Hewlett Packard Enterprise

Technical white paper

Hybrid cloud data protection with HPE, Veeam, and Microsoft Azure

Contents

Executive summary

Having a flexible, efficient, and scalable data protection model is no longer optional in today's data driven economy… it is a requirement. Business continuity and disaster recovery (BCDR) solutions must protect against all disasters, whether small or large, local or regional, natural or man-made. From software bugs and human error to larger natural disasters like earthquakes and floods, BCDR solutions are essential to eliminate data loss and maintain business continuity.

In recent years, there has been a dramatic increase in virtual machine (VM) adoption with VMs now powering much of today's IT and data centers. The use of modern hypervisors adds a layer of complexity that traditional BCDR solutions never addressed, so a BCDR solution that is hypervisor-aware can drastically improve upon traditional bare-metal solutions. Older backup solutions that simply protect data at the LUN, volume, or file system layer make the restoration of VMs a tedious and painful experience. A more sophisticated hypervisor-aware solution like Veeam Backup & Replication can protect at the VM level, taking into account the state, context, and configuration of the entire VM. In this case, restoring a VM can be as simple as pushing a button.

Implementing BCDR using a single solution can be easily achieved in purely homogeneous data center. Unfortunately, homogeneous environments are rare in the real world with siloed, disjoint, and complex solutions much more commonplace. A single BCDR solution that can protect heterogeneous virtual environments without affecting production can protect against disasters, changing business requirements, and even future technological trends. Veeam, HPE, and Microsoft® can provide all this and more with powerful compute, all-flash performance, and advanced data protection for both on-premises (on-prem) and cloud deployments.

Recovery-time objectives (RTOs), recovery-point objectives (RPOs), and the cost of data protection are common metrics used to describe an organization's ability to recover data and associated services. Modern business is driven by applications and services with very different RTO, RPO, and cost requirements, so a "one size fits all" data protection scheme seldom exists. Optimization of cost and function can be achieved by placing data in purpose-built "tiers" designed to suit individual data protection needs. "Hybrid IT" is the future of computing and BCDR is no exception. Today's businesses are faced with a choice of whether to continue the maintenance of data in-house or migrate it to the cloud. The typical outcome is that both on-prem and public cloud tiers are necessary. Businesses often need to store data in their own data centers due to corporate policies, regulations, service level agreements (SLAs), and security considerations. An on-premises backup tier can minimize recovery time, while maximizing the agility and control a company has over its critical data. For applications with less stringent requirements, a public tier can minimize the cost of data protection. With the addition of a public cloud tier, a business can even follow the 3-2-1 rule by having at least three copies of data, in two different formats, with one residing off-site in the cloud.

Traditionally, this tiered protection model has been achieved using several disjoint BCDR solutions, which can add to the complexity, cost, and skills required to protect a company's data. In this reference architecture, HPE, Veeam, and Microsoft provide a single integrated and proven solution that eliminates these silos while optimizing the cost and function of BCDR.

Solution overview

Veeam Backup & Replication software powered by HPE and Microsoft Azure creates a business continuity and disaster recovery (BCDR) solution that is unmatched. Not only can the performance and control of having data on-premises be realized, but also the flexibility and cost advantages of utilizing the cloud. The performance and advanced feature set of HPE Nimble Storage can help businesses scale to meet the needs of their expanding customer base. The efficiency of the HPE StoreOnce with its innovative Catalyst software minimizes data footprint and expedites recovery time of business critical applications. The addition of Microsoft Azure and Veeam Cloud Connect provides businesses the flexibility and cost optimization to meet any current and future requirement.

In this reference architecture, a "Production Tier," an "On-Premises Backup Tier," and a "Public Cloud Tier" are used. Each tier has its own unique hardware, performance, and cost attributes. The production tier consists of HPE ProLiant servers and VMware® ESX® for compute, the HPE SN6000B for SAN infrastructure, and HPE Nimble Storage for scalable primary storage. A Microsoft SQL Server database and client-based I/O generation tool (HammerDB) will be used to demonstrate the backup and recovery of an actual in-use application. The on-premises backup tier is powered by Veeam Backup & Replication and an HPE ProLiant server, all connected to an HPE StoreOnce secondary storage appliance capable of RPOs and RTOs that will exceed any expectation. The public cloud tier consists of Microsoft Azure and Veeam Cloud Connect software, which brings unparalleled off-premises data protection to this already capable solution. Each component and tier will be discussed in-depth in the following pages.

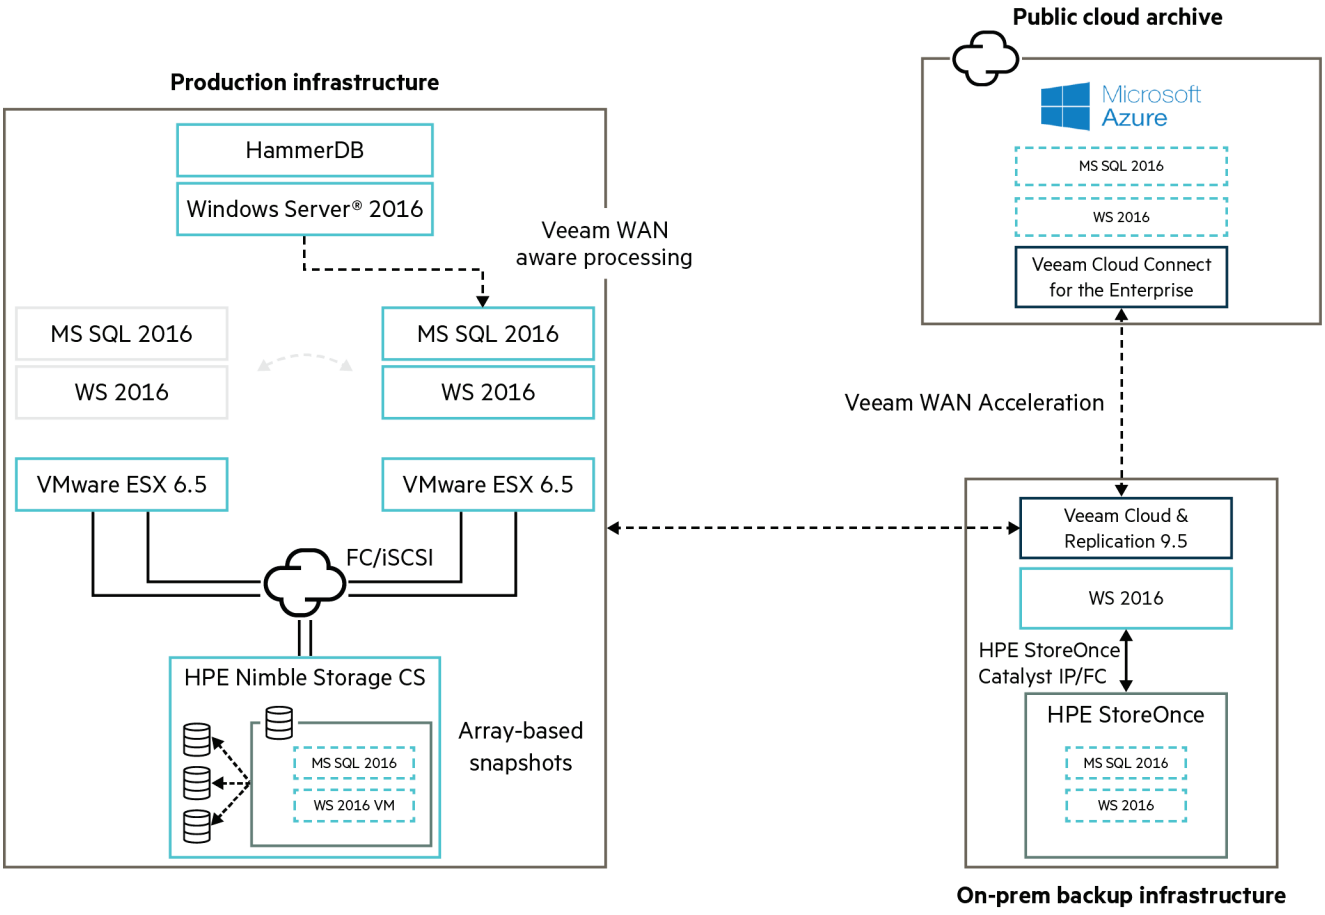

Figure 1. Optimization using multi-tiered backup

Note

Specific hardware and software versions were used in this reference architecture. Although some versions may be detailed, support is not limited to these versions unless explicitly stated. It is encouraged to use the latest HPE, Veeam, and Microsoft products available at the time of implementation as greater performance and benefit is likely.

Solution components

Veeam Backup & Replication

Veeam Backup & Replication can be installed as either a stand-alone product or part of the Veeam Availability Suite. Veeam Backup & Replication is an agentless backup solution that can be used to protect VMware vSphere® and Microsoft Hyper-V virtual machines (VMs). It provides image-level VM backups that can be used to restore entire VMs or individual files within a virtual machine's file system. Veeam Backup & Replication can also replicate VM images for use in on-site or off-site disaster recovery. The storage integration feature of Veeam Backup & Recovery can be used to improve upon its traditional techniques. Veeam storage integration works in conjunction with many of the advanced features of HPE Storage to provide an efficient and bulletproof data protection solution for almost any virtualized environment.

Traditional backup

Veeam Backup & Replication groups all information and processes into a definable unit called a job. A job is where the user will define the configuration, schedule, and type of task to be performed. Each job will contain actions, which can be considered subroutines necessary to accomplish the task.

The process of taking a backup first leverages the native snapshot capabilities of the hypervisor to create a consistent copy of a virtual machine. The technique used to create the initial VM snapshot varies depending on the type of host. Veeam uses native vSphere snapshot functionality with VMware ESX, while using Microsoft Volume Shadow Copy Service (VSS) in Microsoft Hyper-V environments. The resulting copy is then used to back up the virtual machine to a backup repository. Changed block tracking (CBT) is used to minimize the data transfer size of subsequent backups, while deduplication and compression can be used to reduce overall storage consumption. The Veeam Data Mover service is responsible for retrieving VM data from the source and writing to the target or backup repository. The backup repository can be backed by internal server, SAN, or even tape storage media. Veeam provides a rich feature set that includes, among other things, data encryption, WAN acceleration, compression, and deduplication.

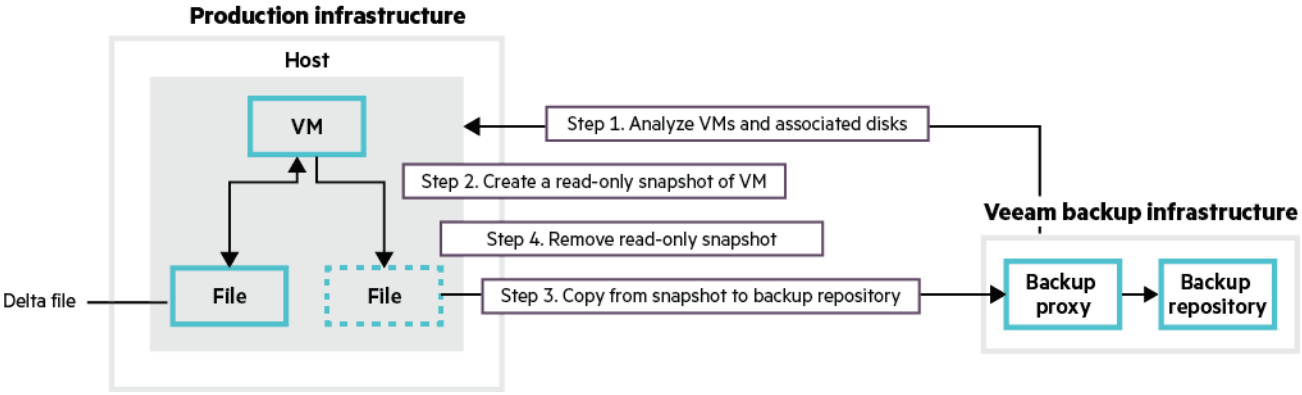

Figure 2. Traditional VM backup using Veeam Backup & Replication

Veeam Backup & Replication is a proven and magnificent data protection option for any virtualized data center solution. Veeam traditional backup capabilities are software-based, which decouples the solution from the underlying infrastructure. Although this flexibility is advantageous, integrating Veeam with the underlying storage infrastructure can provide additional backup performance, recovery-point objectives (RPOs), and recovery-time objectives (RTOs).

Back up with HPE storage integration

There are downsides to using server resources for snapshot creation and data transfer. Compute resources are required to provide services to the end user, so consuming them in backup operations can negatively affect business outcomes and service-level agreements (SLAs). If a host is asked to make a VM snapshot, process ongoing changes, and transfer the backup data, it may have less CPU, memory, and network bandwidth to apply to normal operations. This type of backup is commonly referred to as “on-host” as the backup and data transfer occurs solely on the host infrastructure.

The following are graphs generated by Windows® Performance Monitor (PerfMon) showing processor and network adapter utilization during traditional backup. Notice a spike in utilization during the backup operation.

Process utilization

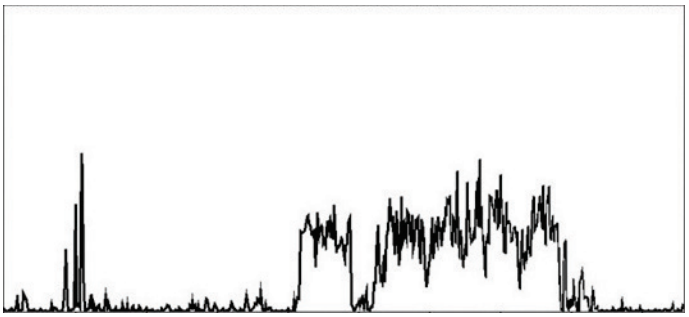

Figure 3. Processor utilization captured on a Microsoft Hyper-V host during traditional backup

Network adapter utilization

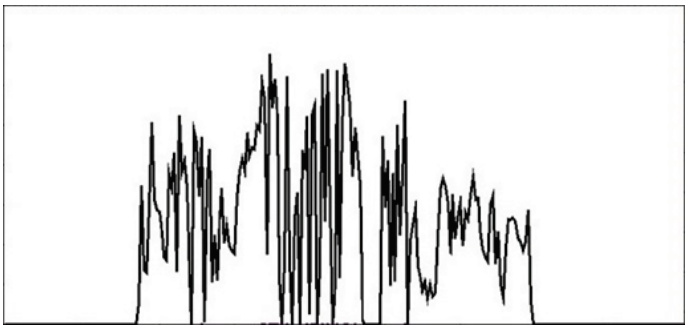

Figure 4. Network utilization captured on a Microsoft Hyper-V host during traditional backup

Veeam integration with HPE Storage minimizes this overhead by using array features to offload most of these traditional backup activities. Using this integration, Veeam Backup & Replication will trigger a hardware assisted storage snapshot taking advantage of the latest HPE Storage data protection features. Using the storage area network (SAN), the storage snapshot is exported directly to the backup infrastructure and used for the data transfer. This type of backup is referred to as "off-host" as the host is offloaded from much of the operation. In this scenario, the snapshot capabilities of the hypervisor are only utilized until the storage snapshot has been taken. Using HPE Storage for fast snapshot creation and an existing SAN for data transfer can greatly reduce the compute and network impact of data protection operations.

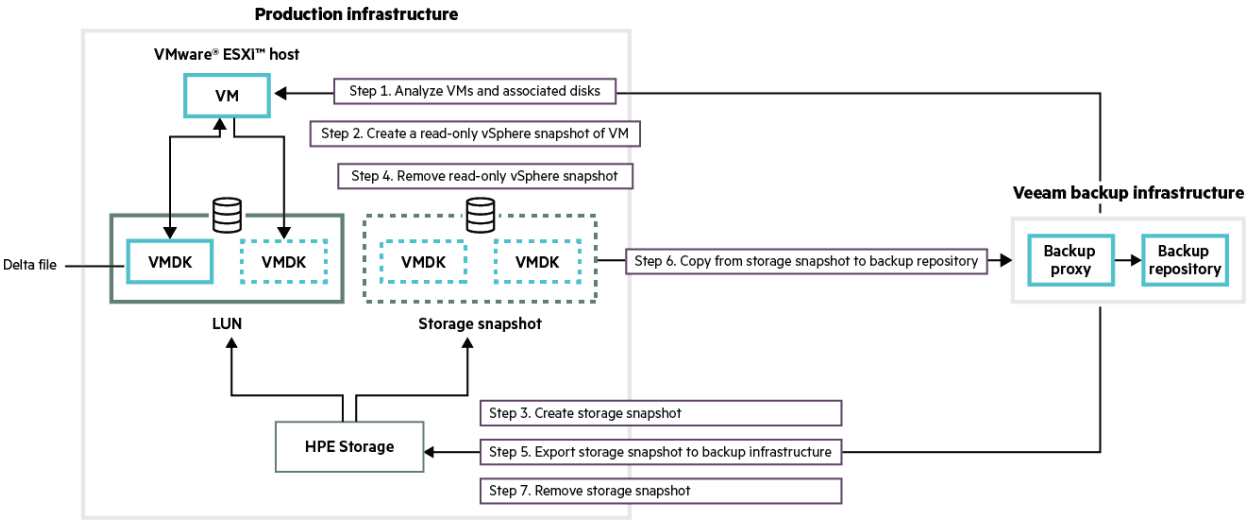

Figure 5. VMware ESX backup using storage snapshots

The following are graphs generated by Windows Performance Monitor (PerfMon) showing processor and network adapter utilization during the backup from an HPE StoreServ Virtual Copy snapshot. Notice that there is little to no utilization during the offloaded backup operation.

Process utilization

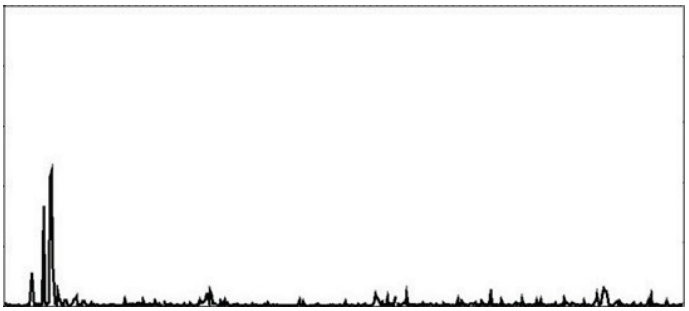

Figure 6. Processor utilization captured on a Microsoft Hyper-V host during the backup from an HPE StoreServ Virtual Copy snapshot

Network adapter utilization

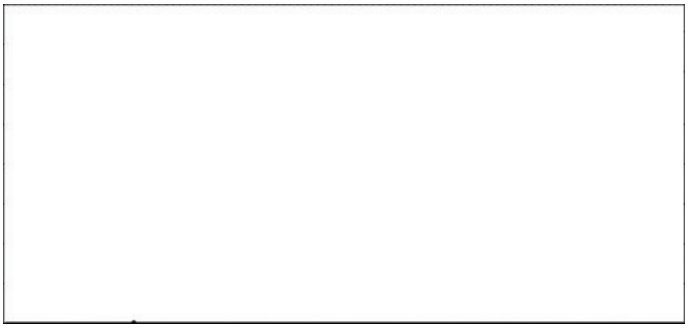

Figure 7. Network utilization captured on a Microsoft Hyper-V host during the backup from an HPE StoreServ Virtual Copy snapshot

As on-host backup activities do not utilize storage snapshots, they are largely storage agnostic and will work similarly with any storage product. This paper focuses on the interoperability and advantages of using off-host backup with Veeam and HPE Storage.

HPE ProLiant DL380 servers

IT needs to operate at the speed of today's business to be an accelerator of new ideas, products, and services. It's all about faster time to value, and today the difference between winning and losing is how fast a company can turn ideas into profit. For companies to be successful, speeding time to value will require a hybrid infrastructure that is fast, flexible, and scalable whether your applications and data are in your data center or a hosted cloud. The reality is that the Right Mix of Hybrid IT infrastructure depends on your workloads and business requirements.

HPE ProLiant servers can help you focus on key areas of IT transformation that will help you increase agility and flexibility, reduce costs, grow revenue and profits, manage risk, and improve your customers and employees' experience.

The data center standard for general-purpose compute, the HPE ProLiant DL380 Gen9 Server delivers the best performance and expandability in the HPE 2P rack portfolio. Reliability, serviceability, and continuous availability, backed by a comprehensive warranty, make it ideal for any environment. Designed to reduce costs and complexity, leveraging Intel®'s latest E5-2600 v4 processors with 21% performance gain, plus the latest HPE 2400 MT/s DDR4 SmartMemory supporting 3.0 TB and up to 23% performance increase. It has flexible storage with up to 24 SFF, 12 LFF or 6 NVMe PCI SSDs and industry-leading technologies like HPE Persistent Memory for databases and analytic workloads.

All HPE ProLiant servers offer a portfolio of options that are tested and qualified by HPE and deliver increased performance, flexibility to grow with the business and complete end-to-end support. Integrated with a simplified, but comprehensive management suite and industry-leading support, the ProLiant DL380 Gen9 helps increase IT staff productivity, and accelerates service delivery.

HPE Nimble Storage

The Power of Predictive—The HPE Nimble Storage platform leverages flash storage and the power of predictive analytics to deliver fast and reliable access to data. This approach closes the app-data gap and radically simplifies operations. HPE InfoSight predictive analytics predict and prevent issues to help deliver greater than 99.9999% measured availability. Utilize a single multicloud architecture to flexibly deploy workloads on flash arrays, converged infrastructure, and the public cloud.

HPE StoreOnce storage

StoreOnce is Hewlett Packard Enterprise's line of secondary storage designed for robust protection of all data types stored on HPE 3PAR, HPE Nimble Storage, HPE SimpliVity and other HPE and third-party primary storage products. The products are tightly integrated with data protection software and business applications to deliver predictable and fast RTO and RPO performance when data is lost or corrupted. StoreOnce Systems are available as purpose built backup appliances and virtual appliances with capacities starting at 5.5 TB and increasing up to 1.7 PB on the largest model. These usable capacities are expanded with class-leading data compaction enabled by StoreOnce deduplication and compression. This can deliver space savings in excess of 95% meaning the maximum effective capacity of StoreOnce Systems is in excess of 30 PB. StoreOnce replication provides deduplication-optimized copies to one or more StoreOnce Systems, at remote locations, for additional, off-site, protection of data.

StoreOnce offers flexible connection to data protection applications. This can be via general-purpose NAS shares, emulations of tape products (VTL) or the data protection optimized Catalyst interface. StoreOnce Catalyst enables backup data to be deduplicated at source and for the data protection application to control StoreOnce replication. This provides users with additional performance and control over traditional NAS and VTL connections. HPE Recovery Manager Central Software (RMC) software is integrated with StoreOnce Catalyst to deliver high-speed data protection for application data stored on HPE Storage arrays by managing array-based snapshots and Express Protect backups from snapshots to StoreOnce Systems. RMC Express Restore enables highly efficient recovery of deleted snapshots on a HPE array from a backup image on StoreOnce. RMC combined with StoreOnce is the fastest way to get application consistent data backed up from and restored to HPE arrays. Catalyst plug-ins are available directly connect Oracle, SAP HANA® and SQL Server applications to StoreOnce Systems for high efficiency protection controlled by the application administrator.

StoreOnce can be evaluated using the StoreOnce VSA (virtual appliance) Freeware and more information is available at hpe.com/Storage/StoreOnce.

Microsoft Azure

Azure is the only consistent hybrid cloud

Build and deploy wherever you want with Azure, the only consistent hybrid cloud on the market. Connect data and apps in the cloud and on-premises—for maximum portability and value from your existing investments. Azure offers hybrid consistency in application development, management and security, identity management, and across the data platform.

- Extend Azure on-premises and build innovative, hybrid apps with Azure Stack.
- Connect on-premises data and apps to overcome complexity and optimize your existing assets.
- Distribute and analyze data seamlessly across cloud and on-premises.

Azure is the cloud for building intelligent apps

Use Azure to create data-driven, intelligent apps. From image recognition to bot services, take advantage of Azure data services and artificial intelligence to create new experiences—that scale—and support deep learning, HPC simulations, and real-time analytics on any shape and size of data.

