

Quest® vRanger® & EMC Data Domain Deduplication: Optimizing VMware Backup & Recovery with DD Boost

Written By Sheldon D'Paiva
Product Manager
Quest Software



Contents

Executive Summary	2
EMC Data Domain	3
Data Domain Boost	4
Quest® vRanger® Pro	5
Introduction	5
Backup Modes	5
LAN-free Backup (Physical Machine Deployment Only)	5
Hot-Add Backup (Virtual Machine or Virtual Appliance Deployment)	7
Network Backup (Physical Machine, Virtual Machine, or Virtual Appliance Deployment)	9
Change Block Tracking (CBT) and Active Block Mapping (ABM)	10
Resource Management	11
Repositories	12
Cataloging	13
Joint Solution: EMC Data Domain with Boost and vRanger	14
Licensing	14
Overview of Installation and Setup	14
Data Domain Boost and vRanger Integration	15
Conclusion	16
For More Information	16
Appendix 1: Test Results	17
Methodology and Setup	17
Results	18

Executive Summary

Backup applications are a critical component of backup, recovery, and disaster preparedness strategies and often direct the operation of media that typically store five to ten times more data than can be found on the primary storage in the same data center. As data continues to grow unabatedly, traditional approaches no longer deliver the efficient data protection and disaster recovery (DR) organizations need. Virtual environments have an additional requirement: data protection and DR approaches must not jeopardize the benefits of virtualization, such as increased server utilization, operational agility, and ease of management. Increasingly, disk-based backup and deduplication products are becoming the solution of choice to manage data protection efficiently in virtual environments.

EMC Data Domain systems with Data Domain Boost technology and Quest® vRanger Pro® combine to deliver an efficient, high-performance, and robust solution for backing up and recovering VMware environments. In this paper, we will show how you can effectively manage data growth in VMware environments by using deduplication. Specifically, we will:

- Review the EMC Data Domain deduplication storage system capabilities
- See how Data Domain Boost further enhances Data Domain systems
- Take a deeper dive into vRanger and how it can be used with Data Domain Boost

EMC Data Domain

EMC Data Domain deduplication storage systems offer an alternative to traditional backup approaches. With a simple NFS/CIFS interface, the Data Domain systems are simple to integrate with existing backup software, and allow users to enjoy the retention and recovery benefits of inline deduplication as well as the offsite disaster recovery protection of replication over the wide area network (WAN).

The Data Domain operating system (DD OS) offers aggressive inline data deduplication for backup and recovery of data, averaging 10–30x data reduction. As data is written to an EMC Data Domain deduplication storage system, it is quickly scanned for patterns that have been stored before. Large patterns across the history of all data stored are identified in the Data Domain filesystem, regardless of application or workload. DD OS stores each unique data sequence only once and saves significant physical storage capacity by substituting small references for each identical redundant sequence. Local compression scans the unique data sequence for small strings across a local small window of comparison, like a tape drive. The combination of both deduplication and compression provides the resulting optimized data reduction.

DD OS supports petabytes of storage for a typical enterprise dataset and backup policy. For environments that need to retain backups for long periods of time to meet compliance and regulatory requirements, the DD990 with Extended Retention Software option supports logical system capacity of up to 65 PB for long-term retention of data. Multiple months of retention on disk is now possible using the same number of “floor tiles” that would normally provide only a couple days of disk staging. Snapshot technology further enables extended local and offsite retention on disk.

Once a backup is performed to a Data Domain system running DD OS, it is easy and cost-effective to use the system for local retention and to limit the use of tape to provide offsite and archive support. DD OS protects the data with network-efficient and encrypted replication, which provides remote office data protection and enables multi-site tape consolidation.

Data Domain systems replicate only the deduplicated and compressed unique changes across any IP network, resulting in up to 99 percent bandwidth reduction. A Data Domain system facilitates a copy of the entire retained dataset, online and disaster-protected. If multiple systems replicate to the same destination, the destination will only store each segment uniquely across all inbound replication streams, further minimizing bandwidth and storage. If confidentiality is required, deduplicated and compressed data can be encrypted in-flight when being replicated between Data Domain systems, independent of the replication topology used.

Data Domain Boost

EMC Data Domain Boost extends the optimization capabilities of Data Domain solutions. DD Boost significantly increases performance by distributing parts of the deduplication process to the backup server or application clients, simplifies disaster recovery procedures, and serves as a solid foundation for additional integration between vRanger and Data Domain systems.

Without DD Boost, the backup server will send all data, unique or redundant, to a Data Domain system for deduplication processing. With DD Boost, parts of the deduplication process are distributed to the backup server or application clients, enabling it to send only unique data segments to a Data Domain system. This dramatically increases the aggregate throughput, up to 31.0 TB/hour, and reduces the amount of data transferred over the network by 80 to 99 percent. These efficiencies can help eliminate future costs by leveraging existing backup servers and Ethernet networks.

DD Boost also increases the speed of restart and completion of failed backups. Since only unique data is sent over the network, once a failed job restarts, the data that has already been sent to the Data Domain system for a given backup job does not need to be sent again. This not only reduces the load on the network substantially, but also improves the overall throughput for the failed backups upon retry.

