



Reference Architecture: Advanced PureStorage FlashArray Data Protection

Using Commvault Software and

ExaGrid Backup Storage

February 2017

Contents

Executive Summary	3
Goals and Objectives.....	3
Audience	3
Pure Storage Overview.....	5
PureStorage FlashArray //m	5
PureStorage FlashArray //m Specifications	5
CommBault backup software Overview	6
Commvault IntelliSnap Technology Overview	7
ExaGrid Overview	8
Solving the Problems of Backup Permanently.....	8
Unique Architecture	8
Adaptive Deduplication.....	9
System Requirements	9
Harnessing Commvault and ExaGrid Data Protection for PureStorage FlashArray	10
Use Case 1: Scaling to a large number of concurrent, application aware, VM-consistent backups.....	11
Typical scenario: Deployment of a new Enterprise Application	11
Use Case 2: Efficient long-term backup retention	16
Use Case 3: Instantly Recovering a Virtual Machine.....	17
Use Case 4: Recover from Remote Location	19
Deployment	21
PureStorage FlashArray & Hypervisor Configuration	22
ExaGrid Deployment	28

Details on Commvault and ExaGrid Configuration31

 Creating Commvault shares 31

Summary 33

References 34

About the Author..... 34

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. Using Pure Storage, Commvault, and ExaGrid together provides the modern data center with more performance and functionality at a lower cost.

Goals and Objectives

Provide technical information about Pure Storage, Commvault backup software, and ExaGrid backup storage appliances:

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

Audience

This paper is written for storage and backup administrators who are interested in understanding how together, Pure Storage Flash, Commvault backup software, and ExaGrid backup storage appliances combine fast, reliable, and cost-effective primary storage along with fast backups to shorten recovery times.

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.

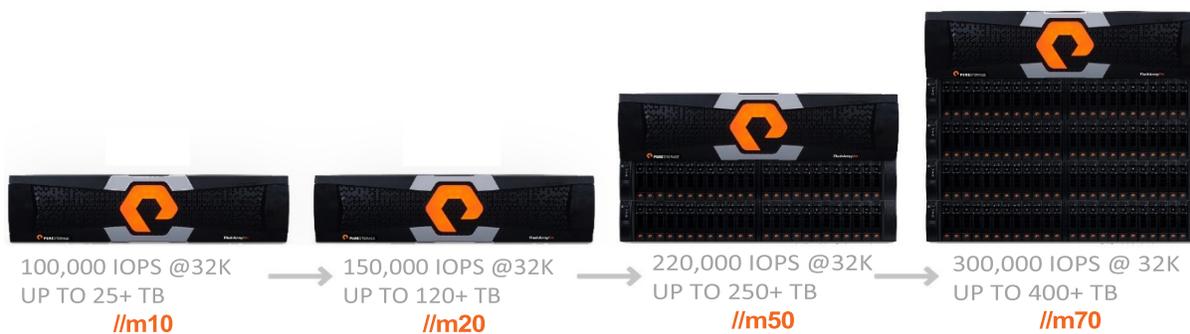
Built on 100% MLC flash, Pure Storage FlashArray //m delivers all-flash enterprise storage that is 10X faster, more space and power efficient, more reliable, and infinitely simpler, and yet typically costs less than traditional performance disk arrays.

PureStorage 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.



