



Technical Report

## E-Series and EF-Series Reference Architecture and Storage Best Practices with Veeam Backup & Replication 9.5

Akash Gupta, NetApp; Shawn Lieu (Americas), Stefan Renner (EMEA), and Michael Cade (Performance), Veeam

February 2018 | TR-4471

In partnership with

**VEEAM**

### **Implementation Overview and Usage Considerations**

This document outlines the reference architecture and best practices when using NetApp® E-Series storage in a Veeam Backup & Replication 9.5 environment.

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary</b>	<b>5</b>
1.1	Introduction	6
1.2	About NetApp	6
1.3	About Veeam	7
<b>2</b>	<b>Reference Architecture Overview</b>	<b>7</b>
2.1	NetApp E-Series as Veeam Backup & Replication Repositories for Backup and Archiving	7
2.2	NetApp E-Series as a Veeam Backup & Replication Backup Repository and an Off-Site Cloud Repository	7
2.3	NetApp E-Series as a Veeam Backup & Replication Backup Repository for NetApp FAS Production Storage	8
2.4	Deployment Scenarios	9
<b>3</b>	<b>NetApp E-Series and NetApp EF-Series Arrays</b>	<b>12</b>
3.1	E-Series Hardware	12
3.2	Key NetApp E-Series and EF-Series Features	20
3.3	Enterprise Reliability and Availability	20
3.4	SANtricity OS 11.40: New Features	20
3.5	SANtricity System Manager vs. Storage Manager	21
<b>4</b>	<b>Veeam Backup &amp; Replication 9.5</b>	<b>22</b>
4.1	Overview	22
4.2	Solution Architecture	22
<b>5</b>	<b>Veeam Backup &amp; Replication 9.5 Requirements</b>	<b>24</b>
5.1	VMware System Requirements	24
5.2	Hyper-V System Requirements	27
5.3	Veeam Sizing Requirements	30
<b>6</b>	<b>Veeam Backup &amp; Replication 9.5 Design Considerations</b>	<b>31</b>
6.1	Direct SAN Access Mode	31
6.2	Data Restore in Direct SAN Access Mode	32
6.3	Off-Host Backup (Hyper-V)	33
6.4	Veeam Backup & Replication 9.5 Components	36
<b>7</b>	<b>NetApp E-Series Volume Configuration Guidelines</b>	<b>37</b>
7.1	NetApp E-Series Storage Configuration Guidelines for Backup & Replication 9.5 Backup Repositories	37
<b>8</b>	<b>NetApp E-Series and EF-Series Host Configuration Guidelines</b>	<b>38</b>
8.1	Host Connectivity	38
8.2	NetApp E-Series Storage Host Mapping Configuration for Direct SAN Access	40

<b>9 CPU and Memory Sizing for Veeam and NetApp E-Series</b> .....	<b>40</b>
9.1 CPU and Memory Sizing of Backup Server .....	40
9.2 CPU and Memory Sizing of Backup Proxy.....	41
9.3 CPU and Memory Sizing of Backup Repository.....	41
<b>10 Capacity Sizing for Veeam and NetApp E-Series</b> .....	<b>41</b>
10.1 Repository Sizing Tool.....	42
<b>11 Network Planning and Sizing for Veeam and NetApp E-Series</b> .....	<b>49</b>
11.1 Test Environment and Setup.....	50
11.2 Volume Configuration for Performance Testing.....	50
11.3 Backup Performance Results.....	51
11.4 Performance Best Practices.....	52
<b>12 NetApp SANtricity Plug-Ins for VMware</b> .....	<b>54</b>
12.1 NetApp SANtricity Plug-In for VMware vCenter .....	54
<b>Summary</b> .....	<b>55</b>
NetApp E-Series Documentation.....	55
Veeam Backup & Replication 9.5 Software Documentation .....	55

## LIST OF TABLES

Table 1) Veeam backup server specifications. ....	25
Table 2) VMware backup proxy server specifications. ....	26
Table 3) VMware backup repository server specifications. ....	27
Table 4) Veeam proxy server specifications (Hyper-V). ....	28
Table 5) Off-host backup proxy server specifications (Hyper-V). ....	29
Table 6) Backup repository server specifications (Hyper-V).....	30
Table 7) Veeam sizing requirements.....	30
Table 8) Veeam Backup & Replication 9.5 components and their description. ....	36
Table 9) Optional host I/O for E-Series.....	39

## LIST OF FIGURES

Figure 1) NetApp E-Series as Veeam Backup & Replication repositories for backup and archiving. ....	7
Figure 2) NetApp E-Series as a Veeam Backup & Replication backup repository and an off-site cloud repository. ....	8
Figure 3) NetApp E-Series as a Veeam Backup & Replication backup repository and off-premises backup repository for FAS production storage .....	8
Figure 4) Simple deployment.....	9
Figure 5) Advanced deployment.....	10
Figure 6) Distributed deployment. ....	11
Figure 7) E2800 hardware overview.....	13

Figure 8) New E5760 hybrid storage array with the front bezel off.....	15
Figure 9) E5724 storage array with the front bezel off.....	16
Figure 10) E5700 controller with the ports identified. ....	17
Figure 11) E5600 hardware overview.....	18
Figure 12) New-generation NetApp EF570 all-flash array with bezel on and off. ....	19
Figure 13) EF560 hardware overview. ....	19
Figure 14) Decision tree to determine which management components to install. ....	22
Figure 15) Data restore in Direct SAN Access mode. ....	33
Figure 16) Process of off-host backup mode.....	35
Figure 19) Backup server mapping. ....	40
Figure 20) Capacity Planning Simulator.....	42
Figure 21) Parameter Input of Example 1 in Veeam Simulator .....	43
Figure 22) Results of Example 1 using Veeam Simulator .....	44
Figure 23) Parameter Input of Example 2 in Veeam Simulator .....	45
Figure 24) Results of Example 2 using Veeam Simulator .....	46
Figure 25) Parameter Input of Example 3 in Veeam Simulator .....	47
Figure 26) Results of Example 3 using Veeam Simulator .....	48
Figure 27) Network Planning Tool.....	49
Figure 28) Test Environment Diagram .....	50
Figure 29) Processing Speeds with Veeam and E2800 .....	51
Figure 30) Restricting concurrent tasks per backup proxy .....	53
Figure 31) Restricting concurrent tasks per backup repository.....	54

# 1 Executive Summary

## The Challenge

With data growing at astounding rates, IT managers depend more and more on reliable data backup and recovery. High-growth businesses require a complete data protection solution that is reliable, flexible, and easy to use. Virtualizing an environment provides increased levels of data availability, but meeting aggressive recovery point objectives (RPOs) and recovery time objectives (RTOs) becomes increasingly difficult.

Traditional backup tools were not created for virtualized environments. That fact makes it hard for many organizations to take full advantage of their virtualized environment, and many IT managers struggle with:

- Unreliable backups
- Recovery that takes too long
- High costs that are associated with managing backup data and secondary storage
- An inability to provide reliable and true backups for compliance purposes
- Lost productivity because of management complexity
- The need to scale backup operations for growth

## The Solution

To meet these challenges, Veeam and NetApp collaborated to offer high-performance storage with reliable data protection that is designed for virtualized environments.

Veeam and NetApp help you modernize your data protection strategy with a solution that is designed to manage large data volumes and to handle the increasing performance and availability demands of a 21st-century infrastructure.

Veeam Backup & Replication unifies backup and replication in a single solution, increasing the value of backup and reinventing data protection for VMware vSphere and Microsoft Hyper-V virtual environments. The Veeam agentless design provides multiple backup options to meet your needs. Features such as source-side deduplication and compression, change block tracking, parallel processing, and automatic load balancing provide fast and efficient backups.

NetApp E-Series and EF-Series storage provides simple and reliable SAN storage that integrates seamlessly with most application environments. Its modular design helps decrease operating expenses while offering many options for connectivity, capacity, and performance that easily scale to meet the demands of a growing backup environment.

Together, Veeam and NetApp create an optimal staging area for backups, reducing backup ingest bottlenecks and providing faster backups through parallel processing.

In addition, Veeam Backup & Replication provides:

- Granular recovery of virtual machines (VMs) and files, including Microsoft Exchange and SharePoint application items
- The ability to automatically verify every backup, VM, and replica every time
- Self-service recovery of VMs and guest files without direct network connection to the VM, user permissions, or the need to deploy costly agents
- Instant VM recovery to recover a failed VM in as little as two minutes
- A choice to back up and recover what you need, where you need it, and when you need it, whether it is on site, on tape, or in the cloud

Veeam and NetApp offer the right solution for performance, flexibility, and reliability, providing an impressive modern disaster recovery solution for your vSphere or Hyper-V environment.

This document is a reference architecture for enabling a collaborative backup and recovery solution on NetApp E-Series with Veeam Backup & Replication 9.5 data protection software.

## 1.1 Introduction

Veeam and NetApp jointly developed this reference architecture to guide successful Backup & Replication 9.5 deployments with E-Series storage and to enable data and application availability.

NetApp E-Series and Veeam Backup & Replication 9.5 combine to offer a data protection and availability solution through this tested reference architecture from industry leaders NetApp and Veeam. This solution is optimized for virtual environments, providing disk-to-disk backup and recovery on high-capacity, flexible, performance-oriented NetApp E-Series storage arrays. This solution provides you with superior data management for virtual environments and high availability while also making your data highly available.

NetApp E-Series arrays provide a high-performing backup repository to house Veeam-created backups. With this capability, the recovery technologies that are enabled through Veeam can satisfy stringent RTOs. Recovery technologies such as instant VM recovery, SureBackup, and on-demand sandbox can leverage backup repositories that are capable of high I/O to achieve their full potential.

These technologies enable you to restore from your backups faster and also enable capabilities such as automated recovery verification. The technologies can also leverage backup data as an ad hoc testing environment. This capability changes the way that users have used backups in the past because more benefits are associated with having backups. No longer do backups sit idle, waiting for an emergency restore; you can apply your backups for many creative uses.

Features include:

- Recovery of a failed VM in as little as two minutes
- Near-continuous data protection with built-in replication
- Fast, agentless item recovery and e-discovery for Microsoft Exchange, SharePoint, and Active Directory, along with transaction-level recovery of SQL Server databases
- Automatic recoverability testing of every backup and every replica, every time
- Off-site backups made up to 50 times faster than the speed of standard file copy with built-in WAN acceleration
- Fast and secure cloud backups with Veeam Cloud Connect
- Deduplication and compression to minimize storage consumption
- Off-site recovery with one-click site failover and support for facilitated data center migrations with zero data loss

## 1.2 About NetApp

NetApp creates innovative products, including storage systems and software that help customers around the world store, manage, protect, and retain one of their most precious corporate assets: their data. We are recognized throughout the industry for continually pushing the limits of today's technology so that our customers don't have to choose between saving money and acquiring the capabilities that they need to be successful.

We find ways to enable our customers to do things that they couldn't do before and at a speed that they never thought possible. We partner with industry leaders to create efficient and cost-effective solutions that are optimized for customers' IT needs and to deliver to and support these customers worldwide. Leading organizations worldwide count on NetApp for software, systems, and services to manage and store their data. Customers value our teamwork, expertise, and passion for helping them succeed now and into the future (<http://www.netapp.com>).

### 1.3 About Veeam

Veeam recognizes the new challenges that companies across the globe face in enabling the always-on business, a business that must operate 24/7/365. To address this challenge, Veeam delivers availability for the modern data center by helping to make sure of recovery time and point objectives (RTPOs) of less than 15 minutes for all applications and data. Information about Veeam is available at [www.Veeam.com](http://www.Veeam.com). Veeam's corporate headquarters are in Baar, Switzerland; a main Americas office is located in Columbus, Ohio.

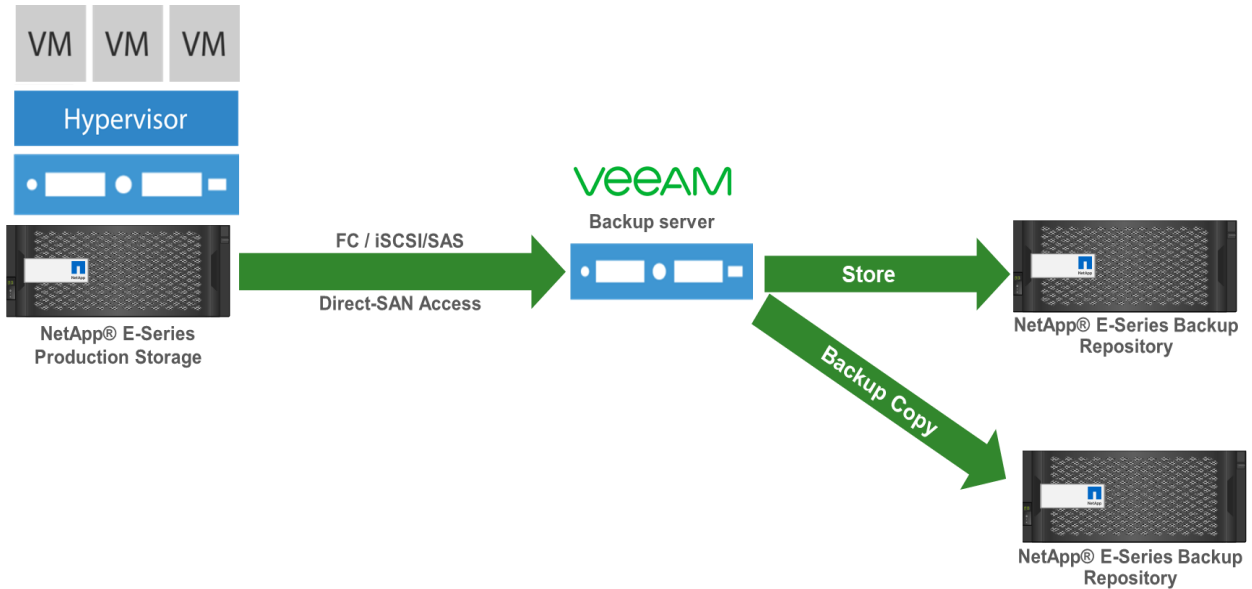
## 2 Reference Architecture Overview

This section details reference architectures that range from those of small environments that protect a few terabytes of data to those in enterprise-size environments with petabytes of data under management.

### 2.1 NetApp E-Series as Veeam Backup & Replication Repositories for Backup and Archiving

Figure 1 gives us a graphical representation of a NetApp E-Series and Veeam setup.

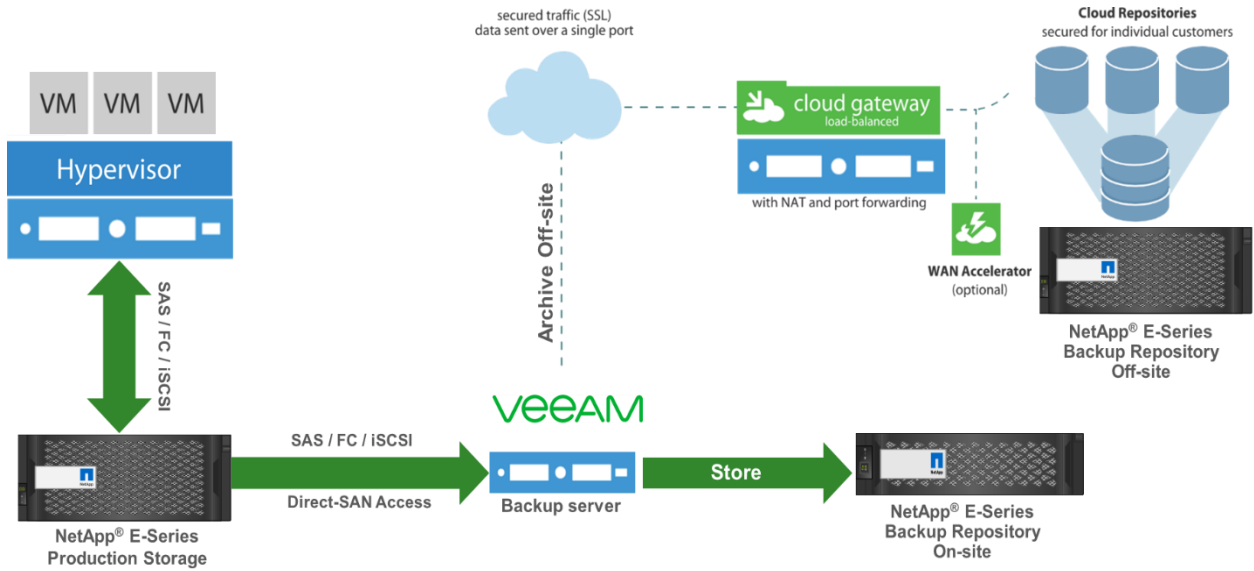
Figure 1) NetApp E-Series as Veeam Backup & Replication repositories for backup and archiving.



### 2.2 NetApp E-Series as a Veeam Backup & Replication Backup Repository and an Off-Site Cloud Repository

Figure 2 shows NetApp E-Series as a Veeam Backup & Replication backup repository and an off-site cloud repository.

