

Reference Architecture



Reference Architecture: Advanced PureStorage FlashArray Data Protection

Using Veeam Backup & Replication and

ExaGrid Backup Storage

October 2016

Contents

Executive Summary	3
Goals and Objectives	3
Audience	3
Pure Storage Overview	5
Purity Storage FlashArray //m.....	5
PureStorage FlashArray //m Specifications	5
Veeam Overview	6
Availability for the Always-On Enterprise	6
High-Speed Recovery	6
Data Loss Avoidance.....	7
Verified Recoverability.....	7
Leveraged Data.....	7
Complete Visibility	8
ExaGrid Overview	8
Solving the Problems of Backup Permanently	8
Unique Architecture	8
Adaptive Deduplication.....	10
Integration with Veeam Backup & Recovery	10
System Requirements	10
Harnessing Veeam and ExaGrid Data Protection for PureStorage FlashArray	11
Use Case 1: Scaling to a large numbers of concurrent, application aware, VM-consistent backups.	12
Typical scenario: Deployment of a new Enterprise Application	12

Use Case 2: Efficient long-term backup retention.....	19
Use Case 3: Easily & regularly prove backups are recoverable	21
Use Case 4: Instantly Recovering a Virtual Machine.....	25
Use Case 5: Recovering granular, application-specific items directly from backups.....	28
Deployment.....	32
PureStorage FlashArray & Hypervisor Configuration	33
ExaGrid Deployment.....	40
Details on Veeam and ExaGrid Configuration	43
Creating ExaGrid shares and Veeam Repositories.....	43
Creating Veeam Backup Jobs.....	47
Summary	49
References	50
About the Author	50

Executive Summary

Organizations are increasingly being pressed to get more from their IT dollars. They have growing data and decreasing budgets. In parallel they are being asked to increase productivity by having faster servers, faster storage, and faster restore, VM boot, and recovery times. The challenge of doing more with less is driving a move to the next generation of storage and backup solutions that increase performance and functionality while decreasing costs. Pure Storage, Veeam, and ExaGrid have formed an alliance around the goal of providing the modern data center with more performance and functionality at a lower cost.

Goals and Objectives

Provide technical information about Pure Storage, Veeam Backup & Replication, and ExaGrid backup storage appliances:

1. Five example use cases illustrating the value of combining all three products.
2. Describe configurations
3. Provide design guidance

Together, this information will enable readers to understand the key benefits of the joint solution and how it delivers “Availability for the Always-On Enterprise”.

Audience

This paper is written for VMware, storage or backup administrators interested in understanding how the combination of Pure Storage Flash, Veeam Backup & Replication, and ExaGrid backup storage appliances combine fast, reliable and cost-effective primary storage along with the fast, low-cost backups to shorten recovery times.

© 2016 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. ESXi, vCenter and vSphere are registered trademarks of VMware in the U.S. and other countries. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc. 650 Castro Street, Mountain View, CA 94041
<http://www.purestorage.com>

Pure Storage Overview

Pure Storage is the leading all-flash enterprise array vendor, committed to enabling companies of all sizes to transform their businesses with flash.


Built on 100% MLC flash, Pure Storage FlashArray //m delivers all-flash enterprise storage that is 10X faster, more space and power efficient, more reliable, and infinitely simpler, and yet typically cost less than traditional performance disk arrays.

Purity Storage FlashArray //m

Who knew that moving to all-flash storage could help reduce the cost of IT? PureStorage FlashArray //m makes server and workload investments more productive, while also lowering storage spend. With PureStorage FlashArray //m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.

PureStorage FlashArray //m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply weren't possible with disk. All by Transforming Your Storage with PureStorage FlashArray //m.

PureStorage FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.



PureStorage FlashArray //m Specifications

Check out the entire FlashArray //m specifications in the below link.

<https://www.purestorage.com/products/flash-array-m/hardware-tech-spec-flash-array.html>

Veeam Overview

Veeam® Backup & Replication™ provides fast, flexible, and reliable recovery of virtualized applications and data. Veeam brings backup and replication together in a single solution with award-winning support for VMware vSphere and Microsoft Hyper-V virtual environments. This powerful, easy-to-use and affordable backup and Availability solution leverages virtualization, storage, and cloud technologies to meet SLAs for recovery time and point objectives (RTPO™).

The demands on today's enterprises — the need to access data and applications 24/7, no patience for downtime or data loss and exponential data growth at 30-50% per year — are not being met by traditional legacy backup tools. In fact, 82% of CIOs report a gap between the level of availability legacy backup solutions provide and what end users demand.

Availability for the Always-On Enterprise

Veeam® bridges this gap by providing customers with a new kind of solution, an Availability solution, delivering recovery time objectives (RTO) and recovery point objectives (RPO), or RTPO® of < 15 minutes for ALL applications and data.

High-Speed Recovery

Rapid recovery of what you want, the way you want it

Veeam Backup & Replication delivers lightning-fast, reliable restore at your fingertips for entire VMs, individual files and application items, giving you low recovery time objectives (RTOs) of < 15 minutes.

- Restore entire VMs in minutes with Instant VM Recovery™
- Recover individual files effortlessly with Instant File-level Recovery
- Quickly and easily restore Microsoft Exchange items with Veeam Explorer™ for Microsoft Exchange
- Recover Active Directory objects, entire containers, OUs and user accounts with Veeam Explorer for Microsoft Active Directory
- Restore entire SharePoint sites; return deleted items back into production with Veeam Explorer for Microsoft SharePoint
- Fast transaction-level recovery and point-in-time restore of Oracle databases with Veeam Explorer for Oracle and SQL Server databases with Veeam Explorer for Microsoft SQL Server

Data Loss Avoidance

Low RPOs and streamlined disaster recovery

Veeam Backup & Replication provides streamlined disaster recovery and simple, secure offsite backups, giving you the ability to achieve recovery point objectives (RPO) of < 15 minutes.

- Create backups from any storage
- 2-in-1: backup and replication™: In addition to backups, maintain image-based replicas either onsite for high-availability or offsite for disaster recovery, and simplify replica failover and failback with little to no business disruption
- Leverage fully-integrated cloud-based disaster recovery and fast and secure cloud backup with Veeam Cloud Connect
- Simplify the management of your backup storage with Scale-out Backup Repository™
- Get backups and replicas off site up to 50x faster with built-in WAN acceleration
- Meet compliance requirements and archiving policies with native tape support
- Back up directly from file-based (NFS) primary storage using Veeam's proprietary NFS client with Direct NFS Access

Verified Recoverability

Guaranteed recovery of every file, application, or virtual server, any time

Veeam Backup & Replication automatically tests every vSphere and Hyper-V backup and replica, every time. Have confidence that your files, applications and virtual servers can be restored when needed.

- Verify the recoverability of every backup with SureBackup®.
- Test every restore point in every VM replica automatically with SureReplica.

Leveraged Data

Using backup data to create an exact copy of your production environment

Veeam Backup & Replication mitigates the risks associated with application deployment by putting your backups and replicas to work in a production-like environment prior to any production rollout.

- Create a Virtual Lab without provisioning additional resources to run your applications from Veeam backups and replicas.
- Test or troubleshoot an application in one or more VMs from an On-Demand Sandbox™ without affecting production workloads.
- Create complete, isolated copies of production from storage snapshots with On-Demand Sandbox for Storage Snapshots

Complete Visibility

Proactive monitoring and alerting of issues before operational impact

With Veeam Backup & Replication and Veeam ONE™, you can take control over your entire backup and virtual environment. By utilizing all of the features and capabilities of Veeam's backup and virtualization technology, you can discover and be alerted to issues before they affect your Always-On Business™.

- Ensure protection, performance and availability with Veeam ONE through:
 - 24x7 real-time monitoring and alerting: Notifies you of backup and performance issues
 - Resource optimization and configuration tracking: Evaluate infrastructure performance and ensure existing configuration meets all known best practices
 - Capacity planning and forecasting: Forecast resource usage and utilization trends with “what if” modeling and resource overcommit tracking for your backup and virtual infrastructures
 - And MORE included in the NEW Veeam ONE v9!
- Gain visibility into the Veeam backup infrastructure from inside Microsoft System Center, Kaseya VSA and LabTech.
- Get enterprise-level usability with a standalone console and Remote Office / Branch Office (ROBO) enhancements
- Scale your business with a distributed architecture and centralized management console.
- Leverage deep VMware integrations with vCloud Director support and a vSphere Web Client Plug-in.

ExaGrid Overview

Solving the Problems of Backup Permanently

ExaGrid's disk-based backup systems help IT organizations solve two of the most pressing issues they face today: how to protect and manage growing data and how to do so at a lower cost.

ExaGrid does so by ensuring that its solutions address four fundamental requirements:

- provide the shortest possible backup window;
- ensure a consistent backup window despite increasing data volumes;
- make it quick and easy to restore or recover data; and
- fix backup permanently.

Unique Architecture

ExaGrid's award-winning scale-out architecture provides customers with a consistent backup window regardless of data growth. Its unique zone-level approach to data deduplication retains the most recent backup in its full unduplicated form, enabling the fastest restores, offsite tape copies, and instant recoveries.

ExaGrid's multiple appliance models can be combined into a "GRID" configuration of up to 2.4PB raw capacity, allowing full backups of up to 1PB with a combined ingest rate of 200 TB/hr. The appliances virtualize into one another when plugged into a switch so that multiple appliance models can be mixed and matched into a single configuration. Each appliance includes the appropriate amount of processor, memory, disk, and bandwidth for the data size, so as each appliance is virtualized into the GRID, performance is maintained and backup times do not increase as data is added. Once virtualized, they appear as a single pool of long-term capacity. Capacity load balancing of all data across servers is automatic, and multiple GRID systems can be combined for additional capacity. Even though data is load balanced, deduplication occurs across the systems so that data migration does not cause a loss of effectiveness in deduplication.

This combination of capabilities in a turnkey appliance makes the ExaGrid system easy to install, manage, and scale. ExaGrid's architecture provides lifetime value and investment protection that no other architecture can match.

Adaptive Deduplication

ExaGrid writes backups directly to a disk landing zone, avoiding inline processing, ensuring the highest possible backup performance resulting in the shortest backup window. “Adaptive” deduplication performs deduplication and replication in parallel with backups while providing full system resources to the backups for the shortest backup window. Available system cycles are utilized to perform deduplication and offsite replication for an optimal recovery point at the disaster recovery site. Once complete, the onsite data is protected and immediately available in its full undeduplicated form for fast restores, VM Instant Recoveries and tape copies while the offsite data is ready for disaster recovery.

Integration with Veeam Backup & Recovery

ExaGrid is the only deduplication appliance that runs the ExaGrid-Veeam Accelerated Data Mover. This software component is Veeam software enabled to run on the ExaGrid appliance, providing:

- 30% faster backup and restore performance compared with CIFS transport
- 6x reduction in duration of Veeam synthetic full backups
- Acceleration of Veeam Backup Copy jobs used for extended retention

ExaGrid and Veeam can instantly recover a VMware virtual machine by running it directly from the ExaGrid appliance in the event that the primary storage VM becomes unavailable. This is possible because of ExaGrid’s “landing zone” – a high-speed cache on the ExaGrid appliance that retains the most recent backups in complete form. Once the primary storage environment has been brought back to a working state, the VM running on the ExaGrid appliance can then be migrated to primary storage for continued operation.

System Requirements

The following requirements exist in order to use Veeam and ExaGrid technology integration with the Pure Storage FlashArray:

- FlashArray 400 series or FlashArray//m
- Purity v4.1.1 (or higher)
- Veeam Backup & Replication v9.5 or later – Enterprise or Enterprise Plus version
- Optionally, Veeam Backup Enterprise Manager
- ExaGrid v4.8.0 or later

Harnessing Veeam and ExaGrid Data Protection for PureStorage FlashArray

Veeam Backup & Replication data protection software plus ExaGrid's backup storage offers an extensive feature set with rich controls and behaviors for protecting Modern Data Centers that are using Pure Storage FlashArrays. The following section will outline the main features of this combination through a set of use cases.

The following functions are typically required by Modern Data Centers:

1. Large numbers of concurrent, application aware, VM-consistent backups
2. Efficient long-term backup retention
3. Easily and regularly prove backups are recoverable
4. Instantly Recover a virtual machine
5. Recover granular, application-specific items directly from backups

The subsequent use cases are not an exhaustive description of all functions and features but rather demonstrate some of the most common ones needed in today's Modern Data Centers.

Use Case 1: Scaling to a large numbers of concurrent, application aware, VM-consistent backups

Modern data centers must be able to scale to meet customer, business and infrastructure demands. Virtualization has enabled an unprecedented level of scale which in turn demands a scale-out architecture across all key infrastructure components:

Scalable Primary Storage

Pure's Evergreen Architecture was designed to fully support the scale out architecture for future hardware. You can flexibly upgrade controller modules or add capacity. You can even seamlessly scale out by mixing different sized flash modules within the same array. Only purchasing what you need, at the best available density, avoids waste and reduces upgrade costs as well as power, cooling and datacenter footprint consumption.

Scalable Data Protection

Veeam Backup & Replication scales out by instantiating sufficient Veeam Backup Proxy Servers in the right infrastructure locations to ensure that application data can be moved to the backup storage at sufficient rates and levels of concurrency to meet demanding backup windows.

