

Quest® vRanger® 7.6.3

**Integration Guide for EMC® Data
Domain® Boost (DD Boost™)**



© 2017 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 website (<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, Foglight, NetVault, vRanger, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. vCenter, vSphere, vMotion, VMware, and ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Active Directory, Hyper-V, SharePoint, SQL Server, Windows, and Windows Server are registered trademarks of Microsoft Corporation in the United States and/or other countries. Data Domain and DD Boost are trademarks or registered trademarks of EMC Corporation in the United States and other countries. 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

-  **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

-  **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Getting started	5
vRanger integration quick start for DD Boost	5
Product overviews	6
What is vRanger?	6
DD Boost overview	6
Understanding DD Boost	8
Introducing DD Boost	8
Supported configurations	8
DD Boost features	8
Distributed segment processing	9
Advanced Load Balancing and Link Failover	9
Managed file replication	10
Installing and configuring DD Boost	12
Installing DD Boost	12
Firewalls and ports	12
Enabling DD Boost on a Data Domain system	13
Configuring the Data Domain system	13
Creating a DD Boost user name	13
Creating Storage Units	14
Configuring distributed segment processing	14
Configuring Advanced Load Balancing and Link Failover	15
Configuring replication	15
Installing vRanger	17
vRanger system requirements	17
Hardware requirements	17
Supported operating systems for installation	18
Installing vRanger	19
Configuring vRanger	19
The DD Boost server must be fully configured and running before adding it as a vRanger Repository. For procedures on how to add a DD Boost server as a repository, see Managing repository replication	19
Adding an EMC Data Domain Boost (DD Boost) repository	19
Managing repository replication	22
.....	24
Administering Data Domain and DD Boost	25
Data Domain system administration	25
Restricting export permissions to DD Boost	25
Modifying an interface group	26
Removing advanced load balancing and link failover configuration	27
Setting network timeouts	27
Using the DD Boost commands	27

About us	34
We are more than just a name	34
Our brand, our vision. Together.	34
Contacting Quest	34
Technical support resources	34

Getting started

- [vRanger integration quick start for DD Boost](#)
- [Product overviews](#)

vRanger integration quick start for DD Boost

NOTE: The information in this topic is intended only to provide an overview of the steps required to configure vRanger and DD Boost. Review all applicable documentation before implementing this solution.

Integrating EMC® Data Domain® Boost (DD Boost™) with vRanger is achieved by adding an EMC® Data Domain® appliance running DD Boost to vRanger as a repository. Backups written to that repository are deduplicated according to your configuration. The following steps describe the integration process at a high level, and provide links to the remaining topics for more exploration. Before implementing this solution in a production environment, read this documentation thoroughly.

- 1 Install DD Boost:** To install and enable DD Boost, perform the following operations:
 - Obtain the license required to enable DD Boost on the Data Domain system. You can purchase a DD Boost license key directly from EMC. The DD Boost license allows you to back up and restore data.
 - Enable and configure DD Boost on the Data Domain system. At a minimum, configuration includes specifying the DD Boost user name and password.
 - Install vRanger, which contains the DD Boost Libraries.
- 2 Configure DD Boost:** After DD Boost is licensed and enabled, perform the following steps, if applicable:
 - **Create DD Boost User:** vRanger uses the DD Boost user credentials to connect to the Data Domain system.
 - **Create Storage Units:** When used with vRanger, a DD Boost repository is configured at the Storage Unit level, meaning each DD Boost repository in vRanger equates to exactly one Storage Unit. You may create storage units through the vRanger user interface (UI), or directly on the Data Domain system, as described in [Creating Storage Units](#).
 - **Configure Distributed Segment Processing:** This step allows parts of the deduplication process to be performed by the DD Boost Library, which avoids sending duplicate data to the DD Boost repository. For more information, see [Configuring distributed segment processing](#).
 - **Configure Advanced Load Balancing and Link Failover:** This feature allows for combining multiple Ethernet links into a group. Only one of the interfaces on the Data Domain system is registered with vRanger. The DD Boost Library negotiates with the Data Domain system on the interface registered with vRanger to obtain an interface to send the data. For instructions, see [Configuring Advanced Load Balancing and Link Failover](#).
- 3 Install vRanger:** Using the vRanger installer, install vRanger on a machine meeting the [vRanger system requirements](#). For an overview of the installation process, see [Installing vRanger](#). For detailed information, see the [Quest vRanger Installation/Upgrade Guide](#).

- 4 **Create DD Boost Repository:** After vRanger is installed, add the DD Boost instance to vRanger as a repository. Any backup written to this repository is deduplicated according to your Data Domain configuration. For more information, see [Managing repository replication](#).

Product overviews

The following topics provide an overview of vRanger and EMC® Data Domain® Boost (DD Boost™), and important information about the licensing required to integrate the two products.

What is vRanger?

vRanger is the leading VMware® data protection solution that also backs up and recovers Windows® physical servers and files with blazing speed and minimal storage requirements. With vRanger, you get comprehensive protection for both virtual and physical environments that you can manage from one intuitive interface.

vRanger has historically supported EMC® Data Domain® systems as repositories, but required that the Data Domain system was configured as a Common Internet File System (CIFS) share. vRanger extends this functionality to support the more advanced functionality offered by EMC® Data Domain® Boost (DD Boost™), and allow direct configuration of DD Boost repositories.

vRanger licensing

For virtual machine (VM) backup, a license for vRanger controls the number of source CPUs that you can configure for backup. For licensing purposes, a multi-core processor is counted as a single CPU. For physical backup, each server protected consumes one physical backup license.

Every VMware® ESXi™ host for which vRanger is expected to provide protection must be properly licensed, both by VMware and in the vRanger Host Licensing tab.

i | **IMPORTANT:** A Quest vRanger license is required for integration with EMC® Data Domain® Boost (DD Boost™). vRanger Standard Edition (SE) and vReplicator are not supported by DD Boost. This documentation uses “vRanger” for simplicity, but only vRanger Backup & Replication is supported for use with DD Boost.

DD Boost overview

EMC® Data Domain® Boost (DD Boost™) enables backup servers to communicate with storage systems without the need for EMC® Data Domain® storage systems to emulate tape. There are two components to the software:

- The DD Boost Libraries for integrating with the DD Boost server running on the Data Domain system.
- The DD Boost server that runs on Data Domain systems. A Data Domain system can be a single Data Domain system or a gateway.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled backup server. By using the DD Boost Libraries, multiple backup servers can use the same storage unit on a Data Domain system as a storage server. Each backup server can run a different operating system, if it is supported by Data Domain.