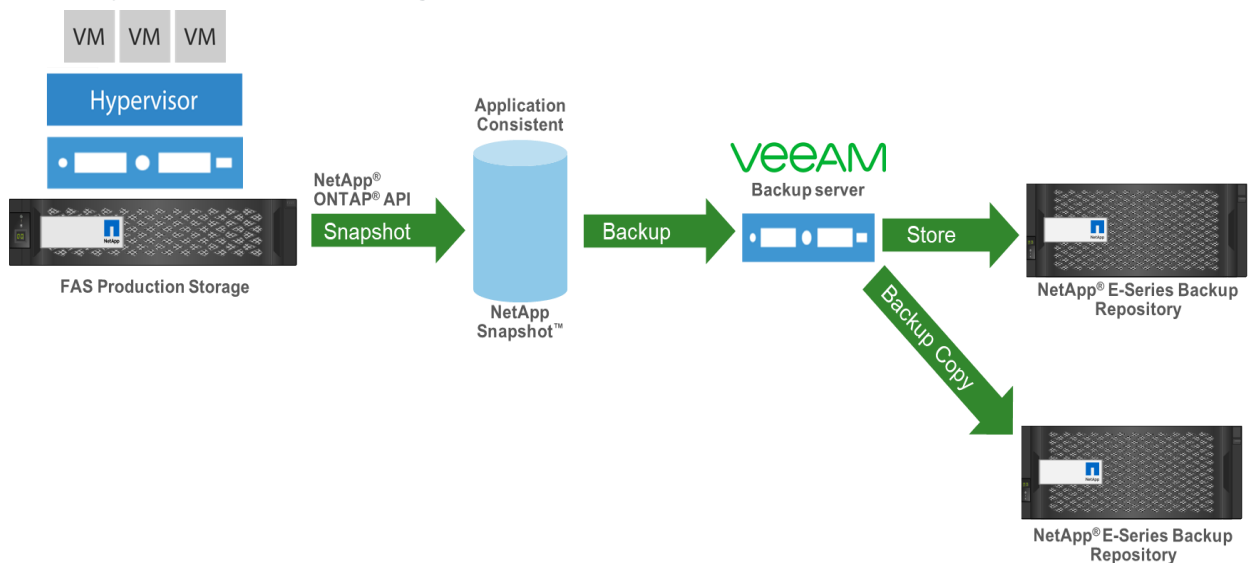
Figure 2) NetApp E-Series as a Veeam Backup & Replication backup repository and an off-site cloud repository.



### 2.3 NetApp E-Series as a Veeam Backup & Replication Backup Repository for NetApp FAS Production Storage

Figure 3 illustrates Veeam integration with a NetApp FAS series production storage array, with newly created backups going to an E-Series array for storage. To provide disaster recovery, backups can also be sent off the premises to another backup repository (another E-Series system). Veeam provides a backup copy job for such scenarios; this job can be leveraged for backups off the premises or for long-term archiving by using Veeam’s built-in grandfather-father-son (GFS)–type retention.

Figure 3) NetApp E-Series as a Veeam Backup & Replication backup repository and off-premises backup repository for FAS production storage.



Leveraging Veeam’s backup copy job architecture is important for achieving that last level of protection. The off-site copy provides safeguards for an entire data center–level disaster. Veeam also provides an optional WAN acceleration component that can help reduce bandwidth utilization. This component can



play a huge role in environments that have active-active sites or that have low available bandwidth to start with. The forever-incremental nature of the backup copy job enables only incremental change data to be transmitted off site after the initial copy. Preseeding options are available for the initial transfer of data for environments that need it.

## 2.4 Deployment Scenarios

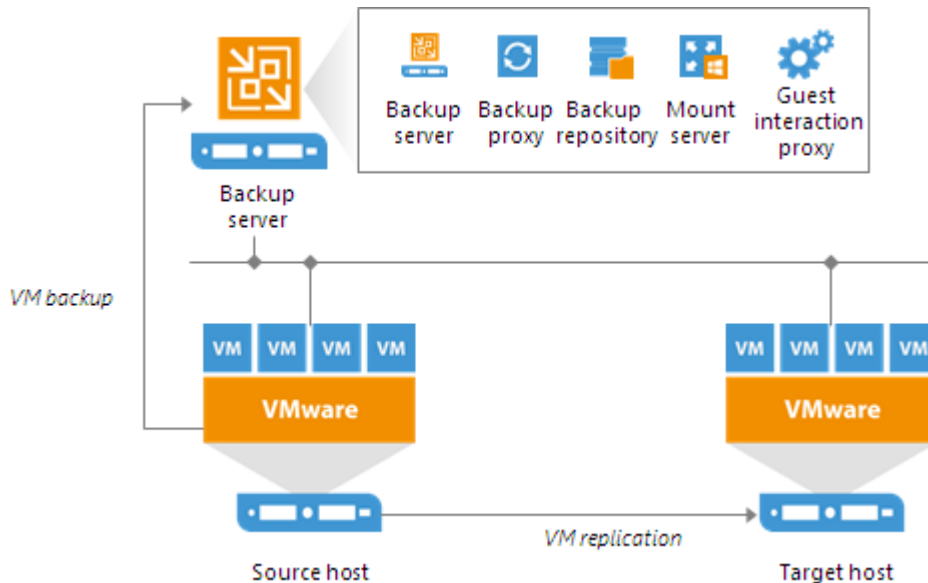
### Simple Deployment

In a simple deployment scenario, one instance of Veeam Backup & Replication is installed on a physical or virtual Windows-based machine. This installation is referred to as a *Veeam backup server*.

*Simple deployment* (Figure 4) implies that the Veeam backup server fills three major roles:

- It functions as a management point, coordinates all jobs, controls job scheduling, and performs other administrative activities.
- It acts as the default backup proxy for handling job processing and for transferring backup traffic. All services that are necessary for the backup proxy functionality are installed on the Veeam backup server locally.
- It is used as the default backup repository. During installation, Veeam Backup & Replication checks volumes of the machine on which you install the product and identifies a volume with the greatest amount of free disk space. On this volume, Veeam Backup & Replication creates the backup folder that is used as the default backup repository.

Figure 4) Simple deployment.



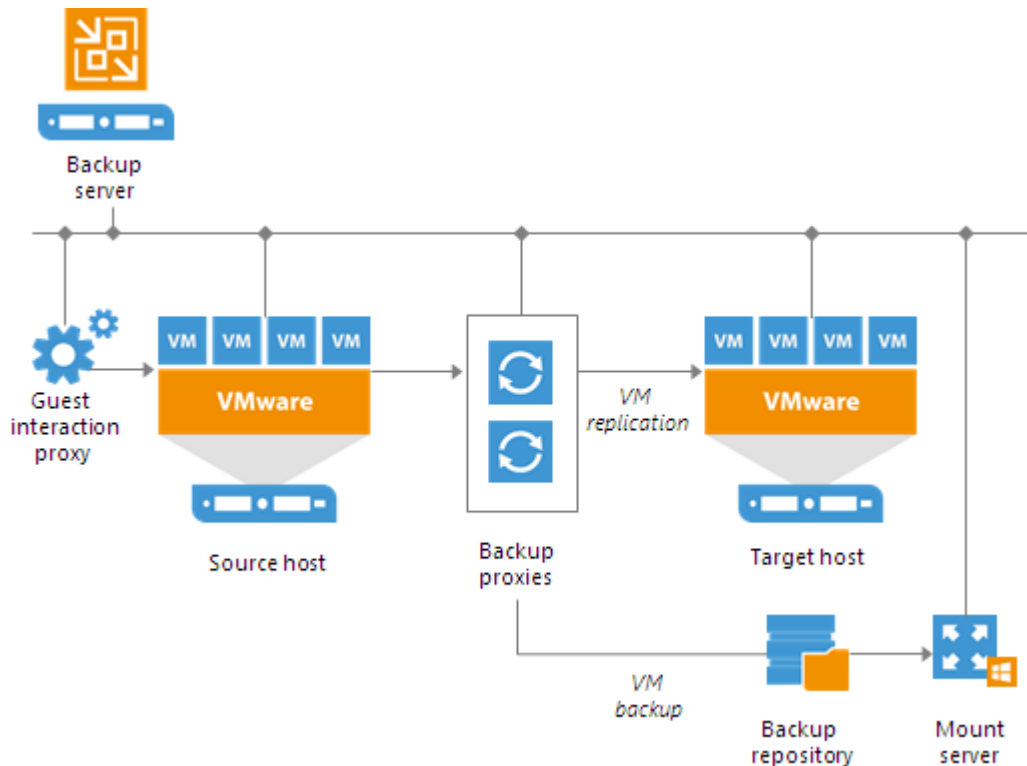
If you plan to back up and replicate only a few VMs or evaluate Veeam Backup & Replication, this configuration is enough to get you started. Veeam Backup & Replication is ready for use right out of the box; as soon as it is installed, you can start using the solution to perform backup and replication operations. To balance the load of backing up and replicating your VMs, you can schedule jobs at different times.

## Advanced Deployment

In large-scale virtual environments with numerous jobs, the load on the Veeam backup server is heavy. In this case, NetApp recommends using the advanced deployment scenario (Figure 5), which moves the backup workload to dedicated backup proxies and backup repositories.

The essence of the advanced deployment is that the backup proxy takes off part of Veeam backup server activities (namely, it collects and processes data and moves backup traffic from the source to the target). In addition, the Veeam backup server no longer acts as a storage location. The backup proxy transports VM data to a dedicated backup repository that keeps backup files, VM copies, metadata, and so on. The Veeam backup server in this scenario functions as a manager for deploying and maintaining backup proxies and repositories.

Figure 5) Advanced deployment.



To deploy a backup proxy or a backup repository, add a server to Veeam Backup & Replication and assign a proxy or a repository role to it, as applicable. Veeam Backup & Replication automatically installs lightweight components and services on these servers. A backup proxy does not require a separate SQL Server database; all settings are stored centrally, within the SQL Server database that the Veeam backup server uses.

With the advanced deployment scenario, you can easily meet your current and future data protection requirements. You can expand your backup infrastructure horizontally in a matter of minutes to match the amount of data that you want to process and the available network throughput. Instead of growing the number of backup servers or constantly tuning job scheduling, you can install multiple backup proxies and repositories and distribute the backup workload among them. The installation process is fully automated, which simplifies deploying and maintaining the backup infrastructure in your virtual environment.

In virtual environments with several proxies, Veeam Backup & Replication dynamically distributes backup traffic among those proxies. You can explicitly map a job to a specific proxy, or you can let Veeam Backup & Replication choose the most suitable proxy. If you opt for the latter, Veeam Backup & Replication checks the settings of available proxies and selects the most appropriate one for the job. The proxy server to be used should have access to the source and target hosts as well as to the backup repository to which files are written.

The advanced deployment scenario can be a good choice for backing up and replicating off site. You can deploy a backup proxy in the production site and another one in the disaster recovery (DR) site, closer to the backup repository. When a job is performed, backup proxies on both sides establish a stable connection, so this architecture also allows efficient data transport over a slow network connection or WAN.

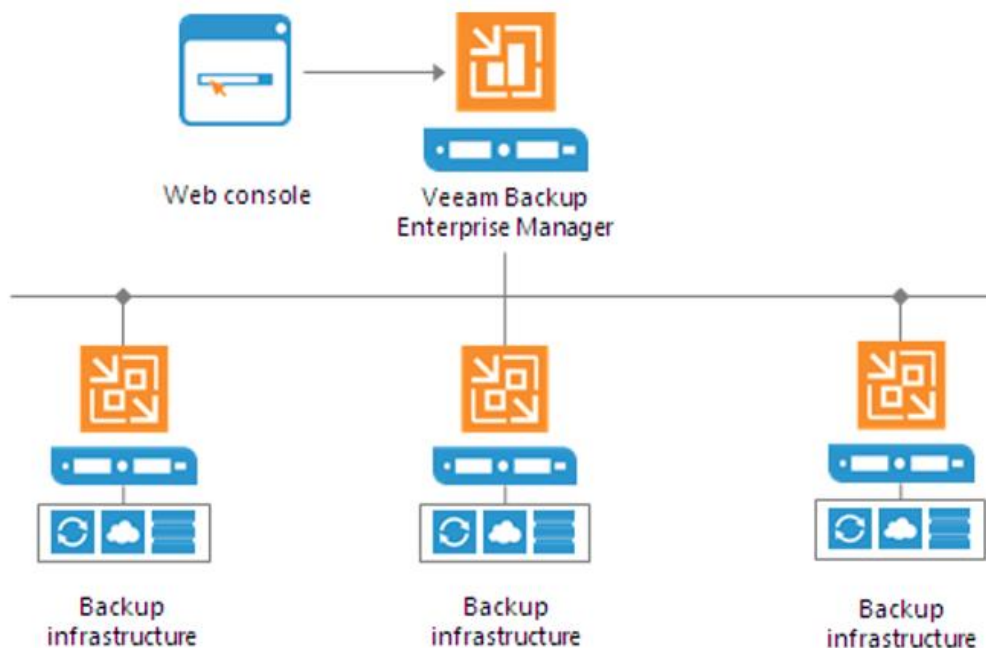
To regulate backup load, you can specify the maximum number of concurrent tasks per proxy and set up throttling rules to limit proxy bandwidth. The maximum number of concurrent tasks can also be specified for a backup repository in addition to the value of the combined data rate for it.

Another advantage of the advanced deployment scenario is that it contributes to high availability. Jobs can migrate between proxies if one of them becomes overloaded or unavailable.

## Distributed Deployment

NetApp recommends the distributed deployment scenario (Figure 6) for large geographically dispersed virtual environments with multiple Veeam backup servers that are installed across different sites. These backup servers are federated under Veeam Backup Enterprise Manager, an optional component that provides centralized management and reporting for these servers through a web interface.

Figure 6) Distributed deployment.



Veeam Backup Enterprise Manager collects data from Veeam backup servers and enables you to run backup and replication jobs across the entire backup infrastructure through a single interface. You can also edit those jobs and clone jobs by using a single job as a template. Enterprise Manager also provides reporting data for various areas (for example, all jobs that were performed within the past 24 hours or 7 days, all VMs that were engaged in these jobs, and so on).

By using indexing data that is consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for VM guest OS files in VM backups that are created on all Veeam backup servers. You can search even if the files are stored in repositories on different sites, and you can recover them in a single click. Searching for VM guest OS files is enabled through Veeam Backup Enterprise Manager; to streamline the search process, you can optionally deploy a Veeam Backup Search server in your backup infrastructure.

With flexible delegation options and security roles, IT administrators can delegate the necessary file restore or VM restore rights to authorized personnel in your organization. For example, they can allow database administrators to restore Oracle or SQL Server VMs.

If you use Veeam Backup Enterprise Manager in your backup infrastructure, you do not need to install licenses on every Veeam backup server that you deploy. Instead, you can install one license on the Veeam Backup Enterprise Manager server, and it is applied to all servers across your backup infrastructure. This approach simplifies tracking license usage and license updates across multiple Veeam backup servers.

In addition, VMware administrators can benefit from the Veeam plug-in for vSphere Web Client, which can be installed by using Veeam Backup Enterprise Manager. Administrators can analyze cumulative information about used and available storage space; view statistics on processed VMs; and review success, warning, and failure counts for all jobs. Administrators can also easily identify unprotected VMs and perform capacity planning for repositories, all directly from vSphere.

## 3 NetApp E-Series and NetApp EF-Series Arrays

### 3.1 E-Series Hardware

The NetApp E-Series hardware portfolio can be divided into three categories: entry-level, midrange, and all-flash storage systems.

#### Entry-Level E-Series Storage Systems

##### NetApp E-Series E2800 Storage System

NetApp E-Series E2800 storage systems address wide-ranging data storage requirements with balanced performance. The E2800 system is equally adept at handling large sequential I/O for video, analytical, and backup applications, as well as small random I/O requirements for small and medium-sized enterprise mixed workloads. The E2800 brings together the following advantages:

- Support for all-flash and hybrid drive configurations
- Modular host interface flexibility (SAS, FC, and iSCSI)
- High reliability (99.999% reliability)
- Intuitive management: simple administration for IT generalists and detailed drill-down for storage specialists

The new entry-level E2800 is a 12Gb SAS 3 system with NetApp SANtricity® 11.30 software. The E2800 introduces new embedded management, including the browser-based SANtricity System Manager 11.30, which features the following new capabilities:


- Embedded web services
- The easy-to-use GUI of SANtricity System Manager
- The ability to store and present up to 30 days of performance data, including I/O latency, IOPS, CPU utilization, and throughput
- The ability to do application and workload tagging

- Easier alert management, including an embedded SNMP agent and MIB
- Embedded NetApp AutoSupport® functionality

Together, these features create an entry-level storage system with the flexibility and performance capabilities to support enterprise workloads without sacrificing simplicity and efficiency. In addition, the E2800 storage system's fully redundant I/O paths, advanced protection features, and extensive diagnostic capabilities deliver a high level of availability, data integrity, and security.

The E2812 and E2824 shelf options support one or two controller canisters, and the E2860 supports only two controller canisters. All shelves support dual power supplies and dual fan units for redundancy (the shelves have an integrated power fan canister). The shelves are sized to hold 12 drives, 24 drives, or 60 drives, as shown in Figure 7.

Figure 7) E2800 hardware overview.



**Note:** For detailed information about the E2800 system, see [TR-4538: Introduction to NetApp E-Series E2800](#).

### Midrange E-Series Storage Systems

The midrange portfolio currently includes the E5700 and E5600 controller pair.