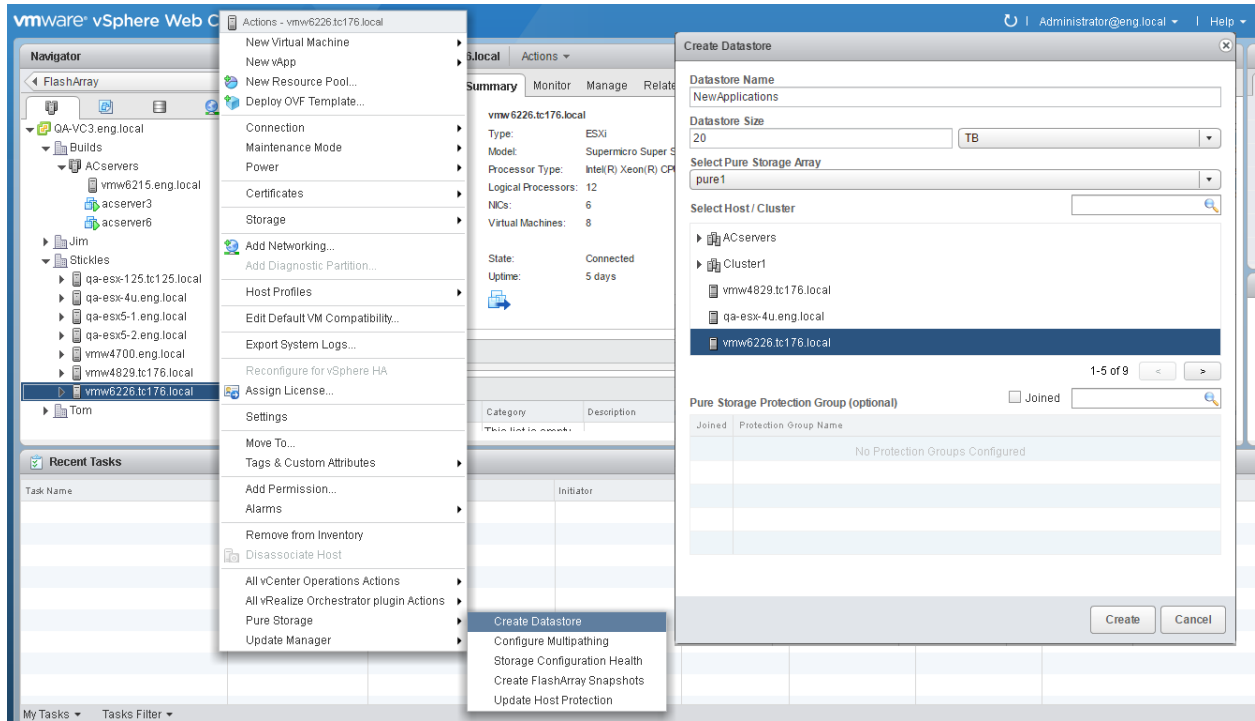
Scalable Backup Storage

ExaGrid scales out by allowing up to 25 appliances in a single GRID, each with its own CPU, memory, network bandwidth and storage. As the volume of backup data grows, appliances are added to the GRID, with the new appliance(s) resources used to keep the backup window constant.

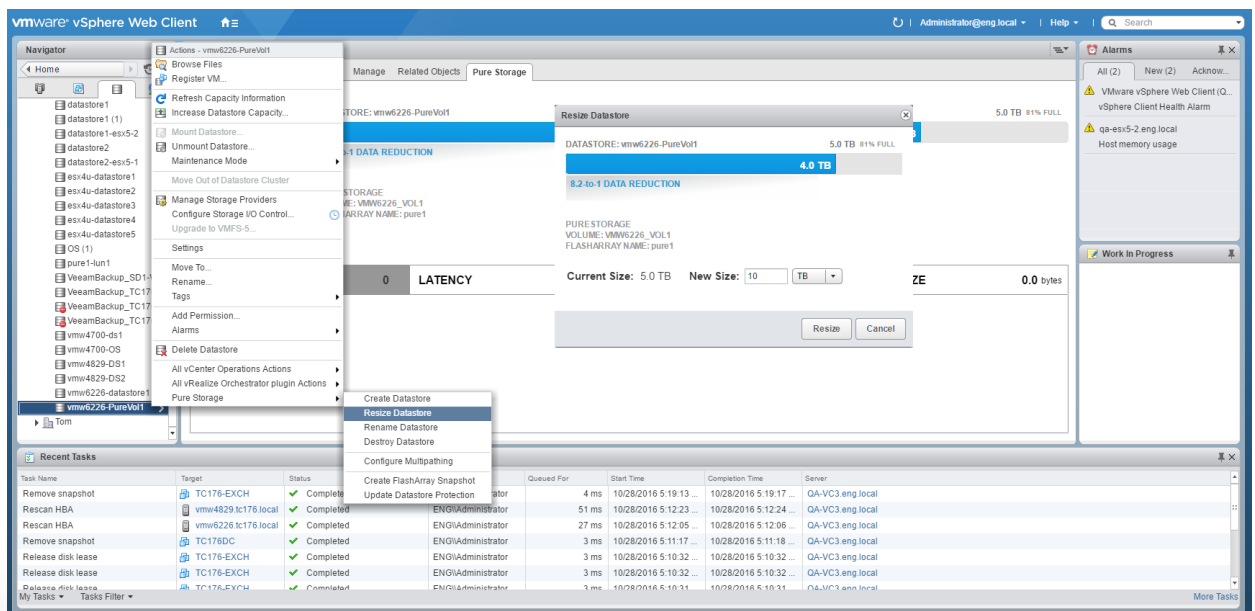
Typical scenario: Deployment of a new Enterprise Application

1. As part of deploying a new enterprise application, a new VMware ESXi host is physically connected to the PureStorage FlashArray //m. New volumes can be created on this PureStorage FlashArray //m which will be seen by all virtual machines (VMs) connected to that VMware ESXi host or VMware ESXi host groups. This host/host groups will be capable of supporting multiple VMs.

2. Datastores can be created on PureStorage FlashArrays in one step on a specific host or host group (clusters) using PureStorage web client plugin for the vsphere web client as shown below.

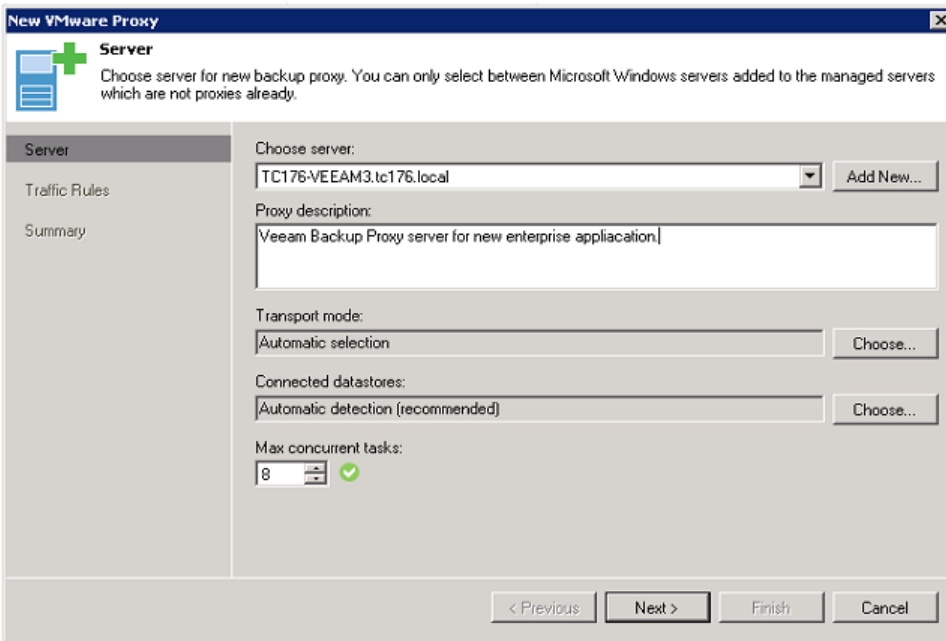


3. The user can resize and gauge the performance of the datastore from within the vsphere web client using PureStorage web client plugin. Datastores on PureStorage FlashArray //m can also be backed up for shorter-term retention by creating array based snapshots on PureStorage FlashArrays. Resizing a datastore can be accomplished in one step using PureStorage web client plugin (as shown below) and it is fast, efficient occupying very little space using PureStorage Data Reduction technology.



4. Because Veeam is already configured with the existing enterprise vCenter, Veeam configuration is limited to modifying an existing Veeam job to add the new VMs. A new, separate Veeam job could be created for the new VMs if the job's schedule or retention settings were different from existing jobs.
5. A new Veeam proxy server is provisioned on the new VMware host. The existing Veeam Backup & Recovery server continues to do "command and control", the new Veeam proxy server moves the backups (and restores) from the new FlashArray //m data stores to the ExaGrid site.
6. Because of the size and retention of this new enterprise application, an additional ExaGrid EX32000E with backup storage capacity for 32TB of weekly full backup is purchased and added to the existing ExaGrid site.

Provision a new Veeam proxy server on the new VMware host. Note Veeam will automatically choose VMware transport mode used during the movement of backup data, or, it can be constrained based on virtualized infrastructure. Concurrent tasks limits depend on CPU and memory configuration of the VM in which the proxy runs.



Create new ExaGrid share, Veeam repository, and add it to Veeam scale-out repository. See **Creating ExaGrid shares and Veeam Repositories** above for a description of the steps to create a Veeam share on the new ExaGrid appliance and adding it to the existing Veeam scale-out repository.

Summary for Site : TC176S1

Landing Space for Site : TC176S1

Server	Actual	Available		Deduplication Status
		% Available for Next Backup	Maximum	
TC176S1	96,000.00 GB	100%	96,000.00 GB	Complete
TC176SP1OB1	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB2	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB3	32,000.00 GB	100%	32,000.00 GB	Complete

Retention Space for Site : TC176S1

Server	Available	Consumed	Shared	Retention Space		Total Capacity	Over Capacity	Replication Status
				% Available				
TC176S1	86,619.29 GB	5,776.33 GB	5%	94%		92,395.62 GB	0.00 GB	Complete
TC176SP1OB1	29,081.24 GB	1,717.30 GB	3%	94%		30,798.54 GB	0.00 GB	Complete
TC176SP1OB2	28,251.80 GB	2,546.74 GB	3%	92%		30,798.54 GB	0.00 GB	Complete
TC176SP1OB3	29,286.25 GB	1,512.29 GB	9%	95%				

Deduplication Summary for Site : TC176S1

Share	Deduplication Ratio	Total Backup Data	Compressed Size	Distinct Size
TC176S1	✓ 14.96 : 1	86,421.03 GB	5,776.33 GB	5,497.55 GB
TC176SP1OB1				
TC176SP1OB1EXT1	✓ 13.91 : 1	26,083.82 GB	1,875.62 GB	1,660.80 GB
TC176SP1OB2				
TC176SP1OB2EXT1	✓ 12.52 : 1	34,068.11 GB	2,721.08 GB	2,459.88 GB
TC176SP1OB3				
TC176SP1OB3EXT1	✓ 16.6 : 1	26,269.09 GB	1,582.79 GB	1,376.87 GB

ExaGrid appliance 1 backup storage

ExaGrid appliance 2 backup storage

ExaGrid appliance 3 backup storage

Reports Manage Logout Help

Summary for Site : TC176S1 Report Generated: 7

Landing Space for Site : TC176S1

Server	Actual	Available		Deduplication Status
		% Available for Next Backup	Maximum	
TC176S1	96,000.00 GB	100%	96,000.00 GB	Complete
TC176SP1OB1	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB2	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB3	32,000.00 GB	100%	32,000.00 GB	Complete

Retention Space for Site : TC176S1

Server	Retention Space				Total Capacity	Over Capacity	Replication Status
	Available	Consumed	Shared	% Available			
TC176S1	86,619.29 GB	5,776.33 GB	5%	94%	92,395.62 GB	0.00 GB	Complete
TC176SP1OB1	29,081.24 GB	1,717.30 GB	3%	94%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB2	28,251.80 GB	2,546.74 GB	3%	92%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB3	29,286.25 GB	1,512.29 GB	9%	95%	30,798.54 GB	0.00 GB	Complete

Deduplication Summary for Site : TC176S1

Share	Deduplication Ratio	Total Backup Data	Compressed Size	Distinct Size
TC176S1	✓ 14.96 : 1	86,421.03 GB	5,776.33 GB	5,497.55 GB
TC176SP1OB1				
TC176SP1OB1EXT1	✓ 13.91 : 1	26,083.82 GB	1,875.62 GB	1,660.80 GB
TC176SP1OB2				
TC176SP1OB2EXT1	✓ 12.52 : 1	34,068.11 GB	2,721.08 GB	2,459.88 GB
TC176SP1OB3				
TC176SP1OB3EXT1	✓ 16.6 : 1	26,269.09 GB	1,582.79 GB	1,376.87 GB

Edit Scale-out Backup Repository

Extents
Add two or more extents (regular backup repositories) to this scale-out repository. Note that you will not be able to use added extents as standalone repositories until you remove them from this scale-out repository.

Name	Extents:												
Extents	<table border="1"> <thead> <tr> <th>Name</th> <th>Add...</th> <th>remove</th> </tr> </thead> <tbody> <tr> <td>TC176SP1OB1EXT1</td> <td></td> <td></td> </tr> <tr> <td>TC176SP1OB2EXT1</td> <td></td> <td></td> </tr> <tr> <td>TC176SP1OB3EXT1</td> <td></td> <td></td> </tr> </tbody> </table>	Name	Add...	remove	TC176SP1OB1EXT1			TC176SP1OB2EXT1			TC176SP1OB3EXT1		
Name	Add...	remove											
TC176SP1OB1EXT1													
TC176SP1OB2EXT1													
TC176SP1OB3EXT1													
Policy													
Summary													

ExaGrid appliance 1 backup storage

ExaGrid appliance 2 backup storage

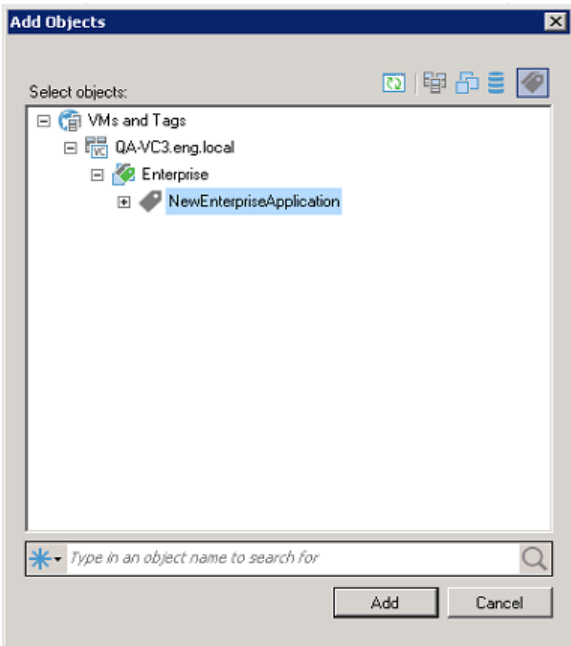
ExaGrid appliance 3 backup storage

Veeam chooses the extent – i.e. ExaGrid appliance – to which backups from Pure Storage-backed VMs are sent – scale out from end to end.