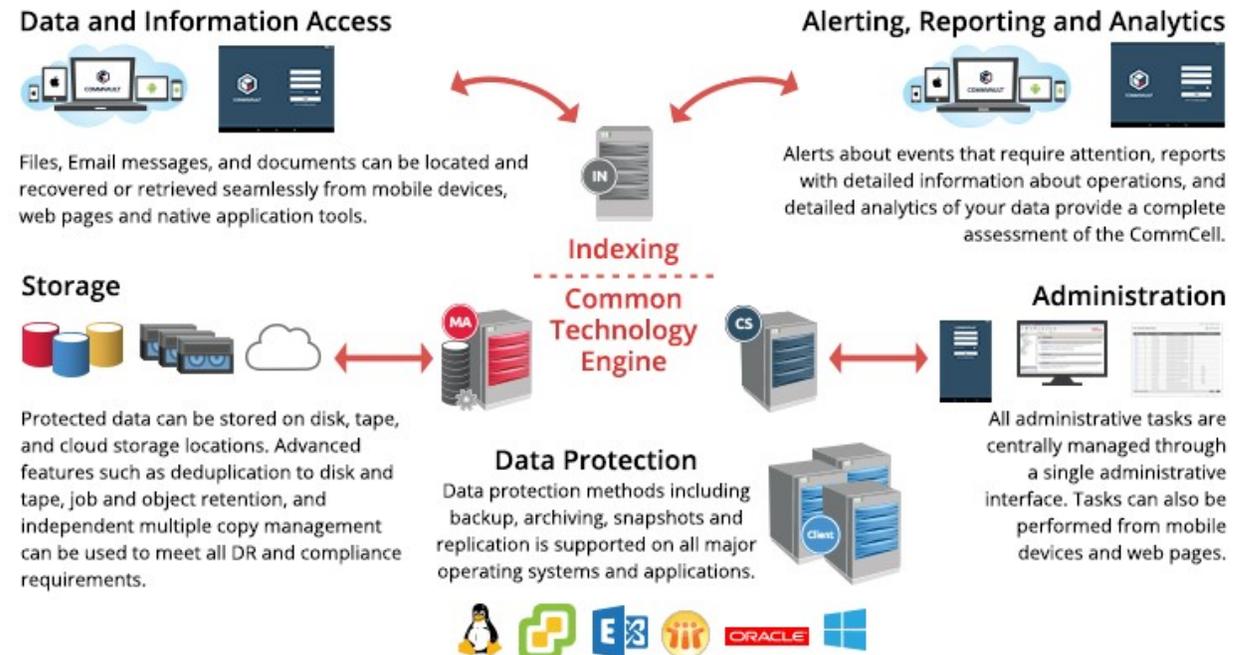
PureStorage FlashArray //m Specifications

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

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

Commvault Backup Software Overview

The Commvault software platform is an enterprise level, integrated data and information management solution, built from the ground up on a single platform and unified code base. All functions share the same back-end technologies to deliver the unparalleled advantages and benefits of a truly holistic approach to protecting, managing, and accessing data. The Commvault software platform contains modules to protect and archive, analyze, replicate, and search your data, which all share a common set of back-end services and advanced capabilities, seamlessly interacting with one another. The Commvault software platform addresses all aspects of data management in the enterprise, while providing infinite scalability and unprecedented control of data and information.



Commvault IntelliSnap Technology Overview

Commvault IntelliSnap technology integrates with native storage array snapshot engines to provide consistent point-in-time recovery copies for large data sets and enterprise applications. IntelliSnap technology quiesces applications or file systems, triggers the storage array-based snapshot, and returns the system to a fully operational state within minutes. By incorporating and linking snapshots with backup and archive operations, software makes more online and offline copies available for recovery while reducing data protection's impact on production systems. IntelliSnap technology harnesses the power of array-based snapshots to accelerate backup and recovery.



Commvault IntelliSnap technology supports Pure Storage arrays – for more details, see <https://www.purestorage.com/resources/type-a/whitepaper-on-intellisnap-on-pure-storage-for-vmware.html>

Combining Pure Storage, Commvault backup software with Intellisnap, and ExaGrid's backup storage provides a complete end-to-end solution for modern data protection requirements.

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 this 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 200TB/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 and ensuring the highest possible backup performance, which results 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 unduplicated form for fast restores, VM Instant Recoveries, and tape copies while the offsite data is ready for disaster recovery.

ExaGrid and Commvault 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 Commvault and ExaGrid technology integration with the Pure Storage FlashArray:

- FlashArray 400 series or FlashArray//m
- Purity v4.1.1 or higher
- Commvault v10 or later
- ExaGrid v4.8.0 or later

Harnessing Commvault and ExaGrid Data Protection for PureStorage FlashArray

Commvault 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. Instantly recover a VM
4. Recover from remote location

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.

