

Backup and Recovery Best Practices with Veeam Backup & Replication™

Revision History

Version	Date	Description	Author
1.2	01/05/2018	Updated Release	Bill Roth
1.1	07/11/2016	Updated Release	Bill Roth
1.0	05/14/2014	Initial Release	Prior Contributor

Table 1 - Revision history

Contents

- Introduction 5
- Intended Audience..... 5
 - Prerequisites..... 5
 - Considerations & Limitations 5
- Consolidated List of Practices..... 5
- Architecture..... 6
 - Backup Server..... 6
 - Backup Proxy..... 6
 - Backup Repository..... 7
 - Backup Infrastructure Deployment Summary..... 8
 - General Options 9
 - Backup Job Configuration Notes..... 10
- VMware Transport Modes.....11
 - Direct Storage Access / Direct NFS Access 13
 - Virtual Appliance / HotAdd 14
 - Network / NBD or NBDSSL..... 15
 - VMware Transport Mode Summary 15
- Veeam vPower 16
 - Instant Recovery17
 - SureBackup..... 18
 - Application Group..... 18
 - Virtual Lab..... 19
- Creating Backup Copies with Veeam 20
- Summary 21
- References 21
- Appendix A: Snapshot Collisions – Tintri Snapshots..... 22
- Appendix B: Tintri VM Level Quality of Service 23
 - Quality of Service Automation with Veeam Backup & Replication..... 24
- Appendix C: Backup Repository Deployment - Tintri System 26
 - Repository Type..... 26
 - Repository Path..... 26
 - Data Reduction & Compression Settings.....27

Introduction

Deploying Veeam Backup & Replication with one or more Tintri systems is straightforward. The products are complimentary in that Veeam is a perfect fit for protecting virtualized environments, and Tintri has been designed for virtualized workloads. This document is focused on VMware deployments protected with Veeam that are hosted on Tintri. Three key areas of interoperability; protecting VMs hosted on a Tintri system, using a Tintri system as a vPower NFS write cache, and using a Tintri system as a virtual lab datastore, are examined within this document. Relevant configuration settings are called out and explained. Recommended usage and solution limitations are highlighted where applicable.

Intended Audience

Focused on building a supported and successful data protection solution, this document targets key best practices and known challenges. Virtualization administrators and staff members associated with architecting, deploying, and administering a Veeam Backup & Replication solution in conjunction with Tintri systems are encouraged to read this document.

Prerequisites

General knowledge of and familiarity with Tintri systems is essential prior to architecting or implementing a data protection solution with Veeam Backup & Replication. Similarly, prior experience and familiarity with Veeam Backup & Replication is also recommended. For additional information about Tintri, please visit <http://www.tintri.com/>. For additional information about Veeam Backup & Replication, please visit <http://www.veeam.com/>.

Considerations & Limitations

Product compatibility and support matrices should be referenced to confirm that a given configuration is supported prior to implementation. This includes but is not limited to Tintri products, and Veeam products. For Tintri support information please visit <http://support.tintri.com>. The Tintri support site requires login credentials.

Descriptions provided and examples depicted within this document are based on Tintri Operating System version 4.3 and higher in conjunction with Veeam Backup & Replication 9.5 update 2 and higher.

This document does not take the place of Tintri product documentation or Veeam product documentation.

The scope of this document is constrained to integrating Tintri systems into a Veeam Backup & Replication environment. This document is not intended as a substitute for formal Tintri or Veeam training.

Consolidated List of Practices

The table below includes the recommended practices in this document. Click the text on any of the recommendations to jump to the section that corresponds to each recommendation for additional information.

- Users experiencing high latency conditions on VMs being backed up with the Direct NFS transport mode are encouraged to upgrade to Veeam Backup & Replication version 9.5 update 2 or higher.
- When using the SCSI HotAdd transport mode, use a backup proxy on the same ESXi host as the VM or VMs being protected.

- When using the SCSI HotAdd transport mode, read Veeam KB 1681 and use the “EnableSameHostHotaddMode” mode when appropriate.
- Use a 10 GbE or faster network connection with the NBD transport mode.
- Reference the Veeam Help Center for additional detail about the use of network mode (NBD).
- Avoid configuring overlapping backup schedules with Veeam Backup & Replication and a Tintri snapshot schedule that takes VM-consistent snapshots of the same VM or VMs.
- Consider using Tintri crash-consistent snapshots in cases where overlapping a Veeam Backup & Replication schedule with a Tintri snapshot schedule cannot be avoided.
- Do not backup VMs residing on a Tintri system to a backup repository residing on the same Tintri system. VMs being protected by Veeam Backup & Replication should always use a backup repository residing on a different storage device.

Architecture

Veeam Backup & Replication is both modular and scalable. It can easily accommodate a large variety of virtual environments and configurations. The infrastructure is comprised of components that fulfill the requirements necessary to perform backups, restores, replication, disaster recovery, and administration. This section provides a basic overview of the primary components and provides insight into how they can be deployed. This section is intended to serve as a review for experienced Veeam Backup & Replication administrators, and as introductory information for data protection administrators that may be deploying Veeam for the first time. If a setting is not explicitly called out as a best practice, it is being discussed for awareness only within the context of this section.

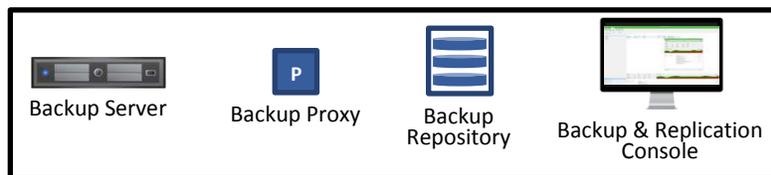


Figure 1 - Basic Veeam Backup & Replication components

A newly deployed Veeam backup server includes by default a backup proxy, backup repository, as well as a backup & replication console.

Backup Server

The backup server coordinates backup, replication, recovery verification, and restore tasks. The backup server also controls job scheduling and resource allocation. It is used to configure and manage backup infrastructure components. It is also used to specify global settings. The backup server includes a local or remotely deployed Microsoft SQL Server database which stores data about the backup infrastructure, jobs, and sessions. There is one backup server in any given Veeam Backup & Replication deployment.

Distributed deployments consisting of multiple Veeam Backup & Replication instances are recommended for geographically dispersed environments. The Veeam Enterprise Manager is an optional component that provides centralized management and reporting by means of a web interface for local and geographically dispersed deployments.

Backup Proxy

The backup proxy is a backup infrastructure component. The backup proxy resides between source data that needs to be protected and a target. A backup proxy can be installed as a standalone entity, or it can

be co-located with other Veeam components. The target can be a backup repository or another backup proxy. The backup proxy processes jobs and delivers backup traffic. The backup server is the point of control for dispatching jobs to one or more backup proxies.

A backup proxy includes a configurable setting that enforces a maximum number of concurrent tasks. The following graphic depicts a VMware backup proxy server.

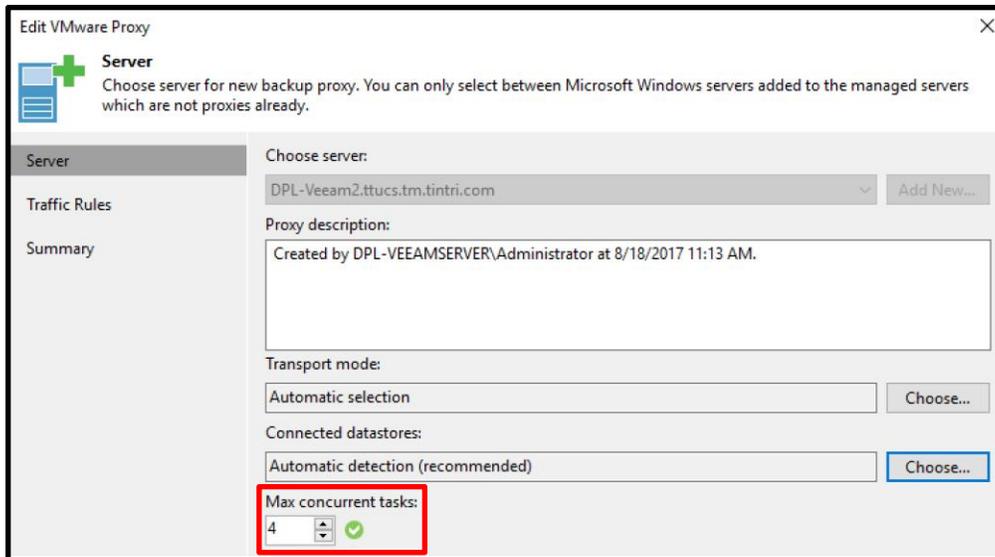


Figure 2 - VMware Proxy - Max concurrent tasks

The default limit is based on the number of CPU cores present on the proxy server, where each concurrent task requires a single CPU core. Adding backup proxy servers to a deployment facilitates scalability such that a higher number of simultaneous tasks can be executed, which may result in greater aggregate data transfer rates.