Overall, DD Boost increases aggregate throughput, substantially reduces backup windows, and improves backup server and application client efficiency. Best of all, DD Boost integrates seamlessly with vRanger because the distributed segment processing is handled by the DD Boost Library on the vRanger server or virtual appliance.

The following are the key benefits of DD Boost:

- Seamless integration and scalability with vRanger
 - Works as any other repository within the vRanger GUI
- Can work with the vRanger server or virtual appliance for enhanced scalability
- Significant reduction in backup time
 - Up to 31.0 TB/hr aggregate throughput performance
 - Distributed deduplication process dramatically increases throughput
 - Reduced network bandwidth utilization
 - Faster restarts of failed backup jobs
- Advanced load balancing and link failover
 - Scalable link aggregation at the application layer
 - Simplified backup application configuration
 - Seamless load balancing of jobs among available ports
 - Link failover keeps backups operational

Quest[®] vRanger[®] Pro

Introduction

vRanger delivers fast, reliable, and easy-to-use image-based backup, replication, and recovery for VMware virtual machines. It is VMware Ready[™]-certified for vSphere 5, which means it seamlessly integrates with and efficiently supports VMware infrastructures. All backup and replication jobs are executed while the protected virtual machine is running, so vRanger jobs do not interrupt critical application services.

An intelligent resource manager improves backup speed by leveraging all available ESX and ESXi hosts and controlling the number of simultaneous jobs using any given resource. To optimize use of disk storage and make multiple recovery points (savepoints) economical, vRanger supports several space-saving mechanisms, including VMware Change Block Tracking (CBT), incremental backup, differential backup, compression, and Active Block Mapping (ABM). These space-saving technologies reduce the size of the backups sent to the Data Domain System, which further reduces the physical size of the backup with its inline deduplication and compression. The vRanger space-reduction technologies and appliance deduplication and compression provide optimal end-to-end data compaction.

For scalability, vRanger allows you to connect to multiple vCenter servers to easily protect larger or distributed environments while consolidating backup data onto the Data Domain System.

Backup Modes

Three backup modes are supported for VMware environments:

- LAN-free backup (physical machine deployment only)
- Hot-Add backup (virtual machine or virtual appliance deployment)
- Network backup (physical machine, virtual machine, or virtual appliance deployment)

Note that all three backup modes below will work with DD Boost; there are differences in how the backup data is read, but all three modes will use DD Boost to leverage DD Boost benefits if vRanger is set up to use a DD Boost repository. For more detailed information, please refer to the [Quest vRanger Deployment Guide](#).

LAN-free Backup (Physical Machine Deployment Only)

Installing vRanger on a physical machine with either a fibre channel or iSCSI SAN connected host environment provides scalability and performance for backing up VMware environments. In this configuration, vRanger can perform LAN-free backups by invoking VMware's SAN transport mode. LAN-free backup can be utilized as the primary backup method, while network backup can be employed for fail-safe redundancy in the event LAN-free fails.

The benefits of this configuration include:

- Isolating backup and restore traffic to the fibre channel or iSCSI network, which completely offloads your production hosts and network from data protection overhead and traffic
- Enabling extremely high performance and scalability for data protection operations
 - Adding HBAs or NICs to the backup machine (one dedicated to reading and another to writing) can further increase performance and scalability

The LAN-free backup configuration is illustrated below. Here vRanger acts as a physical proxy server with storage connections to VMFS and the backup repository. The vRanger machine must be attached to fibre channel or iSCSI SAN environment, and the VMFS volumes containing the protected virtual machines must also be presented to the vRanger machine. vRanger initiates the virtual machine snapshot and reads the VMDKs directly from the fibre SAN or iSCSI LUN that stores the virtual disks. The backup data traffic flows through the vRanger server and is then written to the repository.

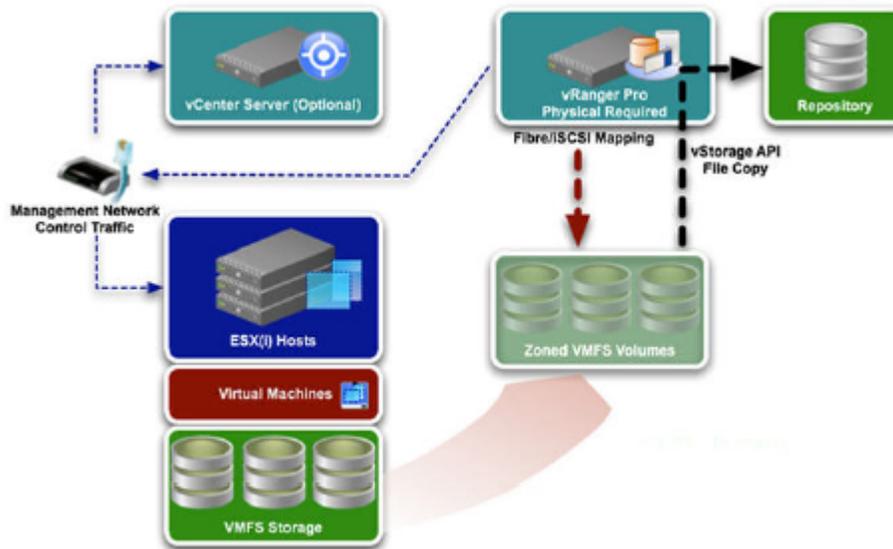


Figure 1: LAN-free backup mode

The steps in LAN-free backup mode can be further broken down as follows:

1. vRanger retrieves the job configuration from the database.
2. vRanger establishes a connection to the repository.
3. vRanger adds a temporary snapshot to the protected virtual machine.
4. vRanger mounts the protected VMDKs to the vRanger physical machine, using vStorage API's proxy mount and SAN advanced transport mode.
5. vRanger reads from the source disk and transfers data from the protected host to the repository.
6. vRanger deletes the temporary snapshot on the protected virtual machine.

There are a number of factors to consider in this configuration, including the following:

- **The number of virtual machines protected by each vRanger physical proxy server** – As a starting point, you should use one proxy server for every 10 hosts or every 100 virtual machines.
- **The number of virtual machines that can be processed concurrently** – The number of concurrent tasks running locally can be set between one through 20 via the Options menu, with the default set to three. As a starting point, you should use one task per CPU core. If your server has eight cores, for example, start with the maximum number of tasks running locally set to eight.
- **Memory and bandwidth** – Scalability will also depend on the amount of memory on the vRanger machine and network bandwidth available to write to repositories.

You should disable automount on the vRanger machine so that Windows does not assign drive letters to protected VMDKs when they're mounted to the vRanger machine. Do not initialize or format unknown or offline disks from the vRanger machine; these represent your VMFS volumes and any changes could potentially corrupt the VMFS volumes.

Hot-Add Backup (Virtual Machine or Virtual Appliance Deployment)

vRanger must use a virtual machine to leverage Hot-Add transport. The Hot-Add transport mode involves an SCSI Hot-Add on the host where vRanger is running or on a virtual appliance as noted at the end of this section. Hot-Add backup works with networked or local storage. You can use vRanger's Hot-Add backup as the primary method and utilize network backup as a secondary mode for redundancy.

The benefits of the Hot-Add backup configuration include:

- Ease of management, since vRanger is deployed in a virtual machine
- Fast read speeds – native virtual machine performance

In this configuration, vRanger is provided with direct access to protected VMDKs through vSphere's I/O stack rather than through the network. The vRanger virtual machine acts as a proxy server since the VMDKs are temporarily mounted directly, and all backup data traffic flows through the vRanger virtual machine to shared or local storage.

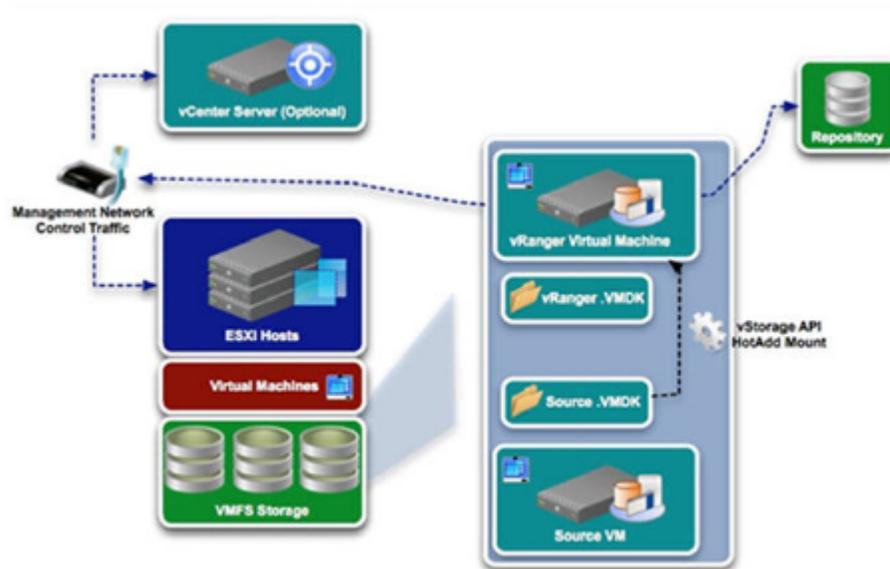


Figure 2: Hot-Add backup

The steps in the Hot-Add backup mode can be further broken down as follows:

1. vRanger retrieves the job configuration from the database.
2. vRanger establishes a connection to the repository.
3. vRanger adds a temporary snapshot to the protected virtual machine.
4. vRanger mounts the protected VMDKs to the vRanger virtual machine, using vStorage API's Hot-Add advanced transport mode.
5. vRanger reads from the source disk, and transfers data from the protected host to the repository.
6. vRanger deletes the temporary snapshot on the protected virtual machine.

In order to employ Hot-Add, the vRanger virtual machine must be able to access the datastores for all protected virtual machines. Make sure the protected virtual disks on the datastore have block sizes that match the block sizes for the vRanger virtual machine datastore. Hot-Add works only for virtual machines with SCSI disks, not IDE disks.

