

# Best Practices for EMC VNX/VNXe and Data Domain with Veeam Availability Suite

**Rick Vanover**

vExpert, MCITP, VCP, Cisco Champion

**Neal Condon**

VMCE, MCSE

**AVAILABILITY**  
for the Always-On Enterprise™

# Contents

|                                                                                    |           |
|------------------------------------------------------------------------------------|-----------|
| <b>Executive overview .....</b>                                                    | <b>3</b>  |
| <b>Introduction.....</b>                                                           | <b>3</b>  |
| <b>It's all about Availability.....</b>                                            | <b>3</b>  |
| <b>Chapter 1 – VNX/VNXe connection .....</b>                                       | <b>4</b>  |
| Connecting Veeam Availability Suite to a VNX or VNXe array .....                   | 4         |
| VNX and VNXe protection schedules .....                                            | 8         |
| Using Veeam Explorer for Storage Snapshots .....                                   | 9         |
| Using Veeam Backup from Storage Snapshots.....                                     | 13        |
| Considerations for the first full backup .....                                     | 15        |
| Considerations for incremental backups.....                                        | 15        |
| Considerations for replication jobs .....                                          | 15        |
| Availability for the Always-On Enterprise made easy with storage integration ..... | 15        |
| <b>Chapter 2 – Data Domain .....</b>                                               | <b>16</b> |
| Data Domain basics.....                                                            | 16        |
| Data Domain protocols .....                                                        | 16        |
| Data Domain data paths .....                                                       | 17        |
| Data Domain administration interfaces.....                                         | 17        |
| CIFS protocol .....                                                                | 17        |
| NFS protocol.....                                                                  | 18        |
| DD Boost protocol .....                                                            | 18        |
| VTL protocol.....                                                                  | 18        |
| Configuring security and firewalls (NFS and CIFS access).....                      | 18        |
| Connecting Veeam Availability Suite to a Data Domain storage appliance .....       | 19        |
| Configuring Data Domain storage without the boost licensing .....                  | 23        |
| Considerations for the first full backup .....                                     | 24        |
| Considerations for incremental backups.....                                        | 24        |
| Considerations for Backup Copy jobs.....                                           | 25        |
| Considerations for Instant Recovery .....                                          | 25        |
| File-level restores .....                                                          | 26        |
| <b>Author information:.....</b>                                                    | <b>27</b> |
| <b>About Veeam Software .....</b>                                                  | <b>27</b> |

## Executive overview

Veeam® Availability Suite™ v9 introduces integration with the EMC VNX and VNXe series of hybrid flash storage. This integration provides data centers with new ways to enhance Availability by unleashing the power of storage snapshots for backup, replication and recovery tasks. This white paper highlights the features and setup of this integration.

In addition to the integration with the EMC VNX and VNXe storage families, Veeam Availability Suite v9 delivers enhanced integration with EMC Data Domain storage devices. This integration has been available since the v8 release and covers the new features in v9.

## Introduction

When it comes to avoiding data loss, one of the best ways to meet this objective is to leverage storage systems to keep Availability levels high. Veeam introduced **Backup from Storage Snapshots** and **Veeam Explorer™ for Storage Snapshots** to address this need. The **EMC VNX and VNXe series of hybrid flash storage** are the latest storage arrays added to this feature. In this document, we'll address what problem is solved by this integration, as well as how to connect a storage array to Veeam Availability Suite and what to look for to see immediate results.

With Veeam's integration with Data Domain Boost and the new enhancements in Veeam Availability Suite v9, you will see improved performance. In this white paper, you'll learn what makes this type of repository different, how to connect a deduplicating storage appliance to Veeam Availability Suite and how to configure for the optimal performance.

## It's all about Availability

When you think about VMware vSphere virtual machines (VMs) running in a data center, you want to keep them available in every sense, including when something doesn't go as expected and in the course of normal operations.

When something doesn't go as expected, you have an opportunity with Veeam Explorer for Storage Snapshots' to provide high-speed recovery techniques directly from the storage array.

In the course of normal operations, you must meet the challenge of keeping service levels high on running VMs and while you are backing up VMs. Specifically for VMware VMs, a widely used framework is the vSphere APIs for Data Protection (VADP). VADP prescribes a sequence of events to read data from the running VM for a specific use case, such as backup jobs or replicated VMs. However, during this time, the mechanism that coordinates writes to the VM has to manage coordination of these changes when the task completes. Specifically, the VM snapshot process will cause a phenomenon called *stun* to coordinate the merging of the writes that were preserved elsewhere during the VADP workflow.

Because of this, IT admins take backups at off-hours when the storage systems aren't under as much stress and the impact of snapshot removal stun isn't as significant.

