


To make local console cable connections for the DR4100 or DR6000 system appliance, complete the following steps:

i NOTE: For the 1-GbE ports, these are two internal LAN on Motherboard (LOM) ports referenced in item 4 above that reside on the motherboard, and two ports on an expansion card referenced in item 5 above. If the system is using the two 10-GbE ports, these reside on an expansion card referenced in item 5 above.

1. Locate the VGA monitor port and the USB ports on the back of your system. See Figure 3 for the VGA and USB port locations and complete steps 5 to 8.
2. Connect the video monitor to the VGA port on the back of your system (see item 2 in the DR4100/DR6000 System Rear Chassis Port Locations table).
3. Connect the USB keyboard to one of the two USB ports on the back of your system (see item 3 in the DR4100/DR6000 System Rear Chassis Port Locations table).

You are now ready to perform initialization using the DR Series system CLI login process. For more information, see [Logging in and Initializing the DR Series System](#).

Figure 9: DR4100/DR6000 System Rear Chassis Port Locations



Item	Indicator, Button, or Connector	Icon	Description
1	iDRAC7 Enterprise port		Dedicated management port for the iDRAC7 Enterprise card (port is available only if an iDRAC7 Enterprise license is installed on your system).
2	Video connector		Connects a VGA display to the system.
3	USB connectors (2)		Connects USB devices to the system. The ports are USB 2.0-compliant.
4	Ethernet connectors (4)		Four integrated 10/100/1000 NIC connectors, or four integrated connectors that include: <ul style="list-style-type: none"> Two 10/100/1000 Mbps NIC connectors Two 100 Mbps/1 Gbps/10 Gbps SFP+/10-GbE T connectors
5	PCIe expansion card slots (3)		Connect up to three full-height PCI Express expansion cards
6	Hard drives (2)		Provides two hot-swappable 2.5-inch hard drives


NOTE: The DR4100/DR6000 system supports up to six 1-GbE ports or up to two 10-GbE ports. For the 1-GbE ports, these are four internal LAN on Motherboard (LOM) ports referenced in item 4 above that reside on the network daughter card (NDC), and two additional ports on a PCI Express expansion card referenced in item 5 above. If the system is using the two 10-GbE ports, these ports reside on the NDC.






DR4300e/DR4300/DR6300



To make local console cable connections for the DR4300e, DR4300, or DR6300 system appliance, complete the following steps:

1. Locate the VGA monitor port and the USB ports on the back of your system. See Figure 3 for the VGA and USB port locations and then complete steps 10 to 12.
2. Connect the video monitor to the VGA port on the back of your system (see item 6 in the DR4300e/DR4300/DR6300 System Rear Chassis Port Locations table).
3. Connect the USB keyboard to one of the two USB ports on the back of your system (see item 7 in the DR4300e/DR4300/DR6300 System Rear Chassis Port Locations table).
4. You are now ready to perform initialization using the DR Series system CLI login process. For more information, see [Logging in and Initializing the DR Series System](#).

Figure 10: DR4300e/DR4300/DR6300 System Rear Chassis Port Locations



Item	Indicator, Button, or Connector	Icon	Description
1	System identification button		The identification buttons on the front and back panels can be used to locate a particular system within a rack. When one of these buttons is pressed, the system status indicator on the back flashes until one of the buttons is pressed again. Press the button to toggle the system ID on and off. If the system stops responding during POST, press and hold the system ID button for more than five seconds to enter BIOS progress mode. To reset iDRAC (if not disabled in F2 iDRAC setup) press and hold the button for more than 15 seconds.
2	System identification connector		Connects the optional system status indicator assembly through the optional cable management arm.
3	iDRAC8 Enterprise port		Dedicated management port for the iDRAC7 Enterprise card (The port is available for use only if the iDRAC8 Enterprise license is installed on your system.)
4	PCIe expansion card slots half height (3)		Allows you to connect up to three half-height PCI Express expansion cards.
5	Serial connector		Allows you to connect a serial device to the system.
6	Video connector		Connects a VGA display to the system.
7	USB connectors (2)		Connects USB devices to the system. The ports are USB 3.0-compliant.
8	PCIe expansion card slots full height (3) (Dell PowerEdge)		Allows you to connect up to three full-height PCI Express expansion cards.

Item	Indicator, Button, or Connector	Icon	Description
	R730xd)		
9	Ethernet connectors (4)		Four integrated 10/100/1000 Mbps Network Interface Card (NIC) connectors or four integrated connectors that include: <ul style="list-style-type: none"> Two 10/100/1000 Mbps NIC connectors Two 100 Mbps/1 Gbps/10 Gbps SFP+/10 GbE T connectors
10	Power supply (PSU1)		AC 495 W, 750 W, or 1100 W or DC 750 W or 1100 W
11	Power supply (PSU2)		AC 495 W, 750 W, or 1100 W or DC 750 W or 1100 W
12	vFlash media card slot		Allows you to insert a vFlash media card.
13	Hard drives (2) (back)		Allows up to two hot-swappable 2.5 inch hard drives.

iDRAC connection

The iDRAC connection requires a network connection between the integrated Dell Remote Access Control (iDRAC) management port on the DR Series system and another computer running the iDRAC remote console session in a supported browser. The iDRAC provides remote console redirection, power control, and the out-of-band (OOB) system management functions for the DR Series system. iDRAC connections are configured using console redirection and the iDRAC6/7 web interface. The login values you can use for making iDRAC connections are:

- Default username: root
- Default password: calvin
- Default static IP address: 192.168.0.120

For information about configuring the iDRAC, see the Dell RACADM Reference Guides at support.dell.com/manuals and the topic, [Accessing iDRAC6/iDRAC7 Using RACADM](#).

When the DR Series System splash screen is displayed, you are ready to begin initialization using the DR Series system CLI logon process. For more information, see [Logging in and Initializing the DR Series System](#).

Accessing iDRAC6/iDRAC7 by using RACADM

You can use SSH-based or Telnet-based interfaces to access iDRAC6/iDRAC7 by using the RACADM utility. RACADM (remote access controller administration) is a command-line utility that allows you to set up and configure the integrated Dell Remote Access Control (iDRAC) interface card to provide an out-of-band management capability.

The iDRAC card contains a controller with its own processor, memory, network connection, and access to the system bus. This gives system or network administrators the capability to configure a system as if they were sitting at the local console using the power management, virtual media access and remote console capabilities, by using a supported web browser or command line interface.

The logon values you can use for making iDRAC connections are:

- Default username: root
- Default password: calvin
- Default static IP address: 192.168.0.120

For more information, see the *RACADM Reference Guides for iDRAC*, the *Integrated Dell Remote Access Controller 6 (iDRAC6) User Guide*, or the *Integrated Dell Remote Access Controller 7 (iDRAC7) User Guide* that are available at support.dell.com/manuals.

Logging on and initializing the DR Series system

Before you can start to use the DR Series system graphical user interface (GUI) for the first time, you must initialize the system.

The Initial System Configuration Wizard lets you configure the following network settings to complete the first-time initialization of your system:

- IP addressing mode
- Subnet mask address
- Default gateway address
- DNS suffix address
- Primary DNS server IP address
- (Optional) Secondary DNS server IP address
- Host name for system

When you initialize the system you will log on to the DR Series system CLI by using a local console KVM (keyboard-video monitor) connection or an iDRAC connection. (For more information, see the topic, [Local Console Connection](#), or [iDRAC Connection](#).) You will then configure your system network settings by using the Initial System Configuration Wizard as described in the steps below.

To log on and initialize the DR Series system, complete the following steps:

1. Launch a terminal emulator application (like PuTTY), and type the default IP address for the DR Series system (if you are not using iDRAC or local console).

2. At the login as: prompt, type administrator, and press **<Enter>**.
3. At the administrator@<system_name> password: prompt, type the default administrator password (**St0r@ge!**), and press **<Enter>**.

Figure 11: Initial System Configuration Wizard Window

```
=====
Initial System Configuration Wizard
=====

You logged in to the machine for the first time.

This wizard will help you in setting up the host name, ip address etc.

Would you like to configure network settings (yes/no/later) ? █
```

4. To configure the network settings, type **y** (for yes), and press **<Enter>**.
5. To configure the use of the default IP address that ships with the system, choose to use static IP addressing.
To do this, at the DHCP prompt, type **no** (this selects static IP addressing), and press **<Enter>**.

i **NOTE:** When you select static IP addressing, you are prompted to type the static IP address (for example, you could use the default IP, 10.77.88.99) for the system, and press **<Enter>**. If your network supports the use of DHCP, type yes at the DHCP prompt, press **<Enter>**, and respond to any prompts.
6. To configure a subnet mask address, type the subnet mask address you want to use (for example, you could use the default subnet mask address, 255.0.0.0), and press **<Enter>**.
7. To configure a default gateway address, type the default gateway address you want to use (for example, 10.10.20.10), and press **<Enter>**.
8. To configure a DNS Suffix, type the DNS suffix you want to use (for example, storage.local), and press **<Enter>**.
9. To configure a primary DNS server IP address, type an IP address you want to use for the primary DNS server (for example, 10.10.10.10), and press **<Enter>**.
10. (Optional) To configure a secondary DNS server IP address, type **y** (for yes), and press **<Enter>**.
If you responded yes, type an IP address you want to use for the secondary DNS server (for example, 10.10.10.11), and press **<Enter>**.
11. To change the default host name (for example, the serial number of the DR Series hardware appliance), type **y** (for yes) and press **<Enter>**.
If you responded yes, type the host name you want to use, and press **<Enter>**. After you configure your host name response, the current system settings are displayed.
12. To accept these settings, type **y** (for yes), and press **<Enter>**.
13. If you want to change any of these settings, type **n** (for no), and press **<Enter>**. Modify the settings as needed, and press **<Enter>**.
When completed, a successful initialization message is displayed.
14. At the prompt, type exit and press **<Enter>** to end the DR Series system CLI session.

You are now ready to log in to the system using the DR Series system GUI.

- i** **NOTE:** Before you log on to the system in the DR Series system GUI, make sure to register it in the local Domain Name System (DNS) for your network so that it is a DNS-resolvable entry.
- i** **NOTE:** At this point, you could modify the bonding mode to use 802.3ad, if this configuration is available in your network.

Logging onto the system GUI for the first time

To log on to the DR Series system GUI for the first time, complete the following steps:

1. In a supported web browser, type the IP address or hostname of your DR Series system, and press **<Enter>**.
 - i** **NOTE:** The DR Series System Login page may display a warning message if the web browser you are using does not properly support the DR Series system. If you are running a Microsoft Internet Explorer (IE) web browser, make sure that you disable the Compatibility View. For more information about disabling the Compatibility View settings, see the topic, [Disabling the Compatibility View Settings](#). For more information about supported web browsers, see the *DR Series System Interoperability Guide*.
 - i** **NOTE:** For best results when using IE web browsers in combination with supported Windows-based servers, ensure that Active Scripting (JavaScript) is enabled on the Windows client. This setting is often disabled by default on Windows-based servers. For more information on enabling Active Scripting, see the topic [Enabling Active Scripting in Windows IE Browsers](#).
2. In the **Username** field, type **administrator**, and in the Password field, type **St0r@ge!** and then click Log in or press **<Enter>**. The Users page is displayed with a Configuration Wizard that guides you to the proper pages in the GUI for configuring your system information.
 - i** **NOTE:** If this is the first time logging on to the system after a software upgrade, you will see an Upgrade notification dialog box, which instructs you to review and verify user roles and logon information for your DR Series system and associated DMAs.

3. Click **Next** or **Previous** to navigate through the Configuration Wizard steps to set up your system. Refer to the topics in the **System Configuration** chapter later in this guide for more detailed instructions on how to complete these steps.
 - a. Step 1: Reset Password Setup & Change Admin Password. Refer to [Modifying password reset options](#) and [Changing a user password](#) for detailed instructions.
 - b. Step 2: Fill in Admin Contact Information & Email Relay Host. See [Modifying a user](#) and [Configuring email notification settings](#) for detailed instructions.
 - c. Step 3: Set the hostname, DNS, and IP Mode. See [Configuring networking settings](#) for detailed instructions.
 - d. Step 4: Join Active Directory, if you need to configure Active Directory settings for your system. See [Configuring Active Directory settings](#) for detailed instructions.
 - e. Step 5: Review the Initial Upgrade Message. You should verify that your system has the latest available updates. Visit support.quest.com/dr-series, and navigate to the Software Downloads page for your specific DR Series model to download the latest upgrade file. Before upgrading, you should exit the Configuration Wizard and then proceed to the DR Series system GUI **Support > SoftwareUpgrade** page. See the topic [Upgrading the DR Series system software](#) in the **Support, Maintenance, and Troubleshooting** chapter later in this guide for detailed instructions on upgrading the software.

When you click **Next** on the final page of the Configuration Wizard, a dialog box is displayed confirming that you have completed the initial configuration of your system.

4. In the dialog box, click **OK**. The Dashboard page is displayed.

Your logon username is displayed at the top of the page in the right corner. If you are logged on as a domain user, the domain is displayed in the format of domain\username. (You can only log on as a domain user after configuring Login Groups under Active Directory. This is a requirement for using Global View.)

Enabling active scripting in IE

To enable Active Scripting (JavaScript) in Microsoft Windows Internet Explorer (IE) web browsers, complete the following steps:

i **NOTE:** This procedure describes how to configure your Windows IE web browser to enable Active Scripting (JavaScript). This setting is often disabled by default on Windows-based servers.

1. Launch the IE web browser, and click **Tools→Internet Options**.
2. Click the **Security** tab, and click **Custom level....**
3. Using the right scroll bar, scroll down the **Settings** choices until you reach **Scripting**.
4. In **Active scripting**, click **Enable**.
5. Click **OK** to enable JavaScript and the **Active Scripting** feature for your web browser.
6. Click **OK**.

Disabling the Compatibility View settings for IE

To disable the Compatibility View settings of the IE web browser, complete the following steps:

i **NOTE:** This procedure describes how to disable the Compatibility View settings to ensure there is no conflict between different versions of the Microsoft Internet Explorer (IE) web browser you use to access the DR Series system. Disabling the compatibility view settings requires that the Display all websites in Compatibility View check box option in the Compatibility View Settings page is cleared, and that there are no DR Series systems or domains associated with these systems listed in the Compatibility View list on this page.

1. Launch the IE web browser, and click **Tools→Compatibility View settings**.
2. If selected, clear the **Display all websites in Compatibility View** check box option.
3. If any DR Series systems are listed in the **Compatibility View** list, select the entry and click **Remove**.
4. Click **Close**.

Configuring the DR Series system settings

In the DR Series system GUI, you can easily view and configure system settings such as, active directory, system date and time, expansion shelf enclosures, licenses, networking, schedules for system operations, SSL certificates, storage groups, and users.

i NOTE: Before you run any DR Series system operations for the first time or edit system settings, you should have first initialized the system. Initializing the DR Series system requires that you configure and manage a number of very important system settings before any other tasks, and Quest recommends that you use the Initial System Configuration Wizard to configure your DR Series system. Changing some of the system settings using the DR Series system GUI (such as bonding, MTU, hostname, IP address, and DNS) can cause issues that can affect your DR Series system GUI access. For more information about initializing the system, see the topic, [Initializing the DR Series System](#).

Viewing system configuration settings

In the DR Series system GUI you can easily current system configuration settings.

To view system settings, complete the following steps.

1. In the left navigation menu, click System Configuration.
2. On the System Configuration page you can view the information about your DR Series system.
 - **Active Directory**—Displays the domain name of the Active Directory Service, if configured.
 - **Date and Time**—Displays the mode for system date and time settings, as either Manual or NTP.
 - **Fibre Channel**—Displays current fibre channel access status, as either Enabled or Disabled.
 - **Networking**—Displays the following networking information for your system:
 - Mode
 - Hostname (FQDN)
 - IP Address
 - Bonding
 - Domain Suffix
 - Primary DNS
 - **Contact Information**—Displays contact information for the administrator of your system, if configured.
 - **Email Relay Host**—Displays the email relay hostname, if configured.
 - **Email Alerts**—Displays the numbers of notification recipients for system alerts, if configured.

Using the Configuration Wizard

At any time you can access and use the Configuration Wizard to guide you in configuring settings on your DR Series system. The Configuration Wizard first appears when you log on to the system for the first time. It gives you instructions for configuring settings and opens the proper pages in the GUI for configuring your system information automatically.

To access and use the Configuration Wizard, complete the following steps.

1. In the left navigation menu, click System Configuration.
2. In the Action menu in the upper right corner of the page, click **Configuration Wizard**.

3. Click **Next** or **Previous** to navigate through the Configuration Wizard steps. Refer to the topics in this System Configuration chapter for more detailed instructions on completing these steps.
 - a. Reset Password Setup & Change Admin Password
 - b. Fill in Admin Contact Information & Email Relay Host
 - c. Set the hostname, DNS, and IP Mode.
 - d. Join Active Directory (as needed).
 - e. Review the Initial Upgrade Message.

You should verify that your system has the latest available updates. Visit support.quest.com/dr-series, and navigate to the Software Downloads page for your specific DR Series model to download the latest upgrade file. Before upgrading, you should exit the Configuration Wizard and then proceed to the DR Series system GUI **Support > SoftwareUpgrade** page. See the topic “Upgrading the DR Series System Software” in the “Support, Maintenance, and Troubleshooting” chapter later in this guide for detailed instructions on upgrading the software.

When you click Next on the final page of the Configuration Wizard, a dialog box is displayed confirming that you have completed the initial configuration of your system.

4. In the dialog box, click **OK**. The Dashboard page is displayed.

Registering with the Support Portal

You can now register your DR Series system with the Quest Support Portal by using the DR Series system CLI `system --support_agent` set of commands. The registered DR Series system will collect certain information that it will transmit to Quest Support. Such information can include operational statistics, performance metrics, diagnostic information and configuration settings of the DR Series system. This enables support personnel to monitor and capture information to proactively help troubleshoot issues with your system.

For more information about using the `--support_agent` set of commands, see the DR Series System Command Line Interface Reference Guide.

Configuring Active Directory settings

You can easily join the DR Series system to your Microsoft Active Directory Services (ADS) domain. This topic describes how to configure Active Directory (AD) settings for the DR Series system, which requires that you direct your DR Series system to join or leave a domain that contains a Microsoft Active Directory Service (ADS). Instructions are provided below to join an ADS domain or to leave an ADS domain. When you join the DR Series system to an ADS domain, this disables the Network Time Protocol (NTP) service and instead uses the domain-based time service.

To configure the DR Series system for a domain using ADS, complete the following steps:

1. In the left navigation menu, click **System Configuration > Active Directory**.
2. Do one of the following:
 - On the Active Directory page, click the **Join** link.
 - On the Action menu, click **Join**.

3. Enter the following AD logon information:

- **Domain Name (FQDN)**—Enter a fully qualified domain name for the ADS; for example, AD12.acme.com. *(This is a required field.)*

i **NOTE:** Supported domain names are limited to 64 characters in length and can only consist of a combination of A-Z, a-z, 0-9, and three special characters: a dash (-), a period (.), and an underscore (_).

- **Username**—Enter a valid user name that meets the user name guidelines for the ADS. *(This is a required field.)*

i **NOTE:** Supported user names are limited to 64 characters in length and can only consist of a combination of A-Z, a-z, 0-9, and three special characters: a dash (-), a period (.), and an underscore (_).

- **Password**—Enter a valid password that meets the password guidelines for the ADS. *(This is a required field.)*

- **Org Unit**—Enter a valid organizational name that meets the organization name guidelines for the ADS. *(This is an optional field.)*

4. Click **Join**.

5. To leave an ADS domain, in the Action menu in the upper right corner of the page, click **Leave**.

- a. In the Leave Active Directory pane, enter the username and password for the ADS domain.
- b. Click **Leave**.

Adding a login group to an ADS domain

After you configure your DR systems within the same ADS domain, you must ensure that a login group exists and add it to the domain.

Adding a login group is only possible when the DR Series system is joined to a domain. Also, you must be logged in as a domain user that is part of an enabled login group.

To add a login group in an ADS domain, complete the following steps:

1. On the left navigation menu, click **System Configuration > Active Directory**.
2. In the **Action** menu in the upper right corner of the page, click **Add Login Group**.
3. In the **Add Login Group** pane, type the name of the login group including the domain name; for example, *Domain\Domain Admins*. If your login group name contains spaces, you must not enclose it in quotation marks. (This differs from the equivalent CLI command.)
4. Click **Add Login Group** to add the login group.

i **NOTE:** Changes made to the login group take effect on the next log in attempt (unlike Windows ADS, no active checking is done on the group).

Removing a login group

To remove a configured ADS login group from the DR Series system, complete the following steps.

1. On the left navigation menu, click **System Configuration > Active Directory**.
2. In the **Action** menu in the upper right corner of the page, click **Remove Login Group**.
3. When prompted to confirm, click **OK**.

i | **NOTE:** Changes made to the login group take effect on the next log in attempt (unlike Windows ADS, no active checking is done on the group).

Configuring system date and time settings

In the DR Series system GUI, you can configure date and time settings, for example, to synchronize with other DR Series systems or clients running in your domain.

i | **NOTE:** System synchronization is critical for proper data archiving and replication service operations.

The Date and Time information settings displayed in the system GUI include:

- Mode
- Time Zone
- Date and Time

To modify the default time and date settings for your DR Series system, complete the following steps:

1. In the left navigation menu, click **System Configuration > Date and Time**. The Date and Time page is displayed, showing the following information:
 - **Mode**—one of two types: Manual and Network Time Protocol (NTP).

i | **NOTE:** You should use NTP when the DR Series system is part of a workgroup and not part of a domain to ensure better replication service operations. By using the NTP mode, you synchronize your system clock whereby NTP ensures that your system has a reliable time stamp. This is critical for successful file exchanges, network log coordination and validation, and resource access requests within a workgroup. When the DR Series system is joined to a domain, such as the Microsoft Active Directory Services (ADS) domain, NTP is disabled and the DR Series system uses the domain time.
 - **Time Zone**—you can select from a list of time zone options around the world.
 - **NTP Servers**—when using NTP mode, you can define up to three NTP servers.
 - **Date and Time**—when in Manual mode, the date and time can be shown as month, day, and time in a 24-hour time format.

2. In the upper right corner of the page, click **Edit**.

i **NOTE:** If the DR Series system is joined to a Microsoft Active Directory Services (ADS) domain, the Edit option will not be available and the Mode, Time Zone, or Date and Time values cannot be changed. If a DR Series system is joined to a domain, the Network Time Protocol (NTP) is disabled and the DR Series system uses the domain-based time service. NTP is used in the Mode setting when the DR Series system is part of a workgroup and not joined to a domain.


3. Next to Mode, select either **Manual** or **NTP**.
 - If you select **Manual**, follow these steps.
 - a. In the **Time Zone** drop-down list, select a time zone.
 - b. In the **Date and Time** drop-down calendar, select the date and then click the arrows to adjust the hour:minute:seconds for the time.
 - If you select **NTP**, follow these steps.
 - a. In the **Time Zone** drop-down list, select a time zone.
 - b. Enter or edit the NTP server information. You can enter information for the following three entries: Ntp Server 0 (required), Ntp Server 1, Ntp Server 2.
4. Click **Save**.

Configuring DR Series system enclosures

To properly set up, add, and connect an expansion shelf enclosure to a DR Series system hardware appliance, you need to complete the following tasks.

1. Power off the DR Series system.
2. Install all cabling that connects the external expansion shelf (or shelves) to the DR Series system (For information, see the topic, "Understanding DR Series System Expansion Shelf Cabling").
3. Power on the external expansion shelf (or shelves), and then power on the DR Series system.
4. Install the license for the expansion shelf enclosures.
5. In the DR Series system GUI, add and activate the expansion shelf enclosure (as described in the steps below).

To view, add and activate expansion shelves in the DR Series system GUI, complete the following steps:

1. In the left navigation menu, click **System Configuration > Enclosures**.
The **Enclosures** page is displayed. (This step assumes that you have completed all expansion shelf enclosure cable connections and that green LEDs are displayed next to the fastplugs on the rear chassis, indicating that cable connections are active.) On this page, you can view information about expansion shelf enclosures in the Physical Storage list, which displays the following information:
 - **Type** — Designates whether the enclosure is internal or external (noted as Enclosure-1, Enclosure-2)
 - **Raw size**—Lists the total capacity size on the device.
 - **Percentage Used**—Lists the percentage of used capacity.
 - **Service Tag**—Lists the service tag number for the enclosure.
 - **Configured**—Lists whether the enclosure has been configured. Possible values are: "Yes," "No," and "Device not Found."
 - **Status**—Displays a status icon as a green check mark or a red icon with a white X.
2. To add an expansion shelf, in the **Physical Storage** list, in the **Action** column, click **Detect** next to the Enclosure (1 or 2) that you want to add.
3. Click **OK** to continue adding the enclosure to the DR Series system.
 **NOTE:** If the enclosure is undetected, wait 5 minutes and try again. If the enclosure still remains undetected, keep the enclosure powered on and reboot your DR Series system appliance.
4. Once the previous step completes, to verify that an enclosure was added, you can view it in the list of Enclosures or click **Dashboard→Health**.

Managing and viewing DR Series system licenses

In the DR Series system you can easily view current system license information, upload and install a new license file, as well as unregister a DR2000v system from a DR Series hardware system (license server). To view and manage licenses in the DR Series system GUI, complete the following steps:

1. In the left navigation menu, click **System Configuration > Licenses**.
2. If you are viewing the Licenses page for any hardware-based DR Series system (that is, not a DR2000v), you can view the following information.
 - **Registered Licenses**—Displays the number of registered licenses for the following categories:
 - Number of DR2000v Licenses Installed
 - Remaining DR2000v Licenses
 - Number of Enclosure Licenses Installed
 - Remaining Enclosure Licenses
 - **Installed Licenses**—Displays a list of installed licenses and the following columns of information.
 - ID—An internal system-related identification number.
 - Expiration Day—Lists N/A unless a DR Series system temporary license is applied in which case it lists the day the license expires.
 - Description—A description of the license, for example, "36TB storage enclosure license."
 - License number—The license number.
 - Status—The license status, such as "enabled" if in use
 - **Registered DR2000V(s)**—Displays a list of the DR2000v systems (if any) registered to the DR Series system and the following columns of information, which is dependent on the name and email information registered on the DR2000v.
 - Customer Name—The name registered on the DR2000v when registering the system.
 - VM Name—The hostname of the VM.
 - Capacity—The capacity of the DR2000v.
 - Email—The email address of the administrator registered on this DR2000v.
 - Service Tag—The randomized service tag number generated by the DR2000v.
3. To upload and install a new license file, at the top of the page under Upload License File, click **Select License File**, and then browse to and select the license file you want to upload

4. If you are viewing the **Licenses** page of a DR2000v, you can view the following details on the license page.
 - **Client information**—Displays information about the current DR2000v.
 - **Model**—The model of the DR Series system.
 - **Company Name**—The company name associated with the administrator user, if configured.
 - **Administrator Name**—The name associated with the administrator user, if configured.
 - **Registered**—The status of the license registration of this DR2000v as Yes or No.
 - **Service Tag**—The Service Tag number of the DR2000v.
 - **System Name**—The hostname of the DR2000v.
 - **Administrator Email**—The email address associated with the administrator user, if configured.
 - **Storage**—The storage capacity of the DR2000v
 - **Comments**—Any comments entered to describe this system.
 - **Server information** —Displays information about the DR Series hardware system to which the DR2000v is registered.
 - **Model**—The model of the DR Series system to which the DR2000v is registered.
 - **Server license**—The IP address of the licensed DR Series system.
 - **Update time**—The timestamp of the last time the DR2000v synced with the DR Series system.
5. On the DR2000v **Licenses** page, to unregister a DR2000v license from its associated DR Series license server, do the following:
 - a. On the **Action** menu in the upper right corner of the page, click **Unregister**.
 - b. In the **Warning** dialog box, click **OK** to unregister the license.
6. To edit the license registration information for a DR2000v, do the following:
 - a. On the bmenu in the upper right corner of the page, click **Edit**.
 - b. In the **Edit Registration** pane, modify the following information as needed, and then click **Save**.
 - **IP Address/Hostname of DR License Server**—The IP address/hostname of the DR Series system/license server to which this DR2000v is registered.
 - **Administrator Name**—The name of the administrator user.
 - **Adminstrator Email**—The email address associated with the administrator user.
 - **Company Name**—The company name associated with the administrator user.
 - **Comment**—Text comments entered to describe this system.

Configuring networking settings

In the DR Series system GUI, you can view and modify the network settings (for example, if you want to change the settings you configured when you first initialized the system)

To view and modify networking settings, complete the following steps.

1. In the left navigation menu, click System **Configuration > Networking**.
2. Under Hostname, you can view the following settings.
 - **Hostname**
 - Hostname (FQDN)
 - iDRAC IP Address
 - **DNS**
 - Domain Suffix
 - Primary DNS
 - Secondary DNS
3. To edit the Hostname or DNS settings, in the Action menu in the upper right corner of the page, click **Edit**.
 - a. In the **Hostname** field, modify the hostname as needed. The hostname entry must meet the following requirements:
 - Alphabetic—allows A-Z, a-z, or a combination of upper and lower case alphabetic characters.
 - Numeric—allows numerals zero (0) through 9.
 - Special characters—allows only the dash (-) character.
 - Length limit—hostnames cannot exceed the maximum length of 19 characters.
 - b. Modify the DNS fields as needed (Domain Suffix, Primary DNS, Secondary DNS).
 - Domain Suffix — enter a domain suffix to use. For example, acme.local. This is a required field.
 - Primary DNS—enter an IP address that represents the primary DNS server for your system. This is a required field.
 - Secondary DNS—enter an IP address that represents the secondary DNS server for your system. This is an optional field.
 - c. Click **Save**.
4. Under Interfaces, you can view the following columns of information for a selected NIC bond or Ethernet port.
 - Device—The name of the device.
 - Enabled—Yes or no.
 - IP Address—The IP address of the NIC bond or Ethernet port.
 - MAC Address—The IP address of the NIC bond or Ethernet port.
 - MTU—The Maximum Transmission Unit setting. (The MTU setting accepts values between 512 and 9000.)
 - Traffic—The type of traffic as Replication, Management, Backup, and Opdup.

You can click the Device to expand to view more details, such as the interface name, mode, member interfaces, netmask address, gateway address, bonding option, NIC, MAC, maximum speed, and duplex.

5. To modify settings for a selected interface (NIC bond or Ethernet port), in the Action column, click the **Edit** icon.

- a. In the Edit Interface pane, for IP Mode, select Static (to set static IP addressing for your system) or DHCP (to set dynamic IP addressing for your system).
- b. If you selected Static for the IP Mode, enter the IP , Netmask, and Gateway addresses. (the system IP address and netmask identify the network to which your system belongs).
- c. For MTU, enter a number to set as the maximum. (The MTU setting accepts values between 512 and 9000.)

i NOTE: Ensure that the value that you enter in MTU is the same for the clients, Ethernet Switch, and the appliance. The connection between the clients, the Ethernet switches, and the appliance will break if the MTU number is not the same on all the components.


i NOTE: In computer networking, jumbo frames are Ethernet frames with more than 1500 bytes of payload (but in some cases, jumbo frames can carry up to 9000 bytes of payload). Many Gigabit Ethernet switches and Gigabit Ethernet network interface cards support jumbo frames. Some Fast Ethernet switches and Fast Ethernet network interface cards also support jumbo frames. Some computer manufacturers use 9000 bytes as the conventional limit for jumbo frame sizes. To support jumbo frames used in an Internet Protocol subnetwork, both the host DR Series system (initiator or source) and the target DR Series system have to be configured for 9000 MTU.

Consequently, interfaces using a standard frame size and those using the jumbo frame size should not be in the same subnet. To reduce the chance of interoperability issues, network interface cards capable of supporting jumbo frames require specific configurations to use jumbo frames.

To verify that the destination system can support a specific frame size, use the DR Series system CLI command `network --ping --destination <IP address> --size <number of bytes>`. For more information, contact Technical Support for assistance.

i NOTE: Make sure that if you are using any Dell network switches that you take full advantage of the latest switch firmware upgrades and application notes. The application notes provide procedures that assist you in performing switch firmware upgrades and saving configuration files (for complete details, see support.dell.com/ and navigate to **Drivers and Downloads** for your system type).


i NOTE: When setting or changing the MTU value, make sure that you verify that the Ethernet network switch is capable of supporting an MTU size that is equal to or larger than the value you are setting. Any mismatch in MTU values between the clients, Ethernet network switch, and the DR Series system appliance will make it inoperable.

- d. In the Bonding Configuration drop-down list, select one of the following bonding configurations.
 - **ALB**—Configures adaptive load balancing (ALB), which is the default setting.
Note: ALB load balancing does not balance the load properly when your backup servers are on a remote subnet. This is because ALB uses the address resolution protocol (ARP) and ARP updates are subnet-specific. Because this is the case, ARP broadcasts and updates are not sent across the router. Instead, all traffic is sent to the first interface in the bond. To resolve this ARP-specific issue, make sure that your data source systems reside on the same subnet as the DR Series system.
 - **802.3ad**—Configures dynamic link aggregation using the IEEE 802.3ad standard.
 -  **CAUTION:** If you change the existing bonding setting, the connection to the DR Series system may be lost unless you are sure that the system can accept this bonding type.
 - e. Click **Save**.
6. Under Bandwidth Limits, you can view the following columns of information, which represent network bandwidth for replication connections:
 - Target System—The DR Series system to which network traffic is being sent.
 - Target IP—The IP address of the DR Series system to which network traffic is being sent.
 - Bandwidth Limit—This limit is shown in KBps, MBps, or GBps and represents the user-set value of the Max Speed.
 7. To modify Bandwidth Limit settings, in the Action column, click the **Edit** icon.
 - a. In the Edit Bandwidth Limit pane, for Bandwidth Speed Rate, select **Unlimited** or **Limited**.
 - b. If you select Limited, for Bandwidth Speed Value, enter a number for the bandwidth speed, and then select **KBps**, **MBps** or **GBps** from the drop-down list.
 - c. Click **Save**.

Understanding the networking page and Ethernet port values

In the DR Series system GUI, the Networking page displays the currently configured multiple Ethernet ports for the DR Series system in the Interfaces list. For 1-Gigabit Ethernet (GbE) ports in the DR4000 system this could be Eth0, Eth1, Eth2, and Eth3, and in the DR4100 system this could be Eth0, Eth1, Eth2, Eth3, Eth4, and Eth5. For 10-GbE/10-GbE SFP+ NICs, this means that the two ports are bonded together into a single interface. For example, the DR Series system port configuration is as follows:

- In a 1-GbE NIC configuration: the DR4000 system supports up to four 1-GbE ports, which consists of up to two internal LAN on Motherboard (LOM) ports and two ports on an expansion card that are bonded together. The DR4100 system supports up to six 1-GbE ports, which consists of up to four internal LOM ports on the network daughter card (NDC) and two ports on a PCI Express expansion card.
- In a 10-GbE or 10-GbE SFP+NIC configuration: the DR4000 system supports up to two 10-GbE or 10-GbE SFP+ ports on an expansion card that are bonded together. The DR4100 system supports up to two 10-GbE or 10-GbE SFP+ ports that reside on the NDC that are bonded together.

