REFERENCE ARCHITECTURE

Veeam and Pure Storage Security Blueprint

Reference architecture for Veeam Backup and Replication with Pure Storage FlashArray//C

Contents

Introduction

Organizations face significant data management challenges, such as balancing rapid data growth with efficient storage utilization and cost management. Ensuring data integrity and availability amid rising ransomware and other cyber attacks—while meeting stringent RPOs and RTOs—adds further complexity. Modern data protection requires advanced software and fast storage media. Pure Storage® FlashArray//C™ delivers an NVMe, all-flash experience at the economics of spinning disks, eliminating compromises on cost, performance, or ransomware mitigation. Combining Veeam Backup & Replication (VBR) with FlashArray//C and SafeMode™ snapshots ensures agility, performance, resilience, and streamlined recovery from ransomware, malware, viruses, and administrative mistakes.

At the core of this architecture is the seamless integration between VBR and FlashArray//C, leveraging Veeam's Universal Storage API V2 (USAPI) to unlock advanced capabilities and simplify the user experience. This integration enables VBR to efficiently orchestrate FlashArray™ snapshots and replication processes. This guide provides detailed insights and best practices to help you optimize your backup and security infrastructure, ensuring resilience and scalability in your data protection strategy.

How to Use This Guide

This guide offers valuable insights to effectively utilize the reference architecture to enhance performance and simplify processes. By following the outlined best practices, you can achieve optimal efficiency and streamline your workflow. The intended audience for this guide includes system architects, systems engineers, IT managers, backup and storage administrators, among others.

Pure Storage and Veeam: An Optimal Solution

Backup and recovery are critical components of modern IT infrastructure. However, organizations face several challenges in effectively managing these processes. Ever-increasing data sizes need more and more backup storage. Backups still have to complete within defined backup windows, and long-running backups can impact the performance of production services. As the last line of defense against ransomware, backups need to be protected from malicious actors but still recoverable quickly in a downtime event. Pure Storage and Veeam address these challenges with an easy-to-use, integrated, high-performance backup and recovery solution that offers advanced data reduction, comprehensive data protection, scalability, and reliability. The advantages of using this optimal solution include:

Optimizing backup storage: As data volumes continue to grow exponentially, the challenge of optimizing storage in backup solutions is escalating, crucial for managing costs and utilizing storage resources.

- FlashArray//C brings always-on data reduction to VBR backup storage, driving efficiency, optimizing storage density and reducing the total cost of ownership (TCO) without compromising performance.

Faster backup and recovery performance: Organizations grapple with meeting stringent backup windows and recovery time objectives (RTOs) while minimizing the impact of backup on production workloads. While backups to disk-based storage may be fast, recovery is often slower due to longer seek times and slower random reads. FlashArray//C delivers both fast backup and recovery, and recovery from FlashArray//C can be even faster than backup.

- Deploying VBR with FlashArray//C enhances performance, enabling backup and recovery at flash speed.
- Veeam Explorers facilitate granular and quick data restoration from FlashArray snapshots, thereby reducing the RTOs.

- Pairing VBR with FlashArray//X™ for primary storage further improves performance, simplifying administration of VMware datastores and seamlessly incorporating hardware snapshots into data protection processes through Veeam's Universal Storage API (USAPI V2). This differentiated capability allows VBR to orchestrate FlashArray snapshots. For example, with the creation of a FlashArray snapshot on the production array, replicate that snapshot to another FlashArray, and then perform the backup from the replicated snapshot. This not only offloads highly-intensive backup operations from production systems, but also increases resiliency, as the snapshot copies can be replicated off-site before being backed up to a VBR Backup Repository.

Ransomware protection: The constant threat of ransomware presents a significant challenge to data integrity and business continuity.

- SafeMode snapshots provide an additional layer of protection against ransomware attacks, ensuring data integrity and security of both the production storage and the backup storage. Leveraging Veeam Hardened Repository with Pure Storage SafeMode immutable snapshots enhances data security and resilience. Furthermore, FlashArray//C is Veeam Ready validated, instilling confidence in its deployment and reliability.

Pairing VBR with Pure Storage solutions provides organizations with optimal solutions for ensuring effective data protection, availability, and cyber resilience in modern IT environments. This integration empowers organizations to safeguard critical data assets with confidence.

Reference Architecture

The environment depicted in Figure 1 illustrates the solution's logical architecture, designed for a datacenter containing 1000 VMs and 100 physical servers. The source VMware vSphere VMs reside on VMFS datastores on FlashArray//X, presented over Storage Area Network (SAN). While other production storage options are supported, pairing VBR with FlashArray//X enhances performance, leveraging flash speed without added complexity.

Veeam Backup Proxies read VM data using one of the Veeam transport modes and transfer it to the Veeam Hardened Repository sitting on SAN-connected FlashArray//C volumes. VBR optionally adds deduplication, compression, in-transit encryption, and inline malware detection. The FlashArray//C70 provides encryption at rest. For scalability, additional scale-out Backup Proxies can be added as needed.

VBR can utilize either the backup data or snapshots for traditional, granular or instant recovery operations.

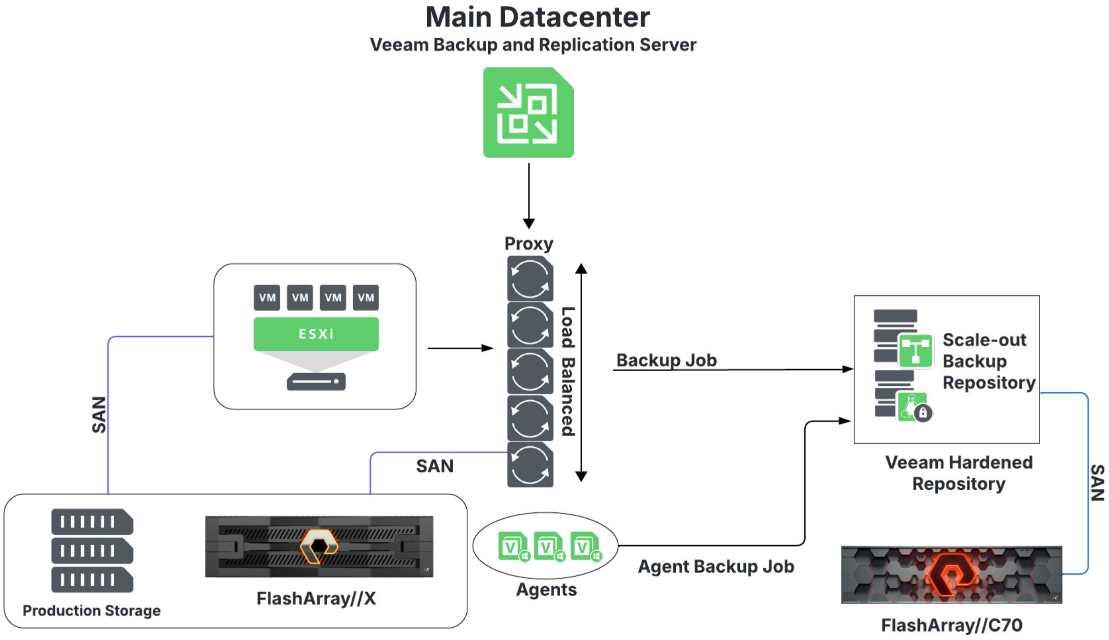

FIGURE 1 Illustration of the logical architecture of the solution

Components Used by the Solution

Pure Storage FlashArray//C

Pure Storage FlashArray//C, the world's pioneering enterprise-grade all-QLC FlashArray, offers NVMe performance, hyper-consolidation, and simplified management essential for your data needs. It now addresses an even wider range of business-critical applications that prioritize high capacity over sub-millisecond latency. With features like DirectFlash® managed QLC with available 75TB modules, inline data reduction, and global deduplication, FlashArray//C optimizes data density, making it an ideal solution for high-performance backup storage. Proven 99.9999% availability, non-disruptive upgrades, and consistent single-millisecond latency make FlashArray//C a reliable home for your value-optimized storage workloads and data. With a modular, upgradable Evergreen™ architecture, you'll never need to do another forklift upgrade or storage migration.