## NetApp E-Series E5700 Storage System

NetApp E-Series E5700 hybrid arrays running SANtricity OS 11.40 have a new modern look, leverage the new 12Gbps DE460C and DE224C drive shelves, support a more secure UI, and deliver significantly higher performance than their predecessor E-Series arrays.

For backup/recovery applications and other high-capacity workloads, the E5760 hybrid array dramatically increases the maximum supported capacity per array footprint from 384 drives to 480 drives, or eight total 4RU (rack unit) shelves. Currently that's up to 4.8PB of ultradense raw capacity in 32RU of rack space, or 150TB/RU, using 10TB NL-SAS drives. This storage density continues to grow every 6 to 12 months as larger capacity drives are qualified for E-Series systems.

As a result, adopting a strategy of purchasing multishelf systems up front with minimum drive counts to start with allows you to continue to increase the footprint density of your E-Series investment over time. This density reduction is achieved by adding sets of higher-capacity drives as you need them to satisfy future storage growth. You can fully accomplish this goal without changing the footprint or planning new power and, most important, without incurring a service disruption.

In fact, growing by drive packs makes the process of growing over time extremely easy and very cost effective. Simply order a new RAID group's or pool's worth of drives as a capacity building block when you need it. Then install the drives without disruption to live systems, map the new capacity to new requirements, or grow existing capacity while maintaining optimal system resiliency and performance over the life of your storage system. This approach drives down your long-term total cost and provides a low-risk and built-in grow-on-demand strategy that can be easily executed in quarter-driven budget cycles.

Figure 8 shows the E5760 array front and rear views. The front view shows just how easy it is to open a drive drawer to install new drives.

**Note:** The front bezel should be installed during normal operating conditions.

Figure 8) New E5760 hybrid storage array with the front bezel off.


### E5760 front bezel



### E5760 front view with drawer 5 open



### E5760 rear view with controllers installed




For workloads that require fast storage such as Splunk and other analytics applications, high-performance databases, and specialty applications that require ultralow latency storage, E-Series E5724 hybrid arrays support up to 192 drives, starting with a base set of 10K RPM SAS drives for the HDD tier. You can add more 10K SAS drives or up to 120 solid-state drives (SSDs) to build a fast tier in the same array. The E5724 also supports 15.3TB SSDs to build a large capacity fast tier (~1.8PB fast, raw capacity). For extreme flexibility, you can add a DE460C expansion drive shelf that supports both SSDs and NL-SAS drives.



**Note:** The DE460C shelf does not support 15.3TB SSDs. The 800GB, 1.6TB, and 3.2TB SSDs are supported in the DE460C shelf.

The system can deliver consistent submillisecond latency response times for small random workloads, or it can deliver up to 21GBps for large sequential read workloads, about 8GBps large sequential write workloads.

Figure 9) E5724 storage array with the front bezel off.



**Note:** E5700 controllers are not offered in the 12-drive DE212C shelf. Only the E-Series E2800 controllers are offered in the 12-drive configuration (that is, E2812).

Each E5700 controller provides two Ethernet management ports for out-of-band management and has two 12Gbps (x4 lanes) wide-port SAS drive expansion ports for redundant drive expansion paths to the drives.

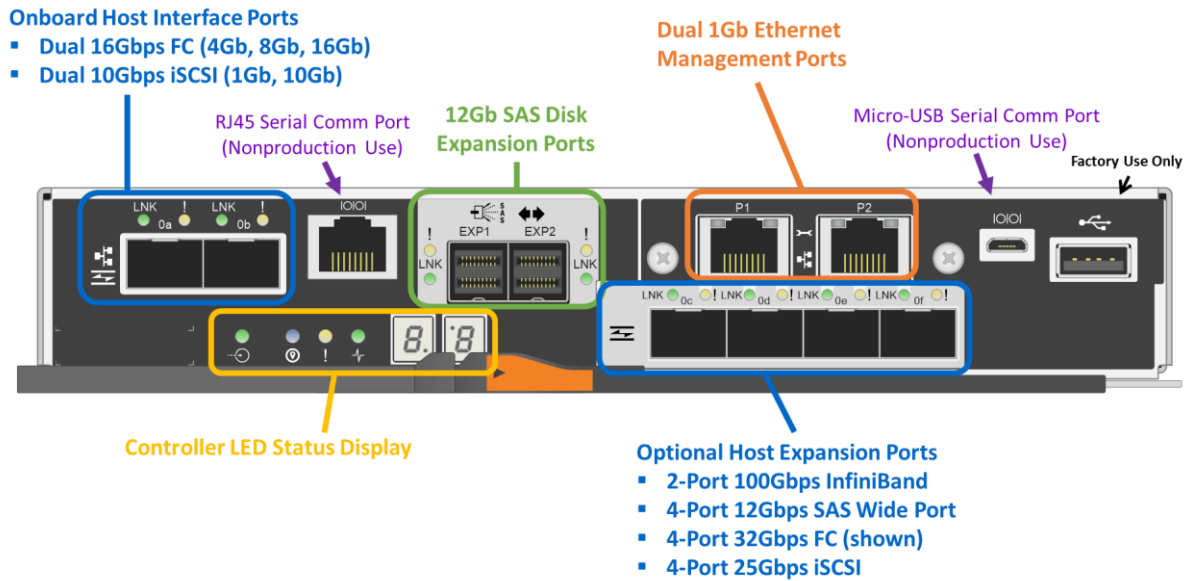
The E5700 controllers also include two built-in host ports, either two 16Gb FC or two 10Gb iSCSI, and your choice of the following optional host interface cards (HICs):

- 4-port 12Gb SAS (mini-SAS 3 connector)
- 4-port 32Gb FC (OM4 fiber required)
- 4-port 25Gb iSCSI (OM4 fiber required)
- 2-port 100Gb InfiniBand ([IB] requires 100Gb-capable cables and HCAs)

**Note:** A software feature pack can be applied in the field to change the host protocol of the optical baseboard ports from FC to iSCSI or from iSCSI to FC. In addition, the IB protocol can be changed to one of three choices: iSER, SRP, and NVMe over Fabrics, which is abbreviated NVMe-oF (IB). Only one IB protocol can be active on an E5700 array.



Figure 10) E5700 controller with the ports identified.



**Note:** You can find further information about the E5700 system at [TR-4627: Introduction to NetApp E5700 Arrays](#) .


## NetApp E-Series E5600 Storage System

The E5600 provides the following benefits:

- Support for wide-ranging workloads and performance requirements
- Fully redundant I/O paths, advanced protection features, and proactive support monitoring and services for high levels of availability, integrity, and security
- Increased IOPS performance by up to 35% compared with the previous high-performance generation of E-Series products
- A winning combination of leading IOPS performance at low latencies and throughput density that makes the E5600 a great choice for high-performance workloads
- A level of price-performance, density, and economics that leads in the industry

The E5600 system is available in three shelf options, which support both HDDs and SSDs. All three shelf options include dual controller modules, dual power supplies, and dual fan units for redundancy (the 12-drive and 24-drive shelves have integrated power and fan modules). The shelves are sized to hold 60 drives, 24 drives, or 12 drives, as shown in Figure 11.

Figure 11) E5600 hardware overview.



**Note:** For detailed information about the E5600 system, visit [TR-4544: Introduction to NetApp E-Series E5600](#).


## All-Flash EF-Series Storage Systems

The all-flash portfolio currently includes the EF570 and EF560 controller pair.

### NetApp E-Series EF570 Storage System

NetApp EF570 arrays have a new modern look, leverage the new 12Gbps DE224C drive shelves, support a more secure user interface, and deliver stunning performance for both mixed random workloads and large sequential workloads in one powerful all-flash array package.

Figure 12) New-generation NetApp EF570 all-flash array with bezel on and off.



The EF570 can deliver consistent submillisecond latency response times for up to 1,000,000 4KB random read IOPS with as little as 24 SSDs, or the same configuration can deliver up to 16GBps large sequential read throughput, about 9GBps cache mirrored large sequential write throughput. If you add an expansion drive shelf and 12 additional SSDs, the EF570 array delivers up to 21GBps throughput for large sequential read workloads and up to 12GBps for full stripe write workloads.


**Note:** You can find further information about the EF560 system at [TR-4637: Introduction to NetApp E-Series EF570 Arrays](#).

### NetApp E-Series EF560 Storage System

The NetApp EF560 all-flash array (Figure 13) provides a robust system for delivering exceptional performance to mission-critical applications. The EF560 flash array leverages the latest in SSD technologies and a strong heritage of handling diverse workloads and providing superior business value through accelerating latency-sensitive and high-I/O applications.

The EF560 is available in a 2U 24-drive shelf that holds up to 24 2.5-inch SSDs and features dual RAID controllers and dual power supplies with integrated fans. The EF560 flash array supports up to 120 SSDs with the addition of four expansion drive shelves. Each of the two EF560 controllers has a combination of onboard storage-side SAS drive expansion ports and add-on host interface ports.

Figure 13) EF560 hardware overview.



**Note:** You can find further information about the EF560 system in [TR-4546: Introduction to NetApp EF560 Flash Array](#).

## 3.2 Key NetApp E-Series and EF-Series Features

Following are some of the important features of NetApp E-Series:

- NetApp E-Series and EF-Series provide highly scalable capacity and performance.
- Maximum density promotes optimal space utilization and reduced power and cooling requirements.
- Dynamic Disk Pools (DDP) technology enables dynamic rebalancing of drive count changes to easily expand storage capacity as backup needs grow while providing added data protection with faster rebuild times if a drive fails.
- You can easily integrate E-Series through multiple host interfaces, drive technologies, and disk shelf options for flexible deployment.
- Full disk encryption offers data security. AES 128-bit encryption from SafeStore encryption services and simplified key management protect data throughout the drive's lifecycle without sacrificing performance.
- No-cost application add-ins enable you to easily manage storage from a single display within the native application.
- Tighter economics provide:
  - A better price per gigabyte
  - A lower total cost of ownership (TCO)
  - Supportability and reliability, availability, and serviceability to minimize the cost of support calls

## 3.3 Enterprise Reliability and Availability

Field-proven technology protects your valuable data:

- E-Series is the right choice for peace of mind because it:
  - Leverages knowledge from 1 million systems
  - Is backed by a worldwide support organization
- E-Series is architected for the highest reliability and availability and includes:
  - A fully redundant I/O path with automated failover
  - Online configuration, expansion, and maintenance
  - Advanced monitoring and diagnostic features that enable fast problem resolution
  - Proactive tracking of SSD wear life and sending of alert messages
- E-Series provides enterprise data protection, including:
  - Robust disaster recovery (sync and async)
  - Local protection with high-efficiency NetApp Snapshot™ copies

## 3.4 SANtricity OS 11.40: New Features

E5700 and EF570 arrays are new with the NetApp SANtricity OS 11.40 release, but there are also significant software enhancements in the release that apply to all the new-generation E-Series arrays, including the E5700, EF570, and entry-level E2800 arrays. An onboard web-based GUI manages these arrays: SANtricity System Manager.

New features in SANtricity OS 11.40 include:

- Support for directory services using Lightweight Directory Access Protocol (LDAP)
- Support for role-based access control (RBAC): five standard roles defined with varying permission levels
- Support for certification authority (CA) and Secure Sockets Layer (SSL) certificates
- Secure CLI implementation: secure when the certificates are installed

- Added support for an external encryption key manager in addition to the legacy E-Series drive security onboard encryption key manager
- Security enhancements that extend to the onboard web services API, where user account passwords are now required

**Note:** If you want to run in the previous security mode with a single administrative password and still use symbols to communicate using API, the new security features can be disabled by the admin user when the storage system is initially set up.

In addition to LDAP and RBAC, there are also enhancements to our most used host multipath functionality that were released in previous SANtricity OS maintenance releases and are now part of the SANtricity OS 11.40 general availability (GA) release.

**Note:** To get in-depth details about SANtricity OS 11.40, see [TR-4627: Introduction to NetApp E5700 Arrays](#).

### 3.5 SANtricity System Manager vs. Storage Manager

E-Series storage systems are managed by using the SANtricity Storage Manager or System Manager, which offers both GUI management and CLI management interfaces for out-of-band and in-band system management. NetApp generally recommends out-of-band management. However, in-band management is useful for cases in which there is no IP network access to remote storage systems but there are in-band traffic connections to the storage system from a local host.

There are two different versions of the storage management software:

- SANtricity System Manager is used to manage individual E2800 storage arrays.
- SANtricity Storage Manager, with its Enterprise Management Window (EMW), provides an aggregated view of all E-Series arrays. The Array Management Window (AMW) of SANtricity Storage Manager is used to manage the E2700, E5600, EF560, and all earlier storage arrays.

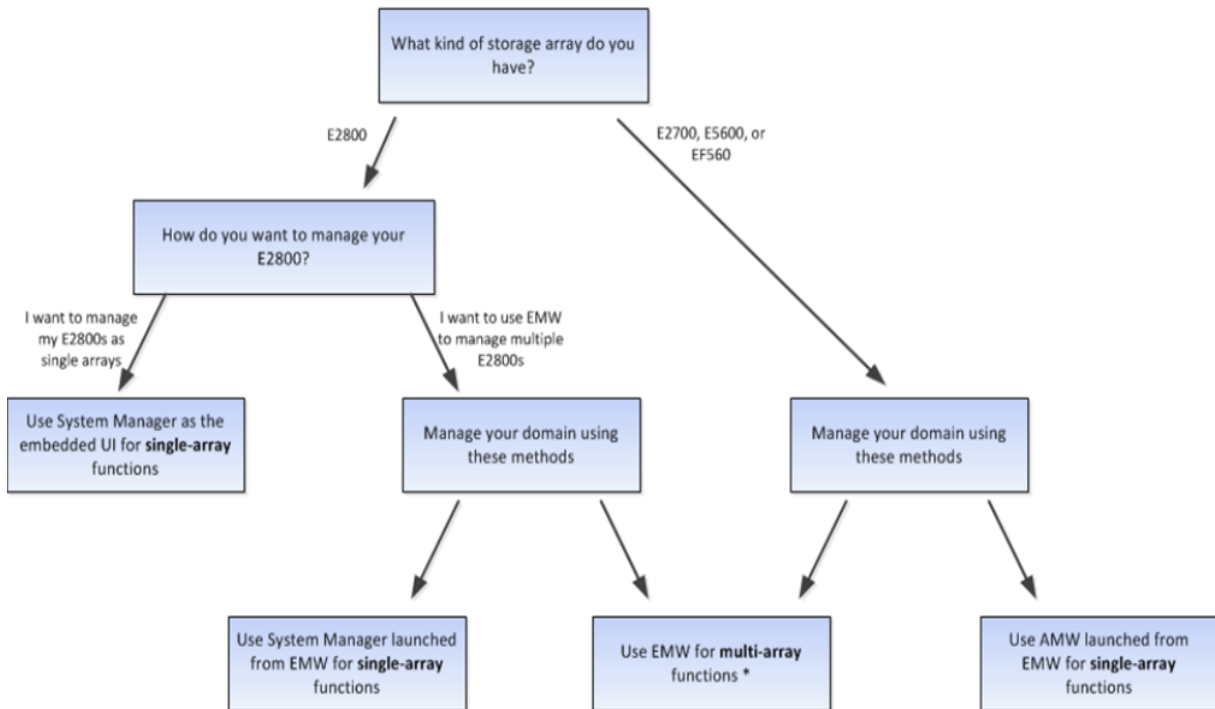
When you choose to manage a storage array from the EMW, the EMW opens the appropriate software (AMW or System Manager), depending on what controller the storage array contains.

The key features of System Manager include the following:

- Runs on box: You do not have to install any storage management software unless you need an aggregated view or use a mirroring feature.
- Displays in a browser and is mobile ready.
- Has a modern look and feel, with a tile-based GUI and an easy-to-use online help system.
- Uses simplified workflows and simplified terminologies.
- Includes new functionality, such as application and workload tagging, enhanced performance data, an embedded monitor, and a graphical view of thin volume usage.
- Includes an embedded RESTful API that can be used for management.

If you have purchased an E2800-based storage array, your decisions about what components to install depend on how you answer the questions in Figure 14.

Figure 14) Decision tree to determine which management components to install.



You can find further information about System Manager in section 3 of the [E2800 technical report](#). You can find more information about Storage Manager in section 2 of the [E5600 technical report](#).

## 4 Veeam Backup & Replication 9.5

### 4.1 Overview

Veeam Backup & Replication is a data protection and disaster recovery solution for VMware vSphere and Microsoft Hyper-V virtual environments of any size or complexity. By combining all the necessary functions in one intuitive interface, Veeam Backup & Replication solves the most critical problems of virtualized infrastructure management. The solution also protects mission-critical VMs from both hardware and software failures.

### 4.2 Solution Architecture

Veeam Backup & Replication is composed of the following three elements:

- Backup server
- Backup proxy
- Backup repository

#### Veeam Backup Server

The Veeam backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the configuration and control center. The Veeam backup server performs all types of administrative activities. This server:

- Coordinates backup, replication, recovery verification, and restore tasks

- Controls job scheduling and resource allocation
- Is used to set up and manage backup infrastructure components and to specify global settings for the backup infrastructure

In addition to its primary functions, a newly deployed Veeam backup server also acts as a default backup proxy and the backup repository (it manages data handling and data storing tasks).