With Backup from Storage Snapshots, the workflow is adjusted in a patented way to reduce the load on the VNX and VNXe family of hybrid flash arrays to allow admins to take backups with no limitations.

Veeam strives to continuously deliver new improvements and capabilities, and our partnership with EMC makes this a reality. Together, we provide fast backups and restores from VNX, VNXe and Data Domain storage for our joint customers.

## Chapter 1 – VNX/VNXe connection

### Connecting Veeam Availability Suite to a VNX or VNXe array

The process of connecting an array to Veeam Backup & Replication™ (part of Veeam Availability Suite) is very easy, but a few things need to happen on both the storage array through the EMC Unisphere interface and in Veeam Backup & Replication. For the most current list of requirements for VNX and VNXe connectivity, consult the [Help Center technical documentation](#).

Logically, the connectivity to the VNX or VNXe is done through one or more Veeam proxies (data movers) that have access to the storage provided on the array. The sequence of steps is as shown:

1. Veeam will analyze which VMs in the job have disks on a configured VNX or VNXe storage array before triggering a snapshot of the storage volume once all VM snapshots have been created
2. Then, Veeam will trigger a vSphere snapshot for all VMs located on the same storage volume (as a part of a vSphere snapshot, Veeam's application-aware processing of each VM is performed normally)
3. The proxy then retrieves the changed block tracking (CBT) information for the VM snapshots created on step 2
4. Next, Veeam will immediately trigger the removal of the vSphere snapshots on the production VMs
5. The Veeam proxy will then mount the storage snapshot to one of the backup proxies connected to the storage fabric
6. The Veeam proxy will then read new and changed virtual disk data blocks directly from the storage snapshot and transport them to the backup repository or replica VM
7. Finally, Veeam will trigger the removal of the storage snapshot once all VMs are backed up

You can see this process below:




Figure 1: The sequenced flow of Backup from Storage Snapshots

You may be wondering how you can set up this type of data mover so that it doesn't interfere with the production vSphere storage resources. We recommend the using simplified steps below (which the rest of this paper will explain) in the following order:

- Install Veeam Backup & Replication
- Determine if you require additional vSphere backup proxies and create and add them to the Veeam Backup & Replication console if necessary
- Connect the proxies to the storage network (iSCSI, Fibre Channel or NFS)
- Add the proxies to Unisphere and configure in Veeam Backup & Replication

Let's walk through an example where a VNXe running VMs on an NFS share will be connected to a Veeam proxy and Unisphere to provide the backup and restore capabilities from the storage snapshots. The Veeam installation is already in place, with regular backups jobs that are not using the Backup from Storage Snapshots technology. First (only for fibre channel proxies), you'll have to add an entry in the hosts section of Unisphere, as shown below:




Figure 2: The hosts section of the EMC Unisphere interface provides access to Veeam proxies

When creating a host, you'll add the selected Veeam proxy through the **Create Host** option in Unisphere. iSCSI connectivity will use IQN and TCP/IP addresses. You should configure the host here in a similar fashion to how the VMware ESXi hosts' connectivity is assigned. For Fibre Channel connectivity, the zoning of the Veeam proxy will be the same as that of an ESXi host. This includes with VMFS datastores are intended for use with Veeam's storage integration with the VNX and VNXe arrays.

The figure below shows how a Veeam proxy (which is also the Veeam Backup & Replication console) is connected to Unisphere:




Figure 3: The host entry connects the Veeam proxy to Unisphere

For all storage protocols, the basic rule is to assign connectivity to the Veeam proxy the same way that you would assign a VMware ESXi host to the storage resource. This is because the Veeam proxy will move the data in the course of a backup or replication job directly from the storage snapshot, which requires storage-level access.

**The Veeam proxy will move the data on behalf of the backup or replication job from the storage snapshot. Aside from the note above about generally giving access similar to that of an ESXi host, the backup infrastructure must also have the following characteristics:**

1. **Add a properly configured VMware backup proxy to the backup infrastructure**
2. **Add the vCenter Server/ESX(i) host with VMs whose disks are located on the storage system to Veeam Backup & Replication**
3. **Add the storage system to Veeam Backup & Replication**

**You can read more specific requirements at [helpcenter.veeam.com](http://helpcenter.veeam.com).**

Once that is complete, it is time to add the VNXe to the Veeam Backup & Replication console. You perform this task in the storage infrastructure section of the user interface. This section

of the user interface is shown below:




Figure 4: Connected arrays will be shown in the storage infrastructure section

To add a storage array to the storage infrastructure associated with this Veeam Backup & Replication console, select which type of array to add (VNXe or VNX as well as block or file storage):




Figure 5: Adding VNX storage as block or file-level storage

The wizard will only inquire for the Unisphere management IP address as well as a username and password. Then, the array is added into Veeam Backup & Replication, and you can use it for

backups, replicas and restore tasks. The storage infrastructure browser will then show the array and storage volumes as well as any snapshots (if present):




Figure 6: A support array is shown with three snapshots on the array. Note the VMs for the volume are shown on the right

## VNX and VNXe protection schedules

You can set a protection schedule in Unisphere to create storage snapshots automatically on the VNX or VNXe. This is incredibly beneficial for the Veeam Explorer for Storage Snapshots use case (explained in the subsequent section) because a storage snapshot is taken automatically. Additionally, many storage administrators set storage snapshots up but don't really know a practical way to use them.