If the vRanger virtual machine can be vMotioned, all hosts that it can be vMotioned to must be able to see the storage for all protected virtual machines as well. Also, be sure to disable automount on the vRanger virtual machine so that Windows does not assign drive letters to protected VMDKs when they're mounted to the vRanger virtual machine.

vRanger can process a significant amount of data to the repository very quickly, provided it has enough available resources. As a guideline, allow for two concurrent jobs per vCPU. We recommend configuring the vRanger virtual machine with a minimum of four vCPUs – one for the vRanger server and three for job activities, which allows for up to six concurrent jobs. If you need to run a higher number of concurrent jobs, simply add additional vCPUs.

You should consider adding a second SCSI controller to the vRanger virtual machine so that it can mount more protected disks at the same time during the Hot-Add operation. vRanger can mount as many virtual disks as vSphere allows.

In vRanger version 5.4 and later, you can also deploy a virtual appliance to offload the vRanger server for increased scalability. You have the flexibility of deploying one virtual appliance per host, or one virtual appliance per cluster depending on the needs of your environment. In this scenario, backups will work the same way and deliver the same benefits as noted above, except that the backups will be processed by the virtual appliance instead of the vRanger VM.

Network Backup (Physical Machine, Virtual Machine, or Virtual Appliance Deployment)

In network backup mode, vRanger backs up virtual machines via each host over the LAN utilizing VMware's LAN transport method.

The main benefit of this backup mode is that it has the least restrictions; it can work with almost any configuration. However, this flexibility comes at the expense of slower performance; therefore, in most environments, you should use network backup mode only as a failover option. To ensure backup data traffic does not flow over the network at any time, you can disable the option to perform a network backup on failure of the other backup modes.

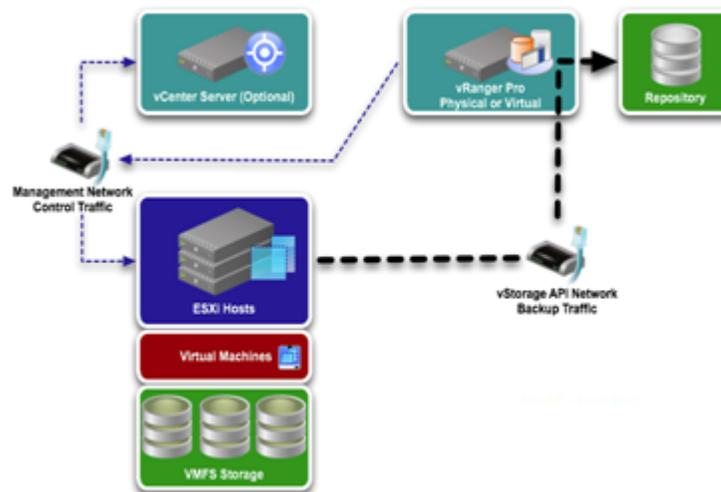


Figure 3: ESXi network backup

As a guideline, allow for two concurrent jobs per vCPU. We recommend configuring the vRanger virtual machine with a minimum of four vCPUs – one for the vRanger server and three for job activities, which allows for up to six concurrent jobs. If you need to run a higher number of concurrent jobs, simply add additional vCPUs.

Change Block Tracking (CBT) and Active Block Mapping (ABM)

vRanger supports VMware Change Block Tracking (CBT), which is a vSphere feature that tracks changed blocks within the virtual machine disk (VMDK) file. With CBT enabled, vRanger efficiently identifies the blocks that have changed since the last backup. Enabling CBT results in fewer blocks backed up, shorter backup times, and reduced storage requirements.

vRanger also includes patented Active Block Mapping (ABM), which can work either with or independently from VMware CBT. With ABM enabled, vRanger will back up only changed blocks with active data, and skip unallocated, deleted, and zero-filled blocks. ABM increases backup efficiency of Windows virtual machines by protecting only the blocks with active data. ABM with CBT can reduce the amount of data backed up by up to one third. ABM can be enabled on a per job basis, as one of the options in the job configuration wizard.

The figures below show how CBT and ABM work together to speed, and reduce the size of, incremental backup.