In all cases, Commvault IntelliSnap technology, which supports Pure Storage arrays, harnesses the power of array-based snapshots to accelerate backup and recovery. For more details, see <https://www.purestorage.com/resources/type-a/whitepaper-on-intellisnap-on-pure-storage-for-vmware.html>

Use Case 1: Scaling to a large number 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

Commvault backup software scales out by supporting multiple media agents, disk libraries, and concurrent backup job execution to meet the data protection needs of today's enterprise data centers. Documented best practices explain how to configure both Commvault and ExaGrid for a high degree of concurrent backup operations that scale as backup data grows.

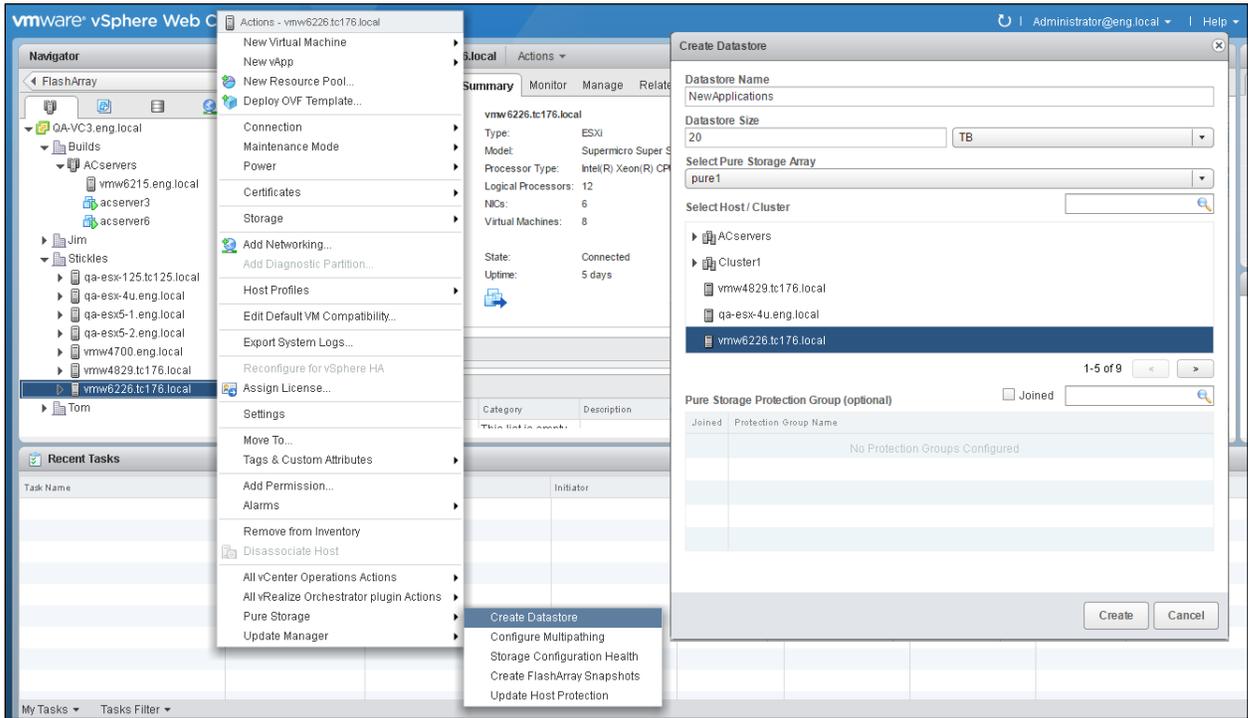
Scalable Backup Storage

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.