- Develop breakthrough apps with built-in AI.
- Build and deploy custom AI models at scale, on any data.
- Combine the best of Microsoft and open source data and AI innovations.

Azure is the cloud you can trust

Ninety percent of Fortune 500 companies trust the Microsoft Cloud. Join them. Take advantage of Microsoft security, privacy, transparency, and the most compliance coverage of any cloud provider.

- Achieve global scale on a worldwide network of Microsoft-managed data centers across 42 announced regions.

- Detect and mitigate threats with a central view of all your Azure resources through Azure Security Center.
- Rely on the cloud with the most comprehensive compliance coverage (50 compliance offerings), and recognized as the most trusted cloud for U.S. government institutions.

Microsoft SQL Server 2016

With SQL Server, you can build intelligent, mission-critical applications using a scalable, hybrid database platform that has everything built-in from in-memory performance and advanced security to in-database analytics. SQL Server has industry-leading performance and security and runs natively on Windows, Linux®, and Docker containers.

Your choice of language and platforms

Build modern applications using the language of your choice, on-premises and in the cloud, now on Windows, Linux, and Docker containers.

Industry-leading performance

Take advantage of breakthrough scalability, performance, and availability for mission-critical, intelligent applications and data warehouses.

Least vulnerable database

Protect data at rest and in motion with the least vulnerable database over the last seven years in the NIST vulnerabilities database.

Real-time intelligence

Gain transformative insights for your business with real-time analytics at up to 1M predictions/second.

End-to-end mobile BI

Turn raw data into meaningful reports that can be delivered to any device—at one-fifth the cost of other self-service solutions.

HammerDB

HammerDB is a free open source load generation and benchmarking tool that supports a variety of database applications, including Microsoft SQL Server 2016. HammerDB's highly customizable workload options allow an Online Transaction Processing (OLTP) workload to be generated that stresses a database and underlying storage. HammerDB is typically installed on an external client that connects to the database to generate load by adding and removing database entries. In addition, it can be run in a master-slave relationship allowing multiple test clients to generate load simultaneously breaking beyond the load generation capabilities of a single system and allowing the load to scale to meet the most demanding needs.

Solution tiers

Production tier

The production tier will accommodate production hardware, applications, and services. These applications drive business critical functions and revenue generation. The hybrid cloud data protection model will be used to protect these assets and maintain business continuity in the case of a disaster. The advanced data protection features of HPE servers, switches, and storage can help protect these assets and ensure high availability. Features like thin provisioning, deduplication, and snapshots can be used to maintain efficiency and safeguard against outages or data loss.

In this reference architecture, the production site will be proven using the following hardware and software. Although this hardware was used in this demonstration, even greater performance and functionality can be realized with the latest HPE servers, switches, and storage.

Table 1. Production tier resource

Role	Hardware	Software	Notes
ESX Server 1	HPE ProLiant DL380 Gen9	VMware ESX 6.5	
ESX Server 2	HPE ProLiant DL380 Gen9	VMware ESX 6.5	
FC Switch	HPE SN6000B	Fabric OS 8.0.2a	
Primary Storage	HPE Nimble Storage	HPE NimbleOS 3.8.1	
SQL Server		Windows Server 2016 Windows SQL Server 2016	
Test Client		Windows Server 2016 HammerDB	

Two HPE ProLiant DL380 Gen9 servers will be used as compute within the production tier. The VMware ESX hypervisor will be used to provision Windows Server 2016 VMs. These VMs will accommodate Microsoft SQL Server 2016, which will act as the production database application. The ESX servers are redundant, so that in the case of a server failure all VMs will remain available while backup operations resume. The production database will be protected and recovered using Veeam Backup & Replication. Veeam Backup & Replication is agentless, which means no software installation is required on the hosts or VMs. Veeam application-aware processing will be used to quiesce I/O and put the database in a consistent state in preparation for backup.

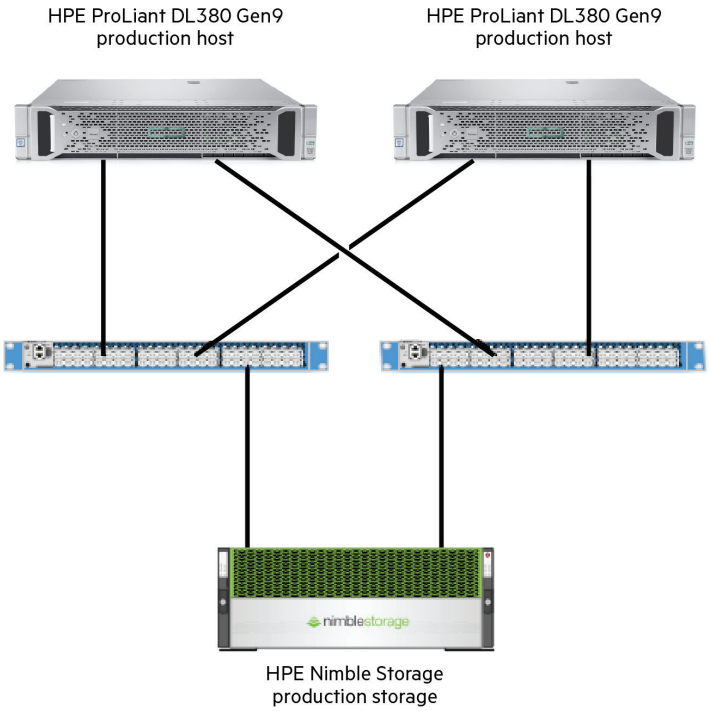

Figure 8. Production tier including HPE ProLiant DL380 Gen9 and HPE Nimble Storage

The storage area network (SAN) consists of redundant SN6000B FC switches. These market-leading 16 Gb Fibre Channel switches provide 48 ports of energy efficient performance in a 1U form factor. All databases, VMDKs, and virtual machines are backed by HPE Nimble Storage. The power and speed of HPE Nimble Storage predictive flash storage can provide extreme performance for the most latency sensitive applications. HPE Nimble Storage can scale up to boost performance, scale deep to increase capacity per node, or scale out to improve both performance and capacity. A single HPE Nimble Storage All Flash Array or Adaptive Flash Array can scale out to a four-node cluster for superior performance, flexibility, and management.

HPE Nimble Storage deduplication, inline compression, and zero-pattern elimination minimizes the data footprint within the production tier without sacrificing performance. HPE Nimble Storage snapshots can be used to protect data locally to prevent downtime or data loss. HPE Nimble Storage snapshots are based on the “redirect on write” (ROW) implementation, which means metadata is initially copied and new writes are redirected to new free space. The ROW implementation results in near instantaneous snapshot creation and zero impact on production applications. HPE Nimble Storage Snapshots can be used to quickly and efficiently protect data within the production tier. This snapshot functionality is also used by the Veeam Storage Integration feature to offload hosts from the impact of Veeam backup jobs. During a backup operation, Veeam will create an HPE Nimble Storage snapshot and use it as the source of the backup. When paired with Veeam Backup & Replication, HPE Nimble Storage snapshots can be used to maximize efficiency, data protection, and recovery times.

Veeam Backup & Replication paired with HPE ProLiant servers, switches, and HPE Nimble Storage can provide the performance, scalability, and simplicity required to make an optimal production tier and BCDR environment.

On-premises backup tier

As part of a hybrid cloud data protection model, an on-premises backup tier serves to quickly resume services in the case of an isolated disaster. This tier can be in the same building or site as the production tier. A key attribute of the on-premises backup tier is that it must maintain a high level of performance and control with minimal RTOs and RPOs. RTOs and RPOs are minimal due to similarities in infrastructure, investment, and geographical location. The on-premises backup tier can be used to backup business-critical and revenue-generating applications that require low RTOs and RPOs. Businesses can maintain a high-level of control over the performance, availability, and management of this on-premises tier.

In this reference architecture, the on-premises backup tier will be proven using the following hardware and software. Along with other platforms, those outlined below were tested and used to demonstrate the data protection capabilities of Veeam Backup & Replication. Even greater performance and functionality can be realized with the latest HPE servers, switches, and storage.

Table 2. On-premises backup tier resources

Role	Hardware	Software	Notes
Veeam Backup Server	HPE ProLiant DL380 Gen9	Veeam 9.5	
Network Switch	HPE 5900		
Veeam Backup Repository	HPE StoreOnce 5100	v3.15.3	StoreOnce Catalyst over IP

In this reference architecture, an HPE ProLiant DL380 Gen9 will be used as a Veeam Backup server within the on-premises backup tier. Individual Veeam components like backup proxies can be hosted on different infrastructure to satisfy performance and scale out requirements, but for the simplicity of this demonstration, they will all be hosted on the same server.

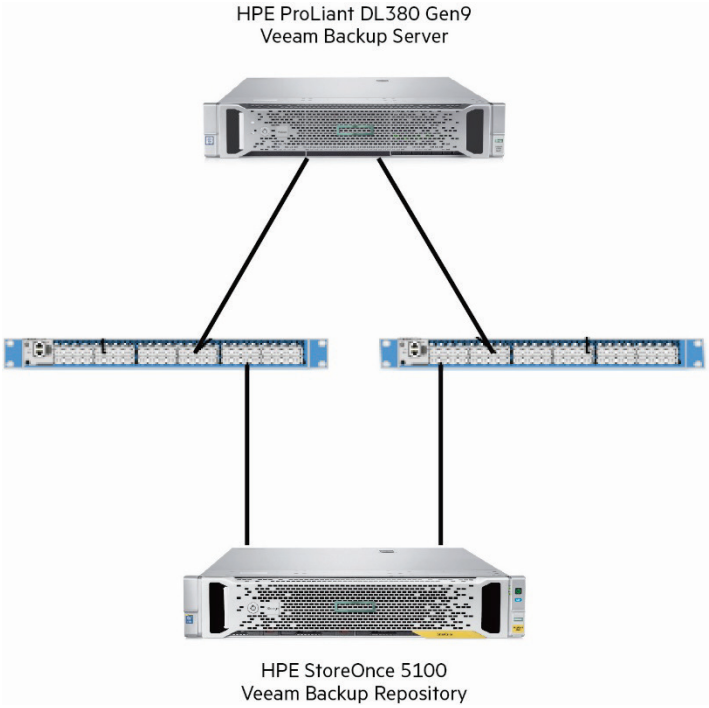

Figure 9. On-premises backup tier including HPE ProLiant DL380 Gen9 servers and StoreOnce 5100

The storage area network (SAN) consists of dual HPE 5900 high-density, ultra-low-latency, top-of-rack (ToR) switches connected to a HPE StoreOnce data protection backup appliance. The HPE StoreOnce Catalyst technology will be used in implementing the backup repository. Veeam Backup & Replication integrates with StoreOnce Catalyst client software and provides source deduplication for maximum efficiency and performance. By deduplicating data on the source-side, HPE StoreOnce Catalyst technology reduces the amount of data that is transferred over the network. Along with network traffic reduction, deduplicating data will minimize the data footprint on the backup appliance. Veeam integration with HPE StoreOnce Catalyst software makes for an extremely powerful and efficient backup solution.

Having top-of-the-line HPE servers, switches, and storage on both the production and on-premises backup tier ensures that isolated failures have minimal impact on business operations. Veeam combined with the performance, efficiency, and features of HPE ProLiant servers, switches, and StoreOnce backup storage can guarantee the smallest impact possible.

Note
The use of HPE StoreOnce in this demonstration is physically on-premises. Using the HPE StoreOnce VSA platform in Microsoft Azure is not supported as a Veeam Cloud Connect repository at this time.

Public cloud tier

A public cloud tier can take the form of a cloud service provider. Cloud service providers commonly offer Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Using Veeam with a capable IaaS offering creates a Disaster Recovery as a Service (DRaaS) environment that is unmatched.

Veeam Cloud Connect can be used to provide cloud repositories for tenants to store their backup files easily with the cost benefits of the public cloud.

In this reference architecture, the public cloud tier will be proven using Microsoft Azure resources. The Veeam Cloud Connect for the Enterprise template is provided by Veeam and available through the Azure Marketplace. A Virtual Network Gateway is created in Azure and used to establish a Point-to-Site VPN from on-prem resources to the Veeam Cloud Connect VM in Azure.

Along with other VPN methods and Azure configurations, those outlined below were tested and used to demonstrate the data protection capabilities of Veeam Backup & Replication. Even greater performance and functionality may be realized with different configurations.

Table 3. Public cloud tier resources

Name	Notes
Microsoft Azure Subscription	
Veeam Cloud Connect for the Enterprise	Available in the Azure Marketplace. Standard D2 v3 VM used (2 vCPUs, 8 GB memory).
Network Security Group	Provisioned automatically with virtual machine
Virtual Network Gateway	Configured for network connectivity

Note:
Refer to "Veeam Cloud Connect Administrator Guide" for more information.

Setup and configuration

1. Production tier deployment

The production tier consists of a two-node VMware ESX 6.5 cluster composed of DL380 Gen9 servers and a virtual machine serving as the Microsoft SQL 2016 application server. The production storage is an HPE Nimble Storage CS500 Array. In this section, we provide an overview of the setup and configuration process for the production tier.

a. VMware ESX 6.5 2-Node Cluster

I. Install 2 ProLiant DL380 Gen9 servers

Refer to the HPE ProLiant DL380 Gen9 Server User Guide

II. Install ESXi 6.5 Update 1 on the 2 ProLiant servers

Recommend using custom HPE ESXi ISO image for installing ESXi 6.5 on ProLiant Servers. These customer ESXi images will provide all required drivers for HPE specific hardware.

Refer to ESXi Image for HPE ProLiant

Refer to HPE **Nimble Storage VMware® Integration Guide** located under documentation on HPE InfoSight

III. Deploy VMware® vCenter Server™ (virtual machine or bare-metal server)

Refer to vSphere Installation and Setup

IV. VMware Cluster Creation

Refer to VMware vSphere documentation, Creating and Using vSphere HA Clusters

b. HPE Nimble Storage

I. Install HPE Nimble Storage Array

Refer to HPE **Nimble Storage Adaptive Flash Array Installation Guide** located under documentation on HPE InfoSight

II. Create an Initiator group for VMware cluster

Refer to HPE **Nimble Storage VMware Integration Guide** located under documentation on HPE InfoSight

Create 3 volumes on HPE Nimble Storage array, selecting the "vSphere Datastore for SQL Server" Performance Policy.

Present the volumes to the ESX hosts and create 3 VMware datastores

VM Boot datastore

SQL Data datastore

SQL Log datastore

Refer to HPE **Nimble Storage VMware Integration Guide** located under documentation on HPE InfoSight

Refer to HPE **Nimble Storage Veeam Integration (KB-000367)** located under documentation on HPE InfoSight

c. Deploy the Microsoft SQL Server Enterprise 2016 (Virtual Machine)

I. Create a virtual machine specifying the "VM Boot datastore" as the installation location

II. Install Windows Server 2016 Data center Edition

III. Add two virtual disk to the virtual machine (one from each of the "SQL Data" and "SQL Log" datastores)

IV. Install Microsoft SQL Server Enterprise 2016 SP1

Refer to the Install SQL Server instruction from Microsoft for details

V. Create SQL Database

2. On-premises backup tier deployment

a. Install StoreOnce 5100

Refer to StoreOnce 3100, 3500, 5100, and 5500 System Installation and Configuration Guide

b. Create a StoreOnce Catalyst Store

c. Deploy Veeam Backup & Replication

I. Install 1 ProLiant DL380 Gen9 server

Refer to the HPE ProLiant DL380 Gen9 Server User Guide

II. Install Windows Server 2016 Data center

III. Install Veeam Backup & Replication

IV. Create Veeam Backup Repository using StoreOnce with Catalyst deduplication
Veeam with HPE StoreOnce Catalyst Configuration Guide

V. Add HPE Nimble Storage array to Veeam Storage Infrastructure

VI. Add VMware vCenter Server to Veeam Backup Infrastructure

3. Azure public cloud tier deployment

a. Azure subscription and resources

An active Azure subscription is required to deploy the Azure Public Cloud Tier. Microsoft offers trial subscriptions as well as pricing and total cost of ownership calculators to help determine the benefit, size, and extent of a business investment in Azure. These tools help businesses quickly assess and estimate the total cost of operating in the cloud.

b. Connectivity to Microsoft Azure

Connectivity between the on-premises Veeam Backup Server and the public cloud provider needs to be established. There are several ways to link an on-premises site to Microsoft Azure including the popular options below.

I. Point-to-Site

II. Site-to-Site

III. ExpressRoute

IV. Third-party VPN solutions

A Point-to-Site VPN is simple and easy to use when connecting a small number of systems to Microsoft Azure. This is typically used for small-scale test and dev scenarios. Point-to-Site VPN uses the Secure Sockets Tunneling Protocol (SSTP), which requires minimal firewall modification on-prem. A Site-to-Site VPN can be ideal for larger environments where connectivity between a greater number of systems is required. Site-to-Site VPNs use an IPSec/IKE VPN tunnel and require both an external facing IP and VPN device on-premises. The use of this method typically scales from test and dev scenarios to small-scale production environments. ExpressRoute can extend your on-prem networks into Microsoft Azure over a private connection. This connection is offered by certain service providers and can provide the best level of service for large production environments. Third-party VPN solutions like SoftEther are also available to gain access to resources in Microsoft Azure. These VPN solutions typically require both server and client side software as well as modifications to a firewall to establish the VPN.

For this reference architecture, a Point-to-Site VPN was used to establish connectivity between the on-prem Veeam Backup Infrastructure and the Microsoft Azure resources required for Veeam Cloud Connect. The following diagram depicts basic connectivity and some of the components required.

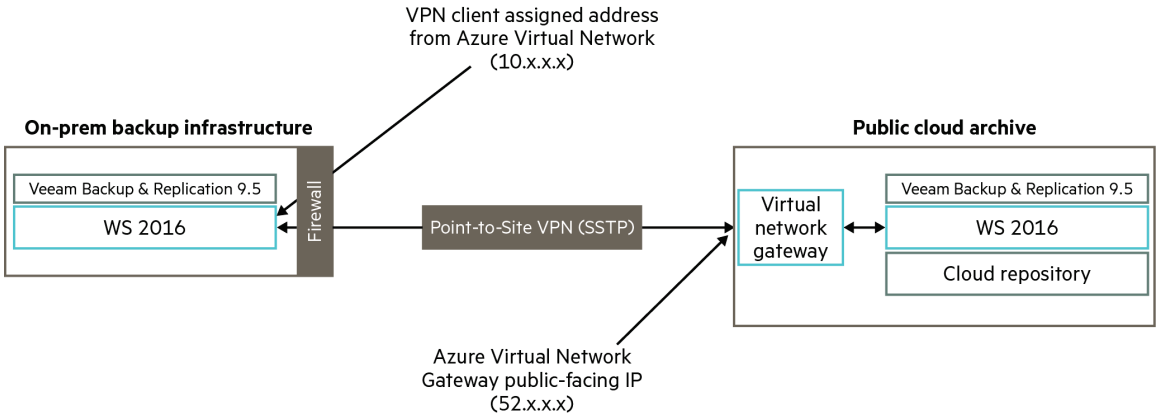

Figure 10. On-premises to Microsoft Azure VPN setup

Note

For more information on VPN methods, refer to the Microsoft Azure planning guides for designing cross-premises connections.

c. Deploy and Configure Veeam Cloud Connect for the Enterprise

I. Provision the Veeam Cloud Connect for the Enterprise image from the Azure Marketplace.

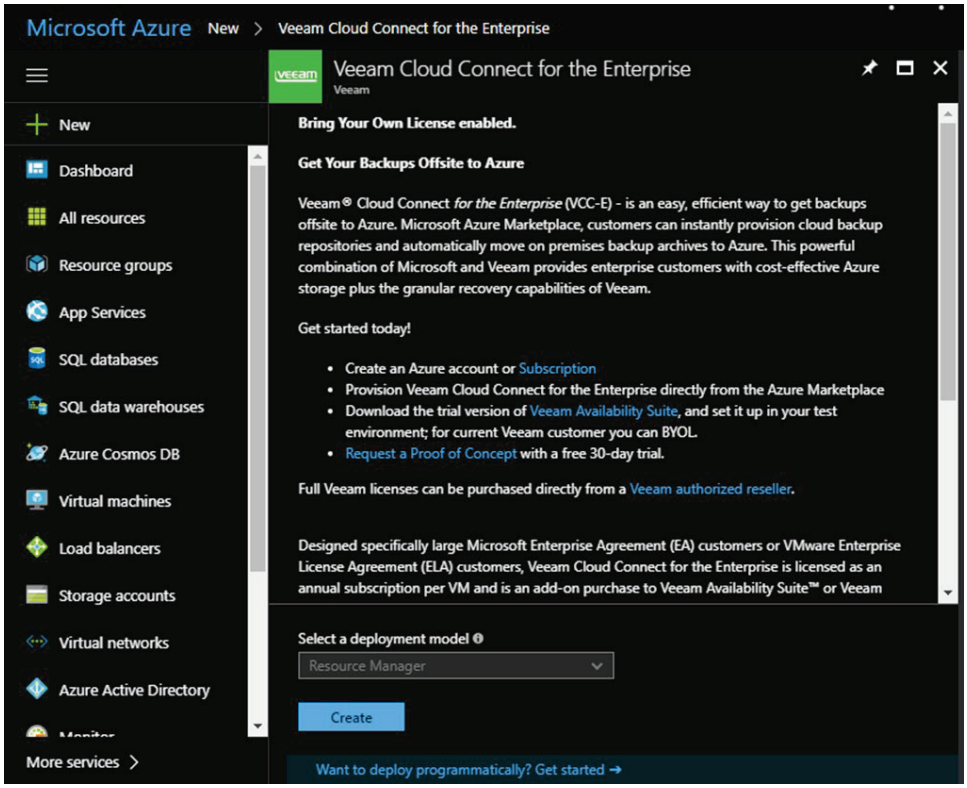

Figure 11. Veeam in Azure Marketplace

II. After specifying the name, user name, password, and resource group, choose the size of the virtual machine. The size of the virtual machine will depend on individual performance and capacity requirements.