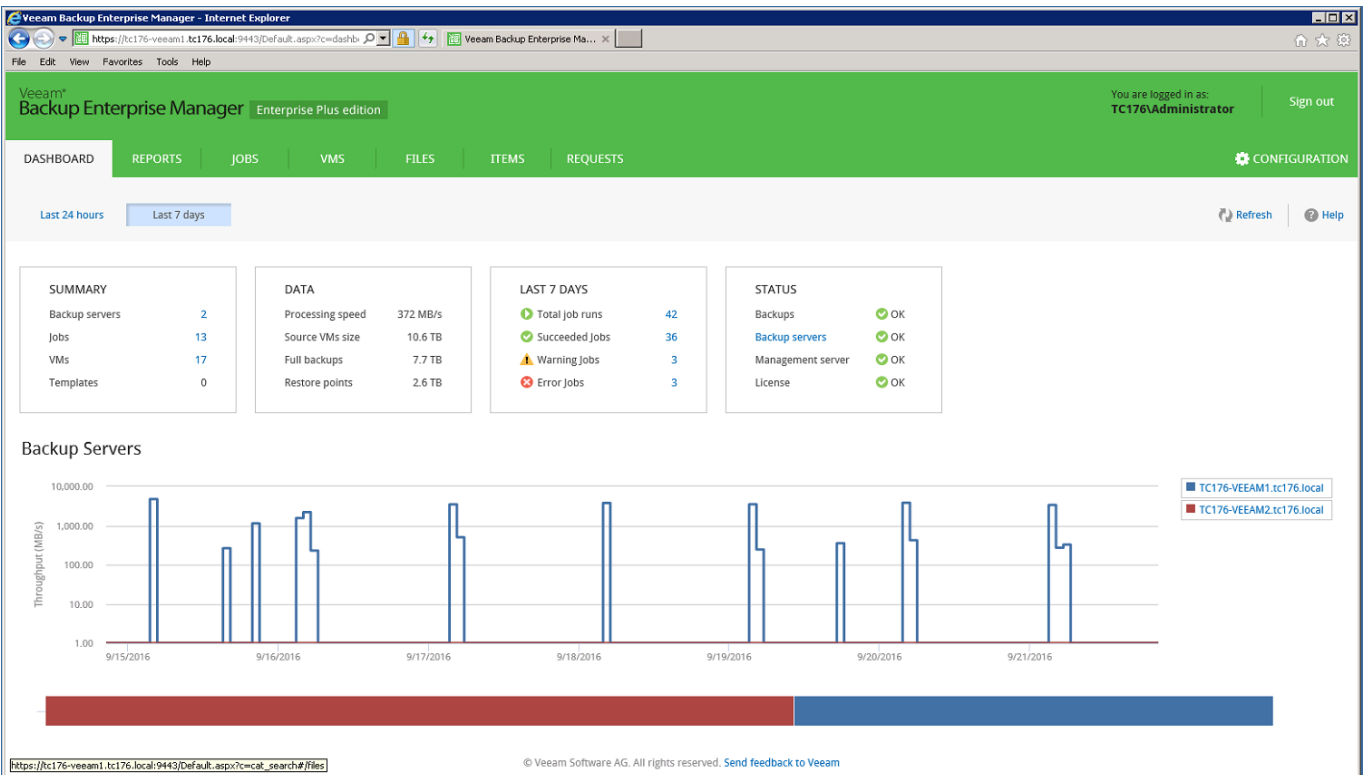
Click Advanced to modify additional scale-out backup repository options. Advanced

< Previous Next > Finish Cancel

Add the new enterprise application VMs to an existing Veeam backup job. Because all the new VMs are tagged in VMware, they can be easily added to an existing (or new) Veeam backup job:



Once these backups are running, further monitoring and reporting is available in Veeam's Backup Enterprise Manager:



And ExaGrid's web-based GUI:

Deduplication Completed Details for Site : TC176S1

Display backup jobs that finished writing on:

- ExaGrid
 - Site: TC176H1
 - TC176H1OB1: EX32000E
 - Site: TC176S1
 - TC176SP1OB1: EX32000E
 - TC176SP1OB2: EX32000E
 - TC176SP1OB3: EX32000E

Name	Size	Completion Time	Completion Time
Site: TC176S1			
Server: TC176SP1OB1			
Share: TC176SP1OB1SC1 (Veeam Backup & Replication™)			
TC176SP1OB1SC1 files from 25 Sep 2016 - 1 Oct 2016 UTC	185.68 GB		
Collection of 2 file(s)	20.16 GB	26 Sep 02:05:08	✓ 26 Sep 03:00:48
Collection of 5 file(s)	71.32 GB	26 Sep 00:07:13	✓ 26 Sep 01:28:46
Server: TC176SP1OB2			
Share: TC176SP1OB2SC1 (Veeam Backup & Replication™)			
TC176SP1OB2SC1 files from 25 Sep 2016 - 1 Oct 2016 UTC	171.13 GB		
Collection of 8 file(s)	85.51 GB	26 Sep 00:36:27	✓ 26 Sep 12:47:12
Server: TC176SP1OB3			
Share: TC176SP1OB3SC1 (Veeam Backup & Replication™)			
TC176SP1OB3SC1 files from 25 Sep 2016 - 1 Oct 2016 UTC	40.25 GB		
Collection of 2 file(s)	20.16 GB	26 Sep 02:05:08	✓ 26 Sep 03:00:43

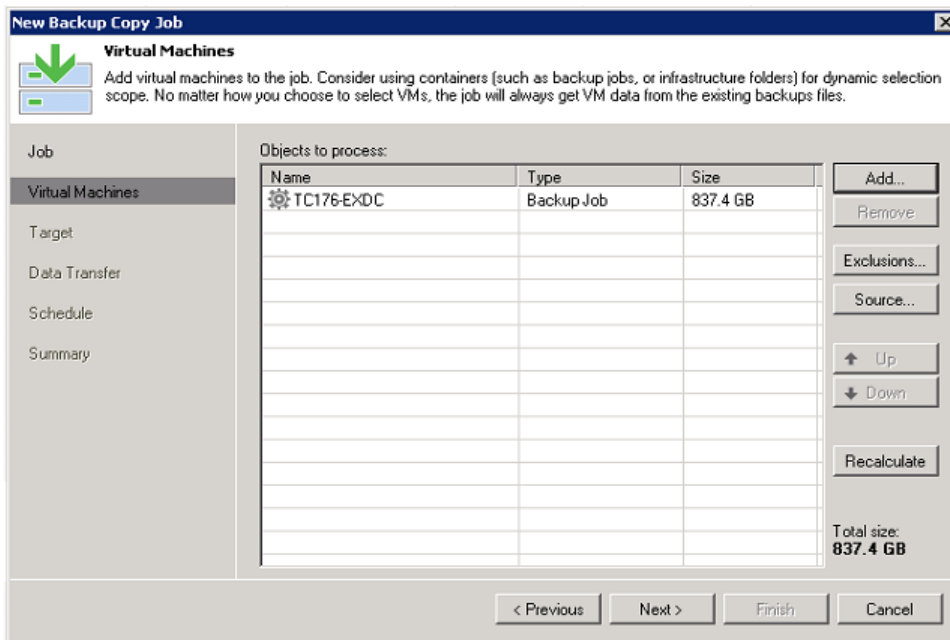
Use Case 2: Efficient long-term backup retention

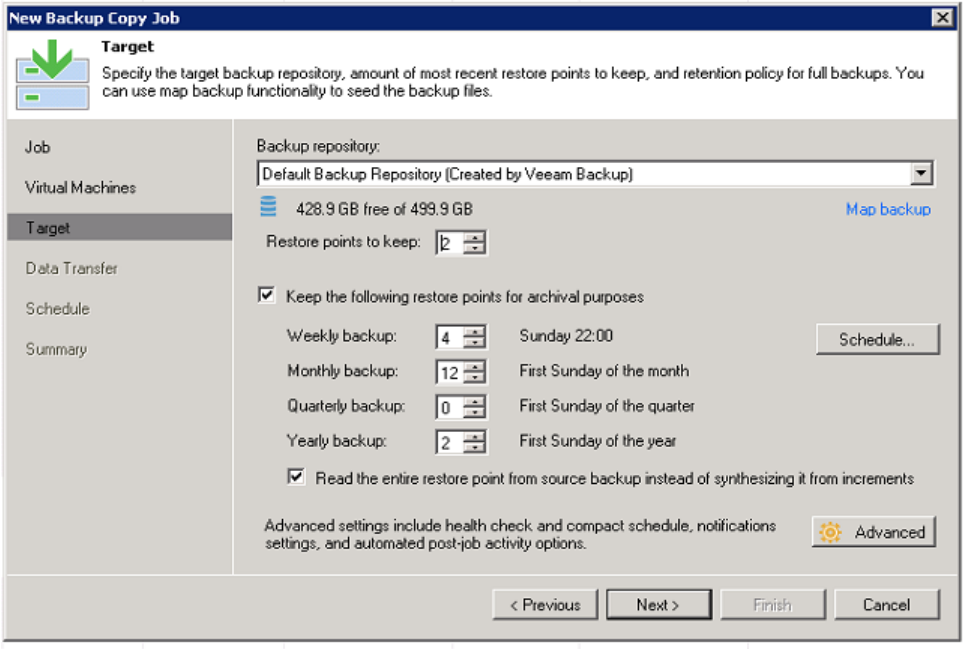
Backup retention periods are increasing, not shrinking.

Since Veeam backup jobs include daily incremental backups in the count of restore points, extended retention of weekly, monthly, quarterly or yearly backups is achieved using Veeam Backup Copy jobs which include a Grandfather-Father-Son retention policy automatically managed by Veeam.

The appropriate Veeam backup is copied to ExaGrid backup target storage and retained by Veeam for an extended time period. ExaGrid's deduplication significantly reduces the cost to retain these backups for extended periods of time.

Here, a Veeam job with a domain controller and exchange server is copied for extended retention, with 1 month of weekly backups, 12 months of monthly backups, and 1 year of yearly backups managed and retained by Veeam and deduplicated on ExaGrid storage.





ExaGrid performs deduplication across all Veeam backups.

Retention Space for Site : TC176S1

Server	Retention Space			% Available	Total Capacity	Over Capacity	Replication Status
	Available	Consumed	Shared				
TC176S1	86,619.29 GB	5,776.33 GB	5%	94%	92,395.62 GB	0.00 GB	Complete
TC176SP1OB1	29,081.24 GB	1,717.30 GB	3%	94%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB2	28,251.80 GB	2,546.74 GB	3%	92%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB3	29,286.25 GB	1,512.29 GB	9%	95%	30,798.54 GB	0.00 GB	Complete

Deduplication Summary for Site : TC176S1

Share	Deduplication Ratio	Total Backup Data	Compressed Size	Distinct Size
TC176S1	✓ 14.96 : 1	86,421.03 GB	5,776.33 GB	5,497.55 GB
TC176SP1OB1				
TC176SP1OB1EXT1	✓ 13.91 : 1	26,083.82 GB	1,875.62 GB	1,660.80 GB
TC176SP1OB2				
TC176SP1OB2EXT1	✓ 12.52 : 1	34,068.11 GB	2,721.08 GB	2,459.88 GB
TC176SP1OB3				
TC176SP1OB3EXT1	✓ 16.6 : 1	26,269.09 GB	1,582.79 GB	1,376.87 GB

Overall deduplication is 30:1 because Veeam deduplicates typically at 2:1

Veeam total backup size after performing 2:1 deduplication.

Requires significantly less storage to retain for 16 weeks on ExaGrid

Veeam Backup Enterprise Manager Enterprise Plus edition

You are logged in as: **TC176Administrator** Sign out

DASHBOARD | REPORTS | JOBS | **VMS** | FILES | ITEMS | REQUESTS | CONFIGURATION

Backup server: ALL Name: Find Restore Fallover plan... Delete View history Export

VM	Backup Server	Job Name	Restore Points	Location	Folder Path	Latest Success
TC176-W2K8-40	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-W2K8	17 points	SCALEOUT1	TC176-VMW6226-W2K8	9/26/2016 12:01:05 am
TC176-W2K8-41	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-W2K8	17 points	SCALEOUT1	TC176-VMW6226-W2K8	9/26/2016 12:01:01 am
TC176-W2K8-42	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-W2K8	17 points	SCALEOUT1	TC176-VMW6226-W2K8	9/26/2016 12:01:02 am
TC176-W2K8-43	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-W2K8	17 points	SCALEOUT1	TC176-VMW6226-W2K8	9/26/2016 12:01:01 am
TC176-W2K8-43	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-BCP	7 points	SCALEOUT1	TC176-VMW6226-BCP	9/16/2016 12:04:08 am
TC176-W2K8-42	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-BCP	7 points	SCALEOUT1	TC176-VMW6226-BCP	9/16/2016 12:03:32 am
TC176-W2K8-41	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-BCP	7 points	SCALEOUT1	TC176-VMW6226-BCP	9/16/2016 12:01:09 am
TC176-W2K8-40	TC176-VEEAM1.tc176.lo...	TC176-VMW6226-BCP	7 points	SCALEOUT1	TC176-VMW6226-BCP	9/16/2016 12:01:09 am
TC176-RH65-1	TC176-VEEAM1.tc176.lo...	TC176-VMW4829-RH65	7 points	SCALEOUT1	TC176-VMW4829-RH65	9/26/2016 02:00:54 am
TC176-RH65-2	TC176-VEEAM1.tc176.lo...	TC176-VMW4829-RH65	7 points	SCALEOUT1	TC176-VMW4829-RH65	9/26/2016 02:00:54 am
TC176-SQL08	TC176-VEEAM1.tc176.lo...	TC176-SQL08	18 points	SCALEOUT1	TC176-SQL08	9/26/2016 12:31:17 am
TC176DC	TC176-VEEAM1.tc176.lo...	TC176-EXDC	15 points	SCALEOUT1	TC176-EXDC	9/26/2016 12:03:47 am
TC176-EXCH	TC176-VEEAM1.tc176.lo...	TC176-EXDC	15 points	SCALEOUT1	TC176-EXDC	9/26/2016 12:04:20 am
TC176-RH65-2	TC176-VEEAM2.tc176.lo...	VMW4829	1 point	TC125H10B4	/home/VeeamUser/TC176VEEAM2/VMW4829/	8/30/2016 04:45:28 pm
TC176-W2K8-21	TC176-VEEAM2.tc176.lo...	VMW4829	1 point	TC125H10B4	/home/VeeamUser/TC176VEEAM2/VMW4829/	8/30/2016 05:17:17 pm

Page 1 of 1 | Displaying items 1 - 21 of 21

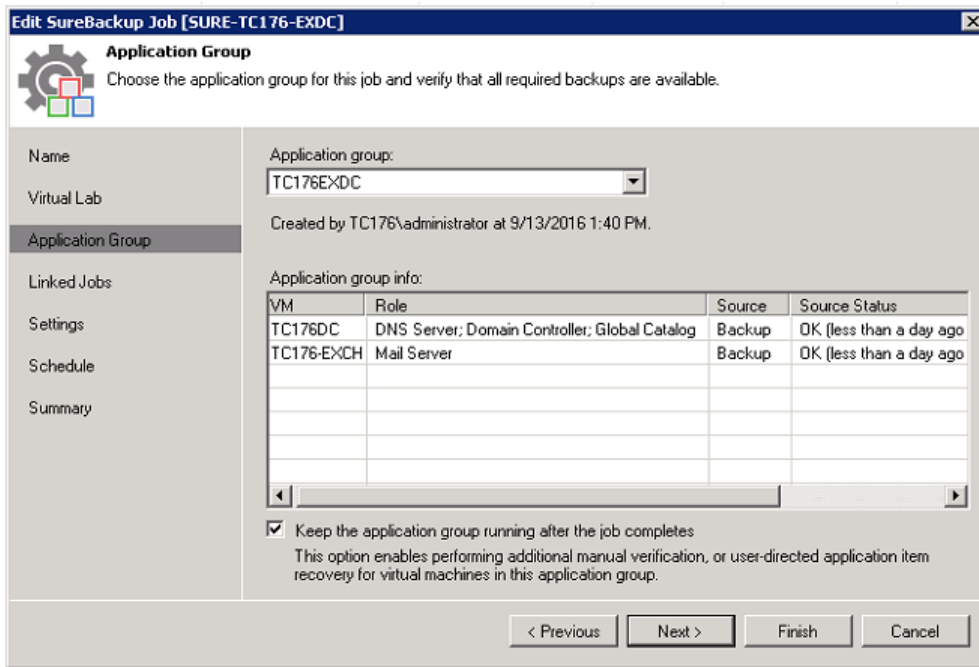
Use Case 3: Easily & regularly prove backups are recoverable

Backups are only as good as their ability to recover data when it's lost. Modern data centers need to regularly test backups and often pass audits of critical backups. Veeam's SureBackup feature simplifies the setup of these backup tests and orchestrates the execution of the tests.