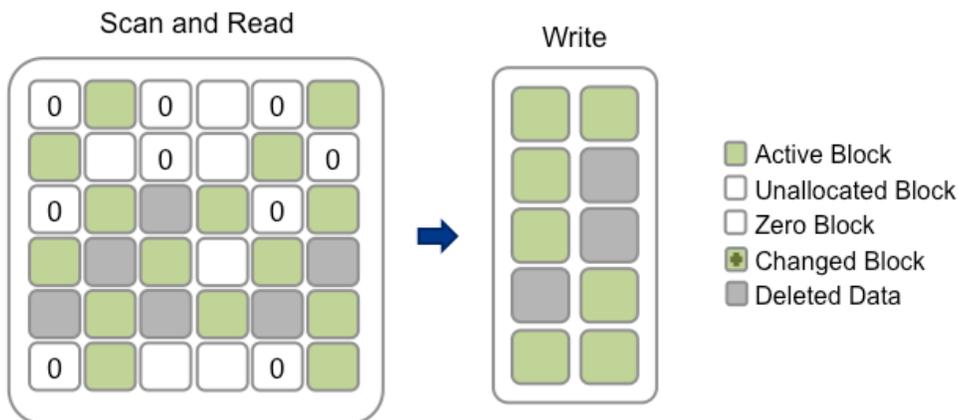


Figure 4: Incremental backup without ABM or CBT

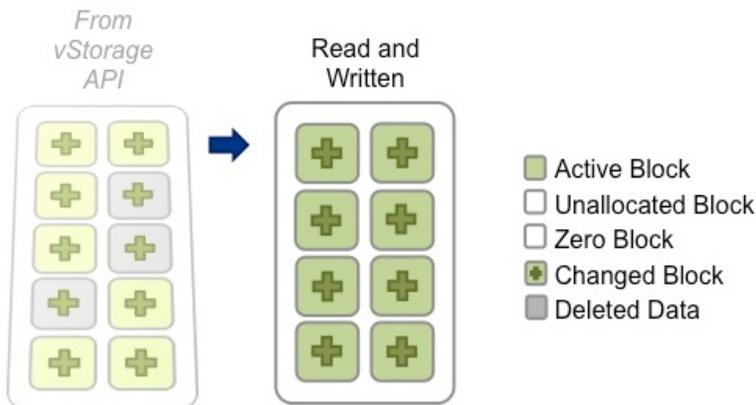


Figure 5: Incremental backup with CBT and ABM

Backups with CBT and ABM may result in lower deduplication ratios achieved in Data Domain compared with backups without CBT and ABM, because the amount of data sent to Data Domain will be reduced.

Resource Management

vRanger's Resource Manager allows control of the number of jobs run concurrently against any given resource. Resource Manager settings can enable more or less throughput to fine tune and optimize vRanger performance in connection with VMware hosts, datastores, and repositories.

The vRanger Resource Manager configuration options include:

- **Maximum number of tasks running on vRanger (>=1)** – This value limits the total number of simultaneous tasks for the vRanger machine. This is affected by machine type, hardware, and deployment architecture. The default value is 100, which does not normally need to be changed since other limits take effect before this value is reached.
- **Maximum number of tasks running off a LUN (1–5)** – In order to avoid storage I/O contention issues, you should limit the number of tasks that can be processed based on their storage location. vRanger defaults to a limit of three backup tasks per LUN. Note that this assumes a 1:1 mapping between the underlying LUN and the datastore.
- **Maximum number of tasks running on a host (1–4)** – This setting should be set fairly low because all of the backup processing (for service console backups) occurs on the host. The default per-host limit is set to one. Due to the amount of CPU and memory used by the host to process backups, it is recommended to use caution when changing this value.
- **Maximum number of tasks running per repository (>=1)** – Bandwidth to the repository is the main limiting factor here, and vRanger defaults this value to 3. For the Data Domain system on a gigabit network, the recommendation is a maximum of 10 concurrent backups. Consider using link aggregation or 10GbE network for increased throughput.
- **Maximum number of tasks running locally (1–20)** – This option limits the number of simultaneous tasks on the vRanger server. The default is set to 3. The recommended maximum is one task per CPU core for LAN-free (SAN) mode and one to two tasks per vCPU for Hot-Add mode (start with one task per vCPU and increase if CPU allows). Monitor CPU utilization and increase number of CPUs if needed to be able to process more backups (see the licensing section below).
- **Maximum number of tasks running per VA (1– 8)** – This option limits the number of simultaneous tasks on a virtual appliance. With the default set to 2, you should have a virtual appliance deployed with 2 CPUs and 1 GB of RAM. You should increase this accordingly if you want to run more simultaneous tasks on a VA; for example, use the recommended setting of 2 CPUs and 2 GB of RAM for 4 concurrent tasks per virtual appliance.

Repositories

vRanger uses a repository to store backup savepoints. A repository consists of a configuration file (GlobalManifest.metadata) and root-level directories for each protected object. Any time you add a repository in vRanger Pro, a GlobalManifest.Metadata XML file is created in the selected folder. It is the presence of that manifest file that tells vRanger that a repository exists in that folder. Whenever you add a repository, that repository contains a global manifest XML file that provides details about the repository.

As you do backups, an entry is placed in the global manifest file. Inside the root-level directory for each object are sub-directories created for each full backup of the object in question. Differential or incremental savepoints will reside within the full backup subdirectory. Inside the sub-directories are archives for every file protected during that task. Also in the subdirectory are two metadata files. Since each savepoint stores the manifest file associated with the backup, the backup can be recovered even if the vRanger database is not available.