Backup Repository

The backup repository is a backup infrastructure component, used by Veeam Backup & Replication to store backups, copies of VMs, and metadata for replicated VMs. A backup repository can be installed as a standalone entity, or it can be co-located with other Veeam components.

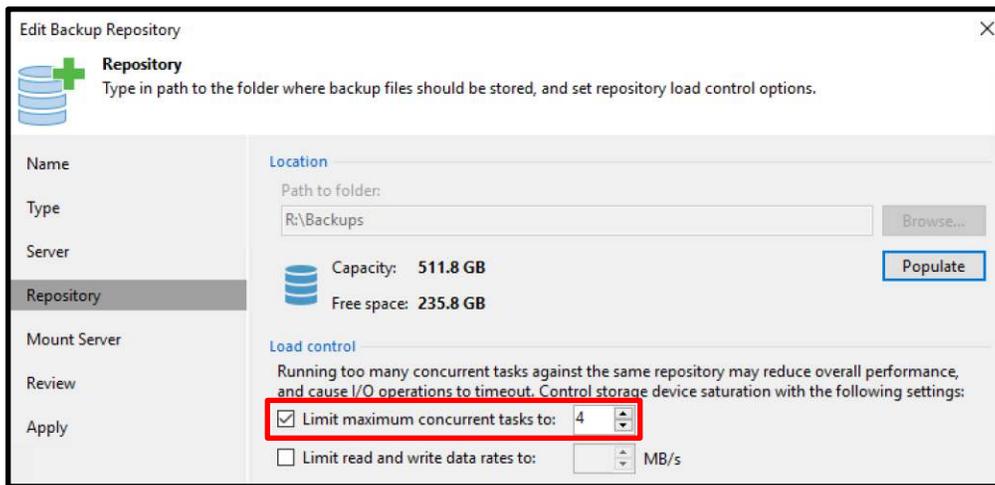


Figure 3 - Backup Repository - Limit maximum concurrent tasks

A backup repository includes a configurable setting that limits the maximum number of concurrent tasks. The use of multiple backup repositories facilitates scaling such that a higher number of concurrent tasks can be executed. The backup repository also includes an optional configuration setting that limits read and write data rates to a user supplied value. The data rate parameter can be set as low as 1 MB/s or as high as 1024 MB/s.

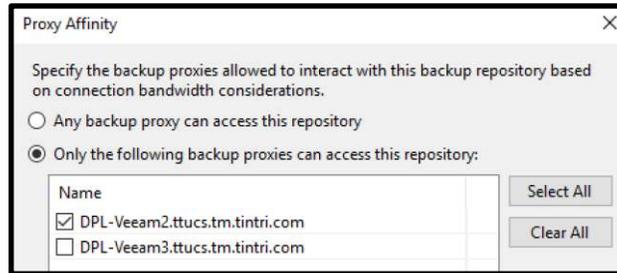


Figure 4 - Proxy Affinity

By default, a backup repository can be used by all backup proxies within a given deployment. Each backup repository includes a configurable setting called “Proxy Affinity”. Proxy affinity enables the ability to control which backup proxies can access a particular backup repository. Example use cases for proxy affinity include:

- The use of a backup repository by a remote backup proxy can be disabled.
- The use of a backup repository by a backup proxy connecting over a slow network can be disabled.

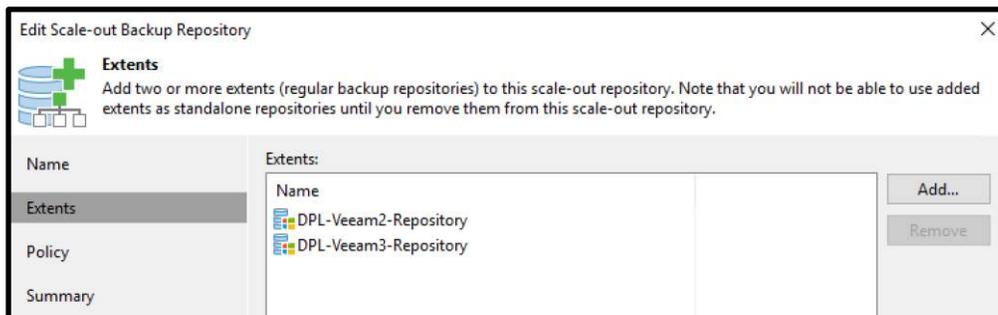


Figure 5 - Scale-out Backup Repository - Extents

Scale-out backup repositories group one or more regular backup repositories into a logical entity. Within a scale-out backup repository, regular backup repositories are listed as extents. The capacity of a scale-out backup repository is represented as the aggregate capacity of its extents. Scale-out backup repository capacity can be expanded by adding one or more extents.

Note that proxy affinity setting cannot be configured directly on a Scale-out backup repository. Instead, proxy affinity settings can be configured at the extent level (regular backup repository level).

Backup Infrastructure Deployment Summary

This subsection provides an overview of where Veeam components can be deployed. It also includes considerations specific to deployment on physical machines, virtual machines, and virtual machines residing on a Tintri system.

Component	Placement	Considerations
-----------	-----------	----------------

Component	Placement	Considerations
Backup Server	<p>A Windows-based physical or virtual machine.</p> <p>When virtualized, the backup server may be deployed on a Tintri system.</p>	<p>Deployment includes a default backup proxy and a default backup repository.</p> <p>Use a backup repository that is not hosted on the backup server for Veeam configuration backups. This may enable to ability to recover the configuration in the event of a backup server outage.</p>
Backup Proxy	<p>A Windows-based physical or virtual machine.</p> <p>The backup proxy can also be deployed in conjunction with a backup repository.</p> <p>When virtualized, a backup proxy may be deployed on a Tintri system.</p>	<p>When the backup proxy is virtualized, Direct NFS transport mode backups and restores will pass through an ESXi host. This may increase ESXi host resource utilization and impact aggregate backup and recovery data transfer rates.</p> <p>When the backup proxy is physical, HotAdd transport mode backups cannot be performed.</p>
Backup Repository	<p>A Backup repository can be a Windows or Linux machine. A backup repository can be a physical or virtual machine.</p> <p>When a backup repository is deployed on a Windows-based machine, it can also be deployed in conjunction with a backup proxy.</p> <p>When virtualized, a backup repository may be deployed on a Tintri system.</p>	<p>When protecting VMs on a given array, use a different array or device as backup target.</p> <p>A virtualized backup repository residing on a Tintri system may impact other VMs residing on the same system when backup, recovery, and copy jobs are being performed.</p>

Table 2 - Component deployment summary

General Options

Enabling parallel processing allows VMs and VM disks within a single job to be processed simultaneously.

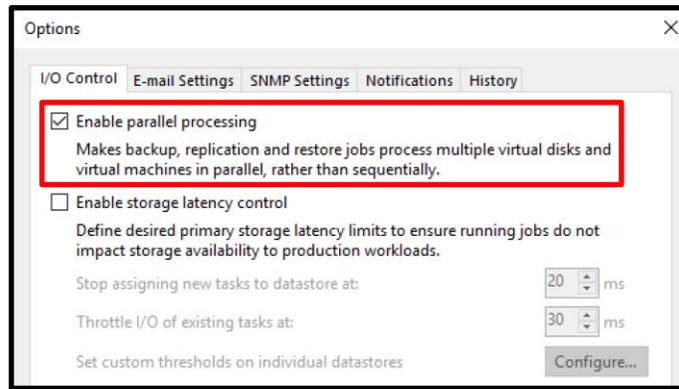


Figure 6 - Options - I/O Control – Enable parallel processing

The “Enable parallel processing” parameter is enabled by default. Shown for reference, the “Options” dialog window is accessed from the Veeam Backup & Replication console. The settings available within this dialog are global in that they affect the entire Veeam instance.