FIGURE 2 FlashArray//C

Veeam Backup and Replication

Veeam Backup and Replication provides seamless backup for virtual and physical environments, including VMs, physical servers, and cloud instances, and replication of VMs for orchestrated failover and failback. Its intuitive interface and ease-of-automation streamline operations, while integration with leading storage platforms ensures compatibility. Widely trusted for its ease of use and high performance, Veeam has continuously evolved VBR with improvements to performance, scalability, and ransomware protection. Included deduplication and compression reduce storage needs and optimize backup performance, while parallel processing maximizes throughput. VBR offers a comprehensive range of recovery options, including instant recovery, full recovery and granular recovery including file-level and application item–level, minimizing downtime and ensuring business continuity with rapid restoration directly from either storage snapshots or backup files.

Components of the Design Guidance

Veeam Backup and Replication Server

The VBR server is the central component of the Veeam Backup and Replication infrastructure, responsible for orchestrating backup, VM replication, and recovery tasks. It serves as the management point for configuring and monitoring backup jobs, managing backup repositories, and controlling data transfer between production VMs, storage, and backup targets. We advocate for running a VBR Server on a virtual machine. Virtualizing the VBR Server offers several benefits including improved flexibility, scalability and resource utilization.

Some recommendations for running the VBR Server within a virtual machine are:

- We recommend VBR version 12.1 and above.
- It's advisable to position the Backup Server within the main data center alongside the infrastructure being protected. This ensures quick response times and facilitates local management traffic.
- The minimum recommended system configuration for the Backup Server is 4 CPU cores and 8GB RAM. Please follow the sizing guidelines on the Veeam Backup & Replication Best Practices Guide.
- To optimize performance, scale the VBR server configuration according to the number of concurrent jobs. As a guideline, reserve 1 CPU core and 4GB of RAM for VBR operations, and allocate an additional 1 CPU core and 4GB of RAM for every 10 concurrent jobs. Concurrent jobs include active backup or replication tasks, as well as jobs with continuous schedules, such as backup copy, database log backups, and tape operations.

By adhering to these guidelines and scaling resources appropriately, you can ensure not just the VBR server, but the entire VBR enterprise infrastructure operates efficiently and effectively manages the workload.

Veeam Agents

Veeam Agents enable the backup and restoration of workloads that cannot use agentless VM snapshots, such as physical computers, certain public cloud VMs, and VMs that cannot have VM snapshots created. VBR allows you to centrally deploy and manage Veeam Agent for Microsoft Windows and Veeam Agent for Linux on computers within your infrastructure directly from the VBR console.

Veeam Agent management involves additional elements in the VBR console, such as:

- **Protection groups**: Container in the VBR inventory designed to group protected computers by type, workload or other criteria.
- **Agent backup jobs:** Scheduled on the Backup Server to process one or more protection groups

Backup Proxies

Backup Proxies serve as the data movers and play a crucial role in achieving efficient backup and restore speeds. Selecting the optimal Backup Proxy server design is crucial for your environment and offers significant control over the impact on the vSphere infrastructure and the flow of backup traffic. Depending on your chosen VMware vStorage APIs for Data Protection (VADP) transport mode, you may need virtual Backup Proxies (such as Hot-Add, also referred to as Virtual Appliance Mode) or physical Backup Proxies (employing Direct SAN (DSAN) Access via iSCSI or FC/Backup from Storage Snapshots).

Additionally, VBR Backup Proxies can provide data reduction, in-flight data encryption, and real-time malware detection.

Transport Mode

A transport mode is a method that is used by the Veeam Data Mover to retrieve VM data from the source and write VM data to the target. VBR supports the following transport modes.

- **Network block device (NBD):** In the NBD mode, the VBR Backup Proxy retrieves VM data over the network from the ESXi host where the VM resides. Subsequently, the data is then transferred to the Backup Repository. While the NBD mode provides a universal approach for data transfer and can be used in various network configurations, it may lead to slower backup and replication speeds compared to other transport modes, especially in larger environments or over WAN connections.
- **Virtual Appliance (HotAdd):** In the HotAdd mode, Veeam leverages VMware's HotAdd technology to attach virtual disks of VMs to a VBR Backup Proxy. The data is then transferred from the VM's virtual disks to the Backup Repository through the VBR backup proxy. Unlike the Direct SAN Access mode, the HotAdd mode operates within the virtual environment and does not require direct access to the underlying storage. While the HotAdd mode can offer good performance, it may introduce additional load on the ESXi hosts and the virtual infrastructure, especially in larger environments with many VMs.
- **Direct storage access (Direct SAN & Direct NFS):** In the direct storage access transport mode, VBR utilizes VMware VADP to directly transport VM data to and from storage systems, using either FC or iSCSI protocol for Direct SAN and NFS protocol for Direct NFS. This method bypasses ESXi hosts, resulting in the fastest data transfer speeds and lowest impact on the production systems.

For more information on transport modes, refer to Veeam user guide on Transport Modes.

When designing Backup Proxies, consider the following best practices:

- Each task processes 1 VM disk at a time, with CPU/RAM resources utilized for inline malware detection, data reduction, and encryption.
- Configure a maximum of two concurrent tasks for each physical core or vCPU.
- Allocate 2GB of RAM for each physical core or vCPU.
- Deploy a minimum of two Backup Proxy servers per site to ensure a baseline level of availability for this critical role.

For more information on how to achieve optimal performance, refer to Veeam's best practices guide on vSphere proxies.