The image below depicts a basic repository structure:

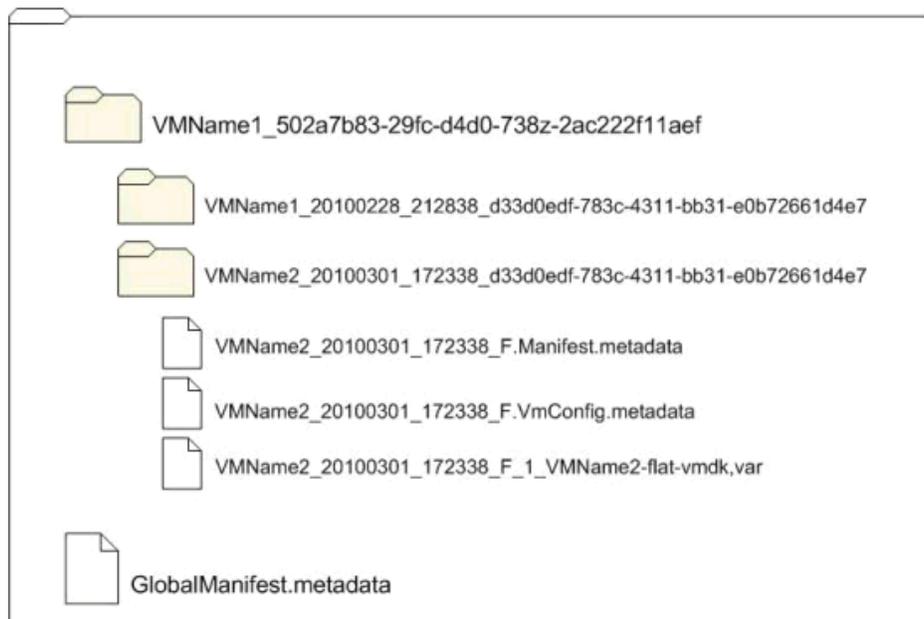


Figure 6: VMware repository structure

vRanger includes an option for encrypting repositories; this option should not be used when backing up to a repository on the Data Domain system because encrypted data cannot be efficiently deduplicated.

Cataloging

vRanger provides the option to index backup data to create a catalog. The catalog provides a simple and fast way to browse and search for file-level recovery across all the repositories. Files are indexed during the backup operation, and indexing does not extend the backup window, so search results are available immediately at recovery time.

Catalog collections are disabled by default, because there may be a slight performance impact. Collection must be enabled globally through vRanger's Tools>Options menu. Collection must also be enabled on a per-job basis. Once enabled, if catalog collections are disabled in the future, the catalog data will be retained in line with retention policies.

Joint Solution: EMC Data Domain with Boost and vRanger

Licensing

Data Domain Boost is licensed separately by EMC, and you must obtain the license required to enable DD Boost on the Data Domain system from EMC. The DD Boost license allows you to back up and restore data from vRanger.

Data Domain Boost can be licensed only with the “Pro” edition of vRanger (vRanger Pro), and not the standard edition (vRanger Standard). Every ESX or ESXi host for which vRanger provides protection must be properly licensed.

vRanger is licensed by the number of source CPUs that you can configure for backup. For licensing purposes, a dual-core processor is counted as a single CPU. Therefore, one ESX(i) host with two dual-core processors would use two CPUs of a license.

vRanger does not support the free version of VMware ESXi. In order to use VMware’s vStorage APIs, a minimum of “vSphere Essentials” licensing is required from VMware.

Overview of Installation and Setup

The overall steps for installing Data Domain Boost are as follows:

1. Obtain the license required to enable DD Boost on the Data Domain system. You can purchase a DD Boost license key directly from EMC.
 - o The DD Boost license allows you to back up and restore data.
 - o 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, and creating storage units.
3. Add a DD Boost Repository to your vRanger installation.

vRanger supports systems with Data Domain OS versions 5.1 and above, and the minimum network requirement for vRanger with DD Boost is a Gigabit Ethernet (GbE) link.

For more information, refer to the product documentation noted at the end of this paper.

Data Domain Boost and vRanger Integration

Data Domain Boost software components are automatically deployed with your vRanger installation, so there are no additional components to install. vRanger tightly integrates with the DD Boost components through the DD Boost API to run parts of the deduplication process so that duplicate data is not sent whenever a DD Boost repository type is used for a vRanger backup job. These components are available on both the vRanger server and any vRanger virtual appliances that you deploy, so that you can have maximum flexibility and scalability when using DD Boost.

The Data Domain system exposes disk volumes called storage units to vRanger. These storage units can be created from within vRanger through the vRanger GUI. While you can create a maximum of 100 storage units on a Data Domain system, a maximum of 14 is recommended so that you do not affect system performance.