DD Boost licensing

EMC® Data Domain® Boost (DD Boost™) software is an optional product that requires a license to operate on the EMC® Data Domain® system. You can purchase a DD Boost software license key for a Data Domain system directly from EMC.

Understanding DD Boost

- [Introducing DD Boost](#)
- [DD Boost features](#)

Introducing DD Boost

EMC® Data Domain® Boost (DD Boost™) enables backup servers to communicate with storage systems without the need for EMC® Data Domain® storage systems to emulate tape. There are two components to the software:

- The DD Boost Libraries for integrating with the DD Boost server running on the Data Domain system.
- The DD Boost server that runs on Data Domain systems. A Data Domain system can be a single Data Domain system or a gateway.

The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled backup server. By using the DD Boost Libraries, multiple vRanger instances can use the same storage unit on a Data Domain system as a storage server. Each backup server can run a different operating system, if it is supported by Data Domain.

Supported configurations

EMC® Data Domain® supports EMC® Data Domain® Boost (DD Boost™) on the following Data Domain (OS) versions:

- OS 5.7—requires vRanger 7.3 or later
- OS 6.0—requires vRanger 7.6.3 or later

The DD Boost Library version must be compatible with your Data Domain system and data protection application configurations.

DD Boost features

This topic provides a technical overview of EMC® Data Domain® Boost (DD Boost™) features. Configuration information for these features is described in [Configuring the Data Domain system](#).

- [Distributed segment processing](#)
- [Advanced Load Balancing and Link Failover](#)
- [Managed file replication](#)
- [Low-bandwidth optimization](#)

Distributed segment processing

Distributed segment processing allows parts of the deduplication process to be performed by the EMC® Data Domain® Boost (DD Boost™) Library, which avoids sending duplicate data to the EMC® Data Domain® system that is configured as a vRanger repository. The option to distribute the processing is enabled or disabled on the Data Domain system; for more information, see [Configuring distributed segment processing](#). The DD Boost Library negotiates with the Data Domain system for the current setting of the option and behaves appropriately. Therefore, the DD Boost Library offers two modes of operation for sending backup data to a Data Domain system: one with distributed segment processing enabled and the other with the distributed segment processing disabled.

Distributed segment processing provides the following benefits:

- Potentially higher throughput, because the DD Boost Library sends only unique data to a Data Domain system versus sending all the data over the LAN. Throughput improvements depend on the redundant nature of the data being backed up, the overall workload on the backup server, and the capability of the backup server. In general, greater throughput is attained with higher redundancy, greater backup server workload, and backup server capability.
- The network bandwidth requirements are reduced since only the unique data is sent over the LAN to the Data Domain systems.

For more information about local compression and its configuration, see the *EMC Data Domain Operating System Administration Guide*.

You manage distributed segment processing by using the **ddboost** option commands, as described in [Using the DD Boost commands](#).

i | **NOTE:** EMC recommends that you use distributed segment processing if your network connection is 1 Gb Ethernet.

Distributed segment processing disabled

In this mode, the EMC® Data Domain® Boost (DD Boost™) Library sends the data directly to the Data Domain system over the LAN. The Data Domain system then segments, deduplicates, and compresses the data before it is written to the disk. The DD Boost Library runs on the backup server.

Distributed segment processing enabled

In this mode, the deduplication process is distributed between the EMC® Data Domain® Boost (DD Boost™) Library and the Data Domain system.

i | **NOTE:** The DD Boost Library runs on the vRanger server. The DD Boost Library's handling of data is transparent to vRanger and the library does not store any data on the backup server.

Parts of the deduplication process are run by the DD Boost Library so that only unique data is sent to data to a Data Domain system over the LAN. The DD Boost Library segments, computes IDs for the segments, checks with the Data Domain system for duplicate segments, compresses unique segments that are not found on the Data Domain system, and sends the compressed data to the Data Domain system. The Data Domain system then writes the unique data to disk.

Advanced Load Balancing and Link Failover

vRanger supports a single host name and IP address per repository. For EMC® Data Domain® systems that require multiple 1 GbE links to obtain full system performance, you must set up multiple storage servers on the Data Domain systems—one per interface—and target the backup policies to different storage servers to spread the load on the interfaces.

The Advanced Load Balancing and Link Failover feature allows for combining multiple Ethernet links into a group. Only one of the interfaces on the Data Domain system is registered with vRanger. The EMC® Data Domain® Boost (DD Boost™) Library negotiates with the Data Domain system on the interface registered with vRanger to obtain an interface to send the data. The load balancing provides higher physical throughput to the Data Domain system compared to configuring the interfaces into a virtual interface using Ethernet level aggregation.

The Data Domain system load balances the connections coming in from multiple vRanger instances on all the interfaces in the group. The load balancing is transparent to vRanger. Because Advanced Load Balancing and Link Failover works at the DD Boost Library software layer, it is seamless to the underlying network connectivity and supports both physical and virtual interfaces.

The data transfer is load balanced based on the number of connections outstanding on the interfaces. Only connections for backup and restore jobs are load balanced.

i | NOTE: The replication connection between the Data Domain systems is not part of the load balancing. A single IP address is used for the target Data Domain system.

i | IMPORTANT: Quest recommends that you exclude one interface from the ifgroup and reserve it for the replication path between the source and target Data Domain systems.

The Advanced Load Balancing and Link Failover feature can be used with other network layer aggregation and failover technologies. The Advanced Load Balancing and Link Failover functionality also works with other network layer functionality on the Data Domain systems, including VLAN tagging and IP aliasing. This functionality allows additional flexibility in segregating traffic into multiple virtual networks, which run over the same physical links on the Data Domain system.

The benefits of Advanced Load Balancing and Link Failover are as follows:

- Eliminates the need to register multiple storage servers—one for each interface—with vRanger, which can potentially simplify installation management.
- If one of the interfaces in the group goes down while the Data Domain system is still operational, the subsequent incoming backup jobs are routed to the available interfaces.
- The backup and restore jobs are automatically load balanced on multiple interfaces in the group, which can potentially result in higher utilization of the links.
- All in-flight jobs to the failed interface transparently are failed over to healthy operational links. From the point of view of vRanger, the jobs continue uninterrupted.

Configuration restrictions

- The interfaces can be added only to the group using an IP address.
- Even though the Advanced Load Balancing and Link Failover feature works with mixed 1GbE interfaces and 10 GbE interfaces in a group, it is not a recommended setup. Quest recommends that you have interfaces with the same link speed in a group.
- Since only one interface group is supported on the EMC® Data Domain® system, for connecting more than one backup server, a switch is needed in the middle.