Backup Repositories

The Backup Repository refers to the storage location where backup files are stored. It can be local storage, network attached storage (NAS), SAN, cloud storage, or deduplicated storage appliances. The Backup Repository server, on the other hand, can be a physical or virtual machine that is responsible for managing the Backup Repository. Veeam recommends whenever possible to use physical machines as Backup Repository servers, in order to maximize performance and enhanced data security and resilience (i.e. Veeam Hardened Repository). Other recommendations include:

- A repository task slot refers to the number of concurrent tasks the Backup Repository can handle. Each task typically represents a single VM disk being processed during backup or restore operations. The number of task slots depends on the repository's hardware resources and can be adjusted to optimize performance and resource utilization.
- Allocate one Backup Repository CPU core per three Backup Proxy cores. VBR Backup Proxies may need additional CPUs to perform compression, deduplication, encryption, and inline malware detection.
- Allocate 4GB of RAM for each CPU core.
- To protect your backup files from loss as a result of malware activity or unplanned actions, you can add to your backup infrastructure a Veeam Hardened Repository based on a Linux physical server. The Veeam Hardened Repository supports immutability and single-use credentials. The role of the Veeam Hardened Repository can be assigned to a Linux machine with local or remotely attached block storage. The Linux distribution must be 64-bit due to Veeam Data Mover requirements.
- Use fast cloning on XFS volumes. Fast Clone is the Veeam Backup and Replication technology where copying a file references existing data blocks on volumes instead of physically copying data blocks between files.
- Use Scale-Out Backup Repositories (SOBR). A SOBR facilitates horizontal scaling for multi-tier storage of data. It consists of one or more Backup Repositories or Object Storage Repositories known as the Performance Tier. There's an option to expand using object storage repositories for long-term and archive storage, known as the Capacity Tier and Archive Tier respectively.
  - The Capacity Tier extends the on-premises Backup Repository to low-cost object storage such as S3-compatible storage, enabling the offloading of older backup files from the Performance Tier. This reduces the load on primary storage while ensuring accessible backups.
  - The Archive Tier is dedicated to long-term data retention and is often linked with more cost-effective storage solutions like AWS Glacier or Azure Archive Storage. It facilitates the transfer of older backup files from the Capacity Tier to the Archive Tier, effectively minimizing storage expenses while maintaining adherence to long-term retention policies.
  - All storage devices and systems within the Scale-Out Backup Repository are integrated into a unified system, combining their capacities.
  - To use fast cloning on XFS with SOBR, you must configure the repository placement policy for “Data locality.”
- Pure Storage recommends deploying FlashArray//C as VBR Backup Repository extents of a SOBR, which offers the flexibility to expand with multiple repository volumes and servers for increased capacity and processing power. Additionally, adding VBR mount servers enhances Instant Recovery performance.
- Pure Storage recommends to use Purity 6.5.3 and above for long term supportability on FlashArray.

For more information, refer to Veeam's online User Guide on Creating Backup Repositories.

Scaling

Regularly monitor and review performance metrics of the production and backup environments, including CPU, memory, network, and storage utilization to identify any constraints or bottlenecks. You can then optimize the VBR infrastructure and configuration accordingly. Adjust your scaling strategy as your business requirements and workload characteristics evolve. When you need to scale the VBR environment, consider the following guidance:

Backup Proxies

- Add Backup Proxies to distribute the backup and recovery workloads, especially in larger environments or with resource-intensive VMs. As long as the Backup Repository can handle more concurrent tasks, adding more Backup Proxies can help shorten the overall time to backup your whole environment
- To handle more concurrent tasks on a Backup Proxy, increase its CPU cores and RAM.

ESXi Hosts

- Consider adding more ESXi hosts to increase compute resources and distribute VM workload effectively.
- Ensure each ESXi host has sufficient network bandwidth and storage resources to support VM density and workload consolidation.

Backup Repositories

- If you need more backup throughput than FlashArray//C can provide, add a second FlashArray//C as Backup Repository storage.
- FlashArray//C can handle more read operations than a single Backup Repository server can drive. If you need faster restores, and if your primary storage can accommodate additional writes, you can add a second Backup Repository server and an additional SOBR extent from the same FlashArray//C.

Pure Storage FlashArray USAPI Plugin

The Pure Storage FlashArray plugin (USAPI V2) for VBR, installed on the VBR server, enables VBR to leverage FlashArray snapshots with direct SAN access transport mode. VBR creates VM snapshots for each VM it is protecting. It then creates FlashArray snapshots of the VMFS datastore volumes, capturing the consistent VM state. VBR then removes the VM snapshots.

To finish protecting the VMs, VBR attaches the FlashArray snapshots to Backup Proxies and performs DSAN transport mode backups.

Using storage snapshots instead of production datastores as the data source for the backup job has the following advantages, especially for workloads with high change rates:

- Reduces performance impact to production VMs
- Reduces load on ESXi hosts
- Efficient and prevents against potential VM-stun during the backup job for larger VMs

For more information, refer to
Advanced Storage Snapshot Integration with Pure Storage FlashArray and Veeam Data Platform v12.

Validation Scenario

The reference architecture was designed to handle 220TB of source data across 1,000 VMs and 100 physical servers. To validate the reference architecture we simulated the environment by testing 200 VMs and 100 agents loaded with 100GB of data each, running at the maximum concurrency. We then extrapolated the results to demonstrate that the architecture can handle the full 1,000 VMs and 100 Agents. The test environment details are as follows.

We built 200 source VMs for agentless backup and 100 VMs with Red Hat Linux and Veeam Agent for Linux installed to simulate physical servers. The VMs were evenly distributed across the ESXi hosts and datastores from three FlashArray//X70 and one FlashArray//X90. Using the FIO data ingestion tool, we generated 100GB of unique random data on all the VMs and maintained a daily change rate of 10% on each VM. The VBR server reported an average data reduction rate of 1.2x. For the target, we used a Veeam Hardened Repository physical server with DISA STIGS applied. Direct-attached Fibre Channel volumes from FlashArray//C70 were mounted on the Backup Repository server. VBR utilized five Backup Proxies for running the backup and recovery jobs effectively.

Covered Use Cases

The following backup and recovery scenarios were tested in our environment.

- NBD, HotAdd, and Direct Storage transport modes.
- USAPI V2 integration
    - Backup from primary storage snapshots on FlashArray//X
    - Backup from replicated snapshots on FlashArray//C70R4.