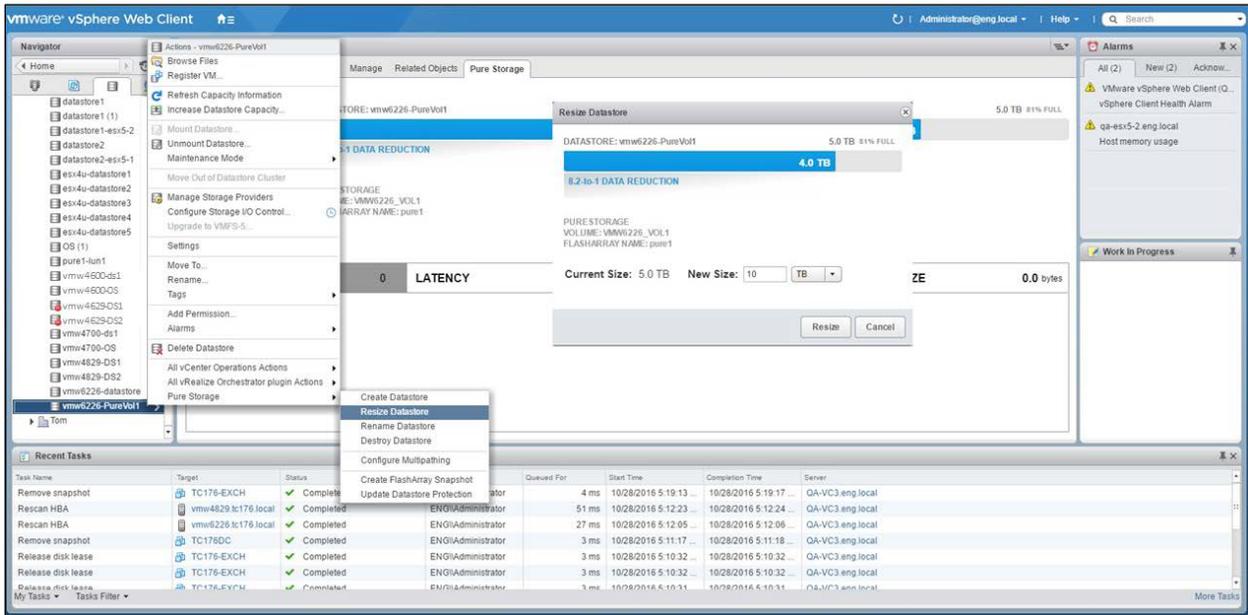
Typical Scenario: Deployment of a New Enterprise Application

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 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 of the size and retention of this new enterprise application, an additional ExaGrid EX32000E appliance with backup storage capacity for 32TB of weekly full backup is purchased and added to the existing ExaGrid site. This new ExaGrid appliance is configured per best practices as a Commvault shared disk library.
5. New Commvault jobs are created to cover the data protection needs of the new VMs and are targeted at the newly configured disk library. These jobs will utilize IntelliSnap technology to harness the power of array-based snapshots to accelerate backup and recovery.

See [Details on Commvault and ExaGrid Configuration](#) below for details on the steps to create a Commvault share on the new ExaGrid appliance and adding it to the Commvault configuration.

After adding the additional ExaGrid EX32000E appliance, the total capacity of the Boston site has increased. All ExaGrid appliances are managed via a single pane of glass.

The screenshot displays the ExaGrid management interface for the BostonHub site. It includes a navigation tree on the left and three main data sections:

Summary for Site : BostonHub

Report Generated: 31 Jan 2017 09:36:28 GMT-0500 [?] Export

Landing Space for Site : BostonHub

Server	Available			Deduplication Status
	Actual	% Available for Next Backup	Maximum	
BostonHub	96,000.00 GB	100%	96,000.00 GB	Complete
Boston1	32,000.00 GB	100%	32,000.00 GB	Complete
Boston2	32,000.00 GB	100%	32,000.00 GB	Complete
Boston3	32,000.00 GB	100%	32,000.00 GB	Complete