You set the protection schedule in Unisphere. The figure below shows an example of a storage snapshot taken every four hours, with the snapshots kept for one day (this isn't a default configuration, but a great way to get more specific restore points throughout the day):





Figure 7: A snapshot schedule will work great with Veeam Explorer for Storage Snapshots

After you've set the schedule on the VNX or VNXe, you can see its effect in the Veeam Backup & Replication interface. This is where you can use Veeam Explorer for Storage Snapshots. The figure

below shows the prescribed storage snapshots kept on the array and are ready for recovery only steps:



*Figure 8: Veeam Backup & Replication reads the snapshots on the VNX or VNXe array*

Note that each storage array snapshot contains the VMs that were present on the array at the time of the snapshot (every four hours in this example). Additionally, note that storage snapshots generated by the storage array have a different nomenclature, generally following a timestamp. The example selected above is named: 2015-11-30\_12:00.00 and corresponds to Nov. 30, 2015, at 12 p.m.

Also, note that based on the settings on the VNX or VNXe, only a certain number of snapshots are kept on the array. From the example below, snapshots are taken every four hours and held for up to 24 hours, so there should be six snapshots present by schedule.

## Using Veeam Explorer for Storage Snapshots

With the array added and snapshots visible in the Veeam Backup & Replication interface, you can perform recovery actions. You should only use this recovery task for VM recovery situations in which you can confirm the integrity of the storage array. For example, if a power failure that turns off the array and the VMs don't come back up correctly, it may not be a good use case for using Veeam Explorer for Storage Snapshots. A better one is if a VM had a script run against it with an option set incorrectly that deleted a lot of data, instead of exporting a lot of data. You can leverage Veeam Explorer for Storage Snapshots and take a backup to a different storage (which this paper explains later on). Make sure to consult the User Guide for more information.

Veeam Explorer for Storage Snapshots provides three categories of restore options, with a total of eight specific options. The three categories are:

- Instant VM Recovery: Boot the VM up from the storage snapshot
- File-level recovery: Provide guest file recovery from both Windows and Linux VMs
- Application recovery: Leverage the Veeam Explorers for SQL Server, SharePoint, Active Directory, Exchange and Oracle

The image below shows the wizard that launches Veeam Explorer for Storage Snapshots:




Figure 9: Veeam Explorer for Storage Snapshots starts here

You need to consider a few things before using Veeam Explorer for Storage Snapshots for a restore. First, you have to determine which host you will leverage to perform the restore as a customization for subsequent tasks. Specifically, you have to select a host, resource pool and VM folder. When this customization is complete, the restore wizard for Veeam Explorer will populate these attributes:




Figure 10: Veeam Explorer for Storage Snapshots settings are customizable

All of the restore tasks will have the storage snapshot mounted to the specified host to perform file-level recovery, application explorers or entire VM recovery. The session log below shows this process:




Figure 11: Mounting the storage snapshot to the host is done transparently

This task presents the storage snapshot to the ESXi host so you can complete the recovery task. It's important to note that the wizard will mount the storage snapshot to the host with parts of the same name as the production datastore. In the example above, the NFS datastore mounted (which is the storage snapshot) is named Veeam\_v9DemoVNX. The actual datastore is named v9DemoVNX. Also, the VM is registered to the host to allow you to view the file system, and it has an extra string appended to the name of it so it does not interfere with the production VM. Fibre channel and iSCSI datastores are preceded with the characters "snap-" before the original datastore name.

After this step, the rest of the process is similar the restore wizard, which will leverage a backup file. The only difference is that the source is the storage snapshot (and the associated steps to mount it to the host). In the end, the file-level recovery wizard is the same user interface that you would see when restoring from a backup file.

Likewise, the wizard for the Instant VM Recovery option from storage snapshots will present a similar set of options to get a VM to the Instant VM Recovery task when running from a backup. The wizard will allow you to get a VM running as quickly as possible. The two steps of the wizard are shown below:




Figure 12: Restoring VMs from storage snapshots takes only a few clicks

Much like the file-level recovery, the storage snapshot mounts directly to the ESXi host by the recovery wizard and the VM inventoried. This presents the VM to the datastore name with Veeam appended to the actual datastore name and the one VM listed as shown below.