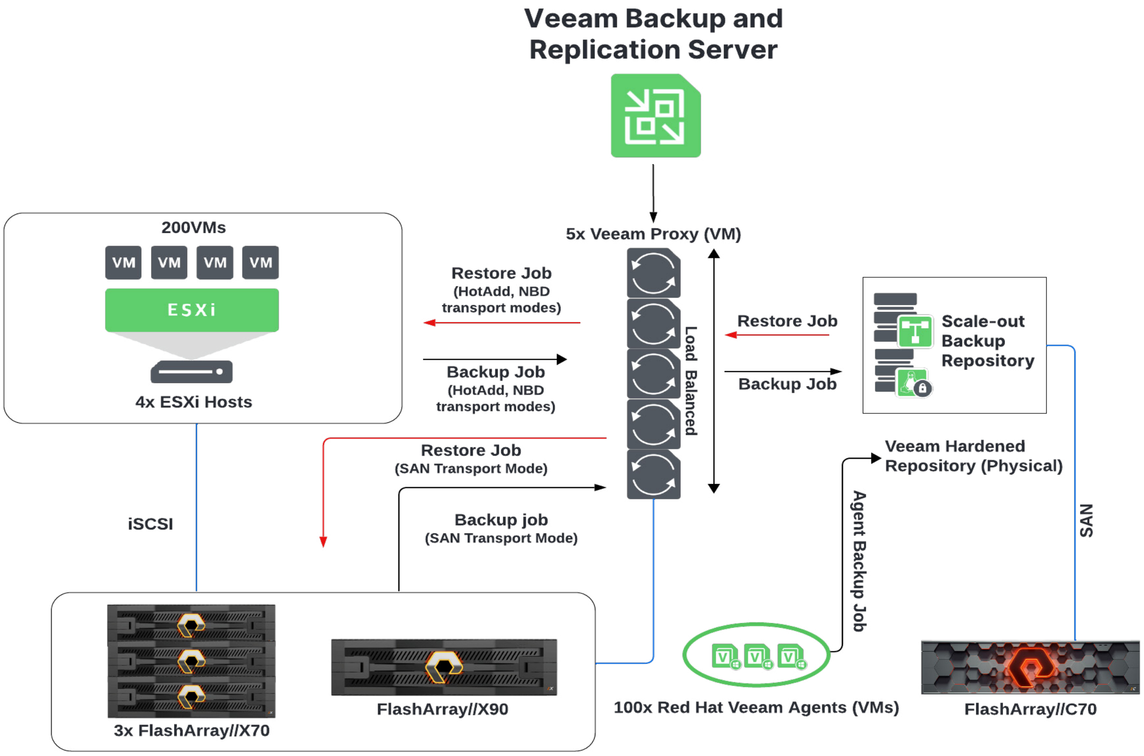
Veeam Backup and Replication Server
5x Veeam Proxy (VM)
Load Balanced
200VMs
VM
VM
VM
VM
ESXi
4x ESXi Hosts
Restore Job (HotAdd, NBD transport modes)
Backup Job (HotAdd, NBD transport modes)
Restore Job
Backup Job
Scale-out Backup Repository
Veeam Hardened Repository (Physical)
Restore Job (SAN Transport Mode)
Backup job (SAN Transport Mode)
iSCSI
Agent Backup Job
SAN
3x FlashArray//X70
FlashArray//X90
100x Red Hat Veeam Agents (VMs)
FlashArray//C70

FIGURE 3 Illustration of the lab environment

Configuration and Optimizations

This section provides details about the setup and optimization used in our lab testing. It involves carefully configuring various components to establish a controlled and effective testing environment, focusing on network configurations, hardware specifications, and software optimizations.

Server

Server Role	CPU	RAM	Networking	Storage	OS
ESXi Host x 4	AMD EPYC 7713P 64-Core Processor, 64CPUs x 2GHz	256GB	Broadcom BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller @ 2×25 Gb/s	Virtual Flash resource - 119.75GB	VMware ESXi, 7.0.3,
Veeam Backup & Replication Server (virtual)	16 vCPUs	28GB	vmxnet3	800GB	Windows Server 2022
Veeam Backup Proxy (virtual) x 5	12 vCPUs	16GB	vmxnet3	100GB	Windows Server 2022
Veeam Hardened Repository server(with DISA STIGS applied)	AMD EPYC 7713P 64-Core Processor, 128 CPUs x 2GHz	256GB	Broadcom BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller @ 2×25 Gb/s Emulex LightPulse Fibre Channel SCSI driver 12.6.0.4, 2-Port 2×32Gb/s	936GB internal storage 1200TB backup storage (FlashArray//C70R4)	Ubuntu Server 20.04

Storage

Storage Role	Array Model	Purity Releases	Physical Storage	Connections
Backup Storage	FlashArray//C70R4	6.6.2	2 × 600TB	12×32 Gb/s FC
Data Source (x3)	FlashArray//X70R3	6.4.5	8 × 50TB	4×25Gb/s iSCSI
Data Source	FalshArray//X90	6.4.10	8 × 50TB	8×32 Gb/s FC

Data Source

Data Source	CPU	RAM	Storage	OS
Linux VMs x 100	2 vCPU	2GB	230GB	Ubuntu 22.04
Microsoft Windows VMs x 100	2 vCPUs	2GB	250GB	Windows 10 Build 19045
Linux VMs with Veeam Agent for Linux x 100	2 vCPUs	2GB	230GB	Red Hat Enterprise Linux 9.1 Veeam Agent for Linux version - 6.1.0.1498

Veeam Backup & Replication Environment Settings

This section details the settings we used during the testing process.

Backup Proxy Settings (VMware VM-based Backup Proxies)

In our environment, Virtual Machine File System (VMFS) datastores are located on the FlashArray//X70 arrays, which are also visible to the VBR Backup Proxy servers on the SAN over iSCSI. We configured the following settings on the Backup Proxies.

- **Connected datastore:** VBR automatically detects datastores to which the VMware Backup Proxy has a direct SAN or NFS connection. For our testing, we used the default automatic detection setting. However, if you prefer the Backup Proxy to work with specific datastores, you can manually configure a list of those datastores.
- **Max concurrent tasks:** Based on available compute resources, the "Max concurrent tasks" setting helps to balance the workload across the backup infrastructure and avoid performance bottlenecks. In our test environment, it was configured to eight for each Backup Proxy.

Backup Repository

For greater flexibility, and performance, Veeam's recommendation for repositories is SOBR. Our lab testing used a SOBR with 2×600TB extents.

- **Limit maximum concurrent tasks:** The "limit maximum concurrent tasks" option restricts the number of concurrent tasks for that repository. Once the limit is reached, VBR will not start a new task until one of the current tasks completes, regardless of limits set on other components. During testing, we have set this option to 16.
- **Use per-machine backup files:** VBR provides an option to save data of each workload into separate backup files. This method of storing backups is more efficient than single-file backup format, especially for operations such as moving backups and launching active full backups for individual workloads. We enabled the "Use per-machine backup files (recommended)" option in the Backup Repository settings.

For more details, please refer to Backup Chain Formats on Veeam help center.

- **Decompress backup file data blocks before storing:** To allow FlashArray to better reduce Veeam backup data, enable the "decompress backup data blocks before storing" option on the Backup Repository. For jobs with data compression enabled, VBR compresses data on the source side, transports it to the target side, decompresses backup data on the target side and then writes uncompressed backup data to the FlashArray. We ran tests with this option enabled.

Backup Job Setup

In our testing, we ran separate jobs for different transport modes with the following settings.

- **Limit processed VM count per snapshot:** By default, when backing up from storage snapshots, VBR creates VMware vSphere snapshots for all VMs defined in the backup job that reside on the same datastore. It then triggers a storage snapshot for the volumes in the datastore. Once the storage snapshot is complete, VBR consolidates and removes the vSphere snapshots. The more VMs residing on the datastore, the longer it takes to create the vSphere snapshots, tracking more changes and increasing consolidation time and impact. Limiting the number of VMs per snapshot shortens the lifetime of vSphere snapshots and lessens the impact. In our tests, we maintained an equal and moderate distribution of VMs across the datastores, so we opted not to enable this setting.
- **Compression level:** Data compression reduces the size of backup files, which can also positively impact backup duration. By compressing data, you decrease network traffic and the disk space needed to store backup files. Compression levels involve a trade-off between CPU usage and compression efficiency: less aggressive compression requires less CPU time.