The Veeam backup server uses the following services and components:

- **Veeam Backup Service** is a Windows service that coordinates all the operations that Veeam Backup & Replication performs, such as backup, replication, recovery verification, and restore tasks. The Veeam Backup Service runs under the local system account or an account that has the local administrator permissions on the backup server.
- **Veeam Broker Service** interacts with the virtual infrastructure to collect and cache the virtual infrastructure topology. It jobs and tasks query information about the virtual infrastructure topology from the Broker Service, which accelerates job and task performance.
- **Veeam Guest Catalog Service** manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog, which is a folder on the backup server. The Veeam Guest Catalog Service running on the backup server works in conjunction with search components that are installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.
- **Veeam Mount Service** mounts backups and replicas for file-level access, browsing the VM guest file system and restoring VM guest OS files and application items to the original location.
- **Veeam backup proxy services:** In addition to dedicated services, the backup server runs a set of data mover services. For details, see the following “Backup Proxy” section.
- **Veeam Backup & Replication Configuration database** stores data about the backup infrastructure, jobs, sessions, and so on. The database instance can be on a SQL Server that is installed either locally (on the same machine where the backup server is running) or remotely.
- **Veeam Backup & Replication console** provides the application UI and allows user access to the application’s functionality.
- **Veeam Backup PowerShell snap-in** is an extension for Microsoft Windows PowerShell 2.0. Veeam Backup PowerShell adds a set of cmdlets that enable you to perform backup, replication, and recovery tasks through the PowerShell CLI or run custom scripts to fully automate Veeam Backup & Replication operation.

## Backup Proxy

When Veeam Backup & Replication is first installed, the Veeam backup server coordinates all job activities and handles data traffic. So, when you run a backup, replication, VM copy, or VM migration job or perform restore operations, VM data is moved from the source to the target through the Veeam backup server. This scenario is acceptable for virtual environments in which few backup jobs are performed. In large-scale environments, however, the workload on the Veeam backup server is significant.

To take the workload off the Veeam backup server, Veeam Backup & Replication uses backup proxies. A backup proxy is an architecture component that sits between the data source and the target and is used to process jobs and to deliver backup traffic. In particular, the backup proxy tasks include retrieving VM data from the production storage. The tasks also include compressing the data and sending it to the backup repository (for example, if you run a backup job) or to another backup proxy (for example, if you run a replication job). As the data handling task is assigned to the backup proxy, the Veeam backup server becomes the point of control for dispatching jobs to proxy servers.

The role of a backup proxy can be assigned to a dedicated Windows Server (physical or virtual) in your virtual environment. You can deploy backup proxies both in the primary site and in remote sites. To optimize the performance of several concurrent jobs, you can use a number of backup proxies. In this case, Veeam Backup & Replication distributes the backup workload between available backup proxies.



By using backup proxies, you can easily scale your backup infrastructure up and down based on your demands. Backup proxies run lightweight services that take a few seconds to deploy. Deployment is fully automated: Veeam Backup & Replication installs the necessary components on a Windows-based server when you add it to the product console. As soon as you assign the role of a backup proxy to the added server, Veeam Backup & Replication starts the required services on it.

The primary role of the backup proxy is to provide an optimal route for backup traffic and to enable efficient data transfer. Therefore, when deploying a backup proxy, you must analyze the connection between the backup proxy and the storage with which it is working. Depending on the type of connection, the backup proxy can be configured in one of the following ways (starting with the most efficient):

- A machine used as a backup proxy should have direct access to the storage on which VMs reside or to the storage to which VM data is written. In this way, the backup proxy retrieves data directly from the datastore, bypassing the LAN.
- The backup proxy can be a VM with HotAdd access to VM disks on the datastore. This type of proxy also enables LAN-free data transfer.
- If neither of the preceding scenarios is possible, you have alternatives. You can assign the role of the backup proxy to a machine on the network that is closer to the source or closer to the target storage with which the proxy works. In this case, VM data is transported over the LAN by using the Network Block Device (NBD) protocol.

Depending on the type of backup proxy and your backup architecture, the backup proxy can use one of the following data transport modes: Direct SAN Access, Virtual Appliance, or Network. If the VM disks are on the SAN storage and the SAN storage is added to the Veeam Backup & Replication console, the backup proxy can also use the Backup from Storage Snapshots mode. You can select the transport mode or let Veeam Backup & Replication automatically choose it.

The backup proxy uses the following services and components:

- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows Server after it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system and installs and upgrades necessary components and services depending on the role that is selected for the server.
- **Veeam Transport** is responsible for deploying and coordinating executable modules that act as data movers and that perform the main job activities on behalf of Veeam Backup & Replication. These activities include communicating with VMware Tools, copying VM files, performing data deduplication and compression, and so on.

## Backup Repository

A backup repository is a location that Veeam Backup & Replication jobs use to store backup files, copies of VMs, and metadata for replicated VMs. Technically, a backup repository is a folder on the backup storage. By assigning different repositories to jobs and by limiting the number of parallel jobs for each one, you can balance the load across your backup infrastructure.

# 5 Veeam Backup & Replication 9.5 Requirements

## 5.1 VMware System Requirements

### Veeam Backup Server

Table 1 lists the specifications and requirements of the Veeam backup server with VMware.



Table 1) Veeam backup server specifications.

Specification	Requirement
CPU	Modern x86-64 processor (minimum 2 cores or vCPUs).
Memory	4GB RAM plus 500MB RAM for each concurrent job.
Disk space	2GB for product installation and 10GB per 100 VMs for a guest file system catalog folder (persistent data). Sufficient free disk space for <i>Instant VM Recovery</i> cache folder (nonpersistent data, at least 10GB is recommended).
Network	1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication are recommended. High latency and reasonably unstable WAN links are supported.
OS	<p>A 64-bit version of the following operating systems is supported:</p> <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1</li> <li>• Microsoft Windows Server 2008 SP2</li> <li>• Microsoft Windows Server 2003 SP2</li> <li>• Microsoft Windows 10</li> <li>• Microsoft Windows 8.x</li> <li>• Microsoft Windows 7 SP1</li> </ul>
Software	<p>During setup, the installer performs a system configuration check to determine whether all prerequisite software is available on the machine in which you plan to install Veeam Backup &amp; Replication. If some of the required software components are missing, the setup wizard offers to install the missing software automatically. This task refers to:</p> <ul style="list-style-type: none"> <li>• Microsoft .NET Framework 4.5.2</li> <li>• Microsoft Windows Installer 4.5</li> <li>• Microsoft SQL Server Management Objects</li> <li>• Microsoft SQL Server System CLR Types</li> <li>• Microsoft Visual C++ 2010 SP1 Redistributable Package</li> </ul> <p>The following software must be installed manually:</p> <ul style="list-style-type: none"> <li>• Microsoft PowerShell 2.0 or later (optional)</li> <li>• Firefox, Google Chrome, Microsoft Edge, or Microsoft Internet Explorer 10.0 or later</li> </ul>
SQL Server database	<p>Local or remote installation of the following versions of Microsoft SQL Server (both Full and Express Editions are supported):</p> <ul style="list-style-type: none"> <li>• Microsoft SQL Server 2016</li> <li>• Microsoft SQL Server 2014</li> <li>• Microsoft SQL Server 2012 (Microsoft SQL Server 2012 Express Edition SP3 is included in the setup)</li> <li>• Microsoft SQL Server 2008 R2</li> <li>• Microsoft SQL Server 2008</li> </ul>

**Note:**

1. You might plan to back up VMs that are running Microsoft Windows Server 2012 R2 or Microsoft Windows Server 2016, and data deduplication is enabled for some VM volumes. In that case, NetApp recommends that you deploy Veeam Backup & Replication on a machine that is running Microsoft Windows Server 2012 R2 or Microsoft Windows Server 2016 with data deduplication enabled. Otherwise, some types of restore operations for these VMs (such as Microsoft Windows File Level Recovery) might fail.
2. Because of its limitations, Microsoft SQL Server Express Edition can be used only for evaluation purposes or for a small-scale production environment. For environments with several VMs, you must install a fully functional commercial version of Microsoft SQL Server.

## VMware Backup Proxy Server

Table 2 lists the specifications and requirements of the VMware backup proxy server.

Table 2) VMware backup proxy server specifications.

Specification	Requirement
CPU	Modern x86 processor (minimum 2 cores or vCPUs). The use of multicore processors improves data processing performance and enables more tasks to be processed concurrently by the backup proxy.
Memory	2GB RAM plus 200MB for each concurrent task. The use of faster memory (DDR3) improves data processing performance.
Disk space	300MB.
Network	NetApp recommends 1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication. High latency and reasonably unstable WAN links are supported.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: <ul style="list-style-type: none"><li>• Microsoft Windows Server 2016</li><li>• Microsoft Windows Server 2012 R2</li><li>• Microsoft Windows Server 2012</li><li>• Microsoft Windows Server 2008 R2 SP1</li><li>• Microsoft Windows Server 2008 SP2</li><li>• Microsoft Windows Server 2003 SP2</li><li>• Microsoft Windows 10</li><li>• Microsoft Windows 8.x</li><li>• Microsoft Windows 7 SP1</li><li>• Microsoft Windows Vista SP2</li></ul>
Software	For a VMware vSphere 5.5 or later backup proxy server running on Microsoft Windows Server 2008 or earlier: Microsoft Visual C++ 2008 SP1 Redistributable Package (x64). The installation package can be downloaded from <a href="http://vee.am/runtime">http://vee.am/runtime</a> .

**Note:** To protect VMs that run on ESXi 5.5 and later, you must deploy backup proxies on machines that run a 64-bit version of Microsoft Windows. VMware Virtual Disk Development Kit (VDDK) 5.5 and later versions do not support 32-bit versions of Microsoft Windows.

## VMware Backup Repository Server

Table 3 lists the specifications and requirements of the VMware backup repository server.

Table 3) VMware backup repository server specifications.

Specification	Requirement
CPU	Modern x86 processor (minimum 2 cores or vCPUs). The use of multicore processors improves data processing performance and enables more tasks to be processed concurrently by the backup proxy.
Memory	4GB RAM plus 2GB RAM (32-bit OS) or 4GB RAM (64-bit OS) for each concurrent job. The use of more memory improves data processing performance for long chains of large backup files on backup repositories that run a 64-bit OS.
Disk space	200MB for Veeam Backup & Replication components and sufficient disk space to store backup files and replicas (NetApp recommends high-RPM drives and a RAID 10 configuration).
Network	NetApp recommends 1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication. High latency and reasonably unstable WAN links are supported.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1</li> <li>• Microsoft Windows Server 2008 SP2</li> <li>• Microsoft Windows Server 2003 SP2</li> <li>• Microsoft Windows 8.x</li> <li>• Microsoft Windows 7 SP1</li> <li>• Microsoft Windows Vista SP2</li> <li>• Linux: Secure Shell (SSH) and Perl are required (check the full list of required Perl modules at <a href="http://www.veeam.com/kb2007">www.veeam.com/kb2007</a>).</li> <li>• 64-bit edition of Linux must be able to run 32-bit programs; pure 64-bit Linux editions are not supported.</li> </ul>
Software	For a VMware vSphere 5.5 or later backup proxy server on Microsoft Windows Server 2008 or earlier: Microsoft Visual C++ 2008 SP1 Redistributable Package (x64). The installation package can be downloaded from <a href="http://vee.am/runtime">http://vee.am/runtime</a> .

## 5.2 Hyper-V System Requirements

### Veeam Backup Server (Hyper-V)

Table 4 lists the specifications and requirements of the Veeam backup server with Microsoft Hyper-V.

Table 4) Veeam proxy server specifications (Hyper-V).

Specification	Requirement
CPU	Modern x86 processor (minimum 2 cores or vCPUs).
Memory	4GB RAM plus 500MB RAM for each concurrent job.
Disk space	2GB for product installation and 10GB per 100 VMs for a guest file system catalog folder (persistent data).
Network	NetApp recommends 1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication. High latency and reasonably unstable WAN links are supported.
OS	<p>64-bit versions of the following operating systems are supported:</p> <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1</li> <li>• Microsoft Windows Server 2008 SP2</li> <li>• Microsoft Windows Server 2003 SP2</li> <li>• Microsoft Windows 10</li> <li>• Microsoft Windows 8.x</li> <li>• Microsoft Windows 7 SP1</li> </ul>
Software	<p>During setup, a system configuration check is performed to determine whether all prerequisite software is available on the machine in which you plan to install Veeam Backup &amp; Replication. If some of the required software components are missing, the setup wizard offers to install missing software automatically. This task refers to:</p> <ul style="list-style-type: none"> <li>• Microsoft .NET Framework 4.5.2</li> <li>• Microsoft Windows Installer 4.5</li> <li>• Microsoft SQL Server Management Objects</li> <li>• Microsoft SQL Server System CLR Types</li> <li>• Microsoft Visual C++ 2010 SP1 Redistributable Package</li> </ul> <p>The following software must be installed manually:</p> <ul style="list-style-type: none"> <li>• Microsoft PowerShell 2.0 or later (optional).</li> <li>• Firefox, Google Chrome, Microsoft Edge, or Microsoft Internet Explorer 10.0 or later.</li> <li>• System Center Virtual Machine Manager (VMM) 2016, 2012 R2, 2012, or 2008 R2 admin UI (optional, to be able to register the System Center Virtual Machine Manager [SCVMM] server with Veeam Backup &amp; Replication infrastructure).</li> <li>• Remote Desktop Protocol (RDP) client version 7.0 and later installed on the backup server (required to open the VM console during SureBackup recovery verification of Microsoft Hyper-V VMs). The RDP client is preinstalled on Microsoft Windows 7 and Windows Server 2008 R2 OS and later. You can download the RDP client from <a href="http://support.microsoft.com/kb/969084/en-us">http://support.microsoft.com/kb/969084/en-us</a>.</li> </ul>

SQL Server database	<p>Local or remote installation of the following versions of Microsoft SQL Server (both Full and Express Editions are supported):</p> <ul style="list-style-type: none"> <li>• Microsoft SQL Server 2016</li> <li>• Microsoft SQL Server 2014</li> <li>• Microsoft SQL Server 2012 (Microsoft SQL Server 2012 Express Edition is included in the setup)</li> <li>• Microsoft SQL Server 2008 R2</li> <li>• Microsoft SQL Server 2008</li> <li>• Microsoft SQL Server 2005</li> </ul>
---------------------	--

**Note:**

1. You might plan to back up VMs that are running Microsoft Windows Server 2012 R2 or Microsoft Windows Server 2016, and data deduplication is enabled for some VM volumes. In that case, NetApp recommends that you deploy Veeam Backup & Replication on a machine that is running Microsoft Windows Server 2012 R2 or Microsoft Windows Server 2016 with data deduplication enabled. Otherwise, some types of restore operations for these VMs (such as Microsoft Windows File Level Recovery) might fail.
2. Because of its limitations, Microsoft SQL Server Express Edition can be used only for evaluation or in a small-scale production environment. For environments with several VMs, you must install a fully functional commercial version of Microsoft SQL Server.

### Off-Host Backup Proxy Server (Hyper-V)

Table 5 lists the specifications and requirements of an off-host backup proxy server with Hyper-V.

Table 5) Off-host backup proxy server specifications (Hyper-V).

Specification	Requirement
CPU	Modern x86 processor (minimum 2 cores or vCPUs). The use of multicore processors improves data processing performance and enables more tasks to be processed concurrently by the backup proxy.
Memory	2GB RAM plus 200MB for each concurrent task. The use of faster memory (DDR3) improves data processing performance.
Disk space	300MB.
Network	NetApp recommends 1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication. High latency and reasonably unstable WAN links are supported.
OS	<ul style="list-style-type: none"> <li>• Microsoft Windows Server 2016 with Hyper-V role enabled</li> <li>• Microsoft Windows Server 2012 R2 with Hyper-V role enabled</li> <li>• Microsoft Windows Server 2012 with Hyper-V role enabled</li> <li>• Microsoft Windows Server 2008 R2 SP1 with Hyper-V role enabled</li> </ul>
Software	A Volume Shadow Copy Service (VSS) hardware provider that supports transportable shadow copies. The VSS hardware provider is typically distributed as part of the client components that are supplied by the storage vendor.

### Backup Repository Server (Hyper-V)

Table 6 lists the specifications and requirements of a backup repository server with Hyper-V.

**Table 6) Backup repository server specifications (Hyper-V).**

Specification	Requirement
CPU	x86 processor (NetApp recommends x86-64).
Memory	4GB RAM plus 2GB RAM (32-bit OS) or 4GB RAM (64-bit OS) for each concurrent job. The use of more memory improves data processing performance for long chains of large backup files on backup repositories that run 64-bit OS.
Disk space	200MB for Veeam Backup & Replication components and sufficient disk space to store backup files and replicas (NetApp recommends high-RPM drives and a RAID 10 configuration).
Network	NetApp recommends 1Gbps LAN for on-site backup and replication and 1Mbps WAN for off-site backup and replication. High latency and reasonably unstable WAN links are supported.
OS	<p>Both 32-bit and 64-bit (recommended) versions of the following operating systems are supported:</p> <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1</li> <li>• Microsoft Windows Server 2008 SP2</li> <li>• Microsoft Windows Server 2003 SP2</li> <li>• Microsoft Windows 10</li> <li>• Microsoft Windows 8.x</li> <li>• Microsoft Windows 7 SP1</li> <li>• Microsoft Windows Vista SP2</li> <li>• Linux: SSH and Perl are required (check the full list of required Perl modules at <a href="http://www.veeam.com/kb2007">www.veeam.com/kb2007</a>).</li> <li>• 64-bit edition of Linux must be able to run 32-bit programs; pure 64-bit Linux editions are not supported.</li> </ul>

### 5.3 Veeam Sizing Requirements

For Veeam sizing requirements, see Table 7.