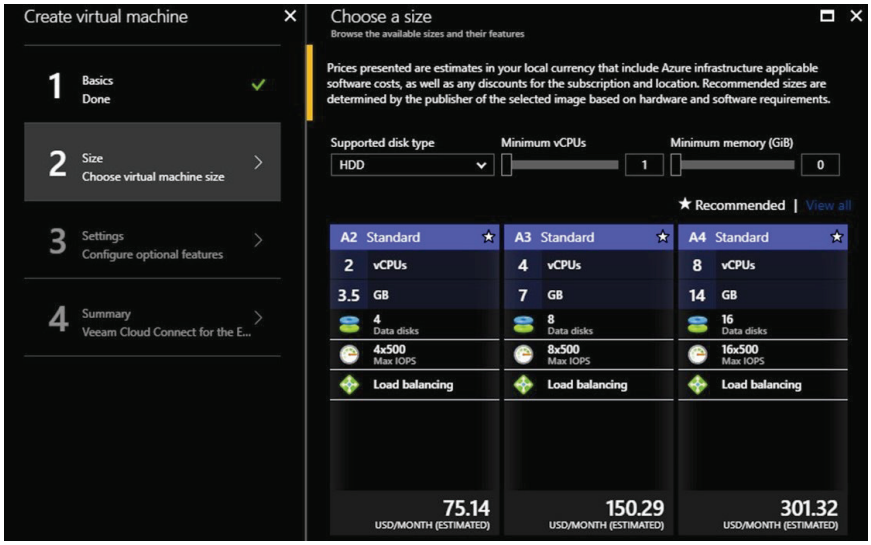

Figure 12. Virtual machine sizing

III. Specify network details, including the virtual network, subnet, network security group, and public IP.

IV. Validate summary and press create.

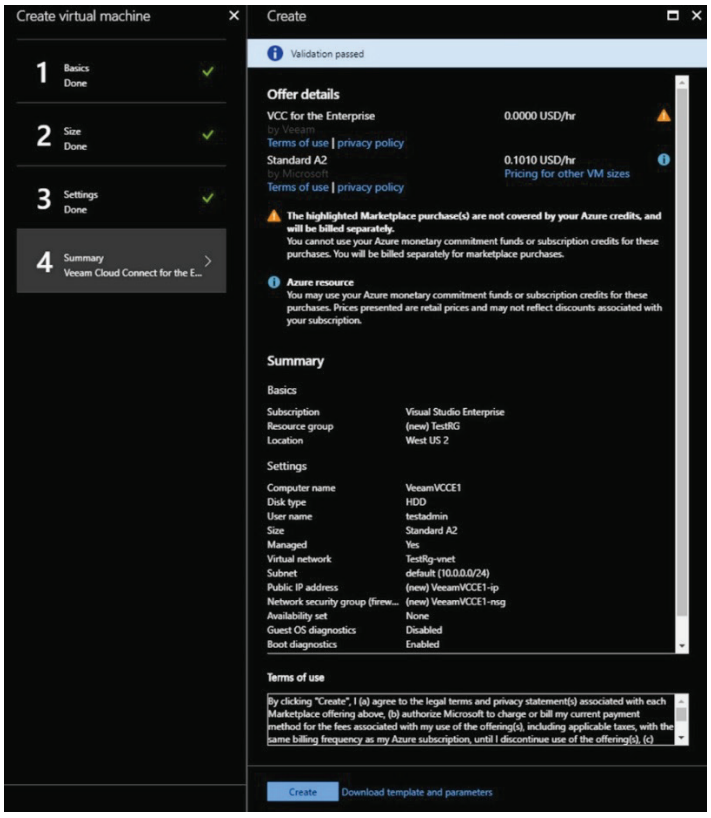

Figure 13. Summary and confirmation

V. Wait for VM to be provisioned within Azure.

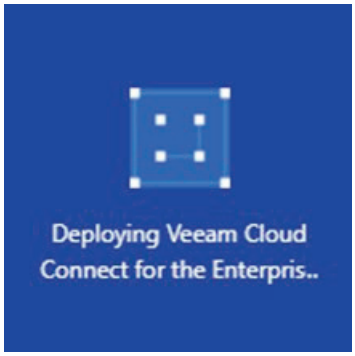

Figure 14. VM creation

VI. Initialize and login to Veeam Cloud Connect.

1. Open the Veeam Backup & Replication Console.
2. Provide a license file for Veeam Cloud Connect.

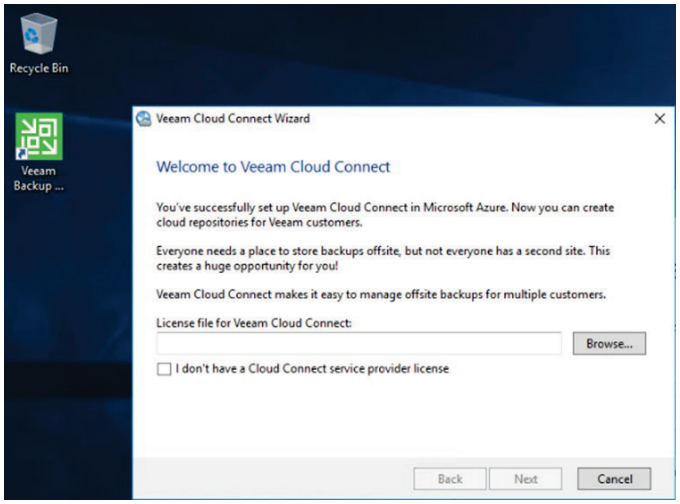

Figure 15. Veeam Cloud Connect Wizard

3. Proceed through the rest of the Veeam Cloud Connect Wizard.

4. Login to the Veeam Backup & Replication Console using Windows session authentication.

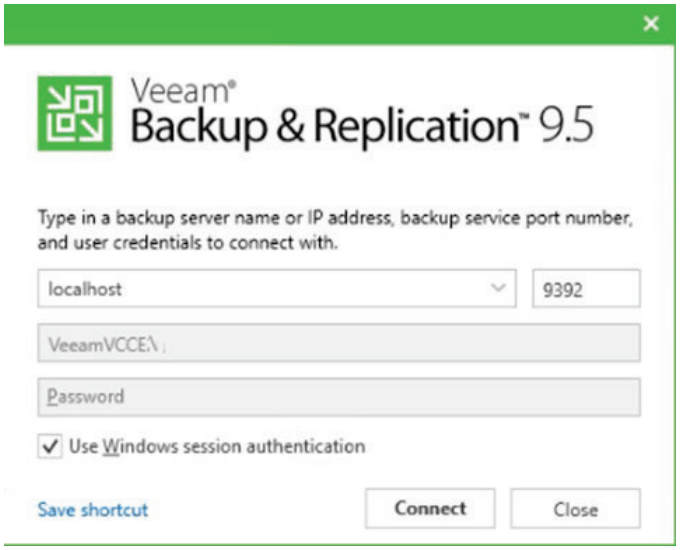

Figure 16. Veeam Backup & Replication login

VII. Create Backup Repository

1. Select the "Backup Infrastructure" View.
2. From the inventory pane, select "Backup Repositories".
3. Under the Repository tab, choose "Add Repository".

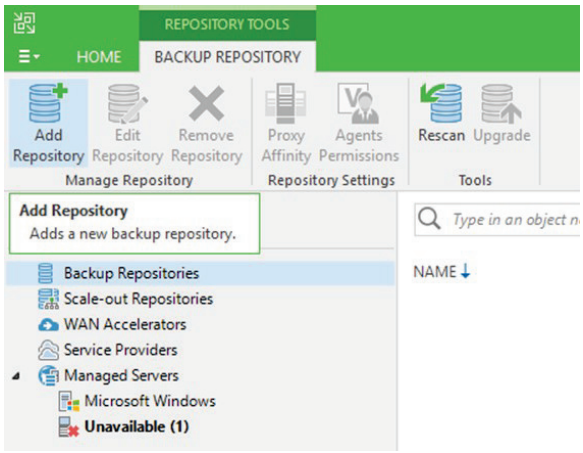

Figure 17. Backup Repository creation

4. Proceed through the “New Backup Repository” wizard. Any type of backup repository can be chosen, but in this case, local storage on the Veeam Cloud Connect server is utilized.

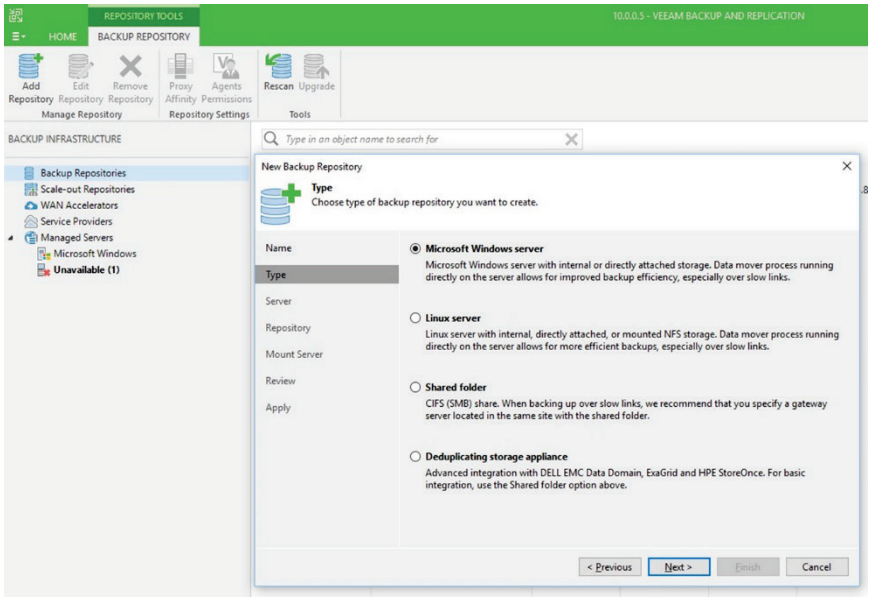

Figure 18. New Backup Repository wizard

Note
Using the HPE StoreOnce VSA platform in Microsoft Azure is not supported as a Veeam Cloud Connect repository at this time.

5. A new backup repository is created.

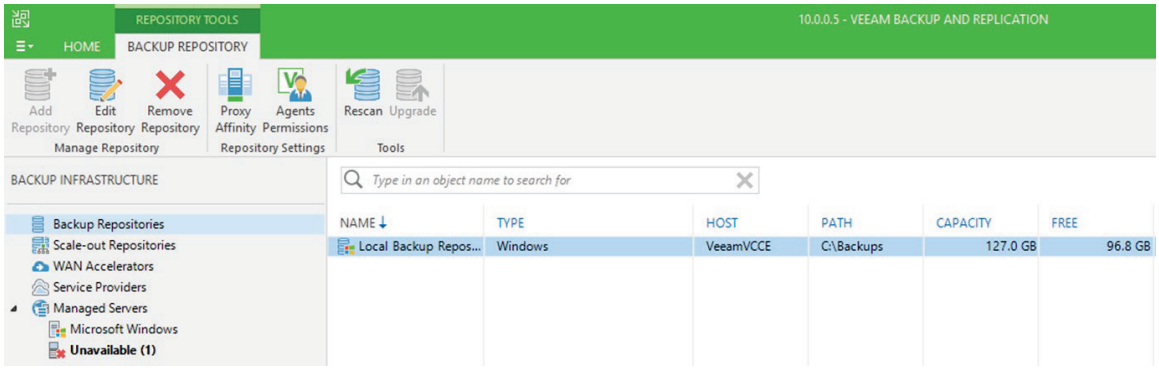

Figure 19. New Backup Repository created

VIII. Create Tenant

1. Select the “Cloud Connect” View.

2. Under the Tenants tab, choose “Add Tenant”.

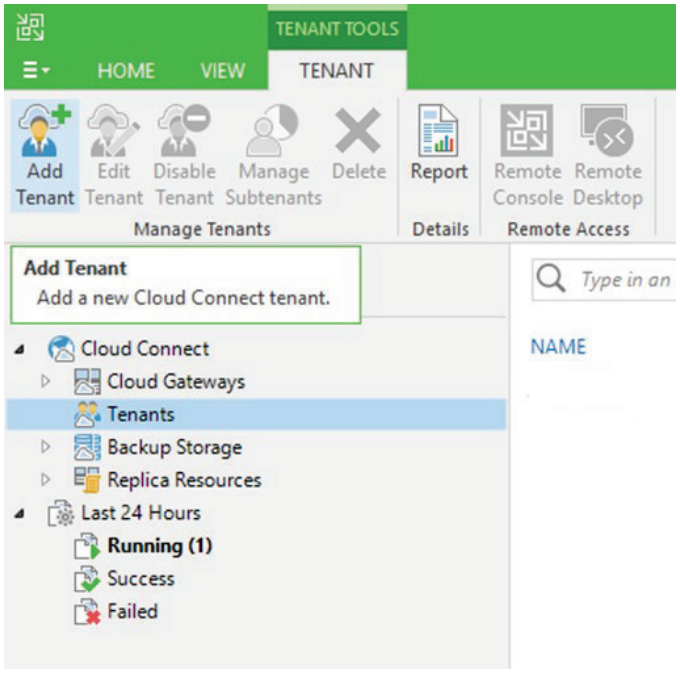

Figure 20. Tenant creation

3. Proceed through the “New Tenant” wizard checking “Backup storage” under “Assigned resources”.

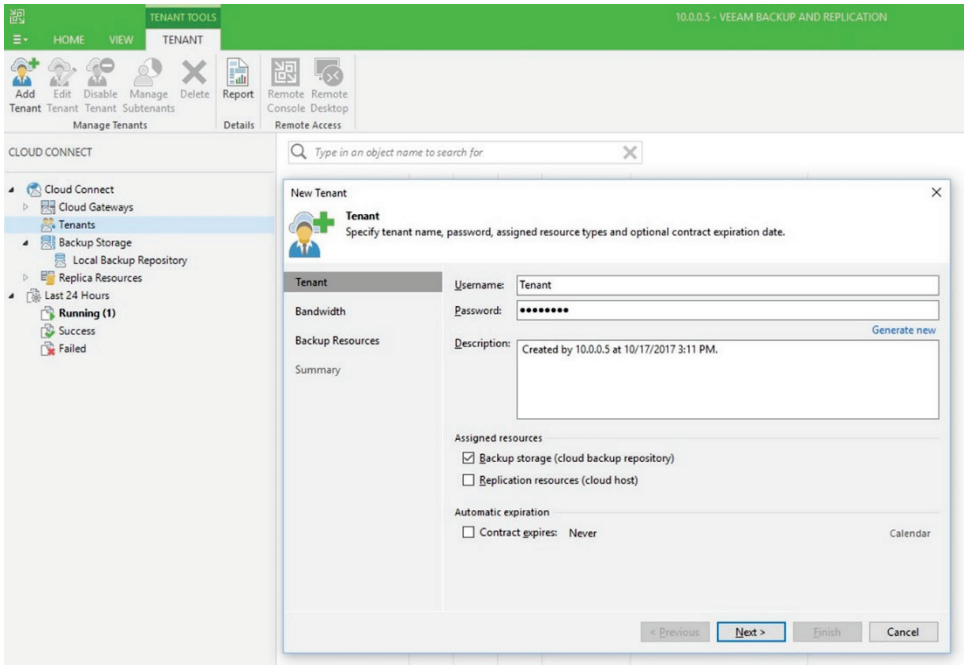

Figure 21. New Tenant wizard

4. Choose the newly created backup repository for the tenant and finish the wizard.

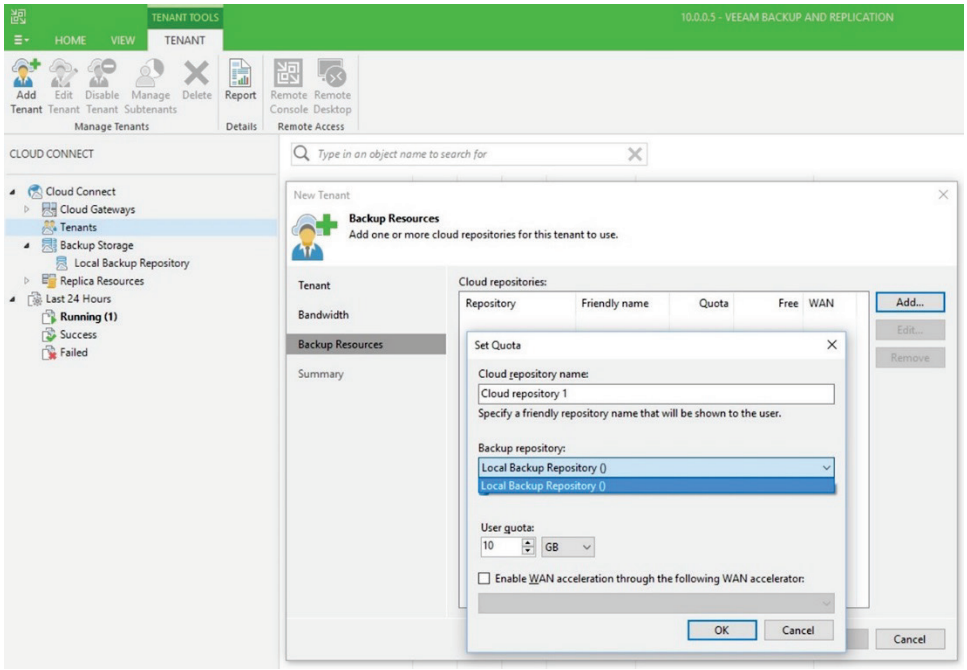

Figure 22. Backup Repository assignment

5. A new tenant is created.

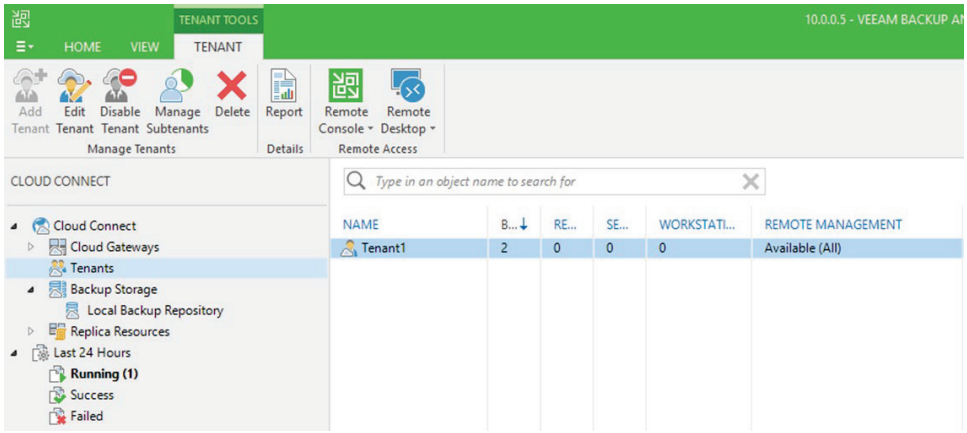

Figure 23. New Tenant created

d. Add Service Provider to On-premises Backup Server

I. Under the Cloud Connect tab, choose “Add Service Provider”.

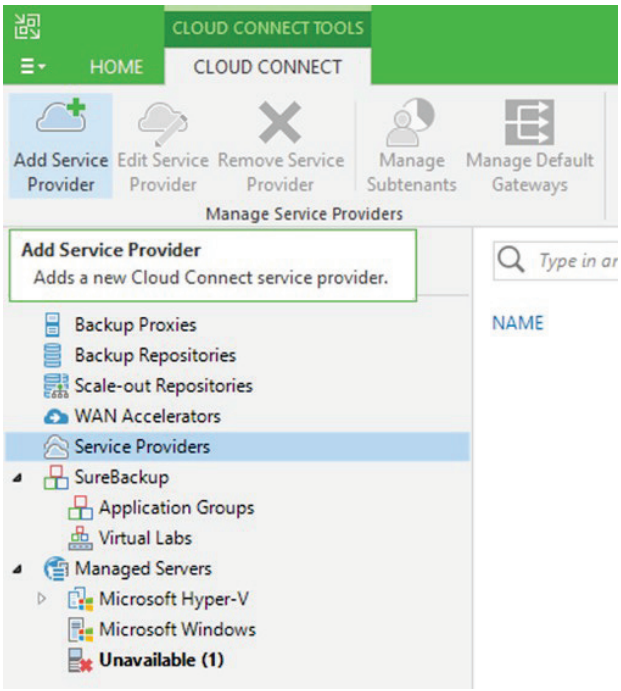

Figure 24. Cloud repository creation

II. Proceed through the Service Provider wizard. Add the DNS name or IP address of the Veeam Cloud Connect instance in Microsoft Azure. Ensure that port 6180 is allowed in the Azure network security group.

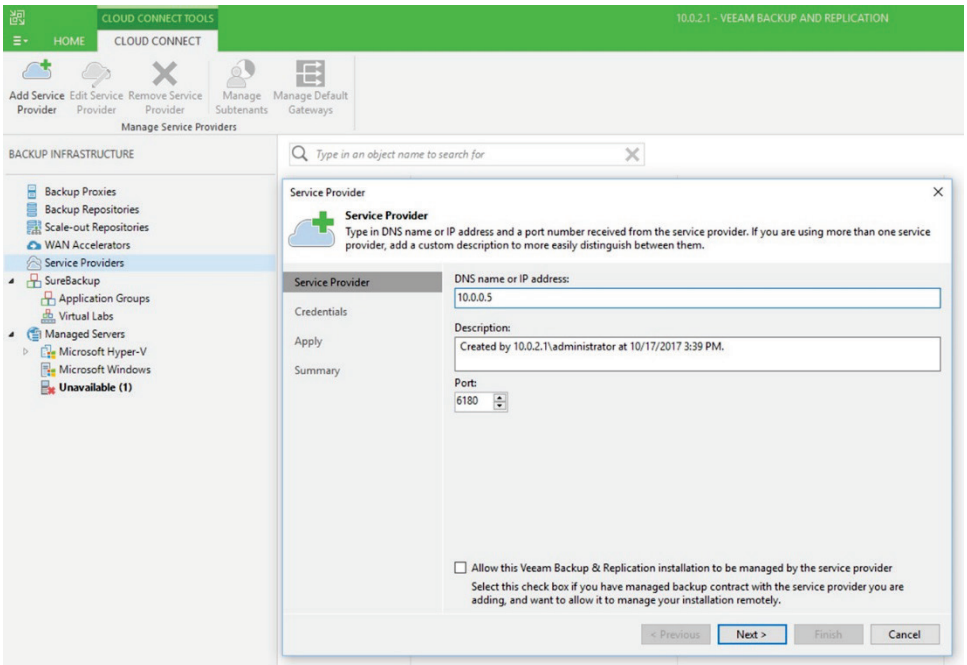

Figure 25. Service Provider wizard

III. The new service provider is created. The backup repository assigned to the tenant can now be used as a cloud repository.

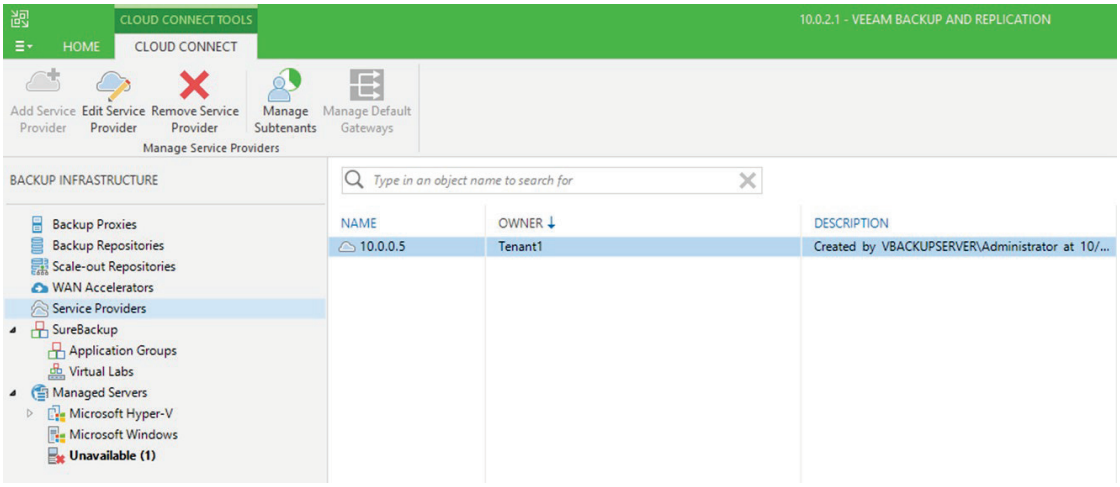

Figure 26. New Service Provider created

Use cases

Use case 1: Back up database from storage-based snapshots

Overview

When your business is down and a full system recovery is needed, time is of the essence. In order to quickly recover and resume business operations, a local backup and recovery solution is critical. In this section, we describe how to configure Veeam to take application consistent snapshots on HPE Nimble Storage arrays. With Veeam's integrated Microsoft SQL Database IO quiesce/flush capabilities, the resulting storage-based snapshots can be used to quickly recover the database without the need to perform any log rollback operations.