  In customer environments where data is typically compressible and physical Backup Proxies are used, backup throughput can be significantly higher, amplifying the benefits of Pure Storage FlashArray//C. In our testing, we selected the "Optimal (recommended)" compression level, which provides the best balance between compression efficiency and performance. This level minimizes CPU usage on the Backup Proxy and ensures the fastest restore times.

- **Storage optimization:** To achieve optimal job performance and storage efficiency, VBR allows users to select the data block size for processing virtual machines. The ideal block size depends on the target storage and file sizes. In our testing, we selected a 1MB block size, which is recommended for SAN environments. Additionally, the always-on data reduction feature across all Veeam Backup Repositories on FlashArray//C enhances backup efficiency, driving high storage density in a small footprint and reducing the total cost of ownership (TCO) without compromising performance.
- **Enable backup from storage snapshots:** Having the Pure Storage FlashArray USAPI V2 plugin installed on the VBR server enables VBR to leverage FlashArray snapshots with direct SAN access transport mode. In our testing, we used this setting while running backup jobs with DSAN transport mode.

Validation Results

Here are the results for backup and recovery jobs in our lab testing.

Backup Results

The backup performance observed in our lab clearly underscores the advantages of Pure Storage FlashArray//C technology. This remarkable throughput was achieved with unique, incompressible data. Key findings include:

- As seen in Figure 4, Pure Storage FlashArray//C demonstrates its exceptional value by delivering an impressive active full backup throughput of 4GB/s across all three transport modes, with an average VBR compression level of 1.2x. This equates to completing an active full backup of the entire 220TB data set in a 20-hour backup window.
- Incremental backups achieve an impressive average throughput of 5GB/s, equivalent to 17.5TB per hour, with a daily data change rate of 10%. This equates to an incremental backup window of less than 1.5 hours.Veeam Agent for Linux backup throughput was 4GB/s.
- As shown in Figure 5, synthetic full backups exhibited faster processing time compared to active full backups. Synthetic full backups are important because they create a complete backup from existing incremental and previous full backups, eliminating the need to read and transfer large volumes of data from the source. This reduces the load on production systems, minimizes backup windows, optimizes storage usage, and lowers overall costs while ensuring comprehensive and up-to-date recovery points.

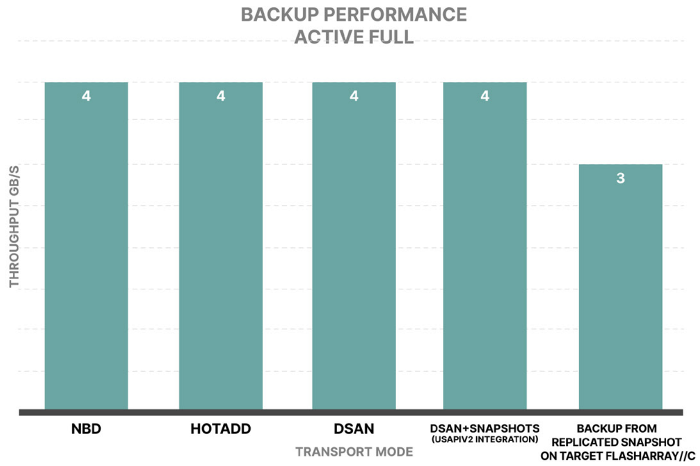

FIGURE 4 Validation results of VM backups: Active full

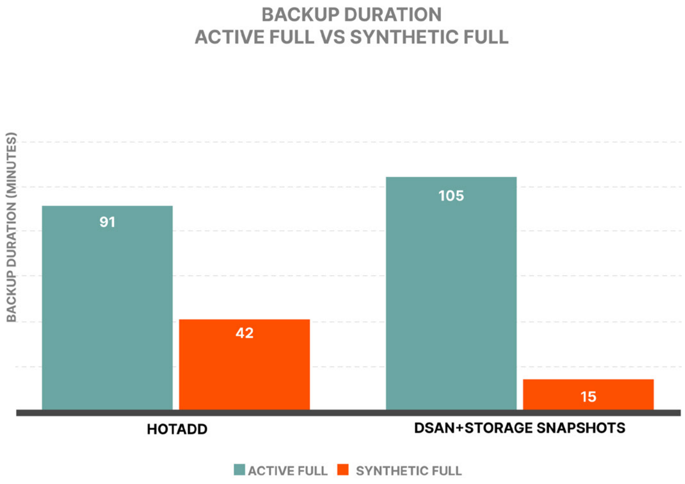

FIGURE 5 Validation results of VM backup duration: Backup duration

Recovery

During the restore process, we conducted multiple full VM restores and achieved impressive throughput of 4.7GB/sec (equivalent to 16.5TB per hour) with HotAdd transport mode for the entire job, as illustrated in Figure 6. In our environment, recovery with HotAdd mode was faster compared to DSAN and NBD modes because it leverages direct datastore access, utilizes high-speed internal data paths, and allows for parallel data processing. This reduces network traffic and overhead, resulting in significantly quicker restores, especially in large-scale environments. DSAN performance was affected by virtual proxies using iSCSI connections and would be different with physical proxies with fibre channel connection.

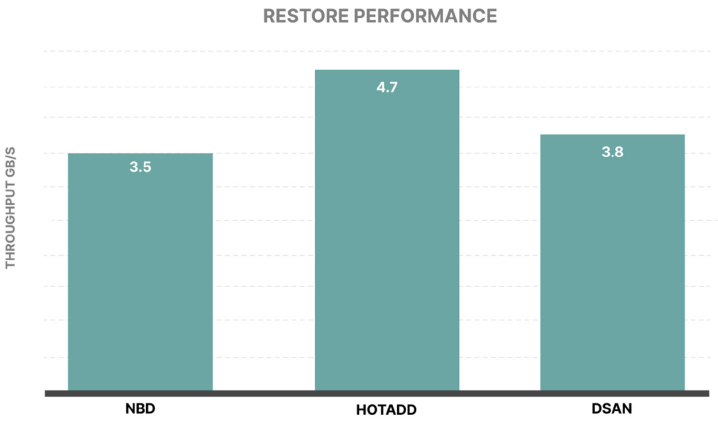

FIGURE 6 Validation results of VM backup duration: Restore performance

Instant Recovery from Storage Snapshots

VBR offers instant Recovery of VMs directly from storage snapshots. This approach significantly speeds up VM recovery, enhances Recovery Time Objectives (RTOs), and minimizes downtime for production VMs. Recovery within FlashArray//X incurs no capacity impact, ensuring excellent same-site performance.

Ransomware Measures

With the surge in malware attacks, particularly ransomware incidents, organizations are facing heightened threats to their data security. The Veeam Data Protection Trends 2024 Report reveals alarming statistics, indicating that 75% of surveyed organizations encountered at least one ransomware attack in the preceding 12 months, with over 50% reporting multiple incidents. In such a landscape, the ability to swiftly restore clean data emerges as a crucial aspect of recovery and operational continuity, especially during a breach. By leveraging VBR with FlashArray//C, you can swiftly restore critical workloads. The all-flash architecture of FlashArray accelerates Instant Recovery, allowing you to run essential services at near-production speeds.