 **NOTE:** For more information on advanced networking options see the *DR Series Command Line Interface Reference Guide* available at support.quest.com/dr-series by selecting your specific DR model and then navigating to Technical Documentation.

The ports for bonded NICs display: MAC address, port speed in megabytes per second (MB/s), maximum speed, and duplex setting. The following example shows Ethernet port values for the four ports in a 1-GbE NIC bonded configuration on a DR4000 system:

Eth0:

- MAC: 00:30:59:9A:00:96
- Speed: 1000Mb/s
- Max Speed: 1000baseT/Full
- Duplex: Full

Eth1:

- MAC: 00:30:59:9A:00:97
- Speed: 1000Mb/s
- Max Speed: 1000baseT/Full
- Duplex: Full

Eth2:

- MAC: 00:30:59:9A:00:98
- Speed: 1000Mb/s
- Max Speed: 1000baseT/Full
- Duplex: Full

Eth3:

- MAC: 00:30:59:9A:00:99
- Speed: 1000Mb/s
- Max Speed: 1000baseT/Full
- Duplex: Full

Understanding system operation scheduling

By scheduling system operations, you can optimize your system resources and achieve the best possible DR Series system performance. The most important thing to remember when scheduling critical DR Series system operations is that you want to ensure that you perform each of these operations at a time when it will not overlap or interfere with the running of any of the other key system operations.

You should carefully plan and schedule time periods in which to perform the following critical system operations:

- Data ingests (which are dependent upon your usage of your DMA(s))
- Replication

- System cleaner (space reclamation)

i **NOTE:** Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from system containers in which files were deleted as a result of deduplication.

In the DR Series system, the main goal in planning and scheduling operations should be to run the Cleaner and Replication operations at times when they do not overlap or interfere with other important system operations. You want to make sure that by properly scheduling and planning, your system can perform each of these key operations independent of the other.

The best practice is to run these two operations during non-standard business hours, so that they do not conflict with any of your other backup or ingest operations.

To view currently scheduled operations and to access the Schedules page to configure system schedules, follow these steps.

- On the left navigation menu, click **System Configuration > Schedules**.
The Schedules page is displayed, showing a daily/hourly calendar representation of currently scheduled operations.

Configuring cleaner schedules

Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from the system. The recommended method is to schedule a time when you can run the Cleaner on your DR Series system with no other planned processes running.

i **NOTE:** Even if no Cleaner schedule is set, but the system detects that there is disk space that can be reclaimed, the Cleaner process runs. However, the Cleaner will not start until the following conditions are met: it detects that there are no active data ingests, that two minutes of system idle time have elapsed since the last data file ingest was completed, and that the Replication process is not running (the Cleaner process runs as a lower system priority operation than the Replication process).

i **NOTE:** Running the Cleaner while ingesting data reduces system performance. Ensure that you schedule the Cleaner to run when backup or replication is not in progress. For more information about best practices for using the system cleaner, refer to the *DR Series System Cleaner Best Practices* white paper for your specific DR Series system at support.quest.com/DR-Series.

To schedule cleaner operations on your system, complete the following steps.

1. On the left navigation menu, click **System Configuration > Schedules**.
The Schedules page is displayed, showing a weekly/daily calendar representation of currently scheduled operations.

2. To add a cleaner schedule, on the Action menu in the upper right corner of the page, select **Add Cleaner Event**.
 - a. In the Cleaner pane, select a day from the **Set Event At Day** drop-down list.
 - b. In the **From** and **To** drop-down lists, select the time (hour) of the day for the cleaner to run.
 - c. Click **Save**.

i NOTE: Even if no Cleaner schedule is set, but the system detects that there is disk space that can be reclaimed, the Cleaner process runs. However, the Cleaner will not start until the following conditions are met: it detects that there are no active data ingests, that two minutes of system idle time have elapsed since the last data file ingest was completed, and that the Replication process is not running (the Cleaner process runs as a lower system priority operation than the Replication process).

i NOTE: You can also use the DR Series system command line interface (CLI) to create, view, and delete the cleaner schedule. The available commands are:

```

                                schedule --add --day <Day of the Week> --
start_time <HH:MM> --stop_time <HH:MM> --cleaner
                                schedule --show --cleaner
                                schedule --delete --day <Day of the Week> --
cleaner

```

For full details on running the cleaner schedule commands, help is available by entering:

```

                                schedule --help

```

3. To run the cleaner now, on the Action menu in the right corner of the page, select **Run Cleaner Now**.

Viewing cleaner status

The current Cleaner Status is shown at the top of the Schedules page as one of the three following states:

- **Pending**—displays any scheduled window for the Cleaner operation.
- **Running**—displays when the Cleaner operation is running.
- **Idle**—displayed only if there is no Cleaner operation running.

On the Dashboard page, you can also view a graph of cleaner statistics.

- In the Action menu in the upper right corner, click **Detailed Graphs**, and then click the **Cleaner** tab.

Viewing cleaner statistics

To view additional detailed cleaner statistics, you can use the DR Series system CLI stats --cleaner command to show the following categories of statistics:

- Last Run Files Processed (number of files processed by Cleaner)
- Last Run Bytes Processed (number of bytes processed by Cleaner)
- Last Run Bytes Reclaimed (number of bytes reclaimed by the Cleaner)
- Last Run Start Time (indicates date and time last Cleaner process started)
- Last Run End Time (indicates date and time last Cleaner process ended)

- Last Run Time To Completion(s) (indicates the number of times that Cleaner process has successfully completed)
- Current Run Start Time (indicates date and time current Cleaner process started)
- Current Run Files Processed (number of files processed by current Cleaner process)
- Current Run Bytes Processed (number of bytes processed by current Cleaner process)
- Current Run Bytes Reclaimed (number of bytes reclaimed by the current Cleaner process)
- Current Run Phase 1 Start Time (indicates date and time for start of current Cleaner process phase 1)
- Current Run Phase 1 Records Processed (lists the number of data records processed in current Cleaner process phase 1)
- Current Run Phase 1 End Time (indicates date and time for end of current Cleaner process phase 1)
- Current Run Phase 2 Start Time (indicates date and time for start of current Cleaner process phase 2)
- Current Run Phase 2 Records Processed (lists the number of data records processed in current Cleaner process phase 2)
- Current Run Phase 2 End Time (indicates date and time for end of current Cleaner process phase 2)
- Current Run Phase 3 Start Time (indicates date and time for start of current Cleaner process phase 3)
- Current Run Phase 3 Records Processed (lists the number of data records processed in current Cleaner process phase 3)
- Current Run Phase 3 End Time (indicates date and time for end of current Cleaner process phase 3)
- Current Run Phase 4 Start Time (indicates date and time for start of current Cleaner process phase 4)
- Current Run Phase 4 Records Processed (lists the number of data records processed in current Cleaner process phase 4)
- Current Run Phase 4 End Time (indicates date and time for end of current Cleaner process phase 4)

For more information about DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Configuring replication schedules

You can configure replication schedules on a weekly basis for individual replication-enabled source containers.

i NOTE: If there is no Replication schedule set, but there is pending data that can be replicated, replication will run when it detects three (3) minutes of idle time for any newly written files in the replicated container.

i NOTE: It is recommended that you do not schedule the running of any Replication operations during the same time period when Cleaner or data ingest operations will be running. If you do not follow this recommendation, the time required to complete the system operations or system performance might be affected.

To configure replication schedules, complete the following steps.

1. On the left navigation menu, click **System Configuration > Schedules**.
The Schedules page is displayed, showing a weekly/daily calendar representation of currently scheduled operations.
2. To view the Replication Schedule, in the Replication Schedule drop-down list, you can select one of the following types of replication schedules to view:
 - None
 - All Containers
 - Source:
 - All
 - <A specific backup container>
 - Target:
 - All
 - <A specific backup container>
3. To add a replication schedule, on the Action menu in the upper right corner of the page, click **Add Replication Event**.
 - a. In the Replication pane, select the replication container for which you want to schedule replication in the **Set Event to Container** drop-down list.
 - b. In the Day drop-down list, select the day for the schedule to run.
 - c. In the **From** and **To** drop-down lists, select the time (hour) of the day for the replication to run.
 - d. Click **Save**.

You can view replication details and status on the Replications page in the DR Series system GUI by selecting a replication and clicking to expand and view details. The current replication status is also shown on the GlobalView page by expanding the detailed view of a selected DR Series system.

Configuring an SSL certificate for your DR Series system

For additional security, you can replace the self-signed, factory-installed certificate with another SSL certificate, for example, with one that is signed by a third-party CA. Once you have obtained your signed certificate, you can install it by using the DR Series system GUI. Only one certificate can be installed on a DR Series system at any given point in time.

- To view the SSL Certificate page, on the left navigation menu, click **System Configuration > SSL Certificate**.

Installing an SSL certificate

To install an SSL certificate, complete the following steps:

1. In the left navigation menu, click **System Configuration**→SSL Certificate.
2. Next to Certificate File Location, click **Browse** to locate and select the SSL certificate on your system that you want to install.

i | **NOTE:** Only .pem formatted SSL certificates are supported.

3. On the SSL Certificate page, click **Install Certificate**.
4. In the Install SSL Certificate dialog box, click **Continue**.
Unless corrupted or expired, certificates files of .pem format type with less than 2048-bit encryption should successfully verify.
5. In the certificate Validation dialog box, click **Continue**.
In the event you see the Certificate Verification Failed dialog box, clicking on “Continue” here will generate a connection reset in the browser. You will still be allowed to continue with certificate installation. Upon successful installation of a certificate, an HTTP server restart is performed, and the browser will move to a connection reset state.

i | **NOTE:** If your browser cannot connect to a DR Series system after a certificate installation, you may need to reset the certificate from the command line interface (CLI) using “maintenance --configuration --reset_web_certificate”. Refer to the *DR Series Command Line Reference Guide* for more information.

6. Click either the page reload icon or the back-arrow on the browser to restore the page.

Resetting the SSL Certificate

If you have installed a custom SSL certificate, you can reset the SSL certificate back to the factory-installed, self-signed certificate. To reset an SSL certificate, complete the following steps:

1. In the left navigation menu, click **System Configuration**→SSL Certificate.
2. On the **Action** menu in the upper right corner of the page, click **Reset SSL Certificate**.

i | **NOTE:** You can also use the command line interface (CLI) command, `maintenance --configuration --reset_web_certificate`. Refer to the *DR Series Command Line Reference Guide* for more information

Installing a CA certificate

You can add a custom Certificate Authority (CA) certificate chain in the DR Series system GUI. To install a Certificate Authority (CA) Certificate file, complete the following steps:

1. In the left navigation menu, click **System Configuration**→SSL Certificate.
2. Next to **Certificate Authority (CA) Certificate File Location**, click **Browse** to locate and select the CA certificate on your system that you want to install.

i | **NOTE:** When adding a CA certificate, you should add from the root of the CA chain first.

Generating a CSR

You can generate a certificate signing request (CSR) from the SSL Certificate page. A certificate authority (CA) can use the CSR to create an SSL certificate for you. This CSR will contain information to be included in the certificate, such as organization name, common name (domain name), locality, and country. It also contains the public key that will be included in the certificate. Generating a new CSR generates a new private key therefore certificates signed with a previously generated CSR will no longer be installable.

To generate a CSR, complete the following steps:

1. In the left navigation menu, click **System Configuration→SSL Certificate**.
2. On the **Action** menu in the upper right corner of the page, click **Generate CSR**. The Generate CSR pane is displayed.
3. Enter the following required information in the form:
 - **Common Name** - The domain to be secured by the certificate.
 - **Organization Name** - The organization's legal business name.
 - **Organization Unit** - A department in the organization.
 - **Locality** - The business location.
 - **State Name** - The state/province of the business location
 - **Country Code** - The country of the business location.
 - **Email** - A contact email address.
 - **Encryption** - Select one of the following options: 2048-bit encryption or 4096 encryption. The default is 2048.

4. Click **Generate**.

The Certificate request output will appear in the window. You can copy and paste the CSR to the CA's web site CSR page, or you can save the CSR to a file

i NOTE: Every time a CSR is generated, a new private key is generated and stored on the DR Series system. When the signed certificate is returned from the CA, and you attempt to install the signed certificate, a verification that the installed signed certificate matches the private key is performed. If the installed certificate does not match the private key, the certificate installation will fail due to private key match failure. You should be careful not to run a subsequent CSR generation while your initial CSR is being signed by a CA, as the returned certificate will no longer match the private key.

5. Click **Save to File** to save it to a file.

Managing storage groups

To organize your data, you can easily create storage groups and create containers within those storage groups on your DR Series system. A storage group allows you to create separate storage policies for different data groups and the different capacities utilized on a single DR Series system. You can also create and organize storage groups for the different organizations in your enterprise, such as Engineering, Sales, Finance, and so on.

After initialization, the DR Series system contains a single default storage group, named DefaultGroup. Refer to these important notes about storage groups.

- Only administrator users can create storage groups.
- Data/containers cannot be moved between storage groups.
- Deduplication is defined at the storage group level and is not global to the appliance.
- Encryption is defined separately for each storage group.
- Compression is defined separately for each storage group.
- The system cleaner cannot be run on a single storage group; only at the system level.
- Before you can delete a storage group, you must delete all containers contained in that storage group.
- A filesystem scan can be run on a single storage group.

i **NOTE:** When upgrading to the DR Series system 4.0 release, existing containers will be placed into the Default Storage Group. You cannot move a legacy container to a newly created storage group after the upgrade. The administrator can create additional storage groups as needed.

To view the Storage Groups page, on the left navigation menu, click **System Configuration > Storage Groups**.

Viewing storage group information

In the DR Series system GUI you can easily view all of your storage groups on the Storage Groups page.

To view storage groups, complete the following steps.

1. In the left navigation menu, click System Configuration→Storage Groups.
2. On the Storage Groups page you can view the following columns of information about your storage groups.
 - **Name**—Displays the name of the Storage Group.
 - **Compression**—Displays the compression type as either Fast or Best.
 - **Encryption**—Displays whether Encryption is turned On or Off.
 - **Containers**—Displays the number of containers in this storage group. You can click this number, which links to the Containers page for the storage group.
 - **Replications**—Displays the number of replications defined in this storage group. You can click this number, which links to the Replications page for the storage group.
3. You can also select a storage group and click to expand it to view more details, including:
 - **Encryption**—Displays the following Encryption details:
 - **Passphrase**—Displays whether a passphrase is Set or Not Set for this storage group.
 - **Encryption Mode**—Displays the mode of encryption applied to the storage group as either None, Static, or Internal.
 - **Key Rotation Interval Days(s)**—Displays the number of key rotation interval days as N/A, or the number that was set for Internal Encryption Mode.

4. To view a chart of throughput activity and detailed storage group statistics, select a storage group, and, in the Actions column, click the Chart icon. A page for the selected storage group is displayed, showing the following:
 - Throughput chart—Displays the number of read data in Mebibytes/per second (MiB/s) based on time (in minutes), and the number of write data in MiB/s based on time (in minutes)
 - Statistics—Lists the following statistics for the storage group.
 - Capacity Used
 - Capacity Used in GB
 - Current Files
 - Post Dedupe Bytes
 - Post Encryption Bytes
 - Bytes Decrypted
 - Compression Status
 - Dedupe Savings
 - Total Inodes
 - Capacity Free
 - Capacity Free in GB
 - Current Bytes
 - Post Compression Bytes
 - Post Encryption Bytes in GiB
 - Cleaner Status
 - Encryption Status
 - Compression Savings
 - Total Savings

Adding a storage group

In the DR Series system GUI you can easily add storage groups on the Storage Groups page. During the process of storage group creation, you can define data storage features such as Storage Optimization/Compression Level and Encryption at Rest.

To add a storage group, complete the following steps.

1. In the left navigation menu, click **System Configuration→Storage Groups**.
2. On the Action menu in the upper right corner of the page, click **Add a Storage Group**. The Add a Storage Group pane is displayed.
3. In the **Group Name** field, enter a name for the storage group.
4. Click **Next**.

5. For Storage Optimization, select a Compression Level from the drop-down list:

- **Fast** — Results in shorter backup time, but with less space savings.
- **Best** — Provides the highest space savings, but with a longer backup time.

6. Click **Next**.

7. If you want to apply Encryption to the data in this storage group, configure the following settings.

i **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale. For more information about recommended guidelines for encryption, see [Understanding Encryption at Rest](#)

- **Passphrase**—Enter the passphrase to be used to encrypt content encryption keys. (The passphrase string can take up to 255 characters. And, alphanumeric and special characters can be entered as part of the passphrase string.)

i **NOTE:** It is mandatory to define a passphrase to enable encryption. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.

- **Confirm Passphrase**—Re-enter the encryption passphrase.
- **Encryption**—Next to Encryption, click **On**.
- **Encryption Mode**—Select the mode of key lifecycle management from one of the following options:
 - None
 - Static
 - Internal
- **Key Rotation Interval (Days)**—If you selected Internal as the mode of key management, select the number of days for key rotation when a new key is to be generated.

i **NOTE:** In Internal mode there is a maximum limit of 1023 keys. The key rotation period is set to 30 days by default when the passphrase is set and/or encryption is turned on. You can later change the key rotation period from 7 days to 70 years for internal mode.

i **NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. Note that encryption is an irreversible process.

8. Click **Next**.

9. Review the Storage Group summary, and then click **Save**.

i **NOTE:** Once you have created a storage group, you can click the link in the Containers column, which opens the Containers page. This makes it easy to add new containers to your storage group.

Modifying a storage group

In the DR Series system GUI you can easily modify certain settings for existing storage groups on the Storage Groups page.

Figure 12: To modify a storage group, complete the following steps.

1. In the left navigation menu, click System Configuration→Storage Groups.
2. In the Storage Group list, select a storage group, and, in the Actions column, click the Edit icon. The Edit Storage Group pane is displayed.
3. In the **Group Name** field, edit the name for the storage group as needed.

i | **NOTE:** You cannot modify the name of the DefaultGroup.

4. For Storage Optimization, select a Compression Level from the drop-down list as needed:

- **Fast**—Results in shorter backup time, but with less space savings.
- **Best** — Provides the highest space savings, but with a longer backup time.

5. Click **Next**.

6. You can modify the following Encryption settings:

i | **NOTE:** For more information about recommended guidelines for setting up encryption, see the topic, [Configuring and Using Encryption at Rest](#).

i | **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

- **Passphrase**—Enter the passphrase to be used to encrypt content encryption keys. (The passphrase string can take up to 255 characters. And, alphanumeric and special characters can be entered as part of the passphrase string.)
- **Confirm Passphrase**—Re-enter the encryption passphrase.
- **Encryption**—Next to Encryption, click **On** or **Off** as needed.
- **Encryption Mode**—Select the mode of key lifecycle management from one of the following options:
 - None.
 - Static—A global, fixed key is used to encrypt all data.
 - Internal—Content encryption keys are generated and rotated on a specified period of days.

i | **NOTE:** The key mode can be changed at any time during the lifetime of the DR Series system; however, changing the key mode can be a significant operation to undertake as all encrypted data must be re-encrypted with the new mode. If you change the mode to **None**, no further data will be encrypted.

- **Key Rotation Interval Day(s)**—If you selected Internal as the mode of key management, select the number of days for key rotation when a new key is to be generated.

i | **NOTE:** In Internal mode there is a maximum limit of 1023 keys. The default key rotation period is set to 30 days by default when the passphrase is set and/or encryption is turned on. You can later change the key rotation period from 7 days to 70 years for internal mode.

i | **NOTE:** After encryption is enabled, all of the data that is backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. Note that encryption is an irreversible process.


7. Click **Next**.
8. Review the Storage Group summary, and then click **Save**.

Deleting a storage group

Before you can delete a storage group in the DR Series system GUI, you must first delete the containers in the storage group.

To delete a storage group, complete the following steps.

1. In the left navigation menu, click **System Configuration→Storage Groups**.
2. In the list of storage groups, select a storage group, and, in the **Actions** column, click the **Delete** icon.

 **NOTE:** You cannot delete the DefaultGroup storage group.

3. When prompted to confirm, click **OK**.

Managing users

The DR Series system gives you the ability to define user roles and assign users to those roles. A user can have more than one role. There are default user roles for the system as well.

The DR Series system has the following types of user roles: CIFS, OST, RDA, NDMP, iSCSI, monitor, email_recipient, and administrator. For the protocol specific user roles, the user is validated with the protocol credentials when the clients connect. Users with the role, email_recipient, will be able to receive email alerts. The monitor user has read only access in the system GUI.

Refer to these important notes about user management in the DR Series system.

- Excluding the administrator, users can have multiple roles.
- However, NDMP, OST, and iSCSI users are exclusive user roles that can only be assigned to one user at a time.
- The administrator user is a special default user; it cannot be deleted and no new administrator roles can be created.
- The maximum number of users that can be created for the system is 64.

To view the Users page, on the left navigation menu, click **System Configuration > Users**.

Viewing users

In the DR Series system GUI you can easily view a list of the system users.

To view users, complete the following steps.

1. In the left navigation menu, click **System Configuration→Users**.
2. On the Users page you can view the following columns of information about users of your system.
 - **Name**—Displays the name of the user.
 - **Role**—Displays the role(s) assigned to this user.


3. You can also select a user and click to expand it to view more details, including:
 - **Full Name**—Displays the user's full name if it has been added.
 - **Phone**—Displays the user's phone number if it has been added.
 - **Email Address**—Displays the email address associated with this user.
 - **Description**—Displays a description about the user if available.

Adding a user

In the DR Series system GUI you can easily add users and assign them specific system roles. The system supports up to 64 users.

To add a user, complete the following steps.

1. In the left navigation menu, click System Configuration→Users.
2. On the Action menu in the upper right corner of the page, click **Add User**. The Add pane is displayed.
3. For User Role, select from the following options.

 **NOTE:** You can select more than one role for a user. However, NDMP, OST, and iSCSI users are exclusive user roles that can only be assigned to one user at a time.

 - **Email Recipient**—Enables the user to receive email notifications and alerts from the system. For more information about configuring the types of notifications and alerts the system sends, see the topic, “Configuring email notification settings.”
 - **Monitor**—Limits the user to read-only access in the DR Series system GUI.
 - **CIFS**—Designates the user as a CIFS protocol user.
 - **RDA**—Designates the user as an RDA protocol user.
 - **OST**—Designates the user as an OST protocol user.
 - **NDMP**—Designates the user as an NDMP protocol user.
 - **iSCSI**—Designates the user as an iSCSI protocol user.
4. Click **Next**.
5. Enter the following information.
 - **Username**—Enter a username between 6 and 32 characters. This setting is required.
 - **Password**—Enter a password between 8 and 16 characters. This setting is required.
 - **Confirm Password**—Re-enter the password. This setting is required.
 - **Full Name**—Enter a name for the user.
 - **Phone**—Enter a phone number for the user.
 - **Email Address**—Enter an email address for the user.
 - **Description**—Enter a description for the user.
6. Click **Save**.

Modifying a user

In the DR Series system GUI you can easily modify settings for an existing user.

To modify a user, complete the following steps.

1. In the left navigation menu, click **System Configuration→Users**.
2. In the list of users, select the user you want to modify, and, in the **Actions** column, click the **Edit** icon. The **Edit User** pane is displayed.
3. You can modify the following optional settings:
 - **Full Name**—Enter the user's full name. This name can be up to 64 characters and must start with an alphabetic character.
 - **Phone**—Enter the user's phone number.
 - **Email Address**—Enter the email address associated with this user.
 - **Description**—Enter a description about the user.
4. Click **Save**.

Changing a user password


To change a user's password for logging in to the DR Series system, including the administrator if you have proper permissions, complete the following steps:

1. In the left navigation menu, click **SystemConfiguration > Users**. The **Users** page is displayed.
2. In the list of users, select the user you want to modify, and, in the **Actions** column, click the **Change Password** icon (which looks like a key). The **Edit User** pane is displayed.
3. In the **Old password** field, type the current password for the user.
4. In the **New password** field, type the new password.
5. In **Confirm password**, retype the new password to confirm.
6. Click **Save**.

Deleting a user

To delete a user, complete the following steps.

1. In the left navigation menu, click **System Configuration→Users**.
2. In the list of users, select the user you want to delete, and, in the **Actions** column, click the **Delete** icon.

 **NOTE:** You cannot delete the **administrator** user.


3. When prompted to confirm, click **OK**.

Modifying password reset options

In the DR Series system GUI, you can configure password reset options, that is, to enable (or disable) users' ability to request to reset their passwords upon login. If enabled, the link, *Forgot Your Password?*, will appear on the logon page. When a user requests to reset their password, an email notification is sent to the administrator of the system.

To enable or disable the password reset option, follow these steps.

1. Select **System Configuration > Users**.
2. In the Action menu in the upper right corner of the page, click **Reset Password Setup**.
3. For the **Enable Reset Password** option, select **No** or **Yes** as appropriate to disable or enable the reset password feature.
4. If you selected **Yes** to enable this feature, enter the following information.
 - **Email Relay Hostname**—Your external mail server (that is, the relay host) fully qualified domain name (FQDN), hostname, or IP address.
 - **Administrator Email**—The email address for the administrator of the system.
5. Under **Password Reset Question 1** and **Password Reset Question 2**, enter security questions and associated answers.

 **NOTE:** You will need these answers to reset your DR Series password.
6. Click **Save**.

Configuring email notification settings

In the DR Series system GUI, you can configure the email notifications/system alerts that are sent to users who have been assigned the Email Recipient role. For more information about the Email Recipient role, see the topic, "Adding a User." You can also send a test email message during this configuration process to verify that the proper email recipients receive the proper system alert messages and notifications.

To configure email notification settings, complete the following steps:

1. Select **System Configuration > Users**.
2. In the **Action** menu in the upper right corner of the page, click **Email Notification**.
3. In the **Email Relay Hostname** field, enter an external mail server (aka relay host) fully qualified domain name (FQDN), hostname, or IP address..

This setting configures the external email relay host to serve your DR Series system if the network email system requires one. The email relay host is typically an external mail server that relays any email alerts from the DR Series system to each of the designated recipient email addresses.

4. Select **Yes** or **No** for the following options to define the type of notifications and alerts to be emailed to the Email Recipient user(s).
 - **Notify me of DR Series appliance alerts**—Sends an email notification for system alerts.
 - **Notify me of DR Series software updates**—Sends an email notification to the designated user when a new software update is released.
 - **Notify me of DR Series daily container statistics reports**—Sends an email notification daily that includes the following statistics about the DR Series system: OS version, service tag, model number, hostname, total storage space used, total storage space available, and container statistics for every container including: container name, bytes ingested, bytes processed, post dedupe bytes, dedupe savings, replication bytes replicated, replication bytes pending.
5. To send a test message, click **Send Test Email Message**.
6. Click **Save**.

Configuring share-level security for CIFS shares

The DR Series system supports setting up share-level permissions for CIFS shares using the standard Microsoft Windows administrative tool, Computer Management. Computer Management is a component that is built into the Microsoft Windows 7, Vista, and XP operating systems.

i NOTE: Any user that is part of BUILTIN\Administrators can edit ACLs on CIFS shares. The local DR Series system administrator is included in the BUILTIN\Administrators group. To add additional domain groups to the BUILTIN\Administrators group, you can use the Computer Manager tool on a Windows client to connect to the DR Series system as Domain administrator and add any groups you want. This capability allows users other than the Domain administrator to modify an ACL as needed.

This administrative tool lets you control access to shares and also configure read-only or read-write access to user groups or individual users within the Active Directory Service (ADS) when joined to an ADS domain.

To implement share-level security on a DR Series system that has been joined to an ADS domain, make sure that you have mapped a drive on the DR Series system using an account with DOMAIN\Administrator credentials (or by using an account that is equivalent to a domain administrator). For more information about joining to an ADS domain, see *Configuring Active Directory Settings*.

i NOTE: If you do not use an account with sufficient privileges, you will not be able to see the shares or you may experience other problems.

1. Click **Start > Control Panel > Administrative Tools > Computer Management**.
2. Click **Action > Connect to another computer...**
3. Click **Another computer**, type the hostname or IP address for this DR Series system, and click **OK**.
4. Click **System Tools**, and click **Shared folders**.
5. Click **Shares** to display a list of the shares managed by the DR Series system.
6. Right-click on the share of interest, and select **Properties**.
7. Click the **Share Permissions** tab in the specified share **Properties** page.

8. To remove existing access permissions to the share, or add additional groups or user that can access the share, complete the following:
 - To add access for a new group or user, click **Add...** to display the Select Users or Groups dialog.
 - Click **Object Types...**, choose the object types you want to select (**Built-in security principals**, **Groups**, or **Users**), and click **OK**.
 - Click **Locations...** and define the root location from which to begin your search, and click **OK**.
 - In the **Enter the object names to select** list box, enter any object name(s) you want to find.

i **NOTE:** You can search for multiple objects by separating each name with a semicolon, and by using one of the following syntax examples: DisplayName, ObjectName, UserName, ObjectName@DomainName, or DomainName\ObjectName.
 - Click **Check Names** to locate all matching or similar object names that are listed in the **Enter the object names to select** list box, by using the object types and directory locations you selected.
9. Click **OK** to add the object to the **Group or user names** list box.
10. In the **Permissions** pane for the selected object, select the **Allow** or **Deny** check box to configure the following permissions:
 - Full Control
 - Change
 - Read
11. Click **OK** to save the selected share permission settings associated with the selected object.

Managing containers

In the DR Series system, data is stored in containers, which are stored in storage groups. Some containers function like a shared file system. These types of containers can be assigned a specific connection type, for example, NFS/CIFS or RDA (including both OST and RDS clients). These containers are then accessed via NFS, CIFS, and RDA protocols. You can also create virtual tape library (VTL) type containers, which can be accessed via NDMP, iSCSI, and Fibre Channel (which is available only on the DR4300 or DR6300 systems).

In the DR Series system GUI you can manage your storage groups and data containers, including viewing storage groups and containers, creating new storage groups and containers, modifying or deleting them, moving data into containers, and viewing current statistics.

i **NOTE:** If only the DefaultGroup storage group exists on your system, all containers you create are automatically added to that group. You can create custom storage groups, and then when you create a container, you can specify that it be added to the custom storage group. For more information about storage groups, see the topic, “Managing Storage Groups.”

Viewing containers

You can easily view a list of containers in your DR Series system. To view the list of containers, follow these steps.

1. In the left navigation menu, click **Containers**, and then select the storage group that has the containers you want to view. (If you only have the DefaultGroup storage group, you will not need to select a group.) The Containers page is displayed.

2. You can view the following columns of information about the containers.
- **Storage Group**—The storage group to which the container belongs.
 - **Container**—The name of the container.
 - **Marker Type**—The marker type that supports your Data Management Application (DMA).
 - **Access Protocol**—the connection type/access protocol for the container:
 - NFS
 - CIFS
 - RDA
 - NDMP
 - iSCSI
 - FC
 - **Connection Status**—The status of the connection as available or unavailable.
 - **Replication**—The current replication status as:
 - Not Configured
 - Stopped
 - Disconnected
 - Trying to Connect
 - Online
 - N/A
 - Marked for Deletion

i **NOTE:** For newly created OST or RDS containers, the Replication status displays **N/A**. When replication data has been deleted from an existing OST or RDS container, the Replication status also displays **N/A**. For existing containers that are in the process of deleting a large amount of data, the Replication status displays **Marked for Deletion** to indicate that the data deletion process has not yet completed.

Creating a container

By default, the DR Series system provides a container named **backup**, which is part of the storage group, DefaultGroup, for your use after you complete the basic system configuration and initialization process. You can create additional storage groups and containers to store your data as needed. For more information about storage groups, see [Managing storage groups](#)

i **NOTE:** The DR Series system does not support container names that begin with a number.

Containers can be accessed using the following connection types:

- **NFS**
- **CIFS**
- **NDMP** (for VTL type containers)
- **iSCSI** (for VTL type containers)

- **FC** (for VTL type containers on the DR4300 and DR6300 platforms)
- **RDA** (Rapid Data Access)
 - **OST** (OpenStorage Technology)
 - **RDS** (Rapid Data Storage)
- **No Access** (an unassigned connection type). Choosing the **No Access** or unassigned connection type lets you create containers that can be configured later as needed.

Adding an NFS or CIFS connection type container

To add an NFS or a CIFS connection type container, complete the following steps:

1. On the left navigation menu, select **Containers**, and then select or the storage group for which you want to create a container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the Action menu in the upper right corner of the page, click **Add Container**.
3. For **Storage Group**, select the storage group to which you want to add this container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
4. For Access Protocol, select NAS (NFS, CIFS).
5. For Container Name, type the name of the container, and then click **Next**.
Container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:
 - A-Z (uppercase letters)
 - a-z (lowercase letters)
 - 0-9 (numbers). Do not start a container name with a number.
 - dash (-) or underscore (_) special characters

i **NOTE:** The DR Series system does not support the use of the following special characters in container names: /, #, or @.
6. For **Access Protocols**, select **NFS** and **CIFS** as appropriate.
(Use NFS to back up UNIX or LINUX clients. Use CIFS to back up Windows clients.)

7. For Marker Type, select the appropriate marker that supports your Data Management Application (DMA).

- **None** — Disables marker detection for the container.
- **Auto** — Automatically detects CommVault, Tivoli Storage Manager (TSM), ARCserve, and HP Data Protector marker types. In addition, select this option if you need to support EMC Networker 2.0.
- **ARCserve**—Supports the ARCserve marker.
- **BridgeHead** — Supports the BridgeHead HDM marker.
- **CommVault**—Supports the CommVault marker.
- **HP DataProtector**—Supports the HP Data Protector marker.
- **Networker** — Supports EMC Networker 3.0. If you need to support EMC Networker 2.0, select Auto.
- **Time Navigator**—Supports the Time Navigator marker.
- **TSM**—Supports the TSM marker.
- **Unix Dump** — Supports the Amanda marker, among others.

Improper marker selection can result in non-optimal savings. As a best practice, if you have only one type of DMA with traffic directed to a container, it is best to select the marker type that supports your DMA (for example, **BridgeHead**, **Auto**, or another). Conversely, as a best practice, if you have traffic from a DMA that is not one of the supported marker types, it is best to disable marker detection for the container by selecting the **None** marker type.

8. Click **Next**.

9. If you selected NFS as the connection type, configure NFS access as follows.

- **NFS Options** — Defines the type of access to the container. Select one of the following options.
 - **Read Write Access** — To allow read-write access to the container.
 - **Read Only Access** — To allow read-only access.
- **Map Root To** — Select one of the following options from the drop-down list to define the user level you want mapped to this container.
 - **Nobody** — to specify a user on the system without root access permissions.
 - **Root** — to specify a remote user with root access to read, write, and access files on the system.
 - **Administrator** — to specify the system administrator.
- **Client Access** — Define the NFS client(s) that can access the NFS container or manage the clients that can access this container by selecting one of the following options.
 - **Open (allow all clients)** — To allow open access for all clients to the NFS container you create. (Select this option *only* if you want to enable access for all clients to this NFS container.)
 - **Create Client Access List** — To define specific clients that can access the NFS container. In the Client FQDN or IP text box, type the IP address (or FQDN hostname) and click the Add icon. The “added” client appears in the Allow Clients list box. (To delete an existing client from this list box, select the IP address (or FQDN hostname) of the client you want to delete, and click the Delete icon. The “deleted” client disappears from the list box.)