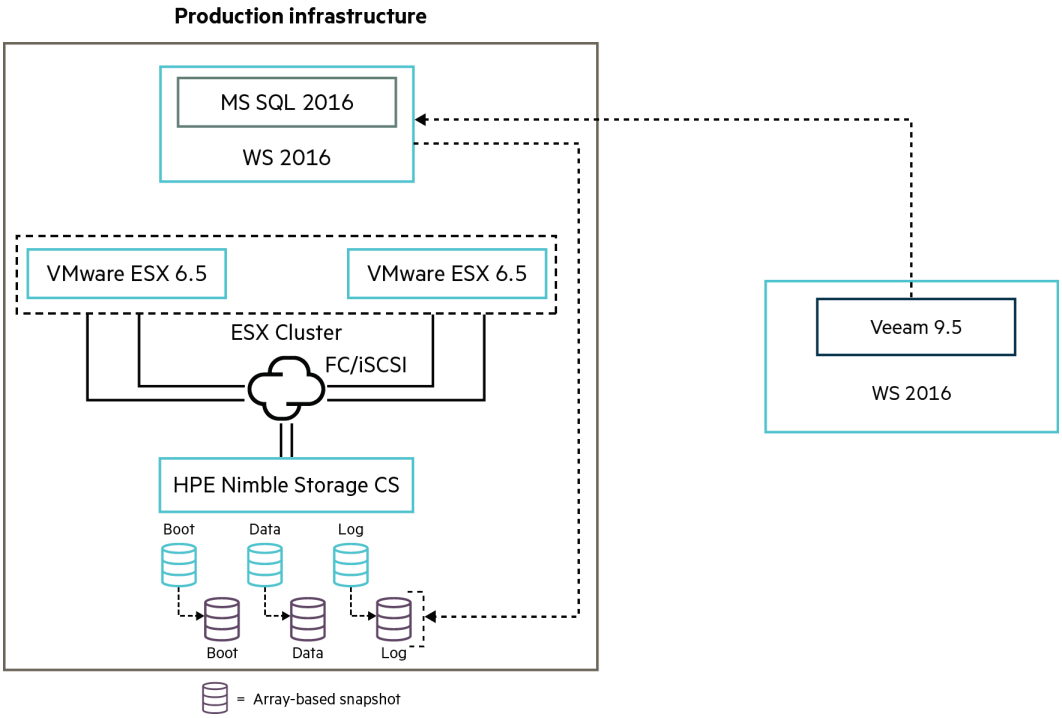

Figure 27. Array-based snapshots orchestrated by Veeam

Steps

1. To launch the “New Backup Job” wizard, from the “Backup & Replication” view, select “Backup Job” from the “Home” tab.

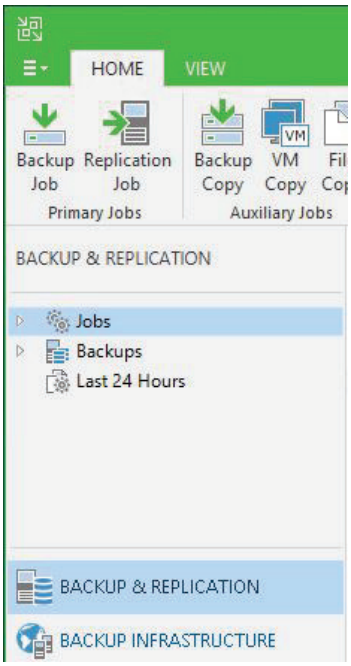

Figure 28. Backup job creation

2. Provide a name for the backup job and click on “Next”.

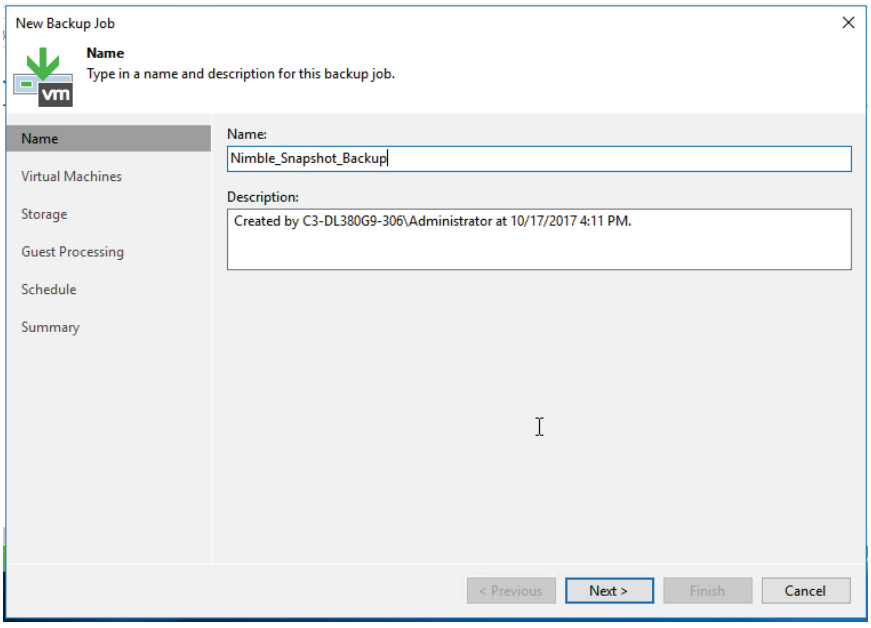

Figure 29. Backup job name

3. From the “Virtual Machines” selection screen, click on “Add …” and navigate through the ESX hierarchy to find the virtual machine(s) to be backed up.

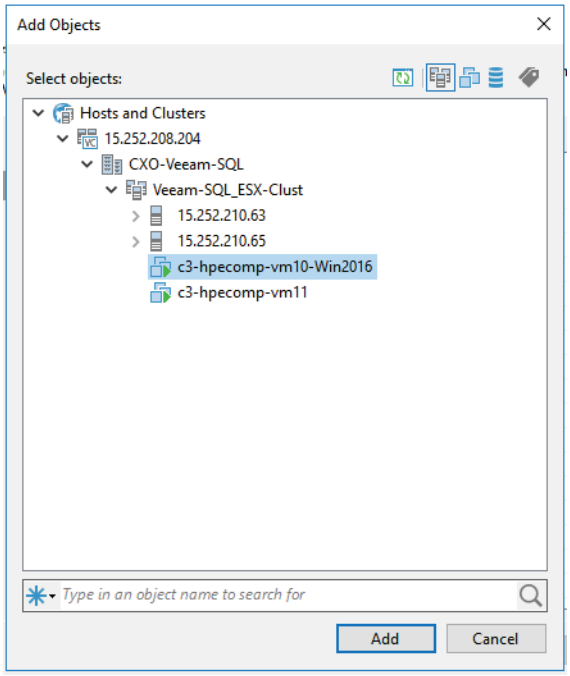

Figure 30. Virtual machine selection

In this case, we have only selected c3-hpecomp-vm10-Win2016 to be backed up. However, a job can be configured to back up more than one VM at a time.

4. Once the VMs have been selected, click on “Next” to proceed to the next screen.

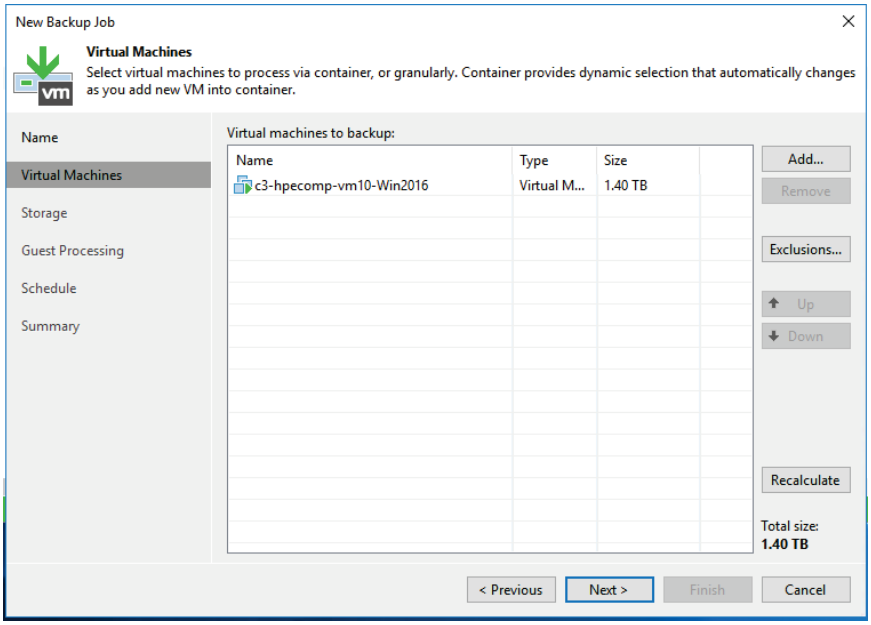

Figure 31. Virtual machine added to backup job

5. From the “Storage” screen, select the HPE Nimble SnapShot (Primary storage snapshot only) backup repository.

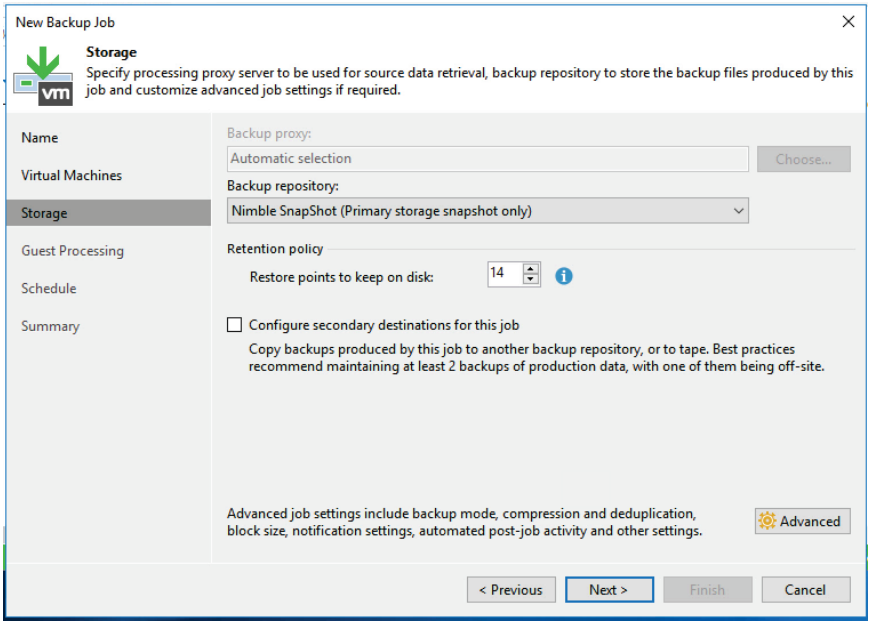

Figure 32. Selection of Backup Repository

6. (Optional) Click on the “Advanced” button to configure advanced options such as: Notifications, vSphere, Integration, and Script options.
7. In the “Guest Processing” window, check the “Enable application-aware processing” box and provide the guest OS credentials for the VMs being backed up.
   a. If the guest OS credentials have not already been set up, click the “Add …” button to create the corresponding credentials.

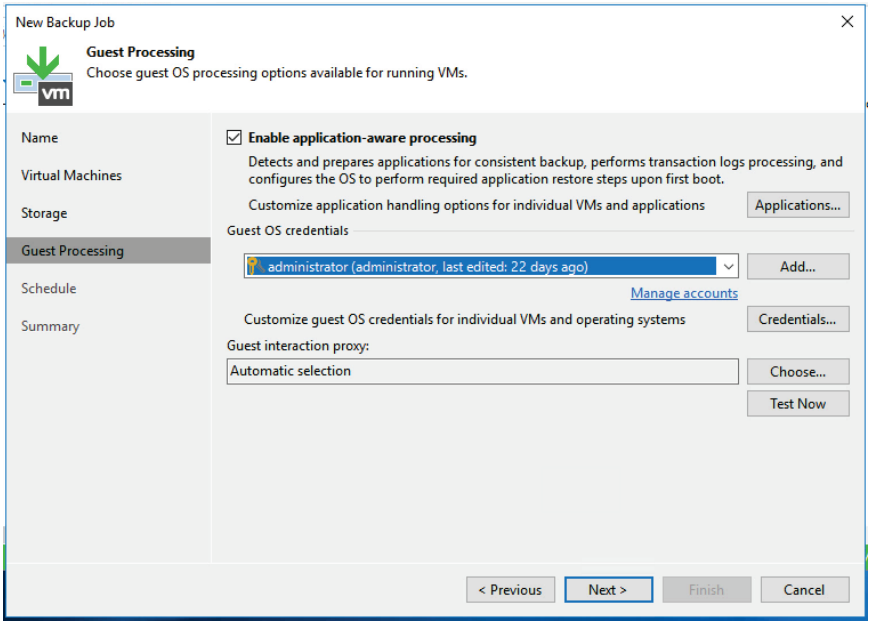

Figure 33. Application-aware processing

8. Click on “Applications” and ensure that the SQL processing options have been enabled.

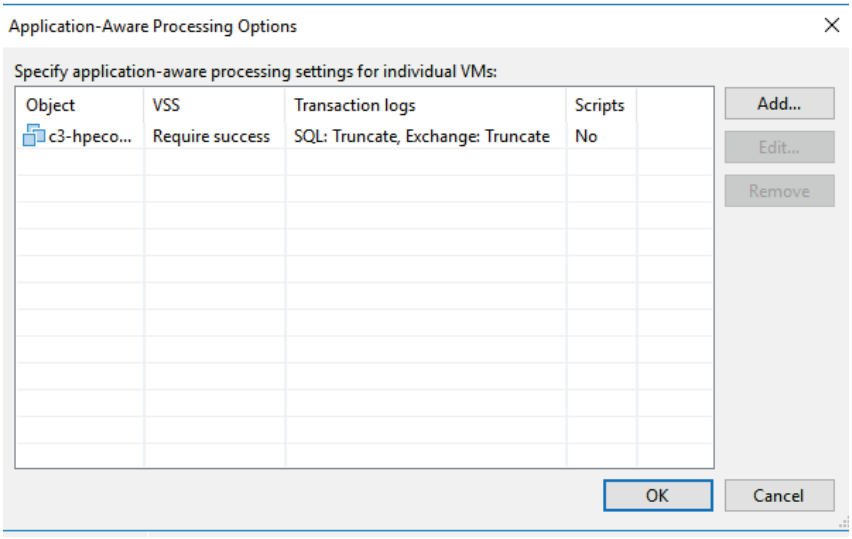

Figure 34. Confirmation of SQL processing options

9. Once ready to proceed, click “OK” from the “Application-Aware Processing Options” screen, and then click “Next” to proceed.
10. In the next screen, the scheduling options are presented. The job can be scheduled to run at predetermined intervals or left to be executed manually. Once the scheduling options have been selected, click “Apply” to proceed.

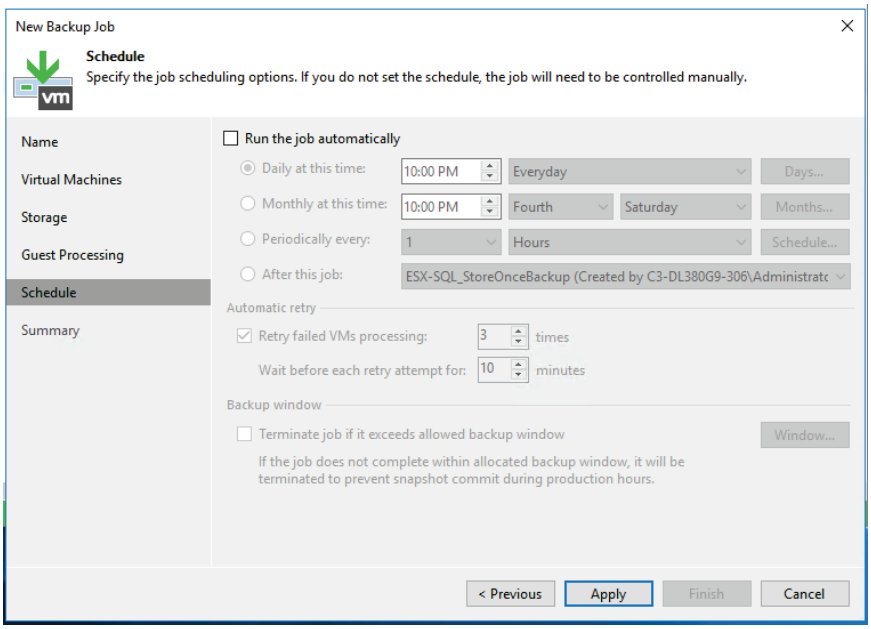

Figure 35. Backup job schedule

11. The next screen provides a summary of the selected options and backup configuration. Review the information provided for accuracy and click “Finish” to complete the job creation process.

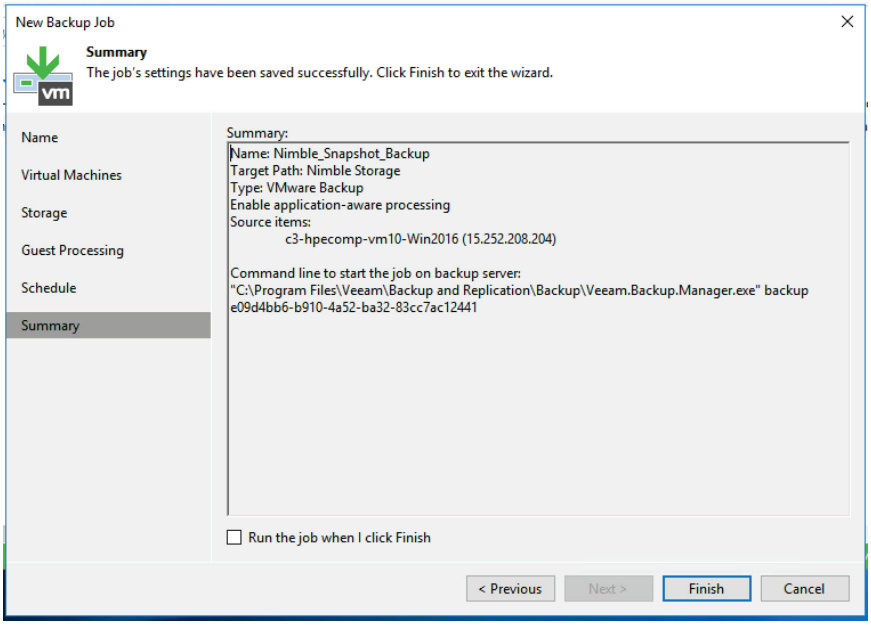

Figure 36. Backup job summary and confirmation

12. The job will now show up in the “Backup” section in Veeam.

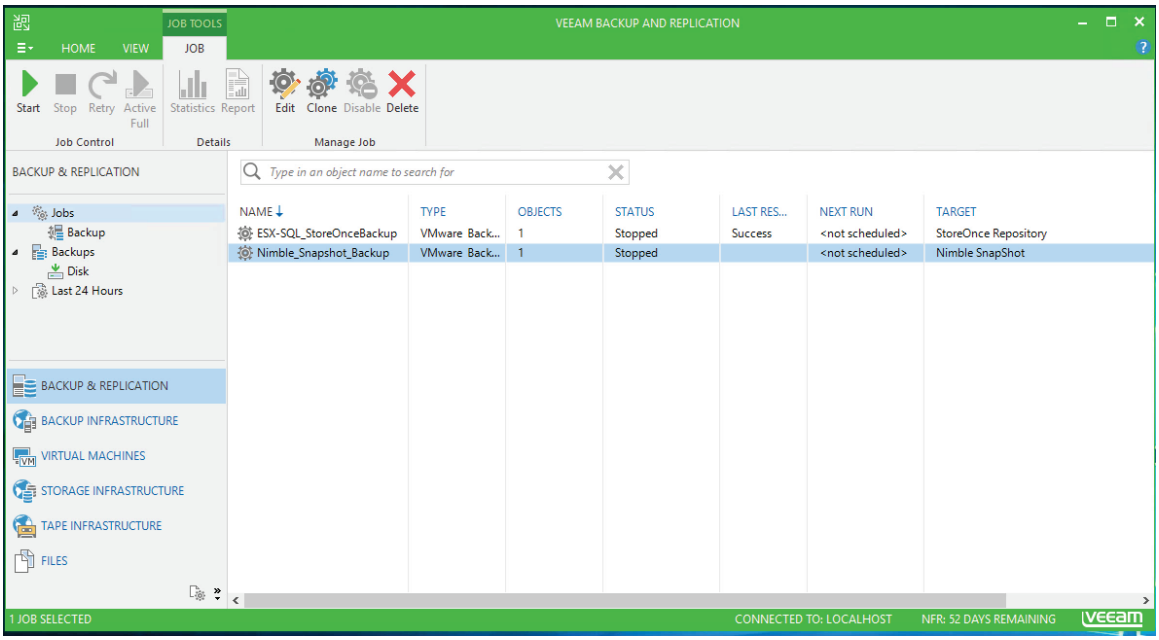

Figure 37. New backup job created

13. If the job was scheduled to be run automatically, it will do so at the specified time. If a schedule was not created, select the job and click on the "Start" button from the "JOB" tab.

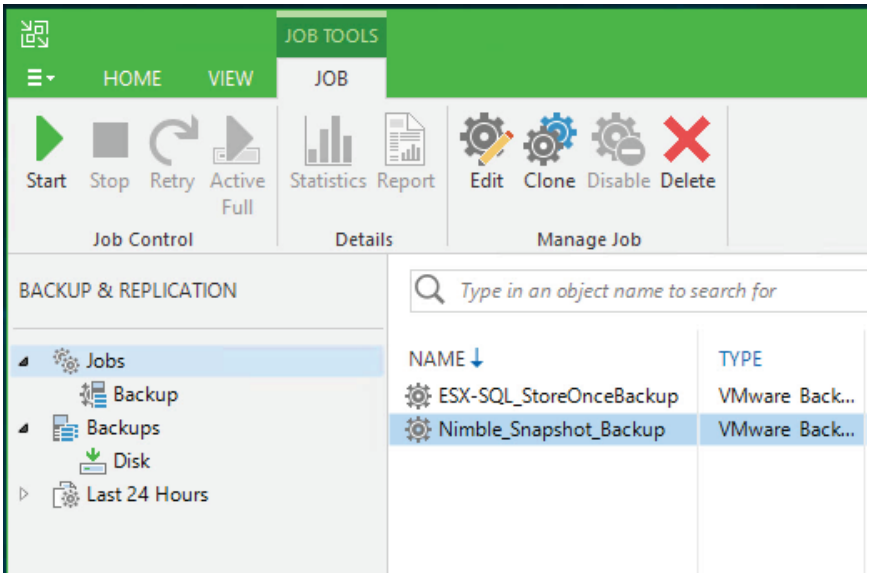

Figure 38. Start backup job

14. The job will start running immediately and a summary of progress is displayed at the bottom of the screen when the running job is selected.
15. If viewing the details of the running job is desired, select the job and click on "Statistics" from the "JOB" tab.

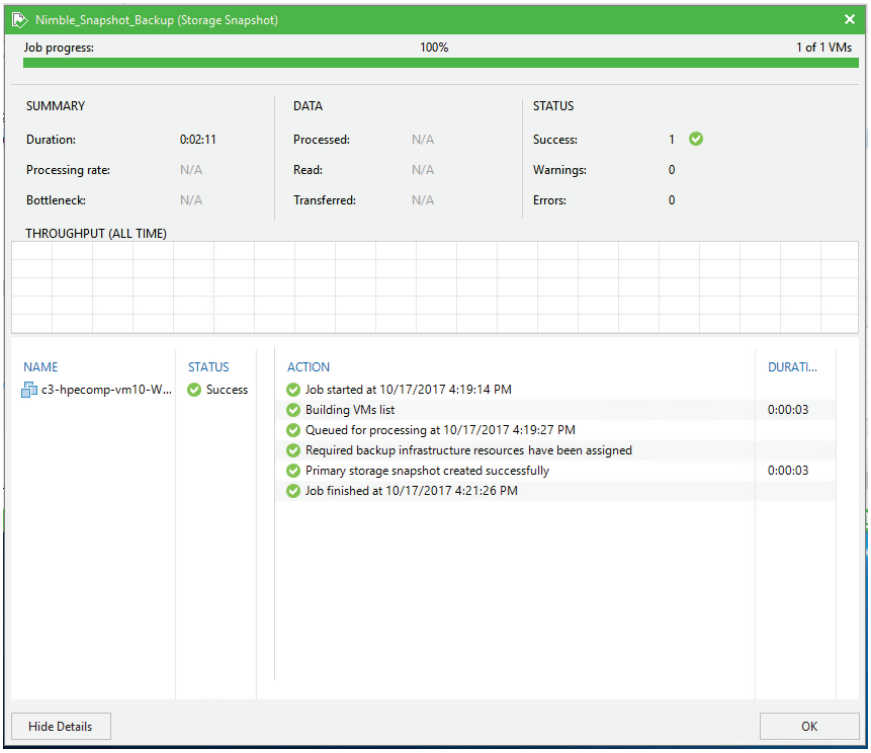

Figure 39. Backup job details

16. Once the job completes, the details of the job will indicate so and the snapshots for all volumes associated with the Virtual Machine being backed up will have been created. Snapshot creation can be confirmed directly from the Nimble Array's user interface or within Veeam from the "Storage Infrastructure" view.