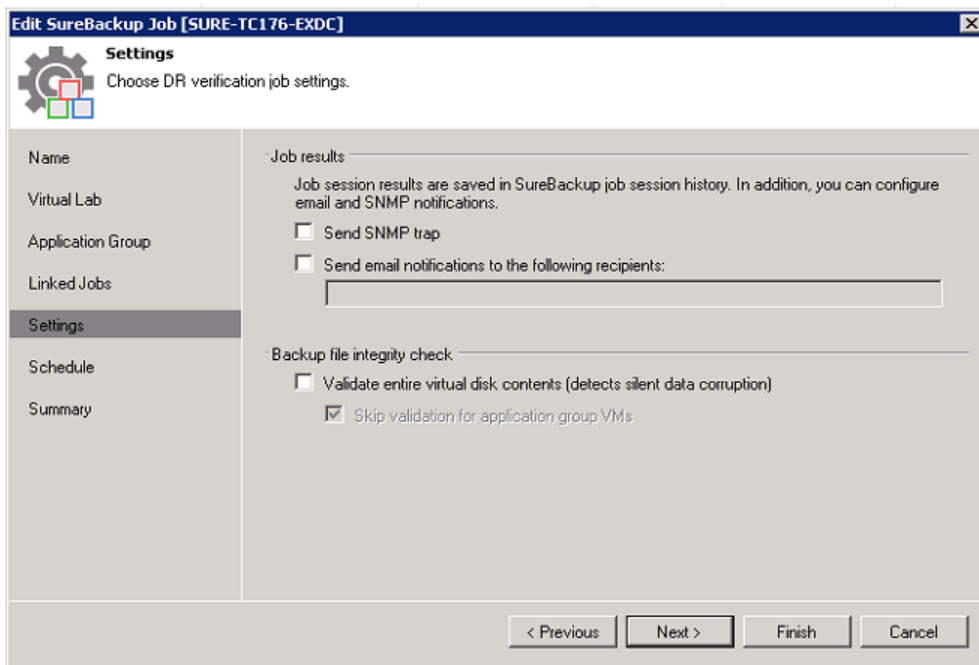
Veeam's SureBackup creates an isolated application environment and stands up one or more VMs in that environment, running the VMs from backups. The combination of Pure Storage, Veeam and ExaGrid excels in the testing of backups because:

1. Veeam's SureBackup makes setup and orchestration of the backup tests simple and easy.
2. ExaGrid's landing zone allows booting the test VMs in seconds to minutes compared to other deduplication appliance architectures. The landing zone is a high-speed disk cache of recent backups in their full, unduplicated form, eliminating the need to rehydrate backup data during backup testing.
3. PureStorage FlashArrays provide high bandwidth and sub millisecond latency fast storage on which Veeam keeps the changes being made while the VMs are running, eliminating a pinch point from slowing down the backup testing and making the backup tests as real-world and realistic as possible. The data on PureStorage Arrays has undergone pattern elimination, deduplication and compression which is key to All FlashArray cost-effectiveness and flash device longevity.

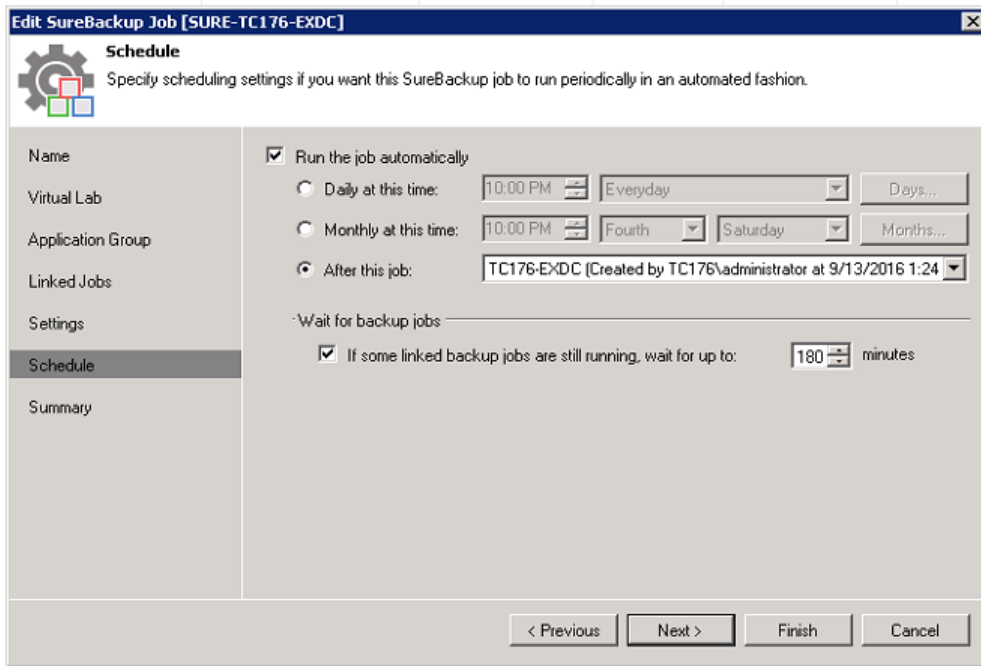
Here, a Veeam application group containing a Domain Controller and Exchange server is used to test recent backups. By specifying VM Role(s), Veeam orchestrates appropriate testing of the powered-on VM, including pings, DNS lookups, AD queries, e-mail operations, etc.



Results can be delivered using SNMP or e-mail as input to auditing processes:



Veeam orchestration of the backup tests:



SureBackup test results allow quick assessment of backup health:

TC176EXDC Session 9/13/2016 1:44:45 PM

VM status:

Name	Status	Heartbeat	Ping	Script	Verification
TC176DC	Success	Success	Success	Success	Disabled
TC176-EXCH	Success	Success	Success	Success	Disabled

Session log:

Message	Duration
✓ Getting virtual lab configuration	
✓ Starting virtual lab routing engine	0:00:43
✓ TC176DC - Publishing	0:00:19
✓ TC176DC - Reconfiguring	0:00:01
✓ TC176DC - Registering	0:00:04
✓ TC176DC - Configuring DC	0:01:14
✓ TC176DC - Powering on	0:05:16
✓ TC176DC - Heartbeat test	0:00:01
✓ TC176DC - Running ping test(s)	0:00:15
✓ TC176DC - Application initialization	0:02:00
✓ TC176DC - Running test scripts	
✓ TC176-EXCH - Publishing	0:00:20
✓ TC176-EXCH - Reconfiguring	
✓ TC176-EXCH - Registering	0:00:05
✓ TC176-EXCH - Powering on	0:03:40
✓ TC176-EXCH - Heartbeat test	0:00:01
✓ TC176-EXCH - Running ping test(s)	0:00:15
✓ TC176-EXCH - Application initialization	0:02:00
✓ TC176-EXCH - Running test scripts	
✓ TC176-EXCH - Powering off	0:00:08
✓ TC176-EXCH - Unregistering	
✓ TC176-EXCH - Cleaning up redo logs	0:00:03
✓ TC176-EXCH - Unpublishing	0:00:05
✓ TC176DC - Powering off	0:00:01
✓ TC176DC - Unregistering	
✓ TC176DC - Cleaning up redo logs	0:00:01
✓ TC176DC - Unpublishing	0:00:05
✓ Stopping virtual lab routing engine	
✓ Job finished	

Stop Session Close

Use Case 4: Instantly Recovering a Virtual Machine

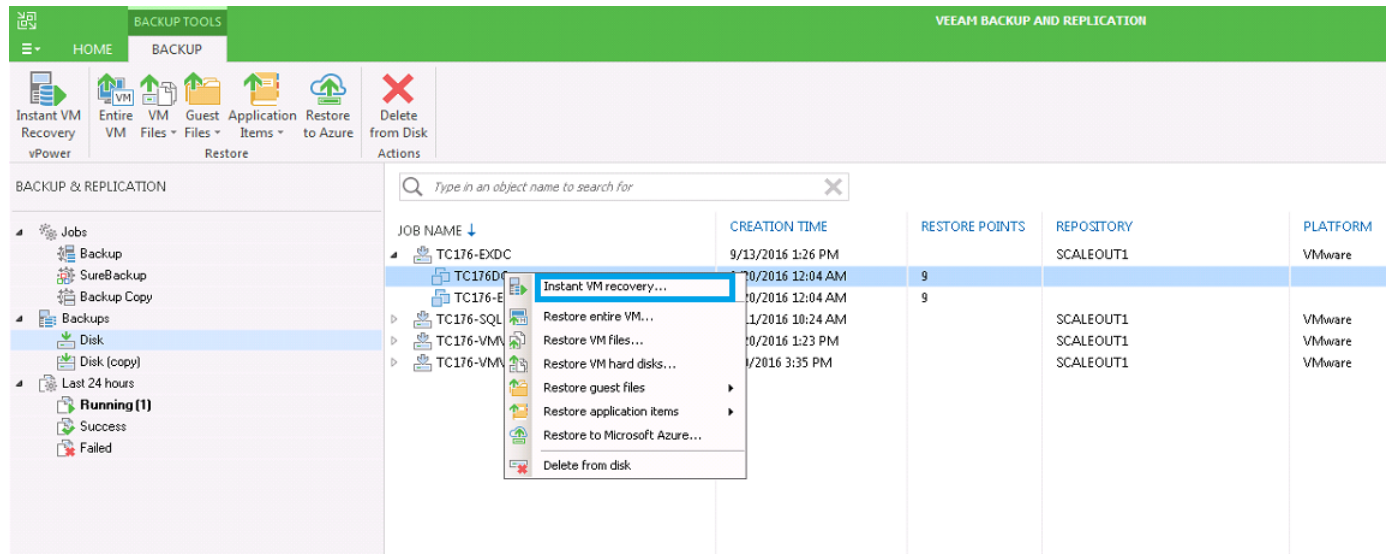
RTOs of applications in modern data centers need to be as short as possible. The hours it can take to fully restore a multi-TB VM is unacceptable. Using Veeam's Instant VM Recovery in conjunction with ExaGrid's landing zone and Pure Storage's FlashArray high performance storage allows a VM to be booted from a backup in seconds to minutes.

ExaGrid's landing zone architecture keeps recent backups on disk in their native form for fastest access and restore speeds. During the Veeam Instant VM Recovery, there is no need to rehydrate deduplicated backups since they are available directly in the landing zone of each ExaGrid appliance.

During Veeam Instant VM Recovery operations, Veeam redirects writes being done to the booting/running VM to PureStorage FlashArray //m storage. The high performance and low latency of the PureStorage FlashArray ensures no storage bottleneck while the VM is booting/running.

Once booted, the VM is typically migrated to production storage. ExaGrid's landing zone facilitates high-bandwidth data transfers through Veeam to high-performance production Pure Storage FlashArrays, assuring the restore to production environment completes as quickly as possible.

Here, a production Active Directory server is booted from a recent backup:



The screenshot shows the Veeam Backup and Replication console interface. The top navigation bar includes 'HOME' and 'BACKUP' tabs. Below the navigation bar, there are several icons for backup and recovery actions: 'Instant VM Recovery vPower', 'Entire VM', 'VM Files', 'Guest Files', 'Application Items', 'Restore to Azure', and 'Delete from Disk Actions'. The main area is divided into a left sidebar and a central table.

The left sidebar, titled 'BACKUP & REPLICATION', shows a tree view of jobs. The 'Running (1)' job is highlighted. The central table displays a list of backup jobs with columns for 'JOB NAME', 'CREATION TIME', 'RESTORE POINTS', 'REPOSITORY', and 'PLATFORM'. A context menu is open over the selected job, with 'Instant VM recovery...' highlighted.

JOB NAME	CREATION TIME	RESTORE POINTS	REPOSITORY	PLATFORM
TC176-EXDC	9/13/2016 1:26 PM		SCALEOUT1	VMware
TC176D...	10/20/2016 12:04 AM	9		
TC176-E...	10/20/2016 12:04 AM	9		
TC176-SQL	11/20/2016 10:24 AM		SCALEOUT1	VMware
TC176-VM...	10/20/2016 1:23 PM		SCALEOUT1	VMware
TC176-VM...	11/20/2016 3:35 PM		SCALEOUT1	VMware

Instant Recovery

Restore Point
Choose restore point you want to recover the selected virtual machine to.

Virtual Machine: VM name: **TC176DC** Original host: **QA-VC3.eng.local**

Restore Point: VM size: **127.0 GB**

Available restore points:

Created	Type
less than a day ago (12:04 AM Tuesday 9/20/2016)	Increment
1 day ago (12:04 AM Monday 9/19/2016)	Increment
2 days ago (12:04 AM Sunday 9/18/2016)	Increment
3 days ago (12:01 AM Saturday 9/17/2016)	Full
4 days ago (4:22 PM Friday 9/16/2016)	Increment
4 days ago (12:01 AM Friday 9/16/2016)	Increment
5 days ago (12:01 AM Thursday 9/15/2016)	Increment
6 days ago (12:01 AM Wednesday 9/14/2016)	Increment
7 days ago (1:26 PM Tuesday 9/13/2016)	Full

< Previous Next > Finish Cancel

Instant Recovery

Destination
Choose ESX server to run the recovered virtual machine on. You can choose to power on VM automatically, unless you need to adjust VM settings first (such as change VM network).

Virtual Machine: Host: Choose...

Restore Point: VM folder: Choose...