**Table 7) Veeam sizing requirements.**

Veeam Sizing Requirements		
Veeam backup proxy	Virtual: 1 vCPU for every 200 VMs (recommended minimum: 2 vCPUs), 4GB RAM, 300MB disk space for installed components, and VMXNET3 network interface card (NIC). No CPU or memory reservations are required.	Physical: 1 CPU core for every 250 VMs, 2GB of RAM for each concurrent virtual machine disk (VMDK) backup

Veeam Sizing Requirements		
Veeam backup server	4GB of RAM plus 500MB of RAM for each concurrent job. Disk space: 2GB for product, plus 10GB per 100 VMs for a guest file system catalog and at least 10GB for a VM recovery cache folder. No CPU or memory reservations are required. Additional sizing considerations should be applied if the back-end SQL Server is deployed on this server. (For additional guidance, see the Veeam documentation.)	
Veeam repository	The Veeam repository can be colocated with the backup server role for small deployments or on a dedicated server (physical or virtual Windows or Linux server, NAS device, or dedicated backup appliance). The repository requires sufficient free space to store all backup job data. If a VM is used, NetApp recommends using the VMware Paravirtual SCSI (PVSCSI) controller for the disk or disks that store backup data. The use of vSphere 5.5 enables the use of disks larger than 2TB, which might be advantageous for a repository server.	

## 6 Veeam Backup & Replication 9.5 Design Considerations

### 6.1 Direct SAN Access Mode

Veeam proxy servers directly connect to the storage fabric because it is the fastest way to perform backups with Backup & Replication. That approach allows a LAN-free backup and provides the possibility to leverage storage Snapshot copies on supported storage systems such as NetApp FAS. Fast backup speed is critical for low RPOs.

An often-discussed point about Direct SAN Access for the Veeam proxy is the possibility that the Veeam proxy can mount VMware Virtual Machine File System (VMFS) LUNs and write a Windows signature onto the LUN. Doing so results in an inaccessible VMFS, which requires VMware support to fix the issue. Veeam Knowledge Base (KB) article #1446 provides instructions on how to configure SAN access for use with Veeam Backup & Replication.

Although Microsoft Windows 2003 had automount enabled by default, the SAN policy in the Windows Server 2008 R2 Enterprise and Datacenter editions is set to Offline Shared by default. However, even if an administrator changed the policy to Online All, Veeam reverts the policy back to Offline Shared. The result is that Windows does not mount and resignature VMFS LUNs unless an administrator changes this setting again or manually mounts VMFS LUNs while ignoring Windows warnings. NetApp recommends the Direct SAN Access transport mode for VMs whose disks are on shared VMFS SAN LUNs that are connected to ESX(i) hosts through FC or iSCSI.

In the Direct SAN Access transport mode, Veeam Backup & Replication leverages VMware VADP to transport VM data directly from and to FC and iSCSI storage over the SAN. VM data travels over the SAN, bypassing ESX(i) hosts and the LAN. The Direct SAN Access transport method provides the fastest data transfer speed and produces no load on the production network.

You can use the Direct SAN Access transport mode for all operations in which the backup proxy is engaged, including:

- Backup
- Replication
- Full VM restore
- VM disk restore



- Replica failback
- Quick migration

## Requirements for the Direct SAN Access Transport Mode

To use the Direct SAN Access transport mode, you must meet the following requirements:

- The backup proxy that uses the Direct SAN Access transport mode must have direct access to the production storage through a hardware or software host bus adapter (HBA). If a direct SAN connection is not configured or is not available when a job or a task starts, the job or the task fails.
- For restore operations, the backup proxy must have write access to LUNs where the VM disks are located.

## Limitations for the Direct SAN Access Transport Mode

The Direct SAN Access transport mode has the following limitations:

- The Direct SAN Access transport mode is not supported for VMs that reside on a VSAN. You can use the Virtual Appliance or Network transport mode to process such VMs. For details about VSAN restrictions, see the VDDK 5.5 Release Notes.
- If at least one VM disk is on a VVol, you cannot use the Direct SAN Access mode.
- You can use the Direct SAN Access transport mode only for the initial run of the replication job. For subsequent replication job runs, Veeam Backup & Replication uses the Virtual Appliance or Network transport mode.
- You can use the Direct SAN Access transport mode to restore only thick VM disks.
- Because of VMware limitations, you cannot use the Direct SAN Access transport mode for incremental restore. You must either disable changed block tracking (CBT) for VM virtual disks for the duration of the restore process or select another transport mode for incremental restore.

For VMware vSphere 5.5 and later, IDE and SATA disks can be processed in the Direct SAN Access transport mode.

For VMware vSphere 5.1 and earlier, you should consider the following:

- IDE disks can be backed up in the Direct SAN Access transport mode. However, restore of IDE disks in the Direct SAN Access transport mode is not supported.
- If a VM disk fails to be processed in the Direct SAN Access transport mode, Veeam Backup & Replication does not fail over to the Network mode.
- If a VM has some disks that cannot be processed in the Direct SAN Access transport mode, Veeam Backup & Replication uses the Network mode for VM disk processing.

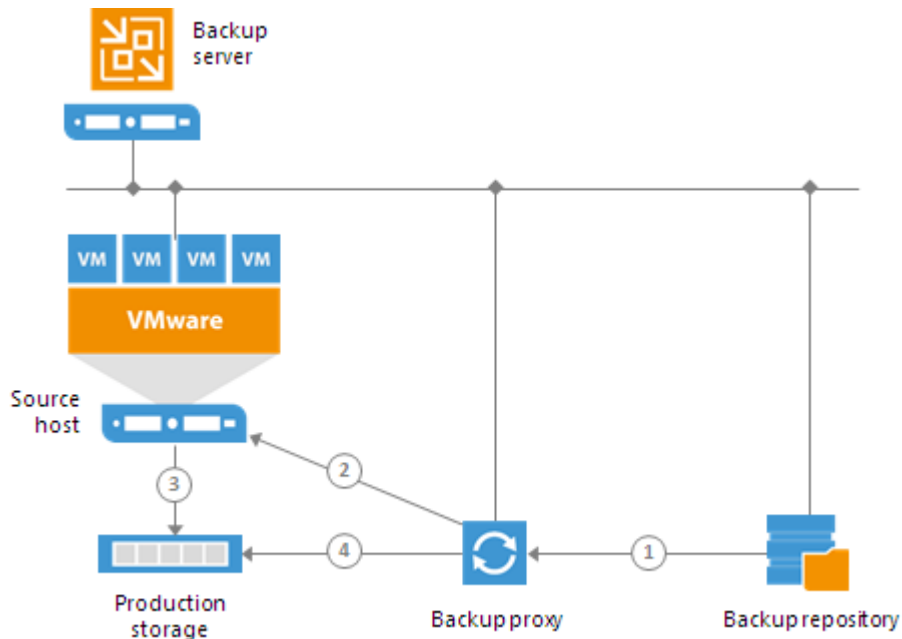
## 6.2 Data Restore in Direct SAN Access Mode

Data restore in the Direct SAN Access transport mode includes the following steps (see Figure 15):

1. The backup proxy retrieves data blocks from the backup repository or a datastore in the target site.
2. The backup proxy sends a request to the ESX(i) host on the source site to restore data to a necessary datastore.
3. The ESX(i) host on the source site allocates space on the datastore.
4. Data blocks that are obtained from the backup proxy are written to the datastore.



Figure 15) Data restore in Direct SAN Access mode.



You can use the Direct SAN Access transport mode to restore VMs only with thick disks. Before VM data is restored, the ESX(i) host must allocate space for the restored VM disk on the datastore:

- When thick disks are restored, the ESX(i) host allocates space on the disk before writing VM data.
- When thin disks are restored, the ESX(i) host attempts to allocate space dynamically as requests for data block restores are received.

As a result, restoring thin disks involves extra allocation overhead when compared with restoring thick disks, which results in decreased performance.

To restore VMs with thin disks, you can use the Virtual Appliance mode or the Network mode. If you plan to process a VM that has both thin and thick disks, select the Direct SAN Access transport mode and choose to fail over to the Network mode if the SAN becomes inaccessible. In that case, Veeam Backup & Replication uses the Direct SAN Access transport mode to restore thick disks and uses the Network transport mode to restore thin disks. Alternatively, you can restore all VM disks as thick.

### 6.3 Off-Host Backup (Hyper-V)

In off-host backup mode, backup processing is shifted from the source Hyper-V host to a dedicated machine: an off-host backup proxy. The off-host backup proxy acts as a data mover. The Veeam transport service that runs on it retrieves VM data from the source datastore, processes it, and transfers it to the destination. This type of backup does not impose load on the Hyper-V host; while resource-intensive backup operations are performed on the off-host backup proxy, production hosts remain unaffected.

To perform off-host backup, Veeam Backup & Replication uses transportable shadow copies. Transportable shadow copy technology lets you create a snapshot of a data volume on one server and import, or mount, it onto another server within the same subsystem (SAN) for backup and other purposes. The transport process is accomplished in a few minutes, regardless of the amount of data. The process is performed at the SAN storage layer, so it does not affect host CPU usage or network performance.

To perform off-host backup, you must meet the following requirements:

1. You must configure an off-host backup proxy. You can assign the role of an off-host backup proxy only to one of the following:
  - A physical Microsoft Windows 2008 Server R2 machine with the Hyper-V role enabled
  - A Windows Server 2012 machine with the Hyper-V role enabled
  - A Windows Server 2012 R2 machine with the Hyper-V role enabled
  - A Windows Server 2016 machine with the Hyper-V role enabled

The version of the Hyper-V host and off-host backup proxy must be the same. For example, you use a Microsoft Windows 2008 Server R2 machine with the Hyper-V role enabled as a Hyper-V host. In that case, you should deploy the off-host backup proxy on a Microsoft Windows 2008 Server R2 machine with the Hyper-V role enabled.

2. In the properties of a backup or replication job, you must select the off-host backup method. If necessary, you can point the job to a specific proxy.
3. The source Hyper-V host and the off-host backup proxy must be connected through a SAN configuration to the shared storage.
4. To create and manage volume shadow copies on the shared storage, you must install and properly configure a VSS hardware provider that supports transportable shadow copies on the off-host proxy and the Hyper-V host. Typically, when configuring a VSS hardware provider, you must specify a server that controls the LUN and disk array credentials to provide access to the array.

The VSS hardware provider is usually distributed as part of the client components that are supplied by the storage vendor. Any VSS hardware provider that is certified by Microsoft is supported. Some storage vendors might require additional software and licensing to work with transportable shadow copies.

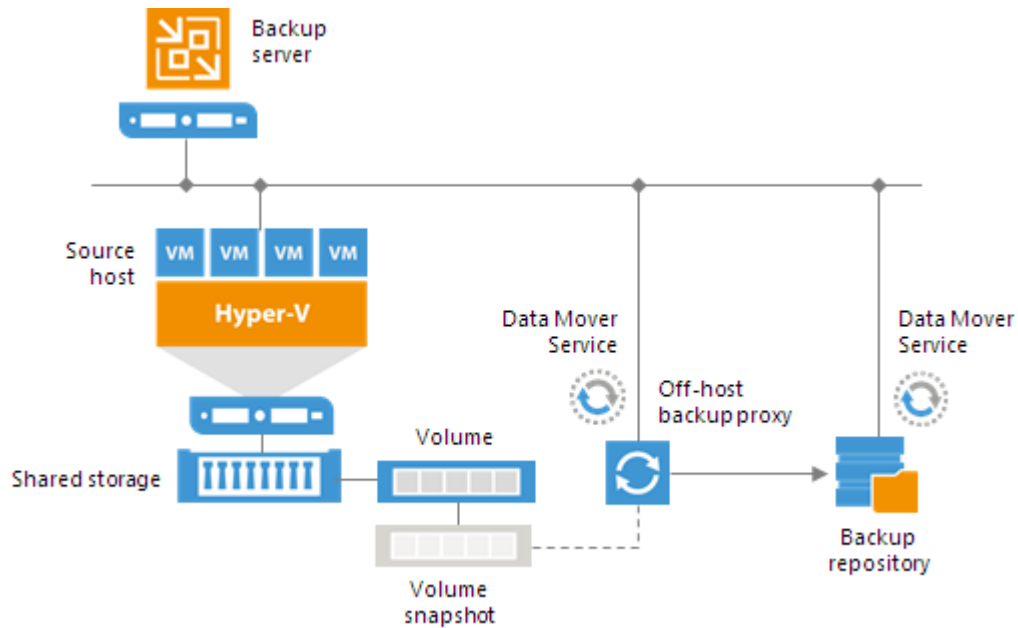
5. If you back up VMs whose disks reside on a CSV with data deduplication enabled, make sure that you:
  - a. Use a Microsoft Windows 2012 R2 or a Microsoft Windows 2016 machine as an off-host backup proxy.
  - b. Enable data deduplication on this off-host backup proxy.

Otherwise, off-host backup fails.

The off-host backup process (Figure 16) includes the following steps:

6. Veeam Backup & Replication triggers a snapshot of the necessary volume on the production Hyper-V host.
7. The created snapshot is split from the production Hyper-V server and is mounted to the off-host backup proxy.
8. The Veeam transport service that runs on a backup proxy uses the mounted volume snapshot to retrieve VM data. The VM data is processed on the proxy server and is copied to the destination.
9. When the backup process is complete, the snapshot is dismounted from the off-host backup proxy and is deleted from the SAN.

Figure 16) Process of off-host backup mode.



**Note:** If you plan to perform off-host backup for a Hyper-V cluster with CSV, deploy an off-host backup proxy on a host that is **not** part of a Hyper-V cluster.

When a volume snapshot is created, this snapshot has the same LUN signature as the original volume. The Microsoft Cluster Service does not support LUNs with duplicate signatures and partition layout. For this reason, volume snapshots must be transported to an off-host backup proxy outside the cluster. If the off-host backup proxy is deployed on a node of a Hyper-V cluster, a duplicate LUN signature is generated, and the cluster fails during backup or replication.

## Off-Host Backup Proxy

By default, when you perform backup, replication, or VM copy jobs in the Hyper-V environment, VM data is processed directly on the source Hyper-V host where the VMs reside. The data is then moved to the target, bypassing the Veeam backup server.

VM data processing can produce unwanted overhead on the production Hyper-V host and can affect the performance of the VMs that run on this host. To take data processing off the production Hyper-V host, use the off-host backup mode.

The off-host mode shifts the backup and replication load to a dedicated machine: an off-host backup proxy. The off-host backup proxy functions as a data mover that retrieves VM data from the source datastore, processes it, and transfers it to the destination.

The machine that performs the role of an off-host backup proxy must meet the following requirements:

- The role can be assigned only to one of the following:
  - A physical Microsoft Windows 2008 Server R2 machine with the Hyper-V role enabled
  - A Windows Server 2012 machine with the Hyper-V role enabled
  - A Windows Server 2012 R2 machine with the Hyper-V role enabled
  - A Windows Server 2016 machine with the Hyper-V role enabled

- The off-host backup proxy must have access to the shared storage that hosts the VMs to be backed up, replicated, or copied.
- To create and manage volume shadow copies on the shared storage, you must install a VSS hardware provider that supports transportable shadow copies on the off-host proxy and the Hyper-V host. The VSS hardware provider is usually distributed as part of the client components that are supplied by the storage vendor.

When you assign the role of an off-host backup proxy to the selected machine, Veeam Backup & Replication automatically installs on it the lightweight components and services that are required for backup proxy functioning. Unlike the Veeam backup server, backup proxies do not require a dedicated SQL Server database; all the settings are stored centrally, within the SQL Server database that the Veeam backup server uses.

To enable a Hyper-V host or a Windows machine to act as an off-host backup proxy, Veeam Backup & Replication installs the following services on it:

- Veeam Installer Service is an auxiliary service that is installed and started on any Windows (or Hyper-V) server after it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system and installs and upgrades necessary components and services.
- Veeam Transport is responsible for deploying and coordinating executable modules that act as data movers and that perform the main job activities on behalf of Veeam Backup & Replication. These activities include performing data deduplication, compression, and so on.
- Veeam Hyper-V Integration Service is responsible for communicating with the VSS framework during backup, replication, and other jobs and for performing recovery tasks. The service also deploys a driver that handles changed block tracking for Hyper-V.

## 6.4 Veeam Backup & Replication 9.5 Components

Table 8 describes each component in Veeam Backup & Replication 9.5.

Table 8) Veeam Backup & Replication 9.5 components and their description.