Figure 13: In the vSphere Web Client, recovered VMs are run from the storage snapshot

It's important to note that the VM can't stay here forever. We recommend that you migrate the VM from the storage snapshot back to a production datastore via either Storage vMotion or Veeam's Quick Migration, both of which can move the storage associated with the VM from the storage snapshot to a production datastore.

**Dismiss the Myth!****Why do a backup if we have all these recovery capabilities from the storage snapshot?**

*It is natural to think that this capability with Veeam Explorer for Storage Snapshots will be enough to satisfy the Availability requirements of your organization. In reality, this is not the case. You will need to ensure your data is recoverable in the event of a complete failure of that storage system. In that situation, the snapshots and the running VMs are not accessible.*

*At Veeam, we promote the 3-2-1 Rule, which is a great mindset to use in this context. The 3-2-1 Rule states that you should have three different copies of your data, on two different media, with one of them stored off site.*

*This is a great mindset as it doesn't lock you into any specific technology and it can address nearly any failure scenario.*

**Using Veeam Backup from Storage Snapshots**

While the storage snapshot is a great recovery technique, a backup to different storage is also needed for the highest levels of Availability. One way to avoid data loss is to leverage the storage snapshot as part of the backup process. Veeam Availability Suite v9 includes the ability to perform backups from the **VNX and VNXe family of hybrid flash arrays**.

Before we highlight the specific functionality of this aspect of the storage integration, it is important to reiterate the problem we are solving, as described in the beginning of the paper. For VMware environments, taking backups can incur a lot of stress on storage. With the storage integration, the impact VADP provides can be minimized. Specifically, it significantly reduces the amount of time that a VMware snapshot is open. By having the VM snapshot open only briefly, the coordination associated with injecting the written blocks is minimized. This allows you to take a backup at virtually any time, even with the busiest systems.

The best part about using the storage snapshot for the backup engine is that it is enabled by default. The figure below shows this setting as part of a backup or replication job:





Figure 14: Storage integration is part of the backup job and is enabled by default

The steps that added the VNX or VNXe to the Veeam Backup & Replication console allow this function to work for Enterprise Plus editions. When a backup or replication job runs from a supported array, the big thing to look for is that the VM snapshot is already removed before the data is copied. This is shown in detail in the figure below:



*Figure 15: Data is transferred after the VMware snapshot has been removed*

Let's explain the importance of the seven selected lines above.

- The first selection is the creation of the VMware snapshot on the VM. This includes any application-aware processing, which will help in subsequent steps and recovery.
- The second line is where the storage snapshot is called on the VNX or VNXe storage array, namely, the collecting disk files location data step.
- The third selection is where the VMware snapshot is removed.
- The fourth through the seventh selections are the task of moving data directly from the storage snapshot for all virtual disks. This is all done after the VMware snapshot has been removed and is application-consistent.

While this VM isn't very large and not all environments are created equal, the notable thing here is that the VMware snapshot was only open for five seconds. There are a number of considerations involved with this technology, which we will outline in subsequent sections.

## Considerations for the first full backup

The first full backup will benefit greatly from leveraging the storage snapshot of the VNX or VNXe array. This significantly reduces, the amount of time the VMware snapshot would be open, especially for larger and busier VMs. Consider a multi-terabyte VM: It may take days to do an active full backup of this type of system. If a VMware snapshot is open for a long amount of time—say 36, 48 or more hours—it can degrade the performance of a running VM significantly when the writes are coordinated.

By leveraging Backup from Storage Snapshots, the VMware snapshot is open for just seconds. All of the data movement from the storage snapshot happens when the VMware snapshot is already removed but the VM is properly prepared for the backup.

## Considerations for incremental backups

The incremental backup also benefits from the storage integration with the VNX or VNXe array, particularly when compared to other solutions that move data from storage snapshots. With Veeam Backup & Replication, the specific point that aids incremental backups is the preservation of VMware's Changed Block Tracking (CBT) for incremental backups.

This will avoid the phenomenon in which the incremental run of a backup from a storage snapshot would have to read the whole disk geometry to find the blocks that have changed. This is called snap and scan, and it slows down incremental backups. Veeam Backup & Replication preserves the CBT data for incremental backups to increase speeds up to 20x faster than competitive products.

## Considerations for replication jobs

Replication jobs will benefit from the storage integration with the VNX or VNXe array as well. This is due to the fact that in most situations, the target of a replication job is a remote data center. Even with the Veeam Built-In WAN Acceleration technology, reducing the time the VMware snapshot is open will significantly help improve service levels on the VM being replicated.

Using the storage integration with the VNX or VNXe arrays in conjunction with Veeam's built-in WAN acceleration technology, building replication jobs off site happens minimal stress on the production VM. This integration is especially beneficial when the option to source a replica from a backup repository (which came in v8) is not used.