Retention Space for Site : BostonHub

Server	Available	Consumed	Shared	Retention Space		Total Capacity	Over Capacity	Replication Status
				% Available				
BostonHub	86,619.29 GB	5,776.33 GB	5%	94%		92,395.62 GB	0.00 GB	Complete
Boston1	29,081.24 GB	1,717.30 GB	3%	94%		30,789.54 GB	0.00 GB	Complete
Boston2	28,251.80 GB	2,546.74 GB	3%	94%		30,798.54 GB	0.00 GB	Complete
Boston3	29,286.25 GB	1,512.29 GB	9%	94%		30,798.54 GB	0.00 GB	Complete

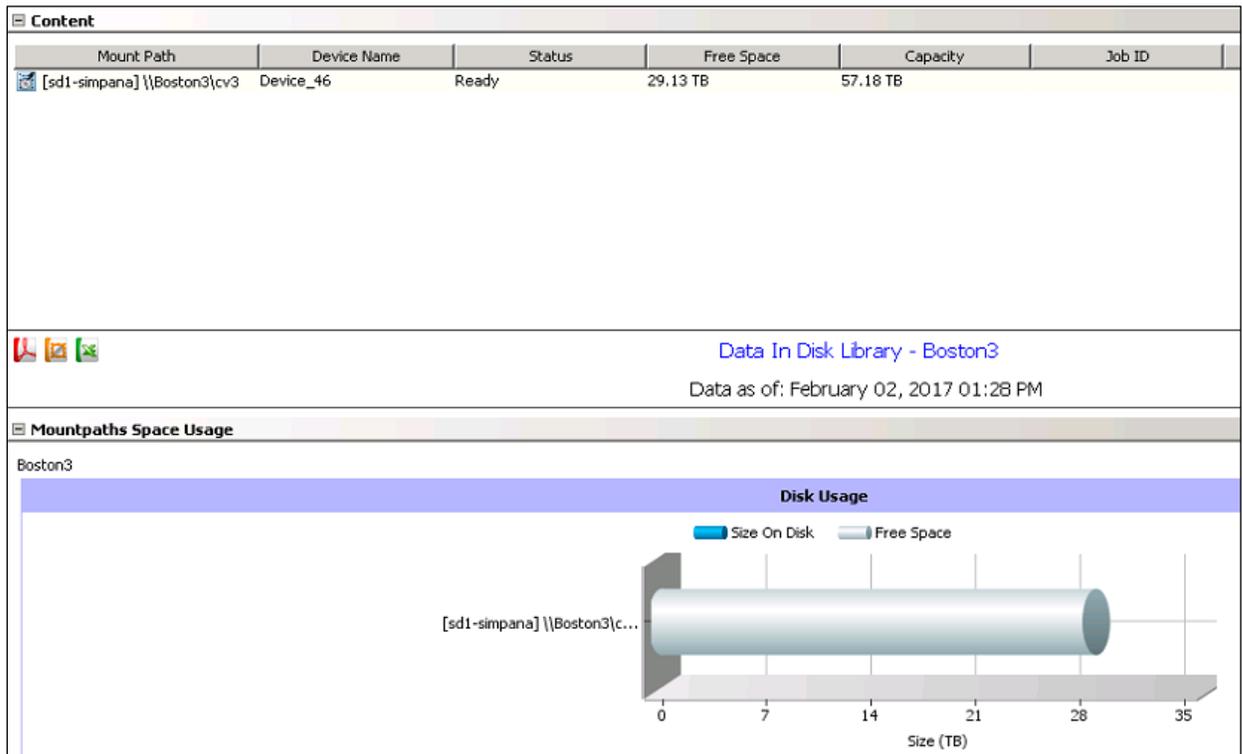
Deduplication Summary for Site : BostonHub

Share	Dedupli				
BostonHub					
Boston1					
cv1	✓ 12.6 : 1	26,082.82 GB	1,875.02 GB	1,660.80 GB	
Boston2					
cv2					
Boston3					
cv3	✓ 16.6 : 1	26,269.09 GB	1,582.79 GB	1,376.87 GB	

Annotations in the screenshot point to the deduplication summary table:

- ExaGrid appliance 1 backup storage (points to cv1)
- ExaGrid appliance 2 backup storage (points to cv2)
- ExaGrid appliance 3 backup storage (points to cv3)

And, after adding the ExaGrid share from the newly-added ExaGrid EX32000E appliance, Commvault is able to utilize the now-increased capacity of the Boston site.