Component	Description
Veeam backup server	<p>The Veeam backup server is a Windows-based physical or virtual machine on which Veeam Backup &amp; Replication is installed. It is the core component in the backup infrastructure that fills the role of the configuration and control center. The Veeam backup server performs all types of administrative activities:</p> <ul style="list-style-type: none"> <li>• Coordinates backup, replication, recovery verification, and restore tasks</li> <li>• Controls job scheduling and resource allocation</li> <li>• Is used to set up and manage backup infrastructure components and to specify global settings for the backup infrastructure</li> </ul> <p>In addition to its primary functions, a newly deployed Veeam backup server also performs the roles of the default backup proxy and the backup repository (it manages data handling and data storing tasks).</p>
Veeam Backup Service	<p>This Windows service coordinates all the operations that are performed by Veeam Backup &amp; Replication, such as backup, replication, recovery verification, and restore tasks. The Veeam Backup Service runs under the local system account or the account that has the local administrator permissions on the Veeam backup server.</p>
Veeam Backup Shell	<p>The Veeam Backup Shell provides the application UI and enables user access to the application's functionality.</p>

Component	Description
Veeam Backup Catalog Service	This Windows service manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog, a folder on the Veeam backup server. The Veeam Backup Catalog Service running on the Veeam backup server works with search components that are installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.
Veeam Backup SQL Server database	This database is used by Veeam Backup Service, Veeam Backup Shell, and Veeam Backup Catalog Service to store data about the backup infrastructure, jobs, sessions, and so on. The database instance can be on a SQL Server installed either locally (on the same machine where the Veeam backup server runs) or remotely.
Veeam Backup PowerShell snap-in	This snap-in is an extension for Microsoft Windows PowerShell 2.0. Veeam Backup PowerShell adds a set of cmdlets to enable you to perform backup, replication, and recovery tasks through the PowerShell CLI or run custom scripts to fully automate Veeam Backup & Replication operation.
Backup proxy services	In addition to dedicated services, the Veeam backup server runs a set of Data Mover Services.
Backup proxy	A backup proxy is an architecture component that sits between the data source and the target and is used to process jobs and deliver backup traffic. The backup proxy retrieves VM data from the production storage. It also compresses the retrieved data and sends it to the backup repository (for example, if you run a backup job) or to another backup proxy (for example, if you run a replication job). As the data handling task is assigned to the backup proxy, the Veeam backup server becomes the point of control for dispatching jobs to proxy servers.
Backup repository	A backup repository is a location that Veeam Backup & Replication jobs use to store backup files, copies of VMs, and metadata for replicated VMs. Technically, a backup repository is a folder on the backup storage. By assigning different repositories to jobs and by limiting the number of parallel jobs for each one, you can balance the load across your backup infrastructure.

## 7 NetApp E-Series Volume Configuration Guidelines

### 7.1 NetApp E-Series Storage Configuration Guidelines for Backup & Replication 9.5 Backup Repositories

NetApp recommends avoiding thin volumes while setting up NetApp E-Series with Veeam.

For **optimal performance**, NetApp recommends that you follow these guidelines:

- Use RAID 6 (8+2) volume groups.
- Create multiple volume groups, having an even number when possible so that you can achieve balance between owning controllers. Make sure of drawer loss protection (DLP) when applicable to reduce the risk of data unavailability and/or data loss.
- Create a single standard (not thin) volume per volume group.

- When creating the volumes, select a 512KB segment size to match the 512KB transfer size that Veeam Backup & Replication presents as a sequential write to the system after sequencing I/O.
- Run multiple backup jobs to each repository.

**Note:** The preceding setup can be configured without any hot spares too, to achieve better performance in case of drive failure. A caveat is that data loss might occur if more than two drives fail at once.

For a **large configuration**, NetApp recommends that you follow these guidelines:

- Use dynamic disk pools (DDP) technology to maximize ease of use and for fast rebuild times.
- Create multiple dynamic disk pools (DDP), having an even number when possible so that you can achieve balance between owning controllers. Make sure of DLP when applicable to reduce the risk of data unavailability and/or data loss.
- Create a single standard (not thin) volume per DDP.
- When creating the volumes with DDP, the default segment size is 128KB. No additional selection is required.
- Run multiple backup jobs to each repository.

**Note:** RAID 10 was also tested. It can also be considered as an option for volume configuration, but it is not recommended due to a single parity and capacity penalty. Also, write throughput performance of RAID 6 and DDP is better in comparison to RAID 10.

## 8 NetApp E-Series and EF-Series Host Configuration Guidelines

### 8.1 Host Connectivity

NetApp E-Series arrays provide multiple options for connectivity.

#### Host Connectivity for E2800

The E2800 controller has the following base hardware features:

- Dual Ethernet ports for management-related activities
- Either two optical FC/iSCSI or two RJ-45 iSCSI baseboard ports for host connection
- SAS drive expansion ports to attach expansion drive shelves

**Note:** Adding an optional HIC is only needed if you want to use the SAS protocol, need more than two host ports per controller, or want to use both FC and iSCSI protocols.

#### Host Connectivity for E5700 and EF570

The E5700 controller has the following base hardware features:

- Dual Ethernet ports for management-related activities
- Dual optical 16Gbps FC or 10Gbps iSCSI baseboard ports for host connection
- Dual 12Gb SAS drive expansion ports to attach expansion drive shelves

**Note:** Adding the new optional HICs for the E5700 controller provides faster host interfaces for iSCSI, FC, and IB.

#### Host Connectivity for E5600 and EF560

The E5600 controller has the following base hardware features:

- Dual Ethernet ports for management-related activities
- SAS, FC, iSCSI, or IB ports for host connection

- SAS drive expansion ports to attach expansion drive

**Note:** One of the host I/O port options must be ordered when the controller is purchased.

Table 9 provides an overview about the different host I/O interfaces for E-Series controllers.

Table 9) Optional host I/O for E-Series.

Controller	Optional Host I/O
<b>E2800</b> <ul style="list-style-type: none"> <li>• Controllers must match</li> <li>• The Base-T iSCSI onboard controller can only use the 2-port Base-T HIC</li> <li>• A software feature pack can be applied to convert the FC HIC ports to iSCSI or to convert iSCSI HIC ports to FC</li> </ul>	2-port 10Gb iSCSI (Base-T) per controller
	2-port 12Gb SAS (wide-port) per controller
	4-port 12Gb SAS (wide-port) per controller
	2-port 10Gb iSCSI (optical)/16Gb FC per controller
	4-port 10Gb iSCSI (optical)/16Gb FC per controller
<b>E5700 and EF570</b> <ul style="list-style-type: none"> <li>• Controllers must match</li> <li>• A software feature pack* can be applied to convert the FC base ports to iSCSI or to convert between the three IB protocols on the optional 100Gb IB HIC</li> </ul>	2-port 100Gb IB (iSER, SRP, or NVMe-oF [IB] based on feature pack installed*)
	4-port 12Gb SAS (wide-port): uses mini-SAS cables
	4-port 32Gb FC (see the Hardware Universe for SFP details: 32Gb SFP supports 32Gbps and 16Gbps but not 8Gbps or slower)
	4-port 25Gb iSCSI (see the Hardware Universe for SFP details: must set port speed for 10Gbps or 25Gbps using SANtricity System Manager: the 25Gbps SFP does work for 25Gb and 10Gb speeds, but the port does not change speeds automatically)
<b>E5600 and EF560</b> <ul style="list-style-type: none"> <li>• Controllers must match</li> <li>• A software feature pack can be applied to convert the 4-port FC and iSCSI HICs to FC-iSCSI split mode</li> </ul>	8-port 12Gb SAS (4-port 12Gb SAS [wide-port] per E5600 controller)
	8-port 16Gb FC (4-port 16Gb FC per E5600 controller)
	8-port 10Gb iSCSI (4-port 10Gb iSCSI per E5600 controller)
	4-port 56Gb IB (2-port 56Gb IB per E5600 controller)

Your optimal design is based on your existing environment and how the data flows from primary storage through the Veeam server to its final destination on the E-Series array. One design approach is to leverage 12Gb SAS for connectivity from the Veeam server to the E-Series array. This approach provides the target backup repository (E-Series) to Veeam as a direct connection. Having a direct connection to the E-Series array prevents having to write the backup file across an existing network.

This configuration might not be ideal for larger environments in which Veeam's distributed architecture is implemented and in which multiple proxy servers process backup data. Having a dedicated network for backup targets might make more sense in those cases. Either way, the options are there for any environment, and Veeam provides a bottleneck detector to help optimize the backup data flow as you progress through your implementation.



## 8.2 NetApp E-Series Storage Host Mapping Configuration for Direct SAN Access

On the E-Series or EF-Series system that hosts the virtual environment's storage, verify that the Veeam backup server is associated with the same host group as the host of the virtual environment. See the example shown for NetApp SANtricity in Figure 17.

Host\_Cluster\_Veeam is the host cluster that has the ESXi Server (Host\_ESXi\_Server) and the Veeam Server (Host\_Veeam\_Server). As seen, all 16 volumes are shared between both the host members.

Figure 17) Backup server mapping.

Name	Type	Associated Objects	Total Assigned Volumes	Reported Capacity (GiB)	Host Type	Edit
– Default Cluster	Default Cluster	1 Host Cluster(s) 2 Host(s)	16	23600.00	Windows	
– Host_Cluster_Veeam	Cluster	Default Cluster 2 Host(s)	16	23600.00	Windows	
Host_ESXi_Server	Host Member	Host_Cluster_Veeam	16	23600.00	Windows	
Host_Veeam_Server	Host Member	Host_Cluster_Veeam	16	23600.00	Windows	

Total rows: 4 |

As a result of hosting both the virtual environment and the backup repository on the E-Series or the EF-Series arrays, Veeam Backup & Replication has resource location awareness in the environment. The resulting data transfers take place over the SAN.

## 9 CPU and Memory Sizing for Veeam and NetApp E-Series

This section covers the CPU and memory considerations when setting up Veeam with NetApp E-Series. Three components need to be considered during this sizing: backup server, repository, and proxy.

The backup server is the core component. Features and component requirements affect your decision how to install the backup server (for example, one data center or multiple locations). It could mean that you choose to install additional backup servers or services in remote locations to optimize the data streams.

Sizing with Veeam is cumulative in respect to configurations. If you want to create an all-in-one appliance (appliance model), add all the resource requirements together (CPU + memory) for backup server, repository, and proxy to understand what in total you will need. The same goes if you only want to have proxy and repository in one host and server independently.

### 9.1 CPU and Memory Sizing of Backup Server

Recommended Veeam backup server configuration is 1 CPU core (physical or virtual) and 5GB RAM per 10 concurrently running jobs. Concurrent jobs include any running backup or replication jobs as well as any job with a continuous schedule such as backup copy jobs and tape jobs. (Concurrent job is a job that processes a single VM with a single virtual disk.) The minimum recommendation is 2 CPU cores and 8GB RAM.

It is recommended to group multiple virtual machines into a single job for better efficiency and resource usage. With default configuration, it is recommended to configure at minimum a VM with a single disk up to around 30 VMs per job.



## 9.2 CPU and Memory Sizing of Backup Proxy

Getting the right amount of processing power is essential to achieving the RTPO defined by the business.

It is a best practice to plan for 1 physical core or 1 vCPU and 2GB of RAM for each of the tasks. A task processes 1 VM disk at a time, and CPU/RAM resources are used for inline data deduplication, compression, encryption and other features that are running on the proxy itself. A task as described earlier should ideally have around 30 VMs.

## 9.3 CPU and Memory Sizing of Backup Repository

In midsized or enterprise environments, the recommended amount of CPU for a repository is 1 core per concurrent job and 4GB RAM per core that processes data on a repository server. At least 2 cores allow for the operating system to be more responsive.

### Example for CPU and Memory Sizing

Consider an all-in-one configuration (physical) set up for 10TB of data or 100 VMs:

- Total number of jobs is 100 VMs; 30 VMs per job = 4 (rounding it up to a higher number)
- CPU and memory for backup server = 1 CPU core and 5GB of RAM (refer to CPU and Memory Sizing of Backup Server)
- CPU and memory for backup proxy = 4 CPU core and 8GB of RAM (refer to CPU and Memory Sizing of Backup Proxy)
- CPU and memory for backup repository = 4 CPU core and 16GB of RAM (refer to CPU and Memory Sizing of Backup Repository)

Total is sum of the three, which gives us 10 CPU cores (rounded up to 10, because 9 is not possible) and around 32GB of RAM (including a couple of gigabytes for operating system and so on, hence 32GB instead of 29GB).

**Note:** Refer to <https://bp.veeam.expert/> for further details.

## 10 Capacity Sizing for Veeam and NetApp E-Series

This section describes the procedure and parameters to be considered while estimating the amount of disk space required.

When estimating the amount of required disk space, you should know the following:

- Total size of VMs being backed up
- Frequency of backups
- Retention period for backups
- Whether jobs use forward or reverse incremental

Also, when testing is not possible beforehand, you should make assumptions about compression and deduplication ratios, change rates, and other factors. The following figures are typical for most deployments; however, it is important to understand the specific environment to figure out possible exceptions:

- Data reduction because of compression and deduplication is usually 2:1 or more; it's common to see 3:1 or better, but you should always be conservative when estimating required space.

- Typical daily change rate is between 2% and 5% in a midsize or enterprise environment. This rate can greatly vary among servers; some servers show much higher values. If possible, run monitoring tools such as Veeam ONE to have a better understanding of the real change rate values.
- Include additional space for one-off full backups.
- Include additional space for backup chain transformation (forward forever incremental, reverse incremental): at least the size of a full backup multiplied by 1.25x.

## 10.1 Repository Sizing Tool

A repository sizing tool that can be used for estimation is available at <http://vee.am/rps>.

**Note:** This tool is not officially supported by Veeam, and it should be used as is, but it's nonetheless heavily used by Veeam architects and regularly updated.

Figure 18) Capacity planning simulator.

# The Restore Point Simulator

Current version : [0.3.3](#)






Feedback via [@tdewin](#) or on [GitHub](#)

RPS heavily relies on some opensource [javascript frameworks](#)

## Quick Presets

Forever Incremental ▾

## Configuration

Style	Incremental ▾	
Used Size GB	1000	
Retention Points	14	
Change Rate	10% Conservative ▾	
Data left after reduction	50% (100GB > 50GB) 2x Conservative ▾	
Interval	Daily ▾	
Time Growth Simulation <input type="checkbox"/>	1 Year ▾	10% ▾ 

## Incremental Specific

Synthetic  MO  TU  WE  TH  FR  SA  SU

Active Full Weekly  MO  TU  WE  TH  FR  SA  SU

Active Full Monthly  Jan  Feb  Mar  Apr  May  Jun  
 Jul  Aug  Sep  Oct  Nov  Dec

## Run

Manual Run  Export  Canvas (experimental)

**Simulate**

Figure 18 shows an overview of the capacity planning simulator by Veeam. Examples of the Ccapacity planning simulator shown later can be used to understand how this tool can be used.

## Examples of the Ccapacity planning simulator

Consider the environment to have the following: -

- 10 VM's, 100 GB each, average used space is around 80 GB.

- The average compression and deduplication rate of 50% is assumed as mentioned in Capacity Sizing for Veeam and NetApp E-Series.
- The change rate is considered to be 5% daily.

**Example 1: Reverse Incremental, Daily Backup with 30-day retention.**

Figure 19 gives us an idea how the simulator would look if all the parameters related to example 1 were input.

Figure 19) Parameter input of example 1 in Veeam simulator.

## The Restore Point Simulator

Current version : [0.3.3](#)  
 Feedback via [@tdewin](#) or on [GitHub](#)  
 RPS heavily relies on some opensource [javascript frameworks](#)

### Quick Presets

Reverse Incremental ▼

### Configuration

Style	Reverse	▼	
Used Size GB	800		
Retention Points	30		
Change Rate	5% Optimistic	▼	
Data left after reduction	50% (100GB > 50GB) 2x Conservative	▼	
Interval	Daily	▼	
Time Growth Simulation <input type="checkbox"/>	1 Year	▼	10%

### Reverse Specific

Active Full Weekly  MO  TU  WE  TH  FR  SA  SU

Active Full Monthly  Jan  Feb  Mar  Apr  May  Jun  
 Jul  Aug  Sep  Oct  Nov  Dec

### Run

Manual Run  Export  Canvas (experimental) Simulate

Results can be seen after we click Simulate.

Figure 20 shows the results of example 1. Required space is around 980GB for this example, along with 420GB of reserve space; total space is ~1400GB.

**Note:** Results can be exported as a link by clicking the Export checkbox while downloading as a .png file by using the Canvas checkbox.

Figure 20) Results of example 1 using Veeam simulator.