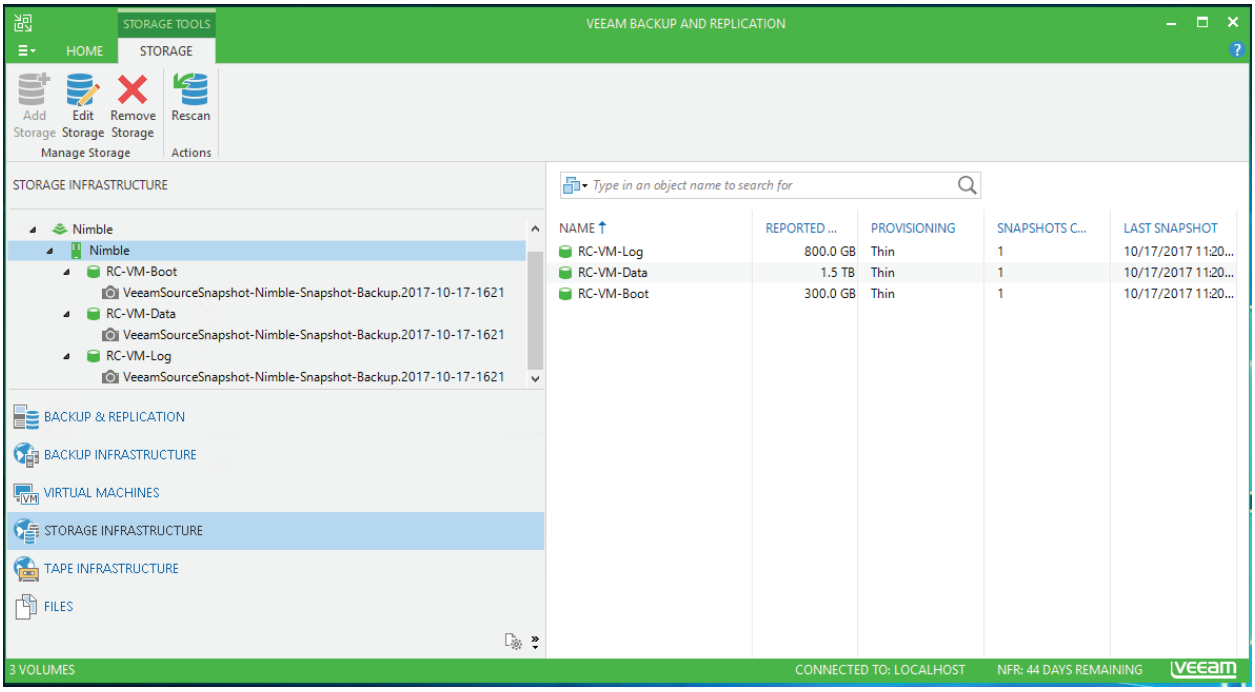

Figure 40. Veeam Storage Infrastructure view

Use case 2: Recover database from storage-based snapshots

Overview

Even though restoration from array-based snapshots is less automated that restoring from a StoreOnce backup, Veeam's storage and application awareness allows for efficient application recovery from synchronized application consistent snapshots. When restoring from storage-based snapshots, the array volume as a whole will be restored to its state at the time the snapshots were taken. If multiple VMs are associated with these volumes, they will all be affected by the restore. To avoid complications with simultaneously restoring multiple VMs, it is recommended to have unique volumes associated with each VM unless application requirements indicate otherwise. In the procedure below, we use an array-based approach to restore our MS SQL 2016 environment. If your database deployment is contained within a single VMware virtual disk, Veeam can also be used to restore, but that procedure is outside the scope of this white paper and not depicted here.

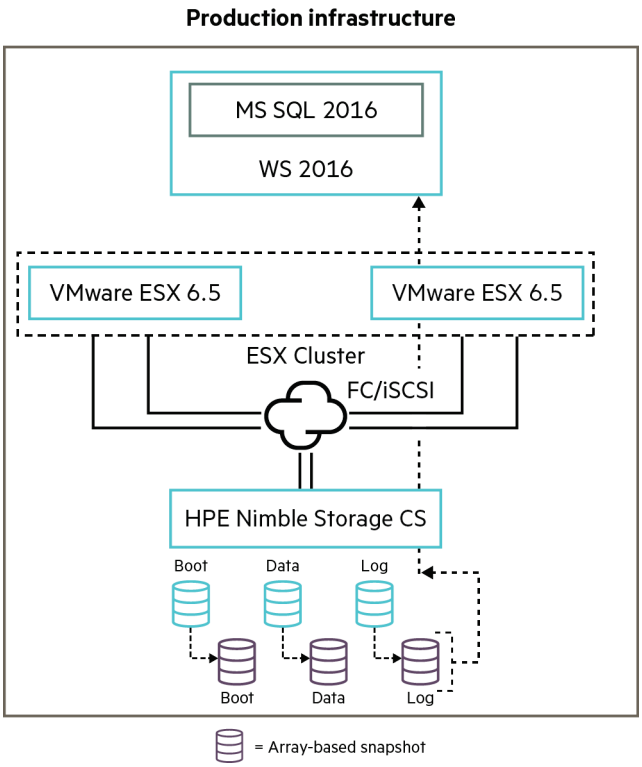

Figure 41. VM and Application restore via **HPE Nimble Storage** snapshots

Steps

1. Ensure that all VMs associated with the volumes being restored are shutdown.
2. Identify the snapshots on the array that will be used for the recovery.
3. From the HPE Nimble Storage Management Interface, click on Manage -> Volumes. Then select one of the volumes associated with the VM.

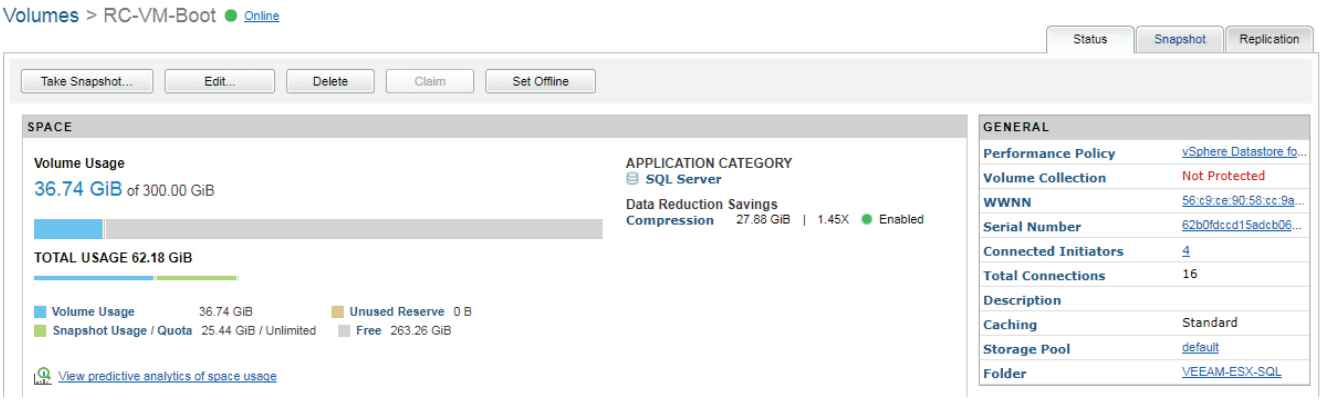

Figure 42. Volume selection on HPE Nimble Storage

4. From the “Snapshot” tab, select the snapshot to revert to and click on “Restore”.

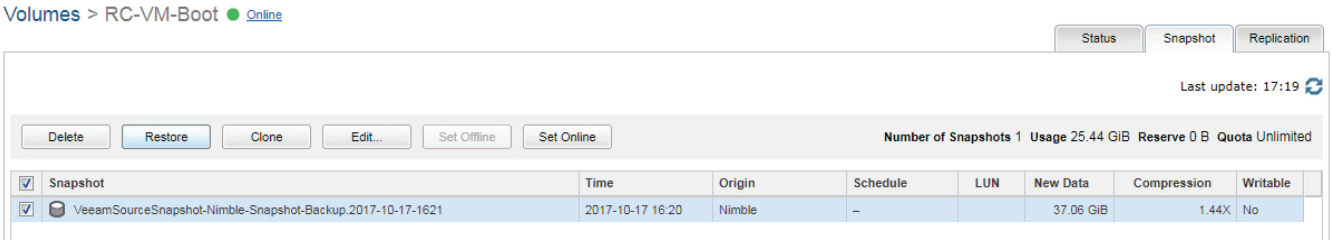

Figure 43. Restore from snapshot

5. From the Warning window, check the “Set volume offline” checkbox and click OK.

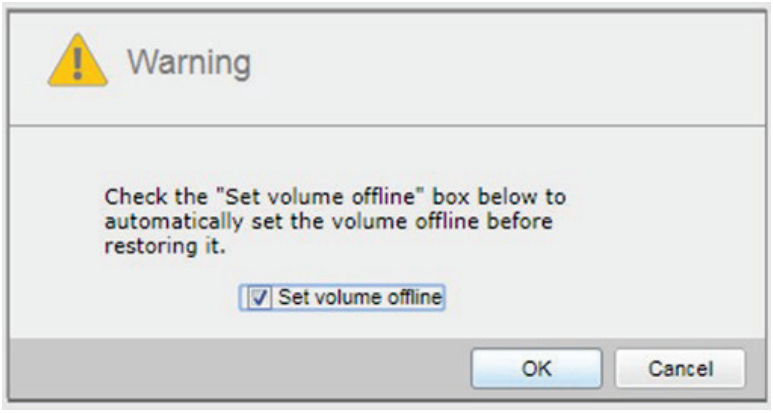

Figure 44. Set volume offline

6. Once the Restore operation completes, click OK in the confirmation box.

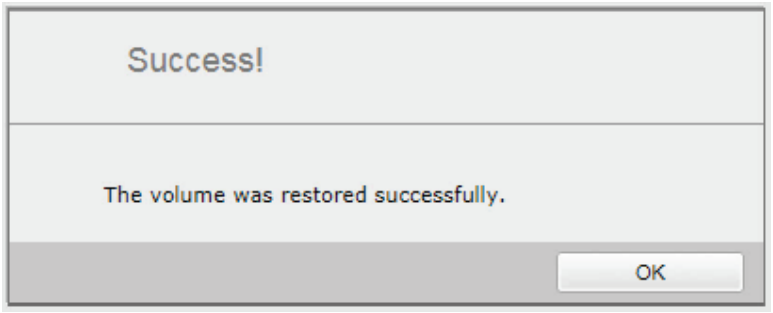

Figure 45. Set volume offline confirmation

7. Repeat steps 3 through 6 for all other volumes associated with the VM or application.
   a. Note that all volumes should be restored to correlating snapshots (that is, snapshots taken at the same time).

8. Bring the restored volumes back online by going to "Manage -> Volumes," selecting the restored volume, and clicking the "Set Online" button.

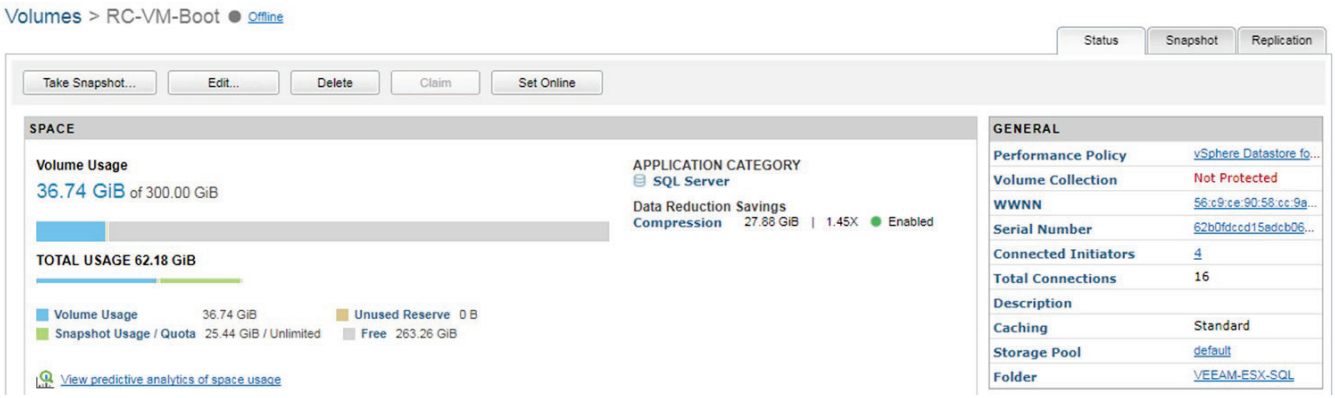

Figure 46. Set volume online

9. Make sure all volumes are back online and healthy. Then power-on the VMs that were previously shutdown.

Use case 3: Back up database to on-premises repository

Overview

With the HPE Nimble Storage, StoreOnce, and Veeam combination, backups can be scheduled to whatever frequency your business continuity plans require. Just as with storage-based snapshots, with Veeam application quiescing capabilities, the resulting StoreOnce backups are application consistent. And should the need arise, full system recovery can be accomplished with a few simple steps. In this scenario, we describe how to perform a backup with Veeam Backup & Replication to protect against a critical system failure. The procedure below assumes that the environment has been properly set up and configured per instructions mentioned earlier.

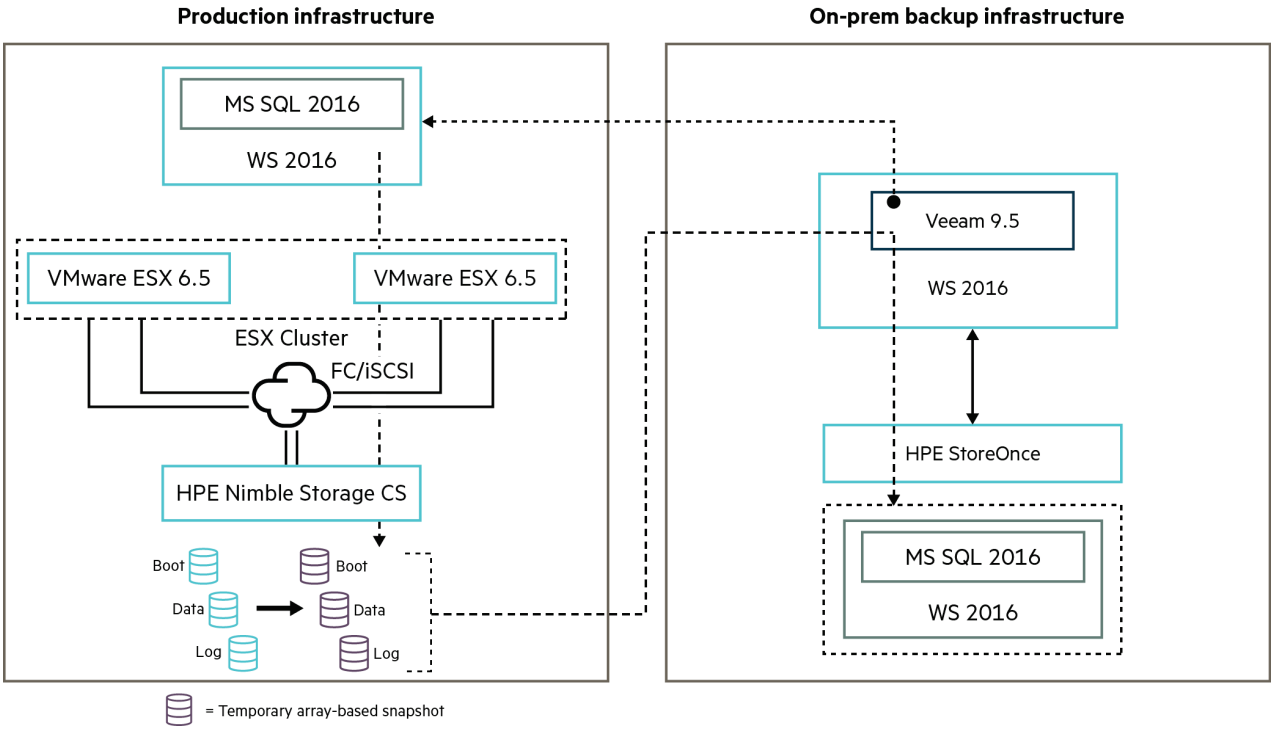

Figure 47. Veeam backup offloading using array-based snapshots

Steps

1. To launch the “New Backup Job” wizard, from the “Backup & Replication” view, select “Backup Job” from the “Home” tab.

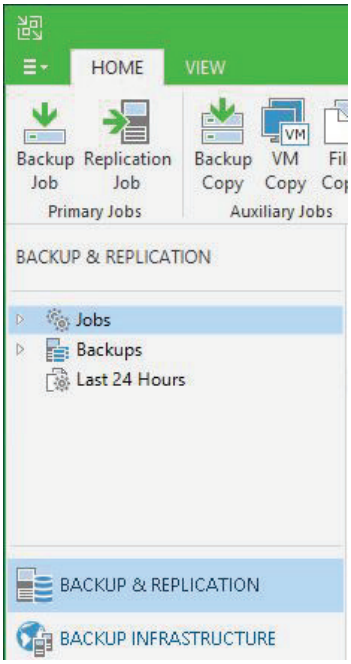

Figure 48. Backup job creation

2. Provide a name for the backup job and click on “Next”.

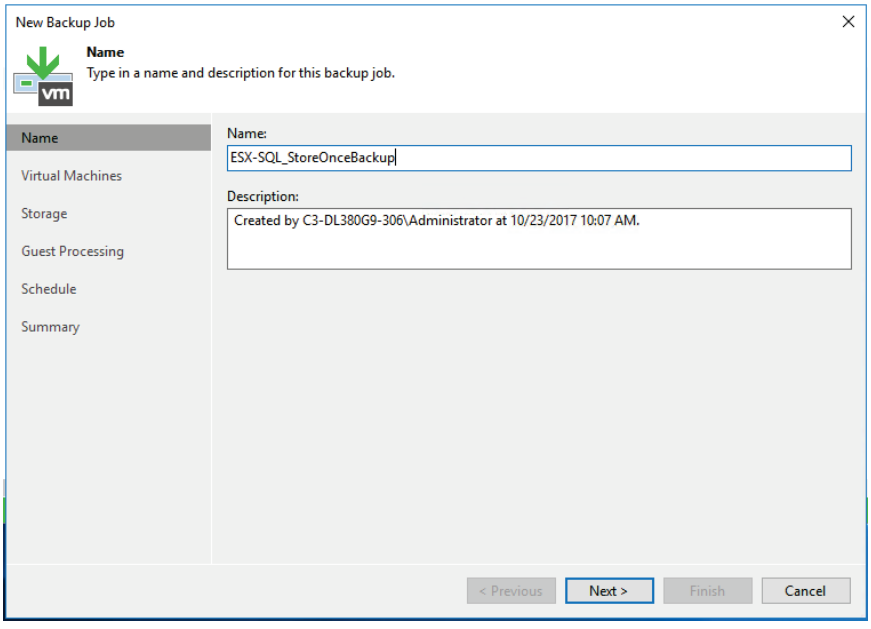

Figure 49. New Backup Job Wizard

3. From the “Virtual Machines” selection screen, click on “Add …” and navigate through the ESX hierarchy to find the virtual machine(s) to be backed up.

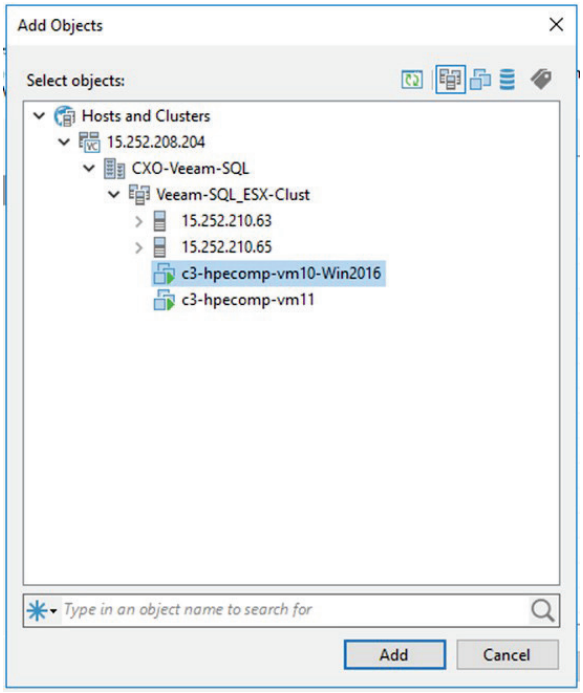

Figure 50. Add virtual machines

In this case, we have only selected c3-hpecomp-vm10-Win2016 to be backed up. However, a job can be configured to back up more than one VM at a time.

4. Once the VMs have been selected, click on “Next” to proceed to the next screen.

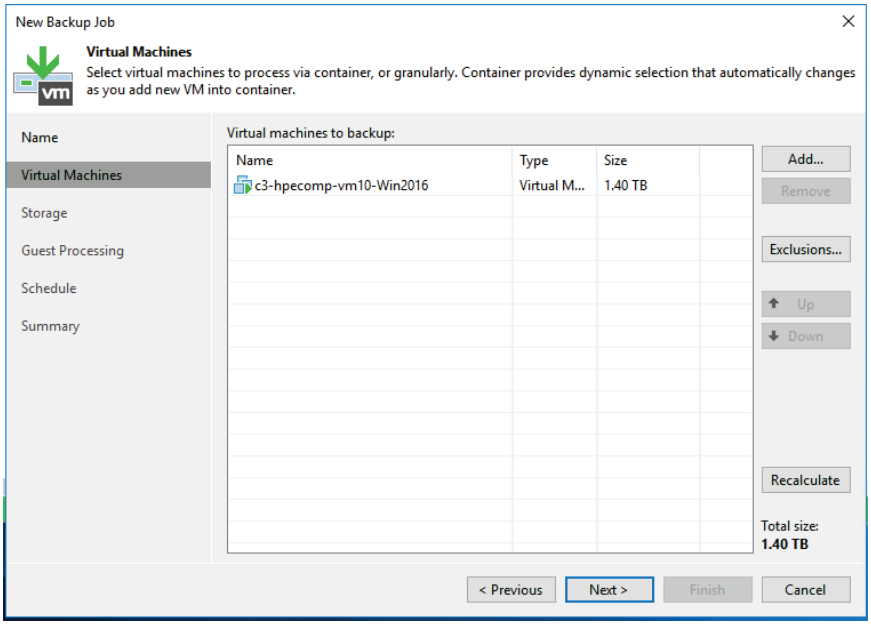

Figure 51. Selected virtual machines

5. From the “Storage” screen, select the StoreOnce backup repository created during the “Setup and Configuration” section.

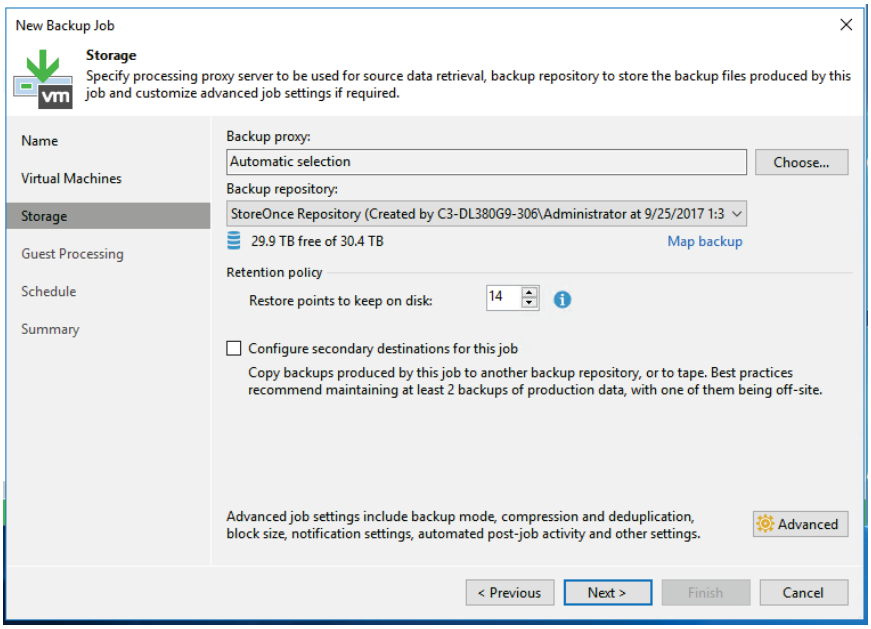

Figure 52. Backup repository selection

6. (Optional) Click on the “Advanced” button to configure advanced options such as: Backup Mode, Maintenance, Notifications, vSphere, Integration, and Script options.
7. Clicking “Next” from the “Storage” screen will pop up a recommended settings window. Select “Yes” to change the storage settings to recommended values.