Use Case 2: Efficient long-term backup retention

Backup retention periods are increasing, not shrinking.

Without deduplication, longer retention periods quickly lead to rapid expansion of backup storage needs.

Achieving high levels of deduplication is computationally expensive, so the trade-off between higher dedup (requiring more CPU/memory) vs. lower dedup (requiring less CPU/memory) is an especially important consideration when longer backup retention requirements exist.

ExaGrid's unique scale-out architecture brings just the right amount of CPU, memory, network and backup storage to achieve that high level of deduplication. Commvault's software-based backup deduplication produces a lower level of deduplication because a higher level of software-based backup deduplication would consume too much backup server/media agent CPU, memory and network.

To achieve the most efficient long-term backup retention, CommVault is deployed without software deduplication, allowing the ExaGrid appliances to reach the highest level of deduplication.

Retention Space for Site : <u>BostonHub</u>							
Server	Retention Space			% Available	Total Capacity	Over Capacity	Replication Status
	Available	Consumed	Shared				
<u>BostonHub</u>	86,619.29 GB	5,776.33 GB	5%	<div style="width: 100%; height: 10px; background-color: green;"></div> 100%	92,395.62 GB	0.00 GB	Complete
<u>Boston1</u>	29,081.24 GB	1,717.30 GB	3%	<div style="width: 100%; height: 10px; background-color: green;"></div> 100%	30,789.54 GB	0.00 GB	Complete
<u>Boston2</u>	28,251.80 GB	2,546.74 GB	3%	<div style="width: 100%; height: 10px; background-color: green;"></div> 100%	30,798.54 GB	0.00 GB	Complete
<u>Boston3</u>	29,286.25 GB	1,512.29 GB	9%	<div style="width: 100%; height: 10px; background-color: green;"></div> 100%	30,798.54 GB	0.00 GB	Complete

ExaGrid performs deduplication across all shares

Deduplication Summary for Site : <u>BostonHub</u>				
Share	Deduplication Ratio	Total Backup Data	Compressed Size	Distinct Size
<u>BostonHub</u>	✓ 14.96 : 1	86,421.03 GB	5,776.33 GB	5,497.55 GB
<u>Boston1</u>				
<u>cv1</u>	✓ 13.91 : 1	26,083.82 GB	1,875.92 GB	1,660.80 GB
<u>Boston2</u>				
<u>cv2</u>	✓ 12.52 : 1	34,068.11 GB	2,721.08 GB	2,459.88 GB
<u>Boston3</u>				
<u>cv3</u>	✓ 16.6 : 1	26,269.09 GB	1,582.79 GB	1,376.87 GB

Overall deduplication 10 to 50:1

Requires significantly less backup storage for long-term retention.

Use Case 3: 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 Commvault's Live 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 Commvault Live VM Recovery, there is no need to rehydrate deduplicated backups since they are available directly in the landing zone of each ExaGrid appliance.

During the Commvault Live VM Recovery operations, writes done booting/running VM are redirected 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. Commvault's Live VM Recovery can start this migration immediately, or after a time delay. During the migration, ExaGrid's landing zone facilitates high-bandwidth data transfers to high-performance production Pure Storage FlashArrays, assuring the restore to production environment completes as quickly as possible.

Restore Options for All Selected Items

General | Job Initiation

vCenter Client: sd-vcenter Proxy Client: sd1-simpana

vCenter Instance Details
vCenter: sd-vcenter.eng.local
User: administrator@vSphere.local

Restore in place

VM and Disk	Change VM display name to	ESX Server	Datastore	vCenter Options
<input checked="" type="checkbox"/> esg-small	esg-small-live	172.21.12.64	qa-esx2:storage6	Configure
- esg-small_2.vmdk			qa-esx2:storage6	
- esg-small_1.vmdk			qa-esx2:storage6	