Managed file replication

File replication replicates a single file. The replication request is initiated by the application. By default, the backup servers set up replication jobs between two EMC® Data Domain® systems after authenticating them using the pre-configured EMC® Data Domain® Boost (DD Boost™) user name and password. If the replication feature is enabled, when an optimized job is set up by backup servers, the session between the source and destination Data Domain systems is encrypted using Secure Sockets Layer (SSL), ensuring all image data and metadata is sent encrypted over the WAN.

Enabling this option on Data Domain system is transparent to the data protection application. When the data protection application requests Data Domain system to perform a replication job, the source and destination systems negotiate automatically to perform encryption without knowledge of the requesting data protection application. Replication uses the ADH-AES256-SHA cipher suite. There is no ability to configure a different suite.

Replication is available to a Data Domain system with an installed Replicator license, and applies to all replication jobs on that system. Both the source and the destination Data Domain systems that are participating in replication jobs must have this option enabled for it to take effect.

Replication can be used with the encryption of data-at-rest feature available on Data Domain operating system with the optional Encryption license. When replication is used with the encryption of data-at-rest feature, the encrypted backup image data is encrypted again using SSL for sending over WAN.

Replication notes

- Both the source and the destination EMC® Data Domain® systems must be running Data Domain OS 5.2 or later to use this feature.
- Enabling this feature does not require restarting the file system on a Data Domain system.
- The low-bandwidth optimization and the encryption options can be used with each other.

Low-bandwidth optimization

The low-bandwidth Replicator option reduces the WAN bandwidth utilization. This option is useful if replication is being performed over a low-bandwidth network (WAN) link. This option provides additional compression during data transfer and is recommended only replication jobs that occur over WAN links that have fewer than 6 Mb/s of available bandwidth.

The low-bandwidth optimization option is available to EMC® Data Domain®s systems with an installed Replicator license. The option is enabled on a Data Domain system and applies to all replication jobs on that system.

Enabling this option on Data Domain system is transparent to the data protection application. When the data protection application requests a Data Domain system to perform a replication job, the source and destination systems automatically perform the additional compression without involving the requesting data protection application.

Each Data Domain system wanting to participate in low-bandwidth replication must have this option enabled.

Installing and configuring DD Boost

- [Installing DD Boost](#)
- [Configuring the Data Domain system](#)

Installing DD Boost

i | **NOTE:** Complete descriptions of commands used in this guide are provided in the *EMC Data Domain Operating System Command Reference Guide*.

The overall steps for installing DD Boost are as follows:

- 1 Obtain the license required to enable DD Boost on the Data Domain system; purchase a DD Boost license key from EMC.
 - The DD Boost license allows you to back up and restore data.
 - A separate replication license enables you to perform replication. You must obtain a replication license for both the source and destination Data Domain systems.

- 2 Enable and configure DD Boost on the Data Domain system.

At a minimum, configuration includes specifying the DD Boost user name and password.

- 3 Install vRanger, which contains the DD Boost Libraries.

i | **NOTE:** This documentation illustrates the DD Boost configuration on Data Domain using commands in Data Domain OS 5.2. If you are using a different version of Data Domain operating system with this version of the DD Boost Library, see the corresponding Data Domain operating system version documentation for CLI details.

Many of the configurations done by command line can also be made using the EMC Data Domain Enterprise Manager web console.

Firewalls and ports

The EMC® Data Domain® system as it is initially configured does not work through a firewall. If you need the Data Domain system to work in the presence of a firewall, contact your network support provider.

The following ports must be open in a firewall for EMC® Data Domain® Boost (DD Boost™) backups and replication to work:

- TCP 2049 (Network File System [NFS])
- TCP 2051 (Replication)
- TCP 111 (NFS portmapper)
- TCP xxx (select a random port for NFS mountd)

Enabling DD Boost on a Data Domain system

Every EMC® Data Domain® system that is enabled for EMC® Data Domain® Boost (DD Boost™) must have a unique name. You can use the Data Domain system's DNS name, which is always unique.

i | IMPORTANT: The user must be configured in vRanger to connect to the Data Domain system.

To enable DD Boost:

- 1 On the Data Domain system, log in as an administrative user.
- 2 Verify that the file system is enabled and running by entering:

```
# fileSYS status
The file system is enabled and running.
```
- 3 Add the DD Boost license using the license key that Data Domain provided:

```
# license add license_key
License "ABCE-BCDA-CDAB-DABC" added.
```
- 4 Set the DD Boost user by entering:

```
# ddboost set user-name ddboost-user
```
- 5 Enable DD Boost by entering:

```
# ddboost enable
DD Boost enabled
```

Configuring the Data Domain system

This topic describes tasks and commands for configuring the EMC® Data Domain® system.

- [Creating a DD Boost user name](#)
- [Creating Storage Units](#)
- [Configuring distributed segment processing](#)
- [Configuring Advanced Load Balancing and Link Failover](#)
- [Configuring replication](#)

Creating a DD Boost user name

An EMC® Data Domain® Boost (DD Boost™) user is an EMC® Data Domain® user. Before setting the DD Boost user access, the username and password must have already been set up on the Data Domain system. vRanger uses the DD Boost user account to connect to the DD Boost repository.

When DD Boost is enabled, a Data Domain OS administrative user can set or change to another DD Boost user. Only one DD Boost user can operate DD Boost at a time.

To set or modify the DD Boost user name:

- 1 In the DD Boost overview pane located under the Data Management tab, click **Add** or **Modify** in the DD Boost User area.
The **Set** or **Modify DD Boost User Name** dialog box appears.
- 2 Do one of the following:
 - To add a user, click the **User Name** button.

- a Enter the user name in the **DD Boost User Name** field.
The user must be configured in the backup application to connect to the Data Domain system.
 - b Enter the password twice in the text fields.
 - c Select the privilege level for this user: admin or user.
- To switch to a user that has already been added, click the **Existing User** button, and select the user name from the drop-down list.
- 3 Click OK.

Creating Storage Units

The EMC® Data Domain® system exposes pre-made disk volumes called storage units to an EMC® Data Domain® Boost (DD Boost™)-enabled backup server. Each Storage Unit is a top-level subdirectory of the **/data/col1** directory; there is no hierarchy among Storage Units. When used with vRanger, a DD Boost repository is configured at the Storage Unit level, meaning each DD Boost repository in vRanger equates to exactly one Storage Unit.

i | **NOTE:** You may also create a Storage Unit from the vRanger UI during the DD Boost repository configuration, or by using the EMC Data Domain Enterprise Manager web console.

- 1 To create a storage unit on the Data Domain system, enter:

```
# ddbboost storage-unit create storage-unit_name
```

i | **IMPORTANT:** Each storage unit name must be unique. For devices running Data Domain OS 5.2 and later, the supported characters are alphanumeric characters and ~!@#\$_+`-={}|,.