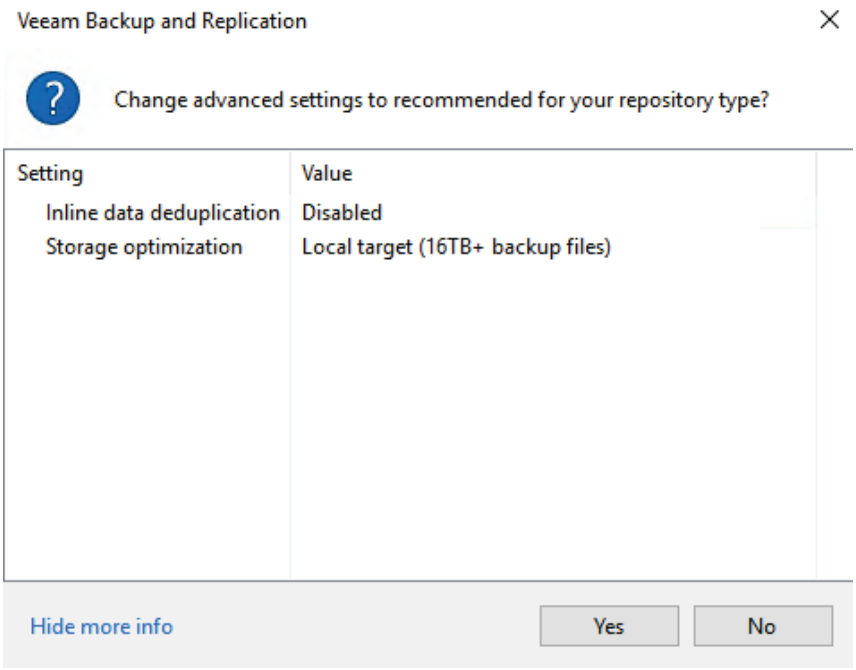

Figure 53. Settings change confirmation

8. In the “Guest Processing” window, check the “Enable application-aware processing” box and provide the guest OS credentials for the VMs being backed up.
   a. If the guest OS credentials have not already been set up, click the “Add …” button to create the corresponding credentials.

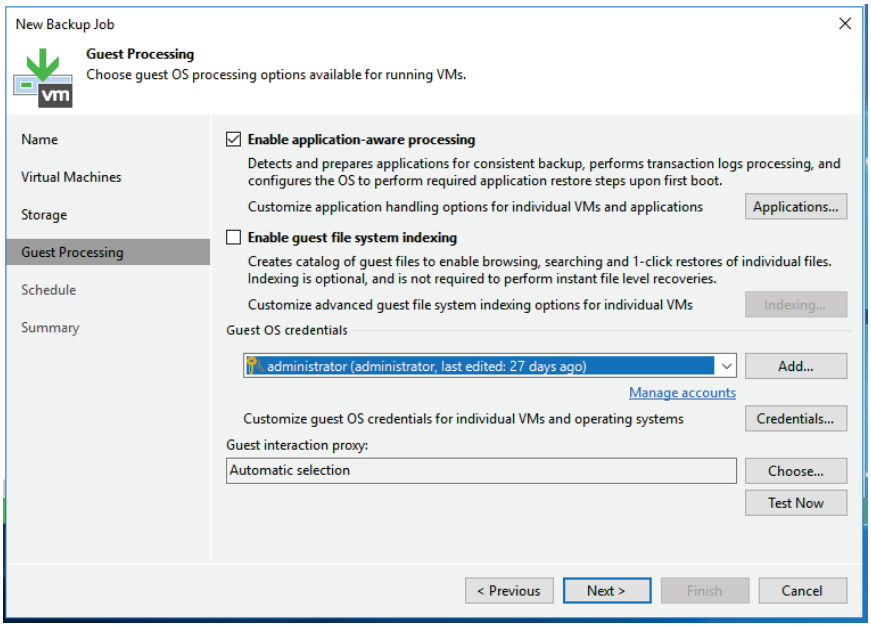

Figure 54. Enable Application-aware processing

9. Once ready to proceed, click “Next”.
10. In the next screen, the scheduling options are presented. The job can be scheduled to run at predetermined intervals or left to be executed manually. Once the scheduling options have been selected, click “Apply” to proceed.

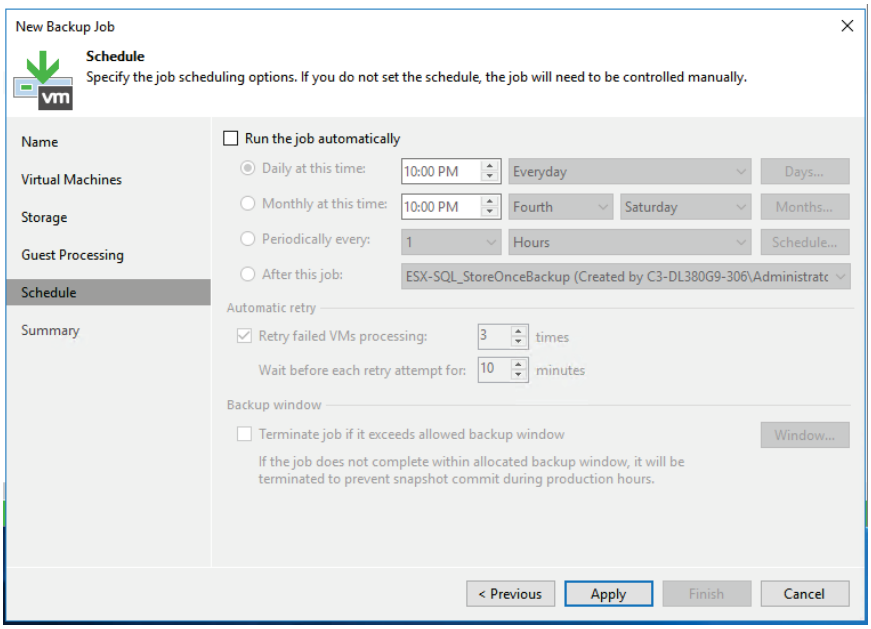

Figure 55. Backup job schedule

11. In the next screen, a summary of the selected options and configuration is presented. Review the information provided for accuracy and click "Finish" to complete the job creation process.

12. The job will now show up in the "Backup" section in Veeam.

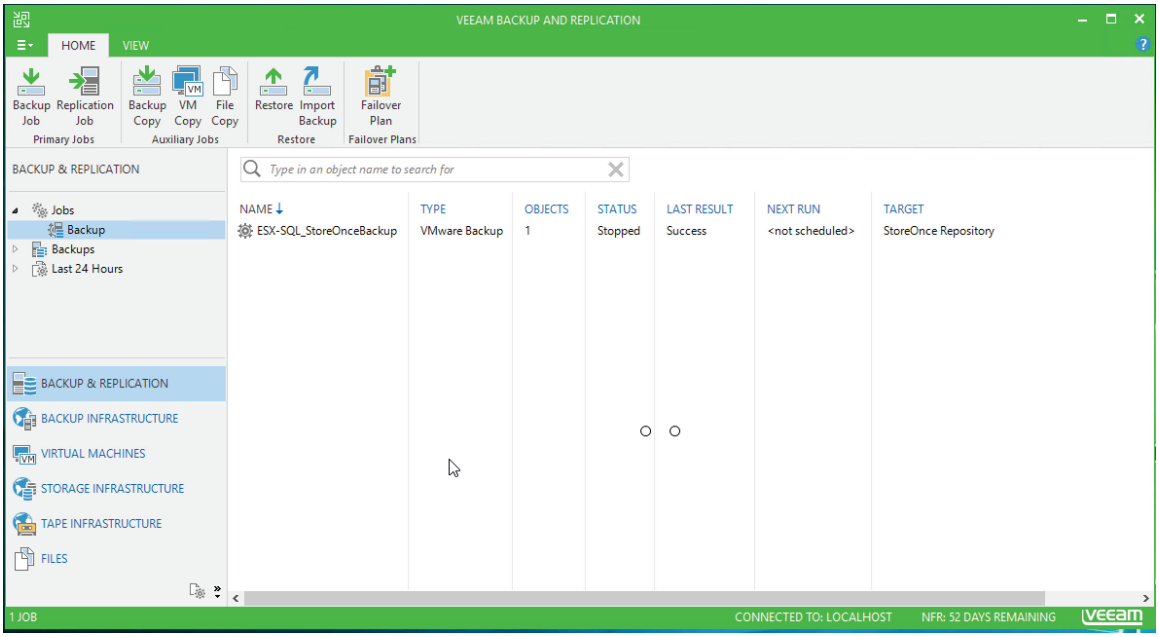

Figure 56. Backup job created

13. If the job was scheduled to be run automatically, it will do so at the specified time. If a schedule was not created, select the job and click on the "Start" button from the "JOB" tab.

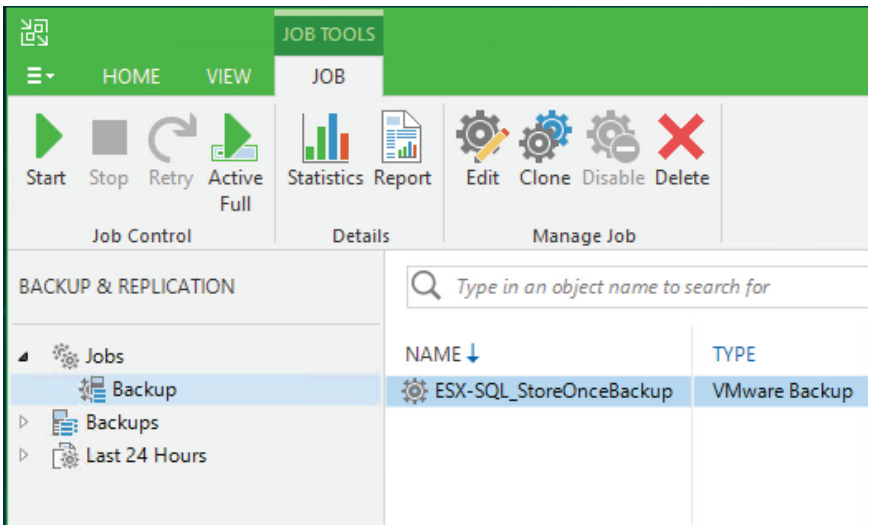

Figure 57. Start job

14. The job will start running immediately and a summary of progress is displayed at the bottom of the screen when the running job is selected.

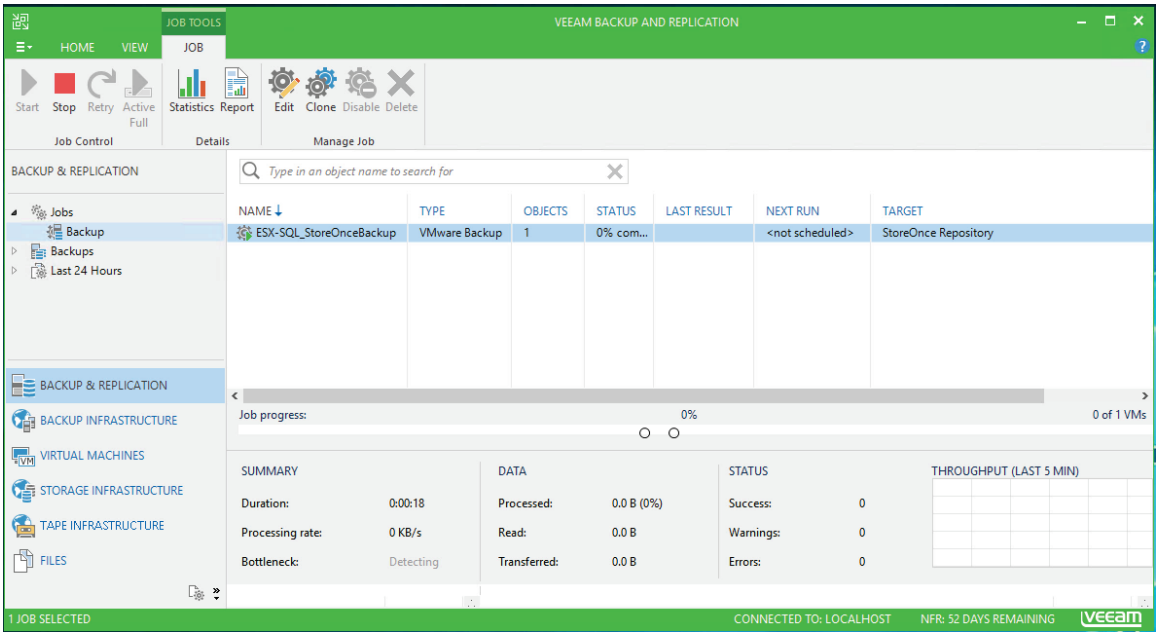

Figure 58. Job progress and summary

15. If viewing the details of the running job is desired, select the job and click on "Statistics" from the "JOB" tab.

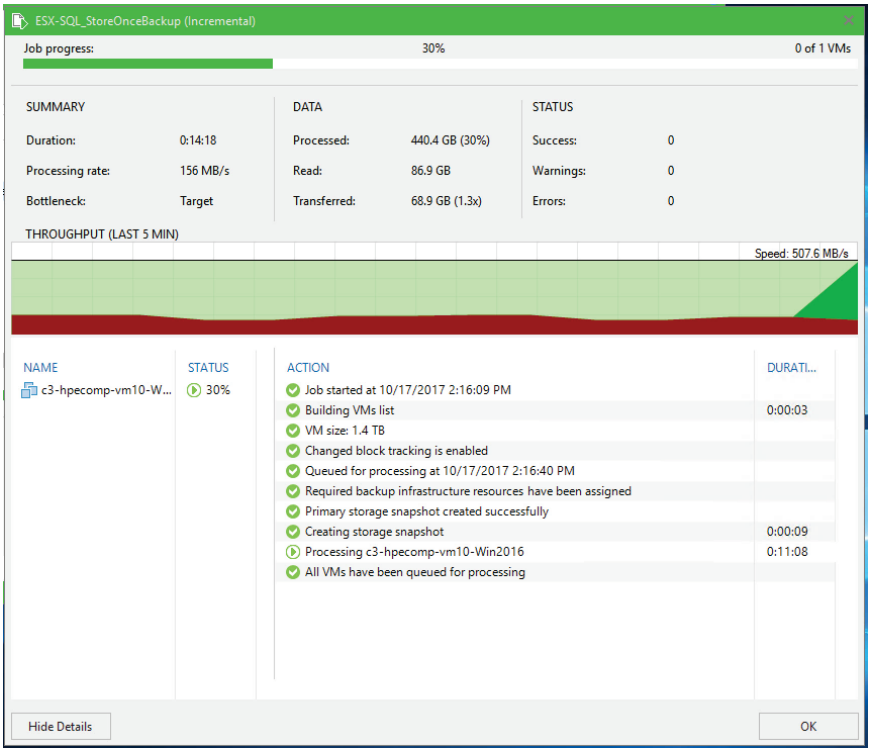

Figure 59. Backup job details

16. Once the job completes, the details of the job will indicate so and a restore point for the VM selected will be created and should be visible from the "Backups" -> "Disk" section.

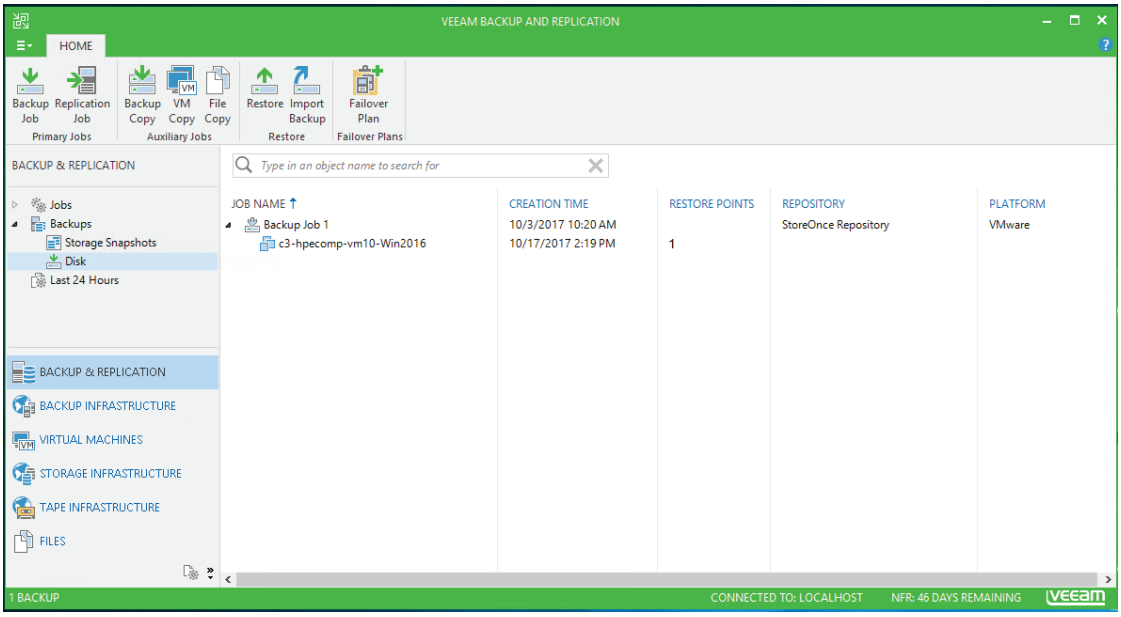

Figure 60. Restore point created upon completion of job

Use case 4: Recover database from on-premises repository

Overview

In the previous Use Cases, we covered how to create StoreOnce centric backups for your VMs and applications. In this section, we discuss how to use those backups to quickly recover your environment should it become unusable.

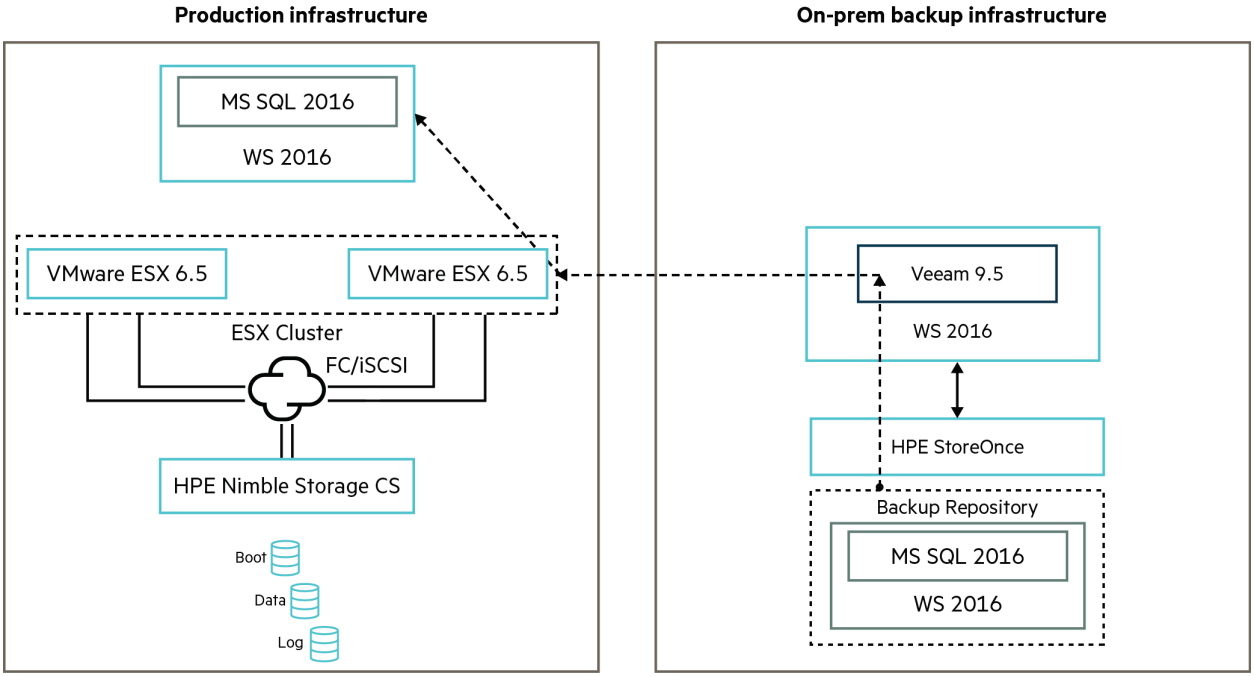

Figure 61. VM restore using Veeam

Steps

1. Turn off the VM to be restored.
   a. If the VM is powered on during the restore process, it will be turned off automatically as part of the process.
2. From the “Backup & Replication” view, select Backups -> Disk and ensure you have at least one restore point for the VM you are looking to recover.

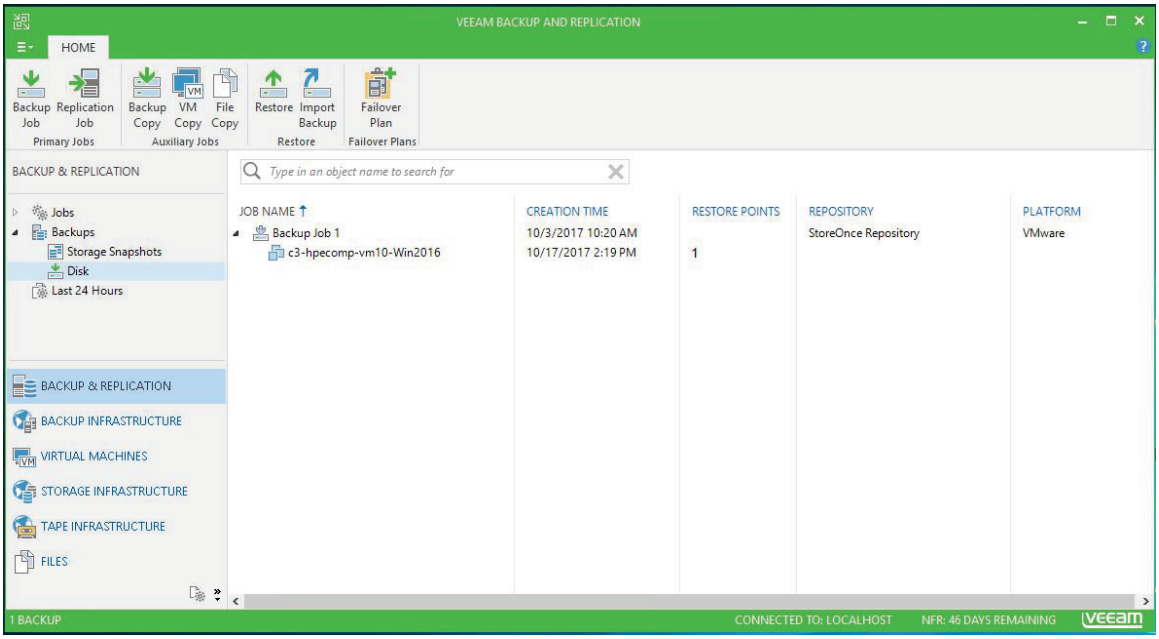

Figure 62. Available restore points

3. Select the Backup job, then from the “Backup” tab, select “Entire VM”.

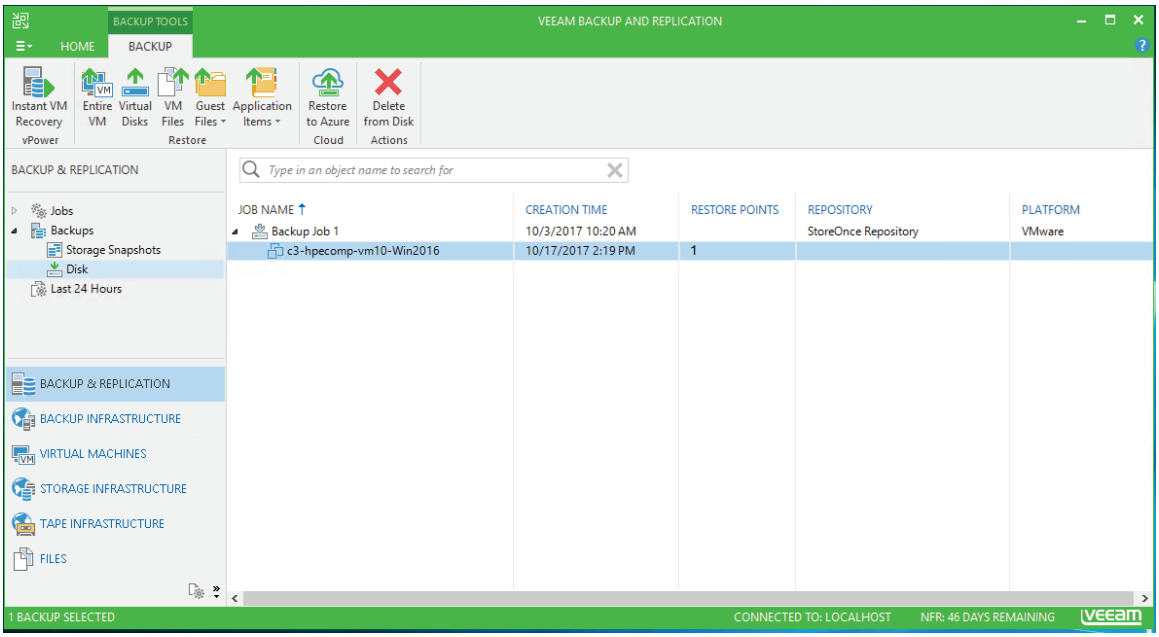

Figure 63. Initiate a Full VM restore

4. The “Full Restore Wizard” is started, displaying the selected VM to be restored. Confirm that the correct VM is selected and click “Next”.