Recovery Mode: Restored VM name:

Destination: Resource pool: **Resources**

Datastore:

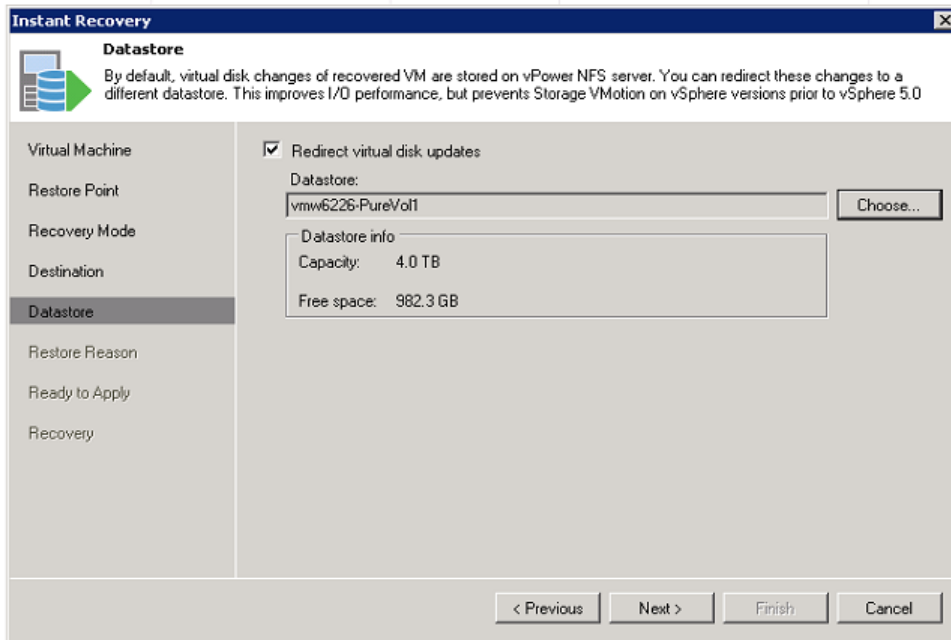
Restore Reason:

Ready to Apply:

Recovery:

< Previous Next > Finish Cancel

If the Veeam vPower NFS server is not using PureStorage FlashArray storage for virtual disk changes, one can be specified here:



The VM is booted from the backup, changes redirected to PureStorage FlashArray, and a migration operation can be initiated in the infrastructure to production storage.



Total RTO – Recovery Time Objective (RTO) is targeted duration of time and a service level within which a running Virtual Machine can be restored from a backup. Total RTO can be drastically reduced due to ExaGrid's Landing Zone architecture and Pure Storage high performance FlashArray storage where all the migration takes place with sub millisecond latency and all the VMs can be restored quickly.

Use Case 5: Recovering granular, application-specific items directly from backups

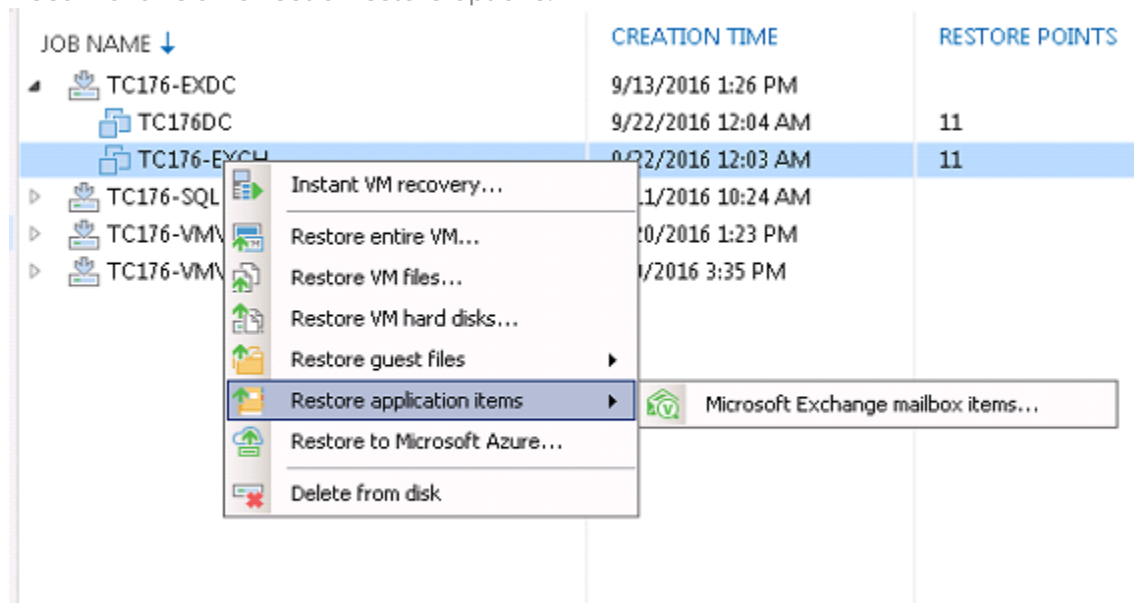
In addition to InstantVM recovery, Veeam provides a comprehensive set of application item recovery mechanisms, allowing the following granular recovery from VM backups without requiring the delay in restoring an entire VM's data or a full VM boot:

- Files
- AD Users, Groups, etc.
- MS SQL databases, tables, etc.
- Oracle databases, tables, etc.
- Exchange e-mails

ExaGrid's landing zone provides the fastest access to backup data, eliminating the need to rehydrate deduplicated backups. As a result, Veeam application item recovery operations such as browsing as well as actual application item restores meet modern data center RTOs.

Using Veeam's Enterprise Manager, these application item restores can be requested directly by users, with access fully controlled by Veeam, freeing up IT staff time for other tasks.

Veeam allows a rich set of restore options:



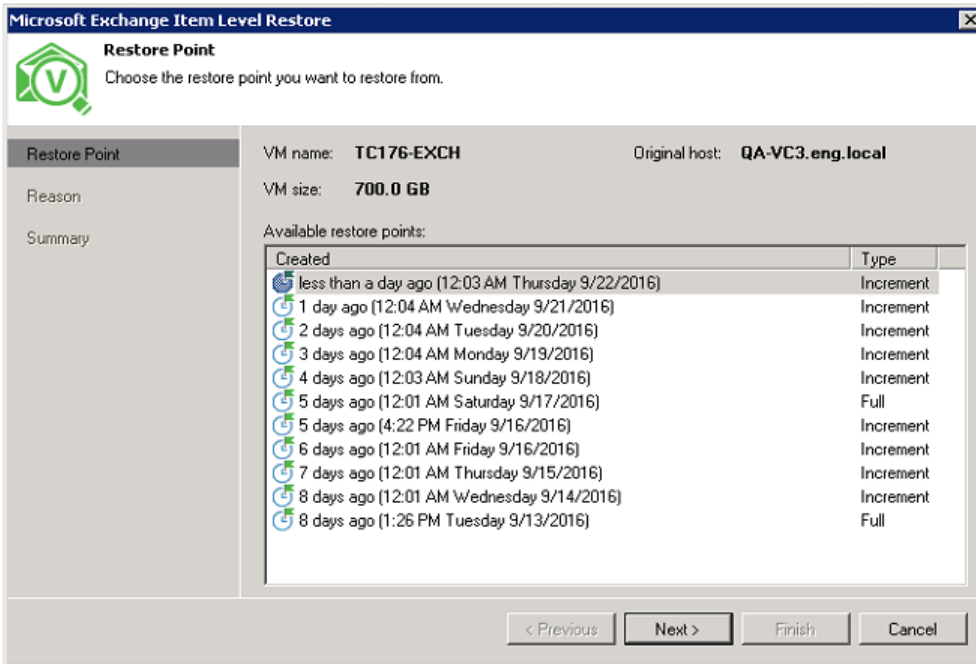
The screenshot shows a table of backup jobs in Veeam Enterprise Manager. The table has three columns: JOB NAME, CREATION TIME, and RESTORE POINTS. A context menu is open over the selected job 'TC176-EXCH', showing various restore options. The 'Restore application items' option is highlighted, and its sub-menu is also visible, showing 'Microsoft Exchange mailbox items...'.

JOB NAME ↓	CREATION TIME	RESTORE POINTS
TC176-EXDC	9/13/2016 1:26 PM	
TC176DC	9/22/2016 12:04 AM	11
TC176-EXCH	9/22/2016 12:03 AM	11
TC176-SQL	9/22/2016 10:24 AM	
TC176-VMX	9/22/2016 1:23 PM	
TC176-VMX	9/22/2016 3:35 PM	

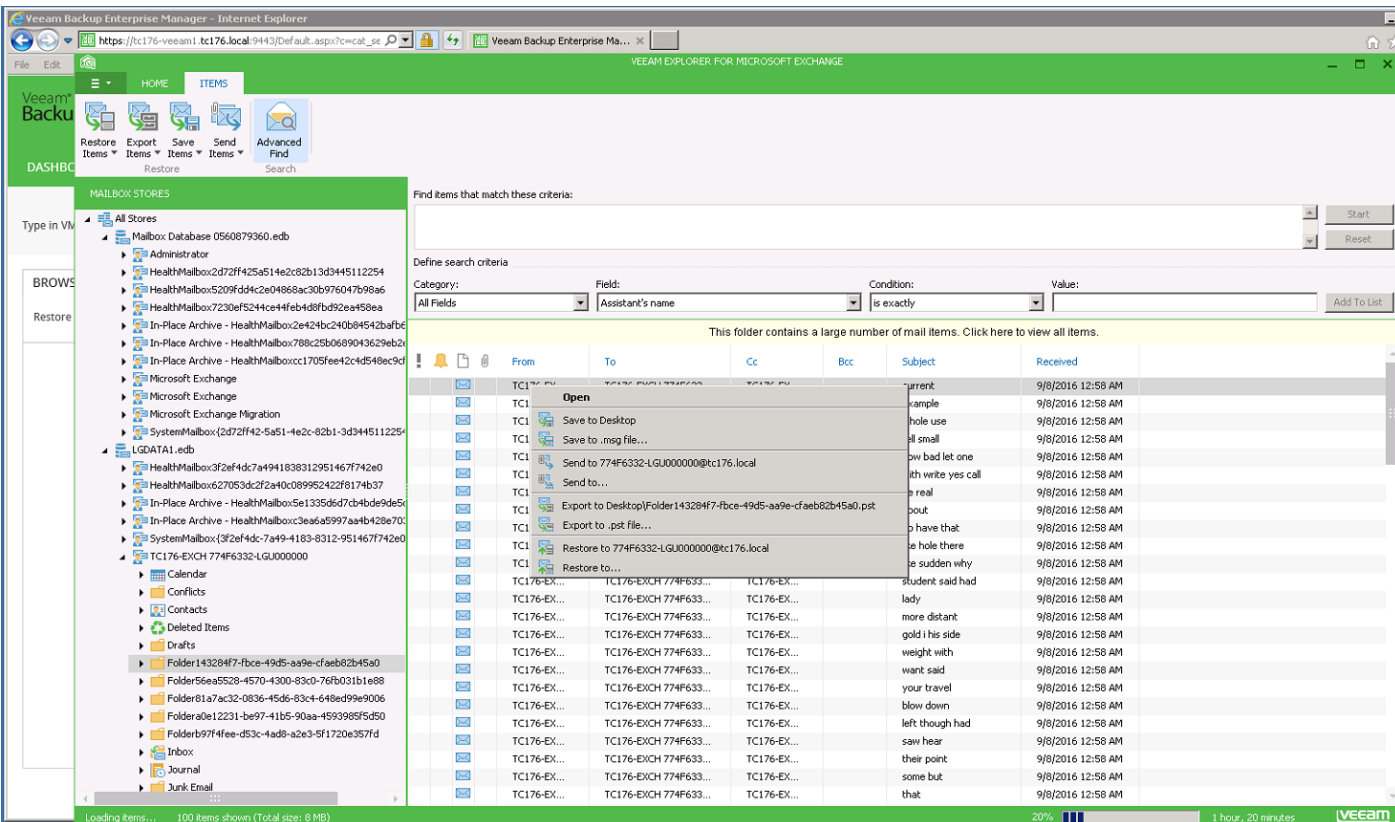
Context menu options for 'TC176-EXCH':

- Instant VM recovery...
- Restore entire VM...
- Restore VM files...
- Restore VM hard disks...
- Restore guest files
- Restore application items
 - Microsoft Exchange mailbox items...
- Restore to Microsoft Azure...
- Delete from disk

Application items can be pulled from any retained backup:



Individual e-mails from a user's inbox or folders can be saved to a variety of locations. "Advanced Find" allows the definition of search criteria to refine results.



Granular item restoration can be extended using Veeam's Enterprise Manager, including self-serve requests:

The screenshot displays the Veeam Backup Enterprise Manager web interface. At the top, it shows the user is logged in as 'TC176\Administrator'. The main content area is titled 'FILES' and indicates that the current view is for 'TC176-EXCH' as of '10/25/2016 12:01:24 am'. Below this, there are 'BROWSE' and 'SEARCH' tabs. A 'Restore point' dropdown is set to '10/25/2016 12:01:24 am', with buttons for 'Restore', 'Download', and 'Add to restore list'. A search bar is also present. The left sidebar shows a file tree with folders like 'Classic .NET AppPool', 'Default', 'Default User', 'Public', 'Windows', 'E:\', 'SRECYCLE.BIN', and 'CDFdata1'. The main pane shows a list of files with columns for 'Name', 'Date Modified', 'Size', and 'Owner'. The file 'Order Details.TXT' is highlighted. At the bottom, there is a pagination control showing 'Page 1 of 1' and 'Displaying 1 - 13 of 13'. A 'View history' link is also visible.

Name	Date Modified	Size	Owner
Categories.TXT	2/8/2006 01:21:56 pm	60.4 MB	Administrators
CustomerCustomerDemo.TXT	3/13/2006 10:49:38 am	5.1 MB	Administrators
CustomerDemographics.TXT	2/8/2006 01:23:25 pm	4.6 MB	Administrators
Customers.TXT	2/8/2006 01:23:15 pm	30.1 MB	Administrators
Employees.TXT	2/8/2006 01:21:15 pm	119.8 MB	Administrators
EmployeeTerritories.TXT	3/13/2006 10:49:44 am	5.1 MB	Administrators
Order Details.TXT	2/8/2006 01:24:47 pm	9.9 MB	Administrators
Orders.TXT	2/8/2006 01:24:40 pm	32.3 MB	Administrators
Products.TXT	2/8/2006 01:23:44 pm	17.4 MB	Administrators
Region.TXT	2/8/2006 01:21:43 pm	7.7 MB	Administrators
Shippers.TXT	2/8/2006 01:21:29 pm	9.6 MB	Administrators

Deployment

The deployments of the PureStorage FlashArray, Veeam Backup & Replication, and ExaGrid all go from box to running in minutes to hours.