Safe Mode Protection

Deploying a Veeam Hardened Repository into the VBR environment reduces the risk of ransomware attackers destroying backup data. SafeMode protected snapshots provide inherent data protection against ransomware threats by consistently protecting VBR Backup Repository files with immutable and indelible snapshots. These snapshots serve as a protective measure, ensuring the integrity of backup data remains intact, impervious to any unauthorized modifications, deletions, or encryption attempts, even in instances of compromised administrative credentials. Enabling SafeMode initiates three key changes to array behavior for administrative users:

- Disabling the ability to eradicate volumes and snapshots from the destroyed items bucket.
- Introducing an adjustable eradication timer ranging from 24 hours to 30 days. Pure Storage recommends setting this timer between 3 to 7 days for an optimal balance of security and data capacity.
- Disabling the ability to shorten the protection group retention period.

Array wide SafeMode

Array-wide SafeMode is preferred as it offers greater protection. However, this protection is also extended to test volumes, superfluous snapshots, etc, consuming incrementally more capacity. You can always switch to object-level SafeMode at a later time.

Note: While SafeMode prevents manual eradication of protection group snapshots, it does not prevent a pgroup from automatically deleting its own snapshots based on retention policy. You must work with Pure Storage support to enable array-wide SafeMode. Support will onboard you into the authorization process, which will be required for changing SafeMode configuration in the future.

Object-level SafeMode

Some customers require both SafeMode protection for critical data and the flexibility to self-manage snapshots for non-critical data. Available in Purity 6.3.0 and higher, object-level SafeMode protects snapshots of individual protection groups (pgroups) that capture the contents of one or more volumes at the same logical instant. Pgroups are ideal for protecting the snapshots of applications that use multiple volumes for data storage.

- By enabling SafeMode at this granularity, you can balance the criticality of recovering various workloads with the capacity budget.
- Protection group SafeMode can only be ratcheted with pgroups that consist solely of volumes. For Pure Storage, “ratcheted” refers to progressively tightening data protection settings, making them increasingly secure and immutable. Once set, protection levels can only become stricter, never relaxed. Administrators cannot ratchet pgroups that include hosts or host groups, nor can they add hosts or host groups to ratcheted pgroups.
- To use object-level SafeMode, start by creating a pgroup and adding the desired volumes. Next, establish a snapshot schedule to ensure regular backups. Finally, enable the SafeMode retention lock with a ratcheted setting

Array administrators can enable object-level SafeMode on their own, without assistance from Pure Storage Technical Services, but trusted representatives must request assistance from the Pure Storage Technical Services group to unlock them, and follow the SafeMode authorization protocol.

Auto-on SafeMode

Auto-on SafeMode is an extension of object-level SafeMode that automatically creates SafeMode-protected snapshots for new volumes. Here is how it works.

- A ratcheted, default protection group will be automatically configured on the array.
- Newly created volumes will be automatically assigned the default protection group unless explicitly configured otherwise.
- New arrays have auto-on SafeMode enabled by default.

Recovering Data From SafeMode

VBR stores backups on disk using a self-describing file-based approach, including the metadata needed for recovery within the backup files. FlashArray snapshots capture all the data and metadata required for a successful recovery. There are a couple of ways to recover VBR backups using SafeMode.

Option 1: Import Repository

In the case of a new VBR server, there is no Backup Repository configuration to import. Therefore, we create a new Backup Repository on the mapped snapshot volume and then scan the Backup Repository for existing backups. From the VBR console, navigate to the "Backup Infrastructure" section and add a new Backup Repository (Figure 7). Follow the wizard to create a Backup Repository as normal.

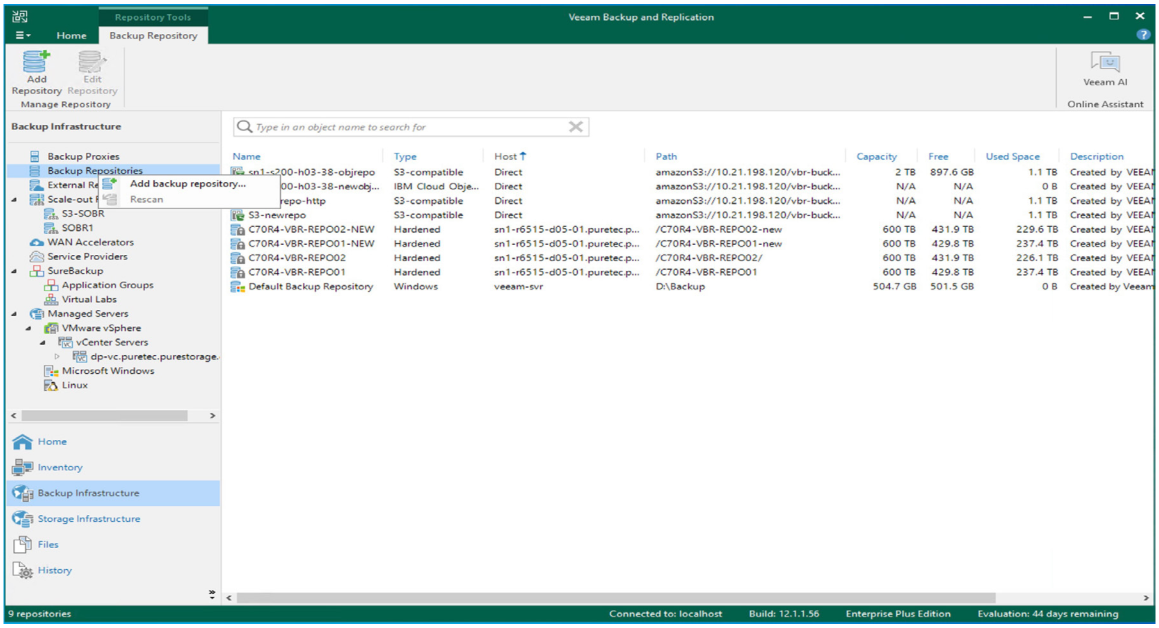

Name	Type	Host	Path	Capacity	Free	Used Space	Description
sn1-s200-h03-38-objrepo	S3-compatible	Direct	amazonS3://10.21.198.120/vbr-buck...	2 TB	897.6 GB	1.1 TB	Created by VEEA
00-h03-38-newobj...	IBM Cloud Obje...	Direct	amazonS3://10.21.198.120/vbr-buck...	N/A	N/A	0 B	Created by VEEA
repo-http	S3-compatible	Direct	amazonS3://10.21.198.120/vbr-buck...	N/A	N/A	1.1 TB	Created by VEEA
S3-newrepo	S3-compatible	Direct	amazonS3://10.21.198.120/vbr-buck...	N/A	N/A	1.1 TB	Created by VEEA
C70R4-VBR-REPO02-NEW	Hardened	sn1-r6515-d05-01.puretec.p...	/C70R4-VBR-REPO02-new	600 TB	431.9 TB	229.6 TB	Created by VEEA
C70R4-VBR-REPO01-NEW	Hardened	sn1-r6515-d05-01.puretec.p...	/C70R4-VBR-REPO01-new	600 TB	429.8 TB	237.4 TB	Created by VEEA
C70R4-VBR-REPO02	Hardened	sn1-r6515-d05-01.puretec.p...	/C70R4-VBR-REPO02/	600 TB	431.9 TB	226.1 TB	Created by VEEA
C70R4-VBR-REPO01	Hardened	sn1-r6515-d05-01.puretec.p...	/C70R4-VBR-REPO01	600 TB	429.8 TB	237.4 TB	Created by VEEA
Default Backup Repository	Windows	veeam-svr	D:\Backup	504.7 GB	501.5 GB	0 B	Created by Veeam

FIGURE 7 Veeam Backup and Repository: Add Backup Repository

When you reach the “Review” screen of the wizard, it’s important to check the option “Search repository for existing backups and import them automatically” checkbox (Figure 8).