10. If you selected CIFS as the connection type, configure CIFS access as follows.
 - **Client Access** — Define the CIFS client(s) that can access the container or manage the clients that can access this container by selecting one of the following options.
 - **Open (allow all clients)** — To allow open access for all clients to the container you create. (Select this option *only* if you want to enable access for all clients to this container.)
 - **Create Client Access List** — To define specific clients that can access the container. In the Client FQDN or IP text box, type the IP address (or FQDN hostname) and click the Add icon. The “added” client appears in the Allow Clients list box. (To delete an existing client from this list box, select the IP address (or FQDN hostname) of the client you want to delete, and click the Delete icon. The “deleted” client disappears from the list box.)

i NOTE: The DR Series system administrator that manages the system has a different set of privileges than does the CIFS administrator user. Only the DR Series system administrator can change the password for the CIFS administrator user. To change the password that allows access for the CIFS administrator user, use the `authenticate --set --user administrator` commands. For more information, see the *DR Series System Command Line Reference Guide*.
11. Click **Next**.

A Configuration Summary of the options you selected for creating the container appears.
12. Click **Save**.

Moving data into a CIFS type container

To move data into an existing CIFS type container, complete the following steps.

1. Click **Start > Windows Explorer > Network**.
2. In the browser **Address bar**, click **Network** to select your DR Series hostname or IP address.

i NOTE: However, if your DR Series system is not listed, you can enter its hostname or IP Address preceded by "https://" and followed by the container name in the **Address bar** to access it (for example in this format, `https://10.10.20.20/container-1`). The DR Series system only supports the Hypertext Transfer Protocol Secure (HTTPS) form of IP addressing.
3. Move data from the source location to the destination container using your regular DMA or backup application process.

i NOTE: If any file ingested by the DR Series system by a DMA or backup application is renamed or deleted without using the DMA or backup application's process, the corresponding catalog must be updated accordingly. Failure to do so may prevent the DMA or backup application from being able to access the data.
4. Verify that the data recently moved now resides in the destination container.

Creating an OST or RDS connection type container

To create an OST or RDS connection type container, follow these steps:

1. On the left navigation menu, click **Containers**, and then select the storage group for which you want to create a container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
 2. In the Action menu in the upper right corner of the page, click **Add Container**.
 3. For **Storage Group**, select the storage group to which you want to add this container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
 4. For **Access Protocol**, select **Rapid Data Storage (RDS)** or **OpenStorage (OST)** as appropriate.
 5. For Container Name, type the name of the container, and then click **Next**.
Container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:
 - A-Z (uppercase letters)
 - a-z (lowercase letters)
 - 0-9 (numbers). Do not start a container name with a number.
 - dash (-) or underscore (_) special characters
- i** | **NOTE:** The DR Series system does not support the use of the following special characters in container names: /, #, or @.
6. If you selected Symantec OpenStorage (OST), for **LSU Capacity**, select one of the following options allowed per container:
 - **Unlimited** — To define the allowed amount of incoming raw data per container (based on the physical capacity of the container). If you selected RDS, by default, Unlimited is selected.
 - **Quota**: To define a set limit in Gibibytes (GiB) for incoming raw data allowed per container.
 7. Click **Next**.
 8. Click **Save**.

Creating a VTL type container

To create a virtual tape library (VTL) type container, complete the following steps.

i | **NOTE:** Only four VTL type containers can be created on a DR Series system.

1. On the left navigation menu, select **Containers**, and then select the storage group for which you want to create a container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the Action menu in the upper right corner of the page, click **Add Container**.
3. For **Storage Group**, select the storage group to which you want to add this container. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
4. For Access Protocol, select Virtual Tape Library (VTL).

5. For Container Name, type the name of the container.

i **NOTE:** The DR Series system does not support spaces or the following special characters in container names: /, #, or @. VTL container names cannot exceed 32 characters in length, must start with a letter, and can be composed of any combination of the following characters:

- A-Z (uppercase letters)
- a-z (lowercase letters)
- 0-9 (numbers). (Do not start a container name with a number.)
- underscore (_) special characters
- hyphen (-) special character

i **NOTE:** iSCSI VTL containers do not support the following characters:

- ASCII CONTROL CHARACTERS and SPACE through ,
- ASCII /
- ASCII ; through @
- ASCII [through `
- ASCII { through DEL

6. Click **Next**.

7. For Robot Model, select the type of virtual tape library for the VTL container.

- STK L700—This is the standard emulation of the StorageTek L700 library.
- DELL DR_L700—This is the Dell OEM version of the StorageTek L700 library.
- QUEST DR_L700 - This is a Quest OEM version of StorageTek L700 library.

i **NOTE:** The Quest and Dell OEM versions of the VTL are supported only with Symantec Backup Exec and Netbackup data management applications (DMAs).

8. For Tape Size, select the size of the tapes for your tape library from one of the following options.

- 800 GB
- 400 GB
- 200 GB
- 100 GB
- 50 GB
- 10 GB

i **NOTE:** Creating a VTL container type creates a tape library of type Storage Tek L700 with 10 tape drives of type IBM Ultrium LTO-4 and 60 tape slots holding 60 tapes. Additional tapes can be added as required. For more information, see the topic later in this guide, "VTL and DR Series Specifications."

9. For VTL Access Protocol, select one of the following options.

- NDMP
- iSCSI
- FC



NOTE: This protocol is only available for the DR4300 or DR6300 systems.

- No Access.



NOTE: Select this option if you are not ready to select a protocol.



NOTE: The DR Series system allows you to create a VTL container type without configuring it with a specific protocol (that is, by selecting No Access). You can configure the container at a later date.

10. For Access Control, do one of the following:

- If you selected NDMP as the access protocol, type the DMA's FQDN or IP address that will access the VTL container.
- If you selected iSCSI as the access protocol, type the FQDN, IQN, or IP address of the iSCSI initiator that can access the VTL container.
- If you selected FC as the access protocol, type the port initiator WWPN(s) for accessing the VTL container.

11. If you selected NDMP as the access protocol, for Marker Type, select the appropriate marker that supports your DMA from one of the following options.

- **None** — Disables marker detection for the container.
- **Unix Dump** — Supports the Amanda marker, among others.

12. If you selected iSCSI as the access protocol, for Marker Type, select the appropriate marker that supports your DMA from one of the following options.

- **None** — Disables marker detection for the container.
- **Auto** — Automatically detects CommVault, Tivoli Storage Manager (TSM), ARCserve, and HP Data Protector marker types. In addition, select this option if you need to support EMC Networker 2.0.
- **ARCserve**—Supports the ARCserve marker.
- **BridgeHead** — Supports the BridgeHead HDM marker.
- **CommVault**—Supports the CommVault marker.
- **HP DataProtector**—Supports the HP Data Protector marker.
- **Networker** — Supports EMC Networker 3.0. If you need to support EMC Networker 2.0, select Auto.
- **TSM**—Supports the TSM marker.
- **Unix Dump** — Supports the Amanda marker, among others.

i **NOTE:** Improper marker selection can result in non-optimal savings. As a best practice, if you have only one type of DMA with traffic directed to a container, it is best to select the marker type that supports your DMA. Conversely, as a best practice, if you have traffic from a DMA that is not one of the supported marker types, it is best to disable marker detection for the container by selecting the None marker type.

13. Click **Next**.

A Configuration Summary of the options you selected for creating the container appears.

14. Click **Save**.

You can add additional tapes to the library after container creation by [editing the container in the GUI](#) or by using the CLI command:

```
vtl --update_carts --name <name> --add --no_of_tapes <number>
```

For Fibre Channel, after a VTL is created with the default of 10 tape drives, you can add or delete drives in the library by using the following CLI command:

```
container --update --name <container name> --no_of_drives <1-60>
```

i **NOTE:** For more information about using the command line interface, see the *DR Series Command Line Reference Guide*.

Viewing VTL tape information

Once you have created a virtual tape library (VTL) type container, you can view the detailed tape information of the VTL. This includes information about the vendor and model information for medium changer and tape drives. To view VTL information, complete the following steps.

1. On the Containers page in the list of containers, select the VTL container for which you want to view detailed information.
2. Click the Tape icon.

3. You can view the following information.

- Library Slot and Tape Cartridges
 - Total Tapes Available
 - Tape Size
 - Max number of Tapes
 - Index
 - Slots
- Library Devices
 - Robot Model
 - Number of Tape Drives
 - Type
 - Vendor
 - Model
 - Serial
 - Info

Editing container settings

To modify the settings of an existing container, complete the following steps:

1. On the left navigation menu, click **Containers**, and then select the storage group that contains the container you want to modify. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the list of containers, select the container you want to modify, and then click the **Edit** icon.
3. Modify the marker and connection type options for the selected container as needed.
 - For information about NFS or CIFS connection type container settings, see the topic, [Creating an NFS or CIFS Connection Type Container](#), and make the corresponding changes.
 - For information about VTL container type settings, see the options available in the topic, [Creating a VTL Type Container](#), and make the corresponding changes.
 - For information about OST or RDS connection type container settings, see the options available in the topic, [Creating an OST or RDS Connection Type Container](#), and make the corresponding changes.

CAUTION: If you are changing the marker type on a DR6000 system and you are using Rapid CIFS, you must remount the share on the client after you change the marker type.

NOTE: The DR Series system administrator who manages the DR Series system has a different set of privileges than the CIFS administrator user. Only the DR Series system administrator can change the password for the CIFS administrator user. To change the password that allows access for the CIFS administrator user, use the DR Series system CLI `authenticate --set --user administrator` command. For more information, see the *DR Series System Command Line Reference Guide*.

4. After the container type settings have been modified, click **Save**.

Deleting a container

To delete an existing container that contains data, complete the following steps.

CAUTION: Before deleting a container, you should first carefully consider whether you need to preserve the data in the container. Before deleting any DR Series container that contains deduplicated data, you should take steps to preserve this data using another means of long-term retention. Once a container is deleted, the deduplicated data cannot be retrieved.

1. On the left navigation menu, click **Containers**, and then select the storage group that contains the container you want to delete. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the list of containers, select the container you want to delete, and then click the Delete icon.
3. In the Warning dialog box, click **Yes** to confirm the deletion.

Viewing container statistics

In the DR Series system GUI, you can view statistics about a selected container. All statistics displayed represent specific information about the backup data, throughput, replication, marker type, and connection type for the selected container. The displayed statistics will vary depending upon the connection type used by the specified container.

To display container statistics for a selected container, complete the following steps.

1. On the left navigation menu, click **Containers**, and then select the storage group that contains the container you want to view. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the list of containers, select the container for which you want to view statistics, and then click the Chart icon.
3. In the Active and Throughput charts, you can view current statistics for the container.
The Active chart displays the number of active files ingested based on time (in minutes), and the number of active bytes ingested based on time (in minutes). The Throughput chart displays the number of read data in Mebibytes/second (MiB/s) based on time (in minutes), and the number of write data in MiB/s based on time (in minutes).

NOTE: The values in the Active and Throughput charts refresh automatically every 15 seconds.

4. In the Backup Data and Throughput panes, click Zoom to select the duration period you want to display.
5. You can view the following additional information about the container.
 - **Marker Type**—the marker type associated with the container.
 - **Connection Type**—the configured connection type for the selected container, which can be NFS, CIFS, NDMP, iSCSI, FC, RDS, or OST. The type of connection information displayed for the container can vary based on connection type.
 - **Total Files**—the total number of files in the selected container.

6. In the Connection information pane, you can view information about the configured connection type for the selected container which can be NFS, CIFS, NDMP, iSCSI, FC, RDS, or OST. The type of information displayed can be different depending on the connection type. For example, the following information is displayed:
 - **NFS Connection Configuration pane**—NFS access path, Client Access, NFS Options, Map root to, and NFS Write Accelerator (DR6000 only).
 - **CIFS Connection Configuration pane**—CIFS share path, Client Access, and CIFS Write Accelerator (DR6000 only).
 - If the container is an RDA connection type container, the Connection Type OST pane or Connection Type RDS pane displays the following information:
 - **Capacity** — displays a Capacity pane with Status, Capacity, Capacity Used, and Total Images
 - **Duplication** — displays a Duplication Statistics pane with Inbound and Outbound statistics in the following categories: Bytes Copied (logical), Bytes Transferred (actual), Network Bandwidth Settings, Current Count of Active Files, and Replication Errors.
 - **Client Statistics** — displays a Client Statistics pane with Images Ingested, Images Complete, Images Incomplete, Images Restored, Bytes Restored, Image Restore Errors, Image Ingest Errors, Bytes Ingested, Bytes Transferred, and Network Savings.

Displaying container statistics by using the CLI

An alternate method for viewing container statistics is by using the DR Series system CLI command: `stats --container --name <container name>`

This command shows the following information:

- Container Name (name of the container)
- Container ID (ID associated with container)
- Total Inodes (total number of data structures in container)
- Read Throughput (read throughput rate in Megabytes or MiB/s for container)
- Write Throughput (write throughput rate in MiB/s for container)
- Current Files (current number of files in container)
- Current Bytes (current number of ingested bytes in container)
- Cleaner Status (current space reclamation process status for the selected container)

For more information on DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Managing replications

In the DR Series system GUI, you can set up and manage data replication operations. Such replication operations include, creating new replication relationships, managing or deleting existing replication relationships, starting and stopping replication, and displaying current replication statistics.

Guidelines and prerequisites for replication

Refer to the following important notes and guidelines for understanding and using replication in the DR Series system.

- **TCP Port Configuration**—If you plan to perform replication operations across a firewall, the DR Series system replication service requires that the following fixed TCP ports be configured to support replication operations:
 - port 9904
 - port 9911
 - port 9915
 - port 9916
- **DMAs and Domain Relationships** — To allow replication storage information to be viewed by a corresponding data management application (DMA), the target DR Series system must reside in the same domain as the source DR Series system in the replication relationship.
- **Replication Limits** — Refer to the *DR Series System Interoperability Guide* for details about the supported system limits for replication per DR Series system model. For a definition of connections and streams, see [Streams_vs_Connections](#).
- **Version Checking** — The DR Series system software includes version checking that limits replication only between other DR Series systems that run the same system software release version. If versions are incompatible, the administrator will be notified by an event, and replication will not continue.

- **Storage Capacity and Number of Source Systems** — Be aware that the storage capacity of the target DR Series system is directly affected by the number of source systems writing to its containers, and also by the amount being written by each of these source systems.
- **MTU Setting** — Primary and secondary replication targets should have the same network maximum transmission unit (MTU) setting. See the topic, “Configuring Network Settings,” for more information about this setting.

i NOTE: Rapid Air Gap is a replication feature in which a secondary target DR Series system in a three-way replication scenario exists in an isolated private network. This solution is available for use with NetVault Backup 11.4.5. The DR Series system secondary target can be hosted with limited set of services. All other services are closed, securing the DR Series system appliance from the external environment. The Secondary target is available for replication for only a specific period of time (when the Air gap closes). And, only during this period of time will replication between the Primary and Secondary target DR systems occur. Commands for using this feature are available in the DR Series system command line interface (CLI). For information about setting up and using the Rapid Air Gap feature, see the *DR Series Air Gap Setup Guide* and the *DR Series System Command Line Reference Guide*.

Viewing replication information

In the DR Series system GUI, the Replication page displays current information about replication relationships for data containers in your DR Series system. To view replication information, follow these steps.

1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to view replication information. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.) The Replications page is displayed.
2. You can view the following columns of information on this page.
 - **Source**—The name of the source container (with IP address or hostname)
 - **Status**—The status of the source container.
 - **Replica**—The name of the target container in the replication process (with IP address or hostname)
 - **Cascaded Replica**—The name of the remote replication container (with IP address or hostname)

3. You can click a selected replication in the list to expand to view the following detailed replication information about the source and target:
- Peer State—Displayed as Online, Offline, Paused, or Disconnected. When replication is started, the Peer State displays the status as Online for the selected container. When stopped, the Peer State initially displays the status as Paused, and then changes to Offline.
 - Peer Bandwidth—Shown as Kibibytes per second (KiB/s), Mebibytes per second (MiB/s), Gibibytes per second (GiB/s), or default (an unlimited bandwidth setting).
 - State—The current peer status as In sync, Paused, or Replicating.
 - Replication Average Transfer Rate
 - Replication Peak Transfer Rate
 - Network Average Transfer Rate
 - Network Peak Transfer Rate
 - Network Bytes Sent
 - Network Bytes Pending
 - Estimated Time to Sync
 - Dedupe Network Savings
 - Compression Network Savings
 - Last INSYNC Time—The last time the system synchronization occurred.
 - Time Until Scheduled Run—The time until the next scheduled run or the value, "In window," meaning the replication is currently in the scheduled window.

i | **NOTE:** These statistics refresh every 30 seconds.

Viewing replication statistics by using the CLI

In addition to using the DR Series system GUI to view replication information, you can also view statistics for a specific replication container by using the DR Series system CLI command: `stats --replication --name <container name>`

This command displays the following statistics.

- Container Name (name of the replication container)
- Replication Source Container (name that identifies the data source)
- Replication Source System (IP address or host name of the data source)
- Peer Status (current status of replication peer; for example, paused)
- Replication State (current state of replication relationship; for example, insync)
- Schedule Status (current status in days, hours, minutes, seconds)
- Replication Average Throughput (in Kibibytes per second, KiB/s)
- Replication Maximum Throughput (in KiB/s)
- Network Average Throughput (average throughput rate in KiB/s)
- Network Maximum Throughput (maximum throughput rate in KiB/s)

- Network Bytes Sent (total network bytes sent in Mebibytes/MiB)
- Dedupe Network Savings (total deduplication network savings in percentage)
- Compression Network Savings (total compression network savings in percentage)
- Last INSYNC Time (date of last sync operation in yyyy-mm-dd hh:mm:ss format)
- Estimated time to sync (time until next sync operation in days, hours, minutes, and seconds)

Data replication history is also displayed on a file-by-file basis, with a replication timestamp, and other file related information.

For more information about DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Adding replication relationships

To add a new replication relationship, complete the following steps.

1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to set up replication. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. In the Actions menu in the upper right corner of the page, click **Add Replication**. The Add Replication pane is displayed.
3. Next to Choose Replication Type, select one of the following options.
 - **Replica Only**
 - **Replica & Cascade**
4. Click **Next**.
5. To define the source container, for Select Container Location, select the **Local** or **Remote** option.
 - If you select Local, select the local container from the drop-down list.
 - If you select Remote, configure the following settings:
 - Username—enter the username for the remote system.
 - Password—enter the password for the remote system.
 - Remote System—enter the domain name of the remote system.
 - Click **Retrieve Remote Containers**.
 - Select Remote Container—Select the remote container from the drop-down list.
6. For **Encryption**, select one of the following encryption options to encrypt the data as it is replicated: Not Enabled, AES 128-bit, or AES 256-bit.

7. Under **Replica Container**, define the target replica container by configuring the following settings.
 - Username—enter the username for the remote system.
 - Password—enter the password for the remote system.
 - Remote System—enter the domain name of the remote system.
 - Click **Retrieve Remote Containers**.
 - Select Remote Container—Select the remote container from the drop-down list.
8. If you specified to set up cascaded replication, define the cascaded replication by doing the following.
 - Under **Cascaded Replica Container**, click **Select a container from the remote system** to select the container you will be using for the cascaded replica.
 - Enter the logon credentials of the remote system.
 - Click the **Retrieve Remote Container** button, and, in the drop-down list, select a remote container from the list of available containers.
 - For **Encryption**, select one of the following encryption options: Not Enabled, 128-bit, or 256-bit.
9. Click **Save**.

i **NOTE:** For information about starting and stopping replication, see the topic, "[Starting and Stopping Replication](#)." For information about scheduling system operations such as replication, see "[Understanding system operation scheduling](#)."

Modifying replication relationships

You can modify the following replication settings: encryption and remote container's IP address/host name settings. To modify settings for an existing replication relationship, complete the following steps.

! **CAUTION:** You should exercise caution when configuring the direction of replication for source and target containers. For example, target containers can have their contents deleted if they contain existing data.

i **NOTE:** Because you cannot modify an existing defined role (source or target replica) for a replication relationship, if necessary, you must delete the existing replication relationship, and then recreate a new relationship with the specific source and target roles that you want.

1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to modify replication. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. From the list, select the replication relationship that you want to modify, and click to expand the details.
3. Click the Edit icon and then, in the Edit Replication pane, modify the settings/values for the Source, Replica, or Cascaded Replica containers as needed.
 - a. For Remote System, modify the IP address/host name and user logon credentials of the source remote system as needed.
 - b. Review the replication details, and then click **Save**.
4. Click **Save**.

Deleting replication relationships

To delete an existing replication relationship, complete the following steps:


1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to delete a replication. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. From the list, select the replication relationship that you want to delete, and click to expand the details.
3. Click the Delete icon and then, in the confirmation dialog box, click **Delete**.

 **NOTE:** If the deletion fails, you can click **Force Delete** to force removal of the relationship.

Starting and stopping replication

To start or stop replication in an existing replication relationship, complete the following steps.

1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to stop or start replication. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. From the list, select the replication relationship that you want to start or stop, and click to expand the details.
3. To stop the replication process, click the **Stop** icon, and, in the confirmation dialog box, click **Yes** to stop replication.
4. To start the replication process, click the **Start** icon, and, in the confirmation dialog box, click **Yes** to start replication.

 **NOTE:** You can also set up replication schedules as needed. For more information see the topic, [Configuring Replication Schedules](#).

Adding a cascaded replica

To add a cascaded replica to an existing replication relationship, complete the following steps.

1. In the left navigation menu, click **Replications**, and then select the storage group that contains the container for which you want to add a cascaded replica. (If you are only using the DefaultGroup storage group in your DR Series system, you do not need to select a group.)
2. From the list, select the replication relationship to which you want to add a cascaded replica, and click to expand the details.
3. Under the **Replica Status** column, click the **Edit** icon.
4. In the Edit Replication pane, select one of the following Encryption options: **None**, **128-bit**, or **256-bit**, and then click **Next** to enter Cascaded Container details.
5. Enter the IP address and the logon credentials of the remote system.

6. Click **Next**.
7. Review the replication details and click **Save**.

Monitoring the DR Series system

This topic describes how you can monitor the current state of DR Series system operations on the **Dashboard** page as well as the Dashboard sub pages for Alerts, Clients, Events, Health, and Usage.

i **NOTE:** If you are using the virtual DR Series system, DR2000v, you might not see all of the monitoring options available as described in this section. For details on using and monitoring the DR2000v, see the *DR2000v Deployment Guide*.

Using the Dashboard page

The Dashboard page contains graphics that show key information about the current state of the DR Series system. This page automatically refreshes every 30 seconds.

To use the Dashboard page, follow these steps.

1. Click **Dashboard** in the left navigation menu of the DR Series system GUI.

i **NOTE:** When you first log on to the DR Series system GUI, you view the Dashboard page by default.

2. You can view the following graphs:
 - **Capacity**—displays total used space, free space, and used and encrypted space in GBs and TBs.
 - **Storage Savings**—displays a total savings in percentage (combining both deduplication and compression) over a time period (for example, every hour, which is the default).
 - **Throughput**—displays the throughput volume (reads and writes) in Mebibytes/second (MiB/s) based on time (for example, every hour, which is the default).
 - **System**—displays information about memory and CPU usage.

3. At the bottom of the Dashboard page, you can also view the System Summary section, which lists key information about the current DR Series system, including:
 - **Total number of files in all containers**
 - **Number of containers**
 - **Number of containers replicated**
 - **Active bytes**—the total bytes before optimization.
 - **Advanced data protection**—the status of the data integrity check.
 - **Cleaner status**—The current cleaner status as one of the following states:
 - **Pending**—displayed when there is any scheduled window set and the current time is outside the scheduled window for the Cleaner operation.
 - **Running**—displayed when the Cleaner operation is running during a scheduled window.
 - **Idle**—displayed only if there is no Cleaner operation running during a scheduled window.
4. To change the time display in the graphs, click **Zoom** within the graph you want to view, and then select the time increment as hour, day, week, month, or year.
5. To quickly view the state of the system, click the System Health Indicator icon, in the upper right corner of the DR Series system GUI. This icon appears green or red depending on the system state and displays the following information:
 - **System State**—displays the status value optimal, warning (a non-critical error has occurred), or actionable state (in which a critical error is detected.) To view more information about the System State, go to the **Dashboard > Usage** page, which displays current detailed status information for the system.
 - **Hardware State**—displays the status value optimal, warning (a non-critical error has occurred), or actionable state (in which a critical error is detected.) To view more information about the Hardware State, go to the **Dashboard > Health** page, which displays current detailed status information for the hardware and expansion shelf enclosures (if installed).
 - **Alerts**—displays the number of alerts that have occurred. You can click this number to view the alert(s).
 - **Events**—displays the number of events that have occurred. You can click this number to view the event(s).

Viewing DR Series system statistics by using the CLI

An alternate method for viewing the current DR Series system statistics is by using the DR Series system CLI command: `stats --system`. This command shows the following categories of system statistics:

- Capacity Used (system capacity used in Gibibytes or GiBs)
- Capacity Free (system capacity free in GiBs)
- Read Throughput (read throughput rate in Mebibytes or MiB/s)
- Write Throughput (write throughput rate in MiB/s)
- Current Files (current number of files in system)

- Current Bytes (current number of ingested bytes in system)
- Post Dedupe Bytes (number of bytes after deduplication)
- Post Compression Bytes (number of bytes after compression)
- Post Encryption Bytes
- Post Encryption Bytes in GiB
- Compression Status (current compression status)
- Cleaner Status (current space reclamation process status)
- Encryption Status
- Total Inodes (total number of data structures)
- Bytes decrypted
- Dedupe Savings (deduplication storage savings by percentage)
- Compression Savings (compression storage savings by percentage)
- Total Savings (total storage savings by percentage)

For more information on DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Monitoring system alerts

You can easily view current system alerts in the DR Series system GUI.

- To monitor DR Series system alerts, on the left navigation menu, click **Dashboard > Alerts**. The Alerts page displays a summary table of alerts listed by index number, timestamp of the system alert, and a brief message describing the alert.

i **NOTE:** For a detailed list of possibly occurring alerts, see the topic, “DR Series system alert and event messages,” in the “Support, maintenance, and troubleshooting” chapter of this guide.

Monitoring clients

You can easily view the current clients that are connection to the DR Series system. The Clients page displays a list of the clients that are connected to the DR Series system, which can include NFS, CIFS, RDA, OST, NDMP, iSCSI, and FC clients.

1. To view client information for the DR Series system, on the left navigation menu, click **Dashboard > Clients**. The total number of currently active clients for a particular type is displayed next to the client name in the individual client panes.

2. To view NFS client information, click the NFS pane to expand it to view the following details:
 - **Name**—the name of the NFS client.
 - **Idle Time**—the amount of idle time (non-activity) for the client.
 - **Connection Time**—the connection time for the client.
3. To view CIFS client information, click the CIFS pane to expand it to view the following details:
 - **Name**—the name of the CIFS client.
 - **Idle Time**—the amount of idle time (non-activity) for the client.
 - **Connection Time**—the connection time for the client.
4. To view RDS or OST client information, click the RDA pane to expand it to view the following details:
 - **Name** — the name of the RDS or OST client.
 - **IP Address**—the IP address of the client.
 - **Type**—the type of RDA client.
 - **Plug-in**—the plug-in type installed on the client.
 - **Backup Software**—the backup software running on the client.
 - **OS**—the operating system of the client.
 - **Idle Time**—the amount of idle time (non-activity) for the client.
 - **Connection**—the number of connections for the client.
 - **Mode**—the current mode type that is set for the client, such as:
 - **Auto:** The DR Series system sets the deduplication as either Dedupe or Passthrough, based on the client's number of cores and whether it is 32- or 64-bit.
 - **Passthrough:** The client passes all data to the DR Series system for deduplication processing (appliance-side deduplication).
 - **Dedupe:** The client processes hashing on data, so that deduplication processing occurs on the server side (client-side deduplication).
5. To view NDMP client information, click the NDMP pane to expand it to view the following details:
 - **ID** — the NDMP session ID.
 - **Duration** — The duration of the current active session.
 - **State** — The current status, for example, Active.
 - **Source** — IP address of the source filer.
 - **Target** — The target tape drive being used for the current NDMP session.
 - **Throughput** — The current and average throughput.
 - **Transfer size** — The total size of data transferred in this backup session.
 - **DMA** — The IP address of the Data Management Application initiating the backup.

The NDMP Completed Sessions Statistics section shows the above information for any completed NDMP sessions.

6. To view iSCSI client information, click the iSCSI pane to expand it to view the following details:
 - **Container Name**—The container name for the iSCSI VTL container.
 - **Container IQN**—the iSCSI Qualified Name for the iSCSI VTL container.
 - **Initiators Connected**—the initiators connected to the iSCSI VTL container.
7. To view iSCSI client information, click the iSCSI pane to expand it to view the following details:
 - **Container Name**—The container name for the iSCSI VTL container.
 - **Container IQN**—the iSCSI Qualified Name for the iSCSI VTL container.
 - **Initiators Connected**—the initiators connected to the iSCSI VTL container.
8. To view Fibre Channel client information, click the FC pane to expand it to view the following details:
 - **Container Name**—The container name of the FC VTL container.

Monitoring system events

You can easily view and filter current system events in the DR Series system GUI. To monitor DR Series system events, follow these steps:

1. On the left navigation menu, click **Dashboard > Events**.
The Events page displays a summary table of events listed by index number, severity (Informational, Warning, or Critical), timestamp of the event, and brief message describing the event.
2. To filter the list of events, expand the Filter pane, and do any of the following:
 - In the Message Contains field, enter a word or string of words for which to search in Event messages. (The system is not case-sensitive)
 - Click the Event Severity menu and select the severity level for which to search as All, Info, Warning, or Critical.
 - For Timestamp From and Timestamp To, select the date and time for which to search.
3. Click **Filter**. Search results appear in the Events summary table.

i | **NOTE:** For a detailed list of possibly occurring events, see [DR Series system alert and event messages](#) in the **Support, maintenance, and troubleshooting** chapter of this guide.

Monitoring system health

You can easily view and monitor the current status of your system hardware in the DR Series system GUI. The Health page displays graphics of the front and rear chassis of the DR Series system. This page also includes a list of hardware components and related system health information. To monitor DR Series system health, follow these steps:

1. On the left navigation menu, click **Dashboard > Health**.
2. In the System pane, you can click individual components in the front and rear chassis graphics to view system health information for the selected hardware component. The front chassis graphic shows the disk drive locations (slots 1–11). The rear chassis graphic shows the NIC ports and power supplies. If you have expansion enclosures installed, there will be tabs across the top of the page, one per enclosure, for you to view enclosure system health graphics.
3. You can view detailed information about the following hardware components in the list. Click to expand each component pane to view details. The status column displays a green check mark for when a component is in a good operational state. A red X with a circle around it is displayed if it is in an unhealthy state.
 - Power Supplies
 - Fans
 - Temperature
 - Storage
 - Voltage
 - NIC
 - CPU
 - DIMM
 - NVRAM
 - HBA

Monitoring system usage

You can easily view and monitor the current DR Series system usage in the DR Series system GUI. The Usage page displays the current system usage status. To monitor DR Series system usage, follow these steps.

1. On the left navigation menu, click **Dashboard > Usage**. The Usage page displays a graph showing overall read and write operations and throughput by date and time.
2. To filter the information that is displayed, expand the Filter pane, set the following parameters as needed, and then click **Filter**:
 - For the Filter Range, select **Latest Range** or **Time Range**.
 - If you select Latest Range, select a time period in the **Display last** fields.
 - If you select Time Range, select a date and time for the **Timestamp From** and **Timestamp To** fields.
 - Next to **Protocols**, select the protocols for which you want to view usage information.
 - Next to **Interfaces**, select the appropriate interfaces for the protocols you are viewing
 - Next to **Disks**, select Disk if you want to view Disk usage information.
 - Next to **Operations**, select the type of operations for which you want to view usage information.
 - Next to **Throughput**, select the type of throughput for which you want to view usage information.

3. In the Action menu in the upper right corner of the GUI, you can also click Detailed Graphs to view details for the following categories of usage.
 - CPU Load
 - System
 - Memory
 - Active Processes
 - Protocols
 - Network
 - Disk
 - Cleaner
 - All

Using GlobalView

This topic describes how to view, add, and navigate to DR Series systems in GlobalView, which provides a real-time view of multiple DR Series systems in the same domain.

- To view the GlobalView page, in the left navigation menu, click **GlobalView**.

The GlobalView page is displayed showing the GlobalView Summary and a list of appliances in your GlobalView.

Understanding GlobalView

GlobalView is a dashboard view providing a holistic, or “global,” picture of all of the DR Series systems that you have added to it. By using GlobalView, you can easily monitor and manage all of the DR Series systems in your enterprise through one view. For example, as an administrator in a headquarters office, you could use GlobalView on your DR Series system to monitor via a single page all of the DR Series systems you have in the different branch offices that replicate to the headquarters office. The GlobalView also provides easy navigation to any DR Series systems in the view.

The following list describes important tips and considerations for using GlobalView:

- For streamlined navigation, your location in the GUI is saved when you navigate between DR Series systems. For example, if you are on the Software Upgrade page in one DR Series system, and you navigate to another DR Series system from the GlobalView page, the Software Upgrade page of the new DR Series system is displayed.
- The GlobalView on a DR Series system is local to that system, and the GlobalView information is maintained in a physical file on the system. If the machine goes down or is otherwise unavailable, the GlobalView is unavailable. In addition, if a factory refresh is performed on the system, the GlobalView information is lost and you must add the systems to the GlobalView dashboard again.

- You can define an identical GlobalView on another DR Series system in your domain to serve as a backup if the DR Series system that contains the original GlobalView is down or otherwise unavailable. For example, suppose you have three DR Series systems: A, B, and C. All of these systems are on the same Active Directory Services (ADS) domain and have identical logon credentials. You log on to DR Series system A, and, on its GlobalView page, you add DR Series systems B and C (resulting in A, B, and C being in the view). Then, you can log on to DR Series system B, and add A and C to its GlobalView page (also resulting in A, B, and C being in the view).
- You cannot import or export a GlobalView dashboard configuration. To create a GlobalView, you must manually define it by adding systems to the GlobalView dashboard. For details, see [Adding a DR Series System to Global View](#).
- The virtual system, DR2000v, can be monitored in GlobalView by the hardware—based DR Series system to which it is registered.

i NOTE: If you are using Internet Explorer 10, it is recommended that you disable the pop-up blocker so that you can open a new browser window when you navigate to another DR Series system from within GlobalView.

Prerequisites for GlobalView

The GlobalView feature is available on all DR Series systems that have version 3.0.0.1 (or later) software installed. The system to which you are currently logged in is automatically included by default in the GlobalView page; however, any other systems must be explicitly added. For details, see [Adding a DR Series System to Global View](#).