Backup Job Configuration Notes

Backup jobs include storage settings that dictate backup proxy selection as well as which backup repository will be used. The following graphic depicts backup proxy selection for a VMware backup job.

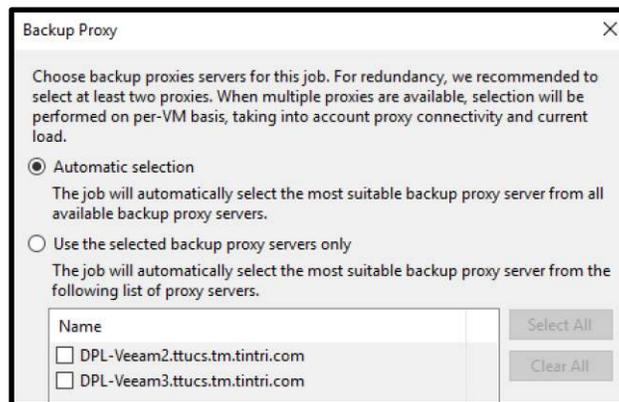


Figure 7 – VMware Backup Proxy –proxy selection

The default backup job backup proxy setting is “Automatic selection”. The “Automatic selection” option enables Veeam Backup & Replication to select the most suitable backup proxy. The default setting can be overridden to use specific backup proxy servers.

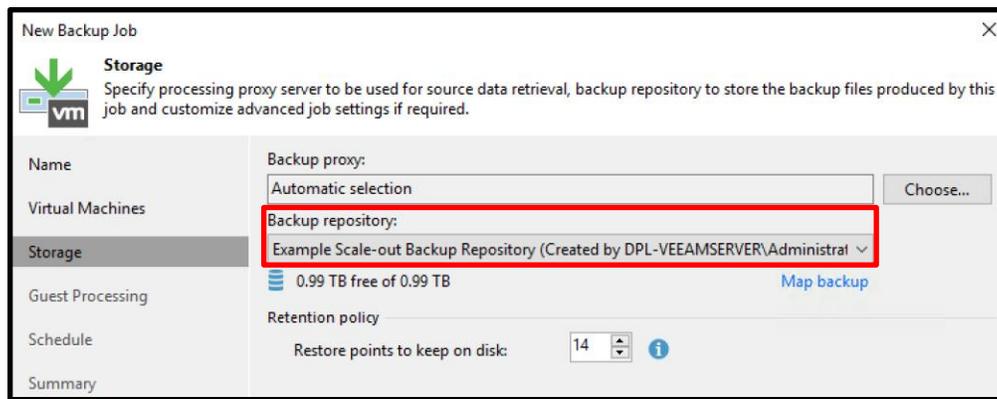


Figure 8 - Backup job - backup repository selection

Backup repository selection is accomplished via a pull-down menu where a single backup repository can be selected. It is important to understand any network hops that may exist between a given backup proxy and the target backup repository, as this may affect performance. Combining the backup proxy and backup repository on the same host eliminates a network hop when a given backup job specifies the use of both the backup proxy and backup repository residing on the same host. A potential consequence of implementing the “no-hop” strategy is that a nonoperational proxy may introduce a single point of failure.

VMware Transport Modes

Veeam Backup & Replication supports the ability to protect VMware with three distinct transport modes. The transport mode used for a given backup job dictates how VM data is retrieved from its source and written to a target backup repository. The VMware vSphere Storage APIs – Data Protection is used by Veeam Backup & Replication for the transport modes discussed in this document. VMware vSphere Storage APIs – Data Protection leverages VMware vSphere snapshots, which enables backup without requiring downtime for virtual machines.

Transport mode settings are independently configurable on each backup proxy.

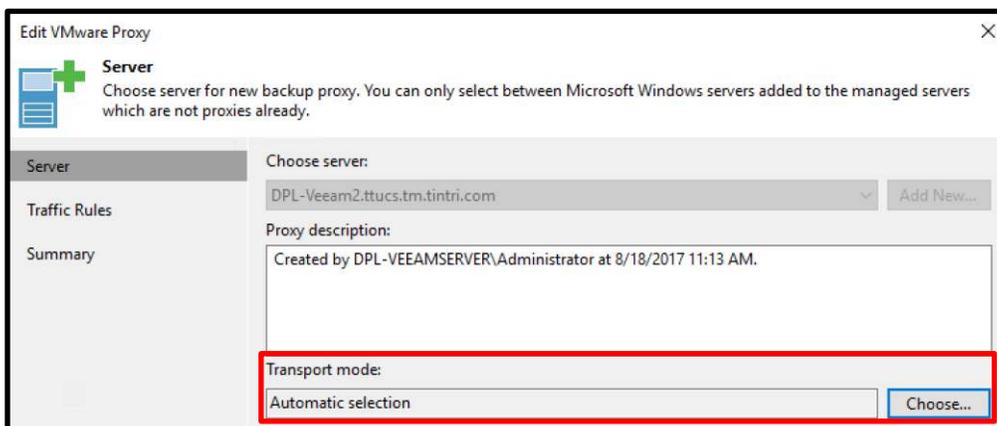


Figure 9 - Proxy transport mode selection

Within the “Edit VMware Proxy” dialog window, clicking the “Choose” button will launch a “Transport Mode” dialog where the backup proxy transport mode can be explicitly selected.

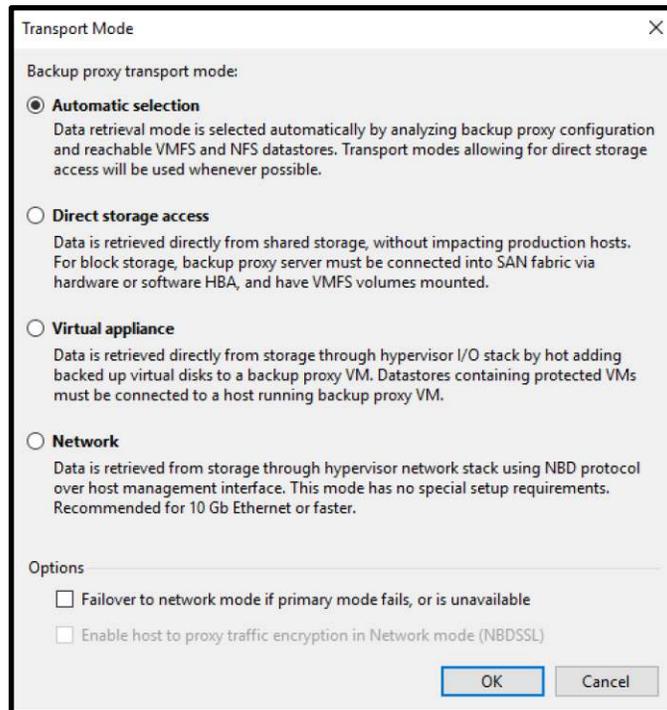


Figure 10 - Transport mode selection

Transport modes consist of “Direct storage access”, “Virtual appliance”, or “Network”. The “Network” mode includes an optional ability to encrypt data transferred between a VM host and backup proxy and is referred to as “NBDSSL”. VMware guests residing on a Tintri system can be protected with any of the available transport modes, dependent on the VMware environment being protected, the configuration of the backup infrastructure, and any specific data protection requirements a given VM may have. Best practice recommendations for each transport mode are covered within the subsection where a given transport mode is detailed.

Using “Automatic selection” within the “Transport Mode” selection dialog window allows Veeam Backup & Replication to automatically select the most efficient backup transport mode by analyzing the backup proxy configuration and the datastore.

The optional setting “Failover to network mode if primary mode fails or is unavailable” can be used in conjunction with the “Automatic selection”, “Direct storage access”, or “Virtual appliance” transport modes. When enabled, this option increases the likelihood that successful backups will occur. This option is enabled by default.

In order of backup efficiency, each transport mode is examined in greater detail in the subsequent subsections. Comprehensive transport mode information, including requirements and limitations, is available in the “Veeam Backup & Replication User Guide for VMware vSphere Environments” document.