- 2 Repeat the previous step for each DD Boost-enabled Data Domain system.

Configuring distributed segment processing

The distributed segment processing option is configured on the EMC® Data Domain® system and applies to all the backup servers and the EMC® Data Domain® Boost (DD Boost™) Libraries installed on them. For more information, see [Distributed segment processing](#).

The option can be configured using the following command:

```
# ddbboost option set distributed-segment-processing {enabled | disabled}
```

i | **NOTE:** Enabling or disabling the distributed segment processing option does not require a restart of the Data Domain file system.

Distributed segment processing is supported with DD Boost Library version 2.2 or later communicating with a Data Domain system that is running Data Domain OS 5.2 or later.

Distributed segment processing is enabled by default on a system initially installed with Data Domain OS 5.2. If a system is upgraded from Data Domain OS 5.0.x to Data Domain OS 5.2 or later, distributed segment processing is left in its previous state.

Configuring Advanced Load Balancing and Link Failover

If an interface group is configured, when the EMC® Data Domain® system receives data from the backup server clients, the data transfer is load balanced and distributed as separate jobs on the private network, providing higher input and output throughput, especially for customers who use multiple 1 GbE connections.

Create an interface group on the Data Domain system by adding existing interfaces to the group and registering the Data Domain system with vRanger, as described in the following steps.

- 1 Verify that the interfaces have been created with the net command.
- 2 Add the interfaces into the group.

```
# ddbboost ifgroup add interface 192.168.1.1
# ddbboost ifgroup add interface 192.168.1.2
# ddbboost ifgroup add interface 192.168.1.3
# ddbboost ifgroup add interface 192.168.1.4
```

i | NOTE: Only one interface group can be created and this group cannot be named.

- 3 Select one interface on the Data Domain system to register with vRanger.

Quest recommends that you create a failover-aggregated interface and register that interface with vRanger.

It is not mandatory to have one of the interfaces in the ifgroup registered with vRanger. An interface that is not part of the ifgroup can also be used to register with vRanger.

i | NOTE: EMC recommends that you register the interface with a resolvable name using DNS or any other name resolution mechanism.

The interface registered with vRanger is used by vRanger and the DD Boost Library to communicate with the Data Domain system. If this interface is not available, backups to that Data Domain system are not possible.

- 4 Enable the feature on the Data Domain system:

```
# ddbboost ifgroup enable
```

- 5 Verify the configuration of the Data Domain system:

```
# ddbboost ifgroup show config
Interface
-----
192.168.1.1
192.168.1.2
192.168.1.3
192.168.1.4
-----
```

After the interface group is set up, you can add or delete interfaces from the group. You can manage Advanced Load Balancing and Link Failover—an interface group—by using the **ddbboost ifgroup** command or from the **Enterprise Manager Data Management > DD Boost** view. For more information, see the *EMC Data Domain Operating System Administration Guide*.

Configuring replication

With EMC® Data Domain® Boost (DD Boost™), vRanger can control replication between multiple EMC® Data Domain® systems and provide backup administrators with a single point of management for tracking all backups and duplicate copies. This configuration lets you efficiently create disaster recovery (DR) copies of your backups over the WAN, and track all the copies using the vRanger catalog for easy recovery.

Enabling replication

When using EMC® Data Domain® Boost (DD Boost™) with vRanger, replication is configured through the vRanger interface. For instructions on enabling replication, see [Configuring repository replication](#).

Enabling low-bandwidth optimization

To enable the low-bandwidth option, enter:

```
# ddboost file-replication option set low-bw-optim enabled
Low bandwidth optimization enabled for replication.
```

Enabling or disabling the low-bandwidth optimization option does not require a restart of the EMC® Data Domain® file system. After enabling low-bandwidth optimization, however, you need to run a full cleaning cycle on the Data Domain system for it to be effective.

Low-bandwidth optimization can be monitored and managed with the command `ddboost file replication` or from the Enterprise Manager Data Management > DD Boost view. For more information, see the *EMC Data Domain Operating System Administration Guide*.

No configuration changes are necessary on the backup server as this feature is transparent to the data protection applications.

Notes:

- Enabling this feature takes additional resources—CPU and memory—on the Data Domain system; Quest recommends that you use this option only when replication is being done over low-bandwidth networks with less than 6 Mbps aggregate bandwidth.
- For more information on this feature, see the *EMC Data Domain Operating System Administration Guide*.
- The low-bandwidth option for replication is supported only for standalone Data Domain systems.

Enabling encrypted file replication

To enable the encrypted file-replication option, enter:

```
# ddboost file-replication option set encryption enabled
```

The output indicates that the encryption you requested was enabled.

For more information, see the command `ddboost file-replication`.

No configuration changes are necessary on the backup server as this feature is transparent to the data protection application. Turning on this feature takes additional resources—CPU and memory—on an EMC® Data Domain® system.

Notes:

- Enabling or disabling the encrypted file replication option does not require a restart of the Data Domain file system.
- For more information on this feature, see the *EMC Data Domain Operating System Administration Guide*.

Installing vRanger

- [vRanger system requirements](#)
- [Installing vRanger](#)
- [Configuring vRanger](#)
- [Managing repository replication](#)

vRanger system requirements

Before installing vRanger and EMC® Data Domain® Boost (DD Boost™), ensure that you read and understand the requirements and limitations of both products. The following topics summarize the system requirements for vRanger.

i | **IMPORTANT:** The following information is a summary. Before implementing vRanger, review the *Quest vRanger Installation/Upgrade Guide*.

For vRanger and DD Boost to integrate, you must use the following versions:

- vRanger
- At least one EMC® Data Domain® device running:
 - Data Domain OS 5.7—requires vRanger 7.3 or later
 - Data Domain OS 6.0—requires vRanger 7.6.3 or later

Hardware requirements

The hardware requirements to run vRanger can vary widely based on several factors. Therefore, you should not do a large-scale implementation without first completing a scoping and sizing exercise.

vRanger: physical machine

The following describes the hardware recommendations for the vRanger physical machine.

Table 1. Hardware requirements for physical machine

CPU	Any combination equaling four cores of CPUs are recommended. Example one quad-core CPU; two dual-core CPUs.
RAM	4 GB RAM is required.
Storage	At least 4 GB free hard disk space on the vRanger machine.
HBA	For LAN-free, Quest recommends that you use two HBAs—one for read operations and one for writing.

vRanger: virtual machine (VM)

The following describes the hardware recommendations for using vRanger in a VM.