The following list describes the prerequisites that must be met so that you can add and view additional DR Series systems in the GlobalView page.

- All DR Series systems must have the same version of 3.x or 4.x software installed. Systems running older software versions cannot be added to the GlobalView page.
- All DR Series systems must have unique service tags.
- When you use GlobalView, you must log on to the DR Series system by using your domain credentials; for example, you must log in as [DOMAIN]Administrator instead of just Administrator.


Viewing and using the GlobalView page

The GlobalView page in the DR Series system GUI displays a convenient view of the operating statistics for all of the DR Series system appliances that you have added. On this page, you can monitor the status of and easily navigate to the DR Series systems that you have added to the GlobalView. Using GlobalView makes it easy to navigate to a different system in your enterprise without having to log out and log on by using new browser sessions.

To view and use the GlobalView page, follow these steps.

1. To view the GlobalView page, in the left navigation menu, click **GlobalView**. The GlobalView page is displayed, showing a GlobalView Summary and a list of appliances that have been added to GlobalView. The information on this page refreshes every 30 seconds.

2. You can easily view summary information in the GlobalView Summary pane.


 **NOTE:** If needed, you can expand the GlobalView Summary pane to view the information.

The following table describes the information that is displayed in the GlobalView Summary.





Column/Item	Description
Appliances	
Configured	Displays the number of appliances that have been added to the GlobalView (including the managing system, which contains the GlobalView dashboard).
Operational	Displays the number of appliances that are currently connected and are operational.
Disconnected	Displays the number of appliances that have been added to the GlobalView but are unable to be reached. To troubleshoot, see the topic, Reconnecting DR Series Systems .
Upgrade in Progress	Displays the number of appliances that have system software currently being upgraded.
Rebooting	Displays the number of appliances that are rebooting.
Initializing	Displays the number of appliances that are currently initializing.
Maintenance Mode	Displays the number of appliances currently in maintenance mode.
Manual Intervention Required	Displays the number of appliances that require manual intervention.
Notifications	
Alerts	Displays the total number of alerts in all appliances in the GlobalView.
Events	Displays the total number of events in all appliances in the GlobalView.
Capacity	
Total	Displays the total physical capacity in all appliances in the GlobalView.
Used	Displays the total physical capacity bytes that are used across all appliances in the GlobalView.
Free	Displays the total physical capacity bytes that are free across all appliances in the GlobalView.
Containers	
Containers	Displays the total number of containers in all appliances in the GlobalView.
Replications	Displays the total number of containers replicated in all appliances in the GlobalView.
Total Files	Displays the total number of files in all containers in all appliances in the GlobalView.
Active Bytes	Displays the total bytes before optimization in all appliances in the GlobalView.

Column/Item	Description
Active Clients	
RDS, NFS, CIFS, OST, NDMP, iSCSI, and FC	Displays the total number of clients configured in all appliances in the GlobalView, organized by container connection type, such as RDS, NFS, CIFS, OST, NDMP, iSCSI, and FC.

3. You can view the list of appliances and associated information under the GlobalView Summary. This appliance list includes all of the appliances in GlobalView and provides a high-level status. The managing DR Series system appears in bold text. By default, appliances are listed alphabetically by appliance name. You can sort the list by clicking the column header, which toggles between ascending and descending order. This sort order is retained if you leave the page and return later.

4.  **NOTE:** You can collapse the GlobalView summary pane for more viewing space on the page, if needed.

The following table describes the information displayed in the appliance list.

Column	Description
Appliance Name	Lists the Active Directory fully-qualified domain name (FQDN), and contains links to each respective DR Series system.
Status	Displays the system operational state as an icon. <ul style="list-style-type: none"> A green check mark icon  indicates that the system is operational. A red X icon  indicates that the system is not connected. This can occur, for example, if the DR Series system is removed from the Active Directory Services (ADS) domain, if it is down, or if it is being rebooted.
Capacity	Displays the used physical storage capacity as a percentage.
Savings	Displays the total savings as a percentage (combining both deduplication and compression).
Alerts	Displays the alert count. You can click the number to navigate to the Alerts page.
Replication	Displays the replication state as an icon. <ul style="list-style-type: none"> A green check mark icon  indicates that replication is operational. A red icon  indicates that replication has failed.
Ingest Rate	Displays the rate of data being written to the DR Series system across your network in MiB/s.

5. You can expand each appliance to view more specific configuration information about the system, including:
 - System State (such as Operational or Manual Intervention Required)
 - Model
 - Service Tag
 - Software Version
 - iDRAC IP
 - Management IP
 - Used Capacity (in bytes)
 - Free Capacity (in bytes)
 - Total Containers
 - Read Throughput (in MiB/s)
 - Configured Replications
 - Connected Replications
 - Disconnected Replications
 - Trying to Connect Replications
 - Stopped Replications
6. To navigate to a different DR Series system from GlobalView, do one of the following:
 - In the left navigation menu, click GlobalView and select the DR Series system you want to view.
 - In the appliance list in the GlobalView page, click the appliance name.
7. To go back to viewing the managing DR Series system, click the **Return to Managing Unit** link, which appears above the GlobalView menu item in the left navigation menu.

Adding a DR Series system to GlobalView

You can add up to 64 machines to GlobalView. This number includes the system to which you are logged on. Before you add a system to GlobalView, you must have logged onto the system by using your domain credentials.

To add a DR Series system to Global View, complete the following steps:

1. In the left navigation menu, click **GlobalView**.
2. On the **Global View** page, click the Action menu in the upper right corner, and then click **Add to Global View**.
The Add to Global View pane appears above the GlobalView Summary.
3. In the field, **DR Unit FQDN or IP address**, add the fully-qualified domain name (FQDN) or IP address of the DR Series system that you want to add.

4. In the **Username** field, enter the domain user name for the DR Series system that you want to add. For example, DOMAIN\administrator. This entry should be identical to the credentials used in all other systems in the GlobalView.
5. In the **Password** field, enter the domain password for the DR Series system that you want to add. This entry should be identical to the credentials used in all other systems in the GlobalView.
6. Click **Add and Connect**.


Removing a DR Series system from GlobalView

You can remove any DR Series system from GlobalView except the system to which you are currently logged on, which contains the GlobalView.

When you remove a DR Series system from GlobalView on one system, it does **not** remove it from any other GlobalViews to which you may have added it on other systems.

To remove a DR Series system from GlobalView, complete the following steps:

1. In the left navigation menu, click **GlobalView**.
2. On the **GlobalView** page, in the appliance list, click the **Delete** icon next to the system you want to delete.

 **NOTE:** No Delete icon appears next to the system that contains the GlobalView; it is not available to be removed.

A warning dialog box is displayed to confirm that you want to delete the system.

3. Click **OK** to confirm.

Configuring and using Rapid NFS and Rapid CIFS

Rapid NFS and Rapid CIFS enable write operation acceleration on clients that use NFS and CIFS file system protocols. Similar to OST and RDS, these accelerators allow for better coordination and integration between DR Series system backup, restore, and optimized duplication operations with Data Management Applications (DMAs) such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of supported DMAs, see the *DR Series System Interoperability Guide*.

Rapid NFS is a new client file system type that ensures that only unique data is written to the DR Series system. It uses user space components and file system in user space (FUSE) to accomplish this. Metadata operations such as file creates and permission changes go through the standard NFS protocol, whereas write operations go through Rapid NFS.

Rapid CIFS is a Windows-certified filter driver that also ensures that only unique data is written to the DR Series system. All chunking and hash computations are done at the client level.

i **NOTE:** The supported DMAs listed in the *DR Series System Interoperability Guide* are the DMAs that have been **tested and qualified** with Rapid NFS and Rapid CIFS. You can use Rapid NFS and Rapid CIFS with other DMAs (such as Veritas products), but those products have not been tested and qualified with Rapid NFS or Rapid CIFS.

Rapid NFS and Rapid CIFS benefits

When Rapid NFS and Rapid CIFS are used with the DR Series system, they offer the following benefits:

- Reduce network utilization and DMA backup time
 - Chunk data and perform hash computation on the client; transfer chunked hash files on the back-end
 - Reduce the amount of data that must be written across the wire
- Improve performance

- Support DMAs such as CommVault, EMC Networker, and Tivoli Storage Manager. For the current list of supported DMAs, see the *DR Series System Interoperability Guide*.
- Compatible with existing NFS and CIFS clients — just need to install a plug-in (driver) on the client
 - Can use Rapid NFS and Rapid CIFS to accelerate I/O operations on any client — including a client that uses home-grown backup scripts
 - Can service multiple and concurrent media server backups

Best practices: Rapid NFS

This topic introduces some recommended best practices for using Rapid NFS operations with the DR Series system.

- Containers must be of type NFS/CIFS

RDA containers cannot use Rapid NFS. If you have existing NFS/CIFS containers, you do not need to create new containers to use Rapid NFS; you can install the plug-in (driver) to existing clients.

- The Rapid NFS plug-in (driver) must be installed on client systems

After the plug-in is installed, write operations will go through Rapid NFS while metadata operations such as file creates and permission changes will go through the standard NFS protocol. Rapid NFS can be disabled by uninstalling the plug-in.

- Markers must be set on the client, not in the DR Series GUI
- If you are using a DMA that supports a marker, should explicitly set it. Your containers should have the marker type of None until you set the marker using the Mount command on the client (after installing the Rapid NFS plug-in).

- For existing containers, re-set the marker by doing the following:

For example, if you wanted to set the CommVault marker (cv):

```
mount -t rd nfs 10.222.322.190:/containers/backup /mnt/backup -o marker=cv
```

Mount command usage:

```
rd nfs [nfs mount point] [roach mount point] -o marker=[marker]
```

where:

nfs mount point = Already mounted nfs mountpoint

roach mount point = A new mount point

marker = appassure, arcserve, auto, cv, dump, hdm, hdpd, nw, or tsm

- Your DR Series system must meet the minimum configuration

Rapid NFS is available on a DR Series system and a client with a minimum of 4 CPU cores running at a minimum of 4 GHz cumulative processing power and 2 GB memory. Kernels must be 2.6.14 or later. For a list of supported operating systems, see the *DR Series System Interoperability Guide*. If you update your operating system, you must update your Rapid NFS plug-in as well. Updates are available on the Support site as well as within the GUI on the Clients page.

- Rapid NFS is stateful

If the DR Series system goes down, the connection will terminate. DMAs will restart from the last checkpoint.

- Rapid NFS and passthrough mode
If Rapid NFS mode fails for any reason, the DR Series system falls back to regular NFS mode automatically. For details, see [Monitoring Performance](#).
- Rapid NFS performance considerations
When using Rapid NFS on your client, Quest recommends that you do not run other protocols to the DR Series system in parallel, as this will adversely affect your overall performance.
- Rapid NFS acceleration constraints
 - Rapid NFS does not support:
 - Direct I/O memory
 - Mapped files
 - File path size greater than 4096 characters
 - File write locks across clients

i | **NOTE:** If the client and server do not have the same times, the times seen will not match typical NFS behavior due to the nature of file system in user space (FUSE).

Best practices: Rapid CIFS

This topic introduces some recommended best practices for using Rapid CIFS operations with the DR Series system.

- Containers must be of type NFS/CIFS
 - RDA containers cannot use Rapid CIFS. If you have existing NFS/CIFS containers, you do not need to create new containers to use Rapid CIFS; you can install the plug-in (driver) to existing clients.
- The Rapid CIFS plug-in (driver) must be installed on client systems
After the plug-in is installed, write operations will go through Rapid CIFS while metadata operations such as file creates and permission changes will go through the standard CIFS protocol. Rapid CIFS can be disabled by uninstalling the plug-in.
- Your DR Series system must meet the minimum configuration
Rapid CIFS is available with a DR Series system and a client with a minimum of 4 CPU cores running at a minimum of 4 GHz cumulative processing power and 2 GB memory. For a list of supported operating systems, see the *DR Series System Interoperability Guide*.
If you update your operating system, you must update your Rapid CIFS plug-in as well. Updates are available on the Support site as well as within the GUI on the Clients page.
- Rapid CIFS is stateful
If the DR Series system goes down, the connection will terminate. DMAs will restart from the last checkpoint.
- Rapid CIFS and passthrough mode
If Rapid CIFS mode fails for any reason, the DR Series system falls back to regular CIFS mode automatically.

- Rapid CIFS acceleration constraints
Rapid CIFS does not support:
 - NAS functionality
 - Optlocks (but supported if a single client is writing)
 - Byte-range locks
 - Optimization of very small files (less than 10 MB). File size can be adjusted using configuration settings.
 - FILE_NO_IMMEDIATE_BUFFERING and FILEWRITE_THROUGH operations (sent via CIFS only).
 - File path size greater than 4096 characters

Setting client-side optimization

Client-side optimization (also known as client-side deduplication) is a process that can contribute to saving time performing backup operations and reducing the data transfer overhead on the network.

To configure dedupe/passthrough before the client makes a connection, you need to use the DR Series system command line interface (CLI).

i NOTE: To update a client in the DR Series system GUI, the client must already be attached (and therefore enabled to appear in the GUI). If a client connection exists, you can select the radio button in the GUI to modify it.

You can turn client-side optimization On or Off by using the CLI commands, `rda --update_client --name --mode`. For more information about DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Installing the Rapid NFS plug-in

The DR Rapid NFS plug-in must be installed on to the media server type you choose (for supported operating systems and DMAs, see the *DR Series System Interoperability Guide*). The plug-in software enables integration between DR Series system data storage operations and the supported data management applications (DMAs). Before you install, make sure you adhere to the Best Practices covered in another topic in this chapter.

The plug-in must be installed on the designated Linux-based media server in the following directory, **/usr/openv/lib/**. The plug-in is installed using a self-extracting installer that installs the Rapid NFS plug-in and all of its related components. The installer supports the following modes, with the default being Help (-h):

- Help (-h)
- Install (-install)
- Upgrade (-upgrade)
- Uninstall (-uninstall)
- Force (-force)

```
$> ./QuestRapidNFS-xxxxx-xxxxx-x86_64.bin -help
Quest plug-in installer/uninstaller
usage: QuestRapidNFS-xxxxx-xxxxx-x86_64.bin [ -h ] [ -install ] [ -uninstall ]
```



```

-h                               : Displays help
-install                         : Installs the plug-in
-upgrade                        : Upgrades the plug-in
-uninstall                     : Uninstalls the plug-in
-force                          : Forces the installation of the plug-in

```

You can download the plug-in installer from the Quest website:

- Go to support.quest.com/dr-series, select your specific DR Series model, and then navigate to Software Downloads.
- Locate the DR Rapid NFS plug-in and download it to your system.

After it is downloaded, follow the steps that follow to run the Plug-In Installer to install the plug-in on your designated Linux-based media server.

i | **NOTE:** The plug-in needs to be installed on client systems to support client-side deduplication.

1. Download QuestRapidNFS-xxxxx-xxxxx-x86_64.bin.gz from the website, as detailed previously.
2. Unzip the package.
`unzip QuestRapidNFS-xxxxx-xxxxx-x86_64.bin.gz`
3. Assign execute bit to change the permission of the binary package:
`chmod +x QuestRapidNFS-xxxxx-xxxxx-x86_64.bin`
4. Install the Rapid NFS package. Before installing, remove the stale NFS entry.
`QuestRapidNFS-xxxxx-xxxxx-x86_64.bin -install`
5. Load the file system in user space (FUSE) module, if not already loaded:
`modprobe fuse`
6. Create a directory on the client. For example:
`mkdir /mnt/backup`
7. Mount Rapid NFS as a file system type using the mount command. For example:
`mount -t rdnfs 10.222.322.190:/containers/backup /mnt/backup`
 If you are using a DMA that supports a marker, set the marker by using -o in the mount command. For example, if you wanted to set the CommVault marker (cv):
`mount -t rdnfs 10.222.322.190:/containers/backup /mnt/backup -o marker=cv`

i | **NOTE:** If you want to do a mount on AIX, you must set the `nfs_use_reserved_ports` and `portcheck` parameters first. The parameters cannot be set to 0. For example: `root@aixhost1 / # nfso -po portcheck=1 root@aixhost1 / # nfso -po nfs_use_reserved_ports=1`

To ensure that the plug-in is running successfully, check the log file at: `tail -f /var/log/oca/rdnfs.log`.

Installing the Rapid CIFS plug-in

The DR Rapid CIFS plug-in must be installed on to the media server type you choose (for supported operating systems and DMAs, see the *DR Series System Interoperability Guide*). The plug-in software enables integration between DR Series system data storage operations and the supported data management applications (DMAs). Before you install, make sure you adhere to the Best Practices covered in another topic in this chapter.

You can download the plug-in installer from the Quest website as follows

- Go to support.quest.com/dr-series, select your specific DR Series model, and then navigate to Software Downloads.
- Locate the DR Rapid CIFS plug-in and download it to your system.

After it is downloaded, follow the steps below to run the plug-in installer to install the plug-in on your designated media server.

NOTE: The plug-in needs to be installed on client systems to support client-side deduplication.

1. On the media server, map a network share to your CIFS-enabled container.
2. Download the plug-in installer from the website, as detailed previously.
3. Open a command prompt with the "Run as Administrator" option selected. To do this using the Windows Start menu, click Start → All Programs → Accessories. Right-click Command Prompt and select Run as Administrator.
This gives all the required privileges to install/copy the driver files to the Windows drivers folder.
4. Run `QuestRapidCIFS-xxxxxx.msi`.
5. Follow the installation prompts.

To ensure that the plug-in is running successfully, check the Windows Event log file.

Determining if your system is using Rapid NFS or Rapid CIFS

Use this procedure to identify whether Rapid NFS or Rapid CIFS is installed and enabled on your DR Series system.

To determine if your system is using the Rapid NFS or Rapid CIFS accelerator:

1. In the GUI, go to the **Dashboard**, and then click **Container Statistics**.
2. In the **Container Name** drop-down list, select a NFS or CIFS container that is associated with your client.
3. In the **Connection Configuration** pane of the statistics page, locate the **NFS Write Accelerator** or **CIFS Write Accelerator** field, depending on the protocol selected.
4. Next to the **Write Accelerator** field is a value. **Active** indicates that the accelerator plug-in is installed and enabled. **Inactive** indicates that the plug-in is not installed or not working correctly.

Viewing the Rapid NFS and Rapid CIFS logs

This topic contains information about locating and reviewing Rapid NFS and Rapid CIFS event logs for troubleshooting purposes

Viewing Rapid NFS Logs

The Rapid NFS log is located at `/var/log/rdnfs.log`. Statistics, throughput, and the plug-in version can be seen on the client by running the `ru` utility on the client, as follows:

```
ru --mpt=[rdnfs mount point] | --pid=[process ID of rdnfs] --show=[name|version|parameters|stats|performance]
```

The configuration file is located `/etc/oca.0/rdnfs.cfg`.

Viewing Rapid CIFS Logs

If you want a high-level view of events and errors for the Rapid CIFS accelerator, open the Windows Event Log.

If you want to view more detailed event messages from Rapid CIFS, you can access a secondary log using the following Rapid CIFS utility command. The utility is located in `Program Files\Quest\Rapid CIFS`.

```
rdcifctl.exe -collect
```

Monitoring performance

This procedure describes how to monitor performance by viewing Rapid NFS and Rapid CIFS usage graphs.

Before you view usage graphs, make sure that the appropriate accelerator is active by viewing the Container Statistics in the DR Series system GUI.

To monitor Rapid NFS and Rapid CIFS performance:

1. In the DR Series system GUI, click **Dashboard > Usage**. The Dashboard page is displayed.
2. In the Actions menu in the upper right corner of the page, click **Detailed Graphs**.
3. Select a time range (if needed) and click **Apply**.
4. Click the **Protocols** tab.
Under **NFS Usage** and **CIFS Usage**, there is an **XWrite** checkbox. This checkbox represents the accelerator activity.
5. In the desired usage graph pane, select the **XWrite** checkbox to view the accelerator performance over time.

If you have Rapid NFS enabled, you can use the command line to view statistics, throughput, and the plug-in version by running the `ru` utility on the client, as follows:

```
ru --mpt=[rdnfs mount point] | --pid=[process ID of rdnfs] --show=[name|version|parameters|stats|performance]
```

If you have Rapid CIFS enabled, you can use the command line to view aggregate statistics (even while a backup job is running) using the following command:

```
\Program Files\Quest\Rapid CIFS\rdcifctl.exe stats -s
```

Uninstalling the Rapid NFS plug-in

Use the following procedure to remove the DR Rapid NFS plug-in from a Linux-based media server. After you uninstall the plug-in, Rapid NFS will be disabled and “inactive” will be shown next to **NFS Write Accelerator** on the **NFS Connection Configuration** pane on the **Container Statistics** page.

i | **NOTE:** You should retain the DR Rapid NFS plug-in installer on the media server in case you need to use it to reinstall the plug-in. It is usually located in `/opt/quest/DR-series/RDNFS/scripts`.

To uninstall the Rapid NFS plug-in on Linux:

1. Stop the Data Management Application (DMA) backup service before using the `-uninstall` option. The Rapid NFS plug-in installer returns an error if the DMA service is running when attempting to uninstall the plug-in.
2. Run the Rapid NFS plug-in installer (usually located in `/opt/Quest/DR-series/RDNFS/scripts`) with the `-uninstall` option, which uninstalls the plug-in, using the following command:

```
$> ./QuestRapidNFS-xxxxx-x86_64.bin -uninstall
```

i | **NOTE:** You must stop the DMA service before uninstalling the Rapid NFS plug-in (you are also required to use the Rapid NFS plug-in installer to uninstall the plug-in).

3. Check that the plug-in is uninstalled by viewing the usage graph in the GUI; it should not indicate any **XWrite** activity.

Uninstalling the Rapid CIFS plug-in

Use the following standard Microsoft Windows uninstall process to remove the Rapid CIFS plug-in from a Windows-based media server. After you uninstall the plug-in, Rapid CIFS will be disabled and “inactive” will be shown next to CIFS Write Accelerator on the CIFS Connection Configuration pane on the Container Statistics page.

Alternatively, if you want to disable (but not uninstall) the plug-in, you can run the following Rapid CIFS utility command. The utility is located in Program Files\Quest\Rapid CIFS.

```
rdcifsctl.exe driver -d
```

i | **NOTE:** Replace this text with a description of a feature that is noteworthy.

To uninstall the Rapid CIFS plug-in on Windows:

1. Click **Start**, and click **Control Panel**.
2. Under **Programs and Features**, click **Uninstall a program**.
3. Locate the Rapid CIFS plug-in in the listed of installed programs, right-click, and select **Uninstall**.
4. Click **Yes** to uninstall the Rapid CIFS plug-in.

Configuring and using Rapid Data Access with NetVault Backup and with vRanger

Overview

Rapid Data Access (RDA) with NetVault Backup and with vRanger provides the logical disk interface that can be used with network storage devices. The DR Series system requires a DR Rapid plug-in to integrate its data storage operations with NetVault Backup and vRanger. The plug-in is installed by default on the NetVault Backup and vRanger servers and the DR Series system when the latest software updates are installed. Using the DR Rapid plug-in, the DMAs can take full advantage of key DR Series system features like replication and data deduplication.

When DR Rapid is used with the DR Series system, it offers the following benefits:

- RDA with NetVault Backup and RDA with vRanger protocols provide faster and improved data transfers:
 - Focus is on backups with minimal overhead
 - Accommodates larger data transfer sizes
 - Provides throughput that is better than CIFS or NFS
- DR Rapid and data management application (DMA) integration:
 - DMA-to-media server software communication
 - DR Series system storage capabilities can be used without extensive changes to DMAs
 - Backup and replication operations are simplified by using built-in DMA policies
- DR Series system and DR Rapid ports and write operations:
 - Control channel uses TCP port 10011
 - Data channel uses TCP port 11000
 - Optimized write operations enable client-side deduplication

- Replication operations between DR Series systems:
 - No configuration is required on the source or target DR Series system
 - Replication is file-based, not container-based
 - Replication is triggered by DMA optimized duplication operation
 - DR Series system transfers the data file (not the media server)
 - After duplication completes, DR Series system notifies DMA to update its catalog (acknowledging the second backup). This makes the DMA aware of the replication location. Restores from either the source or replication target can be used directly from the DMA.
 - Supports different retention policies between source and replica
 - Replication is set up in the DMA itself, not the DR Series system

Guidelines for using RDA with NetVault Backup and with vRanger

For best results, observe the following guidelines for optimal performance with your supported RDA with NetVault Backup and RDA with vRanger operations with the DR Series system:

- Backup, restore, and optimized duplication operations are performed using the RDA with NetVault Backup or RDA with vRanger plug-in.
 - i | NOTE:** The plug-in is installed on client systems to support client-side deduplication.
- Deduplication. The DR Series system supports two deduplication modes:
 - **Passthrough:** When this mode is selected, deduplication processing occurs on the DR Series system. Passthrough writes occur when data is sent from a media server to the DR Series system without applying any optimization to the data.
 - i | NOTE:** Passthrough mode requires at least 200MB of free RAM on the backup client.
 - **Dedupe:** When this mode is selected, dedupe writes occur when data is sent from a media server to the DR Series system after optimization is applied to the data; deduplication processing would occur on a NetVault Backup client, for example.
 - i | NOTE:** Dedupe mode requires at least 4GB of free RAM on the DR Series system.

Best Practices: RDA with NetVault Backup and vRanger and the DR Series System

This topic introduces some recommended best practices for using DR Rapid operations with the DR Series system.

- RDS and non-RDS type containers can exist on the same DR Series system
The DR Series system supports having both RDS and non-RDS containers on the same appliance. However, this can cause incorrect capacity reporting as both container types share the same underlying storage.
- RDS replication and non-RDS replication on the same DR Series system
Non-RDS replication must be configured, and it is replicated on a per-container basis. However, this type of replication will not replicate RDS containers. RDS replication is file-based and is triggered by the DMA.
- Do not change the container connection type from NFS/CIFS to RDS
A non-RDS container must be deleted before this container can then be created as an RDS container using the same name.

Setting client-side optimization

Client-side optimization (also known as client-side deduplication) is a process that can contribute to saving time performing backup operations and reducing the data transfer overhead on the network.

To configure dedupe/passthrough before the client makes a connection, you need to use the DR Series system command line interface (CLI).

i **NOTE:** To update a client in the DR Series system GUI, the client must already be attached (and therefore enabled to appear in the GUI). If a client connection exists, you can select the radio button in the GUI to modify it.

You can turn client-side optimization On or Off by using the CLI commands, `rda --update_client --name --mode`. For more information about DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Adding RDS devices in NetVault Backup

To add RDS devices in NetVault Backup:

1. Start the NetVault Web user interface (UI), and log on to the NetVault Backup Server.
2. Start the configuration wizard by doing one of the following:
 - In the Navigation pane, click **Guided Configuration**, and then, on the NetVault Configuration Wizard page, click b.
 - Alternatively, in the navigation pane, click **Manage Devices**, and then click **Add Device**.
3. Select the **RDA Device** option, and click **Next**.
4. In **Host**, enter the IP address or the system host name of the DR Series system.
5. In **Username**, enter `backup_user`.

i **NOTE:** The Username, `backup_user` is case-sensitive. You can configure RDS containers only while logged on the DR Series system with username `backup_user`.

6. In **Password**, enter the password used to access the DR Series system.
7. In **LSU**, enter the name of the RDS container.
 - i** **NOTE:** The RDS container name in LSU is case-sensitive. Ensure that you enter the RDS container name exactly as it is on the DR Series system.
8. In **Block size**, enter the block size for data transfers. The block size is specified in bytes. The default block size is 131,072 bytes.
9. If the device is already added to another NetVault Backup Server with the same name, select the **Force add** check box. This option can be useful if you have performed a disaster recovery to rebuild the NetVault Backup Server.
10. Click **Next** to add the device.

After the device is successfully added and initialized, a message is displayed.

Removing RDS Devices From NetVault Backup

Refer to the following steps to remove existing RDS devices from NetVault Backup.

- i** **NOTE:** Removing an RDS device from NetVault Backup does not delete the data stored in the RDS container on the DR Series system.
1. Start the NetVault Web user interface (UI), and in the Navigation pane, click **Manage Devices**.
 2. In the list of devices, locate the device, and click the corresponding **Manage Device** icon.
 3. Click **Remove**, and then in the confirmation dialog box, click **Remove** again.
 - i** **NOTE:** Ensure that you remove the RDA device from NetVault Backup before you delete the container from the DR Series system. You must force remove the RDS device from NetVault Backup, if you delete an RDS container from the DR Series system before removing it from the NetVault Backup server.
 4. If NetVault Backup fails to remove the device, select the **Force Removal** check box in the confirmation dialog, and click **Remove**.

The selected RDS device is removed from NetVault Backup. The RDS container can now be removed from the DR Series system.

Backing Up Data on the RDS container using NetVault Backup

You must back up data on the RDS container (available on the DR Series systems) using NetVault Backup. Before you can back up data using the RDS protocol, you must create an RDS container on the DR Series system and then add that container as an RDA device on NetVault Backup. For more information see, [Adding RDS Devices in NVBU](#).

To back up data on the RDS container:

1. Start the NetVault Web user interface (UI), and in the Navigation pane, click **Create Backup Job**.
2. In **Job Name**, type a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress or restoring data. A job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.
3. In the **Selections** list, select an existing Backup Selection Set, or click **Create New**, and select the items that you want to back up. The selection tree is plug-in specific. For more information about selecting data for backups, see the relevant NetVault Backup plug-in user's guide.
4. In the **Plugin Options** list, select an existing Backup Options Set, or click **Create New**, and configure the options that you want to use. These options are plug-in specific. For more information about these options, see the relevant NetVault Backup plug-in user's guide.
5. In the **Schedule** list, select an existing Schedule Set, or click **Create New**, and configure the schedule type and schedule method. For more information about these options, see the *NetVault Backup Administrator's Guide*. To run the job as soon as it is submitted, use the **"Immediate"** set.
6. In the **Target Storage** list, select an existing Target Set, or click **Create New**, and configure the target device and media options for the job.
To use a particular DR Series System, select the **Specify Device** option, and in the list of devices, clear the check marks for the devices that you do not want to use.
For more information about these options, see the *NetVault Backup Administrator's Guide*.
7. In the **Advanced Options** list, select an existing Backup Advanced Options Set, or click **Create New**, and configure the options that you want to use.
For more information about these options, see the *NetVault Backup Administrator's Guide*.
8. To submit the job for scheduling, click **Save & Submit**.

The backup job may take a few minutes to complete depending on the amount of data that is backed up. You can view the progress of the backup job by using the Job Management section of NetVault Backup. For more information about using NetVault Backup, see the *NetVault Backup Administrator's Guide*.

Replicating Data to a RDS Container using NetVault Backup

By using NetVault Backup with the DR Series system, you can run optimized replication jobs. You can replicate data in backup RDS containers on one DR Series system to a target RDS container that is on a different DR Series system. Both the source and target RDS containers must be added to the NetVault Backup server as RDA devices. You can complete optimized replication (or optimized duplication) of backups that you complete using NetVault Backup.

NOTE: You cannot replicate RDS containers using the DR Series system native replication feature.

NOTE: The source or backup container and the target container must use the RDS protocol.

To replicate the data available on the backup RDS container to a target RDS container:

1. In the **NetVault Backup Console**, click **Backup**.
2. From the **Server Location** list, select the relevant NetVault Backup server.

3. In **Job Title**, enter a relevant job title.
4. In the **Selections** tab, select **Data Copy** and then **Backups** or **Backup Sets**, and navigate to the backup job that you want to replicate.
5. Select the **Backup Options** tab, under **Data Copy Options** select the relevant options.

NOTE: Under Copy Type, by default, options are set for Copy and Optimized replication for the DR Series systems.
6. Select the **Schedule** tab, under **Schedule Options** select one of the following:
 - **Immediate** — To start the backup operation as soon as you save the current backup job.
 - **Once** — To run the backup only once at a scheduled time and date.
 - **Repeating** — To run the backup at a scheduled time and date on a daily, weekly, or monthly basis.
 - **Triggered** — To run the backup if the system encounters a pre-specified Trigger name.
7. Under **Job Options** select the relevant options.
8. Select the **Source** tab, under **Device Options** select, **Specify Device**.
9. Select the relevant source RDS device from the list of displayed devices.
You can select more than one device.
10. Select the **Target** tab, under **Device Options** select, **Specify Device**.
11. Select the relevant target RDS device from the list of displayed devices.
You can select more than one device.
12. Under **Media Options** and **General Options**, select the relevant option.
13. Select the **Advanced Options** tab and select the relevant options.
14. To run the optimized replication job, click the **b** icon.

NOTE: For more information on NetVault Backup, see the *NetVault Backup Administrator's Guide*.

Restoring data from a DR Series system by using NetVault Backup

The following steps describe how to use NetVault Backup to restore data from a RDS container on a DR Series system.

To restore data from a DR Series system using NetVault Backup:

1. Start the NetVault Web user interface (UI), and in the Navigation pane, click **Create Restore Job**.
2. In the saveset table, select the saveset that you want to use, and click **Next**.
3. On the **Create Selection Set** page, select the items that you want to restore.
The selection tree is plug-in specific. For more information about selecting data for restores, see the relevant NetVault Backup plug-in user's guide.

4. Click **Edit Plugin Options**, and configure the options that you want to use and then click **Next**.
These options are plug-in specific. For more information about these options, see the relevant NetVault Backup plug-in user's guide.
5. On the **Create Restore Job** page, specify a name for the job. Assign a descriptive name that allows you to easily identify the job for monitoring its progress.
A job name can contain alphanumeric and non-alphanumeric characters, but it cannot contain non-Latin characters. There is no length restriction. However, a maximum of 40 characters is recommended on all platforms.
6. In the **Target Client** list, select the restore target as follows:
 - To restore data to the same client (from which data was backed up), use the default setting.
 - To restore data to an alternate client, select the target client in the list.
 - Alternatively, click **Choose**. In the **Choose the Target Client** dialog box, select the client, and click **OK**.
7. In the **Schedule** list, select an existing Schedule Set, or click **Create New**, and configure the schedule type and schedule method. To run the job as soon as it is submitted, use the **"Immediate"** set.
For more information about these options, see the *NetVault Backup Administrator's Guide*.
8. In the **Source Options** list, select an existing Source Set, or click **Create New**, and configure the source device options. To use a particular DR Series System, select the **Specify Device** option, and in the list of devices, clear the check marks for the devices that you do not want to use.
For more information about these options, see the *NetVault Backup Administrator's Guide*.
9. Click **Submit** to submit the job for scheduling.

i | **NOTE:** For more information about using NetVault Backup, see the *NetVault Backup Administrator's Guide*.