Direct Storage Access / Direct NFS Access

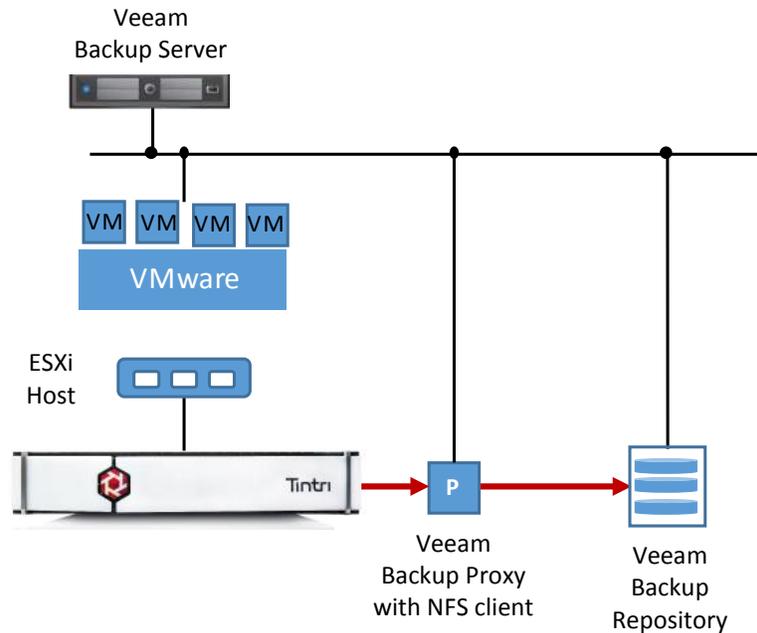


Figure 11 - Direct NFS access

The direct storage access method can function in SAN or Direct NFS transport modes. Direct NFS use is applicable to VMware virtual disks residing on an NFS datastore, such as a Tintri system. This transport mode bypasses the ESXi host and reads or writes data directly from or to an NFS datastore. Veeam Backup & Replication uses its native NFS client on a backup proxy for VM data transport. VM data travels over a LAN connection and does not create a load on the ESXi host.

Using the direct NFS access transport mode with a Tintri system requires that the backup proxy have read/write administrative access to the datastore. By default, a Veeam backup proxy has read/write administrative access to the datastore. The backup proxy can be deployed on a physical or virtual machine. In cases where the backup proxy is virtualized, it should use a VMXNET 3 network adaptor type to connect with the Tintri system data IP subnet. Ideally, both the backup proxy and Tintri system data IP will be configured on the same subnet.

Note that some Veeam Backup & Replication version 9.5 update 1 deployments may have experienced high latency on VMs being backed up with the Direct NFS transport mode. Users are encouraged to upgrade to Veeam Backup & Replication version 9.5 update 2 or higher to circumvent any high latency challenges that may exist with earlier versions of the 9.5 release.

-
- *Users experiencing high latency conditions on VMs being backed up with the Direct NFS transport mode are encouraged to upgrade to Veeam Backup & Replication version 9.5 update 2 or higher.*
-

Virtual Appliance / HotAdd

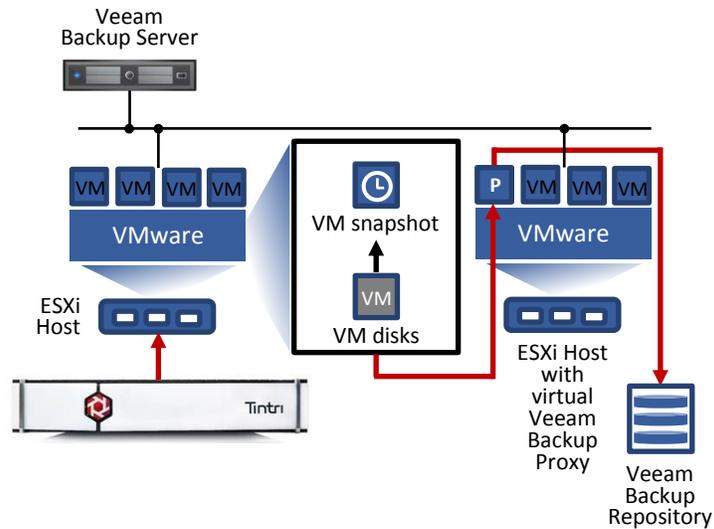


Figure 12 - SCSI HotAdd

The virtual appliance mode uses VMware SCSI HotAdd to attach disks from a backup snapshot to a backup proxy. VM data flows through an ESXi host and is retrieved or written directly from or to the datastore instead of going over the network.

Virtual appliance mode requires the backup proxy role to be deployed on a VM. The ESXi host on which the backup proxy is deployed must have access to the Tintri system hosting the virtual disks of the VMs being processed. Additionally, the backup server and backup proxy must have the latest version of VMware tools installed.

When used with a Tintri system, the SCSI HotAdd transport mode may impact the performance of a protected VM during the vSphere snapshot removal phase of a backup job. This issue may occur when a guest VM and the proxy reside on different ESXi hosts. During the snapshot removal phase of a backup operation the VM may become unresponsive for approximately 30 seconds. Consider a deployment where guest VMs being protected with the HotAdd transport mode are processed using a proxy server residing on the same ESXi host.

Veeam Knowledge Base article 1681 discusses this issue and provides additional information on the challenge, cause, and potential solutions. The article is available at <https://www.veeam.com/kb1681>.

VMware knowledge base article 2010953 also discusses the challenge and is available at https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2010953.

The HotAdd transport mode connects VMDKs to a proxy using a SCSI controller on a proxy server. Multiple simultaneous backup operations may require more than a single SCSI controller on the proxy server.

-
- *When using the SCSI HotAdd transport mode, use a backup proxy on the same ESXi host as the VM or VMs being protected.*
 - *When using the SCSI HotAdd transport mode, read Veeam KB 1681 and use the "EnableSameHostHotaddMode" mode when appropriate.*
-

Network / NBD or NBDSSL

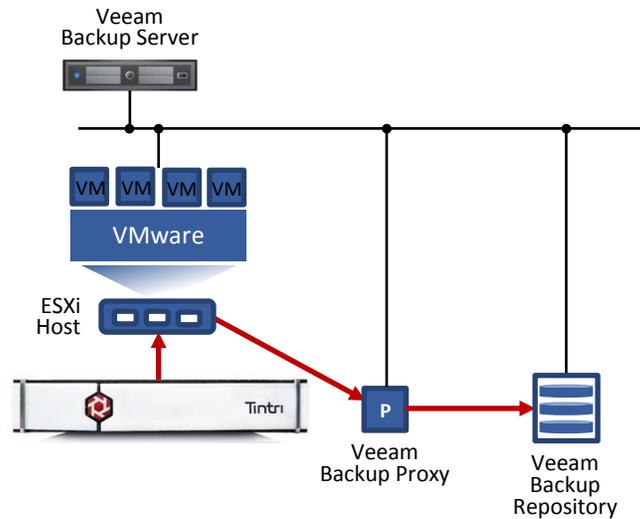


Figure 13 - Network block device

The network mode uses the VMware Network Block Device (NBD) protocol. VM data is retrieved from an ESXi host over a LAN connection by a backup proxy and is written to the target backup repository. The advantage of the network transport mode is that it can be used with any VMware infrastructure configuration. The network mode works best with a high bandwidth network connection. The use of 10 Gigabit Ethernet or faster network connections are recommended.

Additionally, it is important to understand that a given backup proxy will connect to an ESXi host based on DNS name resolution. It is possible to use a specific interface through the use of a hosts file.

The Veeam Help Center provides additional detail about the use of network mode (NBD) at https://helpcenter.veeam.com/backup/80/bp_vsphere/bp_8_network_mode.html.

- Use a 10 GbE or faster network connection with the NBD transport mode.
- Reference the Veeam Help Center for additional detail about the use of network mode (NBD).

VMware Transport Mode Summary

Mode	Data Path	Advantages	Limitations
Direct Storage Access	The backup proxy copies VM data blocks directly from the NFS datastore over the LAN. The backup proxy can be deployed on a physical or virtual machine.	The data path bypasses the ESXi host when using a physical backup proxy. Generally faster when compared to other transport modes.	Cannot be used for VMs that have one or more existing snapshots. Cannot be used with VMware tools quiescence.