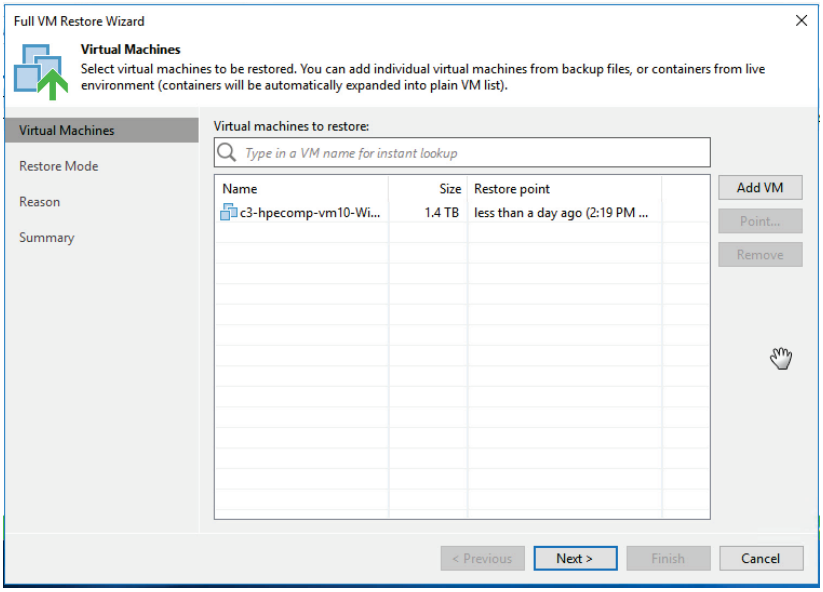

Figure 64. Selected VMs for restoration

5. From the Restore Mode page, select “Restore to the original location” to have the current VM removed and replaced with its backed up version. Alternatively, select “Restore to a new location, or with different settings” to create a new instance of the VM. Note that this will require a manual switchover from the old VM to the new VM.

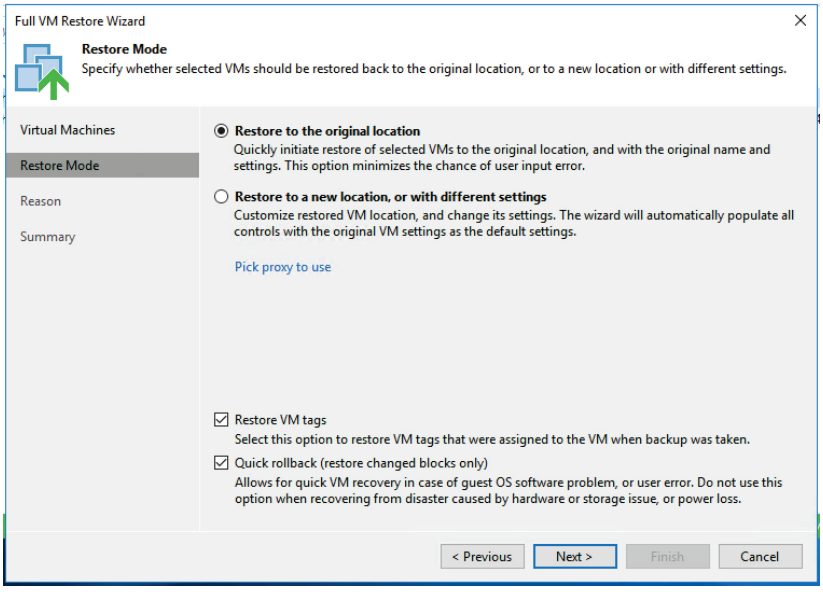

Figure 65. Select Restore to the original location

Note

Ensure that there is sufficient free space available in the restore location. If restoring to the original location, it may be necessary to manually delete the original VM and associated virtual disks from their corresponding datastores if sufficient free space is not available.

6. Also, the “Quick rollback” option can be selected to speed up the VM recovery process if “Quick rollback” requirements are met. Click “Next” when ready to proceed.
7. In the next screen, a migration reason can optionally be provided. Click “Next” to proceed.
8. Veeam then displays the VMs that will be affected by the recover process. Review the list, ensure the list is as expected, and click “OK”.

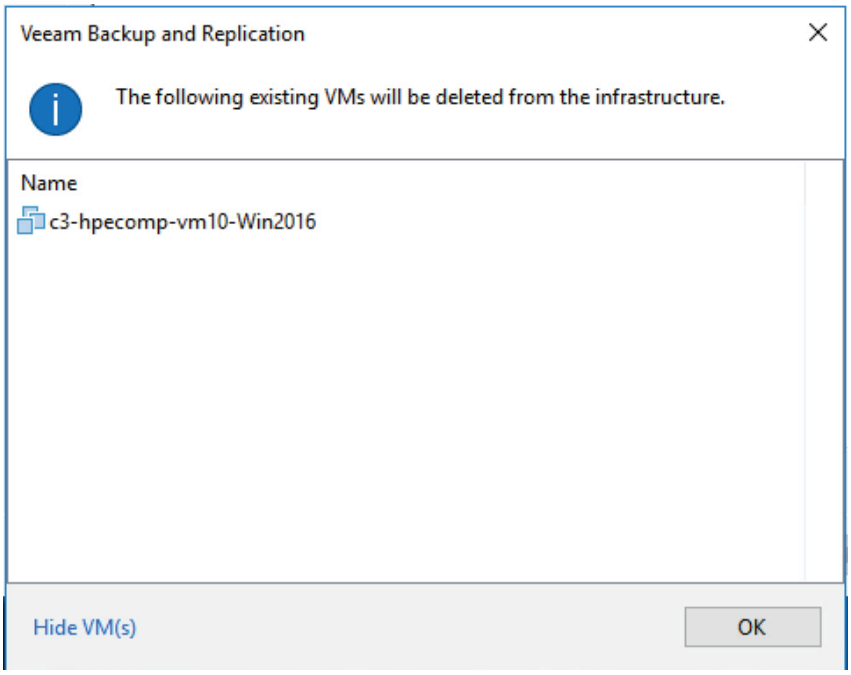

Figure 66. Affected VM confirmation list

9. The Summary screen is displayed. Verify the details shown. Optionally select “Power on target VM after restoring” to have the VM automatically powered on immediately it being successfully restored.

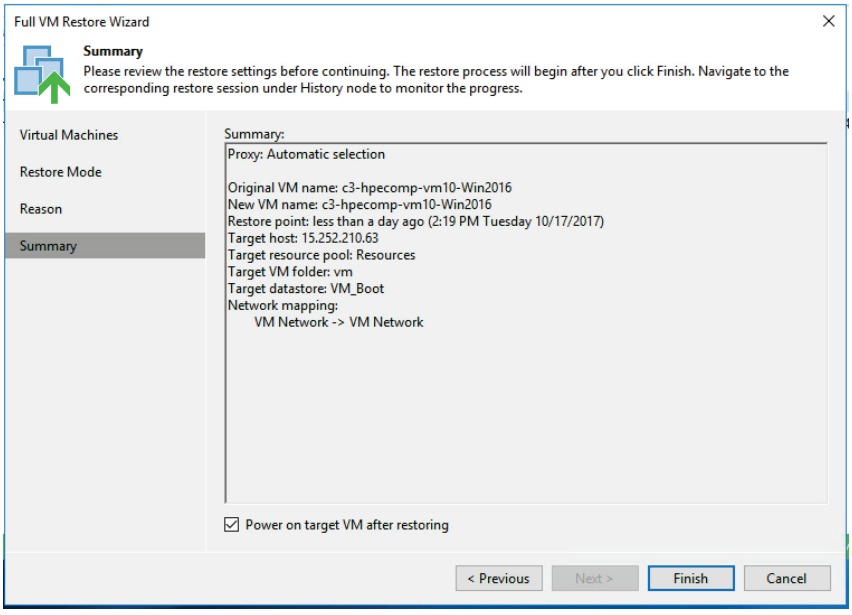

Figure 67. Restore Summary page

10. Click on “Finish” to commence the restoration process.
11. The VM restore details window automatically opens and provides details on the restoration process.

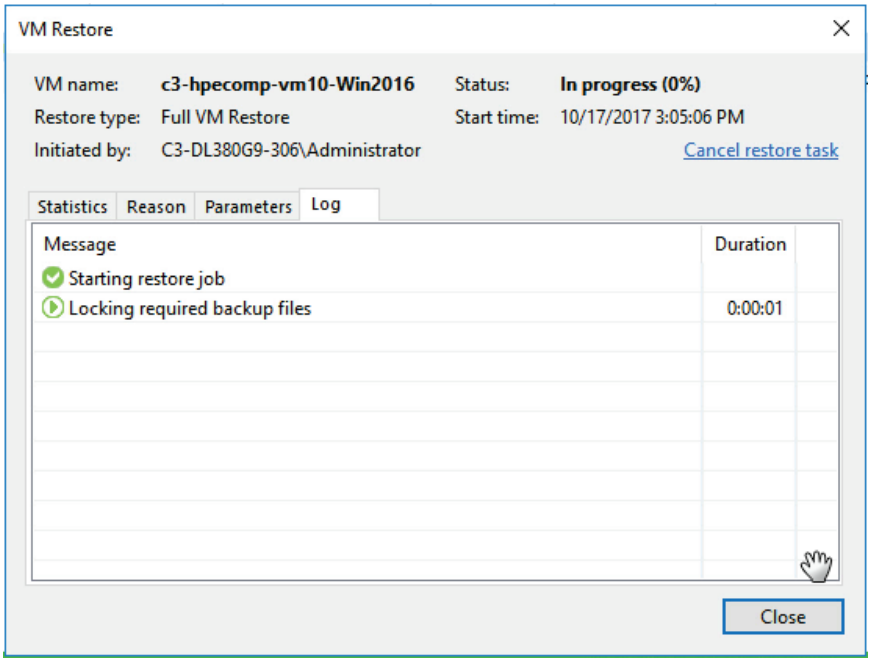

Figure 68. Restoration initiated

12. Once complete, the restore window details indicates success and the VM will reflect the state at the time it was backed up.

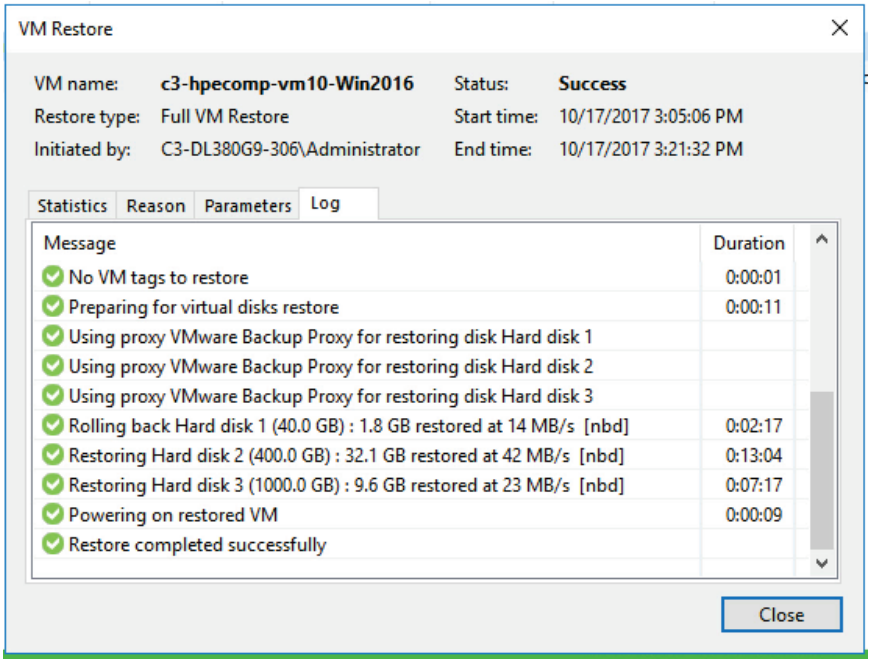

Figure 69. Restoration complete

Use case 5: Create an off-site backup copy on Microsoft Azure

Overview

Creating an off-site backup can protect business assets from most disaster scenarios. A backup copy can be stored in the cloud and even reside in a different region, which can protect it from outages that would otherwise result in data loss. Veeam Backup & Replication can copy backup files from the on-premises backup repository to the cloud, eliminating impact on production infrastructure. In Microsoft Azure, backup copies can be stored in any region or multiple regions, maximizing data protection for the most critical data.

First, a backup job must be created to backup production data from the production tier to the on-premises backup tier. In the previous use case, a virtual machine (VM) hosting a Microsoft SQL Server database was protected using a Veeam Backup Job. This job used the HPE Nimble Storage integration with Veeam to back up the VM from an HPE Nimble Storage snapshot to an HPE StoreOnce Catalyst store. The backup files can now be copied from the on-prem Catalyst store to Microsoft Azure using Veeam Cloud Connect for the Enterprise. The resulting backup files have the same format as the source files, so they can be recovered using any Veeam recovery method. The backup copy job runs on predetermined intervals and incrementally updates the backup copy when a new source restore point is detected on-prem. Since anything after the initial copy is incremental, the amount of data subsequently transferred to the cloud repository is minimal.

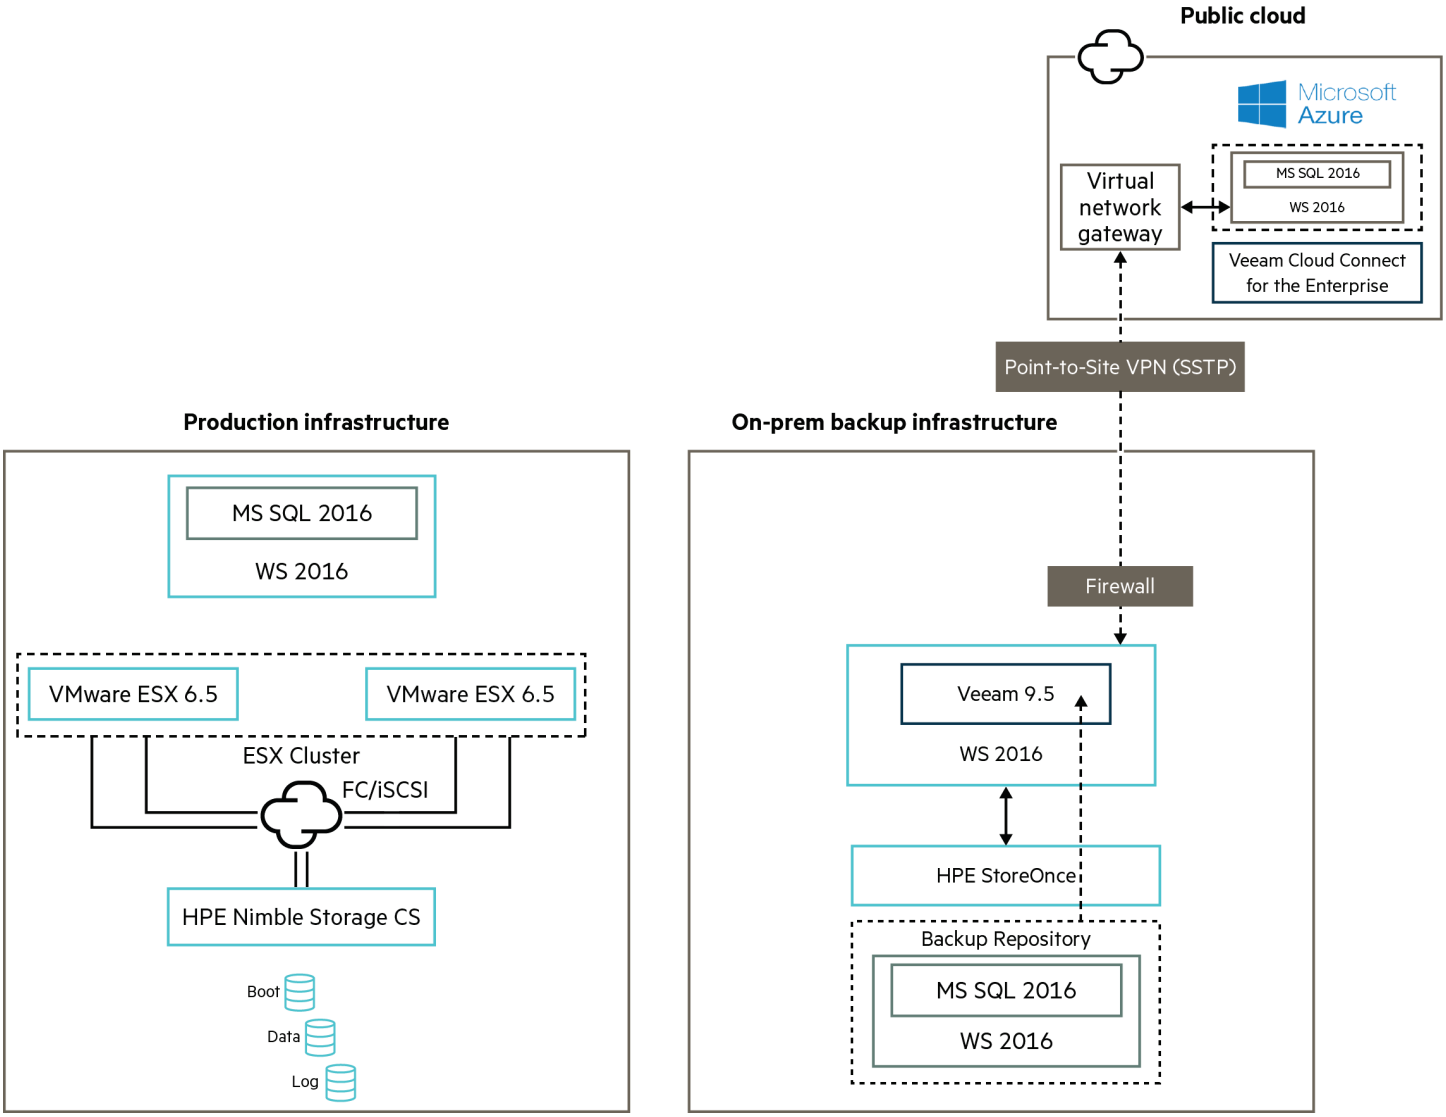

Figure 70. Off-site backup infrastructure

Steps

1. Under the Backup & Replication tab, choose “Backup Copy”.

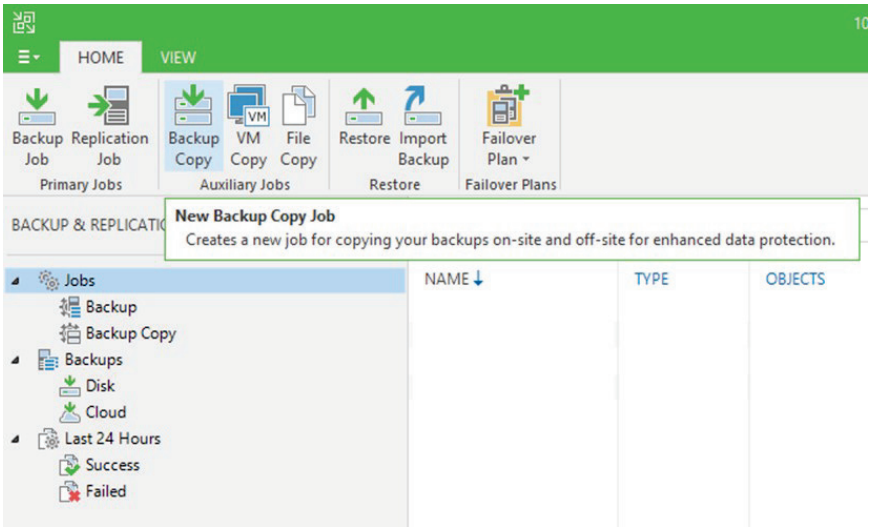

Figure 71. Backup copy creation

2. Specify the name and frequency of the Backup Copy Job.

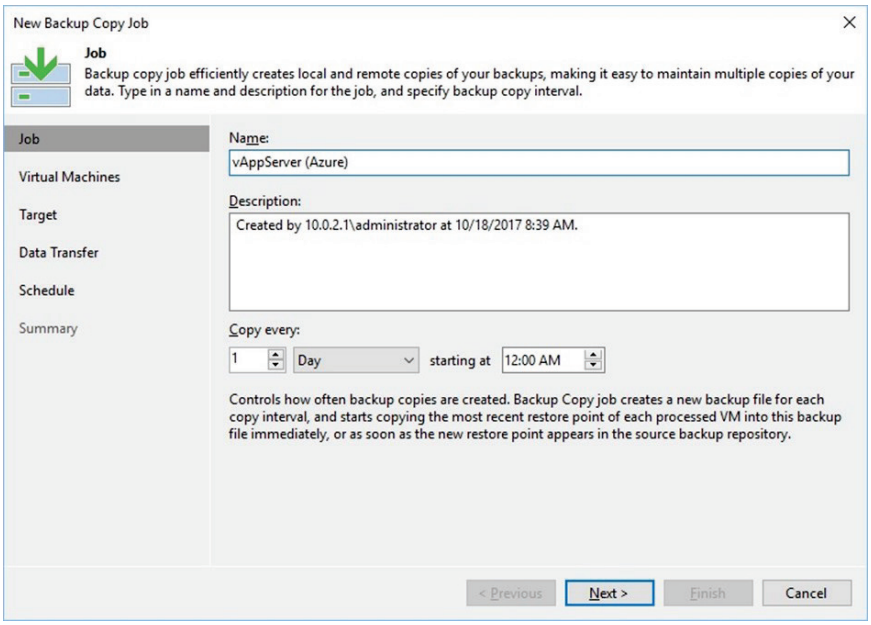

Figure 72. Copy job name and frequency

3. Add a backup job to define the source of the copy. Regardless of how you select the source of the copy, the source will always be existing backup files. Consider using a job for the source, as the backup copy job will not need to change if the source job changes.