Virtual Machine Configuration

Power ON Virtual Machine during restore

Unconditionally overwrite VM with the same name

Disk Provisioning: Original

Transport Mode: Auto

Restore Virtual Machine using Live Recovery (vMotion)

Redirect Writes to Datastore: pure-lun1

Delay migration: 3 hr(s)

OK Cancel Advanced Save As Script Help

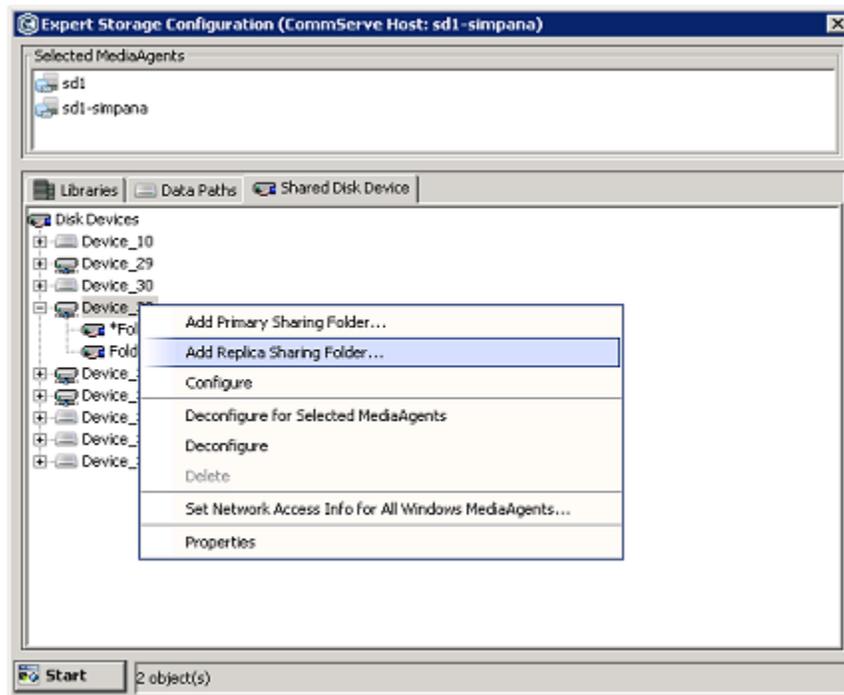


Total RTO – Recovery Time Objective (RTO) is targeted duration of time and a service level within which a running Virtual Machine can be restored from a backup. Total RTO can be drastically reduced due to ExaGrid's landing zone architecture and Pure Storage high-performance FlashArray storage where all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

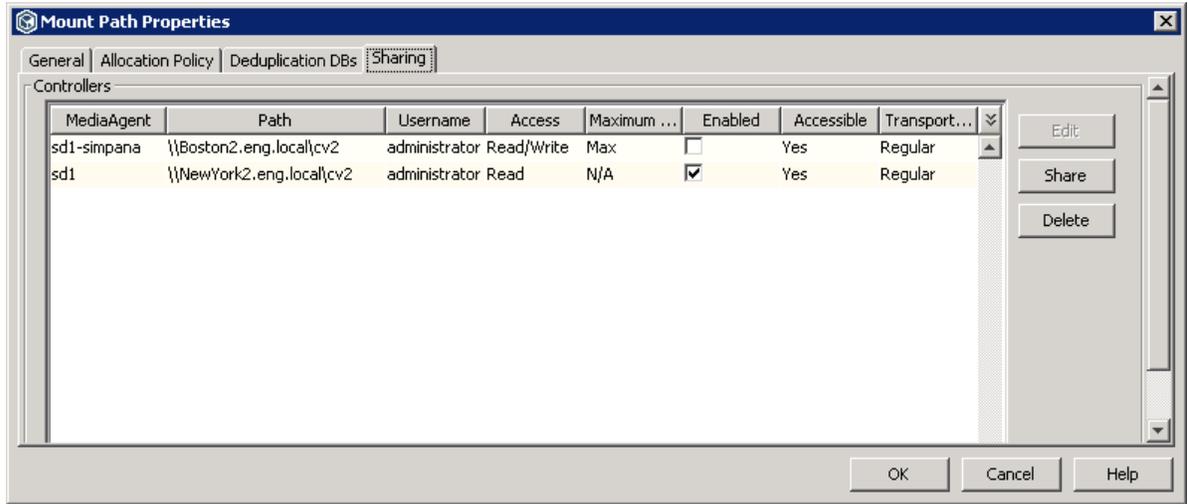
Use Case 4: Recover from Remote Location