Mode	Data Path	Advantages	Limitations
Virtual Appliance	<p>VM disks are attached to the backup proxy and VM data is read from the disks.</p> <p>The proxy must be deployed on a VM running on an ESXi host connected to the datastore.</p>	May provide better performance than the network mode.	VMs may become unresponsive during the snapshot removal phase of a backup in cases where the backup proxy and VM being protected reside on different ESXi hosts.
Network	<p>VM data blocks are copied from production storage through an ESXi host and sent to a backup proxy.</p> <p>The backup proxy can be deployed on any machine in the storage network.</p>	<p>The network block device protocol can be used with any infrastructure configuration.</p> <p>Minimizes VM stuning during the snapshot removal phase of a backup.</p>	<p>Potentially lower data transfer speed over a LAN.</p> <p>Typically uses only 40% of available VMKernel interface bandwidth, limiting aggregate data transfer rates.</p>

Table 3 - VMware transport mode summary

Veeam vPower

Veeam vPower technology enables a number of significant features including recovery verification, instant VM recovery, Universal Application-Item Recovery (U-AIR), and On-Demand Sandbox. An overview of the technology is briefly presented here as a foundation for subsequent topics in this section.

A key component of Veeam vPower technology is the vPower NFS Service, which runs on a Microsoft Windows host as a Windows service.

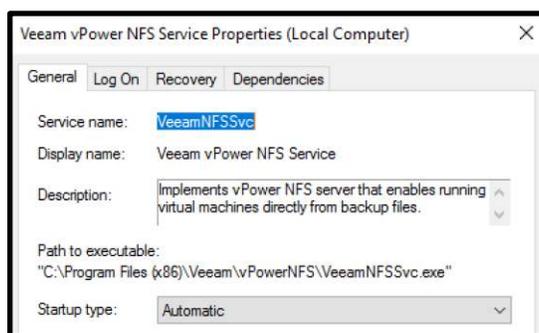


Figure 14 - Veeam vPower NFS service

The “Veeam vPower NFS Service” enables a Microsoft Windows host to act as an NFS Server. The “Enable vPower NFS service on the mount server” is enabled by default when deploying a Microsoft Windows backup repository.

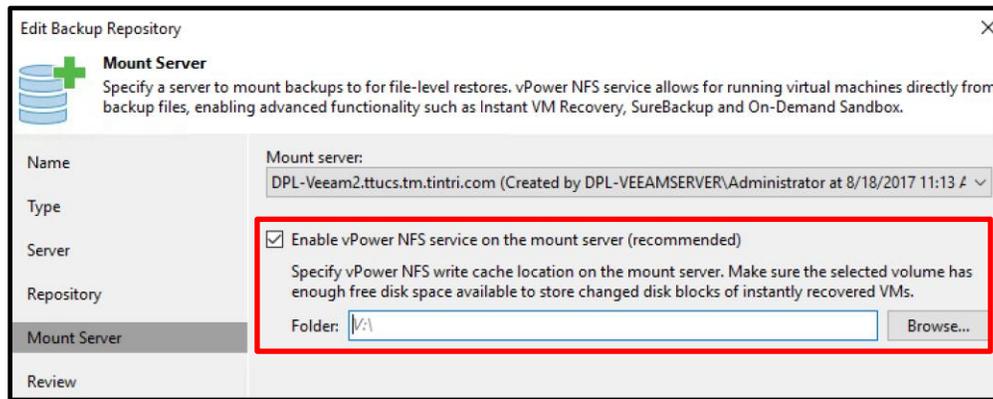


Figure 15 - Backup repository - vPower NFS

An important element of the vPower NFS service is the vPower NFS write cache, which can be deployed on a Tintri system. Deploying the vPower NFS write cache on a Tintri system virtual disk enables the use high performance Tintri storage to store changed disk blocks of an instantly recovered VM. The selected folder must reside on a volume with at least 10 GB of free space.

Instant Recovery

Instant recovery can immediately restore a VM into a production environment by running it directly from a backup file. Veeam vPower technology is used to mount a VM image to an ESXi host directly from a backup file, even when the backup file is compressed and deduplicated. The backup image of the VM remains in a read-only state. All changes that occur on the VMs virtual disk(s) are logged to auxiliary redo logs residing on the NFS server.

In the example provided below, instant recovery is performed on a VMware guest named, “DPL-V-Client1”.

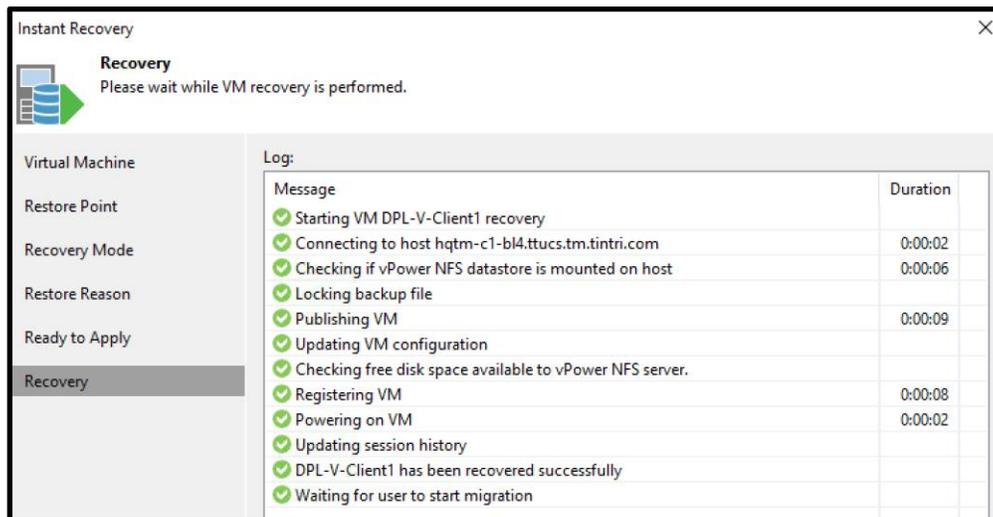


Figure 16 - Instant recovery

During the instant recovery process, the backup file is mounted as a datastore. In the example provided below, the datastore is displayed within the vSphere datastore browser.

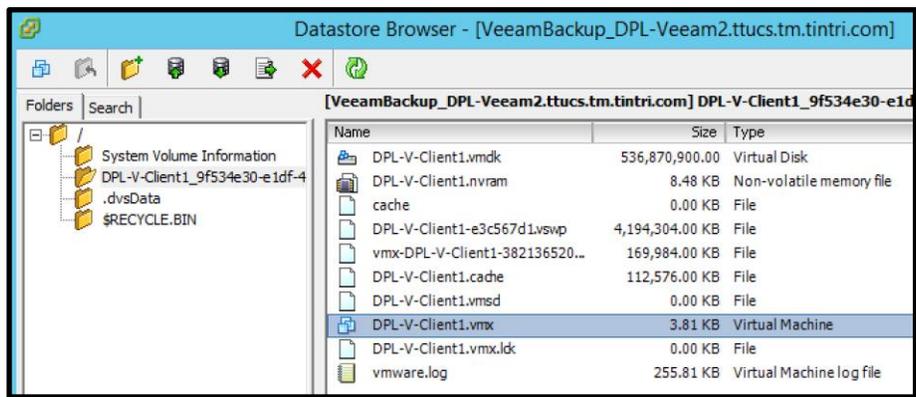


Figure 17 - Instant recovery datastore

At the point where instant recovery has completed, the vPower NFS write cache storage on the backup repository becomes populated. In the example provided below, the vPower NFS write cache is displayed within the Windows file explorer.

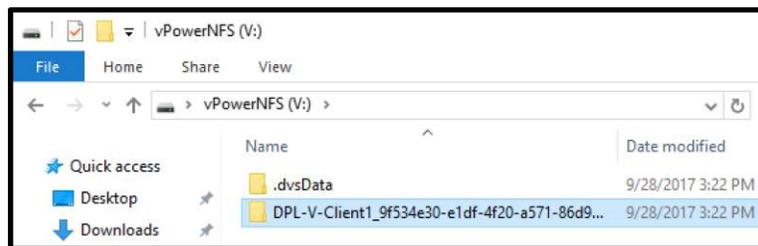


Figure 18 - vPower NFS write cache

The Veeam “Quick Migration Wizard” is then used to migrate the recovered VM. Quick Migration registers the VM on the target host, restores the VM contents from the backup file located on the backup repository and synchronizes the VM restored from backup with the running VM. After the recovered VM has been relocated to a production datastore, the VM backup image is dismounted and the vPower NFS write cache is vacated.

SureBackup

SureBackup recovery verification provides the ability to perform test recoveries from backups. It is comprised of components that are detailed in the subsequent subsections.

Note that SureBackup is available with the Enterprise and Enterprise Plus editions of Veeam Backup & Replication. When using the standard edition, users can perform manual recovery verification in conjunction with Instant VM Recovery.