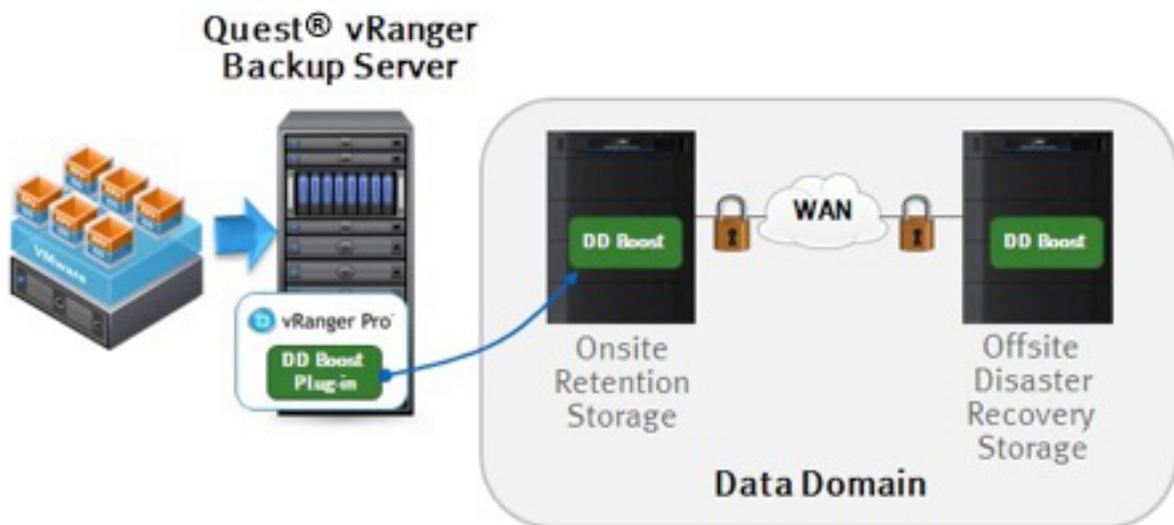


Figure 7: vRanger integrates with DD Boost for efficient backups that target your Data Domain appliance

Note that although vRanger offers native encryption and compression, these options will result in reduced performance with a Data Domain Boost repository. Because of this, these options cannot be enabled from within vRanger when using a Data Domain Boost repository. Note, however, that Data Domain appliances themselves have compression and encryption that you can use.

Conclusion

EMC Data Domain systems with Data Domain Boost technology and Quest vRanger Pro combine to deliver an efficient, high-performance, and robust solution for backing up and recovering VMware environments. The joint solution integrates seamlessly and delivers high-performance deduplicated backup and recovery from disk that is easy to set up and manage. This document described implementation and best practices for an end-to-end solution. The settings should be tuned in a non-production environment before implementation.

The following is a summary of the best-practice recommendations for the joint solution:

- Make sure that you are running vRanger with DD Boost over at least a 1 GB/s link.
- Group the backup of multiple virtual machines into vRanger backup jobs and use its Resource Manager to control the number of concurrent backups per datastore, VMware host, and repository.
- Use vRanger to enable VMware's Change Block Tracking on the virtual machines for the fastest-possible backups.
- Enable vRanger's Active Block Mapping to reduce the size of the backups sent to the EMC Data Domain system.
- Run vRanger in a virtual machine for small- to medium-sized VMware installations where costs must be minimized.
- Run vRanger on a dedicated physical machine in SAN backup mode for larger environments to improve backup performance, or deploy virtual appliances for increased scalability.
- Use the Data Domain Sizing Tool to size the Data Domain component of the solution.
- Ensure the firmware and software in the Data Domain Backup System are fully up to date.
- Remember the backup is to deduplicating disk, not to general-purpose disk, so configure time for the Data Domain housekeeping process to run every day.
- Understand the settings that can be tuned to configure the best-performing protection solution and test the configuration in a non-production environment before implementation.

For More Information

For Data Domain documentation, see <https://my.datadomain.com/documentation>.

For vRanger documentation, visit <http://www.quest.com/vranger/>.

Appendix 1: Test Results

Methodology and Setup

The tests described in this section show the benefit that Data Domain Boost can provide, but note that results can vary depending on the data and type of backups you use in your environment. The test was performed using full backups of eight VMs, first to CIFS and then to DD Boost repositories on the same Data Domain device, with datasets added to or deleted from the VMs between each backup to simulate a certain amount of data change. The eight VMs were spread across multiple hosts and datastores to achieve eight parallel backup streams, with eight savepoints for each VM. The Data Domain system was cleaned of all data when switching from using CIFS to DD Boost repositories to prevent the DD Boost backups from benefiting from the previously done CIFS backups.

The test setup was as follows:

vRanger Configuration

- vRanger version: 5.4 (internal build)
- vRanger system:
 - Operating system: Windows Server 2003 R2 Enterprise
 - CPU: Dual 2.8 GHz Intel Xeon
 - RAM: 2 GB
- Options used: ABM, Cataloging
- LAN: Gigabit Ethernet
- Transport: Fiber SAN to VM datastores

Data Domain Configuration

- Appliance: DD860
- OS Version: 5.1

VMware Environment

- vCenter version: 5.0
- Host versions: 4.0, 4.0i, 4.1, 4.1i, 5.0i
- VM contents
 - Disk size: 11 GB
 - Operating system: Windows Server 2003 R2 Enterprise
 - Dataset: Two known data files (common to all VMs), and two random data files (between 1 to 200 MB)

Results

- Average runtime for full backups of eight VMs using CIFS: **34.75 min**
- Average runtime for full backups of eight VMs using DD Boost: **19.87 min**
- Reduction in runtime: **14.88 min**