Table 2. Hardware requirements for virtual machine

CPU	Four vCPUs.
RAM	4 GB RAM is recommended.
Storage	At least 4 GB free hard disk space on the vRanger machine.

Supported operating systems for installation

The following operating systems are supported for installation of vRanger.

Table 3. Supported operating systems

Operating system	Service pack level	Bit level
Windows Server 2008 R2 ^{ab}	SP1 or later	x64
Windows Server 2012 ^b	All service packs	x64
Windows Server 2012 R2 ^{bc}	All service packs	x64
Windows Server 2016 ^b	All service packs	x64

- Windows 2008 R2 SP1 requires Windows Management Framework 3.0. Refer to Known Issue VR-177 in the vRanger Release Notes for more information.
- The Windows Storage Server edition is not supported as an installation platform for vRanger.
- Before installing vRanger on Windows Server 2012 R2, the updates listed in [Additional required software](#) must be installed.

Additional required software

In addition to a supported version of Windows[®] and a supported VMware[®] Infrastructure, you may need some additional software components, depending on your configuration.

- Microsoft[®] .NET Framework:** vRanger requires the .NET Framework 4.5. The vRanger installer installs it if not detected.
- SQL Server:** [Optional] vRanger utilizes two SQL Server[®] databases for application functionality. vRanger can install a local version of SQL Express 2008 R2 or you can choose to install the vRanger databases on your own SQL instance.
- Windows PowerShell 3 or above.** If you are installing vRanger on Windows 2008 R2 SP1, you will need to install Windows PowerShell 3 or above before installing vRanger
- vRanger virtual appliance (VA):** The vRanger VA is a small, pre-packaged Linux[®] distribution that serves as a platform for vRanger operations away from the vRanger server. vRanger uses the VA for the following functions:
 - Replication to and from VMware[®] ESXi[™] hosts.
 - File-level restore (FLR) from Linux machines.
 - Optionally for backups and restores.
- Updates for Windows Server 2012 R2:** Before installing vRanger on Windows Server 2012 R2, ensure that the Windows updates listed below are installed:
 - KB2939087
 - KB2975061
 - KB2919355

- KB2999226

Installing vRanger

The installation of vRanger has several options. Unless you have a valid reason not to, Quest recommends that you accept the defaults wherever possible. The installation follows this sequence:

- 1 Launch the installer, and accept the vRanger license.
The install process does not continue until the license is accepted.
- 2 Choose an installation directory.
- 3 Enter the credentials under which the vRanger services should run.
Use credentials that have local administrator privileges on the vRanger machine.
- 4 Select the vRanger database.
You may choose to install vRanger with a new instance of SQL Server® Express or on an existing SQL Server. You also need to configure DB credentials now.
The installation completes.

For more information, see the *Quest vRanger Installation/Upgrade Guide*.

Configuring vRanger

vRanger requires some basic configuration before data protection can begin. The bulk of this configuration is driven by the Startup Wizard which starts the first time the application is opened. For more information, see the *Quest vRanger Installation/Upgrade Guide*.

The following topics describe the primary configurations you need to make.

- **Adding vCenters or hosts:** Before you can begin backups, you must add one or more VMware® vCenter™ or ESXi™ servers to the vRanger inventory. To add a vCenter, you must have credentials with Administrator access to the vCenter, along with root-level credentials for each host managed by the vCenter. vRanger provides the option to exclude hosts from the vRanger inventory, which also excludes them from licensing.
- **Adding repositories:** Repositories are where vRanger stores the savepoints created by each backup job. For the purposes of this integration, an EMC® Data Domain® Boost (DD Boost™) repository should be used.

The DD Boost server must be fully configured and running before adding it as a vRanger Repository. For procedures on how to add a DD Boost server as a repository, see [Managing repository replication](#).

Adding an EMC Data Domain Boost (DD Boost) repository

The following process describes how to add an EMC® Data Domain® Boost (DD Boost™) repository. The following procedure assumes that:

- You have a properly configured Data Domain appliance that is accessible to the vRanger machine.
- DD Boost is licensed and enabled on the Data Domain device.
- You have designated a DD Boost User account.

i | **NOTE:** For information on setting up the preceding configurations, see your EMC Data Domain documentation.

To add an EMC Data Domain Boost Repository:

- 1 In the **My Repositories** pane, right-click anywhere, and click **Add > EMC Data Domain Boost (DDB)**.
- 2 In the **Add EMC Data Domain Boost Repository** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.
 - **DNS Name or IP:** Enter the IP or FQDN for the Data Domain device.
 - **User name and Password:** Enter the credentials for the account selected as the **DD Boost User** on the Data Domain device.

i | **NOTE:** For information on creating a DD Boost user name, see your Data Domain documentation.

- **Storage Unit:** This field specifies the Storage Unit configured on the Data Domain device. Select one from the drop-down menu, or enter the name for a new Storage Unit and vRanger creates one for you.

i | **IMPORTANT:** The supported characters are alphanumeric characters and ~!@#\$%^_+`-={}[].,.

- 3 Click **OK**.

The connection to the repository is tested and the repository is added to the **My Repositories** pane and the **Repository Information** dialog box.

vRanger checks the configured repository location for existing manifest data to identify existing savepoints.

- 4 If vRanger finds existing savepoints, click the applicable button:
 - **Import as Read-Only:** To import all savepoint data into the vRanger database, but only for restores, click this button. You cannot back up data to this repository.
 - **Import:** To import all savepoint data into the vRanger database, click this button. vRanger is able to use the repository for backups and restores. vRanger requires read and write access to the directory.
 - **Overwrite:** To retain the savepoint data on the disk and not import it into vRanger, click this button. vRanger ignores the existence of the existing savepoint data and treats the repository as new.

Configuring repository replication

Replication is configured through the My Repositories pane of the vRanger interface. The following procedure assumes that:

- You have configured at least one DD Boost repository in vRanger.
- At least one other Data Domain device with DD Boost enabled is available in your environment to serve as the replication target. You do not have to add this device to vRanger.
- All devices used in replication have a Replicator license installed.

To configure replication:

- 1 In the **My Repositories** pane, right-click the Data Domain Boost repository, and select **Configure Replication**.
- 2 In the **Configure Replication** dialog box, complete the following fields:
 - **Repository Name:** Enter a descriptive name for the repository.
 - **Description:** [Optional] Enter a long-form description for the repository.

- **DNS Name or IP:** Enter the IP or FQDN for the Data Domain device.
- **DD Boost User name and Password:** Enter the credentials for the account selected as the **DD Boost User** on the Data Domain device.
- **Storage Unit:** This field specifies the Storage Unit configured on the Data Domain device. Select one from the drop-down menu, or enter the name for a new Storage Unit and vRanger creates one for you.