**Result**

Retention		File	Size	Modify Date	Point Date
30		reverse.vrb	<u>20 GB</u>	2017-12-21 Th 22	2017-12-20 We 22
29		reverse.vrb	<u>20 GB</u>	2017-12-22 Fr 22	2017-12-21 Th 22
28		reverse.vrb	<u>20 GB</u>	2017-12-23 Sa 22	2017-12-22 Fr 22
27		reverse.vrb	<u>20 GB</u>	2017-12-24 Su 22	2017-12-23 Sa 22
26		reverse.vrb	<u>20 GB</u>	2017-12-25 Mo 22	2017-12-24 Su 22
25		reverse.vrb	<u>20 GB</u>	2017-12-26 Tu 22	2017-12-25 Mo 22
24		reverse.vrb	<u>20 GB</u>	2017-12-27 We 22	2017-12-26 Tu 22
23		reverse.vrb	<u>20 GB</u>	2017-12-28 Th 22	2017-12-27 We 22
22		reverse.vrb	<u>20 GB</u>	2017-12-29 Fr 22	2017-12-28 Th 22
21		reverse.vrb	<u>20 GB</u>	2017-12-30 Sa 22	2017-12-29 Fr 22
20		reverse.vrb	<u>20 GB</u>	2017-12-31 Su 22	2017-12-30 Sa 22
19		reverse.vrb	<u>20 GB</u>	2018-01-01 Mo 22	2017-12-31 Su 22
18		reverse.vrb	<u>20 GB</u>	2018-01-02 Tu 22	2018-01-01 Mo 22
17		reverse.vrb	<u>20 GB</u>	2018-01-03 We 22	2018-01-02 Tu 22
16		reverse.vrb	<u>20 GB</u>	2018-01-04 Th 22	2018-01-03 We 22
15		reverse.vrb	<u>20 GB</u>	2018-01-05 Fr 22	2018-01-04 Th 22
14		reverse.vrb	<u>20 GB</u>	2018-01-06 Sa 22	2018-01-05 Fr 22
13		reverse.vrb	<u>20 GB</u>	2018-01-07 Su 22	2018-01-06 Sa 22
12		reverse.vrb	<u>20 GB</u>	2018-01-08 Mo 22	2018-01-07 Su 22
11		reverse.vrb	<u>20 GB</u>	2018-01-09 Tu 22	2018-01-08 Mo 22
10		reverse.vrb	<u>20 GB</u>	2018-01-10 We 22	2018-01-09 Tu 22
9		reverse.vrb	<u>20 GB</u>	2018-01-11 Th 22	2018-01-10 We 22
8		reverse.vrb	<u>20 GB</u>	2018-01-12 Fr 22	2018-01-11 Th 22
7		reverse.vrb	<u>20 GB</u>	2018-01-13 Sa 22	2018-01-12 Fr 22
6		reverse.vrb	<u>20 GB</u>	2018-01-14 Su 22	2018-01-13 Sa 22
5		reverse.vrb	<u>20 GB</u>	2018-01-15 Mo 22	2018-01-14 Su 22
4		reverse.vrb	<u>20 GB</u>	2018-01-16 Tu 22	2018-01-15 Mo 22
3		reverse.vrb	<u>20 GB</u>	2018-01-17 We 22	2018-01-16 Tu 22
2		reverse.vrb	<u>20 GB</u>	2018-01-18 Th 22	2018-01-17 We 22
1		full.vbk	<u>400 GB</u>	2018-01-18 Th 22	2018-01-18 Th 22
			980 GB		
		Work Space	+420 GB		
			<u>1400 GB</u>		

**Example 2: Forward Incremental, Daily Backup with 30-Day Retention Along with a Weekly Full**

Figure 21 gives us an idea how the simulator would look if all the parameters related to example 1 were input.

Figure 21) Parameter input of example 2 in Veeam simulator.

## The Restore Point Simulator

Current version : 0.3.3






Feedback via @tdewin or on [GitHub](#)

RPS heavily relies on some opensource [javascript frameworks](#)

### Quick Presets

Incremental Weekly Active Full

### Configuration

Style	Incremental	
Used Size GB	800	
Retention Points	30	
Change Rate	5% Optimistic	
Data left after reduction	50% (100GB > 50GB) 2x Conservative	
Interval	Daily	
Time Growth Simulation	<input type="checkbox"/>	1 Year 10% 

### Incremental Specific

Synthetic  MO  TU  WE  TH  FR  SA  SU

Active Full Weekly  MO  TU  WE  TH  FR  SA  SU

Active Full Monthly  Jan  Feb  Mar  Apr  May  Jun  
 Jul  Aug  Sep  Oct  Nov  Dec

### Run

Manual Run  Export  Canvas (experimental)

Simulate

Figure 22 shows the results of example 2. Required space is around 3000GB for this example, along with 420GB of reserve space; total space is ~3420GB.

Figure 22) Results of example 2 using Veeam simulator.

## Result

Retention		File	Size	Modify Date	Point Date
36 (30)		full.vbk	<u>400 GB</u>	2017-12-23 Sa 22	2017-12-23 Sa 22
35 (30)		incremental.vib	<u>20 GB</u>	2017-12-24 Su 22	2017-12-24 Su 22
34 (30)		incremental.vib	<u>20 GB</u>	2017-12-25 Mo 22	2017-12-25 Mo 22
33 (30)		incremental.vib	<u>20 GB</u>	2017-12-26 Tu 22	2017-12-26 Tu 22
32 (30)		incremental.vib	<u>20 GB</u>	2017-12-27 We 22	2017-12-27 We 22
31 (30)		incremental.vib	<u>20 GB</u>	2017-12-28 Th 22	2017-12-28 Th 22
30		incremental.vib	<u>20 GB</u>	2017-12-29 Fr 22	2017-12-29 Fr 22
29		full.vbk	<u>400 GB</u>	2017-12-30 Sa 22	2017-12-30 Sa 22
28		incremental.vib	<u>20 GB</u>	2017-12-31 Su 22	2017-12-31 Su 22
27		incremental.vib	<u>20 GB</u>	2018-01-01 Mo 22	2018-01-01 Mo 22
26		incremental.vib	<u>20 GB</u>	2018-01-02 Tu 22	2018-01-02 Tu 22
25		incremental.vib	<u>20 GB</u>	2018-01-03 We 22	2018-01-03 We 22
24		incremental.vib	<u>20 GB</u>	2018-01-04 Th 22	2018-01-04 Th 22
23		incremental.vib	<u>20 GB</u>	2018-01-05 Fr 22	2018-01-05 Fr 22
22		full.vbk	<u>400 GB</u>	2018-01-06 Sa 22	2018-01-06 Sa 22
21		incremental.vib	<u>20 GB</u>	2018-01-07 Su 22	2018-01-07 Su 22
20		incremental.vib	<u>20 GB</u>	2018-01-08 Mo 22	2018-01-08 Mo 22
19		incremental.vib	<u>20 GB</u>	2018-01-09 Tu 22	2018-01-09 Tu 22
18		incremental.vib	<u>20 GB</u>	2018-01-10 We 22	2018-01-10 We 22
17		incremental.vib	<u>20 GB</u>	2018-01-11 Th 22	2018-01-11 Th 22
16		incremental.vib	<u>20 GB</u>	2018-01-12 Fr 22	2018-01-12 Fr 22
15		full.vbk	<u>400 GB</u>	2018-01-13 Sa 22	2018-01-13 Sa 22
14		incremental.vib	<u>20 GB</u>	2018-01-14 Su 22	2018-01-14 Su 22
13		incremental.vib	<u>20 GB</u>	2018-01-15 Mo 22	2018-01-15 Mo 22
12		incremental.vib	<u>20 GB</u>	2018-01-16 Tu 22	2018-01-16 Tu 22
11		incremental.vib	<u>20 GB</u>	2018-01-17 We 22	2018-01-17 We 22
10		incremental.vib	<u>20 GB</u>	2018-01-18 Th 22	2018-01-18 Th 22
9		incremental.vib	<u>20 GB</u>	2018-01-19 Fr 22	2018-01-19 Fr 22
8		full.vbk	<u>400 GB</u>	2018-01-20 Sa 22	2018-01-20 Sa 22
7		incremental.vib	<u>20 GB</u>	2018-01-21 Su 22	2018-01-21 Su 22
6		incremental.vib	<u>20 GB</u>	2018-01-22 Mo 22	2018-01-22 Mo 22
5		incremental.vib	<u>20 GB</u>	2018-01-23 Tu 22	2018-01-23 Tu 22
4		incremental.vib	<u>20 GB</u>	2018-01-24 We 22	2018-01-24 We 22
3		incremental.vib	<u>20 GB</u>	2018-01-25 Th 22	2018-01-25 Th 22
2		incremental.vib	<u>20 GB</u>	2018-01-26 Fr 22	2018-01-26 Fr 22
1		full.vbk	<u>400 GB</u>	2018-01-27 Sa 22	2018-01-27 Sa 22
			<u>3000 GB</u>		
		Work Space	<u>+420 GB</u>		
			<u>3420 GB</u>		

**Note:** As shown in Figure 22, sometimes Veeam has more retention points than specified. This result is because restore point 31 is an incremental backup, and it depends on the next full backup. Veeam best practices recommend to account for that additional space due to the extra restore points and not manipulate the simulator to avoid the few additional restore points.

### Example 3: Forward Incremental, Daily Backup with 30-Day Retention Along with a Monthly Full

Figure 23 gives us an idea how the simulator would look if all the parameters related to example 3 were input.

Figure 23) Parameter input of example 3 in Veeam simulator.

## The Restore Point Simulator

Current version : [0.3.3](#)

Feedback via [@tdewin](#) or on [GitHub](#)

RPS heavily relies on some opensource [javascript frameworks](#)

### Quick Presets

Incremental Monthly Active Full

### Configuration

Style	Incremental
Used Size GB	800
Retention Points	30
Change Rate	5% Optimistic
Data left after reduction	50% (100GB > 50GB) 2x Conservative
Interval	Daily
Time Growth Simulation	<input type="checkbox"/> 1 Year 10%

### Incremental Specific

Synthetic  MO  TU  WE  TH  FR  SA  SU

Active Full Weekly  MO  TU  WE  TH  FR  SA  SU

Active Full Monthly  Jan  Feb  Mar  Apr  May  Jun  
 Jul  Aug  Sep  Oct  Nov  Dec

### Run

































































Manual Run  Export  Canvas (experimental)

Simulate

Figure 24 shows the results of example 3. Required space is around 2420GB for this example, along with 420GB of reserve space; total space is ~2840GB.



Figure 24) Results of example 3 using Veeam simulator.

64 (30)		full.vbk	400 GB
63 (30)		incremental.vib	20 GB
62 (30)		incremental.vib	20 GB
61 (30)		incremental.vib	20 GB
60 (30)		incremental.vib	20 GB
59 (30)		incremental.vib	20 GB
58 (30)		incremental.vib	20 GB
57 (30)		incremental.vib	20 GB
56 (30)		incremental.vib	20 GB
55 (30)		incremental.vib	20 GB
54 (30)		incremental.vib	20 GB
53 (30)		incremental.vib	20 GB
52 (30)		incremental.vib	20 GB
51 (30)		incremental.vib	20 GB
50 (30)		incremental.vib	20 GB
49 (30)		incremental.vib	20 GB
48 (30)		incremental.vib	20 GB
47 (30)		incremental.vib	20 GB
46 (30)		incremental.vib	20 GB
45 (30)		incremental.vib	20 GB
44 (30)		incremental.vib	20 GB
43 (30)		incremental.vib	20 GB
42 (30)		incremental.vib	20 GB
41 (30)		incremental.vib	20 GB
40 (30)		incremental.vib	20 GB
39 (30)		incremental.vib	20 GB
38 (30)		incremental.vib	20 GB
37 (30)		incremental.vib	20 GB
36 (30)		incremental.vib	20 GB
35 (30)		incremental.vib	20 GB
34 (30)		incremental.vib	20 GB
33 (30)		incremental.vib	20 GB
32 (30)		incremental.vib	20 GB
31 (30)		incremental.vib	20 GB
30		incremental.vib	20 GB
29		full.vbk	400 GB
28		incremental.vib	20 GB
27		incremental.vib	20 GB
26		incremental.vib	20 GB
25		incremental.vib	20 GB
24		incremental.vib	20 GB
23		incremental.vib	20 GB
22		incremental.vib	20 GB
21		incremental.vib	20 GB
20		incremental.vib	20 GB
19		incremental.vib	20 GB
18		incremental.vib	20 GB
17		incremental.vib	20 GB
16		incremental.vib	20 GB
15		incremental.vib	20 GB
14		incremental.vib	20 GB
13		incremental.vib	20 GB
12		incremental.vib	20 GB
11		incremental.vib	20 GB
10		incremental.vib	20 GB
9		incremental.vib	20 GB
8		incremental.vib	20 GB
7		incremental.vib	20 GB
6		incremental.vib	20 GB
5		incremental.vib	20 GB
4		incremental.vib	20 GB
3		incremental.vib	20 GB
2		incremental.vib	20 GB
1		full.vbk	400 GB
Total			2420 GB
Working Space			+ 420 GB
Grand Total			2840 GB

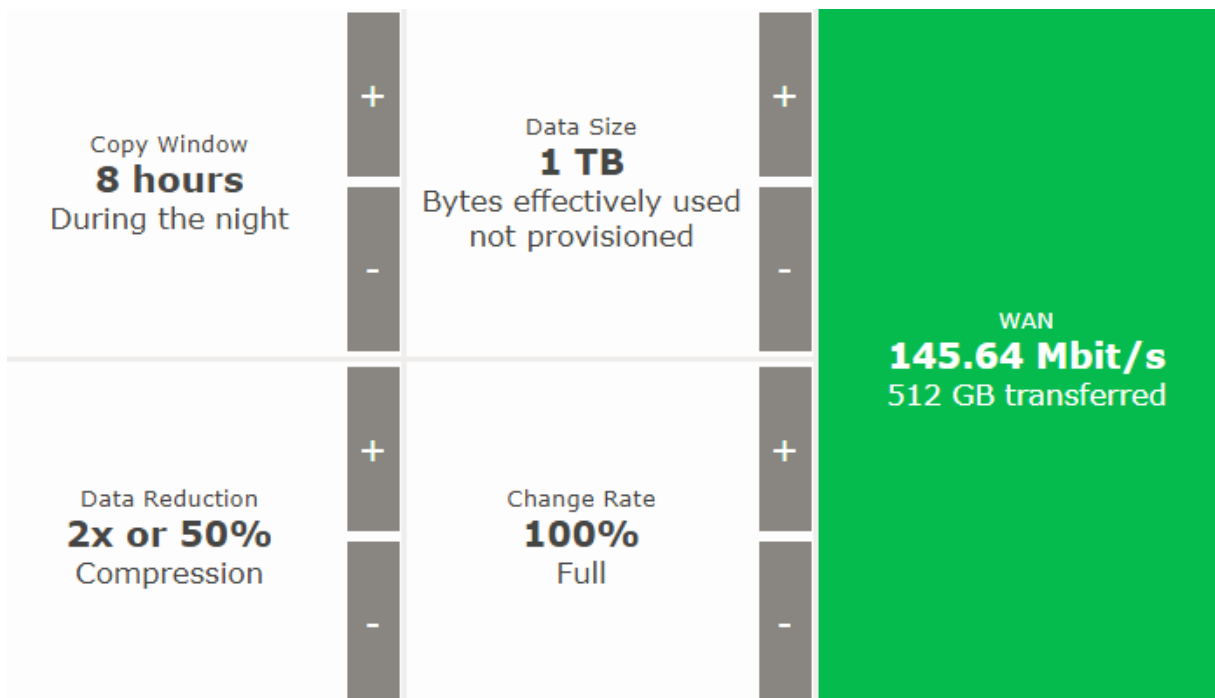
## 11 Network Planning and Sizing for Veeam and NetApp E-Series

Veeam has an unofficial tool that can assist in network planning. It can be accessed at <http://rps.dewin.me/bandwidth/>.

Figure 25 gives a snapshot of the tool. Overall it has four parameters that need input, based on which it gives an approximate bandwidth required.

- Copy Window: This parameter is divided into four categories; select the input based on the expected copy window:
  - Short Slot: Your copy window is below 6 hours.
  - During the night: Your copy window is below 8 to 9 hours.
  - During business hours: Your copy window is below 24 hours.
  - Over multiple days: Your copy window can span multiple days.
- Data Size: Make sure you input the bytes effectively used, not provisioned.
- Data Reduction: This parameter is divided into four categories:
  - Reduction Disabled: No compression or data reduction is carried out.
  - Compression: From 40% to 80% data is transferred after compression and reduction.
  - Extreme Compression: From 33% to 36% data is transferred after compression and reduction.
  - Veeam WAN Acceleration: WAN acceleration is used to reduce the data transferred below 30%. To learn more about WAN acceleration, visit [https://bp.veeam.expert/resource\\_planning/wan\\_acceleration.html](https://bp.veeam.expert/resource_planning/wan_acceleration.html).
- Change Rate: This parameter should be selected based on the data change rate expected. Average is ~10%. Databases have a change rate ~20% to 50%.

Figure 25) Network planning tool.




Performance with Veeam and NetApp E-Series

The performance testing carried out with NetApp E2800 and Veeam proved to be a robust solution that meets the backup and restore needs of modern data centers. The testing objective was to determine the best practices in addition to those highlighted in Veeam Backup & Replication 9.5 Design Considerations and NetApp E-Series Volume Configuration Guidelines.

## 11.1 Test Environment and Setup

Figure 26) Test environment diagram.



The test setup is divided into the following components:

- E-Series production storage:
  - Production storage was on a NetApp E5600 with 24 SSDs.
  - Four Dell R720 ESXi servers hosted a total of 96 VMs.
  - Each server has 24 VMs with 60GB of capacity.
  - All VMs had a test file to make sure that real-world test scenarios were replicated.
- Veeam Backup Server:
  - One Dell R720 server was used to run Veeam Backup & Replication. This server was also used as a backup proxy.
- E-Series backup repository:
  - NetApp E2800 was used to host the backup repositories.
  - System consisted of a mix of 4TB and 8TB drives.
  - The E2800 was directly connected to the Veeam Backup Server using SAS.