Supported DR Series system CLI commands for RDS

The following are the supported DR Series system CLI commands for RDS operations:

```
administrator@DocTeam-SW-01 > rda
```

Usage:

```
rda --show [--config]

                [--file_history] [--name <name>]
                [--active_files] [--name <name>]
                [--clients]
                [--limits]

rda --setpassword

rda --delete_client --name <RDA Client Hostname>
```

```

    rda --update_client --name <RDA Client Hostname>--mode <auto|passthrough|dedupe>
    rda --limit --speed <<num><kbps|mbps|gbps> | default> --target <ip address
| hostname>

    rda --help

```

rda <command> <command-arguments>

<command> can be one of:

```

--show Displays command specific information.

--setpassword      Updates the Rapid Data Access (RDA) user password.

--delete_client    Deletes the Rapid Data Access (RDA) client.

--update_client     Updates attributes of a Rapid Data Access (RDA) client.

--limit            Limits bandwidth consumed by Rapid Data Access
(RDA) when replicating over a WAN link.

```

For command-specific help, please type

```

rda --help <command>
    eg:
        rda --help show

```

i **NOTE:** The **--files** in the **rda --show --file_history** command represents replicated files that were processed via the DMA optimized duplication operation. This command displays only up to the last 10 such files. The **--name** in the **rda --show --name** command represents the RDA container name. For more information about RDA-related DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Configuring and using RDA with OST

This topic introduces key RDA with OST tasks and provides links to other related topics that contain procedures that describe how to perform these tasks:

- Configuring the DR Series system for use with OST and the supported DMAs; for more information, see [Configuring the DR Series System Using the Backup Exec GUI](#), and [Configuring DR Series System Information Using NetBackup](#)
- Configuring the Logical Storage Unit (LSU) using the DR Series system GUI; for more information, see [Configuring an LSU](#)
- Installing the RDA with OST plug-in to a supported media server (Linux or Windows)
- Using supported DMAs to perform backup and restore operations; for more information, see
 - [Backing Up Data from a DR Series System Using NetBackup](#)
 - [Restoring Data from a DR Series System Using NetBackup](#)
 - [Duplicating Backup Images Between DR Series Systems Using NetBackup](#)
 - [Creating Backups on the DR Series System Using Backup Exec](#)
 - [Restoring Data from a DR Series System Using Backup Exec](#)
 - [Optimizing Duplication Between DR Series Systems Using Backup Exec](#)

i **NOTE:** This capability to use RDA with OST, also known as DR Rapid, adds tighter integration with backup software applications, such as the following OpenStorage-enabled backup applications: NetBackup and Backup Exec.

Understanding RDA with OST

OpenStorage Technology (OST) provides the logical disk interface that can be used with network storage devices, and the DR Series system appliance requires RDA with OST plug-in software to integrate its data storage operations with supported data management applications (DMAs). For details on the applications supported, see the *DR Series System Interoperability Guide*.

The DR Series system integrates with supported DMAs using the RDA with OST plug-in, through which DMAs can control when backup images are created, duplicated, and deleted. Via the plug-in, the DMAs can take full advantage of key DR Series system features like replication and data deduplication.

The DR Series system accesses the OpenStorage API code through the RDA with OST plug-in, which can be installed on the supported media server platform type that you choose (Windows or Linux). When RDA with OST is used with the DR Series system, it offers the following benefits:

- RDA with OST protocol provides faster and improved data transfers:
 - Focus is on backups with minimal overhead
 - Accommodates larger data transfer sizes
 - Provides throughput that is significantly better than CIFS or NFS
- RDA with OST and DMA integration:
 - OpenStorage API enables the DMA-to-media server software communication
 - DR Series system storage capabilities can be used without extensive changes to DMAs
 - Backup and replication operations are simplified by using built-in DMA policies
- DR Series system and RDA with OST ports and write operations:
 - Control channel uses TCP port 10011
 - Data channel uses TCP port 11000
 - Optimized write operations enable client-side deduplication
- Replication operations between DR Series systems:
 - No configuration is required on the source or target DR Series system
 - Replication is file-based, not container-based
 - Replication is triggered by DMA optimized duplication operation
 - DR Series system transfers the data file (not the media server)
 - Once duplication completes, DR Series system notifies DMA to update its catalog (acknowledging the second backup)
 - Supports different retention policies between source and replica
 - Replication is set up in the DMA itself, not the DR Series system

Guidelines

For best results, observe the following guidelines for optimal performance with your supported RDA with OST operations with the DR Series system:

- Backup, restore, and optimized duplication operations need to be performed via the RDA with OST plug-in

i **NOTE:** The RDA with OST plug-in needs to be installed on client systems to support client-side deduplication.

- Backup:
 - Passthrough writes: Passthrough writes are when data is sent from a media server to the DR Series system without applying any optimization to the data.
 - Optimized writes: Optimized writes are when data is sent from a media server to the DR Series system after optimization is applied to the data.
- Minimum client memory requirements:
 - Minimum number of CPUs — Four (4) cores
 - Minimum amount of free physical memory — 4 GB

Terminology

This topic introduces and briefly defines some basic RDA for OST terminology used throughout the DR Series system documentation.

Term	Description
BE	Veritas DMA, Backup Exec (BE)
DMA/DPA	Data Management Application (also known as Data Protection Application), which are terms for the role played by the backup applications used with RDA with OST; for example, Veritas NetBackup or Backup Exec.
LSU	Logical Storage Unit, which from the DR Series system perspective, represents any container created for data storage. <i>LSU</i> is a common storage term while <i>container</i> is a common term in DR Series systems that represents a location for storing data.
media server	This is the host running the DMA media server and is where the RDA with OST plug-in is installed. The RDA with OST plug-in can also be installed on a client.
NBU	Veritas DMA, NetBackup (NBU)
OST	The OpenStorage Technology from Symantec, which allows storage devices to deliver backup and recovery solutions with NetBackup. RDA with OST uses the OpenStorage API and a plug-in installed on either a Linux or a Windows-based media server platform.

Supported RDA with OST software and components

For the list of supported DMAs and DR Rapid plug-ins, see the *DR Series System Interoperability Guide*.

The DR Series system licensing is all-inclusive, so that no additional licensing is required to use RDA with OST or the optimized duplication capability. The RDA with OST plug-in that gets installed on a supported Linux or Windows media server platform is a free download. However, if you are using Veritas backup applications, you may need to purchase additional licensing to enable OpenStorage Technology; refer to your Veritas documentation.

Best Practices: RDA with OST and the DR Series System

This topic introduces some recommended best practices for using RDA with OST operations with the DR Series system.

- OST and non-OST containers can exist on the same DR Series system. The DR Series system supports having both OST and non-OST containers on the same appliance. However, this can cause incorrect capacity reporting as both container types share the same underlying storage.
- OST replication and non-OST replication on the same DR Series system. Non-OST replication needs to be configured, and it is replicated on a per-container basis. However, this type of replication will not replicate OST containers. OST replication is file-based and is triggered by the DMA.
- Do not change the container connection type from NFS/CIFS to OST. A non-OST container must be deleted before this container can then be created as an OST container using the same name.

Setting client-side optimization

Client-side optimization (also known as client-side deduplication) is a process that can contribute to saving time performing backup operations and reducing the data transfer overhead on the network.

To configure dedupe/passthrough before the client makes a connection, you need to use the DR Series system command line interface (CLI).

i NOTE: To update a client in the DR Series system GUI, the client must already be attached (and therefore enabled to appear in the GUI). If a client connection exists, you can select the radio button in the GUI to modify it.

You can turn client-side optimization On or Off by using the CLI commands, `rda --update_client --name --mode`. For more information about DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Configuring an LSU

You can configure a logical storage unit (LSU) as an OpenStorage Technology (OST) connection type container for data storage by using the DR Series system GUI. To configure an LSU as an OST connection type container, log in to the DR Series system and complete the following:

1. Navigate to the **Containers** page (in the Dashboard navigation panel).
2. Click **Create** to create a new container.
3. In **Container Name**, enter a name for the container.
4. In **Marker Type**, select the **None** marker type.
For OST operations, only the NetBackup and Backup Exec media servers are supported.
5. In **Connection Type**, set the container type to **Rapid Data Access (RDA)**.
6. In the **RDA** pane, set the RDA Type to **Symantec OpenStorage (OST)**.

7. In **Capacity**, select either the **Unlimited** or **Size** options to set the capacity for the OST connection type container.
If you select **Size**, make sure to define the desired size in Gibibytes (GiB).
8. Click **Create a New Container** (or click **Cancel** to display the **Containers** page).

- i** **NOTE:** For general information about creating DR Series system containers, see [Creating Containers](#), and for creating an OST connection type container, see [Creating an OST or RDS Connection Type Container](#).
- i** **NOTE:** The capacity option in this command example sets the quota on the LSU. This is the maximum number of bytes (ignoring optimization) that can be written to an LSU and it is listed in the gigabytes (GB). If the capacity option is not specified (or if 0 is specified for the capacity), then the LSU will not have a quota. If this is the case, then this means that the amount of data that can be written to the LSU is limited only by the amount of free space on the disk.

Installing the RDA with OST plug-in

Before you can start the installation process for the RDA with OST plug-in, you need to understand its role. The plug-in must be installed on to the media server type you choose. (For details on supported platforms, see the *DR Series System Interoperability Guide*.) The RDA with OST plug-in software enables integration between DR Series system data storage operations and the supported data management applications (DMAs).

Understanding the RDA with OST Plug-in (Linux)

The plug-in must be installed on the designated Linux-based media server running the support Linux server operating system software in the following directory: `/usr/openv/lib/ost-plugin-ins`. The RDA with OST plug-in is installed using a self-extracting installer that installs the plug-in and all of its related components. The installer supports the following modes, with the default being Help (-h):

- i** **NOTE:** If no option is selected, the Help mode is displayed by default.

- Help (-h)
- Install (-install)
- Upgrade (-upgrade)
- Uninstall (-uninstall)
- Force (-force)

```
$> ./QuestOSTPlugin-xxxxxx-x86_64.bin -help
Quest plug-in installer/uninstaller
usage: QuestOSTPlugin-xxxxxx-x86_64.bin [ -h ] [ -install ] [ -uninstall ]
-h                               : Displays help
-install                         : Installs the plug-in
-upgrade                       : Upgrades the plug-in
-uninstall                     : Uninstalls the plug-in
-force                         : Forces the installation of the plug-in
```

You can download the RDA with plug-in installer from the Quest website as follows:

- Go to support.quest.com/dr-series, select your specific DR Series model and then navigate to Software Downloads.
- Locate the RDA with OST plug-in for Linux and download it to your system.

After it is downloaded, run the RDA with OST plug-in installer to install the plug-in on your designated Linux-based media server.

i | **NOTE:** The RDA with OST plug-in must be installed on client systems to support client-side deduplication.

Understanding the RDA with OST Plug-in (Windows)

The RDA with OST plug-in must be installed in the following directory on the designated Windows-based media server running the supported Microsoft Windows server operating system software: \$INSTALL_PATH\VERITAS\Netbackup\bin\ost-plugin-ins for NetBackup installations, and \$INSTALL_PATH\Veritas\Backup Exec\bin\ for Backup Exec installations. After it is downloaded, you can use **SETUP** to install the RDA with OST plug-in.

i | **NOTE:** The RDA with OST plug-in must be installed on client systems to support client-side deduplication.

Installing the RDA with OST plug-in for Backup Exec on Windows

This topic describes how to install the RDA with OST plug-in within a Microsoft Windows environment for performing DR Series system operations via the plug-in.

Make sure that you meet all of the following prerequisites before installing the plug-in:

- The Backup Exec installation must be running on one of the supported Windows media server platforms. For information on the supported versions of Backup Exec and operating systems, see the *DR Series System Interoperability Guide*, available at support.quest.com/dr-series.
- The Windows RDA with OST installer must be downloaded. If not, download the Windows installer (QuestOSTPlugin-xxxxx.msi), which is available at support.quest.com/dr-series, to a network directory location you can access.

To install the RDA with OST plug-in, complete the following steps:

1. Launch the **Backup Exec Administrator** console, select **Tools**, and **Backup Exec Services....**
2. Select the server on which you want to install the RDA with OST plug-in, and select **Stop all services**.
3. Click **OK**.
4. Launch the **Storage Plug-In for OST Setup Wizard** (and accept all default values).
5. In the **Welcome** page, click **Next** to continue.
6. Click **I accept the terms in the License Agreement**, and click **Next**.
7. In the **Destination Folder** page, accept the default destination location, and click **Next**.

8. In the **Ready to Install Storage Plug-In for OST** page, click **Install**.
When the plug-in has been installed, the **Completed the Storage Plug-In for OST Setup Wizard** page is displayed.
9. Click **Finish** to exit the wizard.

Installing the RDA with OST plug-in for NetBackup on Windows

This topic describes how to install the RDA with OST plug-in on a media server running the supported Microsoft Windows server operating system software (and using the NetBackup DMA).

Ensure that you have downloaded the RDA with OST plug-in installer into the correct directory on the designated media server. The plug-in installer is saved as QuestOSTPlugin64-xxxxx.msi (for 64-bit operating systems), or QuestOSTPlugin-xxxxx.msi (for 32-bit operating systems). Ensure that the correct plug-in is downloaded to support your 64-bit or 32-bit system.

1. Stop the NetBackup services if they are running, by using the following command:

```
$INSTALL_PATH\VERITAS\NetBackup\bin\bpdown.exe
```
2. Run **SETUP** to install the plug-in.
3. Check that the plug-in is installed by using the following NetBackup command on the Windows media server:

```
$INSTALL_PATH\VERITAS\NetBackup\bin\admincmd\bpstsinfo.exe -pi
```

This NetBackup command lists the plug-in details, as shown in the following example:

- Plug-In Name: libstspiocaMT.dll
- Prefix: Quest
- Label: OST Plug-in that interfaces with the DR Series system
- Build Version: 9
- Build Version Minor: 1
- Operating Version: 9
- Vendor Version: Quest OST plug-in 10.1

4. Start the NetBackup services by using the following command:

```
$INSTALL_PATH\VERITAS\NetBackup\bin\bpup.exe
```

Uninstalling the RDA with OST plug-in for Windows

Use the following steps if you need to uninstall the RDA with OST plug-in from a Windows-based media server. Use the standard Microsoft Windows uninstall process to uninstall the RDA with OST plug-in from a Windows-based media server.

i **NOTE:** You should retain the RDA with OST plug-in installer on the media server in case you need to use it to reinstall the plug-in.

1. Click **Start**, and click **Control Panel**.
2. Under **Programs and Features**, click **Uninstall a program**.
3. Locate the RDA with OST plug-in in the listed of installed programs, right-click and select **Uninstall**.
4. Click **Yes** to uninstall the plug-in.

Installing the RDA with OST plug-in for NetBackup on Linux

This topic describes how to install the RDA with OST plug-in on a media server running the supported Red Hat Enterprise Linux or SUSE Linux server operating system software (using the NetBackup data management application).


Ensure that you have downloaded the RDA with OST plug-in installer into the correct directory on the designated media server. The plug-in installer is saved as QuestOSTPlugin-xxxxx-x86_64.bin.gz, where xxxxx represents the build number.

1. Unzip the RDA with OST plug-in installer file using the following command:

```
$> /bin/gunzip QuestOSTPlugin-xxxxx-x86_64.bin.gz
```
2. Configure the executable bit on the plug-in installer using the following command:

```
$> /bin/chmod a+x QuestOSTPlugin-xxxxx-x86_64.bin
```
3. Stop the NetBackup nbrmms service before using the -install option.
The plug-in installer returns an error if the NetBackup nbrmms service is running when attempting to install the plug-in.
4. Run the plug-in installer using the -install option, and install the plug-in using the following command:

```
$> ./QuestOSTPlugin-xxxxx-x86_64.bin -install
```

 **NOTE:** The location for installing the plug-in is not user-configurable.

5. After the RDA with OST plug-in installer has stopped running, and the system prompt returns, verify that the plug-in has loaded properly by checking the output using the following NetBackup command on the Linux media server:

```
$> /usr/openv/netbackup/bin/admincmd/bpstsinfo -plugininfo
```

This NetBackup command lists the plug-in details as shown:

- Plug-In Name: libstspiocaMT.so
- Prefix: QUEST
- Label: Quest OpenStorage (OST) Plug-in
- Build Version: 10
- Build Version Minor: 1
- Operating Version: 10
- Vendor Version: (EAR-2.0.0) Build: 41640

6. Retain the plug-in installer on the media server so you can use it if needed to uninstall the plug-in.

Uninstalling the RDA with OST plug-in for Linux

Use the following process if you need to uninstall the RDA with OST plug-in from a Linux-based media server:

1. Stop the NetBackup nbrmms service before using the -uninstall option.
(The plug-in installer returns an error if the NetBackup nbrmms service is running when attempting to uninstall the OST plug-in.)
2. Run the RDA with OST plug-in installer with the -uninstall option, which uninstalls the plug-in, using the following command:

```
$> ./QuestOSTPlugin-xxxxx-x86_64.bin -uninstall
```

3. Check that the plug-in is uninstalled by using the following NetBackup command on the Linux media server:

```
$> /usr/opensv/netbackup/bin/admincmd/bpstsinfo -plugininfo
```

i | **NOTE:** If the `-plugininfo` command returns any plug-in details, this means that the plug-in has not been uninstalled.

4. Retain the plug-in installer on the media server in case you need to use it to reinstall the plug-in.

Configuring DR Series system information using NetBackup

The topic introduces the concept of configuring the DR Series system information using the NetBackup media server command line interface (CLI) commands and graphical user interface (GUI) menus, tabs, and options. The NetBackup CLI commands and GUI menus, tabs, and options allow you to configure both the Linux or Windows media servers. In this guide, you will find specific topics that address operations for using the NetBackup CLI, such as adding the DR Series system name to NetBackup on each Linux and Windows media server you intend to use with the DR Series system, using the NetBackup GUI to configure it to work with the DR Series system via OST, using the NetBackup GUI to configure disk pools from logical storage units (LSUs) on the DR Series system, and using the NetBackup GUI to create storage units using the disk pools on the DR Series system.

Using NetBackup CLI to add DR Series system name (Linux)

This topic describes how to use the NetBackup CLI to add the DR Series system name to each Linux-based media server you plan to use with the DR Series system.

1. Add the DR Series system name to NetBackup by using the following command:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -creatests  
-storage_server servername -stype QUEST -media_server mediaservername
```

2. Log in to and authenticate with DR Series system by using the following command (for details, see [Configuring an LSU](#)).

```
/usr/opensv/volmgr/bin/tpconfig -add -storage_server servername -stype QUEST -sts_user_id backup_user -password password
```

i | **NOTE:** On the DR Series system, only one user account exists, and the user ID for that account is backup_user. You can only change the password for this account; you cannot create a new account nor can the existing account be deleted.

Using NetBackup CLI to add DR Series system name (Windows)

This topic describes how to use the NetBackup CLI to add the DR Series system name to each Windows-based media server you plan to use with the DR Series system.

1. Add the DR Series system name to NetBackup by using the following command:

```
$INSTALL_PATH\VERITAS\NetBackup\bin\admincmd\nbdevconfig  
-creatests -storage_server servername -stype QUEST -media_server mediaservername
```

2. Log in to and add the valid credentials for authentication by the DR Series system by using the following command (for details, see [Configuring an LSU](#)).

```
$INSTALL_PATH\VERITAS\Volmgr\bin\tpconfig -add -storage_server servername -stype QUEST -sts_user_id backup_user -password password
```

Configuring NetBackup for the DR Series system

Use the NetBackup graphical user interface (GUI) to configure it to work with the DR Series system via RDA with OST. This process is essentially the same type of operation for either the Linux or Windows platforms.

Log in to NetBackup, and complete the following:

1. In the main window of the **NetBackup Administrator** console, click **Configure Disk Storage Servers** to launch the **Storage Server Configuration Wizard**.
2. Select **OpenStorage** to choose the type of disk storage that you want to configure in this window, and click **Next**.

3. Enter the following values to configure a storage server:
 - In Storage server type, enter **DELL** or **QUEST** as appropriate.
 - In Storage server name, enter the name of the DR Series system.
 - In the **Select media server** drop-down list, select the desired media server (the server on which you are configuring RDA with OST).
 - Enter values for the credential needed to authenticate with the DR Series system:
 - **User name**
 - **Password**
 - **Confirm password**

The credentials should be the same as the credentials that are required for the DR Series system. For more information, see [Configuring an LSU](#).
4. Click **Next**.
5. Click **Next**.
6. Click **Next** and click **Finish** to close the **Storage Server Configuration Wizard**.

Configuring NetBackup for optimized synthetic backups

This procedure describes how to configure NetBackup so that it supports optimized synthetic backups. Optimized synthetic backups use RDA with OST to share data between images and synthesize the backup directly on the DR Series system without data being read to and written from the backup server. This saves time, expense, and space.

The DR Series system supports optimized synthetic backups with NetBackup 7.1 and 7.5. The NetBackup storage server inherits the Optimized Image attribute during storage server configuration (nbdevconfig -creatests).

To configure NetBackup to use optimized synthetic backups:

1. Use the following command to add the OptimizedImage flag to each NetBackup storage server that needs to support optimized synthetic backups:

```
nbdevconfig -change sts -stype PureDisk -storage_server ss_name -setattribute OptimizedImage
```

For `ss_name`, make sure to type the name of the storage server as you configured it in NetBackup.

2. Use the following command to add the OptimizedImage flag to each NetBackup disk pool that needs to support optimized synthetic backups:

```
nbdevconfig -changedp -stype PureDisk -dp dp_name -setattribute OptimizedImage
```

For `dp_name`, make sure to type the name of the disk pool as you configured it in NetBackup. Make sure to add the `OptimizedImage` flag to the storage server first, and then to the disk pool.

Creating disk pools from LSUs

Use the NetBackup graphical user interface (GUI) to configure disk pools from logical storage units (LSUs) on the DR Series system.

Log in to NetBackup, and complete the following steps:

1. In the main window of the **NetBackup Administrator** console, click **Configure Disk Pools** to launch the **Disk Pool Configuration Wizard**.
2. In the Welcome to the Disk Pool Configuration Wizard page, click **Next**.
3. In Type, select **OpenStorage**, and click **Next**.
4. In the **Storage server** list, select a server, and click **Next**.
5. Select the LSUs (volumes) to include from the list, and click **Next**.
6. Enter a **Disk pool** name, and click **Next**.
7. Verify the disk pool configuration in the Summary page, and click **Next** to configure the disk pool you created.

The **Performing required task** page is displayed, with the status being: Configuration completed successfully. You have several options available at this point:

- Clear the **Create a storage unit** for the disk pool.
- Click **Finish** and close the **Disk Pool Configuration Wizard**.
- Click **Next** to create the storage unit with this disk pool.

i **NOTE:** If you create the storage unit using the Disk Pool Configuration Wizard, you can skip the step where you create storage units using a disk pool.

8. Click **Next** to continue with creating a storage unit using this wizard.
9. Enter a **Storage unit name**, and click **Next**.
10. Click **Finish**.

To display the disk pool you created, click **Devices → Disk Pools** in the left navigation pane in the **NetBackup Administrator** console.

Creating storage units by using the disk pool

Use the NetBackup GUI to create storage units by using the disk pools on the DR Series system.

Log in to NetBackup, and complete the following tasks:

1. In the main window of the **NetBackup Administrator** console, click **Storage** in the left navigation pane, and select **Storage Units**.
2. In the NetBackup Administrator console main window, right-click and select **New Storage Unit** from the drop-down list.
3. In the **New Storage Unit** page, enter a name in **Storage unit name**, and select the OST disk pool that you created in the Disk pool drop-down list.
4. Click **OK** to create the new storage unit.

Backing up data from a DR Series system with NetBackup

This topic describes how to use NetBackup to back up data from a DR Series system.

Before backing up data, you first need to configure a policy that creates a backup on the OST logical storage unit (LSU). This type of policy is similar to what is done for network-attached storage (NAS) shares, except that when defining policy attributes, you need to select the LSU that contains the OST disk pool.

To back up data from a DR Series system using a policy, complete the following:

1. Log into the **NetBackup Administrator** console.
2. Click **NetBackup Management** in the left navigation pane, and select **Policies**.
3. In the **All Policies** main window, right-click **OST**, and select **Change Policy** from the drop-down list.
4. In the **Change Policy** page, click the **Attributes** tab, and select the settings for the policy you want to create.
5. Click **OK** to create the policy, which displays under OST in the main window.
6. Right-click the policy, and select **Manual Backup** from the drop-down list.
7. In the **Manual Backup** page, enter the name of the media server in **Server**, and click **OK**.

To monitor the status of any backup operation, click **Activity Monitor** in the left navigation pane of the **NetBackup Administrator** console, and select the backup job you are interested in to view details about the operation.

Restoring data from a DR Series system with NetBackup

This topic describes how to use NetBackup to restore data from a DR Series system. The process for restoring data from OST logical storage units (LSUs) is similar to how restores are performed from any backup device.

To restore data from a DR Series system, complete the following:

1. Log into the **NetBackup Administrator** console.
2. Click **Backup, Archive, and Restore** in the left navigation pane.
3. In the **Restore** main window, click the **Restore Files** tab.
4. Select the data that you want to restore, and click **OK**.

To monitor the status of any restore operation, click **Activity Monitor** in the left navigation pane of the **NetBackup Administrator** console, and select the restore job you are interested in to view details about the operation.

Duplicating backup images between DR Series systems with NetBackup

By using NetBackup with the DR Series system, you can duplicate backup images from a disk pool on one DR Series system to a target disk pool (or a storage unit derived from it) that could be on the same DR Series system or on a different DR Series system.

To duplicate backup images between DR Series systems using NetBackup, complete the following steps:

1. Log into **NetBackup Administrator** console.
2. Click **NetBackup Management** in the left navigation pane, and select **Catalog**.
3. In the **Catalog** main window, select **Duplicate** from the **Action** drop-down list, and click **Search Now**.
4. Right-click to select the image in the **Search Results** pane that you would like to duplicate, and select **Duplicate** in the drop-down list.
5. In the **Setup Duplication Variables** page, select the LSU that is the target DR Series system in the **Storage unit** drop-down list, and click **OK**.
6. To monitor the status of any duplicate image operation, perform the following:
 - a. Click **Activity Monitor** in the left navigation pane of the **NetBackup Administrator** console.
 - b. Select the data duplication job in which you are interested.
 - c. View the operation details.

Using Backup Exec with a DR Series system (Windows)

This topic introduces the RDA with OST plug-in and describes the installation prerequisites for Backup Exec within a Microsoft Windows environment. After it is installed, Backup Exec can perform DR Series system operations via the plug-in.

RDA with OST Plug-In and Supported Versions

For details on the supported Backup Exec versions and media server operating systems, see the *DR Series System Interoperability Guide*, available at support.quest.com/dr-series.

Installation Prerequisites for the RDA with OST Plug-in for Backup Exec

This topic introduces the installation prerequisites for installing the plug-in for Backup Exec on Windows media servers. Ensure that you meet the following prerequisites prior to installing the plug-in:

- The Backup Exec installation must be running on one of the supported Windows platforms.
- You should have created and configured an OST container on your DR Series system appliance.
- The RDA with OST plug-in must be downloaded. If not, download the Windows installer (QuestOSTPlugin-xxxxx.msi or QuestOSTPlugin64-xxxxx.msi), which is available at support.quest.com/dr-series, to a network directory location you can access.
- The plug-in needs to be installed in the following directory on the designated Windows-based media server running the supported Microsoft Windows operating system software (\$INSTALL_PATH\VERITAS\NetBackup\bin\ost-plugins) for NetBackup installations.

Configuring the DR Series system with the Backup Exec GUI

Backup Exec only supports the use of its own graphical user interface (GUI) for configuring the DR Series system. There is no supported Backup Exec command-line interface (CLI) for using Backup Exec 2010 version.

To configure the DR Series system using the Backup Exec GUI, complete the following steps:

1. Launch the **Backup Exec Administrator** console, select **Tools**, and **Backup Exec Services....**
2. Select the server that you want to configure in the Backup Exec Services Manager page, and select **Start all services**.
3. Verify that all services have been started, and click **OK**.
4. In the **Connect to Media Server** page, log into the media server, and enter a **User name**, a **Password**, and click **OK**.
5. In the Backup Exec Administrator page, click **Network**, and click **Logon Accounts**.
6. Click **New** to create a new logon account.
7. In the **Account Credentials** pane, enter the **User name** and **Password** account credentials for the DR Series system, and click **OK** (for example, the default user name is **backup_user**).
8. In the **Backup Exec Administrator** page, click the **Devices** tab, and right-click on the local system name that is listed as the root node.
9. Select **Add OpenStorage** in the drop-down list.
10. Configure the **Add OpenStorage Device** page with the following information, and click **OK**:
 - **Server**—enter the host name or IP address of the DR Series system.
 - **Logon account**—select the account from the drop-down list, which has credentials for accessing the DR Series system.
 - **Server type**—select the type of plug-in from the drop-down list (QUEST OST plug-in).
 - **Logical storage unit**—enter the LSU (DR Series system container) name to use.
11. Click **Yes** in response to the prompt about making the new device the default destination for new jobs.
12. Close the **Add OpenStorage Device** page.
13. Click **Restart Now** to restart the Backup Exec services.

Creating backups on the DR Series system using Backup Exec

This topic describes how to use Backup Exec to create backups on the DR Series system.

To create backups on the DR Series system using Backup Exec, complete the following steps:

i **NOTE:** This procedure documents this process using Backup Exec 2010. The procedure for Backup Exec 2012 is different. For specific details and procedures, see the product-specific documentation from Veritas for the specific DMA product and version you are using.

1. Launch the **Backup Exec Administrator** console, and select the **Job Setup** tab.
2. Click **Backup Tasks** in the left navigation panel, and select b.
3. In the left navigation pane of the **Backup Job Properties** page, select **Source**, and select **Selections**.
4. Select the system or node name in the center pane of the **Selections** page, and click the check boxes that correspond to the files you want backed up.
5. In the left navigation pane of the **Backup Job Properties** page, select **Destination**, and select **Device and Media**.
6. In the **Device** pane in the **Device and Media** page, select the OST device in the drop-down list, and click **Run Now** to start the backup job.
7. Click the **Job Monitor** tab to view the progress of the backup job you created.

Optimizing duplication between DR Series systems using Backup Exec

Backup Exec can replicate backups between two DR Series systems that are part of a defined source and target replication pair. This process uses the deduplication and replication features of the DR Series system via RDA with OST.

Using RDA with OST, backed up data is catalogued which makes it available from the designated media server so that a seamless restore can be performed from either the target or source DR Series system. This is considered an integrated replication, where the appliance does the replication. It is considered to be “optimized” because the data flows from the local appliance directly to the remote appliance in a deduplicated format, and it does not travel through the media server.

When the data is in a deduplicated format (in an optimized form), only new or unique data is copied between the two DR Series systems. Because the duplication job is initiated by Backup Exec, there are two entries in its catalog: one entry is for the source file, while the other entry is for the target file. The backup administrator can restore backup data from either appliance in case of data loss or disaster.

To optimize duplication between DR Series systems, create an additional OST device that points to the target DR Series system, and complete the following steps:

1. Launch the **Backup Exec Administrator** console, select the **Devices** tab, and right-click the target DR Series system.
2. Select **Add OpenStorage** in the drop-down list.

3. Configure the **Add OpenStorage Device** page with the following information:
 - **Server**—enter the host name or IP address of the DR Series system.
 - **Logon account**—select the account from the drop-down list (or click ... and browse to the account location), which has credentials for accessing the DR Series system.
 - **Server type**—select the type of server from the drop-down list (**DELL** or **QUEST**).
 - **Logical storage unit**—enter the name of the logical storage unit (LSU), also known as a DR Series system container, to use.
4. Click **Yes** in response to the prompt if you want to make the new device the default destination for new jobs.
5. Close the **Add OpenStorage Device** page.
6. Click the **Job Setup** tab.
7. In the left navigation pane, select **Backup Tasks**, and click **New job** to duplicate backup sets.
8. Select **Duplicate existing backup sets**, and click **OK**.
9. Click the **View by Resource** tab in the **Selections** page, and select the dataset you want copied.
10. In the left navigation pane, select **Destination**, and select **Device and Media**.
11. In **Device**, select the destination device from the drop-down list (that was created in this procedure), and click **Run Now** to start the replication operation between the two DR Series systems.
12. Click the **Job Monitor** tab to view the progress of the replication operation you created.

Restoring data from a DR Series system with Backup Exec

This topic describes how to use Backup Exec to restore data from a DR Series system.

To restore data from a DR Series system using Backup Exec, complete the following steps:

1. Launch the **Backup Exec Administrator** console, and select the **Job Setup** tab.
2. In the left navigation pane, select **Restore Tasks**, and click **New job**.
3. Click the **View by Resource** tab in the **Selections** pane, and select the dataset to be restored.
4. Click **Run Now** to start the restore job.
5. Click the **Job Monitor** tab to view the progress of the restore job operation you created.

Understanding the OST CLI commands

The `--mode` command in the DR Series system command line interface (CLI) command supports three values, which represent optimized writes done via:

- **deduplication** (`--mode dedupe`) — The client will process hashing on the data, so that deduplication processing occurs on the server side (client-side deduplication).

- **passthrough** (`--mode passthrough`) — The client will pass all data to the DR Series system for deduplication processing (appliance-side deduplication).
- **auto** (`--mode auto`) — The DR Series system will set the deduplication to Dedupe or Passthrough, based on the client's number of cores and whether it is 32- or 64-bit.

These OST commands are used in the following format: `ost --update_client --name --mode`.

i NOTE: If a RDA with OST client has four or more CPU cores, it is considered to be “dedupe-capable.” However, the client operating mode depends upon how it is configured in the DR Series system (**Dedupe** is the default RDA with OST client mode). If the administrator did not configure a client to operate in a specific mode and it is dedupe-capable, it will run in the **Dedupe** mode. If a client is not dedupe-capable (meaning the client has less than four CPU cores), and the administrator sets it to run in the **Dedupe** mode, it will only run in the **Passthrough** mode. If a client is set to run in **Auto** mode, the client will run in the mode setting determined by the media server. The following table shows the relationship between the configured client mode types and the supported client mode based on client architecture type and corresponding number of CPU cores.

Table 1: Supported RDA with OST Client Modes and Settings

Client Mode Settings	32-Bit Client (4 or more CPU cores)	64-Bit Client (4 or more CPU cores)	32-Bit Client (Less than 4 CPU cores)	64-Bit Client (Less than 4 CPU cores)
Auto	Passthrough	Dedupe	Passthrough	Passthrough
Dedupe	Not Supported	Supported	Not Supported	Not Supported
Passthrough	Supported	Supported	Supported	Supported

Supported DR Series System CLI Commands for RDA with OST

The following are the supported DR Series system CLI commands for RDA with OST operations:

```
administrator@acme100 > ost
Usage:
  ost --show [--config]
              [--file_history] [--name <name>]
              [--clients]
              [--limits]

  ost --setpassword
  ost --delete_client --name <OST Client Hostname>

  ost --update_client --name <OST Client Hostname>
      --mode <auto|passthrough|dedupe>

  ost --limit --speed <<num><kbps|mbps|gbps> | default>
      --target <ip address | hostname>

  ost --help

ost <command> <command-arguments>
```

<command> can be one of:

--show	Displays command specific information.
--setpassword	Updates the OST user password.
--delete_client	Deletes the OST client.
--update_client	Updates attributes of the OST client.
--limit	Limits bandwidth consumed by ost.

For command-specific help, please type `ost --help <command>`

For example:

```
ost --help show
```

i **NOTE:** The `--files` in the `ost --show --file_history` command represents replicated files that were processed via the DMA optimized duplication operation. This command displays only up to the last 10 such files. The `--name` in the `ost --show --name` command represents the OST container name.

i **NOTE:** For more information about OST-related DR Series system CLI commands, see the *DR Series System Command Line Reference Guide*.

Understanding RDA with OST Plug-In Diagnostic Logs

You can collect diagnostic logs for supported DMAs with the RDA with OST plug-in.

i **NOTE:** The directory location, `C:\ProgramData`, is considered to be a hidden directory on Windows-based systems. However, you can copy and paste `C:\ProgramData\Quest\DRVog\` into your Internet Explorer **Address bar** or you can enter this into the Windows command prompt window (**Start→All Programs→Accessories→Command Prompt**).

For more information about RDA with OST plug-ins and logs, see the topics that follow.

Rotating RDA with OST Plug-In Logs for Windows

By default, the Windows log rotation size is set at 10 megabytes (MB). Once a log file has been reached this size, the RDA with OST plug-in automatically renames the existing log file, `libstspioca.log` to `libstspioca.log.old`, and creates a new log.

Modifying Log Rotation Size

To modify the log rotation size, you can edit the following registry key value:

```
HKLM\Software\QUEST\OST\LogRotationSize
```

Immediately after modifying this value, the new rotation size value takes effect (meaning that you do not have to restart the backup process).

Collecting diagnostics by using a Linux utility

You can use a Linux utility called ocaDRDiags to collect diagnostics from Linux-only clients. This Linux utility is installed by the OST plug-in installer in the /opt/oca directory. The tool collects the following types of information:

- var/log/libstspioca.log.*
- usr/opensv/netbackup/logs
- usr/opensv/logs/nbemm/
- usr/opensv/logs/nbrmms/

The diagnostics file is written to the following location: /var/log/diags_client location.

The following example shows the process for collecting the RDA with OST diagnostic logs (the root user account shown represents one that resides on the media server, and is not to be confused with a root user account on the DR Series system):

```
[root@QA-OST-rh6-64 admincmd]# cd /opt/oca
[root@QA-OST-rh6-64 oca]# ./ocaDRDiags --collect
Collecting diagnostics...Done
Diagnostics location: /var/log/diags_client//QA-OST-rh6-64_2017-11-01_03-20-56.tgz
```

The default log level can be modified via the DR Series system CLI or GUI.

i | **NOTE:** For backward compatibility the /opt/dell folder also exists on the client.

Rotating RDA with OST plug-in logs for Linux

If you set the RDA with OST plug-in log level to Debug, this can cause the plug-in log to grow in size quickly. The best practice for preventing any issues with log sizes is to rotate the plug-in logs using the logrotate utility that is commonly available on Linux-based systems.

To configure log rotation, complete the following:

1. Create a file in /etc/logrotate.d/, name it "ost", and add the following entries:

```
/var/log/libstspioca.log {
    rotate 10
    size 10M
    copytruncate
}
```

2. Create a file in /etc/cron.hourly/, name it "ost_logrotate.cron", and add the following entries:

```
#!/bin/bash
/usr/sbin/logrotate /etc/logrotate.d/ost
```

The logrotate utility runs every hour, and rotates the logs whenever the log file size exceeds 10 megabytes (MB). This procedure is automated as part of the plug-in installation.

Guidelines for gathering media server information

In addition to the DR Series system diagnostics log file bundles and core files that you can collect for history and troubleshooting purposes, if you have run any RDA with OST operations, Quest recommends that you also gather some important media server-related files. This topic introduces some of these key media server files that reside on Linux and Windows platforms .

NetBackup on Linux Media Servers

For NetBackup running on a Linux media server, Quest recommends gathering the following files:

- RDA with OST plug-in configuration files and log files from the media server
 - Location: `/var/log/libstspioca.log.*`
- NetBackup backup job logs and command logs from the media server:
 - Location: NetBackup log files reside in `/usr/opensv/netbackup/logs/`. For each process in NetBackup, there is a subdirectory in the logs directory. Quest is interested in the following process-related logs: `bptm`, `bpdm`, `bprd`, `bpcd`, `bpbrm`.
 - Be aware that these five directories may not exist by default, so only collect these logs if they exist on your media server. If they were created, the locations where these log files reside are as follows: `/usr/opensv/netbackup/logs/bptm`, `/usr/opensv/netbackup/logs/bpdm`, `/usr/opensv/netbackup/logs/bpcd`, `/usr/opensv/netbackup/logs/bprd`, and `/usr/opensv/netbackup/logs/bpbrm`.
 - Quest recommends that you collect logs from the following directories: `/usr/opensv/logs/nbemmm` and `/usr/opensv/logs/nbrmms/`.
- Check for any core files that were generated on the NetBackup media server or on the DR Series system that can include:
 - Core files on a Linux NetBackup media server reside in the `/usr/opensv/netbackup/bin` directory. Most of the NetBackup binaries that link with the RDA with OST plug-in are in this directory.
 - The location of the core files on the client is not a fixed location. Verify if the core files reside in following directories: `/`, `/root/`, or the directory mentioned in the `/proc/sys/kernel/core_pattern`. For example, if the following is a `core_pattern` from a DR Series system (`/var/cores/core.%e.%p.%t`), then all the core files would reside in `/var/cores`.

Quest recommends that if `core_pattern` on the client is set by NAT to a specific directory, then the diagnostics script has to look into that directory for any related cores.

NetBackup on Windows Media Servers

For NetBackup running on a Windows media server, Quest recommends gathering the following files:

- RDA with OST plug-in configuration files and log files from the media server:
 - Location: `%ALLUSERSPROFILE%\Quest\OST\log\libstspioca.log*`

- NetBackup job logs and command logs from the media server, with log files from following directories:
 - C:\Program Files\Veritas\NetBackup\logs\bptm (if it exists)
 - C:\Program Files\Veritas\NetBackup\logs\bpdm (if it exists)
 - C:\Program Files\Veritas\NetBackup\logs\bpbrm (if it exists)
 - C:\Program Files\Veritas\NetBackup\logs\bpdr (if it exists)
 - C:\Program Files\Veritas\NetBackup\logs\bpdc (if it exists)
 - C:\Program Files\Veritas\NetBackup\logs\nbemm
 - C:\Program Files\Veritas\NetBackup\logs\nbrmms
- Any core files generated on the NetBackup media server or on the DR Series system.
- If a server failure is involved (which could be an inapparent or silent failure), the Windows media server event log for the application could be collected by using **Administrative Tools → Event Viewer**. Next, check the **Windows Logs → Application**. Typically, the last entry marked with **Error** is the one for which you are searching.
 - Copy and paste this text in the window, as shown in the following example:


```
Faulting application bptm.exe, version 7.0.2010.104, time stamp 0x4b42a78e,
faulting module libstspiocaMT.dll, version 1.0.1.0, time stamp 0x4f0b5ee5,
exception code 0xc0000005, fault offset 0x000000000002655d, process id
0x12cc, application start time 0x01cccf1845397a42.
```
 - If the system is unresponsive, force the crash of bptm.exe and complete the following:
 1. Click to open **Task Manager**.
 2. Locate the process.
 3. Right-click, and select **Create Dump File**.
 4. Retrieve the dump file from the location specified in the dialog that displays after the dump file is created.

Backup Exec on Windows Media Servers

For Backup Exec running on a Windows media server, Quest recommends gathering the following files:

- RDA with OST plug-in configuration files and log files from the media server:
 - Location: %ALLUSERSPROFILE%\Quest\OST\log\libstspioca.log*
- Backup Exec job logs and command logs from the media server.
- Any core files generated on the Backup Exec media server or on the DR Series system.
- If a crash is involved, collect any mini-dump file(s) that reside in %ProgramFiles%\Veritas\Backup Exec\BEDBG.

- If the system is unresponsive, force the crash of pvlsrv.exe and bengine.exe, and complete the following:
 1. Open Task Manager.
 2. Locate the process.
 3. Right-click, and select **Create Dump File**.
 4. Retrieve the dump file from the location specified in the dialog that displays after the dump file is created.

Configuring and using VTL

This topic introduces Virtual Tape Libraries (VTLs) and related concepts and tasks. Refer to the subsequent topics and procedures in this section for more information.

Understanding VTL

A Virtual Tape Library (VTL) is an emulation of a physical tape library on a disk-based deduplication and compression system such as the DR Series system. The tape library is exposed to a Data Management Application (DMA) as if it is a physical library with tape drives and cartridges, which the application uses for backup. Because a VTL completely emulates a standard library, the introduction of virtual tape is seamless and transparent to existing tape backup/recovery applications. The management of the library, including the drives and tapes, is done by the DMA using SCSI commands. For details on the applications supported, see the *DR Series System Interoperability Guide*.

Terminology

This topic introduces and briefly defines some basic VTL terminology used throughout the DR Series system documentation.

Term	Description
Library	A library is an emulation of a physical tape library and shares the same characteristics such as media changer, tape drives, and slots (cartridge slots).
Tape Drive	A Tape drive is a logical unit which is part of the emulated library. The media or cartridge is loaded in the Tape drives to be accessed by the Data Management application.
Tapes/Media/Cartridges	Tapes are represented as files and are units within the VTL where data is actually written. Tapes are loaded into a Tape Drive before being accessed.
Slots	Tapes are parked in Slots before they are retrieved by the data management application for access.

Supported virtual tape library access protocols

The DR Series system supports the following virtual tape library (VTL) tape access protocols.

- Network Data Management Protocol (NDMP)
- Internet Small Computer System Interface (iSCSI)
- Fibre Channel (FC)

NDMP

The Network Data Management protocol (NDMP) is used to control data backup and recovery between primary and secondary storage in a network environment. For example, a NAS server (Filer) can talk to a tape drive for the purposes of a backup.

You can use the protocol with a centralized data management application (DMA) to back up data on file servers running on different platforms to tape drives or tape libraries located elsewhere within the network. The protocol separates the data path from the control path and minimizes demands on network resources. With NDMP, a network file server can communicate directly to a network-attached tape drive or virtual tape library (VTL) for backup or recovery.

The DR Series system VTL container type is designed to work seamlessly with the NDMP protocol.

iSCSI

iSCSI or **Internet Small Computer System Interface** is an Internet Protocol (IP)-based storage networking standard for storage subsystems. It is a carrier protocol for SCSI. SCSI commands are sent over IP networks by using iSCSI. It also facilitates data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over LANs or WANs.

In iSCSI, clients are called *initiators* and SCSI storage devices are *targets*. The protocol allows an *initiator* to send SCSI commands (*CDBs*) to the *targets* on remote servers. It is a storage area network (SAN) protocol, allowing organizations to consolidate storage into data center storage arrays while providing hosts (such as database and web servers) with the illusion of locally attached disks. Unlike traditional Fibre Channel, which requires different cabling, iSCSI can be run over long distances using existing network infrastructure.

iSCSI is a low-cost alternative to Fibre Channel, which requires dedicated infrastructure except in FCoE (Fibre Channel over Ethernet). Note that the performance of an iSCSI SAN deployment can be degraded if not operated on a dedicated network or subnet.

The VTL container type is designed to work seamlessly with the iSCSI protocol. For details, see the topic, [Creating Storage Containers](#).

Fibre channel

Fibre Channel (FC) is a high-speed network technology primarily used to connect computer data storage to servers in storage area networks (SAN) in enterprise storage. Fibre Channel networks are known as a Fabric because they operate in unison as one big switch. Fibre Channel mainly runs on optical fiber cables within and between data centers. Virtual tape libraries (VTLs) can ingest data over a Fibre Channel interface, which enables seamless integration with many existing backup infrastructures and processes.

The DR Series system VTL container type is designed to work seamlessly with the FC interface. With FC, the DR Series system can direct attach to NAS filers or Fibre Channel switches and supports SAN devices.

A FC VTL container on a DR Series system supports multiple initiators, making it possible for the VTL to be shared across multiple clients of a Data Management Application (DMA).

i | **NOTE:** VTL access via FC is only available on the DR4300 and DR6300.

VTL and DR Series specifications

This topic describes key specifications of VTL support in the DR Series hardware systems.

i | **NOTE:** The use of VTL is not supported on the DR2000v.

- **Supported VTL Types** — The DR Series system supports two types of virtual tape libraries.
 - Standard emulation of StorageTek L700 library
 - Dell OEM version of the StorageTek L700 library
 - Quest OEM version of the StorageTek L700 library

i | **NOTE:** The Dell and Quest OEM type VTL is supported only with VeritasBackup Exec and Netbackup data management applications (DMAs).

i | **NOTE:** Refer to the documentation for your specific DR Series system, which includes DMA best practices whitepapers and the latest *DR Series System Interoperability Guide*, for a complete list of the supported DMAs. Visit the following site and select your specific DR Series system to download documentation: support.quest.com/dr-series.

- **Number of Tape Drives** — Each tape library contains 10 tape drives of the type IBM-LTO-4 ('ULT3580-TD4')

- **Tapes or Media Sizes**— Each library initially is created with 60 slots housing 60 tape media of the default size of 800GiB, which is the equivalent of an LTO4 tape. You can add additional tapes to the library as needed by editing the container in the GUI or by using the following CLI command:

```
vtl --update_carts --name <name> --add --no_of_tapes <number>
```

i **NOTE:** For more information about using the CLI, see the *DR Series System Command Line Interface Reference Guide*.

A library can only contain tapes of the same size. For example, if the library is originally created with 10 tapes of size 10GiB, additional tapes of size 10GiB can only be added.

Tapes of the following capacity are supported:

Tape	Size	Max number of slots supported
LTO-4	800GiB	2000
LTO-4	400GiB	4000
LTO-4	200GiB	8000
LTO-4	100GiB	10000
LTO-4	50GiB	10000
LTO-4	10GiB	10000

- **Maximum Number of DMAs or Initiators Supported** — A tape library can be accessed by one DMA or iSCSI initiator at a time.

i **NOTE:** A Fibre Channel (FC) VTL container on a DR Series system supports multiple initiators, making it possible for the VTL to be shared across multiple clients of a DMA.

Guidelines for configuring VTL

The overall steps and recommended guidelines for using and configuring a virtual tape library (VLT) with the DR Series system are described below.

Plan your Environment

Determine the following before creating a container of type VTL.

- Identify the Data Management Application (DMA) that you will be using to back up data. Refer to the *DR Series System Interoperability Guide* for a complete list of the supported DMAs.
- For the NDMP protocol, determine the filer that will be backed up using NDMP Refer to the *DR Series System Interoperability Guide* for a list of the supported Filers and Operating systems.
- For the iSCSI protocol, determine the iSCSI initiator's properties – This is the DMA IP, hostname or IQN of the software initiator on the operating system.

- For the FC protocol, determine the initiator WWPN, and create the FC switch zone and enable it. (Refer to the administrator documentation for your switch for more information.) FC zoning is required to be "single target single initiator" zoned.

i **NOTE:** Point to point cabling is not supported (directly attaching the DR Series system to another system rather than using a switch), and multi-pathing is not currently supported.

- Assess the estimated size of full and incremental backups and retention periods.

i **NOTE:** The size of the full and incremental backups will determine the tape capacity size that you set. You should use a larger tape size for full backups and a smaller size for incremental backups that have smaller retention periods. Note that faster expiration periods of incremental backups residing on smaller tapes results in the release of space back to the system for future backups.

Create Containers of Type VTL

- Determine the VTL library type (NDMP, iSCSI, or FC) that you should be using as per the suggested type in the best practices guide of your preferred DMA.
Refer to the DR Series system documentation, which includes best practices whitepapers for the supported DMAs for your specific DR Series system at: support.quest.com/dr-series
- When creating the container in the GUI or by using the CLI, you will need to set the connection type of either NDMP, iSCSI, or FC. You need to provide either the DMA IP/hostname for NDMP, the IP/hostname or IQN for an iSCSI connection type, or the initiator WWPN for FC.

Refer to the topics, [Creating Storage Containers](#) and [Creating a VTL Type Container](#), for detailed instructions about creating containers. Refer to the *DR Series System Command Line Interface Guide* for details about the CLI commands for creating containers.

Authentication/User Management Considerations

- You can use the following commands to view user information and manage passwords for the iSCSI user: `iscsi_user`, and NDMP user: `ndmp_user`.
 - `iscsi --show`
 - `ndmp --show`
 - `iscsi --setpassword`
 - `ndmp --setpassword`

Refer to the *DR Series System Command Line Reference Guide* for more details about using these commands.

- For iSCSI, you need to set the system-wide CHAP account for the DR Series system. You can add this user on the Users page in the DR Series system GUI. See the topic, [Adding a User](#), for instructions for adding an iSCSI user and password.

- For NDMP, you can set the password for ndmp_user by using the CLI or on the Users page in the DR Series system GUI. These credentials are needed for configuring the NDMP-VTL in the DMA. See the topic, [Adding a User](#), for instructions for adding an NDMP user and password.

Verify the Tape Library Creation

You can easily check that the library has been created and is available for use by using the following commands.

- Check the container properties by executing the following command:

```
container --show -verbose
```

 - Upon initial addition of the connection, the NDMP/iSCSI connection status shows as 'Added'. At this time, the library is not officially created.
 - After a few minutes, the NDMP/iSCSI connection status changes to "Available" . This status indicates that the library is online, and the tape drives and media is available for usage.
- To check the status of the virtual tape library and all the tapes in the library, you can execute one of the following commands:
 - ```
vtl -show
```
  - ```
vtl --show --name <container_name> --verbose
```

Configure the Library in the DMA

See the DR Series system documentation, which includes DMA best practices whitepapers for your specific DR Series system at:
support.quest.com/dr-series.

Configuring and Using Encryption at Rest

This chapter introduces the concept of Encryption at Rest as used by the DR Series system as well as related concepts and tasks.

i **NOTE:** Due to export regulations, the encryption at rest feature is not available in certain markets, and, therefore, may not be available in your locale.

Understanding Encryption at Rest

Data that resides in the DR Series system can be encrypted. When encryption is enabled, the DR Series system uses the Industry standard FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) encryption algorithm for encrypting and decrypting user data. The content encryption key is managed by the key manager, which operates in either a Static mode or an Internal mode. In Static mode, a global, fixed key is used to encrypt all data. In internal mode, key lifecycle management is performed in which the keys are periodically rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days. A user-defined passphrase is used to generate a pass phrase key, which is used to encrypt the content encryption keys. It is mandatory to define a passphrase to enable encryption. The system supports up to a limit of 1023 different content encryption keys. All streams of a data-store are encrypted or re-encrypted with the same content encryption key. DR Series system statistics report the amount of data encrypted and decrypted bytes consistently.

Encryption at Rest Terminology

This topic introduces and briefly defines some basic encryption at rest terminology used in the DR Series system documentation.

Term	Description
Passphrase	A passphrase is a sequence of words or other text used to control access to data, similar to a password in usage, but is generally longer for added security. In the DR Series system, the passphrase is user-defined and is used to generate a passphrase key that encrypts the file in which the content encryption keys are kept. The passphrase is a human readable key, which can be up to 255 bytes in length. It is mandatory to define a passphrase to enable encryption.
Content encryption key	The key used to encrypt the data. The content encryption key is managed by the key manager, which operates in either a static mode or an internal mode. The system supports up to a limit of 1023 different content encryption keys.
Key management mode	The mode of key lifecycle management as either static or internal.
Static mode	A global mode of key management in which a fixed key is used to encrypt all data.
Internal mode	A mode of key lifecycle management in which the keys are periodically generated and rotated. The minimum key rotation period before the content encryption key can be rotated and a new key is generated is 7 days. This rotation period is user-configurable and can be specified in days.

Encryption at Rest and DR Series Considerations

This topic describes key features and considerations of using Encryption at Rest in the DR Series system.

- **Key Management** — In internal mode there is a maximum limit of 1023 keys. By default when encryption is enabled on the system, the key rotation period is set to 30 days. Users can later change the key rotation period from 7 days to 70 years, while configuring internal mode of encryption.
- **Performance Impacts** — Encryption should have minimal to zero impact on both backup and restore workflows. It should also have no impact on the replication workflows.
- **Replication** — Encryption must be enabled on both the source and target DR Series systems to store encrypted data on the systems. This means that encrypted data on the source does not automatically imply that when it is replicated to the target it will be encrypted unless encryption is explicitly turned 'ON' on the target DR Series system.
- **Seeding** — Encryption must be enabled on both the source and target DR Series systems to store encrypted data on the systems. If seeding is configured for encryption, then the data will be re-encrypted and stored. When the data stream is imported onto the target from the seed device, the stream will be encrypted as per the target policy and stored.
- **Security Considerations for Passphrase and Key Management** —
 - A passphrase is very important part of the encryption process on the DR Series system as the passphrase is used to encrypt the content encryption key or keys. If the passphrase is compromised or lost, the administrator should change it immediately so that the content encryption keys do not become vulnerable.
 - The administrator should closely consider security requirements to drive the decision for selecting the mode of key management for the DR Series system.
 - The Internal mode is more secure than the Static mode since the keys are periodically changed. Key rotation can be set to 7 days minimum.
 - Key modes can be changed at any time during the lifetime of the DR Series system; however, changing the key mode is a significant operation to undertake as all encrypted data must be re-encrypted.
 - Content encryption keys are stored in their encrypted form in a primary keystore, which is maintained on the same enclosure as the data-stores. For redundancy purposes, a backup copy of the primary keystore is stored on the system in the root partition, separate from the data-store partitions.

Understanding the encryption process

The overall steps for how Encryption at Rest is enabled and used in the DR Series system are described below.

1. **Enabling encryption.**

Encryption is disabled by default on a factory-installed DR Series system (running version 3.2 software or later) or a DR Series system that has been upgraded to version 3.2 from a previously released version. An administrator can enable encryption by using the GUI or CLI.

Encryption is set at the storage group level.

2. **Setting a passphrase and setting the mode.**

When defining encryption for a storage group, a passphrase is set. This passphrase is used to encrypt the content encryption keys, which adds a second layer of security to the key management. At this time, the mode is also set. The default key management mode is “internal” mode, in which key rotation happens periodically as specified by the set key rotation period.

3. **Encryption process.**

After encryption is enabled, the data in the storage group that gets backed up is encrypted and is kept encrypted until it is expired and cleaned by the system cleaner. Note that the encryption process is irreversible.

4. **Encryption of pre-existing data.**

Any pre-existing data will also be encrypted using the currently set mode of key management. This encryption occurs as part of the system cleaner process. Encryption is scheduled as the last action item in the cleaner workflow. You must launch the cleaner manually using the maintenance command to reclaim space. It then encrypts all pre-existing unencrypted data. The cleaner can also be scheduled as per the existing pre-defined cleaner schedule.

i NOTE: The cleaner can take some time to start the encryption process if the system is nearing full system capacity. Encryption starts only after the cleaner processes data slated for cleaning and the related logs. This ensures that space reclamation is prioritized when free space is low and also ensures that data stores are not redundantly encrypted.

Refer to the *DR Series System Command Line Interface Reference Guide* for information about the CLI commands used for encryption.

Support, maintenance, and troubleshooting

The DR Series system GUI provides various information and tools that can help you better understand the current state of your system and that provide basic, support, maintenance, and troubleshooting functionality.

Using the DR Series system support options

In the DR Series system GUI, the Support page provides valuable support options and functionality to help you maintain the state of your DR Series system. These Support options include Diagnostics and Software Upgrades. From the Support page, you can also reboot or shut down the system.

Viewing support information

The **Support** page displays valuable support information for your system. This information can be useful if you need to contact Technical Support for any technical assistance.

To view the Support page, follow these steps.

1. In the left navigation menu, click **Support**.
2. On the Support page, you can view the following information:
 - **Product Name**—The DR Series system product name.
 - **Software Version**—DR Series system software version installed.
 - **Service Tag**—DR Series system appliance bar code label.
 - **Last Diagnostic Run**—Timestamp of latest diagnostics log file.
 - **BIOS Version**—current version of installed BIOS.
 - **MAC Address**—current address in standard two-digit hexadecimal grouping format.
 - **iDRAC IP Address**—current IP address of iDRAC (if applicable).
 - **NIC Ports**—displays information about bonded ports only (if the 10-GbE NICs are installed, it only displays information about the two supported 10-GbE ports):
 - Eth0 MAC address and port speed
 - Eth1 MAC address and port speed
 - Eth2 MAC address and port speed in
 - Eth3 MAC address and port speed in
 - This page also displays information about the FC cards on PCI slot.

Viewing system diagnostic log files

A DR Series system diagnostics log file is a bundle that contains a variety of file types that record the latest system settings and saves them in a compressed .Izip file format.

In the DR Series system GUI, the Diagnostics page allows you to generate diagnostic logs that capture the state of your system. You can also download these log files or delete them as needed.

To view the system diagnostics page, follow these steps.

1. In the left navigation menu, click **Support > Diagnostics**.
2. You can view the following columns of information on the Diagnostics page for the diagnostics logs that have been generated.

- **File name**—in this format, `<hostname>_<date>_<time>.lzip`, as in this example: **acme-sys-19_2012-10-12_13-51-40.lzip**



NOTE: Diagnostic log file names are limited to 128 characters.

- **Size**—in Megabytes.
- **Time**—timestamp of when the log file was created.
- **Reason**—describes the reason the log file was generated (for example, [admin-generated]: generated by Administrator).



NOTE: Diagnostic reason descriptions are limited to 512 characters, and the descriptions can only be added using the DR Series system CLI.

- **Status**—indicates the status of the log file (for example, Completed or In Progress).

Understanding diagnostics collection

The Diagnostics service in the DR Series system lets you collect and manage your system's diagnostic log file bundles. The Diagnostics service works by collecting all the system-related information that could help when diagnosing a problem or error condition in the system. Each diagnostic log file bundle provides:

- A current snapshot of system operations
- System-related information that assists in understanding system operations
- A record of system operations in case Technical Support needs to provide technical assistance

Diagnostics runs as a service during system startup, and this process listens for incoming requests. There are two modes in which the diagnostics collection process is started:

- **Admin-Generated mode:** when a DR Series system CLI or DR Series system GUI request is made by the administrator (and the default reason that is listed is admin-generated).
- **Auto-Generated mode:** when a process or service failure is reported, the DR Series system starts collecting system-related information. After it completes the auto-generated collection, it generates a system event.

When the diagnostics log directory exceeds the maximum storage capacity, any log older than one hour is automatically deleted. The DR Series system GUI lets you download and save diagnostics log files to other systems on your network. The DR Series system also maintains a separate archive logs directory that collects other system-related information, and these archive logs are also automatically deleted when they exceed a maximum capacity. When you generate a diagnostics log file bundle, it contains all of the DR Series system information that you need when contacting Technical Support for technical assistance. The diagnostics log file bundle collects the same type of hardware, storage, and operating system information that is collected when using the Dell System E-Support Tool (DSET) and the DR Series system CLI commands (diagnostics --collect --dset). When a diagnostics log file bundle is generated, this process also collects all the previous auto-generated diagnostics and deletes them from the system.

The Diagnostics service collection process observes the following guidelines:

- DR Series system triggers an automatic diagnostic log collection of the DR Series system status for any system process or service failures.
- All automatic diagnostic collection requests are queued and executed sequentially.

The DR Series system GUI provides options to display existing diagnostics logs, generate new diagnostics logs, download and save copies of existing diagnostics logs, or delete existing diagnostics logs. The DR Series system CLI also provides the means for managing, generating, or downloading the diagnostics log files. For more information, see the *DR Series System Command Line Reference Guide*.

Generating a diagnostics log file

A DR Series diagnostics log file is a bundle that contains a variety of file types that record the latest system settings, and saves them in a compressed .zip file format. When you generate a diagnostics log file bundle, it contains all of the DR Series system information that may be needed when contacting Technical Support for technical assistance. This also includes all the previous auto-generated diagnostics log files, which are then deleted from the DR Series system. The diagnostics log file bundle collects the same type of hardware, storage, and operating system information collected by the Dell System E-Support Tool (DSET) from the DR Series system appliance hardware:

- To collect a DSET log file, use the DR Series system CLI command, `diagnostics --collect --dset`.
- To collect the comprehensive DR Series system diagnostics log file bundle (which also includes DSET information), use the DR Series system CLI command, `diagnostics --collect`.

To generate a diagnostics log file bundle for your system, complete the following steps:

1. In the left navigation menu, click **Support > Diagnostics**.
2. In the **Action** menu in the upper right corner of the page, click **Generate**.
3. In the **Select Diagnostics Type** drop-down list, select one of the following options.
 - basic
 - dset
 - logs
 - cores
 - tcpdump
 - all
4. To verify that a new diagnostics log file is being generated, check the status of the diagnostics log file on the **Diagnostics** page.

Once completed, the new diagnostics log file resides at the top of the File Name column in the table. To verify, check its timestamp (using its date and time), to ensure this is the latest diagnostics file created.

Downloading diagnostics log files

To download an existing diagnostics log file, complete the following steps:

1. In the left navigation menu, click **Support > Diagnostics**.
2. In the list, select the diagnostics log file you want to download, and click the **Download** icon.
3. Download and save the file as needed.

Deleting a Diagnostics Log File

To delete an existing diagnostics log file from the Diagnostics summary table on the Diagnostics page, complete the following:

1. Select **Support > Diagnostics**.
2. Click **Select** to select the diagnostics file you want to delete, and click **Delete**.
3. Click **OK** to delete the selected diagnostics log file (or click **Cancel** to display the **Diagnostics** page).

Rebooting the DR Series system

If needed, you can reboot the DR Series system.

To reboot your system, follow these steps:

1. In the left navigation menu, click **Support**.
2. In the **Action** menu in the upper right corner of the page, click **Reboot**.
3. In the **Confirmation** dialog box, click **Yes** to proceed with rebooting the system.

Shutting down the DR Series system

If needed, you can shut down the DR Series system in the DR Series system GUI.

CAUTION: You should fully understand what this action means to system operations before attempting to shut down the system. A system shutdown powers off the appliance on which the DR Series system software is installed. Once powered Off, you can only power it On again at its physical location, or you must use an iDRAC connection to the DR Series system.

NOTE: To shutdown the DR Series system using a UPS after a power loss, refer to the following article for information on how to do this using the shutdown command in the IPMI interface:
<http://www.dell.com/downloads/global/power/ps4q04-20040204-murphy.pdf>.

To shutdown your DR Series system, complete the following:

1. In the left navigation menu, click **Support**.
2. In the **Actions** menu in the upper right corner of the page, click **Shutdown**.
3. In the **Confirmation** dialog box, click **Yes** to proceed with shutting down the system.

Viewing the software upgrade page

On the Software Upgrade page you can view the currently installed version of the DR Series system software as well as the upgrade history. .

To view the Software Upgrade page, follow these steps.

1. On the left navigation menu, click **Support > Software Upgrade**.
The Software Upgrade page is displayed.
2. On this page you can view the Current Version of the DR Series system software that is installed as well as the Upgrade History of previously installed software versions.

Upgrading the DR Series system software

You can obtain the latest DR Series system software binary package from the Quest website. You would then upload this file in the DR Series system GUI for the system software upgrade.

i **NOTE:** The DR Series system only supports the copying of upgrade images and diagnostics files to and from the system using WinSCP. The DR Series system does not support the copying or deleting of any other file types using WinSCP. To use WinSCP to copy DR Series software upgrade and diagnostics log files, ensure that the File Protocol mode is set to SCP (Secure Copy) mode.

i **NOTE:** You can use other SCP tools with the DR Series system, but you cannot use these other SCP tools to copy other types of files to or from the DR Series system.

To upgrade the DR Series system software, complete the following steps.

1. Using your browser, go to support.quest.com/dr-series, select your specific DR model and then navigate to **Software Downloads**.
2. Under **Upgrade File**, locate the version you want to download, and click the **Download** button. If you are not logged in, you may be prompted to log in with your registered Quest account.
3. Save the latest system software upgrade file to a network location accessible by the DR Series system that is running the browser session started by the DR Series administrator.
4. In the DR Series system GUI, in the left navigation menu, click **Support > Software Upgrade**.
5. Click the button, **Select DR Binary Package**.
6. Browse to the location of the file you downloaded from quest.com, select the file, and click **Open**.
The system verifies that the file is the proper format.

7. Once the file has been uploaded, on the Software Upgrade page, click Start Upgrade.

The administrator who initiated the software upgrade sees a System Information pane that displays an alert and upgrade status, as well as the Current Version and Upgrade History versions of the DR Series system software listed.

The first time you log on after a software upgrade, you will see an Upgrade notification dialog box, which instructs you to review and verify user roles and logon information for your DR Series system and associated DMAs.

There are only three possible outcomes for a DR Series system software upgrade operation:

- Upgrade has completed successfully—no reboot is required.
- Upgrade has completed successfully—but a reboot is required (click Reboot in the Software Upgrade page).
- Upgrade has failed.

i **NOTE:** If the DR Series system software upgrade operation fails, you can reboot the system and attempt another software upgrade operation using the DR Series system GUI. If this is unsuccessful, you can use the DR Series system CLI `show` command to view the current System State status. DR Series system software upgrades can also be performed using the DR Series system CLI. For details, see the *DR Series System Command Line Reference Guide*. If both the DR Series system GUI and CLI attempts are unsuccessful, contact Technical Support for assistance.

Restore Manager (RM)

The **Restore Manager** (RM) utility can be used to restore the DR Series system software. RM can be used when a non-recoverable hardware or software failure prevents the DR Series system from functioning correctly.

RM can also be used to reset the system back to its initial factory settings when moving it from a test environment to a production environment. RM supports the following two modes:

- **Recover Appliance**—in Recover Appliance mode, RM reinstalls the operating system and attempts to recover the prior system configuration and the data residing in the containers.

i **NOTE:** To use the Recover Appliance mode, you must use an RM build that is compatible with the DR Series system software version that was running before the OS reset was attempted.

- **Factory Reset**—in a Factory Reset mode, RM reinstalls the operating system and resets the system configuration back to the original factory state. It is important to note that when doing a factory reset, all of the containers and the data in the containers gets deleted.

! **CAUTION:** Using the Factory Reset mode deletes all of the DR Series system data. The Factory Reset mode must only be used when the container data is no longer needed.

Downloading the Restore Manager

The **Restore Manager** (RM) utility runs from a USB boot key that contains the RM image, which must first be downloaded from the Quest Support site.

1. Using your browser, go to support.quest.com/dr-series, select your specific DR model and then navigate to **Software Downloads**.
2. Under **Restore Manager**, locate the version you want to download, and click the **Download** button. If you are not logged in, you may be prompted to log in with your registered Quest account.

Creating the Restore Manager USB key

To create a Restore Manager (RM) USB key, you must first download the RM image (.img) file from the Quest Support site, and then transfer it to a USB key. The USB key must be a minimum of 4 GB (Gigabytes) in size or larger. Windows USB image tools can be used to transfer the RM image when they meet the following conditions:

- Support using the .img file format
- Support using a direct block-to-block device copy to ensure that the USB key is bootable

To transfer the RM image to the USB key on a Linux or Unix system, perform the following:

1. Copy the downloaded RM image file to a Linux or Unix system.
2. Insert the USB key into an available USB port on the Linux or Unix system.
Make note of the device name that is reported by the operating system (for example, /dev/sdc4).
3. Do not locally mount the USB device to a file system at this time.
4. Copy the RM image to the USB key using the **dd** command:

```
dd if=<path to .img  
file> of=<usb device> bs=4096k
```

For example:

```
dd if=/root/DR-RM-1.05.03.313-2.1.0851.2.img  
of=/dev/sdc4 bs=4096
```

Running the Restore Manager (RM)

To run the **Restore Manager** (RM) utility, boot the DR Series system using the RM USB key created in [Creating the RM USB key](#).

1. Insert the RM USB key into an available USB port on the system.
You can also use the virtual media option of iDRAC to remotely load the RM USB key. For more information, see *Configuring and Using Virtual Media in the Integrated Dell Remote Access Controller 6 (iDRAC6) User Guide* at support.dell.com/support/edocs/software/smdrac3/.
2. Boot the DR Series system using the RM USB key.
3. During the time when the Power-On Self-Test (POST) screen displays, press **F11** to load the Boot Manager.
4. Within the Boot Manager, navigate to the system hard drive (C:), select the USB key as the boot device, and press **<Enter>**.
5. After a few minutes, Restore Manager loads and displays its main screen.

6. Select the desired Restore mode (either **Recover Appliance** or **Factory Reset**).
7. Enter the confirmation string, and press **<Enter>** to proceed.

CAUTION: The Factory Reset mode deletes all DR Series data. The Factory Reset mode should only be used when the container data is no longer needed.

NOTE: After Restore Manager completes, only the administrator account will remain enabled. To re-enable the root or service accounts, see the DR Series system CLI user `--enable --user` command in the *DR Series System Command Line Reference Guide*.

NOTE: If you had previously joined the DR Series system to any Active Directory Services (ADS) domain before running Restore Manager, after it completes you will need to manually rejoin the desired ADS domain. For information about joining an ADS domain, see [Configuring Active Directory Settings](#).

Resetting the Boot LUN Setting in PERC H700 BIOS After Running RM

In the event that both of the 2.5-inch 300 GB 10K RPM 6 GB/s SAS internal drives (OS) in RAID1 are replaced, you must run the Restore Manager (RM) utility to recover the DR Series system OS drives.

Following the RM recovery process, the boot logical unit number (LUN) has to be reset to VD0 RAID1. The DR Series system unsuccessfully attempts to boot from RAID6 instead of RAID1.

To resolve this issue, reset the PERC H700 BIOS to revise the proper boot order setting to configure the proper boot LUN to be RAID1. To reset the proper LUN boot order, complete the following steps:

1. Start **Restore Manager**.
2. Select **Option 1 > Recover My Appliance**.
3. Click **Proceed**.
4. Click **Reboot**, and during reboot, press **Ctrl+R** to enter the PERC BIOS.
5. Select **Controller 0: PERC H700** in the list.
6. Press **Ctrl+N** twice to select the **Ctrl Mgmt (Controller Management)** tab.
7. Select **Ctrl Mgmt**, click **Select bootable VD**, and select VD 0 as the VD0 RAID1.
8. Click **Apply**, and reboot the DR Series system.

Hardware removal or replacement

To properly remove or replace any DR Series system hardware, you must observe and use the best practice shut down and start up procedures. For a comprehensive set of removal and replacement procedures with step-by-step instructions, see the *DR Series System Owner's Manual*.

For more information about the best practices, see [DR Series System: Proper Shut Down and Start Up and Shutting Down the DR Series System](#).

Proper shutdown and startup of the system

Before you attempt to remove or replace any hardware component in the DR Series system, ensure that you observe the following best practices for properly shutting down and starting up the system.

1. Power off the DR Series system by selecting **Shutdown** in the **Support** page of the DR Series system GUI.
For more information, see [Shutting Down the DR Series System](#). Another method you can use to shut down the system is the DR Series system CLI command, `system --shutdown`.
2. Allow the DR Series system to complete its power-down process.
When the power-down process has completed, the Power Supply status indicator is unlit.
3. Disconnect the DR Series system power cables from the electrical power outlet.
4. Wait an additional period of time (up to 10 minutes), and/or verify that all of the green and amber NVRAM LEDs on the rear panel of the system chassis are unlit.

i **NOTE:** If you do not allow the NVRAM super capacitor sufficient time to discharge, the NVRAM status will report **DATA LOSS** when the DR Series system is subsequently powered up.
5. Unlatch the latch release lock and slide the DR Series system cover back and away to gain entry to the appliance internal components.
To gain entry to the interior of the DR Series system, remove the cover. For more information, see the *DR Series System Owner's Manual*.
6. Remove and replace the system hardware components as needed.
7. Replace the cover, and reconnect the system power cables to the electrical power outlet.
8. Power on the DR Series system by pressing the power-on indicator/power button.

DR Series System NVRAM

NVRAM is a field replaceable unit (FRU) in the DR Series system (*Note:* NVRAM is not part of the DR4300e Core system). The super capacitor that powers the NVRAM double-data rate (DDR) memory must be able to move its contents to the solid-state drive (SSD) during a power loss.

This data transfer requires maintaining the power to run the system for 3 minutes of operation (normally, it only takes approximately one minute). If a problem occurs during the data backup to the SSD, the subsequent system reboot detects this. NVRAM can experience backup failure when the following occurs:

- The NVRAM failed to backup the data during the power shutdown
- The super capacitor did not maintain sufficient power to backup the DDR contents into the SSD.
- The NVRAM/SSD encountered an end-of-line (EOL) or another error.

If any of these conditions occur, the NVRAM requires either a failure recovery or a replacement.

i **NOTE:** Quest recommends the following supported method for flushing DR Series system data from the NVRAM to the RAID6 before replacing the NVRAM by using either of the following DR Series system CLI commands: `system --shutdown` or `system --reboot`.

i **NOTE:** If you need to remove or replace the NVRAM in the DR Series system, see the topics, [Shutting Down the DR Series System](#) and [NVRAM Field Replacement](#).

NVRAM Backup Failure Recovery

After you have physically replaced the NVRAM card in a PCIe x4 (or x8) slot in the DR Series system chassis, you can recover from an NVRAM backup failure by completing the following tasks:

- ! CAUTION:** You must wait a minimum of 20 minutes after powering on the DR Series system before using the DR Series system CLI maintenance `--hardware --reinit_nvram` command. This 20-minute post power-on waiting period allows the NVRAM card, the super capacitor calibration, and all solid state drive (SSD) processes to fully complete, which are necessary for the proper operation of the DR Series system.

During Maintenance mode, the DR Series system determines, detects, and repairs the data loss. During the system reboot process, it ensures that no valuable data remains on the NVRAM.

1. Enter the following DR Series system CLI command: `maintenance --hardware --reinit_nvram`. This formats the SSD and clears all of the backup and restore logs, by reinitializing the NVRAM.
2. Verify that the DR Series system enters its Maintenance mode.
For more information about replacing the NVRAM, see [NVRAM Field Replacement](#) and [DR Series System: Proper Shut Down and Start Up](#).

NVRAM Field Replacement

Whenever the DR Series system NVRAM is replaced in the field, you must observe this best practice procedure:

- ! CAUTION:** You must wait a minimum of 20 minutes after powering on the DR Series system before using the DR Series system CLI command: `maintenance --hardware --reinit_nvram`. This post power-on waiting period allows the NVRAM card, the super capacitor calibration, and the SSD processes to fully complete, which are necessary for the proper operation of the DR Series system.

- i NOTE:** For more information, see [DR Series System: Proper Shut Down and Start Up](#).

1. Verify that the DR Series system software detects the NVRAM as being new to the system.
2. Enter the following DR Series system CLI command: `maintenance --hardware --reinit_nvram`. This command initializes the NVRAM, creates new partitions, and updates information used internally by the DR Series system software.
3. Verify that the DR Series system enters Maintenance mode.
If properly initialized, the DR Series system will automatically enter Maintenance mode. The filesystem checker examines every blockmap and datastore to determine how much data was lost due to the failed NVRAM.

Troubleshooting error conditions

To troubleshoot error conditions that disrupt your normal DR Series system operations, complete the following:

1. Generate a DR Series system diagnostics log file bundle if one has not already been automatically created.
For more information, see [Generating a Diagnostics Log File](#).

2. Check the system alert and system event messages to determine the current status of your DR Series system.
For more information, see [DR Series System Alert and Event Messages](#), [Monitoring System Alerts](#), and [Monitoring System Events](#).
3. Verify if the DR Series system has recovered or whether it has entered into Maintenance mode.
For more information, see [About the DR Series System Maintenance Mode](#).
4. If you cannot resolve the issue using the information in this DR Series system documentation, contact Technical Support.

DR Series system alert and event messages

The DR Series system provides a variety of system alert and system event message types that describe the current state of the system. You can review these messages, and see if there are any actions you can perform to resolve any reported issue.

You may be able to resolve any basic issues using the information presented in the DR Series system documentation. You should refer to the material in this and other related topics before:

- Making any attempt to troubleshoot your DR Series system.
- Contacting Technical Support for technical assistance.

Some alert and event messages are purely informational, and provide general system status. Other alert and event messages display specific status or component information or suggest a specific task you can perform to resolve an issue or to verify that a condition exists.

There are still other alert and event messages that direct you to contact Technical Support for assistance, where Technical Support intervention may be required.

- Table 1 lists the DR Series System Alert Messages by system alert type: general system, system chassis, NVRAM, and PERC-specific alert messages that could be displayed during the course of backup and deduplication-related operations.
- Table 2 lists the DR Series System Event Messages by system event type (type 1 through 7): event messages that could be displayed during the course of backup, replication, deduplication, diagnostics, cleaner, DataCheck, maintenance, and OpenStorage Technology (OST) operations.

Table 2: DR Series System Alert Messages

Alert Message	Description/Meaning or Action
General System Alerts	
Filesystem scan requested.	System is switching to Maintenance mode. Filesystem has read-only access.
NVRAM not detected.	Ensure that the NVRAM card is seated properly.
NVRAM capacitor is disconnected.	Contact Technical Support for possible support assistance or intervention.
NVRAM capacitor has degraded.	Contact Technical Support for possible support assistance or intervention.

Alert Message	Description/Meaning or Action
NVRAM solid-state drives (SSD) are disconnected.	Contact Technical Support for possible support assistance or intervention.
NVRAM has failed to backup or restore data during the last boot.	Contact Technical Support for possible support assistance or intervention.
NVRAM hardware failure.	Contact Technical Support for possible support assistance or intervention.
Data volume is not present. Check that all drives are installed and powered up.	Contact Technical Support for possible support assistance or intervention.
File server failed to start after multiple attempts.	Contact Technical Support for possible support assistance or intervention.
File server failed multiple times. Entering Maintenance mode.	Contact Technical Support for possible support assistance or intervention.
Insufficient disk space exists.	The filesystem is now read-only.
Unable to detect filesystem type on the Data volume.	Contact Technical Support for possible support assistance or intervention.
Unable to detect filesystem type on the Namespace volume.	Contact Technical Support for possible support assistance or intervention.
Filesystem scan discovered inconsistencies.	Please check the filesystem report, and perform the suggested action. Contact Technical Support for possible assistance or intervention.
Replication peer network disconnected.	Check access to remote site.
NVRAM does not match the data volume.	If this is a newly replaced NVRAM, use the <code>maintenance --hardware --reinit_nvram</code> command to reinitialize the NVRAM. For more information, see the <i>DR Series System Command Line Reference Guide</i> .
Storage usage is approaching the system capacity.	Clean up the filesystem. If issue persists, contact Technical Support for possible assistance or intervention.
Replication re-sync cannot proceed.	Namespace limit has reached its maximum.
Out of space on replication target.	Clean up the filesystem. If issue persists, contact Technical Support for possible assistance or intervention.
The filesystem has reached the maximum allowable limit for files and directories. Creating new files and directories will be denied.	Clean up the filesystem. If issue persists, contact Technical Support for possible assistance or intervention.
Appliance available storage level reached VTL threshold, unload all drives, expire old backups and schedule filesystem cleaner. Run <code>"vtl --set_rw ..."</code> to set the containers IO mode back to Read-Write.	Unload all the drives and ensure no carts are loaded. Clean up the filesystem. To use the library again, set the VTL containers to read-write mode using the CLI command <code>"vtl --set_rw"</code> , after the filesystem comes back to read-write mode.
System Chassis Alerts	
Power Supply <number> detected a failure.	<ul style="list-style-type: none"> Reconnect the power cable to the designated power supply unit if it is disconnected.

Alert Message	Description/Meaning or Action
	<ul style="list-style-type: none"> • Ensure that there is input AC power at the power cable. • Use a different power cord. <p>If this does not resolve the issue, replace the designated power supply.</p>
Power Supply <number> is missing or has been removed.	<ul style="list-style-type: none"> • The power supply might not be making a proper connection. • Try reseating the power supply in the power supply slot. • Reconnect the power cable to the designated power supply unit if it is disconnected. • Ensure that there is input AC at the power cable. • Use a different power cord. <p>If this does not resolve the issue, replace the designated power supply.</p>
Power Supply <number> is unplugged.	<ul style="list-style-type: none"> • Reconnect the power cable to the designated power supply unit if it is disconnected. • Ensure that there is input AC power at the power cable. • Use a different power cord.
Fan <number> failed.	<ul style="list-style-type: none"> • Verify that the designated cooling fan is present and is installed correctly. • Verify that the designated cooling fan spins up and runs. <p>If this does not resolve the issue, replace the designated cooling fan.</p>
Fan <number> is missing.	Attach or replace the designated missing cooling fan.
Abnormal network errors detected on Network Interface Controller <number>.	<p>The Network Interface Controller errors could be caused by network congestion or by packet errors.</p> <ul style="list-style-type: none"> • Check your network. If that does not resolve the problem, then replace the NIC. • If the NIC is embedded, the DR Series system appliance requires service.
Network Interface Controller is missing.	<ul style="list-style-type: none"> • Remove and reinsert the NIC. • If this does not resolve the problem, replace the NIC.
Network Interface Controller <name> is disconnected.	Connect it to a network and/or check your network switches or routers for any network connectivity issues.
Network Interface Controller <name> is disabled.	Enable the port on the designated NIC.
Network Interface Controller <name> driver is bad.	Upgrade the DR Series system appliance (in the Software Upgrade page, and click Start Upgrade).

Alert Message	Description/Meaning or Action
CPU <name> failed.	Replace the designated failed processor.
CPU <name> is missing.	Reinsert or replace the designated missing processor.
DIMM <name> failed.	Replace the designated failed DIMM (Dual In-line Memory Module) device.
DIMM <name> is missing.	<ul style="list-style-type: none"> • Reinsert or replace the designated DIMM device. • The memory capacity of the storage appliance is below the minimum required for correct operation. • The storage appliance requires service.
Temperature probe <name> failed.	The storage appliance requires service.
Voltage probe <name> failed.	The storage appliance requires service.
Temperature probes have recorded temperatures in the failed range.	<ul style="list-style-type: none"> • Check the Events page in the DR Series system for specific temperature events and the location of the temperature probes. • Check the data center air conditioning, ventilation, and internal system cooling fans for any problems. • Ensure there is proper air flow through the storage appliance, and as needed, clean the cooling vents.
Voltage probes have recorded readings in the failed range.	<ul style="list-style-type: none"> • Check the Events page in the DR Series system for specific voltage events and the location of the voltage probes. • Check the power supplies. If there are no issues with the power supplies, have a service technician check the DR Series system appliance to see if it requires any servicing.
Storage Controller <number> failed.	Replace the RAID controller in the DR Series system.
Storage Controller <number> is missing.	Reinsert or replace the RAID controller in the DR Series system.
Storage Controller <number> has an illegal configuration.	<p>The expected number of virtual drives is <number>, and the actual number of virtual drives found was <number>. Run the Restore Manager (RM) utility to repair the drive configuration mismatch.</p> <p>The expected number of enclosures is <number>, and the actual number of enclosures found was <number>.</p> <ul style="list-style-type: none"> • Check the SAS cable connections between the storage controller and all its enclosures. • Check the power cable connections to the enclosure power supplies.
Physical disk <number> failed.	Replace the physical disk that failed.
Physical disk <number> is missing, removed, or it cannot be detected.	Reinsert or replace the physical disk.

Alert Message	Description/Meaning or Action
Physical disk <number> predictive failure reported.	Replace the physical disk. Note: Even though the disk may not have failed yet, the recommended best practice is to replace the disk.
Physical disk <number> is an unsupported type.	This disk type is unsupported and cannot be used in this configuration. Replace the unsupported physical disk with a supported SAS physical disk.
Physical disk <number> has been manually set to offline with a configuration command.	Remove the physical disk and reinsert it (the drive is non-operational in this state).
Physical disk <number> is foreign.	This can occur when a storage controller has been replaced or all drives have been migrated from another system. In such cases, the foreign configuration should be imported. If this is seen on a single physical disk, the foreign configuration should be cleared. Note: This condition can also be seen when a drive is removed and reinserted while a rebuild is still in progress.
Virtual Disk <number> failed.	Replace any failed or missing physical disk(s) and run the Restore Manager (RM) utility.
Virtual Disk <number> has an invalid layout.	Run the Restore Manager (RM) utility to repair this installation.
Virtual Disk Virtual Disk <number> redundancy is degraded. Replace the physical disk(s) with a supported SAS physical disk(s).	One or more physical disks have failed and therefore Virtual disk redundancy is degraded. Once the failed disks are replaced, the system will attempt to rebuild the redundancy automatically.
<device> failed.	<ul style="list-style-type: none"> Verify that the device is present, and then check that the cables are properly connected. For more information, see the <i>DR Series System Owner's Manual</i> to verify the system cabling is correct. Check the connection to the controller battery and the status of battery health. If none of these steps resolve the problem, replace the storage controller battery.
<device> is missing.	<ul style="list-style-type: none"> Verify that the device is present, and then check that the cables are properly connected. For more information, see the <i>DR Series System Owner's Manual</i> to verify the system cabling is correct. Check the connection to the controller battery and the status of battery health. <p>Note: A battery with a weak or depleted charge can cause this warning.</p>
Storage <device> has failed.	Check cable connections between the storage controller and the enclosure or backplane.
Storage <device> is missing.	Perform the following:

Alert Message	Description/Meaning or Action
	<ul style="list-style-type: none"> • Check SAS and power cable connections between the storage controller and the enclosure or backplane. • Check the external enclosure management modules (EMM) and PERC status LEDs.
NVRAM Alerts	
NVRAM PCI Controller failed.	Replace the NVRAM PCI Controller.
NVRAM PCI Controller is missing.	Reinsert or replace the NVRAM PCI Controller.
Super Capacitor on the NVRAM PCI Controller failed.	Replace the NVRAM PCI Controller.
Super Capacitor on the NVRAM PCI Controller is missing.	Replace the NVRAM PCI Controller.
Failed to check software compatibility.	Upgrade the DR Series system appliance (in the Software Upgrade page, click Start Upgrade).
The system software package is incompatible with the current software stack.	Upgrade the DR Series system appliance (in the Software Upgrade page, click Start Upgrade).
PERC Alerts	
The storage appliance failed to gather the system diagnostics.	<ul style="list-style-type: none"> • Resolve all issues in the DR Series system diagnostics log bundle. • Re-attempt to collect the diagnostics log bundle. • Contact Technical Support for assistance.
Storage Appliance Critical Error: BIOS System ID is incorrect for correct operation of this storage appliance.	<ul style="list-style-type: none"> • The DR Series system appliance requires service. • Contact Technical Support for assistance.
Seeding Alerts	
Seeding device became full.	Add a new seeding device to continue.
Seeding cannot contact the target device.	Check to make sure that the target device is available and write-enabled. Then, remove and re-add the target device.
Seeding process complete.	Informational message. No user intervention is required.
System has reached space full condition, seeding will be stopped.	
Seeding failed to create Zero log entries.	Switch to maintenance mode to correct the issue.
Found corrupted stream on seeding device. This error will be rectified during replication re-sync done on this seed data.	Informational message. No user intervention is required.
Seeding device metadata info file missing, unable to import.	
Seeding device mount not accessible.	
Seeding export paused as the device contains	Cleanup the device and re-add to continue seeding.

Alert Message	Description/Meaning or Action
data from another seeding job.	
Seeding encountered error.	
Unable to decrypt the Seeding data.	Check that the "password" and "encryption type" matches the Seeding export job.
System diagnostics partition is running low on space.	Copy out the old diagnostics bundles and delete for future auto diagnostics collection.
Appliance available storage level is below the set threshold.	Schedule filesystem cleaner or expire older backups.
Primary Keystore corruption detected.	Run filesystem scan with data verification check.

Table 3: DR Series System Event Messages

System Event Message	Description/Meaning or Action
System Event = Type 1	
System requires a Restore Manager (RM) recovery.	
System failed basic initialization.	
HTTP Service failed. Web services will be unavailable.	
HTTP Service started.	Informational message. No user intervention is required.
HTTP Service is available now.	Informational message. No user intervention is required.
Diagnostics collection service failed.	
Diagnostic collection service started.	Informational message. No user intervention is required.
Diagnostics collection service re-started.	Informational message. No user intervention is required.
Configuration Service started.	Informational message. No user intervention is required.
Configuration Service is not healthy.	
Configuration Service is healthy.	Informational message. No user intervention is required.
Configuration Service failed to start.	
Unsupported RAID Configuration detected.	
No Fault Tolerant RAID configuration found.	
Data volume not present.	Check all drives are inserted and powered up.
Unable to detect filesystem type on data volume.	
Non certified disk drive detected. Disk needs to be pulled out for the system to become operational.	Disk needs to be pulled out for the system to become operational.
NVRAM devices not found.	Check card is seated properly.
Invalid/Unsupported Network Configuration	Use CLI "network --restart" to re-configure network cards.

System Event Message	Description/Meaning or Action
detected.	
Some of the network cards are not part of the bond configuration.	
No IP address has been assigned to the system.	
No valid hostname has been assigned to the system.	Use "system --setname" to set hostname.
No valid system name found in configuration Database.	Use "maintenance --configuration --restore" to recover from backup configuration.
No valid system configuration file(s) found in configuration database.	Use "maintenance --configuration --restore" to restore from backup configuration.
Data volume filesystem is not yet initialized.	
Backup configuration file is missing.	Contact Technical Support.
Working configuration file is missing.	Use "maintenance --configuration --restore" to restore configuration from backup.
Working configuration file is corrupted.	Use "maintenance --configuration --restore" to restore configuration from backup.
NVRAM signature is missing.	If NVRAM device was replaced use "maintenance --hardware --reinit_nvram" to initialize NVRAM.
Windows Active Directory client module failed to start. Active Directory support will not be available.	
Windows Active Directory client module started.	Informational message. No user intervention is required.
Windows Server module started.	Informational message. No user intervention is required.
Windows Server module re-started.	Informational message. No user intervention is required.
Windows Server module is down. Windows client access will be disrupted.	
Windows Server module has been disabled because of multiple crashes.	
System initialization is required.	Informational message. No user intervention is required.
Filesystem server maintenance requested.	
Filesystem server re-started.	Informational message. No user intervention is required.
Filesystem server started.	Informational message. No user intervention is required.
Filesystem server re-started, in Read-Only mode.	Informational message. No user intervention is required.
Filesystem server started, in Read-Only mode.	Informational message. No user intervention is required.
Filesystem server is not healthy. Client access will be interrupted.	
Filesystem scan triggered.	

System Event Message	Description/Meaning or Action
Filesystem check re-started.	Informational message. No user intervention is required.
Filesystem check continued from previous boot.	Informational message. No user intervention is required.
Filesystem checker is not healthy, will be re-started.	
Filesystem checker terminated with unexpected error.	
Filesystem checker crashing multiple times, entering support mode.	Please contact Technical Support.
Diagnostics collection module failed to start.	Reboot the system to recover. If problem persists contact Technical Support.
Hardware Health Monitor module failed to start.	Reboot the system to recover. If problem persists contact Technical Support.
System is exiting Support Mode.	Informational message. No user intervention is required.
De-dupe engine dictionary is corrupted.	Use "maintenance --configuration --reinit_dictionary" to re-init.
Not enough memory to validate NVRAM contents.	System reboot is required.
Failed to complete basic system initialization.	
Unable to detect filesystem type on the Name Space Volume.	
Name Space Volume is not mounted.	
iSCSI server started.	Informational message. No user intervention is required.
iSCSI server re-started.	Informational message. No user intervention is required.
iSCSI server is not healthy.	
iSCSI server is crashing repeatedly.	Contact Technical Support.
NDMP tape server started.	Informational message. No user intervention is required.
NDMP tape server re-started.	Informational message. No user intervention is required.
NDMP tape server is not healthy.	
NDMP tape server has crashed repeatedly.	Contact Technical Support.
Virtual Tape Library daemons started successfully.	Informational message. No user intervention is required.
Virtual Tape Library daemons re-started successfully.	Informational message. No user intervention is required.
Virtual Tape Library daemons are not healthy.	
Virtual Tape Library daemons have crashed repeatedly. All Virtual Tape functionality will not be available.	
Failed to process deleted files and containers.	Contact Technical Support.

System Event Message	Description/Meaning or Action
Internal failure processing ingest log.	Contact Technical Support for assistance or intervention.
Hardware Health Monitor Database is corrupted.	Use "maintenance --hardware --restore_hw_db".
Unable to communicate with Hardware Health Monitor.	Informational message. No user intervention is required.
Unable to communicate with NVRAM device. Check hardware.	Verify that the NVRAM card is seated properly in the DR Series system appliance. Contact Technical Support for assistance or intervention.
Capacitor is disconnected from NVRAM. If problem persist after reboot, replace NVRAM card.	Contact Technical Support for assistance or intervention.
SSD is disconnected from NVRAM device. If problem persist after reboot, replace NVRAM card.	Contact Technical Support for assistance or intervention.
NVRAM capacitor is not charging. If problem persist after 5 minutes of power shutdown, replace NVRAM card.	Contact Technical Support for assistance or intervention.
NVRAM has failed to backup or restore data during the last boot.	Contact Technical Support for assistance or intervention.
NVRAM is not yet ready to accept write commands.	Wait for NVRAM to become ready.
NVRAM hardware has failed.	Contact Technical Support for assistance or intervention.
Filesystem server is crashing repeatedly. Entering Maintenance mode to run filesystem scan utility.	
System is not initialized.	Use "system --init" to initialize the system.
NVRAM does not match the data volume.	If this is a newly replaced NVRAM, use "maintenance --hardware --reinit_nvram" to initialize.
Software upgrade is in progress.	Informational message. No user intervention is required.
Upgrade did not complete.	Retry upgrade after rebooting the appliance.
Upgrade completed successfully. Reboot required.	Reboot the system.
Upgrade completed successfully. System coming online.	Informational message. No user intervention is required.
Unable to set Filesystem Scan Markers.	Reboot the system. If problem persist contact Technical Support.
Not enough space to run Filesystem Scan.	Please clean-up older diagnostics file(s) and reboot the system. On reboot, execute "maintenance --filesystem --start_scan" to start filesystem scan. If the file system scan fails with not enough space, please contact Technical Support.
Filesystem server is crashing repeatedly in Maintenance Mode.	Please contact Technical Support.

System Event Message	Description/Meaning or Action
One or more software package is incompatible, please upgrade the appliance to rectify the issue.	Please upgrade the system appliance to rectify the issue. Upgrade the DR Series system appliance (in the Software Upgrade page, click Start Upgrade).
NVRAM Controller detected a memory failure	
NVRAM Health check in progress, please wait for it to complete before using the system.	Informational message. No user intervention is required.
NVRAM Health check is required, system will perform a quick health check	Informational message. No user intervention is required.
Failed to start NVRAM Health check, please reboot the appliance to recover from this state	Reboot the system.
Appliance encountered O/S issues. Please reboot the appliance to recover from this condition.	Reboot the system.
High system memory usage detected, system performance will be sluggish.	
System memory usage has returned to an optimal level.	Informational message. No user intervention is required.
A high level of system process usage has been detected, if it persists, please collect system diagnostics.	Informational message. No user intervention is required.
System process usage has returned to an optimal level.	Informational message. No user intervention is required.
A high-temperature reading has been detected on the NVRAM PCI controller. System will operate only in a read-only mode. Please check system airflow.	Informational message. No user intervention is required.
A high-temperature reading has been detected on the NVRAM PCI controller. System will not become operational until the temperature reduces to an ambient value of 55 degrees Celsius (131 degrees Fahrenheit).	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
The next NVRAM capacitor health check is scheduled for <variable>.	Informational message. No user intervention is required.
Windows Active Directory client is unable to contact the Active Directory domain server.	Informational message. No user intervention is required.
Active Directory domain server connectivity is restored.	Informational message. No user intervention is required.
Storage enclosure <variable> is authorized.	Informational message. No user intervention is required.
Storage enclosure <variable> is de-commissioned.	Informational message. No user intervention is required.
The system IP address has changed from <variable> to <variable>.	Informational message. No user intervention is required.
Refresh NHM Inventory	

System Event Message	Description/Meaning or Action
One or more storage enclosures have gone offline, please power-off the appliance, fix the connectivity issues and power-on the appliance.	Please power off the appliance, check whether all required storage enclosure(s) are powered-up, fix connectivity issues, and power on the appliance.
Data Volume has become in-accessible.	Contact Technical Support.
Data Volume has become read-only.	Contact Technical Support.
Namespace Volume has become in-accessible, please call Technical support.	Contact Technical Support.
Namespace Volume has become read-only.	Contact Technical Support.
Core Volume has become in-accessible.	Contact Technical Support.
Invalid storage appliance memory configuration.	
Storage Enclosure with Service Tag <variable> added successfully.	Informational message. No user intervention is required.
One of the storage enclosure has become offline.	Power-off the appliance, fix the connectivity issues and power-on the appliance.
Data Volume has become read-only.	Contact Technical Support.
Upgrade did not complete. Retry upgrade.	Retry to upgrade. If the problem persists, contact Technical Support.
Filesystem scan completed, restarting filesystem for normal operation.	Informational message. No user intervention is required.
Storage enclosure license(s) are missing.	If Restore Manager (RM) recovery was performed recently, please re-apply the license(s) and reboot.
BIOS System ID is incorrect for correct operation of this storage appliance. The storage appliance requires service.	
System clock has drifted more than 24 hours, from the last filesystem start.	Please check your clock settings and reboot.
This DR4x00 Virtual Machine usage time limit has expired.	Please contact your DR4x00 Sales representative to get the Hardware Version.
This DR4x00 Virtual Machine is for evaluation purpose only. Evaluation period ends on <variable>.	Informational message. No user intervention is required.
This DR4x00 Virtual Machine requires an evaluation license.	Please contact your DR4x00 Sales representative.
This DR4x00 Virtual Machine is designed to work only with 4 CPU(s) and 8GB of memory.	Informational message. No user intervention is required.
This DR2000v requires a license to operate.	Please install an evaluation license or register the DR2000v with a DR4000/DR4100/DR6000 series hardware appliance.
This DR2000v is unable to contact the license server to validate the license usage.	Please rectify the connectivity issues and reboot the system.

System Event Message	Description/Meaning or Action
This DR2000v Virtual appliance usage time limit has expired.	Contact your DR Series Sales representative to get a permanent license.
This DR2000v Virtual appliance usage time limit will expire on <variable>.	Informational message. No user intervention is required.
System Asset Tag information has non-printable characters.	Please use the iDRAC interface console and fix the issue.
This DR2000v has been deleted at license server.	Register using the CLI command "virtual_machine --register" again.
This DR4300e requires a license to operate, please install storage usage license.	DR4300e comes with a 30-day trial license. Then you need to install a capacity license. Contact Technical Support for assistance.
Internal storage license(s) are missing, if Restore Manager (RM) recovery was performed recently, please re-apply the license(s).	Re-add the old capacity license(s) after recovering the system. Contact Technical Support for assistance.
DR4300e data storage expanded successfully.	Informational message. Storage on DR4300e has now been expanded to 9TB after the addition of a second 4.5TB license. No user intervention is required.
Filesystem scan requested. Switching to Maintenance Mode. Filesystem has read-only access.	Informational message. No user intervention is required.
NVRAM not detected.	Ensure card is seated properly.
NVRAM capacitor is disconnected.	Contact Technical Support.
NVRAM capacitor has degraded.	Contact Technical Support.
NVRAM SSD is disconnected.	Contact Technical Support.
NVRAM has failed to backup/restore data during last boot.	Contact Technical Support.
NVRAM hardware failure.	Contact Technical Support.
Data volume is not present. Check that all drives are inserted and powered up.	Contact Technical Support for assistance or intervention.
Filesystem server failed to start after multiple attempts.	Contact Technical Support for assistance or intervention.
Filesystem server crashed multiple times. System is now entering Maintenance mode.	Contact Technical Support for assistance or intervention.
Insufficient disk space. Filesystem switched to read-only mode.	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
Unable to detect filesystem type on the Data Volume.	Contact Technical Support for assistance or intervention.
Unable to detect filesystem type on the Namespace Volume.	Contact Technical Support for assistance or intervention.
Filesystem scan discovered inconsistencies.	Please check report and take the recommended action.

System Event Message	Description/Meaning or Action
	Contact Technical Support for assistance or intervention.
NVRAM does not match data volume.	If this is a newly replaced NVRAM device, use the CLI maintenance --hardware --reinit_nvram command. For more information, see the <i>DR Series System Command Line Reference Guide</i> .
Storage usage is approaching the DR Series system capacity.	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
Replication re-sync cannot proceed because the Namespace depth has reached its maximum.	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
Filesystem has reached the maximum allowable file(s) and directories limit. New file and directory creation will be denied until sufficient space exists.	Please clean up the filesystem. If issue persists, contact Technical Support for assistance or intervention.
Filesystem is reaching the maximum allowable file(s) and directories limit. New file and directory creation will be denied after the limit has been reached.	Please clean up the filesystem. If issue persists, contact Technical Support for assistance or intervention.
Replication has encountered an unexpected error.	Contact Technical Support for assistance or intervention.
DataCheck has detected a potential corruption.	Run data consistency checks at the first available opportunity. If this issue persists, contact Technical Support for assistance or intervention.
Datacheck detected potential namespace inconsistency.	Run filesystem scan as soon as possible. ("maintenance --filesystem --start_scan")
Datacheck detected inconsistency in Isu image.	Run filesystem scan as soon as possible. ("maintenance --filesystem --start_scan verify_rda_metadata")
Datacheck detected potential corrupt Isu info.	Run filesystem scan as soon as possible. ("maintenance --filesystem --start_scan verify_rda_metadata")
Temperature warning detected on NVRAM PCI controller.	Please check the data center air conditioning, rack ventilation, and internal cooling fans for any issues. Ensure that there is proper air flow through the system appliance, and clean the system cooling vents as needed. If issue persists, contact Technical Support for assistance or intervention.
Filesystem Name Space partition has reached its maximum allowable limit.	Please delete any old, unused file(s) or disable replication (s). If issue persists, contact Technical Support for assistance or intervention.
Filesystem Name Space partition is reaching its maximum allowable limit.	New replication resynch(s) will be stopped. If issue persists, contact Technical Support for assistance or intervention.
One or more software package is incompatible.	Please upgrade the appliance to rectify the issue.
Filesystem volume has become in-active.	Please contact Technical Support for assistance or intervention.
Filesystem server response time exceeded max	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
threshold.	
The memory capacity of the storage appliance is below the minimum required for correct operation. The storage appliance requires service.	
An OST container quota is exceeded.	Check the event for container details.
One of the storage enclosure has become offline.	Please power-down the appliance and rectify the issue.
One or more storage enclosure(s) are missing/offline.	Please check whether the storage enclosure(s) are powered-up and connected to the appliance.
Storage enclosure license(s) are missing.	If Restore Manager (RM) recovery was performed recently, please re-apply the license(s) and reboot.
System has a huge backlog of book keeping work. Filesystem cleaner will be enabled outside of schedule setting and performance impact will be observed.	Informational message. No user intervention is required.
System clock has drifted more than 24 hours, from the last filesystem start.	Please check your clock settings and reboot.
Replication is disconnected on one or more containers.	Please check event log or replication stats for details.
One or more replication target systems are running low in space.	Please check event log or replication stats for details.
Filesystem scan completed with no inconsistencies. Switching back to operational mode.	Informational message. No user intervention is required.
Replication detected potential inconsistency.	Run filesystem scan with data verification check as soon as possible. ("maintenance --filesystem --start_scan verify_data")
Seeding device became full.	Add new seeding device to continue.
Seeding cannot contact the target device.	Check to make sure that the target device is available and write-enabled. Then remove and re-add the target device.
Seeding process complete.	Informational message. No user intervention is required.