- PureStorage FlashArray//m20 has been deployed with iSCSI connections to the top of the rack switch. There are four 10 Gb NICs which are used in active-active configuration in the front end of PureStorage FlashArray//m20. All the 4 physical iSCSI connections are connected to the four 10Gb ports on the top of rack switch. There are 2 active-active controllers processing data. The performance for PureStorage FlashArray //m10 is about 100,000 IOPS and 3GB/s bandwidth at 32KB block size with sub millisecond latency for 100% Read workload.
- Veeam Backup & Replication deploys (generally) on VMs, with a backup server first, then additional backup proxy servers added based on size of environment and access to primary storage.
- After rack mounting, ExaGrid appliances are initialized via a standard web browser, joined together in a scale-out GRID, and provisioned with at least one Veeam share per appliance.
- A Veeam Backup Repository is configured for each ExaGrid appliance/share, and then these Veeam repositories are combined into a “Scale-Out Backup Repository” – SOBR.
- Veeam jobs are created or changed to use the scale-out repository, providing Veeam access to all ExaGrid appliances/shares.
- When Veeam jobs run, Veeam selects which ExaGrid appliance receives the individual per-VM backups. This “scale out everywhere” provides flexibility and efficiency for future growth. As VMs grow, new VMs are added, and ExaGrid storage is expanded, there is virtually no Veeam configuration re-work required.

The following sections contain key concepts and screen shots of typical deployment steps.

PureStorage FlashArray & Hypervisor Configuration

Configure PureStorage FlashArray volumes following best practices for the hypervisor in use. The following parameters need to be set as part of the best practices guide on the ESX hypervisor.

1. **VMware Native Multipathing Plugin:** Round Robin path selection policy (PSP) must be used on all hosts connected to the array.
2. **Set IO Operation Limit to 1**
 - 2.1 Round Robin can also be set on a per-device, per-host basis using the vSphere Web Client. The procedure to setup Round Robin policy for a Pure Storage volume is shown in the below figure. Note that this does not set the IO Operation

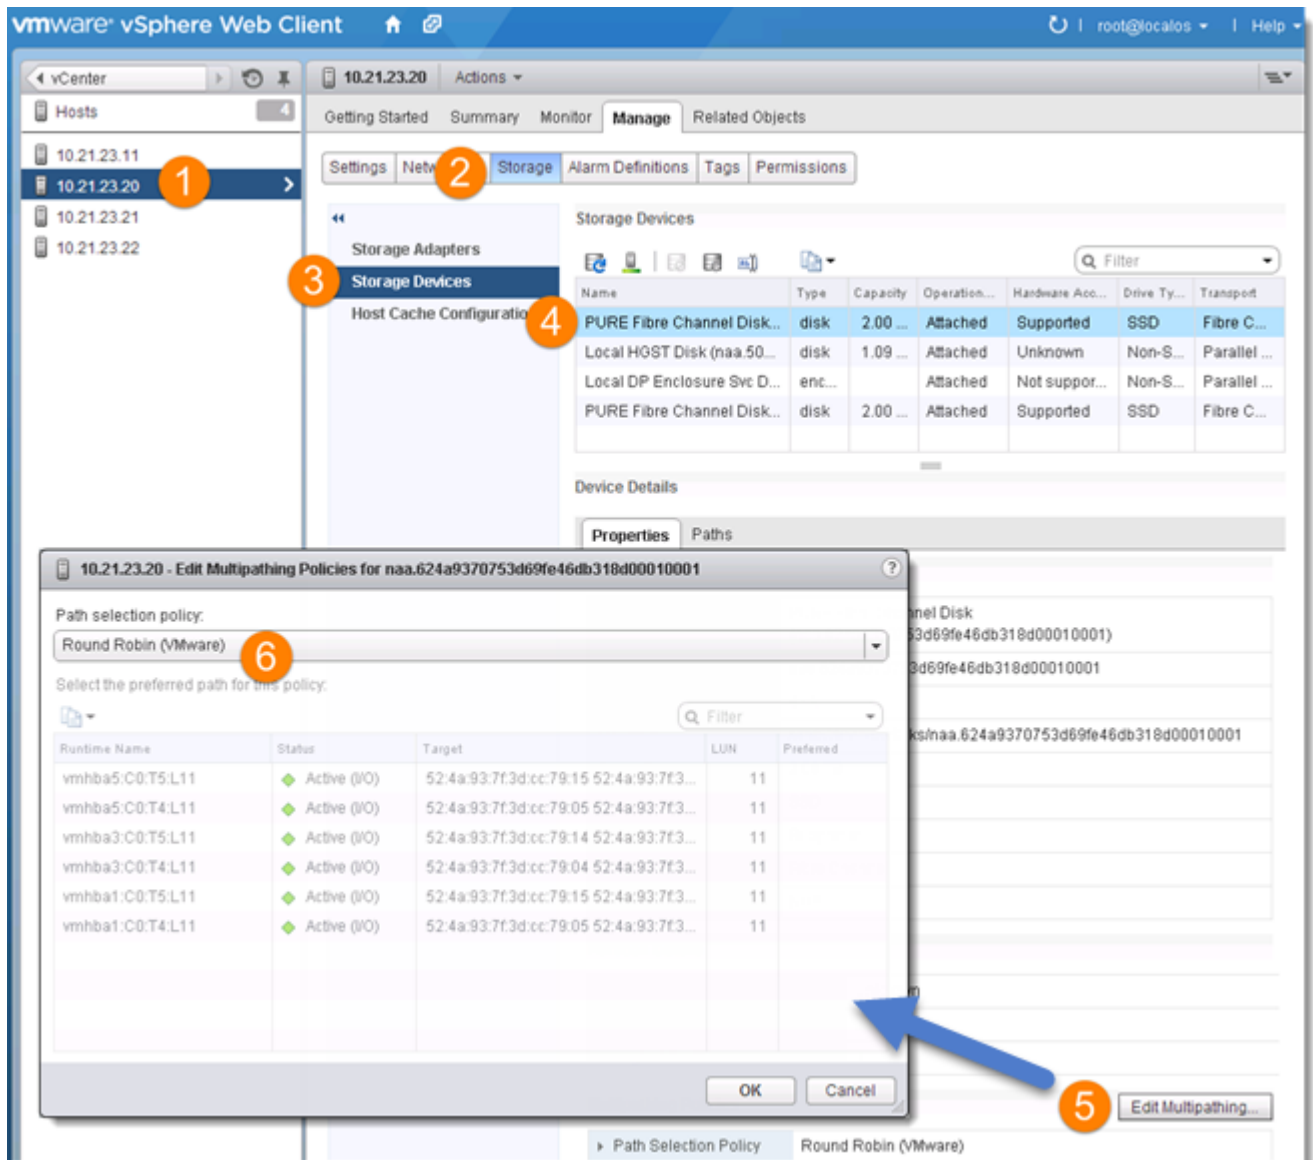


Figure: Setting Round Robin on a device with the vSphere Web Client

Limit to 1 which is a command line option only.

To set a device that is pre-existing to have an IO Operation limit of one, run the following command:

```
esxcli storage nmp psp roundrobin deviceconfig set -d naa.<naa address>
-I 1 -t iops
```

3. **Performance:** Eagerzeroedthick virtual disks perform the best followed by Zeroedthick and then thin virtual disks.
4. **Protection against space exhaustion:** Zeroedthick and Eagerzeroedthick virtual disks are not susceptible to VMFS logical capacity exhaustion because the space is reserved on the VMFS upon creation.
5. **Virtual Disk Density:** It should be noted that while all virtual disk types take up the same amount of physical space on the FlashArray due to FlashReduce technology, they have different requirements on the VMFS layer. Thin virtual disks can be oversubscribed (more capacity provisioned than the VMFS reports as being available) allowing for far more virtual disks to fit on a given volume than either of the thick formats. This provides a greater virtual machine to VMFS density and reduces the number of volumes that are required to store them.
6. **XCOPY Performance:** Eagerzeroedthick and Zeroedthick virtual disks copy significantly faster than thin virtual disks when cloned or moved using VAAI XCOPY.
7. **Time to create:** the virtual disk types also vary in how long it takes to initially create them. Since thin and Zeroedthick virtual disks do not zero space until they are actually written to by a guest they are both created in trivial amounts of time—usually a second or two. Eagerzeroedthick disks, on the other hand, are pre-zeroed at creation and consequently take additional time to create. If the time-to-first-IO is paramount for whatever reason, thin or zeroedthick is best.
8. **Verifying that VAAI is enabled**

In ESXi 5.x hosts, to determine if VAAI is enabled using the service console in ESXi or the vCLI in ESXi, run these command to check if Int Value is set to 1 (enabled):

```
esxcli system settings advanced list -o
/DataMover/HardwareAcceleratedMove
esxcli system settings advanced list -o
/DataMover/HardwareAcceleratedInit
esxcli system settings advanced list -o
/VMFS3/HardwareAcceleratedLocking
```

You will see an output similar to:

```
Path: /VMFS3/HardwareAcceleratedLocking
Type: integer
Int Value: 1           ← Value is 1 if enabled
Default Int Value: 1
Min Value: 0
Max Value: 1
String Value:
Default String Value:
Valid Characters:
Description: Enable hardware accelerated VMFS locking (requires
compliant hardware)
```

8.1 To enable atomic test and set (ATS) AKA hardware accelerated locking:

```
esxcli system settings advanced set -i 1 -o  
/VMFS3/HardwareAcceleratedLocking
```

8.2 To enable Hardware accelerated initialization AKA WRITE SAME:

```
esxcli system settings advanced set --int-value 1 --option  
/DataMover/HardwareAcceleratedInit
```

8.3 To enable Hardware accelerated data move AKA XCOPY (full copy):

```
esxcli system settings advanced set --int-value 1 --option  
/DataMover/HardwareAcceleratedMove
```

9. iSCSI Tuning

Just like any other array that supports iSCSI, Pure Storage recommends the following changes to an iSCSI-based vSphere environment for the best performance. Note that these changes are not uncommon as most vendors have similar recommendations.

9.1 Set login timeout to a larger value

For example, to set the login timeout value to 30 seconds, use commands similar to the following:

```
vmkiscsi-tool -W -a "login_timeout=30" vmhba37
```

```
esxcli iscsi adapter param set -A vmhba37 -k LoginTimeout -v 30
```

The default login timeout value is 5 seconds and the maximum value that you can set is 60 seconds.

This can be also done via the vSphere Web Client, by clicking on the software iSCSI adapter properties and then selecting advanced settings.

9.2 Disable Delayed ACK

Disabling Delayed Ack in ESXi 5.x

1. Log in to the vSphere Web Client and select the host under Hosts and Clusters.
2. Navigate to the Manage tab.
3. Select the Storage option.
4. Under Storage Adapters Select the iSCSI vmhba to be modified.

5. Modify the delayed Ack setting using the option that best matches your site's needs, as follows:

Case 1: Modify the delayed Ack setting on a discovery address (recommended).

1. Select Targets.
2. On a discovery address, select the Dynamic Discovery tab.
3. Select the iSCSI server.
4. Click Advanced.
5. Change DelayedAck to false.

Case 2: Modify the delayed Ack setting on a specific target.

1. Select Targets.
2. Select the Static Discovery tab.
3. Select the iSCSI server.
4. Click Advanced.
5. Change DelayedAck to false.

Case 3: Modify the delayed Ack setting globally for the iSCSI adapter.

1. Select the Advanced Options tab.
2. Click Advanced.
3. Change DelayedAck to false.

9.3 Configure end-to-end jumbo frames

In many environments with iSCSI it is recommended to enable Jumbo Frames for performance boosts or to adhere with the external network configuration. Enabling Jumbo Frames is a cross-environment change so careful coordination is required to ensure proper configuration. It is important to work with your networking team and Pure Storage representatives when enabling Jumbo Frames.