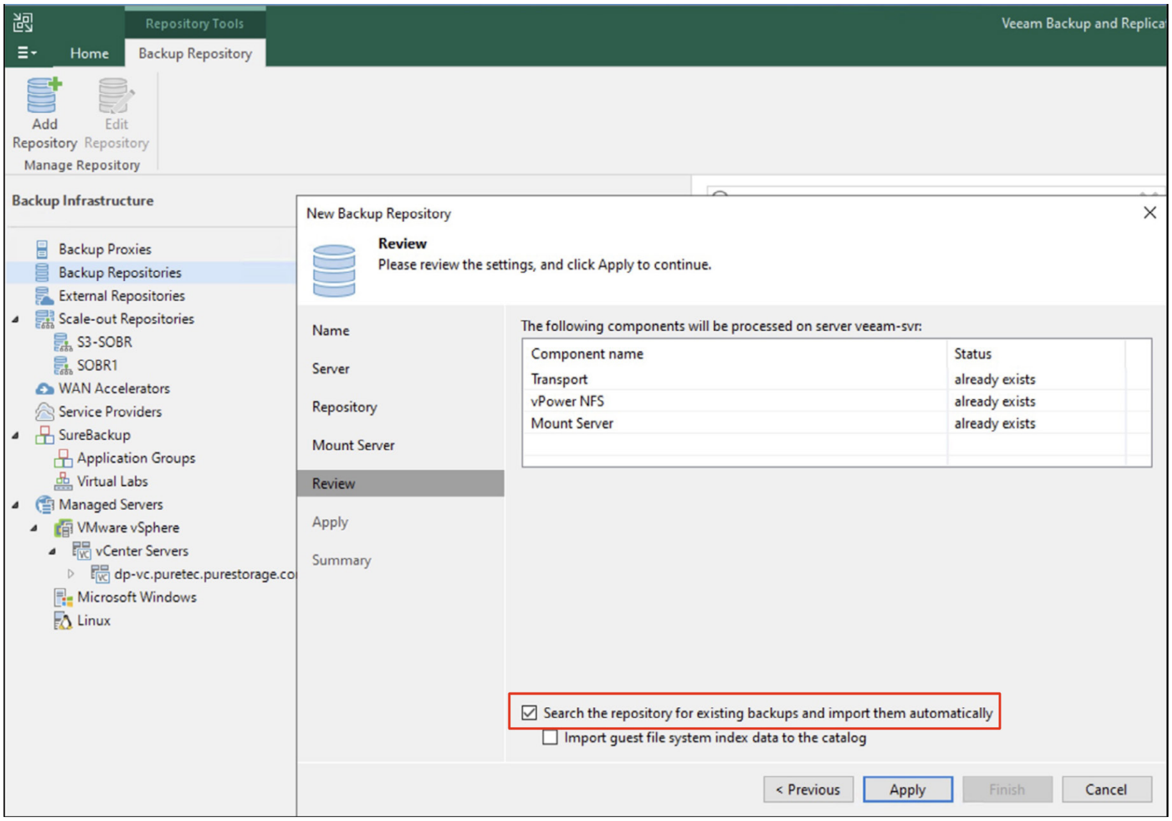

New Backup Repository

Review

Please review the settings, and click Apply to continue.

The following components will be processed on server veeam-svr:

Component name	Status
Transport	already exists
vPower NFS	already exists
Mount Server	already exists

☑ Search the repository for existing backups and import them automatically

☐ Import guest file system index data to the catalog

FIGURE 8 Add Backup Repository: Review

After the wizard is complete, the new Backup Repository is added to the backup infrastructure.

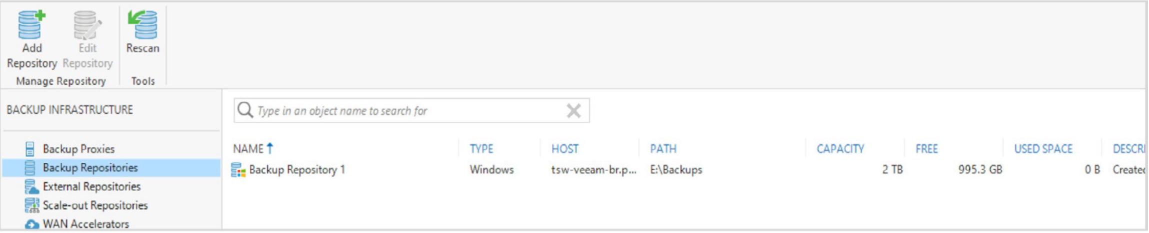

NAME ↑	TYPE	HOST	PATH	CAPACITY	FREE	USED SPACE	DESCRI
Backup Repository 1	Windows	tsw-veeam-br.p...	E:\Backups	2 TB	995.3 GB	0 B	Createc

FIGURE 9 Backup Repository wizard

Option 2: Import Backups and Secure Restore

Since scanning the Backup Repository volume for viruses can be time-consuming, you may initially need to prioritize restoring a few critical services, such as DNS, Active Directory, and email, to ensure continuity. For this purpose, we can import individual backup jobs containing the necessary VMs or physical server backups. Subsequently, we can initiate the restore process for these systems while ensuring that the antivirus scan remains enabled, allowing restoration to occur concurrently with scanning. This functionality is referred to as a Secure Restore. Access it through the Veeam management interface under the "Home" section and then the "Import Backup" action.