3 Click **OK**.

The connection to the device is tested and the device is added as a repository is added to the **My Repositories** pane and the **Repository Information** dialog box.

After replication is configured for a repository, the **Configure Replication** option is disabled for that repository.

i | **IMPORTANT:** After a repository is configured for replication, you must select a synchronization method before replication occurs. For information on manual and automatic synchronization options, see [Managing repository replication](#).

Editing a replication configuration

You may edit an existing replication configuration to update credentials or timeout values. You may also use the **Edit Repository Details** dialog box to view free space for the repository.

- 1 In the **My Repositories** pane, right-click the **EMC Data Domain Boost** repository, and select **Edit Replication Configuration**.
- 2 In the **Edit EMC Data Domain Boost Repository** dialog box, edit any of the following fields:
 - **Repository Name**
 - **User Name**
 - **Password**
- 3 Alternatively, view the **Free Space** field for up-to-date information about this repository.
- 4 If you edited any of the fields, click **Update**.

i | **NOTE:** For information on disabling repository replication, see [Disabling repository replication](#).

Removing an EMC DD Boost repository

The process for removing an EMC® Data Domain® Boost (DD Boost™) repository is the same as removing any other repository type, except that you also have the option of removing the storage unit used by the repository. There is a hard limit on the number storage units that can be configured on a Data Domain device, which makes removing unused storage units an important maintenance effort.

To remove a DD Boost repository:

- 1 In the **My Repositories** pane, right-click the applicable EMC Data Domain Boost repository, and select **Remove**.

The **Remove Repository** dialog box appears, showing the savepoints in the selected repository.

When removing a repository, you have the option of keeping the savepoints on disk or deleting them. To remove the storage unit associated with the repository, you need to remove the savepoints. If replication is configured for this repository, you are also given the option to delete the savepoints in the replicated repository.

- 2 Select **Delete all savepoints in this repository**.

! | **CAUTION:** This step deletes the savepoints from the repository disk, not just the vRanger database. Exercise caution when deleting savepoints.

- 3 To delete the savepoints in a replicated repository, select **Delete all savepoints in replication repository**.

- 4 Click **OK**.
- 5 When the **Delete Savepoints** dialog box appears, select **Delete DD Boost storage unit**, and click **OK**.
The savepoints are deleted, along with the storage unit associated with the repository.

Managing repository replication

vRanger supports repository replication. Repositories configured for replication can be synchronized in one of three ways:

- Automatically, after a successful backup task to that repository.
- As a separate job on a scheduled basis.
- Manually, using the Synchronize option.

Configuring a repository for automatic replication

When a managed repository is configured for savepoint replication, and a backup task completes successfully, each savepoint is also replicated to the replication repository.

To enable automatic savepoint replication:

- 1 From the **Tools** menu on the vRanger toolbar, select **Options**.
- 2 Under the **Repositories** node, select **Replication**.

i | **TIP:** You may also right-click the target repository in the **My Repositories** pane, and select **Repository Replication Options**.

- 3 Select **Enable savepoint replication for a successful backup job task**, and click **Ok**.

Scheduling repository synchronization

When using scheduled repository synchronization, savepoints are replicated to the replication repository according to a configured schedule.

To schedule a repository synchronization:

- 1 From the **Tools** menu on the vRanger toolbar, select **Options**.
- 2 Under the **Repositories** node, select **Replication**.
- 3 In the **Repository Replication Configuration** pane, find the **Repository Name** column, and select the applicable repository.

The repositories listed are the target repositories.

- 4 Select **Schedule repository synchronization**.
- 5 Configure the replication schedule as desired, using the following information as a guide.
 - a **Start:** In the drop-down list, select the time for the replication task to begin.
 - b **Recurrence Pattern:** Establish how often the changes should be synchronized. There are five options within this section:
 - **Daily:** The daily option can be scheduled to synchronize the repository every weekday or every x number of days.
 - **Weekly:** Repository synchronization can be configured to run on weekly intervals, from every week to every 99 weeks. The day of the week on which to run synchronization tasks can be configured.

- **Monthly:** The monthly option offers the following configurations:
 - **Day [x] of every [y] month:**

x can be any value from 1 to 31. This value determines the day of the month on which the synchronization job occurs.

y can be any value from 1 to 99. This value determines the monthly interval—for example, every two months sets the job to run every other month.
 - **The [f] [d] of every [y] month(s):**

f can be either: first, second, third, fourth or last.

d can be: weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.

y can be any value from 1 to 99. This value determines the monthly interval—for example, every two months sets the synchronization task to run every other month.
- **Yearly:** The yearly option offers the following configurations:
 - **Every [m] [x]:**

m is any month of the year. This value determines the month of the year in which the synchronization occurs.

x can be any value from 1 to 31. This value determines the day of the month on which the synchronization occurs.
 - **The [f] [d] of [m]:**

f can be either: first, second, third, fourth, or last.

d can be: day, weekday, weekend day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.

m is any month of the year. This value determines the month of the year in which the synchronization occurs.
- **Interval:** The interval option lets you select the number of days, hours, and minutes that should pass between backup jobs. The interval selected must be greater than or equal to five minutes.

6 Click **OK** to schedule the job.

Synchronizing a repository manually

- 1 In the **My Repositories** pane, right-click the source or replication repository, and select **Synchronize**.
- 2 When the confirmation prompt appears, click **Yes**.

Disabling repository replication

To disable repository replication:


- 1 Do one of the following:
 - From the **Tools** menu on the vRanger toolbar, select **Options**. Under the **Repositories** node, select **Replication**.
 - Right-click the replication repository, and select **Repository Replication Options**.
- 2 Clear the check box for **Enable savepoint replication for a successful backup job task** or **Schedule repository synchronization**, or both.
- 3 Click **OK**.

Replication is disabled as indicated by a red circle icon containing a white exclamation point.

Changing a replication repository to a primary repository

If your primary repository becomes corrupt or otherwise unavailable, you can quickly change your replication repository to a primary repository to continue backup and recovery operations.

To change a replication repository to a primary repository:

- 1 In the **My Repositories** pane, right-click the replication repository, and click **Remove**.
 | **IMPORTANT:** Ensure that **Delete all savepoints in this repository** is not selected.
- 2 Click **OK**.
- 3 In the **My Repositories** pane, click **Add**.
- 4 Select the applicable repository type.
- 5 Complete the **Add Repository** dialog box, and click **OK**.
- 6 When vRanger detects that the repository being added contains savepoint data and displays the **Warning: Existing Repository Found** dialog box, click **Import** to reconfigure the repository as a primary repository.