1. Configure Jumbo frames on Pure using the GUI or CLI.

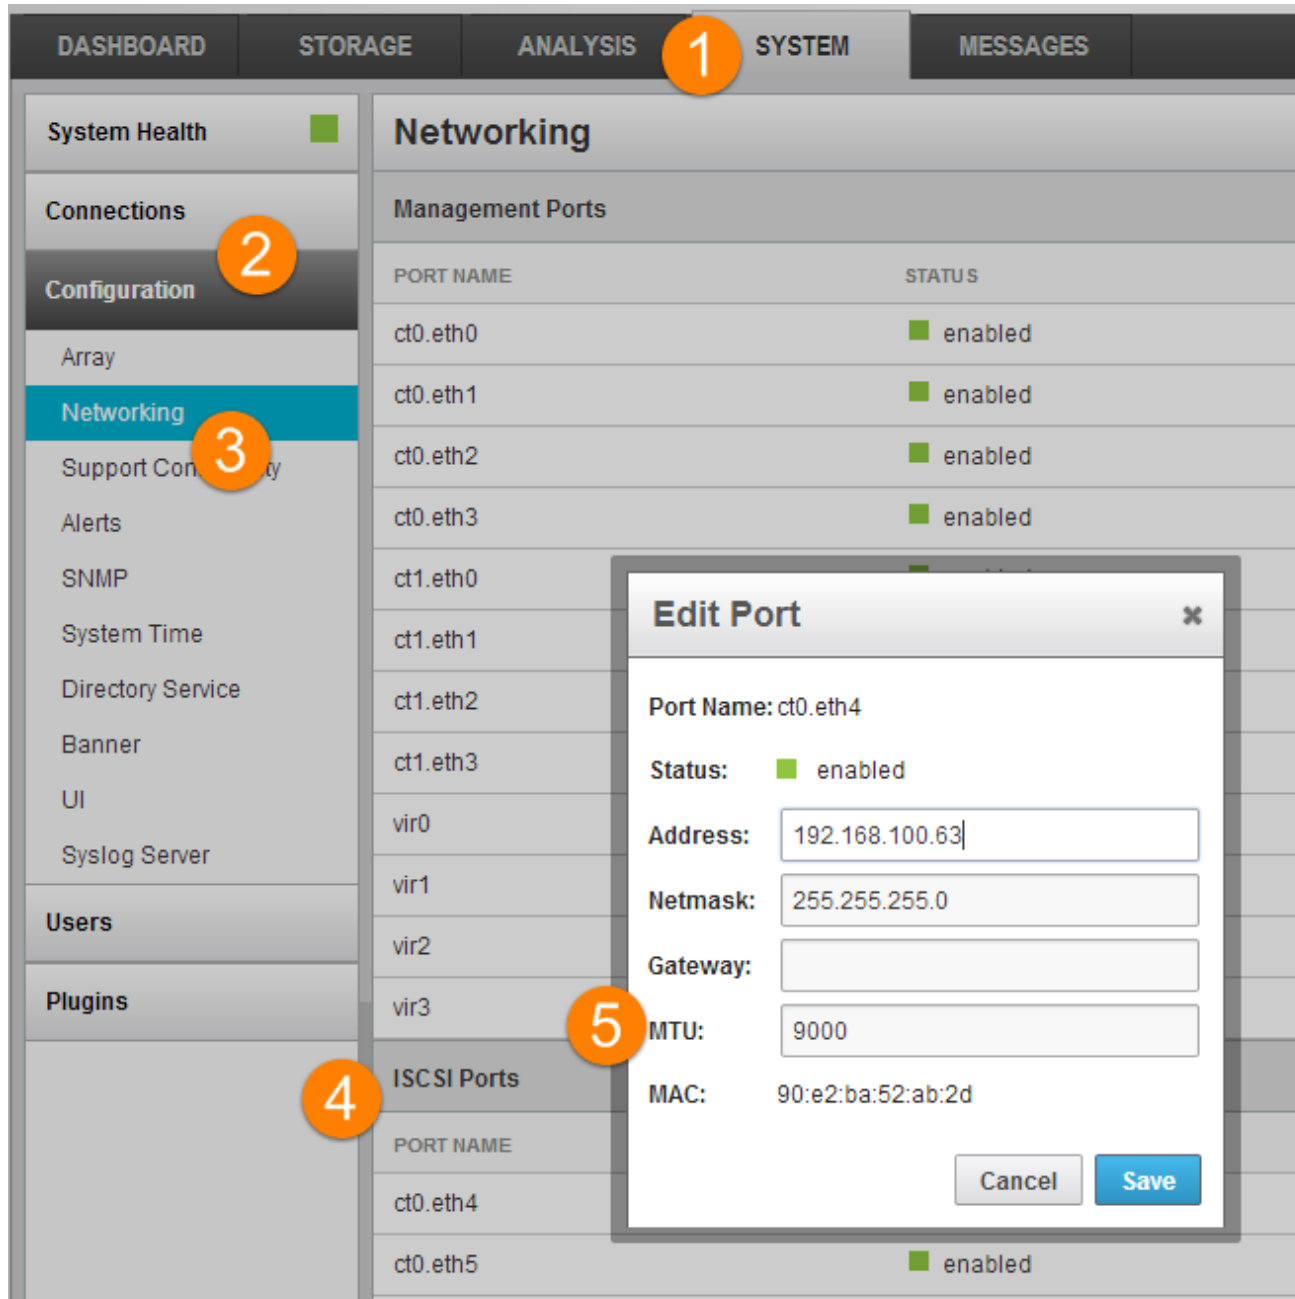


Figure: Configuring Jumbo Frames on the Pure Storage FlashArray

```
purenetwork -setattr -mtu <MTU> <Ethernet-interface>
```

2. Configure jumbo frames on the network switch for each port using the relevant switch CLI or GUI.

3. Configure jumbo frames on ESX. (Refer to Figure above in Step 5)
 - A. Browse to a host in the vSphere Web Client navigator.
 - B. Click the Manage tab, and select Networking > Virtual Switches.
 - C. Select a standard switch from the virtual switches table.
 - D. Click the name of the VMkernel network adapter.
 - E. Click Edit.
 - F. Click NIC settings and set the MTU to 9000.

Once jumbo frames are configured, verify end-to-end jumbo frame compatibility. To verify, try to ping a route with vmkping.

```
vmkping -s 9000 <ip address of Pure Storage iSCSI port>
```

This ping operation shouldn't hang otherwise jumbo frames are not properly configured at some point in the network.

9.4 Tuning iSCSI for a single host

When using iSCSI, there is a performance limitation of 16 I/Os in flight (~30K IOPS with latency of 0.5 ms) per iSCSI session, per port. This doesn't mean that the port cannot go over 30K IOPS, but if you have only 1 host with one iSCSI session that's what you'll get. If more hosts are connected to the same port, then the port can handle more than 16 I/Os and will go over 30K IOPS without a problem.

So it is important to note that if you have only one host connected you will get:

Connected to 2 iSCSI target ports: 32 I/Os in flight = ~60K IOPS @ 0.5 ms

Connected to 4 iSCSI target ports: 64 I/Os in flight = ~120K IOPS @ 0.5 ms

There are 3 ways to work around this limitation when doing performance testing over iSCSI:

1. Use as many target ports as possible.
2. Use more than one host to drive I/O.
3. If the above are not possible, configure the host to create more than 1 iSCSI session per port.

Here's how to create multiple iSCSI sessions per port (you can create 4 sessions per port to increase BW)

```
esxcli iscsi session list
```

Note down the session ID, so say the ID is 00023d000008.

```
esxcli iscsi session add -s 00023d000008 -A vmhba32 -n iqn.2010-06.com.purestorage:flasharray.3efb9d4xxxxxxxxxx97
```

10. Hypervisor Considerations

Avoid provisioning physical RDMs or independent disks. Use virtual mode for any RDM.

Veeam's vPower technology enables the following Veeam features:

1. Instant VM Recovery – boot a VM in seconds to minutes from backups
2. SureBackup – ensure integrity of backup including multiple VMs making up an application
3. VirtualLab – isolated VMs running from backups for upgrade testing, devops, etc.

All of these vPower features require Veeam to store changes to the running VMs in a separate datastore. For best vPower performance, ensure the datastore used for Veeam vPower operations uses FlashArray storage. This can be done by:

1. Selecting folders/volumes on the Veeam mount server that you know are backed by FlashArray storage, or
2. Provisioning FlashArray storage and hypervisor datastore specifically for use by vPower

ExaGrid Deployment

After physically racking the ExaGrid appliances, a web-based GUI is used to initialize each appliance and to join them together into a site. The following screen shots provide an example of part of this initialization:

EXAGRID™

Network Configuration (2 of 2)

Use this page to define the ExaGrid Server's name and network configuration. You will need to click *Apply* after making changes. Now is the right time to make sure all network cables are connected to the appropriate switch ports.

Name	Type	Link	VLAN	Enabled	IP MTU	DHCP	IP Address	Mask	Gateway	ExaGrid	Traffic	Backup
NIC1	1Gb	✓	none	✓	1500	✓	172.21.0.131	/ 22	default	All Sites	▼	✓

ExaGrid Server Name:

Default Gateway: Static
 From DHCP

DNS Servers:

DNS Search Path: eng.local

EXAGRID™

Site Configuration

Do you wish to join this server to an existing ExaGrid Site or to create a new ExaGrid Site?

- Add this ExaGrid Server to an existing ExaGrid Site.
 - To join an existing ExaGrid Site, you will need the ExaGrid Site Administration Key from the other site.
- Use this Server as the first in a new ExaGrid Site.
 - An ExaGrid Site contains one or more interconnected servers.

Click **Next** to begin initializing this ExaGrid Server.



ExaGrid Site To Join

The ExaGrid Server being initialized will be part of an existing site.

- An ExaGrid System consists of one or more interconnected ExaGrid Sites in a *hub and spoke* model.
- See the *Assembling and Initializing Your ExaGrid System* guide for more on configuration.

Existing Site's IP Address:

Existing Site's ExaGrid Site Administration Key:

To obtain a copy of the ExaGrid Site Administration Key:

- From a new browser, log into the existing ExaGrid Site.
- From the ExaGrid administration interface main menu, select **Manage > Security > Site Administration Key**. The Site Administration Key will be displayed.
- Copy and paste the ExaGrid Site Administration Key from the page into the space provided above. **Note:** The key is case sensitive.

Click **Next** to begin initializing this ExaGrid Server.



< Back

Next >

Once initialization is complete, ExaGrid's web-based GUI provides a series of management pages and reports. The following is an example of a two site topology, one site containing three EX32000E appliances and the other containing a single EX32000E.

EXAGRID

Reports Manage Logout Help

Summary for Site : TC176S1 Report Generated:

Landing Space for Site : TC176S1

Server	Available			Deduplication Status
	Actual	% Available for Next Backup	Maximum	
TC176S1	96,000.00 GB	100%	96,000.00 GB	Complete
TC176SP1OB1	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB2	32,000.00 GB	100%	32,000.00 GB	Complete
TC176SP1OB3	32,000.00 GB	100%	32,000.00 GB	Complete

Retention Space for Site : TC176S1

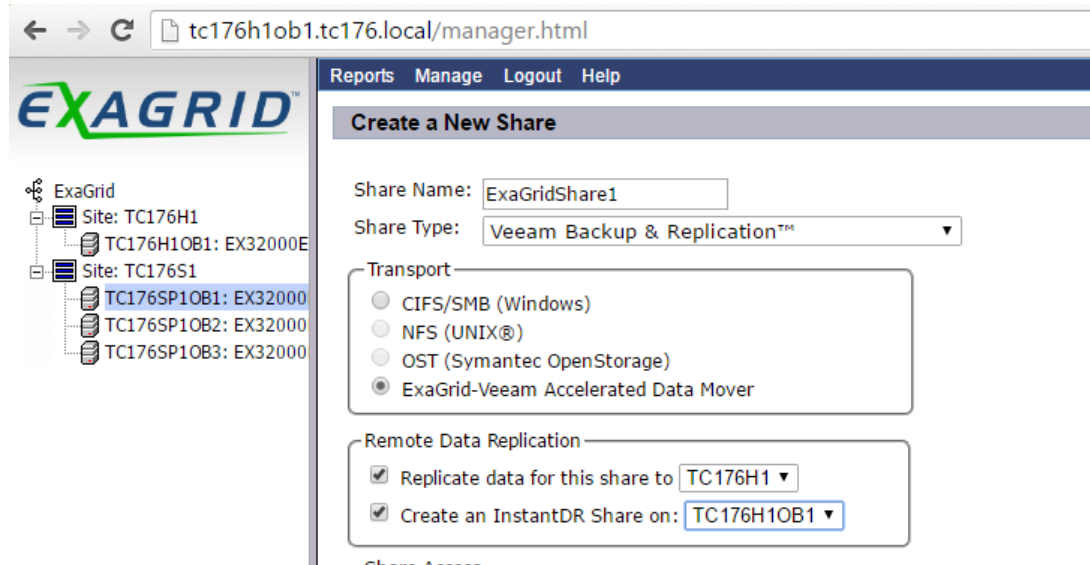
Server	Retention Space				Total Capacity	Over Capacity	Replication Status
	Available	Consumed	Shared	% Available			
TC176S1	86,619.29 GB	5,776.33 GB	5%	94%	92,395.62 GB	0.00 GB	Complete
TC176SP1OB1	29,081.24 GB	1,717.30 GB	3%	94%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB2	28,251.80 GB	2,546.74 GB	3%	92%	30,798.54 GB	0.00 GB	Complete
TC176SP1OB3	29,286.25 GB	1,512.29 GB	9%	95%	30,798.54 GB	0.00 GB	Complete

Details on Veeam and ExaGrid Configuration

The screen shots below show more details on the Veeam and ExaGrid configuration. ExaGrid provides best practices for this configuration and Veeam has several knowledge base articles covering best practices.

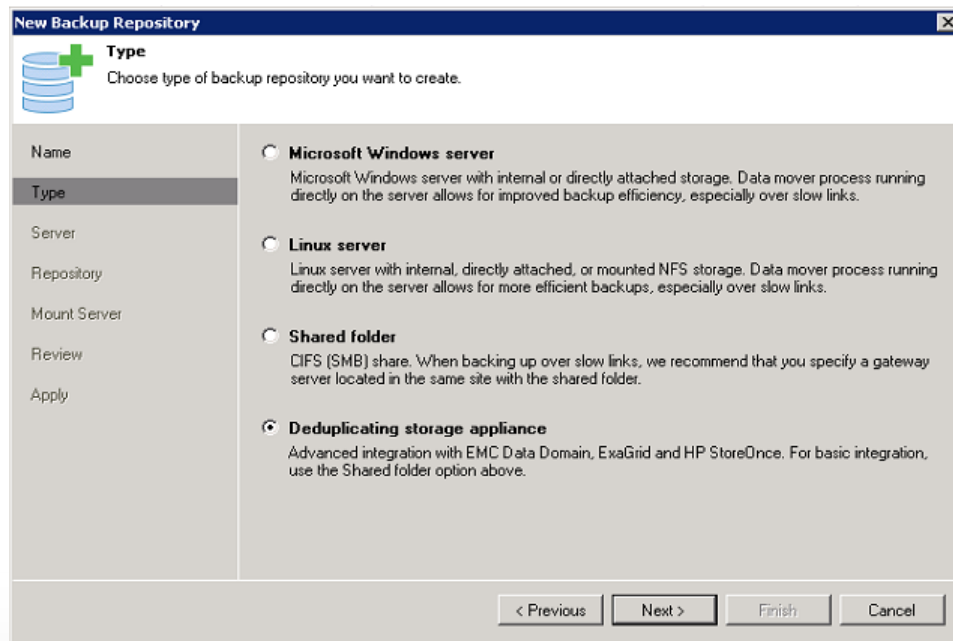
Creating ExaGrid shares and Veeam Repositories

At least one Veeam share is created on each ExaGrid appliance:



Note options for replicating this share to a 2nd ExaGrid site.

A Veeam repository is created for each ExaGrid share:



New Backup Repository ✕

Deduplicating Storage
Specify the deduplication storage appliance you want to add.

Name
Type
Deduplicating Storage
Server
Repository
Mount Server
Review
Apply

EMC Data Domain
DD Boost license, and DD OS 5.5 or later are required. If your Data Domain storage does not meet these requirements, go back and choose the Shared Folder option instead.

ExaGrid
ExaGrid firmware version 4.7 or later is required. If your ExaGrid storage does not meet these requirements, go back and choose the Shared Folder option instead.

HP StoreOnce
StoreOnce Catalyst license and StoreOnce firmware version 3.13.1 or later. If your StoreOnce appliance does not meet these requirements, go back and choose the Shared Folder option instead.

< Previous Next > Finish Cancel

New Backup Repository ✕

Server
Choose server backing your repository. You can select server from the list of managed servers added to the console.