Application Group

An application group defines the virtual machine(s) running a production application and any services the production application may be dependent on. The group typically contains at least a domain controller, DNS server and DHCP server. The application group includes configurable settings that define what verification tests will be performed when a SureBackup job is executed within a virtual lab:

- The role of each VM; DNS Server, Domain Controller, Global Catalog, Mail Server, SQL Server, or Web Server.

- Startup options for each VM including memory allocation, maximum boot time, an application initialization timeout value, and boot verification based on heartbeat presence or ping response.
- Test scripts that by default are based on the role of each VM, which can be optionally configured to execute a user customized test script.
- Access credentials for each VM consisting of a user ID and password.

Virtual Lab

A virtual lab is an isolated, fenced off lab environment used to verify VM recovery based on the configuration of an application group in conjunction with a SureBackup job. A virtual lab includes configurable settings for a number of user specified variables:

- The ESXi host on which the virtual lab will reside.
- The datastore on which to store redo logs, the temporary files where virtual disk changes are accumulated while VMs are running from read-only backup files.
- The optional proxy appliance, which is required for automated recovery verification.
- Networking settings that accommodate single or multiple production networks.

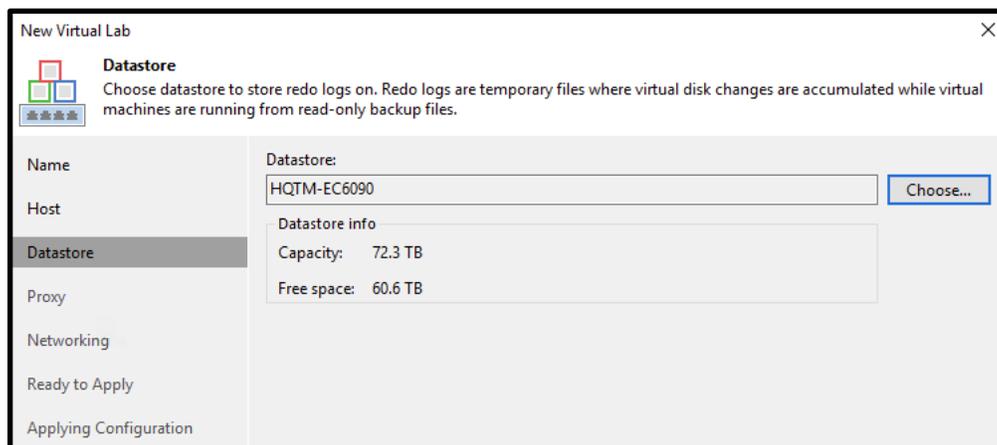


Figure 19 - Virtual lab datastore

A Tintri system can be used for hosting the datastore used by a virtual lab. Deploying the datastore used by a virtual lab on a Tintri system enables the use high performance primary storage to store redo logs, the temporary files where virtual disk changes are accumulated while VMs are running from read-only backup files.

The option to create a proxy appliance is presented during virtual lab configuration. The proxy appliance provides the Veeam backup server with access to the virtual machines running in the virtual lab.

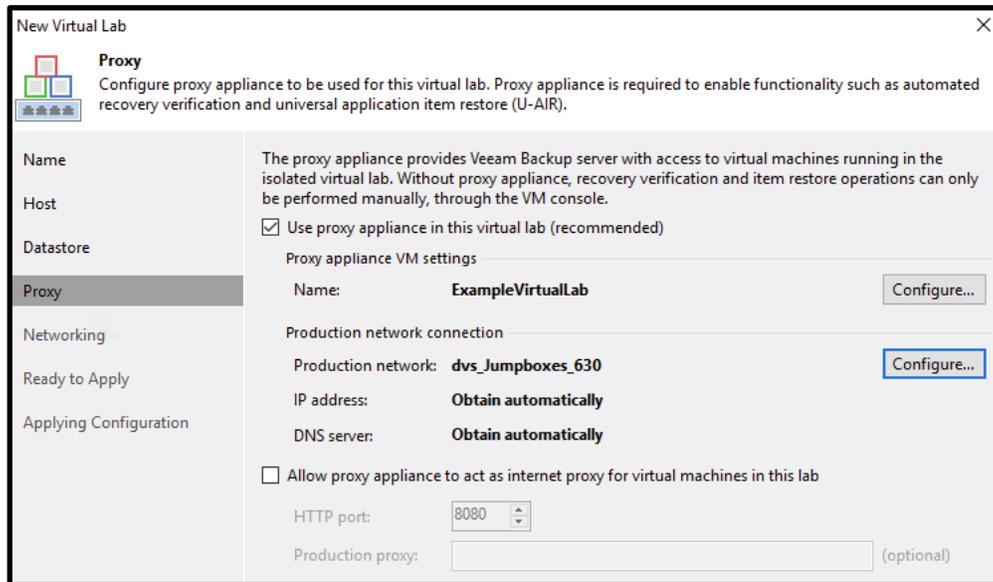


Figure 20 - Virtual lab proxy

The proxy appliance enables communication between the production environment and the isolated network(s) in the virtual lab. The proxy appliance is a Linux-based VM that is deployed on the ESXi host where the virtual lab is deployed.

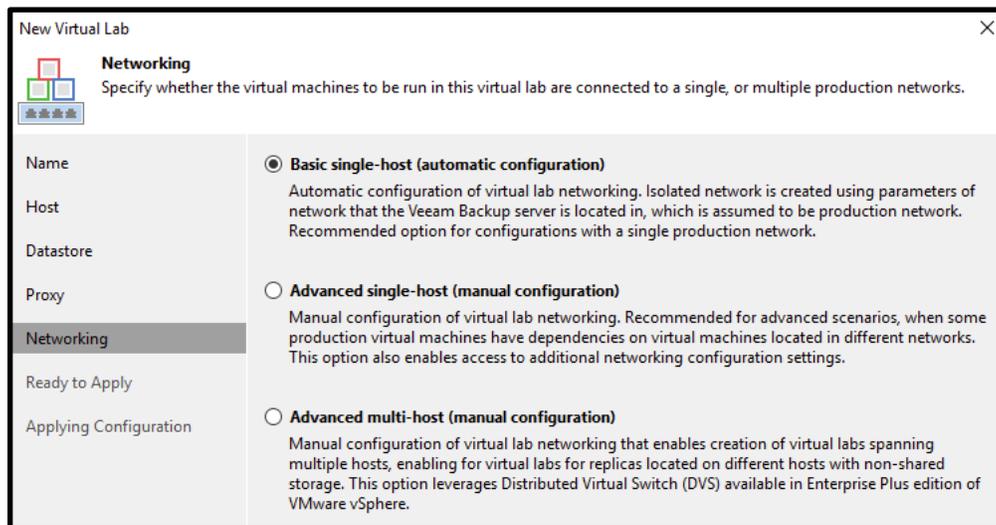


Figure 21 - Virtual lab networking

A virtual lab is fenced off from the production environment and provides advanced networking deployment options. The network configuration of the virtual lab mirrors the network configuration of the production environment. For additional information see the “SureBackup Recovery Verification” section of the “Veeam Backup & Replication User Guide for VMware vSphere Environments”.

Creating Backup Copies with Veeam

A comprehensive data protection strategy includes the creation of additional copies of backups that can be retained offsite, and backup copies on different media types. Veeam suggests following a simple “3-2-1” rule where 3 copies of important data are retained, storing the data on 2 different media types, and keeping 1 backup copy offsite. Veeam provides a variety of choices to assist in adhering to the “3-2-1” rule:

- Veeam backup copy jobs with WAN acceleration – Protect Veeam backups by copying them to an offsite backup repository using WAN accelerators to minimize replication network bandwidth utilization. This serves to create an additional copy of backups offsite. If the repository storage types are different, this solution may result in the creation of 2 backup copies on different media types as well as 1 offsite copy. Note that Veeam WAN acceleration is available in the Enterprise Plus edition of Veeam Backup & Replication.
- Veeam tape copy jobs – Copy Veeam backups to tape for offsite archiving. This solution fulfills the goal of having 2 different media types and 1 offsite copy if tapes are being stored in an offsite location.
- Veeam Cloud Connect – Use a Veeam Cloud Connect Service Provider partner to copy backups to offsite hosted backup repositories. This solution fulfills the goal of having 1 offsite copy. Dependent on the Cloud Connect media type, the solution may also fulfill the goal of creating backups on 2 different media types.