## Availability for the Always-On Enterprise made easy with storage integration

In this paper, we have toured the usability of VNX and VNXe storage integration with Veeam Availability Suite. The goal of this integration is to enable high-speed recovery through Veeam Explorer for Storage Snapshots and avoid data loss by leveraging Backup from Storage Snapshots technology. The integration with Veeam Availability Suite v9 and the VNX/VNXe arrays easily deliver these benefits.

By unleashing the power of storage snapshots from the VNX and VNXe arrays, new options emerge for backups, replication jobs and recovery tasks.

## Chapter 2 – Data Domain

The process of connecting the **Data Domain storage** to Veeam Backup & Replication—part of Veeam Availability Suite—is very easy. You just have to make sure a few things happen on both the Data Domain storage device via the Data Domain System Manager and in Veeam Backup & Replication. For the most current list of requirements for Data Domain connectivity, review the [Help Center technical documentation](#).

### Data Domain basics

In the next few sections, we will review the protocols, data paths and general configurations of the Data Domain storage device.

### Data Domain protocols

NFS — Network file system (NFS) clients can have access to Data Domain system directories and Mtrees

CIFS — Common Internet Filesystem (CIFS) clients can have access to Data Domain system directories and Mtrees

VTL — Virtual tape library (VTL) protocol enables backup applications to connect to and manage Data Domain system storage, as if it were a tape library

All of the functionality generally a physical tape library supports is available with a Data Domain system configured as a VTL

Backup software (not the Data Domain system) manages the movement of data from a system configured as a VTL to a physical tape

The VTL protocol is used with Fibre Channel (FC) networking

DD Boost — The DD Boost protocol enables backup servers to communicate with storage systems without the need for Data Domain systems to emulate tape

There are two components to DD Boost:

One component that runs on the backup server

One component that runs on a Data Domain system

NDMP — If the VTL communication between a backup server and a Data Domain system is through NDMP, FC is not required.

When you use NDMP, all initiator and port functionalities do not apply

## Data Domain data paths

Data Domain data paths over FC networks – DD Boost, VTL

Data Domain data paths over Ethernet networks – NFS, CIFS and DD Boost

## Data Domain administration interfaces

The Data Domain administration interfaces include:

The Data Domain System Manager, which is the graphical user interface (GUI)

The command line interface (CLI) – Access CLI via SSH, serial console, telnet, keyboard and monitor

Data Domain Initial Configuration using Data Domain System Manager – Configuration Wizard:  
(Command Line – “config setup” Command)

The configuration wizard consists of these sections:

1. Licenses
2. Network
3. File system
4. System
5. CIFS
6. NFS

## CIFS protocol

### Enter the following:

1. Authentication: Workgroup: Enter the CIFS server name, if not using the default
2. Active Directory: Enter the full realm name for the system as well as a domain joining credential user name and password. You also have the option to enter the organizational unit name, if you're not using the default
3. Share: Enter the share name and directory path. Enter the client name, if you're not using the default
4. Summary: Review the summary carefully

## NFS protocol

### Enter the following:

1. Export: Enter a path name for the export. Enter the NFS client server name you want to add to / backup, if you're not using an existing client. Then, select NFS options for the client. These clients receive the default permission — which are read and write permissions with root squashing turned off and mapping of all user requests to the anonymous UID/GID turned off, and it is secure
2. Summary: Review the summary carefully

## DD Boost protocol

### Enter the following:

1. Storage Unit: You have the option change the Storage Unit name. Either select an existing user or create a new user by entering a user name, password and minimum management role, which can be:  
  
Backup (backup-operator): In addition to user privileges, this lets you create snapshots, import and export tapes to a VTL and move tapes within a VTL.  
  
None (none): This is intended only for EMC DD Boost authentication, so you cannot monitor or configure a Data Domain system.  
  
security (security): In addition to user privileges, this lets you set up security-officer configurations and manage other security-officer operators.  
  
sysadmin (admin): This role lets you configure and monitor the entire Data Domain system.  
  
user (user): This lets you monitor Data Domain systems and perform the fast copy operation.
2. Fibre Channel: If DD Boost is to be supported over FC, select the option to configure it. Enter a unique name for the Access Group. (Duplicate access groups are not supported.) Select one or more initiators. You also have the option to replace the initiator by entering a new one. The devices to be used are listed.

## VTL protocol

### Enter as follows:

1. Library: Enter the library name, number of drives, drive model, number of slots and CAPs, changer model name, starting barcode, and, optionally, tape capacity.
2. Access Group: Enter a unique name for the Access Group. (Duplicate access groups are not supported.) Select one or more initiators. You also have the option to replace the initiator name by entering a new one. The devices (drives and changer) to be used are listed.
3. Summary: Review the summary carefully.