Administering Data Domain and DD Boost

- [Data Domain system administration](#)
- [Using the DD Boost commands](#)

Data Domain system administration

This topic discusses the following subtopics:

- [Restricting export permissions to DD Boost](#)
- [Modifying an interface group](#)
- [Removing advanced load balancing and link failover configuration](#)
- [Setting network timeouts](#)

Restricting export permissions to DD Boost

By default, the DD Boost service is accessible to all backup server clients. By default, an asterisk (*) wildcard is used. To restrict access to a certain set of backup servers, remove the wildcard from the list and add individual clients to specify only those backup servers.

The list of backup-server clients can contain both FQDNs or short names. The FQDN of the backup host must be correctly configured for reverse lookup in DNS.

To delete all clients from the DD Boost clients list, enter:

```
# ddboost clients delete client-list
```

Optionally, to delete all clients previously added and reset the DD Boost clients list, enter:

```
# ddboost client reset
```

Clients can be added as both FQDNs and short names. To add clients to the DD Boost clients list, enter:

```
# ddboost clients add client-list [encryption-strength {medium | high}
authentication-mode {one-way | two-way | anonymous}]
```

Example:

```
# ddboost clients add ddboost-dl.emc.com ddboost-dl
Addedddboost-dl : Added
```

To view the DD Boost clients list, enter:

```
# ddboost clients show config
Client                Encryption Strength  Authentication Mode
-----
*                      none                 none
*.corp.emc.com        medium               anonymous
```

```
rtp-ost-ms02.domain      high      anonymous
rtp-ost-ms02.domain.com  high      anonymous
```

Notes:

- Disabling DD Boost disables data access to all backup servers.
- Administrators need to ensure that no backup jobs are running to the Data Domain system during these operations.
- An FQDN, an IP address, or a DNS resolvable name of the client must be specified during the modification of the client access list.

Modifying an interface group

After the interface group is set up, you can add or delete interfaces from the group. The following example shows how to remove an interface from the configured interface group on the EMC® Data Domain® system.

Make sure that no jobs are active from vRanger to the Data Domain system on the interface that you want to remove from the group. You can complete this step from the Data Domain system by checking the status of existing connections in the interface group by entering the following command:

```
# ddboost show connections
```

Sample Output:

```
Client Connections:
Max Client Connections: 149
ifgroup status: Enabled

Interface          ifgroup  ----- Connections -----
                  member   Backup   Restore   Control *  Total
-----
192.168.54.138     yes      15        0          2          17
192.168.52.192     yes      15        0          3          18
192.168.55.206     yes      16        0          3          19
-----
Total Connections          46         0          8          54

* Control connections for file-replication, file operations
```

Remove the interface from the group on Data Domain system:

```
# ddboost ifgroup del default interface 192.168.1.3
```

After this step, the interface 192.168.1.3 is released from the group and is no longer used by the DD Boost Storage Server for any jobs from the backup servers.

i | **NOTE:** Removing the interface registered with vRanger makes the Data Domain system inaccessible to the backup servers. The configuration of the ifgroup on the Data Domain system is not deleted.

To change the interface that is added to the interface group on the Data Domain system at the network layer, remove the interface from the group and add it back.

If you use the **net** command to change the interface, such as enabling an interface that is configured for ifgroup, run the **ddboost show connections** command to update the load balancing view. This updating allows the ifgroup to use the interface. For more information about the **ddboost show connections** command, see [Show connections](#).

Removing advanced load balancing and link failover configuration

The following example illustrates removing a configured interface group on the EMC® Data Domain® system.

- 1 Make sure that no jobs are active from vRanger to the Data Domain system.

You can complete this step from the Data Domain system by checking the status of existing connections in the interface group with the following command:

```
# ddbboost ifgroup show connections
```

- 2 Ensure that there are no pending jobs from any of the connected backup servers to this system.
- 3 Disable the feature on the Data Domain system:

```
# ddbboost ifgroup disable
```

- 4 Reset the interface group:

```
# ddbboost ifgroup reset
```

All the interfaces are released from the group. However, backup servers can still access the DD Boost storage server on the Data Domain system on the interface registered with vRanger. In the preceding example, the Data Domain system is still registered with vRanger using 192.168.1.1.

Setting network timeouts

Backup and restore jobs often take a long time to complete. Although the EMC® Data Domain® Boost (DD Boost™) Library can recover from temporary network interruptions, the operating system on vRanger might terminate a job prematurely if vRanger timeouts are set too low.

EMC recommends setting timeouts to at least 30 minutes (1800 seconds).

i | **NOTE:** After losing a network connection, administrators should issue the `ddbboost reset stats` command to clear job connections.

Using the DD Boost commands

This topic describes the EMC® Data Domain® Boost (DD Boost™) commands that you use to manage the DD Boost feature on the EMC® Data Domain® system.

i | **NOTE:** For information about configuring DD Boost using the graphical user interface (GUI)-based Enterprise Manager, see the *EMC Data Domain Operating System Administration Guide*.

The **ddbboost** command manages the integration of Data Domain systems and disk backup devices. DD Boost is a licensed feature. If basic options do not work, verify that the proper licensing has been implemented on your Data Domain system. The **ddbboost** command includes the following options.

Access

Table 1.

Description	Manage DD Boost access for clients.
Syntax	<ul style="list-style-type: none">• ddboost access add clients: Add clients to a DD Boost access list. <code>ddboost access add clients client-list</code>• ddboost access delete clients: Delete clients from DD Boost access list. <code>ddboost access del clients client-list</code>• ddboost access reset: Reset DD Boost client access list to factory default. <code>ddboost access reset</code>• ddboost access show: Show DD Boost client access list. <code>ddboost access show</code>

Destroy

Description	Administrators use this command to delete all storage units and their contents from the Data Domain system. The command permanently removes all the data (files) contained in the storage units. You must also manually remove (expire) the corresponding catalog entries.
Syntax	<code>ddboost destroy</code>

Disable

Description	Administrators use this command to disable DD Boost for the Data Domain system.
Syntax	<code>ddboost disable</code>

Enable

Description	Administrators use this command to enable DD Boost for the Data Domain system. Whenever the user, user ID (UID), or group ID (GID) changes, the Data Domain system updates all files and storage units the next time that the <code>ddboost enable</code> command is issued.
Syntax	<code>ddboost enable</code>