Summary

Veeam Backup & Replication is a comprehensive data protection solution for virtualized environments hosted on one or more Tintri systems. Veeam Backup & Replication can easily leverage multiple virtual disk transport modes, enabling a variety of data protection strategies. High performance Tintri systems are also an excellent choice for use as vPower NFS write cache storage, as well as virtual lab datastores.

References

1. Veeam references:

“Veeam Backup & Replication User Guide for VMware vSphere Environments”

“Veeam Best Practices for Deployment and Configuration (VMware)”

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

2. VMware vSphere Documentation Center:

“Virtual Disk Transport Methods”

https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vddk.pg.doc_50%2FvddkDataStruct.5.5.html

3. Tintri references:

“VMstore Overview”

“Managing VM Data with Tintri”

“NFS Storage Best Practices for IT Administrators”

“Data Protection Overview and Best Practices with Tintri VMstore and Tintri Global Center”

<http://www.tintri.com/resources>

“Tintri VMstore System Administration Manual”

“Tintri Automation Tool Kit Quick Start & Overview Guide”

<https://support.tintri.com/>

Appendix A: Snapshot Collisions – Tintri Snapshots

This appendix item considers what may happen if a given VM is protected with both Veeam Backup & Replication backups and Tintri snapshot backups.

It is important to understand the implications of deciding to co-mingle Tintri snapshots and Veeam Backup & Replication together to protect the same VM or VMs. Both data protection methods can potentially use vSphere snapshots. In cases where a Veeam requested vSphere snapshot occurs at about the same point in time as a Tintri requested VM-consistent snapshot, a snapshot collision may occur. It is possible that one of the two snapshots will fail. If deciding to protect one or more VMs with both data protection methods, schedule them to occur such that they do not overlap. If an overlapping schedules cannot be avoided, consider using Tintri crash-consistent snapshots as they will not invoke a vSphere snapshot.

-
- *Avoid configuring overlapping backup schedules with Veeam Backup & Replication and a Tintri snapshot schedule that takes VM-consistent snapshots of the same VM or VMs.*
 - *Consider using Tintri crash-consistent snapshots in cases where overlapping a Veeam Backup & Replication schedule with a Tintri snapshot schedule cannot be avoided.*
-

In cases where a Veeam Backup & Replication vSphere snapshot has been created, and a Tintri VM-consistent or crash-consistent snapshot occurs, the Tintri snapshot will contain the Veeam backup temporary snapshot. Although neither backup task failed, recovery from or cloning of the Tintri snapshot will also recover the Veeam backup temporary snapshot. The temporary snapshot is effectively orphaned should this occur.

A potential alternative to co-mingling Tintri snapshots with Veeam Backup & Replication to protect the same VM or VMs, is to employ a strategy where some VMs are protected with Veeam, and the remaining VMs are protected with Tintri snapshots.

Appendix B: Tintri VM Level Quality of Service

Tintri systems feature the ability to configure QoS (Quality of Service) settings at a VM level. Normalized maximum and minimum IOPS values can be specified. The settings apply to all VMDKs on a given VM. The maximum IOPS setting, when configured, can limit the IOPS available for VM I/O which may impact backup data transfer rates.

VM level QoS is easily configured from within the Tintri system user interface. In the example provided below, maximum normalized IOPS has been set to a value of 1000 resulting in an effective MBps data transfer rate of 8.2.

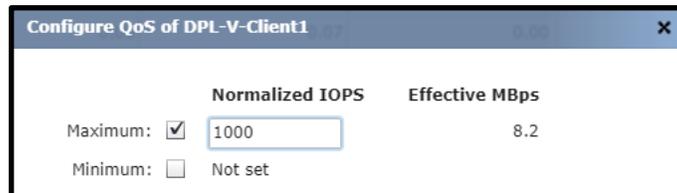


Figure 22 – Tintri system - Configure QoS

Within the Tintri system user interface it is easy to see if QoS limits have been set on a VM. In the example provided below, the VM named “DPL-V-Client1” has been configured with a maximum normalized IOPS value of 1000.

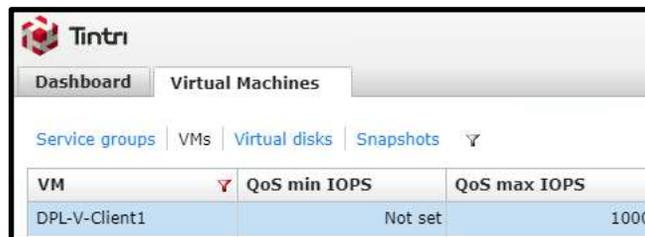


Figure 23 - Viewing QoS settings

VM level QoS can also be configured and viewed from within the Tintri Global Center user interface. In the example provided below, maximum normalized IOPS has been set to a value of 1000 resulting in an effective MBps data transfer rate of 8.2.

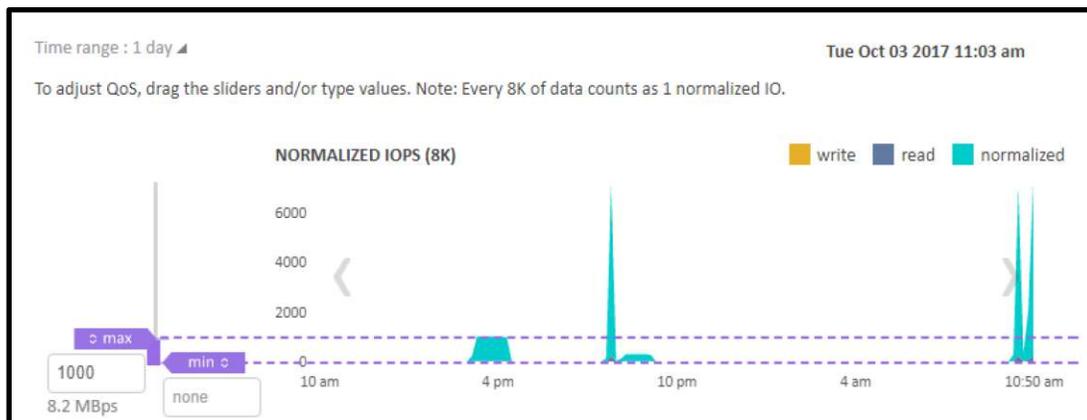


Figure 24 - Tintri Global Center - Configure QoS

Note that Tintri Global Center policy management settings may impact the ability to view QoS settings that have been configured using the Tintri system user interface. In the example provided below, Tintri

Global Center policy management has been configured to accept changes that have been applied from within the Tintri system user interface.

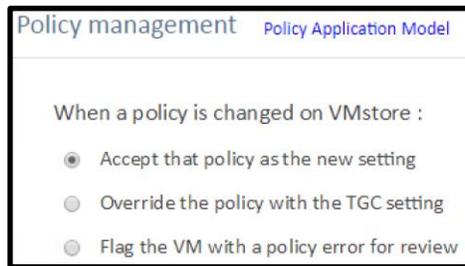


Figure 25 - Tintri Global Center - Policy management

When the “override” Tintri Global Center policy management setting is selected, QoS settings applied from within the system user interface will not take effect, but will instead be overridden by Tintri Global Center settings.

When the third option, “flag the VM with a policy error” is selected, QoS settings applied from within the Tintri system user interface will be applied. However, these QoS settings will not be reflected from within the Tintri Global Center user interface.

Quality of Service Automation with Veeam Backup & Replication

VM level QoS maximum normalized IOPS settings may reduce the data transfer rate that can be achieved when a VM is being backed up. Limiting maximum normalized IOPS also limits the effective data transfer rate at which data can be read when a backup is being performed. The net result is that backup duration may elongate, which may introduce additional challenges.

Users that have configured QoS on one or more VMs may wish to temporarily remove any QoS settings when backups execute. Ideally, the removal of QoS settings should occur only on VMs that are being backed up, not on all VMs resident on a given Tintri system. Additionally, the original QoS settings should be re-applied when the VM backup process completes.

The Tintri PowerShell Toolkit provides a granular suite of cmdlets that can be used in conjunction with the Veeam PSSnapin to automate the removal of QoS settings on a VM when the VM is being backed up. The same suite of cmdlets can be used to reapply the original QoS settings when the VM backup has completed.