## Configuring security and firewalls (NFS and CIFS access)

You should configure the firewall so that only required and trusted clients have access to the Data Domain system.

- By default, anonymous users from known CIFS clients have access to the Data Domain system
- For security purposes, change this option from disabled (the default) to enabled: # cifs option set restrict-anonymous enabled

The following tables show the TCP and UDP ports the Data Domain system uses for inbound and outbound traffic and the services that make use of them.

Please refer to your administrative guide for ports used by Data Domain systems for inbound traffic.

| DD OS SW version              | DD OS SW version | 5.0    | 5.1    | 5.2    | 5.4    | 5.5    | 5.6    | 5.7    |
|-------------------------------|------------------|--------|--------|--------|--------|--------|--------|--------|
| <b>Boost</b>                  |                  | 2.3    | 2.4    | 2.5    | 2.6    | 3.0    | 3.1    | 3.2    |
| <b>SW DA Date</b>             |                  | Dec-10 | Now-11 | May-12 | Jul-13 | May-14 | Apr-15 | Now-15 |
| <b>SW release TSSR notify</b> |                  | Jul-15 | Jul-15 | Jul-15 | Jul-15 | May-16 | Apr-17 | Now-17 |
| <b>SW release TSSR</b>        |                  | Jul-16 | Jul-16 | Jul-16 | Jul-16 | May-17 | Apr-18 | Now-18 |

Table 1: Data Domain OS & Boost matrix

Veeam supports the following Data Domain Boost plug-ins and Data Domain OS versions:

| DD Boost plug-in version | DD OS accompanying version | Veeam supported version |
|--------------------------|----------------------------|-------------------------|
| 2.6                      | 5.2-5.5                    | 8                       |
| 3                        | 5.4-5.6                    | 9                       |

Table 2: Veeam version support matrix

## Connecting Veeam Availability Suite to a Data Domain storage appliance

The process of connecting the deduplicating storage appliance to Veeam Backup & Replication (part of Veeam Availability Suite) is simple: you just need to make sure a few things happen on both the storage through the EMC Data Domain System Manager interface and in Veeam Backup & Replication.

Make sure the following steps are completed prior to creating your first backup repository:

1. Install Veeam Backup & Replication
2. Determine if additional VMware backup proxies are required and create and add them to the Veeam Backup & Replication console, if necessary
3. Connect the proxies to the storage network (iSCSI, Fibre Channel or NFS)
4. Configure Data Domain system manager
5. Configure Veeam Backup & Replication

For the complete administrative guide, please visit our website: [vee.am/documentation](http://vee.am/documentation).

***Due to the way Data Domain actively deduplicates, stores and rehydrated block operations within Veeam Backup & Replication may be limited by the overhead that these activities require. This can affect operations designed for production storage and are highly transactional or require random I/O by their nature.***

***Operations that utilize the Data Domain as a replacement for production storage will be limited by the speed at which the Data Domain can actively deduplicate inbound blocks.***

***If DDBoost is not licensed then the Data Domain appliance must be added as a repository under conventional methods. It is advised to use a Linux server with the volume mounted via NFS as a relay server to help improve performance. Using SMB share will provide acceptable, but lower performance.***

A DD Boost storage unit is the local object that will become a target for the Veeam backup jobs. Within the Data Domain System Manager console, navigate to: Data Management | DD Boost | Storage Units and click Create, as shown below:




Figure 16: Creating a storage device using the Data Domain System Manager

Take note of the newly created storage unit name provisioned because it is required during the initial Veeam repository setup.

The connectivity to the Data Domain is done through one or more Veeam proxies (data movers) that have access to the storage provided. The following steps will walk you through the process:

1. Within the Veeam Backup & Replication server console, locate backup repositories. You can choose the upper left blue highlighted section of the GUI to add repository, or you can right click on the right pane window and select the same there.

2. Name your device.

The screenshot shows the 'New Backup Repository' wizard at the 'Name' step. The left sidebar has 'Name' selected. The main area has a 'Name' field with 'Data Domain' and a 'Description' field with 'Created by HPDEMO\veeamsupport at 12/9/2015 2:58 PM.' The 'Next >' button is highlighted.

3. Select **Deduplicating storage appliance**.

The screenshot shows the 'New Backup Repository' wizard at the 'Type' step. The left sidebar has 'Type' selected. The main area shows four radio button options: 'Microsoft Windows server', 'Linux server', 'Shared folder', and 'Deduplicating storage appliance'. The 'Deduplicating storage appliance' option is selected and highlighted. The 'Next >' button is highlighted.