Ifgroup add interface

Description	<p>Administrators use this command to add an IP address to a private network that is to process data transfer. The IP address must be configured on the Data Domain system, and its interface enabled.</p> <p>The command checks if the IP address is valid. If not, a message reports the IP address is invalid—either because it is not configured on the Data Domain system, or its interface is not enabled. Use the <code>net show settings</code> or <code>net show config</code> commands to view available IP addresses.</p> <p>If the IP is valid, the message IP added to list is displayed. After adding an IP address as an interface, you need to enable Advanced Load Balancing and Link Failover using the <code>ddboost ifgroup enable</code> command.</p> <p>Notes:</p> <ul style="list-style-type: none">• If an interface—or a network interface card (NIC) that has multiple interfaces—fails, all the in-flight jobs to that interface transparently fail over to a healthy interface in the ifgroup. Any jobs that are started after the failure are routed to the healthy interfaces. You can add public or private IP addresses for data transfer connections.• Quest does not recommend that you use 1 GbE and 10 GbE connections in the same interface group.• You can add public or private IP addresses for data transfer connections.
Syntax	<pre>ddboost ifgroup add interface <IP address></pre>

Ifgroup del interface

Description	<p>Administrators use this command to remove a valid IP address added as an interface.</p> <p>If you are attempting to delete the last IP address interface, you are informed that doing so disables the ifgroup, and you are given the option of terminating this command.</p>
Syntax	<pre>ddboost ifgroup del interface <IP address></pre>

Ifgroup disable

Description	<p>Administrators use this command to disable the Advanced Load Balancing and Link Failover option for new jobs. Currently running jobs are not affected and continue to run.</p> <p>If DD Boost is disabled, the message disabled is displayed.</p>
Syntax	<pre>ddboost ifgroup disable</pre>

Ifgroup enable

Description	<p>Administrators use this command to enable Advanced Load Balancing and Link Failover. At least one IP address must have already been added.</p>
Syntax	<pre>ddboost ifgroup enable</pre>

Ifgroup reset

Description	Administrators use this command to remove the interfaces for Advanced Load Balancing and Link Failover and to disable the ifgroup. The ifgroup reset command is equivalent to issuing the ddboost ifgroup disable command followed by issuing multiple ddboost ifgroup del interface <IP address> commands.
Syntax	<code>ddboost ifgroup reset</code>

Ifgroup show config

Description	Display IP addresses added to the interface group.
Syntax	<code>ddboost ifgroup show config</code>

Option reset

Description	Administrators use this command to reset distributed segment processing to the default, which is enabled.
Syntax	<code>ddboost option reset distributed-segment processing</code>

Option set

Description	Administrators can use this command to enable or disable the distributed segment processing feature on the Data Domain operating system. By default, it is enabled. If this feature presents any problem for a backup server, use this command to disable the feature on the Data Domain system. Distributed segment processing is supported only if the installed version of the DD Boost Library is version 2.2 or later, and the feature is enabled on the Data Domain operating system. Notes: <ul style="list-style-type: none">• Distributed segment processing is enabled by default.• If a system is upgraded from Data Domain OS 4.9 to Data Domain OS 5.0 and the DD Boost option was disabled before the upgrade, the distributed-segment-processing option is disabled after the upgrade.
Syntax	<code>ddboost option set distributed-segment-processing {enabled disabled}</code>

Option show

Description	Use this command to show whether the distributed segment processing option is enabled or disabled.
Syntax	<code>ddboost option show distributed-segment-processing</code>

Reset stats

Description	Administrators use this command as follows: <ul style="list-style-type: none">• To reset all statistics when DD Boost is enabled, or• As a network recovery procedure to clear job connections after the network connection has been lost.
Syntax	<code>ddboost reset stats</code>

Reset user-name

Description	Delete and reset the DD Boost user name.
Syntax	<code>ddboost reset user-name</code>

Set user-name

Description	Set the DD Boost user name when DD Boost is enabled.
Syntax	<code>ddboost set user-name <user-name></code>

Show connections

Description	Use this command to show the number of active clients, and for each active client: <ul style="list-style-type: none">• The client's name.• Whether the client is idle (yes/no).• The number of installed CPUs.• The amount of memory in MiB.• The installed Library version number.• The name and version of the operating system.• vRanger version number.• The number of connections that a system uses for DD Boost, and whether these connections are spread across interfaces using Advanced Load Balancing and Link Failover.• The number of connections used for a given group, even when the interface has been removed from the ifgroup.• An overview of interfaces available for DD Boost.
Syntax	<code>ddboost show connections</code>

Show histogram

Description	<p>Use this command to display a DD Boost histogram for the Data Domain system.</p> <ul style="list-style-type: none">• op: The name of the NFS operation.• mean-ms: The mathematical mean time for completion of the operations.• stddev: The standard deviation for time to complete operations, derived from the mean time.• max-s: The maximum time taken for a single operation.• 2,4,6,8, or 10ms: The number of operations that took less than the specified number of milliseconds (ms).• 100ms: The number of operations that took between 10 ms and 100 ms.• 1s: The number of operations that took between 100 ms and one second.• 10s: The number of operations that took between 1 second and 10 seconds.• >10s: The number of operations that took over 10 seconds.
Syntax	<pre>ddboost show connections</pre>

Show stats

Description	<p>Use this command to show the output of previous show stats command, the number of bytes written to and read from files contained in storage units, and the number of files created and deleted from storage units. The number of errors encountered for each operation is also shown. Optionally, you can show statistics for a specific number of seconds. The count displays the number of lines equal to the count value.</p> <p>If distributed segment processing is enabled, the number of bytes transferred through distributed segment processing and the amount of compression achieved is displayed.</p>
Syntax	<pre>ddboost show stats [interval seconds] [count count]</pre>

Show user-name

Description	Display the current DD Boost user.
Syntax	<pre>ddboost show user-name</pre>

Status

Description	Enter this command to determine whether DD Boost is enabled or disabled.
Syntax	<pre>ddboost status</pre>

Storage-unit create

Description	Administrators use this command to create a storage unit with a given name.
Syntax	<pre>ddboost storage-unit create storage-unit-<name></pre>

Storage-unit delete

Description	Administrators use this command to delete a specified storage unit and all its contents. You must also manually remove (expire) the corresponding catalog entries.
Syntax	<code>ddboost storage-unit delete storage-unit-<name></code>

Storage-unit show

Description	<p>Use this command to display the names of all storage units or, optionally, the names of all files in a specified storage unit.</p> <p>To display the compression for all storage units (the original byte size, global compression, and local compression for all storage units), enter the compression option.</p> <p>NOTE: To interrupt the output of this command, press Ctrl+C.</p>
Syntax	<code>ddboost storage-unit show [compression] [storage-unit-name]</code>

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 <https://www.quest.com/company/contact-us.aspx> or call +1-949-754-8000.

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.