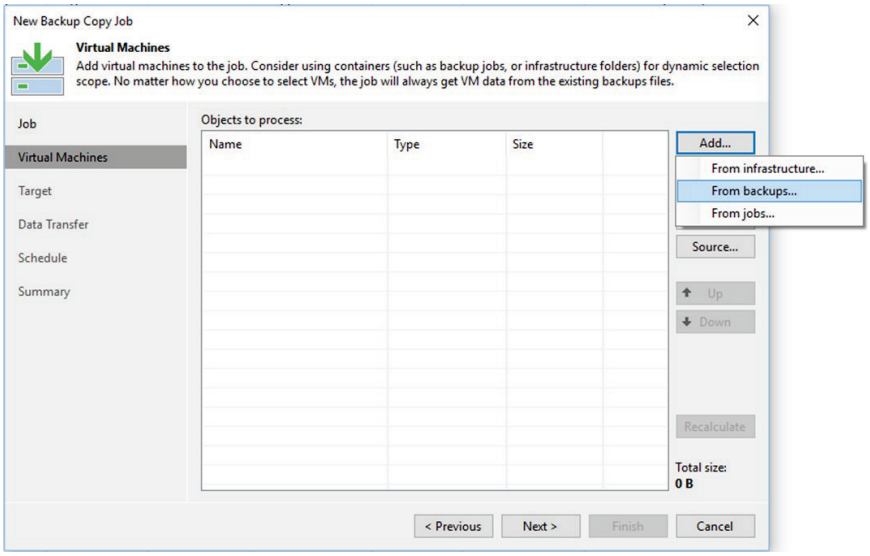

Figure 73. Create copy from backups

4. Choose the Cloud Repository as the backup repository.

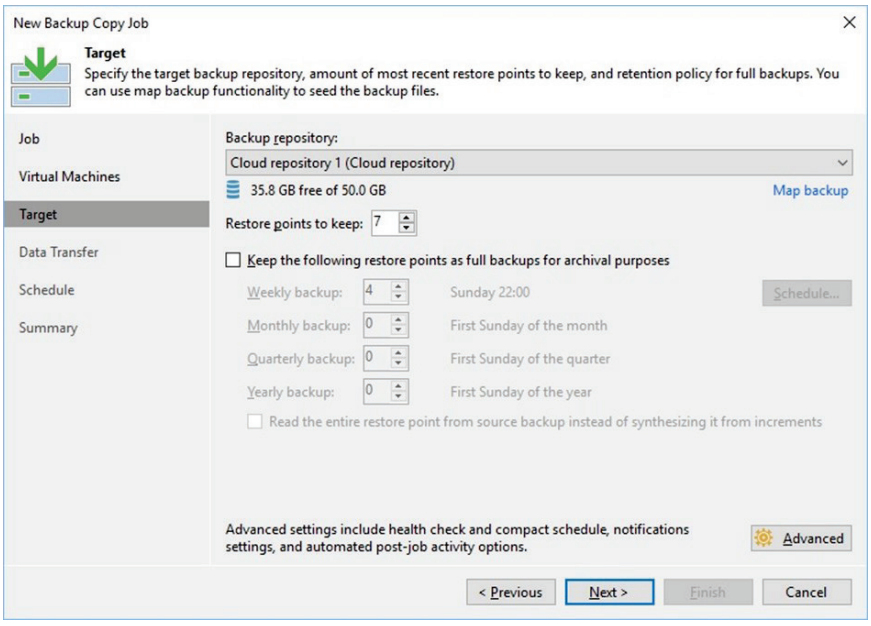

Figure 74. Cloud repository selection

5. Determine when the transfer of data is allowed over the network.

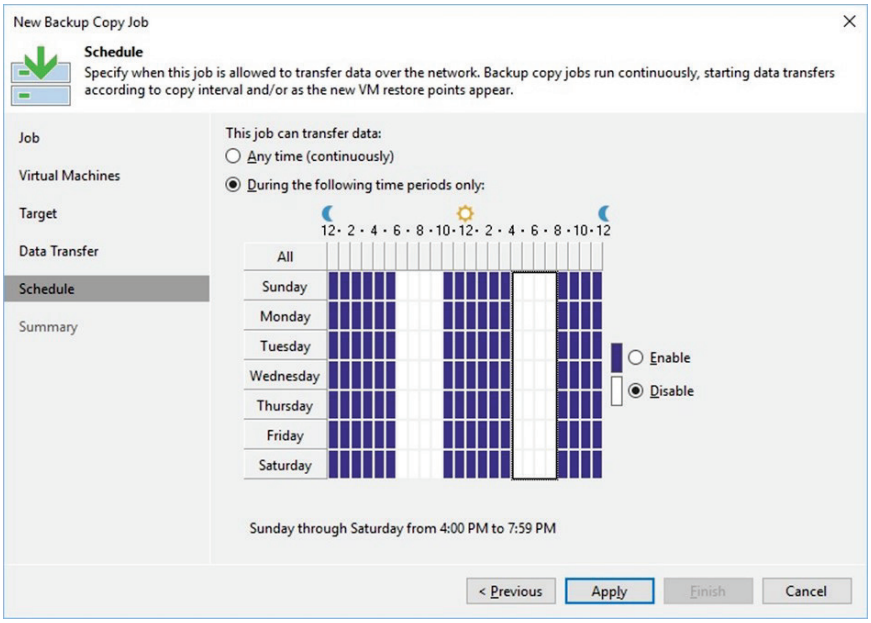

Figure 75. Data transfer schedule

6. Finish the New Backup Copy Job Wizard.

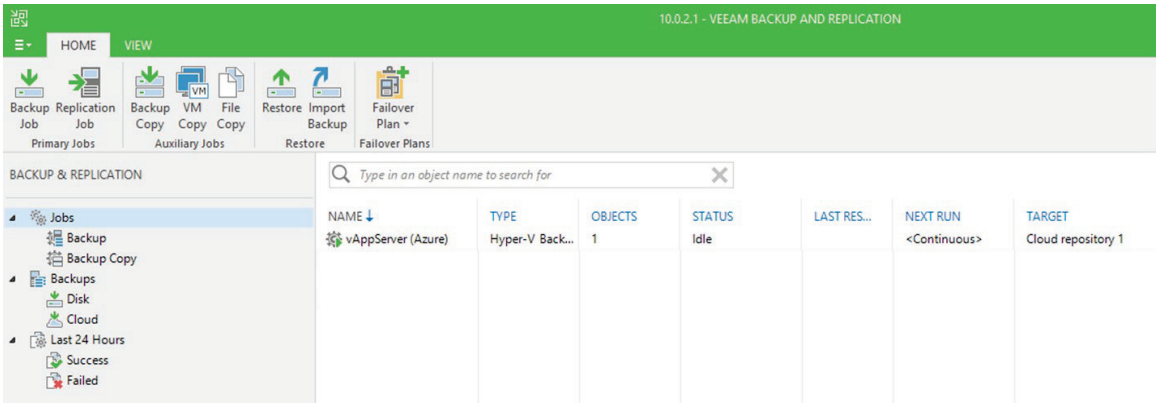

Figure 76. New backup copy created

Use case 6: Recover database from off-site backup copy on Microsoft Azure

Overview

Off-site backup copies can be restored as easily as they were created. Backup files that were copied to Microsoft Azure using Veeam Cloud Connect can be restored to the original or even different infrastructure. In the case of a disaster, the original infrastructure may be unavailable for a significant amount of time, if not indefinitely. Recovering VMs from the public cloud can assure business critical data is recovered as quickly as possible.

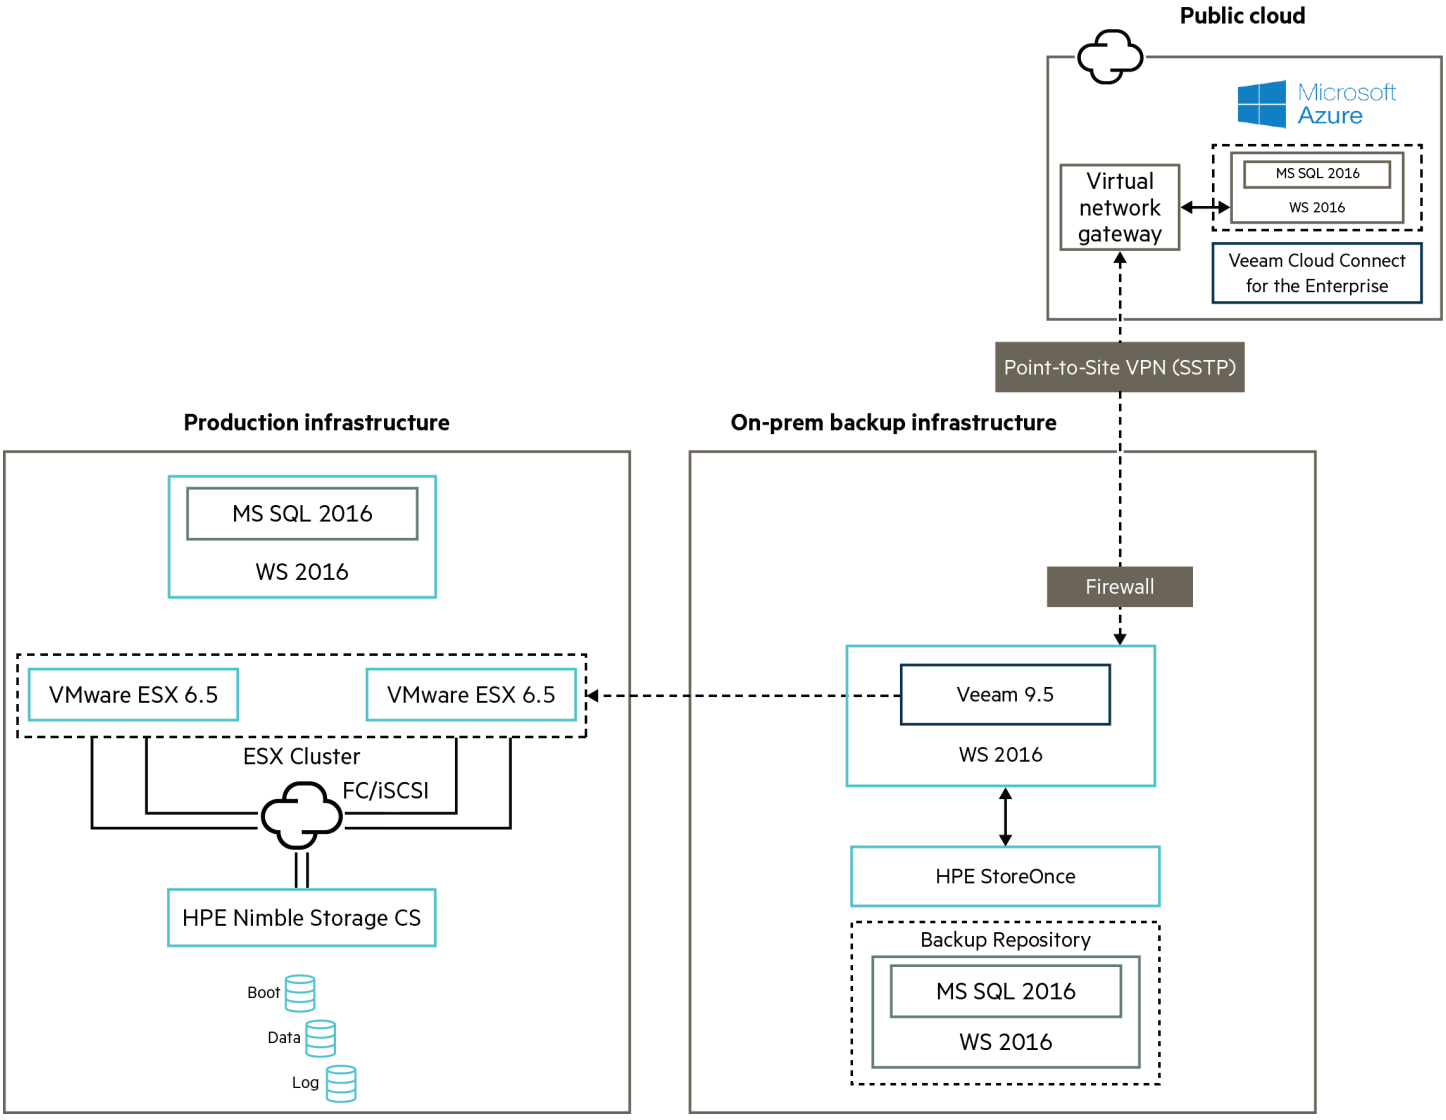

Figure 77. Recovery from an off-site backup copy on Azure

Steps

1. Under the Backup & Replication tab, choose “Restore”.

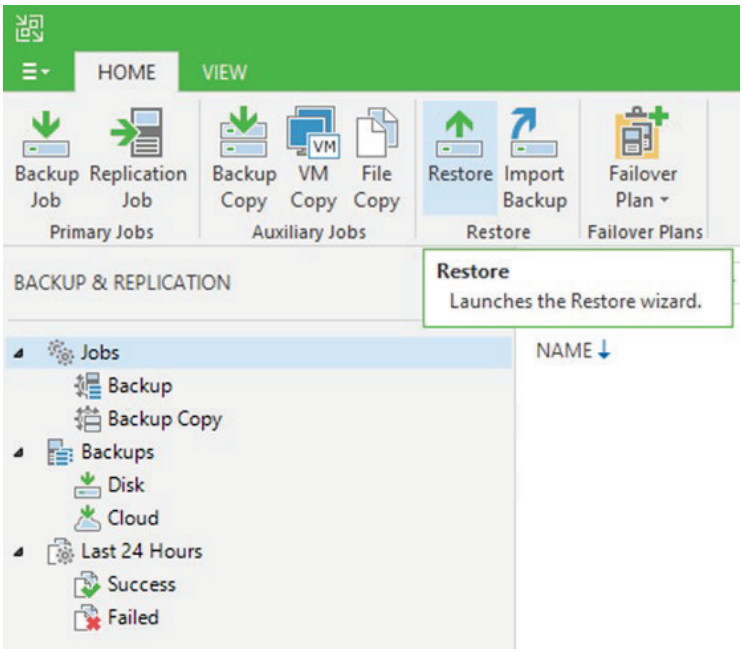

Figure 78. VM restore

2. Choose “Entire VM”.

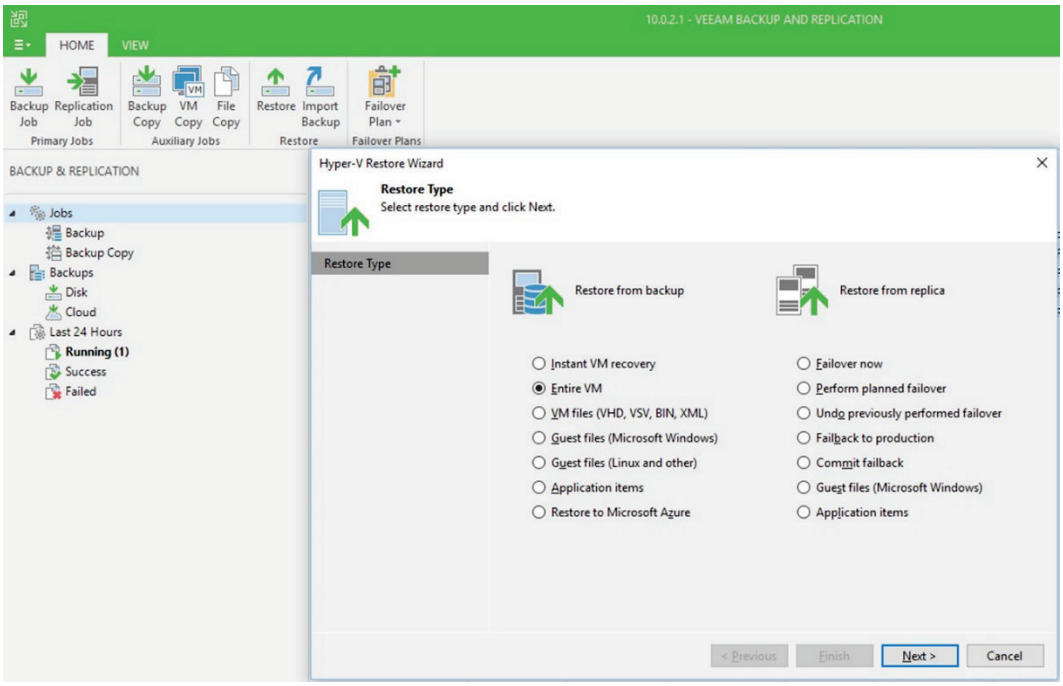

Figure 79. Restore type

3. Choose the backup to restore and validate the proper restore point is chosen.

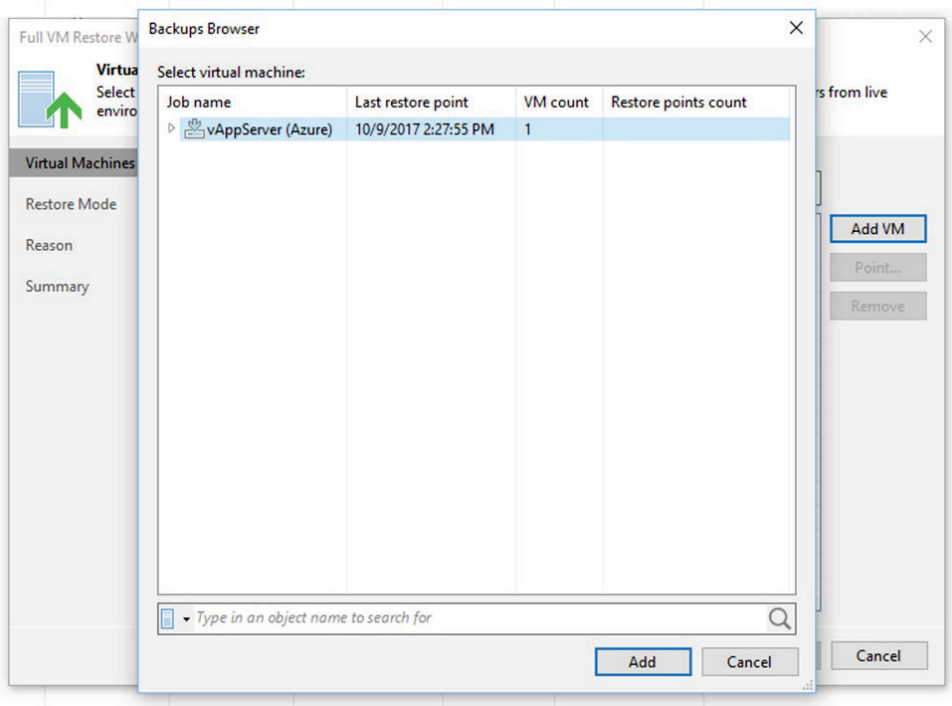

Figure 80. Backup browser

4. Choose the restore location.

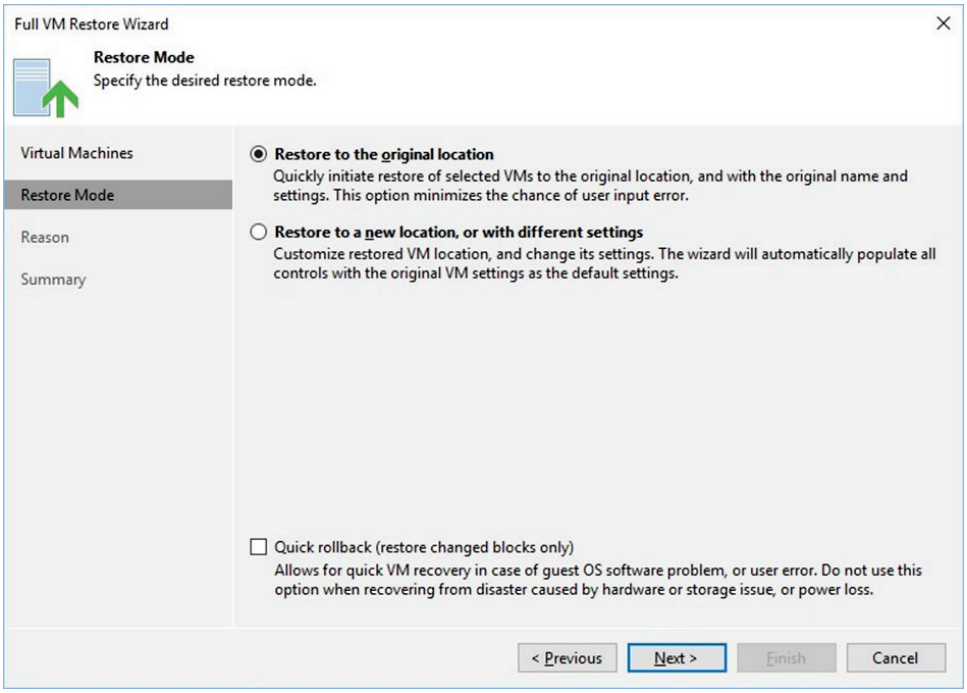

Figure 81. Restore mode

5. Validate summary and finish the Full VM Restore Wizard.

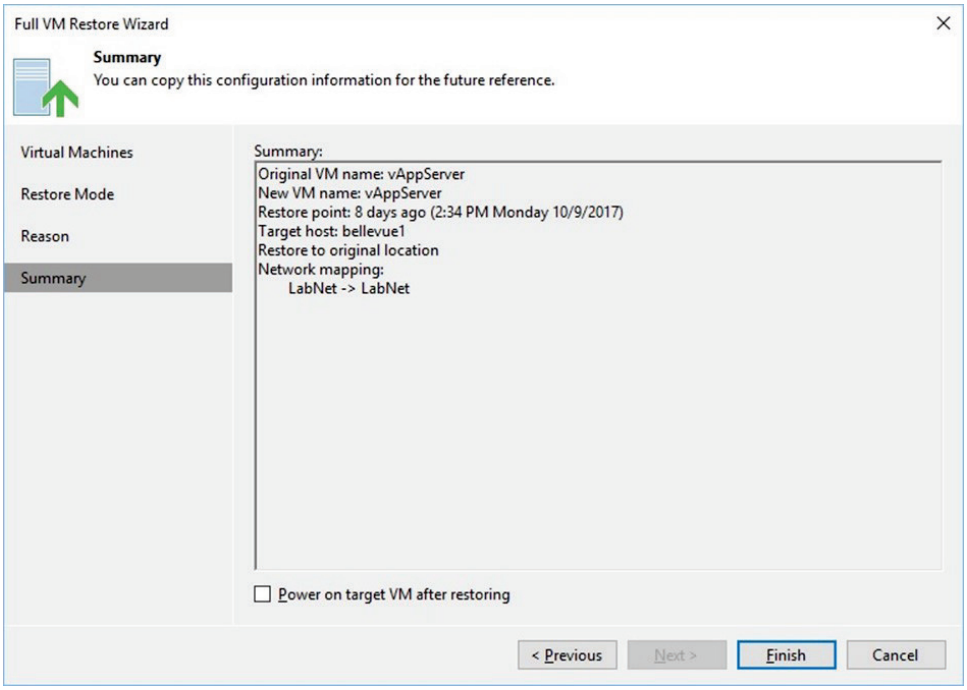

Figure 82. Full VM restore summary

Conclusion

As demonstrated in this reference architecture, hybrid cloud data protection can efficiently meet performance, RTO/RPO, and regulatory requirements, while optimizing the cost of data protection. Whether it's for performance or regulatory concerns, data can be stored on-premises. For other types of data that lack the same stringent requirements, a public cloud option improves flexibility and cost of any BCDR solution. The hybrid cloud model can future-proof a BCDR solution and allow businesses to move data from on-prem to the cloud and from the cloud to on-prem with ease.

The power of HPE ProLiant servers, scalability of HPE Nimble Storage, and efficiency of HPE StoreOnce can provide the base for any production and BCDR solution. The native integration of these HPE products with Veeam Backup & Replication, makes for an on-premises BCDR solution that is unmatched. Using Microsoft Azure and Veeam Cloud Connect for cloud backup bolsters the solution and provides "next level" flexibility and cost optimization for modern business.

Combining the on-premises prowess of HPE, data protection of Veeam, and cloud leadership of Microsoft Azure provides a BCDR solution for all data, now, and into the future.

Learn more at

HPE Complete Veeam Software

Sign up for updates

a00035571ENW, November 2017