4. Choose **EMC Data Domain** as your deduplicating storage.

The screenshot shows the 'New Backup Repository' wizard at the 'Deduplicating Storage' step. The left sidebar has 'Deduplicating Storage' selected. The main area shows three radio button options: 'EMC Data Domain', 'ExaGrid', and 'HP StoreOnce'. The 'EMC Data Domain' option is selected and highlighted. The 'Next >' button is highlighted.

5. Now enter the Data Domain server name, check the box if you're connected via FC and enter the credentials that have administrative rights to the target storage.


**Note:** Use the special naming convention for your Fibre channel connection.

The system name is the DNS name or TCP/IP address of the Data Domain server that will house the DD Boost Storage Unit connected in the previous steps. The communication that leverages EMC Data Domain Boost can go over the LAN or FC as indicated earlier. If FC connectivity is to be used, select the option to leverage the FC. This requires that the FC be correctly zoned for the Veeam gateway server to access the storage targets.



6. Additionally, the credentials set in the configuration phase of the DD Boost Storage Unit are specified here in the backup repository wizard.

**Note:** The account must have Data Domain Boost privileges, not just administrative rights.



You also have the choice to select your gateway server specifically or leave it set to automatic selection. A critical decision in this wizard is the selection of a gateway server. You should place the gateway server in all situations from a network perspective as close as possible (from a latency perspective) to the actual Data Domain server in all situations. This is because the gateway server will send the traffic to the Data Domain server and communicate with the DD Boost Storage Unit. Backup (or backup copy job) data flow may originate in other sites or networks, but the gateway server role for the Data Domain server should be as close as possible to the appliance.

## Configuring Data Domain storage without the boost licensing

If your Data Domain is not licensed for DDBoost, you have two primary options for connecting to the appliance.

### CIFS

1. Launch the creation of a new Repository. On the Type tab, select CIFS.  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_launch.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_launch.html)  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_type.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_type.html)
2. On the next tab, configure the path to which the Repository will write and set credentials to access that share.  
For more information, see the second section named **Shared Folder** on this page:  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_server.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_server.html)
3. On the Repository tab in the advanced section, enable **Decompress backup data blocks before storing**.
4. Unless your environment requires you to specify a different vPower NFS server, you can use the default settings for the last steps in the repository configuration.

### NFS

You will need to configure the Data Domain for NFS access and a Linux server to mount the volumes from the Data Domain via NFS. Please refer to the following links for further information regarding connecting Linux to the Data Domain via NFS:

<http://forums.veeam.com/veeam-backup-replication-f2/veeam-datadomain-and-linux-nfs-share-t8916.html>  
<http://tsmith.co/2014/veeam-and-datadomain>

1. Launch the creation of a new repository. On the Type tab, select Linux.  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_launch.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_launch.html)  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_type.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_type.html)
2. On the next tab, select the Linux server that you need to connect to. If it is not present in the list, select **Add New** ....
3. On the Repository tab, specify the path on the Linux server that leads to where you mounted the Data Domain via NFS. In the advanced section on this tab, enable **Decompress backup data blocks before storing**.  
[http://helpcenter.veeam.com/backup/80/vsphere/repository\\_repository.html](http://helpcenter.veeam.com/backup/80/vsphere/repository_repository.html)
4. Unless your environment requires you to specify a different vPower NFS server, you can use the default settings for the last steps in the repository configuration.

## Considerations for the first full backup

When data is sent to the Data Domain storage, it uses source-side deduplication, depending upon the Distributed Segment processing option enabled in the Data Domain server. Blocks from the virtual disks of VMs that already reside on the device will not be transferred. A great visual representation of this can be seen in our backup windows statistics:

- This is important for a number of reasons. Most importantly, this is where you will see the bottleneck analysis. When Data Domain Boost is enabled, we see that the target is likely not the bottleneck.

**Note: this is not the case when running on high performance systems.**

- Veeam built-in deduplication and compression may be happening by default, but don't worry — when we set up a Data Domain Boost enabled repository, we will ensure that the data is decompressed as it is written to the deduplicated target.
  - In the data section, the read count is a metric of the source data mover. Here, the transferred count is a metric of data transfer between source and target data movers. Data Domain Boost kicks in between target data mover and the Data Domain system. Since it directly affects “Target” bottleneck detector it will indirectly affect the Transferred counter. However, you’ll see backup times are quicker with Data Domain Boost.
  - We have a number of different lab environments here at Veeam, but I noticed that this environment started performing better with Data Domain Boost as I used it more. Why? Because more data has been ingested. There is now more source data that is deduplicated, and as more deduplication hits come in, they are deduplicated before they are even transferred.
  - Whether there are VMs that are in more than one backup job or if VMs are deployed from a template, similar blocks will benefit from Data Domain Boost over time. Either way, this feature will help on ingest to the Data Domain over time, making backups perform better.
  - Get up to 10x faster backup performance with the new per-VM backup file chains option, which enables multiple write streams by leveraging parallel VM processing in Veeam Availability Suite v9.