**Note:** FC or iSCSI could have been used instead of SAS.

## 11.2 Volume Configuration for Performance Testing

The section on NetApp E-Series Volume Configuration Guidelines highlights the various recommended configuration options. During our performance testing, it was decided to run spot tests on each of the potential configuration options:

- RAID 6, RAID 10, and DDP were tested.
- An even number of volume groups was used per test as per the NetApp E-Series Volume Configuration Guidelines.
- One volume per volume group was created for each of the RAID levels to have an equal number of volumes per controller.

### 11.3 Backup Performance Results

The following were the observations made from the testing we carried out:

- Spot checks on each of the configuration options showed consistent performance between all RAID levels, including DDP.
- The RAID level the customer chooses should be based on the level of data protection and availability desired. The amount of usable capacity varies based on the RAID level selected.

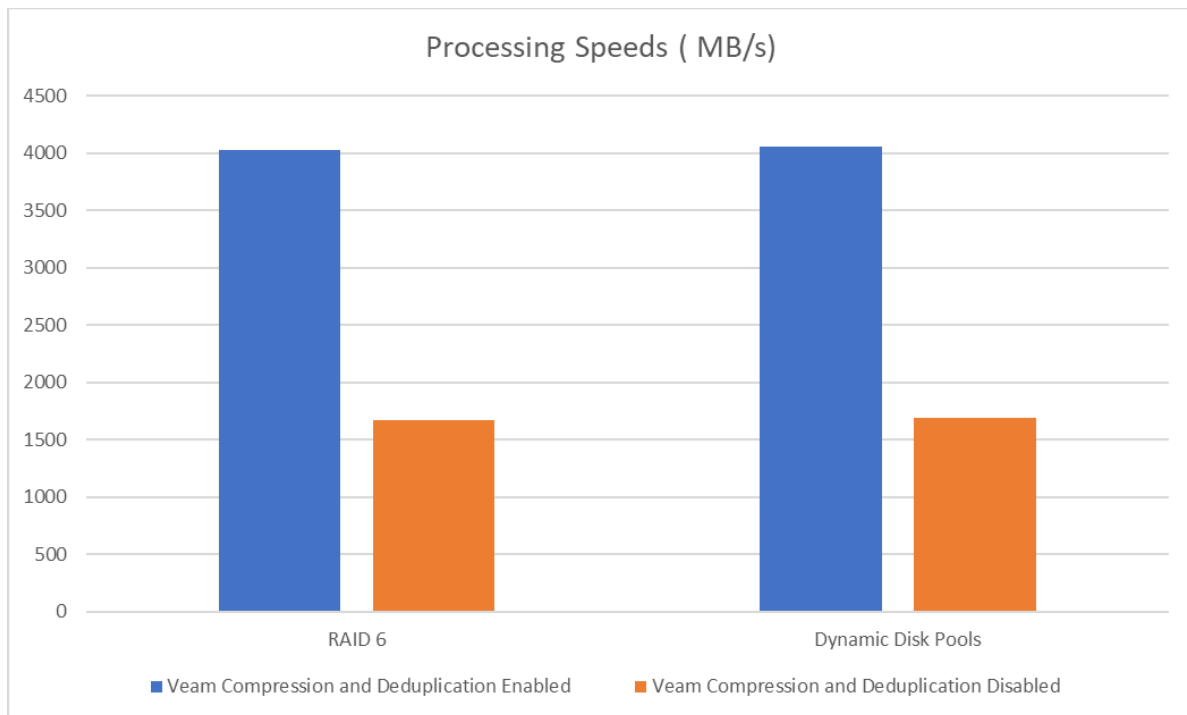
#### Performance Results with E2800

Figure 27 shows the processing speeds achieved when NetApp E2800 is used for backup repositories with Veeam Backup & Replication:

- The chart shows speeds for RAID 10, RAID 6, and Dynamic Disk Pools (DDP) with compression/deduplication enabled and disabled.
- Consistent performance is achieved with all three volume configurations.
- With compression and deduplication enabled, processing speeds of ~4GBps were achieved. This result is not the actual throughput achieved because it includes Veeam compression and deduplication.
- The processing speeds achieved with compression and deduplication disabled are ~1.6GBps. This result can be considered as the actual throughput achieved with this solution because it only accounts for the actual data transferred.

**Note:** Processing speeds vary depending on the compression and deduplication ratio achieved with Veeam. It is fair to assume it will be greater than ~1.6 GB/s as that was achieved without compression and deduplication.

Figure 27) Processing Speeds with Veeam and E2800



**Note:** Scale out backup repository can be used with Veeam to take advantage of NetApp E-Series hybrid capability. To learn more about scale out backup repository use the [link](#).

## 11.4 Performance Best Practices

Based on the performance tests carried out, a list of best practices was created. These help setup and tune the system to achieve excellent backup/restore performance.

1. When you start a data protection or disaster recovery job, Veeam Backup & Replication analyzes the list of VMs added to the job, and creates a separate task for every disk of every VM to be processed. Veeam Backup & Replication then defines what backup infrastructure components must be used for the job, checks what backup infrastructure components are currently available, and assigns necessary components to process the created job tasks.

If you use the parallel data processing mode and/or schedule several jobs to run in parallel, backup infrastructure components typically process several tasks at the same time. You can limit the number of tasks that backup infrastructure components must process concurrently. Task limitations helps you balance the workload across the backup infrastructure and avoid performance bottlenecks.

Veeam Backup & Replication lets you limit the number of concurrent tasks for the following backup infrastructure components:

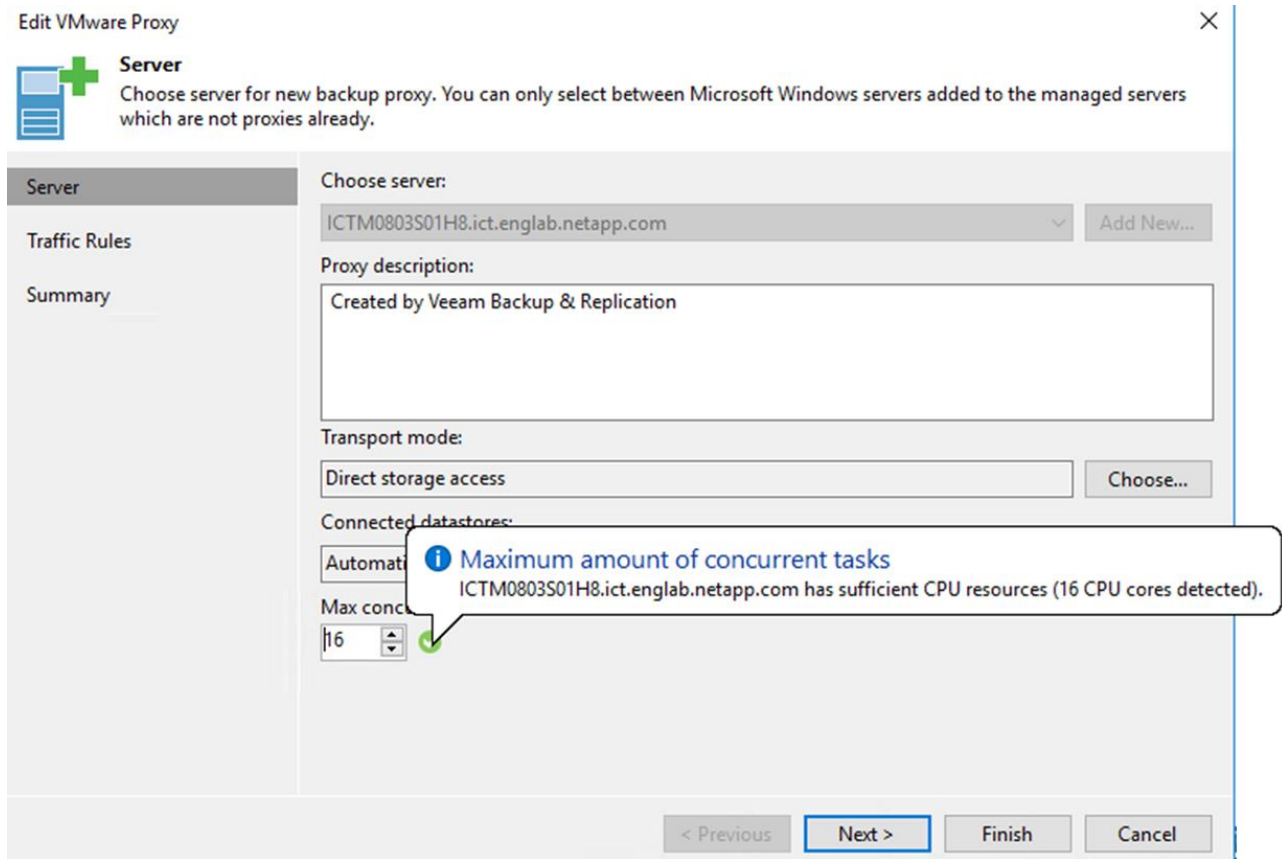
**Note:** Task limits set for backup infrastructure components influence the job performance. For example, you add a VM with 4 disks to a job and assign a backup proxy that can process maximum 2 tasks concurrently for the job. In this case, Veeam Backup & Replication will create 4 tasks (1 task per each VM disk) and start processing 2 tasks in parallel. The other 2 tasks will be pending.

- a. Backup Proxies

To limit the number of concurrent tasks on a backup proxy, you must define the Max concurrent tasks setting for the backup proxy. For on host backup, you must define the Task limit setting for the source Microsoft Hyper-V host.

The maximum number of concurrent tasks depends on the number of CPU cores available on the backup proxy. It is strongly recommended that you define task limitation settings using the following rule: 1 task = 1 CPU core. For example, if a backup proxy has 4 CPU cores, it is recommended that you limit the number of concurrent tasks for this backup proxy to 4.

Figure 28) Restricting concurrent tasks per backup proxy



#### b. Backup Repositories

To limit the number of concurrent tasks on a backup repository, you must enable the Limit maximum concurrent tasks to <N> option on the backup repository and define the necessary task limit.

The maximum number of concurrent tasks depends on the number of CPU cores available on the backup repository. It is strongly recommended that you define task limitation settings using the following rule: 1 task = 1 CPU core.

It is recommended to configure 2 GB RAM per core. In case of shared folder backup repositories, the same amount of resources is required for gateway servers.

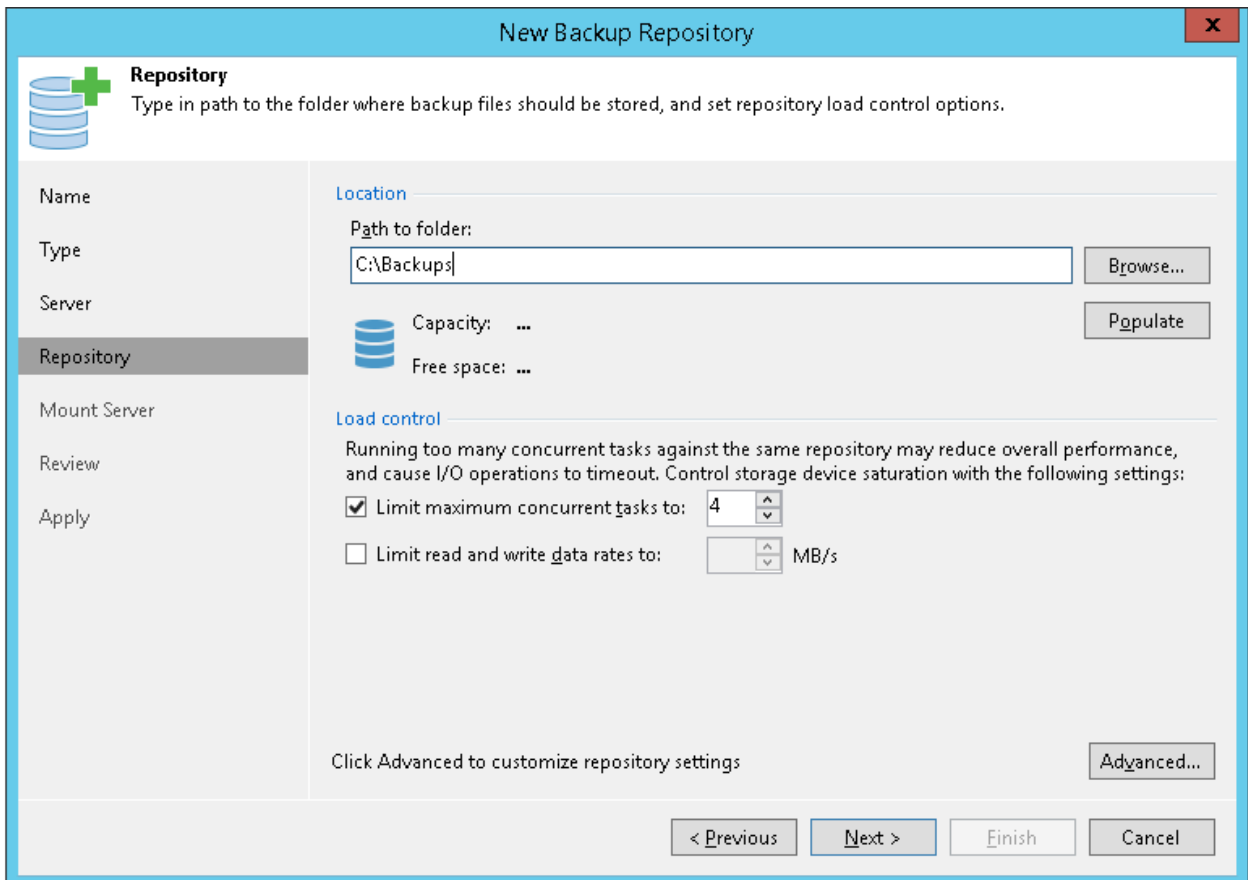
Synthetic operations performed on the backup repository (such as synthetic full backup, backup files merge and transform) are also regarded as tasks. The number of tasks performed during these operations depends on the type of backup chains stored on the backup repository:

- For regular backup chains, Veeam Backup & Replication creates 1 task per job.
- For per-VM backup chains, Veeam Backup & Replication creates 1 task per every VM chain (that is, every VM added to the job).

If you use backup repositories for backup copy jobs, you must also consider tasks for read operations.

**Note:** When you limit the number of tasks for the backup repository, bear in mind the storage throughput. If the storage system is not able to keep up with the number of tasks that you have assigned, it will be the limiting factor. It is recommended that you test components and resources of the backup infrastructure to define the workload that they can handle.

Figure 29) Restricting concurrent tasks per backup repository



2. Adjust the number of active snapshots per datastore, to be in line with the number of concurrent tasks per backup repository. By default, it is restricted to 4 as a “protection method” in order to avoid filling up a datastore. If left at default value, your number of tasks per backup repository will be restricted to 4.
  - The default 4 active snapshot per datastore value can be modified by creating a registry DWORD value in ‘HKEY\_LOCAL\_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication\’ called MaxSnapshotsPerDatastore and use the appropriate hex or decimal value.
  - Any number can be selected but ensure you have enough capacity on your backup repository.
3. Use the recommended guidelines with regards to Direct SAN Access Mode. This is the recommended transfer mode.

## 12 NetApp SANtricity Plug-Ins for VMware

### 12.1 NetApp SANtricity Plug-In for VMware vCenter

The NetApp SANtricity Plug-In for VMware vCenter is a VMware vCenter Server plug-in that provides integrated management of E-Series storage arrays and EF-Series flash arrays from within a VMware vSphere Web Client. The vSphere Web Client is a single management interface that you can use to manage the VMware infrastructure and all your day-to-day storage needs.

The plug-in enables you to perform the following tasks:

- Configure ESXi hosts to NetApp SANtricity E-Series storage arrays and EF-Series flash arrays.



- Provision new and existing storage array volumes.
- Map storage array volumes to ESXi hosts and host groups.
- Manage synchronous and asynchronous mirroring and storage array snapshots.
- View vCenter datastores that are on E-Series and EF-Series storage volumes.

For the latest download and details, go to [support.netapp.com](http://support.netapp.com).

## Summary

Veeam Backup & Replication offers powerful and cutting-edge capabilities in the data protection industry, but without a proper repository, backup windows and recoveries can be negatively affected. Veeam provides more recovery options and faster restoration capabilities, but to realize these benefits, the storage system must have the necessary performance profile. With technologies such as instant VM recovery, you can run an application directly from your backup file, but how is that application going to perform? NetApp E-Series arrays offer the performance that you need when recovering one or more applications and give you confidence that the data that you backed up is protected and is available when you need it.

### NetApp E-Series Documentation

Go to the following link for NetApp E-Series documentation:

<https://mysupport.netapp.com/info/web/ECMP1658252.html>

### Veeam Backup & Replication 9.5 Software Documentation

Go to the following link for Veeam Backup & Replication 9.5 software documentation:

<http://www.veeam.com/documentation-guides-datasheets.html>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright Information**

Copyright © 2018 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### **Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.