An example Windows PowerShell script is provided at:

<https://github.com/Tintri/tintri-powershell-examples/tree/master/Veeam>

The script, “Veeam_Backup_Tintri_QoS.ps1” is provided “as is” as an example that can be customized to meet user requirements.

The example script output provided below depicts QoS settings being discovered and removed from VMs that are being actively backed up.

```

2017-10-03 10:42:15 [DEBUG] VeeamClients-A: Running (968bc0ca-7cfc-4507-b19f-634888c89d52)
2017-10-03 10:42:15 [Info ] Managing running job: VeeamClients-A - 968bc0ca-7cfc-4507-b19f-634888c89d52
2017-10-03 10:42:15 [DEBUG] VeeamClients-A already present
2017-10-03 10:42:15 [DEBUG] Returning task sessions
2017-10-03 10:42:15 [Info ] Manage 2 running tasks
2017-10-03 10:42:15 [Info ] DPL-V-Client1 - InProgress - cbce16d5-7598-499a-93dc-366d47607e92
2017-10-03 10:42:16 [Info ] Clearing QoS DPL-V-Client1 for back-up from 150, 1500
2017-10-03 10:42:16 [DEBUG] DPL-V-Client1 is in progress
2017-10-03 10:42:16 [Info ] DPL-V-Client2 - InProgress - f6a6d9da-6fb6-4a99-9aa9-f9b4b92f10b6
2017-10-03 10:42:16 [Info ] Clearing QoS DPL-V-Client2 for back-up from 200, 2000
2017-10-03 10:42:16 [DEBUG] DPL-V-Client2 is in progress

```

Figure 26 - Clearing QoS from active VM backups

The example script output provided below depicts detection of job completion and reapplies QoS settings.

```

2017-10-03 10:50:05 [DEBUG] VeeamClients-A: Running (968bc0ca-7cfc-4507-b19f-634888c89d52)
2017-10-03 10:50:05 [Info ] Managing running job: VeeamClients-A - 968bc0ca-7cfc-4507-b19f-634888c89d52
2017-10-03 10:50:05 [DEBUG] VeeamClients-A already present
2017-10-03 10:50:05 [DEBUG] Returning task sessions
2017-10-03 10:50:05 [Info ] Manage 2 running tasks
2017-10-03 10:50:05 [Info ] DPL-V-Client1 - Success - cbce16d5-7598-499a-93dc-366d47607e92
2017-10-03 10:50:05 [Info ] Set DPL-V-Client1 QoS to 150, 1500
2017-10-03 10:50:05 [Info ] DPL-V-Client1 is stopped
2017-10-03 10:50:05 [Info ] DPL-V-Client2 - Success - f6a6d9da-6fb6-4a99-9aa9-f9b4b92f10b6
2017-10-03 10:50:06 [Info ] Set DPL-V-Client2 QoS to 200, 2000
2017-10-03 10:50:06 [Info ] DPL-V-Client2 is stopped

```

Figure 27 - Reapplying QoS to completed VM backups

Appendix C: Backup Repository Deployment - Tintri System

While not necessarily a best practice recommendation, this section provides information and considerations related to deploying a Veeam backup repository on a Tintri system.

To begin with, the following guideline should be well understood:

- When protecting VMs that reside on a Tintri system, they should not be backed up to the same Tintri system configured as a backup repository within Veeam Backup & Replication. Avoiding this scenario is highly recommended. In the unlikely event of a Tintri system outage, VMs and the backups of those VMs become inaccessible when both entities are stored on the same array.

-
- *Do not backup VMs residing on a Tintri system to a backup repository residing on the same Tintri system. VMs being protected by Veeam Backup & Replication should always use a backup repository residing on a different storage device.*
-

Repository Type

Two different backup repository types can be created on a Tintri system; Microsoft Windows server, or Linux server.

Windows based repositories are typically preferred because they can also be configured to function as vPower NFS servers. In this use case, Veeam Backup & Replication will run the Veeam vPower NFS service directly on the backup repository and provide ESXi hosts with transparent access to backed up VM images stored on the repository.

Repository Path

A backup repository includes a storage location where backups are stored. When creating a backup repository on a Tintri system, that path can point to a virtual disk that also resides on a Tintri system, or to other storage. Examples of the storage that can be used include a local virtual disk, a virtual disk residing on a non-Tintri datastore, or an in-guest iSCSI LUN.

Note that deciding to use a virtual disk residing on a non-Tintri datastore or an in-guest iSCSI LUN precludes the ability to use Tintri native snapshots to protect the backup repository. Protecting a backup repository residing on a Tintri system with native Tintri snapshots is not specifically recommended as a best practice.

A number of significant storage compatibility settings can be configured on a backup repository. The settings deployed should be based on the type of storage being used for the backup repository storage location.

Tintri systems are available in both “Hybrid-Flash” and “All-Flash” models.

If the backup repository deployment is using Tintri Hybrid-Flash storage as a repository storage location, the following settings are recommended:

- “Align backup file data blocks” should be disabled. Tintri Hybrid-Flash storage does not deduplicate data stored on disk. This setting should not be enabled, as it may increase space usage and data fragmentation.
- “Decompress backup data blocks before storing” should be enabled on Tintri T800 series products. By design, the T800 series of Tintri Hybrid-Flash storage compresses all stored data. On Tintri T540

and T600 series products “Decompress backup data blocks before storing” should be disabled. The T540 and T600 series of Tintri Hybrid-Flash storage do not provide compression.

- “This repository is backed by rotated hard drives” should be disabled. A Tintri system virtual disk is attached to the backup repository VM as a VMDK, not as a physical hard drive that may be rotated in and out of service.
- “Use per-VM backup files”, this setting should be carefully considered. When this setting is enabled, a backup job will use a separate write stream for every VM in the job and store its data to a separate backup file. As a result, resources of the storage device will be used more efficiently, and job performance may increase. However, Veeam Backup & Replication will not deduplicate data between VMs within a given job. Enabling or disabling this setting amounts to a tradeoff between performance and storage space utilization.

If the backup repository deployment is using Tintri All-Flash storage as a repository storage location, the following settings are recommended:

- “Align backup file data blocks” should be enabled. Tintri All-Flash storage deduplicates data, enabling this setting may increase deduplication ratios.
- “Decompress backup data blocks before storing” should be enabled. By design, Tintri All-Flash storage compresses all stored data.
- “This repository is backed by rotated hard drives” should be disabled. A Tintri system virtual disk is attached to the backup repository VM as a VMDK, not as a physical hard drive that may be rotated in and out of service.
- “Use per-VM backup files” should be enabled. When this setting is enabled, a backup job will use a separate write stream for every VM in the job and store its data to a separate backup file. As a result, resources of the storage device will be used more efficiently, and job performance may increase.

If the backup repository deployment is using another storage vendor’s product as a repository location, please consult with the appropriate vendor to determine what settings should be used.

Data Reduction & Compression Settings

When a new backup job is created, inline data deduplication is enabled by default. Also by default, the compression level is set to “Optimal”. The effect of these settings typically decreases network traffic and disk space consumption on a backup repository.

If the target backup repository storage location for the job is using a storage device that supports hardware compression, deduplication, or both, these default settings should be altered. For instance, the “Enable inline data deduplication” setting can be disabled and the “Compression level” setting can be set to “None”.

For example, if the backup repository storage location is a Tintri All-Flash product, the inline deduplication setting should be disabled, and the compression level should be set to “None”. Setting the compression level to “None” assumes that any required network link between a backup proxy and backup repository will have adequate available bandwidth so as not to impact backup or recovery performance.

If the backup repository storage location is a Tintri Hybrid-Flash series T800 product, the compression level setting can be set to “None” because the Tintri system compresses stored data by default. Setting the compression level to “None” assumes that any required network link between a backup proxy and backup repository will have adequate available bandwidth so as not to impact backup or recovery performance.

If the backup repository storage location is not a Tintri product, please consult with the appropriate vendor to determine what settings should be used.

© 2017 Tintri, Inc. All rights reserved. Tintri, Tintri VMstore, Tintri Global Center, ReplicateVM, SecureVM, and SyncVM are trademarks of Tintri, Inc., and may be registered in the U.S. Patent and Trademark Office and in other jurisdictions. All other marks appearing in this publication are the property of their respective owners.

Tintri believes the information in this document is accurate as of its publication date. The information in this publication is provided as is and is subject to change without notice. Tintri makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose.



303 Ravendale Drive
Mountain View CA 94043
+1 650.810.8200
info@tintri.com