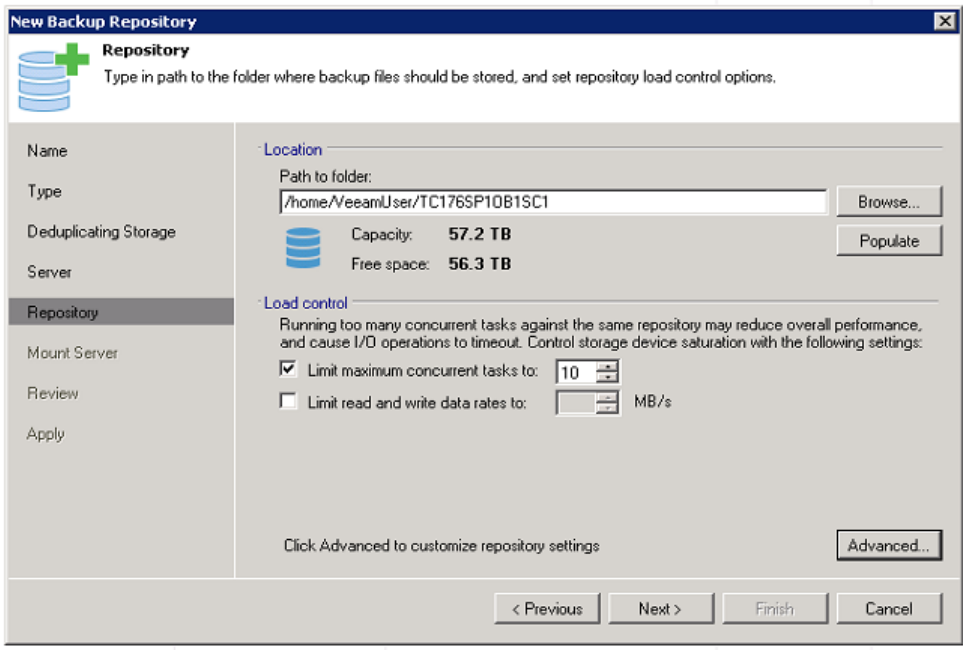
Name
Type
Deduplicating Storage
Server
Repository
Mount Server
Review
Apply

Repository server:
TC176SP10B1 (Created by TC176\administrator at 9/9/2016 3:05 PM.) Add New...

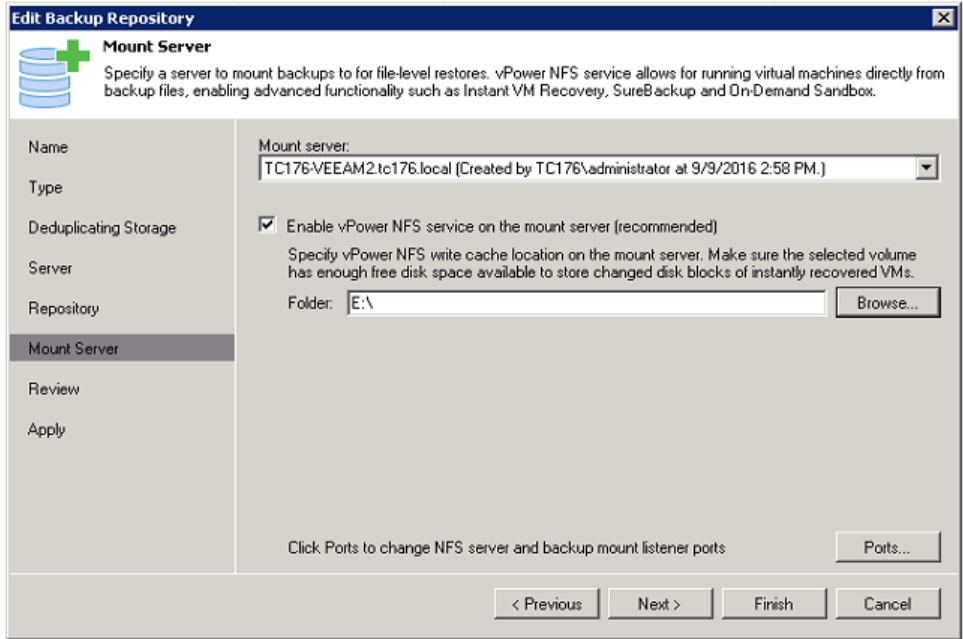
Path ^	Capacity	Free
📁 /home/VeeamUser/TC176SP10B1SC1 (/dev...	57.2 TB	56.3 TB

Populate

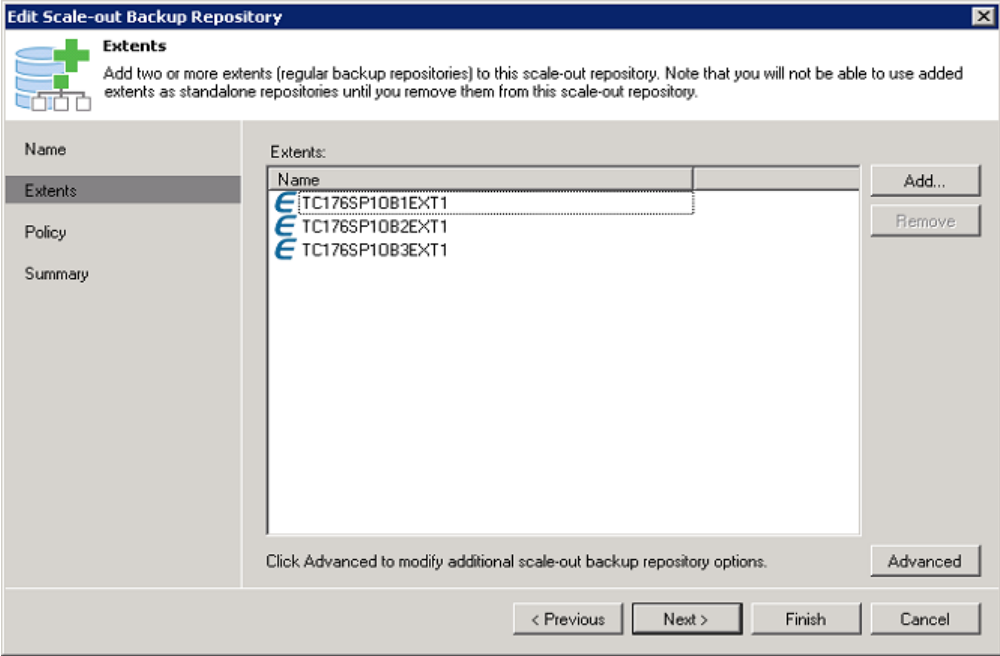
< Previous Next > Finish Cancel



Specify a mount server for use in vPower operations like Instant VM recovery, SureBackup, Virtual Lab, etc. Specify a folder that is backed by FlashArray storage for best performance.



Combine the Veeam repository created for each ExaGrid Veeam share into a single scale-out repository:



Creating Veeam Backup Jobs

Create Veeam backup jobs, sending backups through the Veeam scale-out repository to the ExaGrid site (multiple appliances). Here's an example of a Veeam job containing a domain controller and Exchange VM, with 14 restore points (two weeks):

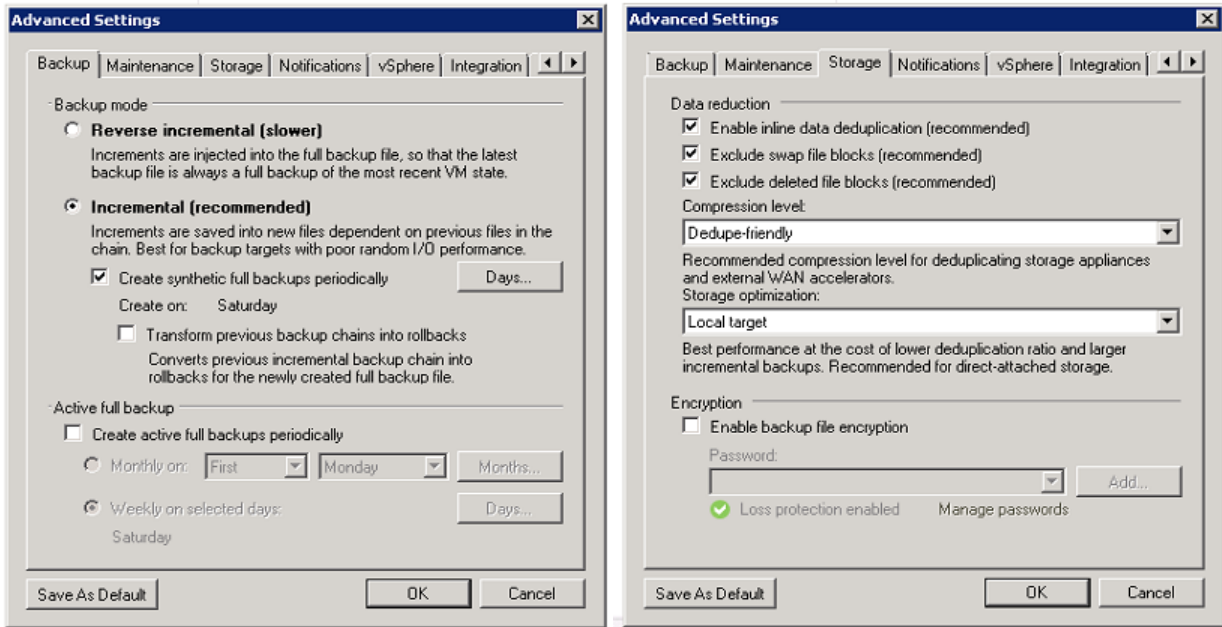
The screenshot shows the 'Edit Backup Job [TC176-EXDC]' window with the 'Virtual Machines' tab selected. The window title is 'Edit Backup Job [TC176-EXDC]'. The 'Virtual Machines' section is active, showing a list of virtual machines to be backed up. The list has columns for Name, Type, and Size. Two VMs are listed: TC176DC (Virtual Machine, 127.0 GB) and TC176-EXCH (Virtual Machine, 700.0 GB). The total size is 827.0 GB. Navigation buttons include '< Previous', 'Next >', 'Finish', and 'Cancel'. A 'Recalculate' button is also present.

Name	Type	Size
TC176DC	Virtual Machine	127.0 GB
TC176-EXCH	Virtual Machine	700.0 GB

The screenshot shows the 'Edit Backup Job [TC176-EXDC]' window with the 'Storage' tab selected. The window title is 'Edit Backup Job [TC176-EXDC]'. The 'Storage' section is active, showing backup proxy and repository settings. The backup proxy is set to 'Automatic selection'. The backup repository is 'SCALEOUT1 (Created by TC176\administrator at 9/16/2016 2:37 PM.)' with 169.1 TB free of 171.5 TB. The retention policy is set to 14 restore points to keep on disk. There is an option to 'Configure secondary destinations for this job' which is currently unchecked. An 'Advanced' button is visible at the bottom right. Navigation buttons include '< Previous', 'Next >', 'Finish', and 'Cancel'.

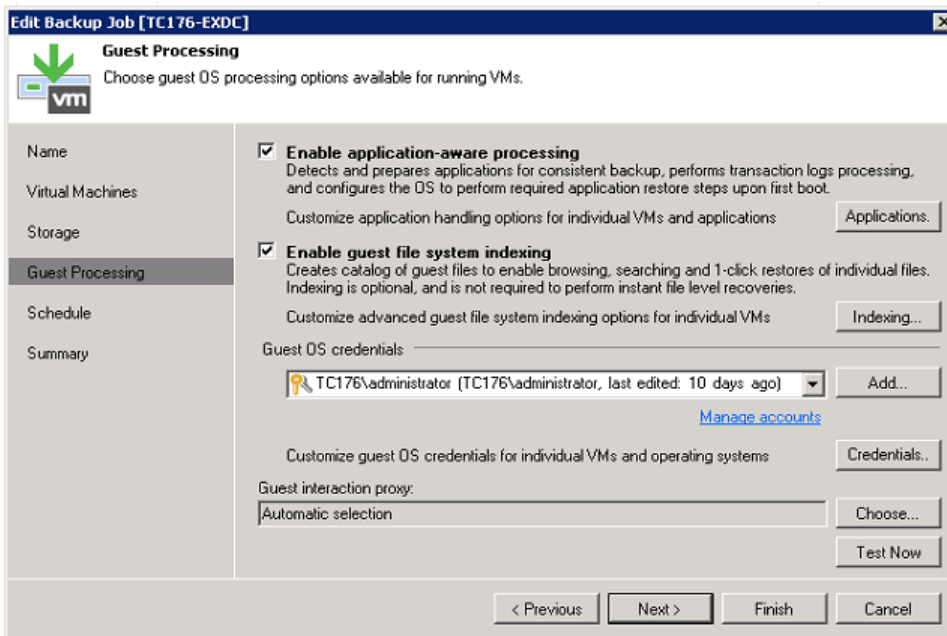
Advanced settings include:

1. Daily incremental backups with a weekly synthetic full
2. Enabling Veeam's deduplication to reduce backup traffic
3. Enabling Veeam's "dedup-friendly" compression to reduce backup traffic
4. Storage optimization of "Local Target" to take full advantage of ExaGrid's storage bandwidth



Guest processing settings include:

1. Enabling application-aware processing for the domain controller and Exchange server including application-consistent snapshots, log processing, etc.
2. Enabling guest file system indexing so multiple backups can be searched for files to restore



Summary

The powerful combination of Pure Storage all-flash storage, Veeam Backup & Replication software, and ExaGrid backup storage provides superior end-to-end reliability and performance. This integrated solution combines Pure's fast and low latency storage with Veeam's backup protection and ExaGrid's backup storage to cost-effectively deliver the fastest backups and shortest recovery times, reducing the cost of IT across primary storage and data protection tiers. This solution is particularly appropriate for the always-on, do-more-with-less modern data center that requires the highest possible performance levels to store, back up, and recover application data.

References

<https://community.purestorage.com/>

https://www.veeam.com/pdf/datasheet/veeam_top_10_reasons_choose_veeam_9_0_en.pdf

<http://www.exagrid.com/exagrid-products/product-architecture/why-exagrid-versus-traditional-inline-disk-based-backup-storage-appliances/>

<http://www.exagrid.com/exagrid-products/product-line/>

About the Author

Tom Gillispie
ExaGrid
Director, Application Interop & Product Management

Tom has 35 years of software development and engineering management experience, working in the storage and storage management field for over 16 years. Prior to his 10 years at ExaGrid, Tom lead storage product engineering teams at Revivio, Sun Microsystems, and High Ground Systems. Tom also has a background in networking protocols and product development, working in network software teams at several small and large companies. Tom has a BSEE from Carnegie-Mellon University and has privately developed and sold home and light commercial automation software and systems. When not working, Tom enjoys boating, adding to his automated home and restoring old pinball machines.

Saurabh Talwar (Sunny)
PureStorage
Solutions Architect (VMware)

Sunny has extensive experience in defining and developing technical product solutions & marketing strategy for enterprise products. Sunny has worked in release engineering, marketing and Quality Engineering departments. Sunny has worked closely with Engineering to understand Product Architecture and came from HGST (Subsidiary of Western Digital) where he was designing solutions for the SSDs. Sunny was doing performance analysis with competitor SSDs to do comparative analysis and write reports which can be given to the product marketing team. Sunny was also developing demos with real world applications like MySQL, Aerospike and showing the performance of our PCIe Cards and also answering any questions about the product at various conferences like Vmworld and FlashMemory Summit. Apart from work, Sunny likes to swim and do Crossfit.