System has reached space full condition, seeding will be stopped.	Informational message. No user intervention is required.
Seeding failed to create Zero log entries.	Switch to maintenance mode to correct the issue.
Found corrupted stream on seeding device. This error will be rectified during replication re-sync done on this seed data.	Informational message. No user intervention is required.
Seeding device metadata info file missing, unable to import.	
Seeding device mount not accessible.	
Seeding export paused as the device contains	Cleanup the device and re-add to continue seeding.

System Event Message	Description/Meaning or Action
data from another seeding job.	
Seeding encountered error.	
Unable to decrypt the Seeding data.	Please check that the "password" and "encryption type" matches the Seeding export job.
System diagnostics partition is running low on space.	Please copy out the old diagnostics bundles and delete for future auto diagnostics collection.
Appliance available storage level is below the set threshold.	Please schedule filesystem cleaner or expire older backups.
Primary Keystore corruption detected.	Run filesystem scan with data verification check.
System Event = Type 2	
Data check configuration successful.	Informational message. No user intervention is required.
Successfully <variable> system marker.	Informational message. No user intervention is required.
<variable> OPDUP encryption updated to <variable>	Informational message. No user intervention is required.
System storage usage alert has been set at <level>.	Informational message. No user intervention is required.
Successfully <variable> container <variable> with the following marker(s) "<markers>".	Informational message. No user intervention is required.
Container <name> created successfully.	Informational message. No user intervention is required.
Container <name> marked for deletion.	For more information, see Deleting Containers . Use the DR Series system CLI maintenance --filesystem --reclaim_space command to recover this storage space.
Container <name> has been deleted.	Informational message. No user intervention is required.
Successfully renamed container <name> as <name>.	Informational message. No user intervention is required.
Container <name> is configured to access over <variable> by the following clients: <clients> (* means access for everyone).	Informational message. No user intervention is required.
Container <name> is updated to access over <variable> by the following clients: <clients> (* means access for everyone).	Informational message. No user intervention is required.
Disabled access for Container <name> over <variable> for the following clients: <clients> (* means disabled access for everyone).	Informational message. No user intervention is required.
Successfully added connection entry for container <name>; type <variable> clients <variable>.	Informational message. No user intervention is required.
Successfully updated connection entry for container <name>; type <variable> clients <variable>.	Informational message. No user intervention is required.
Successfully deleted connection entry for	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
container <name>: type <variable> clients <variable>.	
Replication entry created successfully for container <name>: role <variable> peer <variable> peer container <variable>.	Informational message. No user intervention is required.
Replication configuration updated successfully for container <name>: role <variable> peer <variable>.	Informational message. No user intervention is required.
Replication marked for deletion for Container <name>: peer <variable> peer container <name>.	Informational message. No user intervention is required.
Replication deleted for container <name>.	Informational message. No user intervention is required.
Successfully initiated replication re-sync on Container <name>.	Informational message. No user intervention is required.
Replication <variable> defaults successfully updated: role <variable> peer <variable>.	Informational message. No user intervention is required.
Successfully updated replication bandwidth limit for <variable> to <variable>.	Informational message. No user intervention is required.
Successfully removed replication bandwidth limit for <variable>.	Informational message. No user intervention is required.
Successfully set <variable> replication bandwidth limit.	Informational message. No user intervention is required.
Replication enabled for container <name> with role <role>.	Informational message. No user intervention is required.
Replication disabled for container <name> with role <role>.	Informational message. No user intervention is required.
Snapshot <variable> → <variable> created successfully.	Informational message. No user intervention is required.
Snapshot <variable> → <variable> successfully updated.	Informational message. No user intervention is required.
Snapshot <variable> → <variable> successfully deleted.	Informational message. No user intervention is required.
Client <client> authorized to access NDMP Tape Server.	Informational message. No user intervention is required.
Successfully updated NDMP to use port <number>.	Informational message. No user intervention is required.
De-authorized NDMP client - <client>.	Informational message. No user intervention is required.
NDMP password successfully updated.	Informational message. No user intervention is required.
OST password updated successfully.	Informational message. No user intervention is required.
OST state updated successfully.	Informational message. No user intervention is required.
OST client <variable> with mode <variable>	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
added successfully	
OST client <variable> deleted successfully.	Informational message. No user intervention is required.
OST client <variable> with mode <variable> updated successfully.	Informational message. No user intervention is required.
OST client <variable> deleted successfully.	Informational message. No user intervention is required.
OST client <variable> with mode <variable> updated successfully.	Informational message. No user intervention is required.
Successfully updated <variable> schedule.	Informational message. No user intervention is required.
System compression level set to <variable>.	Informational message. No user intervention is required.
System configuration backup failed.	
Rapid Data Access (RDA) password updated successfully.	Informational message. No user intervention is required.
Rapid Data Access (RDA) state updated successfully.	Informational message. No user intervention is required.
Rapid Data Access (RDA) client <variable> with mode <variable> added successfully.	Informational message. No user intervention is required.
Rapid Data Access (RDA) client <variable> deleted successfully.	Informational message. No user intervention is required.
Rapid Data Access (RDA) client <variable> with mode <variable> updated successfully.	Informational message. No user intervention is required.
DR2000v with UUID <variable> IP Address <variable> Hostname <variable> registered successfully.	Informational message. No user intervention is required.
DR2000v with UUID <variable> IP Address <variable> Hostname <variable> unregistered successfully.	Informational message. No user intervention is required.
System Event = Type 3	
System is entering Maintenance mode.	Informational message. No user intervention is required. Contact Technical Support for assistance or intervention.
System entering Support Mode.	Contact Technical Support for assistance or intervention.
Internal failure—OFS client initialization failure.	Contact Technical Support for assistance or intervention.
Internal failure—mtab initialization failure for container <variable>.	Contact Technical Support for assistance or intervention.
Internal failure—cannot initialize node mtab.	Contact Technical Support for assistance or intervention.
Internal failure retrieving configuration for container ID <variable>.	Contact Technical Support for assistance or intervention.
Internal failure deleting container ID <variable>.	Contact Technical Support for assistance or intervention.
Internal failure stopping container ID <variable>.	Contact Technical Support for assistance or intervention.

System Event Message	Description/Meaning or Action
Internal failure adding connection <variable> for container ID <variable>.	Contact Technical Support for assistance or intervention.
Internal failure deleting connection <variable> for container ID <variable>.	Contact Technical Support for assistance or intervention.
Name space volume nearing low space condition. To prevent faster exhaustion of space, snapshot required for replication seeding for container <variable> will be paused until Name space volume recovers from low space conditions.	
Replication started as per schedule, will be active until <variable>.	Informational message. No user intervention is required.
Replication stopped as per schedule, will restart at <variable>.	Informational message. No user intervention is required.
Container replay failed for container <variable>.	Informational message. No user intervention is required. Contact Technical Support for assistance or intervention.
Internal failure—Name Space subsystem initialization failed.	Informational message. No user intervention is required. Contact Technical Support for assistance or intervention.
Inconsistencies were found in the Name Space.	Please schedule a filesystem consistency check using the DR Series system CLI maintenance --filesystem --start_scan command.
System entering Maintenance mode—Name Space log replay failed.	Contact Technical Support for assistance or intervention.
System entering Maintenance Mode—Name Space transaction failure.	Contact Technical Support for assistance or intervention.
Internal failure—failed to commit Name Space transaction.	Contact Technical Support for assistance or intervention.
Filesystem has reached the maximum supported number of Name Space entries.	Please clean up the filesystem to allow new file and directory create operations. If this condition persists, contact Technical Support for assistance or intervention.
Filesystem has recovered from a lack of available Name Space entries.	Filesystem create operations will now be allowed. Contact Technical Support for assistance or intervention.
Internal attributes of some files were found to be corrupt. The DR Series system will not allow the setting or removing of Attributes or ACLs on files that have corrupt attributes.	To find all files with corrupt attributes and to clear the state, please perform a maintenance scan using the DR Series system CLI maintenance --filesystem --start_scan command. Contact Technical Support for assistance or intervention.
System entering maintenance mode - Name Space Log Rotation failed	
File/Directory Statistics table out of sync. Switching to maintenance mode.	
Root inode of container, id <variable>, was found to be inconsistent. Fixed the attribute, ACL on root inode needs to be manually verified and fixed.	

System Event Message	Description/Meaning or Action
Replication re-sync started for container <variable>.	Informational message. No user intervention is required.
Replication internal re-sync started for container <variable>.	Informational message. No user intervention is required.
Replication re-sync completed for container <variable>.	Informational message. No user intervention is required.
Replication internal re-sync completed for container <variable>.	Informational message. No user intervention is required.
Internal failure creating replication snapshot for container <variable>.	If condition persists, reduce number of inodes, or contact Technical Support for assistance or intervention.
Internal failure deleting replication snapshot for container <variable>.	If condition persists, reduce number of inodes, or contact Technical Support for assistance or intervention.
Replication client connected for container <variable>.	Informational message. No user intervention is required.
Replication client disconnected for container <variable>.	Verify that the ports for replication (9904, 9911, 9915, and 9916) and OST (10011 and 11000) operations have been enabled. If issue persists, contact Technical Support for assistance or intervention.
Replication server connected for container <variable>.	Verify that the ports for replication (9904, 9911, 9915, and 9916) and OST (10011 and 11000) operations have been enabled. If issue persists, contact Technical Support for assistance or intervention.
Replication server disconnected for container <variable>.	Informational message. No user intervention is required.
Replication Name Space oplog full for container <variable>.	Verify that the ports for replication (9904, 9911, 9915, and 9916) and OST (10011 and 11000) operations have been enabled. If issue persists, contact Technical Support for assistance or intervention.
Replication switching to re-sync due to corrupt replication Name Space oplog for container <variable>.	
Replication data operations log (oplog) full for container <variable>.	The DR Series system should self-correct itself. If condition persists, reduce number of inodes, or contact Technical Support for assistance or intervention.
Replication switching to re-sync due to corrupt replication data oplog for container <variable>.	
Replication transmit log (txlog) full for container <variable>.	The DR Series system should self-correct itself. If condition persists, reduce number of inodes, or contact Technical Support for assistance or intervention.
System entering Maintenance mode due to corrupt replication txlog for container <variable>.	Collect a diagnostics log file, and open a Support record with Technical Support for assistance.
System entering Maintenance mode due to replication txlog commit error <variable> for container <variable>.	Collect a diagnostics log file, and open a Support record with Technical Support for assistance.

System Event Message	Description/Meaning or Action
Unable to make progress on filesystem replication for container <variable>.	Collect a diagnostics log file, and open a Support record with Technical Support for assistance.
Replication syncmgr exited for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication syncmgr event for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Name Space replicator exited for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication data replicator exited for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication protocol version mismatch for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication protocol version mismatch detected for container <variable>. Replication will continue with downgraded source protocol version.	
Replication delete cleanup failed for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication target system <variable> is running low on space. Replication cannot proceed further on container <variable>.	Informational message. Contact Technical Support for assistance or intervention.
Replication misconfiguration detected for container <variable>. Replication relationship might have been deleted forcibly on target system <variable>.	Informational message. Contact Technical Support for assistance or intervention.
Replication failed for container <variable> error <variable>.	Collect a diagnostics log file bundle, and open a Support record with Technical Support for assistance.
Replication server failed to commit blockmap for container <variable>. System is entering Maintenance mode.	The DR Series system should self-correct itself. If condition persists, reduce number of inodes, or contact Technical Support for assistance or intervention.
Container <variable> replication is paused, cleaner on replica is reclaiming space.	Run the Cleaner on the replica container. If condition persists, contact Technical Support for intervention or assistance.
Found mismatch in system software version with peer <variable>. Replication on source container <variable> would be stalled.	
Replication stalled on source container <variable> due to a mismatch in system software version or network issue with peer <variable>.	
Found mismatch in system software version with peer <variable>. Backup or replication on some or all target containers would be stalled.	
Received a garbled message from peer <variable>. Connection would be dropped.	

System Event Message	Description/Meaning or Action
Container <variable> replication encountered encryption setup error.	
NFS client successfully mounted <variable>.	Informational message. No user intervention is required.
Maximum NFS connection limit <variable> reached, active NFS connections <variable>.	You have reached the threshold limit. Reduce the number of connections.
NFS client <variable> successfully unmounted <variable>.	Informational message. No user intervention is required.
NFS client <variable> successfully unmounted all containers.	Informational message. No user intervention is required.
CIFS client successfully connected to container <variable>.	Informational message. No user intervention is required.
CIFS client <variable> session successfully disconnected from container <variable> .	Informational message. No user intervention is required.
Maximum <variable> connection limit <variable> reached.	You have reached the threshold limit for the specified protocol. Reduce the number of connections.
CIFS server failed to start <variable>.	Reboot the DR Series system. If issue persists, contact Technical Support for assistance or intervention.
CIFS client connected <variable> times to container <variable>.	Reboot the DR Series system. If issue persists, contact Technical Support for assistance or intervention.
CIFS server started successfully.	Informational message. No user intervention is required.
NFS server started successfully.	Informational message. No user intervention is required.
Storage usage approaching system capacity.	Informational message. No user intervention is required.
Online data verification (DataCheck) started.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) suspended.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) stopped.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) resumed.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) detected <variable> corruption.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) detected <variable> corruptions.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Online data verification (DataCheck) failed to start.	Informational message. If issue persists, contact Technical Support for assistance or intervention.
Seeding device became full.	Add new seeding device to continue.
Seeding cannot contact the target device.	Check to make sure that the target device is available and write-enabled. Then remove and re-add the target device.
Seeding process complete.	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
System has reached space full condition, seeding will be stopped.	
Seeding failed to create Zero log entries.	Switch to maintenance mode to correct the issue.
Found corrupted stream on seeding device. This error will be rectified during replication re-sync done on this seed data.	Informational message. No user intervention is required.
Seeding device metadata info file missing, unable to import.	
Seeding device mount not accessible.	
Seeding export paused as the device contains data from another seeding job.	Clean up the device and re-add to continue seeding.
Seeding encountered error.	
Unable to decrypt the Seeding data.	Please check that the "password" and "encryption type" matches the Seeding export job.
Seeding device deleted.	Informational message. No user intervention is required.
Seeding device added.	Informational message. No user intervention is required.
Seeding started.	Informational message. No user intervention is required.
Seeding stopped.	Informational message. No user intervention is required.
Container <variable> added to seeding.	Informational message. No user intervention is required.
Container <variable> is removed while seeding is in progress.	Informational message. No user intervention is required.
Container <variable> removed from seeding.	Informational message. No user intervention is required.
Seeding job created.	Informational message. No user intervention is required.
Seeding job deleted.	Informational message. No user intervention is required.
Seed space reclamation triggered.	Informational message. No user intervention is required.
Unable to use old seed dict. Creating a new dict.	Informational message. No user intervention is required.
Unable to read bmap scid. Retry seeding after running filesystem scan.	Retry seeding after running filesystem scan.
Unable to read DS scid. Retry seeding after running filesystem scan.	Retry seeding after running filesystem scan.
Seeding device mount not accessible. Check the CIFS mount and re-add the device to continue.	Check the CIFS mount and re-add the device to continue.
System Event = Type 4	
Internal Error. Unable to load deduplication dictionary <variable>.	Use the DR Series system CLI maintenance --configuration --reinit_dictionary command. If this issue persists, contact Technical Support for assistance or intervention.
Internal Error. Unable to locate deduplication	Use the DR Series system CLI maintenance --configuration

System Event Message	Description/Meaning or Action
dictionary <variable>.	--reinit_dictionary command. If issue persists, contact Technical Support for assistance or intervention.
Filesystem cleaner run <variable> started.	Informational message. No user intervention is required.
Filesystem cleaner run <variable> completed in <variable> milliseconds (ms).	Informational message. No user intervention is required.
Filesystem cleaner process encountered input/output (I/O) errors.	Maintenance-based issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands. Contact Technical Support as needed.
Failure to sync NVRAM <variable>.	NVRAM hardware issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands.
Failure reading from NVRAM <variable>.	NVRAM hardware issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands.
Failure writing to NVRAM <variable>.	NVRAM hardware issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands.
Failure to write sync NVRAM <variable>.	NVRAM hardware issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands.
Internal Error. Datastore <variable> length mismatch <variable>.	Maintenance-based issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands. Contact Technical Support as needed.
Data volume capacity threshold reached.	Informational message. No user intervention is required.
Out of space. Rollback of updates on object <variable> failed. Restarting file server.	Maintenance-based issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands. Contact Technical Support as needed.
Failure reading from data volume.	Maintenance-based issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands. Contact Technical Support as needed.
Failure writing to data volume.	Maintenance-based issue with the DR Series system. Check status using Maintenance mode or DR Series system CLI commands. Contact Technical Support as needed.
Checksum verification on metadata failed.	Contact Technical Support for assistance or repair the filesystem. For repairs, see About The DR Series Maintenance Mode .
Internal Error. Optimization engine log replay failed.	Contact Technical Support for assistance or repair the filesystem. For repairs, see About The DR Series Maintenance Mode .
Decompression of datastore failed <variable>.	Contact Technical Support for assistance or intervention.

System Event Message	Description/Meaning or Action
Internal Error. Failed to clean active datastore <variable>.	Contact Technical Support for assistance or intervention.
Internal Error. Negative reference on datastore <variable>. Record type: <variable>. Count: <variable>.	Contact Technical Support for assistance or repair the filesystem. For repairs, see the topic, "About the DR Series Maintenance Mode."
Internal Error. Data store <variable> contains negative stream reference count. Record type: <variable>. Count: <variable>.	Informational message. No user intervention is required.
Internal Error. Data store <variable> total reference count reached threshold. Record type: <variable>. Count: <variable>.	Informational message. No user intervention is required.
Data store[%s] stream(s) marked bad.	A chunk of data on the system failed consistency check. An attempt will be made to correct this chunk of data.
Internal Error. Entering Maintenance mode due to failure in processing logs.	Contact Technical Support for assistance or intervention.
Internal Error. Failed to acquire optimizer pipeline. Error: <variable>.	Contact Technical Support for intervention or assistance.
Internal Error. Failed to create optimizer event. Type: <variable>, Error: <variable>.	Contact Technical Support for intervention or assistance.
Internal Error. Task execution in fiber <variable> timed out after <variable> milliseconds (ms). Restarting file server.	Filesystem restarted. Collect diagnostics log file bundle, and upload diagnostics log file bundle to Technical Support.
Internal Error. Memory allocation failure.	Collect diagnostics log file bundle.
Background compression started.	Informational message. No user intervention is required.
Background compression completed.	Informational message. No user intervention is required.
Optimization initialized on container <variable>.	Informational message. No user intervention is required.
Optimization terminated on container <variable>.	Informational message. No user intervention is required.
Cleaner aborted at <variable>.	The DR Series system should enter Maintenance mode, and Cleaner process will restart.
Internal Error. Moving data from NVRAM to disk failed. System is entering its Maintenance mode.	Informational message. No user intervention is required.
System entering Maintenance Mode due to corrupt encryption keystore. Triggering key import.	Run filesystem scan with verify data enabled
Key rotation successful in internal mode	Informational message. No user intervention is required.
Key limit reached, reusing the last key	Informational message. No user intervention is required.
Filesystem encryption setting changed	Informational message. No user intervention is required.
Filesystem Cleaner process started as per schedule (will be active until <variable>).	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
Filesystem Cleaner process stopped as per schedule (will restart at <variable>).	Informational message. No user intervention is required.
Filesystem cleaner is paused, to speed up disk maintenance (e.g. Rebuild / Background Init) activities.	Informational message. No user intervention is required.
System entering Support Mode due to keystore repair failure, both primary and backup keystore are corrupt	
System entering Support Mode due to keystore empty failure, both primary and backup keystore are empty or removed	
System Event = Type 5	
System shutdown initiated by administrator.	Informational message. No user intervention is required.
System reboot initiated by administrator.	Informational message. No user intervention is required.
Start system upgrade to version <variable>.	Informational message. No user intervention is required.
System name changed to <variable>.	Informational message. No user intervention is required.
System date changed to <variable>.	Informational message. No user intervention is required.
System time zone changed to <variable>.	Informational message. No user intervention is required.
Password changed for user: administrator.	Informational message. No user intervention is required.
NTP server <variable> added.	Informational message. No user intervention is required.
NTP server <variable> deleted.	Informational message. No user intervention is required.
NTP service enabled.	Informational message. No user intervention is required.
NTP service disabled.	Informational message. No user intervention is required.
User data destroyed using CLI command.	Informational message. No user intervention is required.
User <variable> enabled.	Informational message. No user intervention is required.
User <variable> disabled.	Informational message. No user intervention is required.
Networking interfaces restarted.	Informational message. No user intervention is required.
DHCP enabled: IP address assigned by DHCP.	Informational message. No user intervention is required.
Static IP address <variable> assigned.	Informational message. No user intervention is required.
Network interface bonding mode set to <variable>.	Informational message. No user intervention is required.
Network MTU size set to <variable>.	Informational message. No user intervention is required.
System name set to <variable>.	Informational message. No user intervention is required.
Email relay host set to <variable> for email alerts.	Informational message. No user intervention is required.
Recipients for email alerts set to <variable>.	Informational message. No user intervention is required.
Recipient <variable> added to receive email	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
alerts.	
Recipient <variable> is no longer receiving email alerts.	Check whether email recipient still exists, or if mailbox is full.
Administrator information set to <variable> for email alerts.	Informational message. No user intervention is required.
Test email sent.	Informational message. No user intervention is required.
Joined the Windows Active Directory domain <variable>.	Informational message. No user intervention is required.
Left the Windows Active Directory domain <variable>.	Informational message. No user intervention is required.
System diagnostics package <variable> deleted.	Informational message. No user intervention is required.
All diagnostic packages deleted.	Informational message. No user intervention is required.
System diagnostic package <variable> is copied off the system.	Informational message. No user intervention is required.
System statistics reset by administrator.	Informational message. No user intervention is required.
System diagnostic package <variable> is collected.	Informational message. No user intervention is required.
System diagnostics space usage exceeded threshold. Auto cleaning oldest package: <variable>.	Informational message. No user intervention is required.
Internal Error. CIFS server cannot access file service.	Contact Technical Support for intervention or assistance. Collect diagnostics log file bundle, and upload to Technical Support.
Host <variable> added to SNMP alert recipient list.	Informational message. No user intervention is required.
Host <variable> deleted from SNMP alert recipient list.	Informational message. No user intervention is required.
Host <variable> enabled for SNMP alerts.	Informational message. No user intervention is required.
Host <variable> disabled for SNMP alerts.	Informational message. No user intervention is required.
User <variable> logged into the system.	Informational message. No user intervention is required.
CIFS user <variable> added.	Informational message. No user intervention is required.
CIFS user <variable> deleted.	Informational message. No user intervention is required.
Password changed for CIFS user <variable>.	Informational message. No user intervention is required.
System upgrade completed <variable>.	Informational message. No user intervention is required.
Cleared foreign configuration on disk <variable>.	Informational message. No user intervention is required.
User <variable> logged into the system.	Informational message. No user intervention is required.
Disk <variable> configured as hot spare.	Informational message. No user intervention is required.

System Event Message	Description/Meaning or Action
Telnet service enabled.	Informational message. No user intervention is required.
Telnet service disabled.	Informational message. No user intervention is required.
DNS settings updated with primary <variable>, secondary <variable>, and suffix <variable>.	Informational message. No user intervention is required.
System initialized successfully.	Informational message. No user intervention is required.
<variable> added with entitlement id <variable>.	Informational message. No user intervention is required.
Security privilege(s) changed for <variable>.	Informational message. No user intervention is required.
User <variable> logged into the administrative web interface.	Informational message. No user intervention is required.
Network interface(s) <variable> enabled.	Informational message. No user intervention is required.
Network interface(s) <variable> disabled.	Informational message. No user intervention is required.
SMBD backup traffic interface(s) <variable> do not have an IP.	
DR2000v registered successfully.	Informational message. No user intervention is required.
DR2000v unregistered successfully.	Informational message. No user intervention is required.
DR2000v data storage expanded by 1 TiB.	Informational message. No user intervention is required.
Miscellaneous Invalid/Last Event.	Informational message. No user intervention is required.
System Event = Type 6	
File system check started.	Informational message. No user intervention is required.
File system check completed successfully. No inconsistencies were found.	Informational message. No user intervention is required.
File system check found some inconsistencies.	The DR Series system Maintenance mode repair process should resolve this. If the problem persists, contact Technical Support for assistance or intervention.
File system repair started.	Informational message. No user intervention is required.
File system repair completed.	Informational message. No user intervention is required.
File system check stop requested.	Informational message. No user intervention is required.
One (or more) file(s) were deleted as part of the repair process.	Informational message. No user intervention is required. To verify, please use the DR Series system CLI maintenance -filesystem --repair_history verbose command.
One or more file(s) were deleted as part of the repair process for container <variable>. Replication will be stopped for this container.	Informational message. No user intervention is required.
One or more file(s) were deleted as part of the repair process for container <variable>. Re-sync has been initiated for this container.	Informational message. No user intervention is required.

System Event = Type 7

System Event Message	Description/Meaning or Action
RDA server started successfully.	Informational message. No user intervention is required.
RDA server failed to start.	Restart the RDA server. If issue persists, contact Technical Support for assistance or intervention.
RDA server stopped successfully.	Informational message. No user intervention is required.
<Variable> client authentication failed.	Retry the OST client authentication. If issue persists, contact Technical Support for assistance or intervention.
<Variable> Logical Storage Unit (LSU) quota exceeded <variable>.	Informational message. Reduce the number of LSUs. If issue persists, contact Technical Support for assistance or intervention.
<Variable> backup failed <variable>.	Retry the OST backup operation. If issue persists, contact Technical Support for assistance or intervention.
<Variable> Opdup failed <variable>.	The OST optimized duplication process failed. If issue persists, contact Technical Support for assistance or intervention.
<Variable> Restore failed <variable>.	The OST restore process failed. If issue persists, contact Technical Support for assistance or intervention.
RDA connections exceeded the maximum limit; count: <variable>, maximum limit: <variable>.	Informational message. Reduce the number of OST connections. If issue persists, contact Technical Support for assistance or intervention.
Connection from the <variable> client <variable> aborted.	Informational message. No user intervention is required.
RDA client protocol version is not supported.	Informational message. No user intervention is required. Check for the supported OST client versions in the <i>DR Series System Interoperability Guide</i> .
System is entering the Maintenance mode: <variable> LSU information file is corrupted.	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
System is entering the Maintenance mode: <variable> image information is corrupted.	Informational message. No user intervention is required. If issue persists, contact Technical Support for assistance or intervention.
<variable> client connection was reset.	Informational message. No user intervention is required.
System is entering the Maintenance mode: RDA meta directory is corrupted.	Informational message. No user intervention is required.
RDA server initialization failed.	Informational message. No user intervention is required.
RDA server initialization was successful.	Informational message. No user intervention is required.
System entering Maintenance Mode - RDA txlog full, LSU <variable>.	Informational message. No user intervention is required.
System entering Maintenance Mode - RDA txlog operation error <variable>, LSU <variable>.	Informational message. No user intervention is required.
System entering Maintenance Mode - RDA txlog roll-forward error <variable>, LSU <variable>.	Informational message. No user intervention is required.

About the DR Series system maintenance mode

In general, the DR Series system enters the **Maintenance** mode whenever the file system has encountered an issue that prevents it from operating normally.

i **NOTE:** You can use the Reason code information available in the **Maintenance** mode to contact Technical Support. All maintenance must be conducted using the DR Series systems Command Line Interface.

When in its **Maintenance** mode, the filesystem is in a read-only state, and the system runs the following maintenance-based operations:

i **NOTE:** Whenever the DR Series systems enters or exits from the **Maintenance** mode state, all communication via protocols is lost.

- Runs an internal filesystem check.
- Generates a filesystem status report (if the filesystem check finds no issues, the DR Series system switches back to **Operational** mode without user intervention).

If the filesystem check finds issues, you can choose to make repairs (using Confirm Repair Filesystem) or ignore the detected issue (using Skip Repair Filesystem), at which point the system switches back to **Operational** mode.

The **Maintenance** mode process displays a number of stages, indicated on the Maintenance Mode progress bar, which include:

- Preparing for Filesystem Check
- Scan in Progress
- Completed Generating Report

i **NOTE:** If the Filesystem Check detects any repairable files, it generates a Repair Report that identifies these reported files. The Maintenance Mode progress bar halts at the Completed Generating Repair stage, and remains in **Maintenance** mode until you click Confirm Repair Filesystem. The DR Series system does not advance to the Switching to Operation Mode stage until the filesystem repair is completed.

- Switching to Operational Mode
- Operational Mode (Normal State)

The **Maintenance Mode** page provides the following information:

- Maintenance Mode Progress bar:
 - Displays the five stages of **Maintenance** mode
 - Updates the progress bar as each stage completes

i **NOTE:** If an alert displays above the Maintenance Mode progress bar, this indicates that the filesystem check has completed, and it has generated a report on the repairable files (which are displayed in the Repair Report pane under the Maintenance Mode progress bar). To repair all of the reported files listed in the Repair Report, you must click Confirm Repair Filesystem.

- Repair Report:
 - Displays a list of repairable filesystem files that were detected in the Filesystem Check.
 - Identifies the repairable files by Container ID, File/Inode/Directory location, and a brief reason for failure.
 - Provides a search capability that allows you to click prev or next to display the previous or next page in the Repair Report, or lets you display a specific page number of the Repair Report by entering this number in the Goto page and click go.
- System Information pane:
 - **System Name**
 - **Software Version**
 - **Current Date/Time**
 - **iDRAC IP Address**
- Support Information
 - **Service Tag**
 - **Last Diagnostic Run**
 - **BIOS Version**

i NOTE: When in Maintenance mode, the DR Series system navigation panel displays the following options that are links to display the correspond page in the DR Series system GUI:

- **Alerts**
- **Events**
- **Health**
- **Usage**
- **Diagnostics**
- **Software Upgrade**

After the DR Series system enters **Maintenance** mode, there can only be two possible outcome states:

- **Operational** mode (Normal State): where the filesystem check was successful, and no system files need to be repaired (Filesystem Check: successful).
- **Maintenance** mode has halted: where the filesystem check detected one or more repairable files (Filesystem Check: unsuccessful).

Filesystem Check — Successful: when the **Maintenance** mode successfully completes all of its stages, the DR Series system displays its status as having entered **Operational** mode (Normal State). Only after the **Maintenance** mode has successfully completed its internal check can it return to an **Operational** mode.

To return to the **Operational** mode, click Go to Dashboard on the Maintenance Mode page options bar. Go to Dashboard is only active when all of the internal system checks have completed and the progress bar indicates that all stages have been completed.

i **NOTE:** You may encounter issues when using data management agents (DMAs) such as NetBackup with expired backup images when the DR Series system is in its **Maintenance** mode.

i **NOTE:** When in **Maintenance** mode, image expiration fails because the DR Series system is in a read-only state. If this occurs, the DMA assumes that the backup images have expired. However, the DR Series system administrator may be unaware that the backup data images still reside on the DR Series system.

Filesystem Check — Unsuccessful: when the **Maintenance** mode halts at the Completed Generating Report stage, this indicates that the filesystem check detected some repairable files, and listed them in the Repair Report pane on the **Maintenance Mode** page.

To return to the **Operational** mode, click Confirm Repair Filesystem on the Maintenance Mode page options bar to repair the files listed in the Repair Report. Confirm Repair Filesystem is the only active option you can select when the progress bar indicates that some filesystem files are in need of repair.

Supported Ports in a DR Series system

The following table lists the application and service ports found on a normally operating DR Series system. There may be other ports that are not listed here, that an administrator may need to open and enable to support specific operations across the network. Be aware that the ports listed in the following table may not reflect your specific network environment, or any planned deployment. While some of these DR Series system ports may not need to be accessible through the firewall, this information is made available when deploying the DR Series system in your own network because it indicates supported ports that may need to be exposed.

Table 4: Supported DR Series System Ports

Port Type	Number	Port Usage or Description
DR Series System Application Ports		
TCP	20	File Transfer Protocol (FTP)—for transferring files.
TCP	23	Telnet—remote terminal access protocol for unencrypted text communications.
TCP	80	Hypertext Transfer Protocol (HTTP)—unencrypted protocol communications.
TCP	443	HTTPS—combination of the HTTP with Secure Socket Layer (SSL)/Transport Layer Security (TLS).
TCP	1311	Hardware Health Monitor (Note: this is not used on the DR2000v)
TCP	9901	Watcher
TCP	9904	Configuration Server (needed for replication operations)
TCP	9911	Filesystem Server (needed for replication operations)
TCP	9915	MetaData Replication (needed for replication operations)
TCP	9916	Data Filesystem Server (needed for replication operations)
TCP	9918	Diagnostics Collector
TCP	9920	Data path used for OST or RDS replications
TCP	10011	Control channel (needed for OST or RDS operations)
TCP	11000	Data channel (needed for OST or RDS operations)
DR Series System Service Ports		
TCP	22	Secure Shell (SSH)—used for secure logins, file transfers like SCP (Secure Copy) and SFTP (Secure File Transfer Protocol)

Port Type	Number	Port Usage or Description
TCP	25	Simple Mail Transfer Protocol (SMTP)—used for routing and sending email
TCP	139	SMB daemon—used for SMB protocol-related processes
TCP	199	SNMP daemon—used by Simple Network Management Protocol (SNMP) requests
TCP	801	NFS status daemon

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece — you — to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product