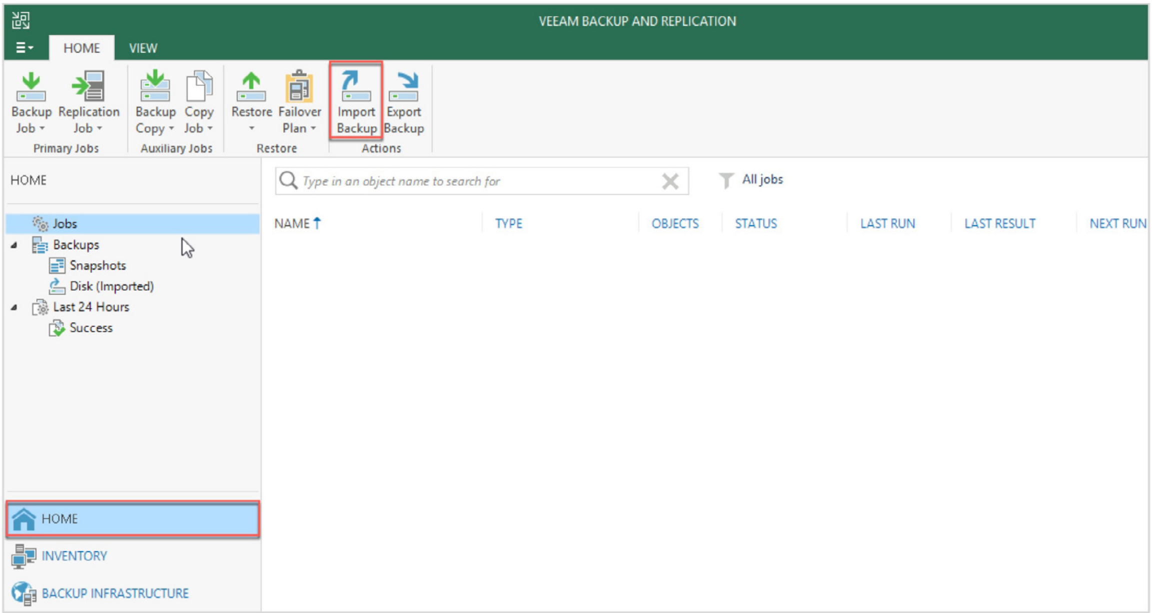

FIGURE 10 Import backup

You will be prompted to specify which backup to import. Since this option allows you to point directly to the backup job, you may need to repeat this process for each backup job (Figure 11).

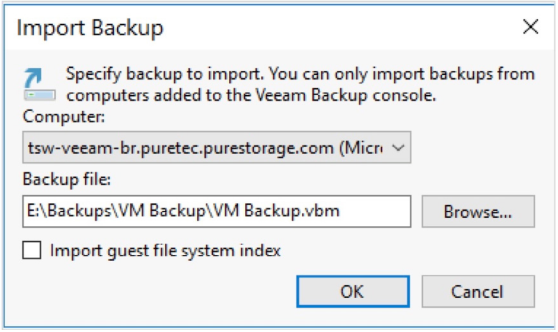

Import Backup

Specify backup to import. You can only import backups from computers added to the Veeam Backup console.

Computer: tsw-veeam-br.puretec.purestorage.com (Micr

Backup file: E:\Backups\VM Backup\VM Backup.vbm

Browse...

☐ Import guest file system index

OK Cancel

FIGURE 11 Choosing a backup file

Once Veeam scans the disk, a new shortcut for the recovered imported backups is displayed.

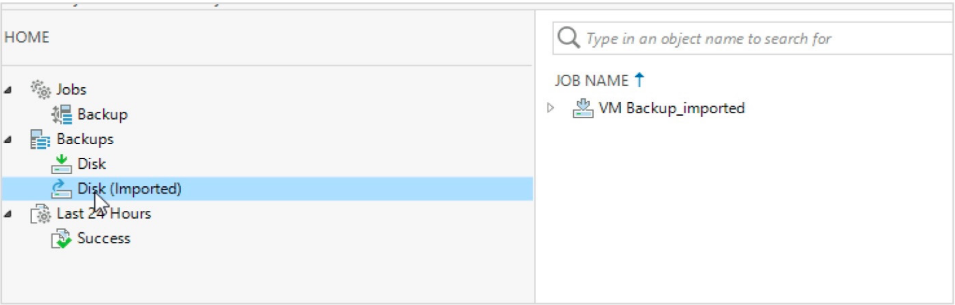

FIGURE 12 Imported backups

Run the antivirus scanning tools during the restoration process. If this is a VM, you can disable network adapters for additional checks if necessary.

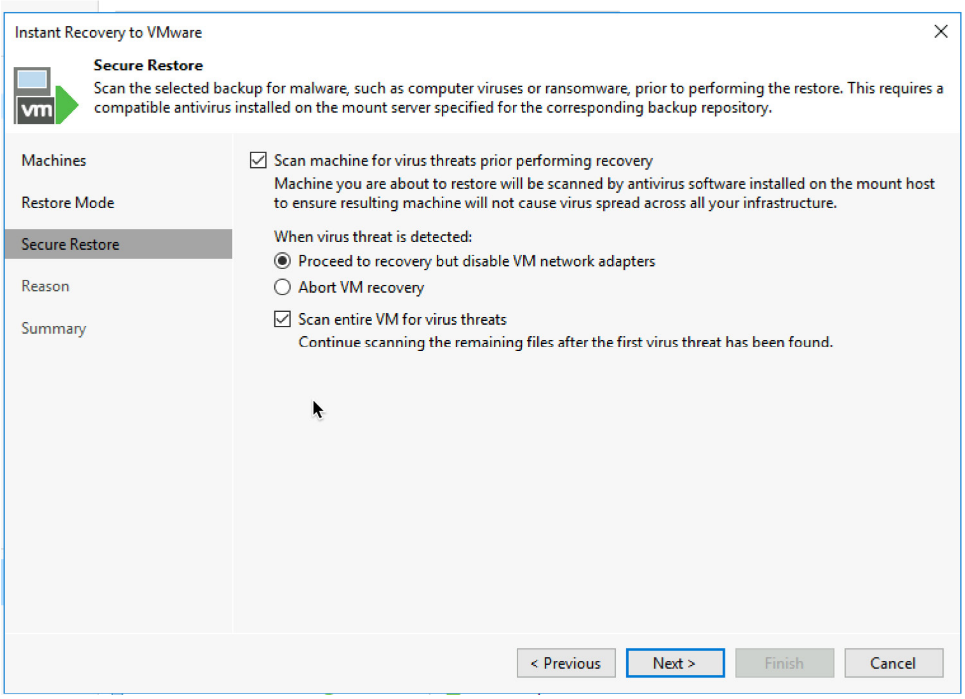

Secure Restore

Scan the selected backup for malware, such as computer viruses or ransomware, prior to performing the restore. This requires a compatible antivirus installed on the mount server specified for the corresponding backup repository.

☑ Scan machine for virus threats prior performing recovery

Machine you are about to restore will be scanned by antivirus software installed on the mount host to ensure resulting machine will not cause virus spread across all your infrastructure.

When virus threat is detected:

● Proceed to recovery but disable VM network adapters

○ Abort VM recovery

☑ Scan entire VM for virus threats

Continue scanning the remaining files after the first virus threat has been found.

FIGURE 13 Scan for virus

At its core, SafeMode is focused on privileged user attack protection. It's not just about snapshots, object locks, or any specific feature; it is a comprehensive governance mechanism that encompasses these elements to provide robust protection. SafeMode safeguards against the worst-case scenarios posed by threat actors, ensuring the security and integrity of your data.

For more detailed information on recovery procedure from SafeMode snapshots, please refer to Ransomware Remediation Veeam.

References

- Veeam Help Center
- Veeam Backup & Replication Best Practices
- Veeam VMware vSphere backup proxy.
- User Guide on Creating Backup Repositories.
- Ransomware Remediation with Veeam White Paper
- Pure Storage FlashArray//C

purestorage.com

800.379.PURE

PURESTORAGE®
Uncomplicate Data Storage, Forever