## Considerations for incremental backups

When we talk about backup jobs specifically, you should note that reverse incremental backups are not supported.

Use of the forever incremental backup method is only advisable in situations where you have less than 30 restore points. If you plan to have more restore points, you should enable the Active Full option and configure it to create a new active full monthly. This will reduce restore times by segmenting each month into its own backup chain.

## Considerations for Backup Copy jobs

You can create a traditional Veeam replication job or ideally leverage the Veeam backup copy job, because it provides more granular control than just copying the entire backup files. For example, you can choose individual VMs to copy (not the entire backup files), do many-jobs-to-one-backup-file type of copy and so on. It is also storage agnostic, so you can use it to get data to and from Data Domain easily.

Now of course, we also have customers who are using native Data Domain replication to copy the entire backup files as they are, but we do not control this native replication in our upcoming product release.

You can improve local backup copy performance and reduce the load on deduplication appliances by eliminating data rehydration with new support for active full backups with backup copy jobs in Veeam Availability Suite v9.

## Considerations for Instant Recovery

The aforementioned limitations of a Data Domain appliance can adversely affect Instant Recovery, and, depending on the type of VM being restored, highly transactional VMs will require more IOPS from the Data Domain than others. With this in mind, you can expect to only be capable of running only from one to a few Instant Recoveries simultaneously, depending on Data Domain mode. Instantly recovered VMs started from a backup file stored on a Data Domain may react or start slowly because the majority of their read operations will be hindered by the limited rehydration performance of the Data Domain.

For vSphere users, it is highly advised that the user select to have virtual disk updates redirected to a high-performance datastore when performing an Instant Recovery. This will assist in improving some performance.

We also advise advised that VMware users start the migration steps as soon as possible if they intend to make the VM permanent.

Performing instant VM recoveries will consume Data Domain system resources, which may affect performance of other processes (i.e. ingest, replication, cleaning, etc.).

For better Instant Recovery performance, we recommend you use the smallest block size (under the storage optimization setting):

[http://helpcenter.veeam.com/backup/80/vsphere/backup\\_job\\_advanced\\_storage\\_vm.html](http://helpcenter.veeam.com/backup/80/vsphere/backup_job_advanced_storage_vm.html).

- Please note that for a backup job targeted at Data Domain, we recommend that you have the largest block size for better performance.
- For better Instant Recovery performance, we recommend that you have fewer VMs in the job.
- The longer an incremental chain is the more read operations are performed when restoring from an incremental restore point.

## File-level restores

The backup browser may take longer than usual to open if an increment is selected. Furthermore, that increment's distance from the full restore point can also cause a delay when opening the backup browser.

Navigating between folders within the backup browser may take additional time because each folder's content must be retrieved from the backup file in order to display it.

For more information regarding backup job settings, please refer to <http://www.veeam.com/kb1745>.

In conclusion, the new integration with EMC Data Domain Boost in Veeam Availability Suite v9 enables source-side deduplication and data in-flight encryption over the WAN for faster, more secure backups to off-site EMC Data Domain appliances.

## Author information:



**Rick Vanover** (vExpert, MCITP, VCP, Cisco Champion) is a senior product strategy manager for **Veeam Software** based in Columbus, Ohio. Rick is a popular blogger, podcaster and active member of the virtualization community. Rick's IT experience includes system administration and IT management; with virtualization being the central theme of his career recently. Follow Rick on Twitter [@RickVanover](#) or [@Veeam](#).



**Neal Condon** (VMCE, MCSE) is a Solutions Architect for **Veeam Software** based in Sturbridge, Massachusetts. He has spent more than 27 years in the information technology field. Follow Neal on Twitter [@NealCondon](#) or [@Veeam](#).

## About Veeam Software

**Veeam**<sup>®</sup> recognizes the new challenges companies across the globe face in enabling the Always-On Business<sup>™</sup>, a business that must operate 24/7/365. To address this, Veeam has pioneered a new market of *Availability for the Always-On Enterprise*<sup>™</sup> by helping organizations meet SLAs for recovery time and point objectives (RTPO<sup>™</sup>) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified recoverability, leveraged data and complete visibility. **Veeam Availability Suite**<sup>™</sup>, which includes **Veeam Backup & Replication**<sup>™</sup>, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 37,000 ProPartners and more than 183,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.




**AVAILABILITY**  
for the Always-On Enterprise™

---

# **NEW Veeam®** **Availability** **Suite™ v9**

**RTPO™ <15 minutes for**  
**ALL applications and data**

---



**Learn more and preview**  
**the upcoming v9 release**

**[vee.am/v9](https://vee.am/v9)**