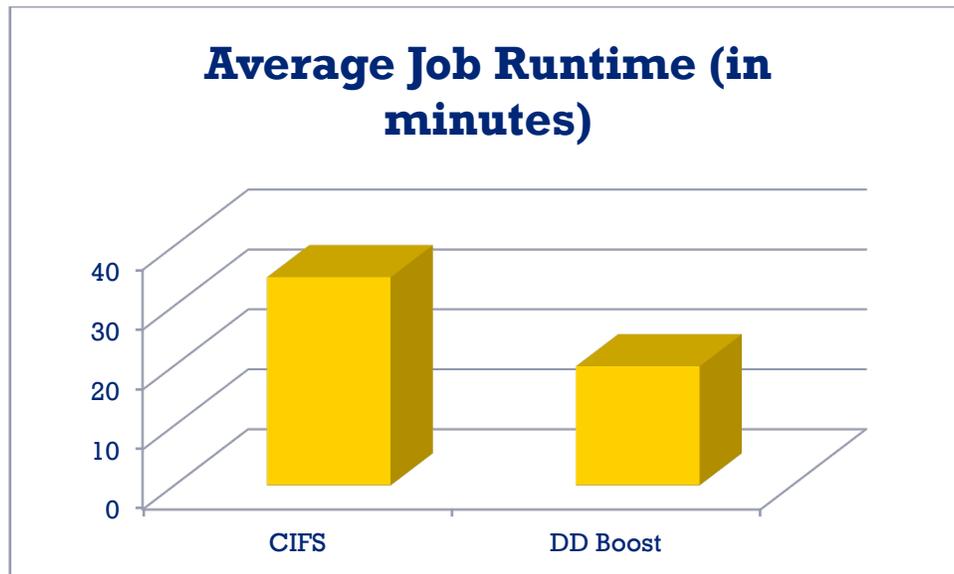


Figure 8. Test results: average job runtime

On average, using vRanger with DD Boost **reduced job runtime by 42.8 percent.**

© 2012 Quest Software, Inc.

ALL RIGHTS RESERVED.

This document contains proprietary information protected by copyright. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose without the written permission of Quest Software, Inc. ("Quest").

The information in this document is provided in connection with Quest 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 products. EXCEPT AS SET FORTH IN QUEST'S TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST 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 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 HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest 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 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 World Headquarters

LEGAL Dept
5 Polaris Way
Aliso Viejo, CA 92656
www.quest.com
email: legal@quest.com

Refer to our Web site for regional and international office information.

Trademarks

Quest, Quest Software, the Quest Software logo, AccessManager, ActiveRoles, Aelita, Akonix, AppAssure, Benchmark Factory, Big Brother, BridgeAccess, BridgeAutoEscalate, BridgeSearch, BridgeTrak, BusinessInsight, ChangeAuditor, ChangeManager, Defender, DeployDirector, Desktop Authority, DirectoryAnalyzer, DirectoryTroubleshooter, DS Analyzer, DS Expert, Foglight, GPOAdmin, Help Desk Authority, Imceda, IntelliProfile, InTrust, Invirtus, iToken, IWatch, JClass, Jint, JProbe, LeccoTech, LiteSpeed, LiveReorg, LogAdmin, MessageStats, Monosphere, MultSess, NBSpool, NetBase, NetControl, Npulse, NetPro, PassGo, PerformaSure, Point,Click,Done!, PowerGUI, Quest Central, Quest vToolkit, Quest vWorkSpace, ReportAdmin, RestoreAdmin, ScriptLogic, Security Lifecycle Map, SelfServiceAdmin, SharePlex, Sitraka, SmartAlarm, Spotlight, SQL Navigator, SQL Watch, SQLab, Stat, StealthCollect, Storage Horizon, Tag and Follow, Toad, T.O.A.D., Toad World, vAutomator, vControl, vConverter, vFoglight, vOptimizer, vRanger, Vintela, Virtual DBA, VizionCore, Vizioncore vAutomation Suite, Vizioncore vBackup, Vizioncore vEssentials, Vizioncore vMigrator, Vizioncore vReplicator, WebDefender, Webthority, Xaffire, and XRT are trademarks and registered trademarks of Quest Software, Inc in the United States of America and other countries. Other trademarks and registered trademarks used in this guide are property of their respective owners.

Updated—MAY 2012

About Quest Software, Inc.

Quest Software (Nasdaq: QSFT) simplifies and reduces the cost of managing IT for more than 100,000 customers worldwide. Our innovative solutions make solving the toughest IT management problems easier, enabling customers to save time and money across physical, virtual and cloud environments. For more information about Quest solutions for application management, database management, Windows management, virtualization management and IT management, go to www.quest.com.

Contacting Quest Software

PHONE 800.306.9329 (United States and Canada)

If you are located outside North America, you can find your local office information on our Web site.

EMAIL sales@quest.com

MAIL Quest Software, Inc.
World Headquarters
5 Polaris Way
Aliso Viejo, CA 92656
USA

Contacting Quest Support

Quest Support is available to customers who have a trial version of a Quest product or who have purchased a commercial version and have a valid maintenance contract.

Quest Support provides around-the-clock coverage with SupportLink, our Web self-service.

Visit SupportLink at <https://support.quest.com>.

SupportLink gives users of Quest Software products the ability to:

- Search Quest's online Knowledgebase
- Download the latest releases, documentation and patches for Quest products
- Log support cases
- Manage existing support cases

View the Global Support Guide for a detailed explanation of support programs, online services, contact information and policies and procedures.

TBV-EffBUREcVMwareEnviron-US-VG