When configuring the Commvault disk devices, the ExaGrid InstantDR share at a remote (DR) site is added to the Commvault configuration using a media agent located at the remote site. The InstantDR share is continually updated with the deduplicated backups, ready for restores should the primary site fail or for DR rehearsals and audits. Normally, backups and restores use the primary path to the ExaGrid primary site share. Should the primary path be lost, a restore can be done using the secondary path to the remote site ExaGrid. There is no need to discover the ExaGrid remote site and no time-consuming cataloging operations. DR rehearsal and audits are done by disabling the primary path in Commvault, and doing normal Commvault restores – full, granular, etc. – using the media agent at the remote site.

Here, the remote (DR) site’s ExaGrid InstantDR share is added to Commvault as a “Replica Sharing Folder,” creating a secondary, read-only path to the backup data at the remote (DR) site.



During the restore of a client, the alternate media agent and secondary path are selected by Commvault.



Deployment

The deployments of the PureStorage FlashArray 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.
- 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 Commvault share per appliance.
- The ExaGrid shares are configured as Commvault disk libraries, and Commvault client policies/jobs are modified or created to utilize the ExaGrid disk library.

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

PureStorage FlashArray and 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 Limit to 1 which is a command line option only.

The screenshot displays the VMware vSphere Web Client interface. The left-hand navigation pane shows a list of hosts, with host 10.21.23.20 selected (callout 1). The main content area shows the 'Storage' tab for this host, with 'Storage Devices' selected (callout 2). A table of storage devices is visible, with the 'PURE Fibre Channel Disk...' device selected (callout 3). The 'Device Details' section shows the selected device's properties (callout 4). A modal dialog titled '10.21.23.20 - Edit Multipathing Policies for naa.624a9370753d69fe46db318d00010001' is open, showing the 'Path selection policy' set to 'Round Robin (VMware)' (callout 6). The dialog also displays a table of active paths and a 'Preferred' column. A blue arrow points to the 'Edit Multipathing...' button (callout 5).

Name	Type	Capacity	Operation...	Hardware Acc...	Drive Ty...	Transport
PURE Fibre Channel Disk...	disk	2.00 ...	Attached	Supported	SSD	Fibre C...
Local HGST Disk (naa.50...	disk	1.09 ...	Attached	Unknown	Non-S...	Parallel ...
Local DP Enclosure Svc D...	enc...		Attached	Not suppor...	Non-S...	Parallel ...
PURE Fibre Channel Disk...	disk	2.00 ...	Attached	Supported	SSD	Fibre C...

Runtime Name	Status	Target	LUN	Preferred
vmhba5:C0:T5:L11	Active (IO)	52:4a:93:7f:3d:cc:79:15:52:4a:93:7f:3...	11	
vmhba5:C0:T4:L11	Active (IO)	52:4a:93:7f:3d:cc:79:05:52:4a:93:7f:3...	11	
vmhba3:C0:T5:L11	Active (IO)	52:4a:93:7f:3d:cc:79:14:52:4a:93:7f:3...	11	
vmhba3:C0:T4:L11	Active (IO)	52:4a:93:7f:3d:cc:79:04:52:4a:93:7f:3...	11	
vmhba1:C0:T5:L11	Active (IO)	52:4a:93:7f:3d:cc:79:15:52:4a:93:7f:3...	11	
vmhba1:C0:T4:L11	Active (IO)	52:4a:93:7f:3d:cc:79:05:52:4a:93:7f:3...	11	

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> -l 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.

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

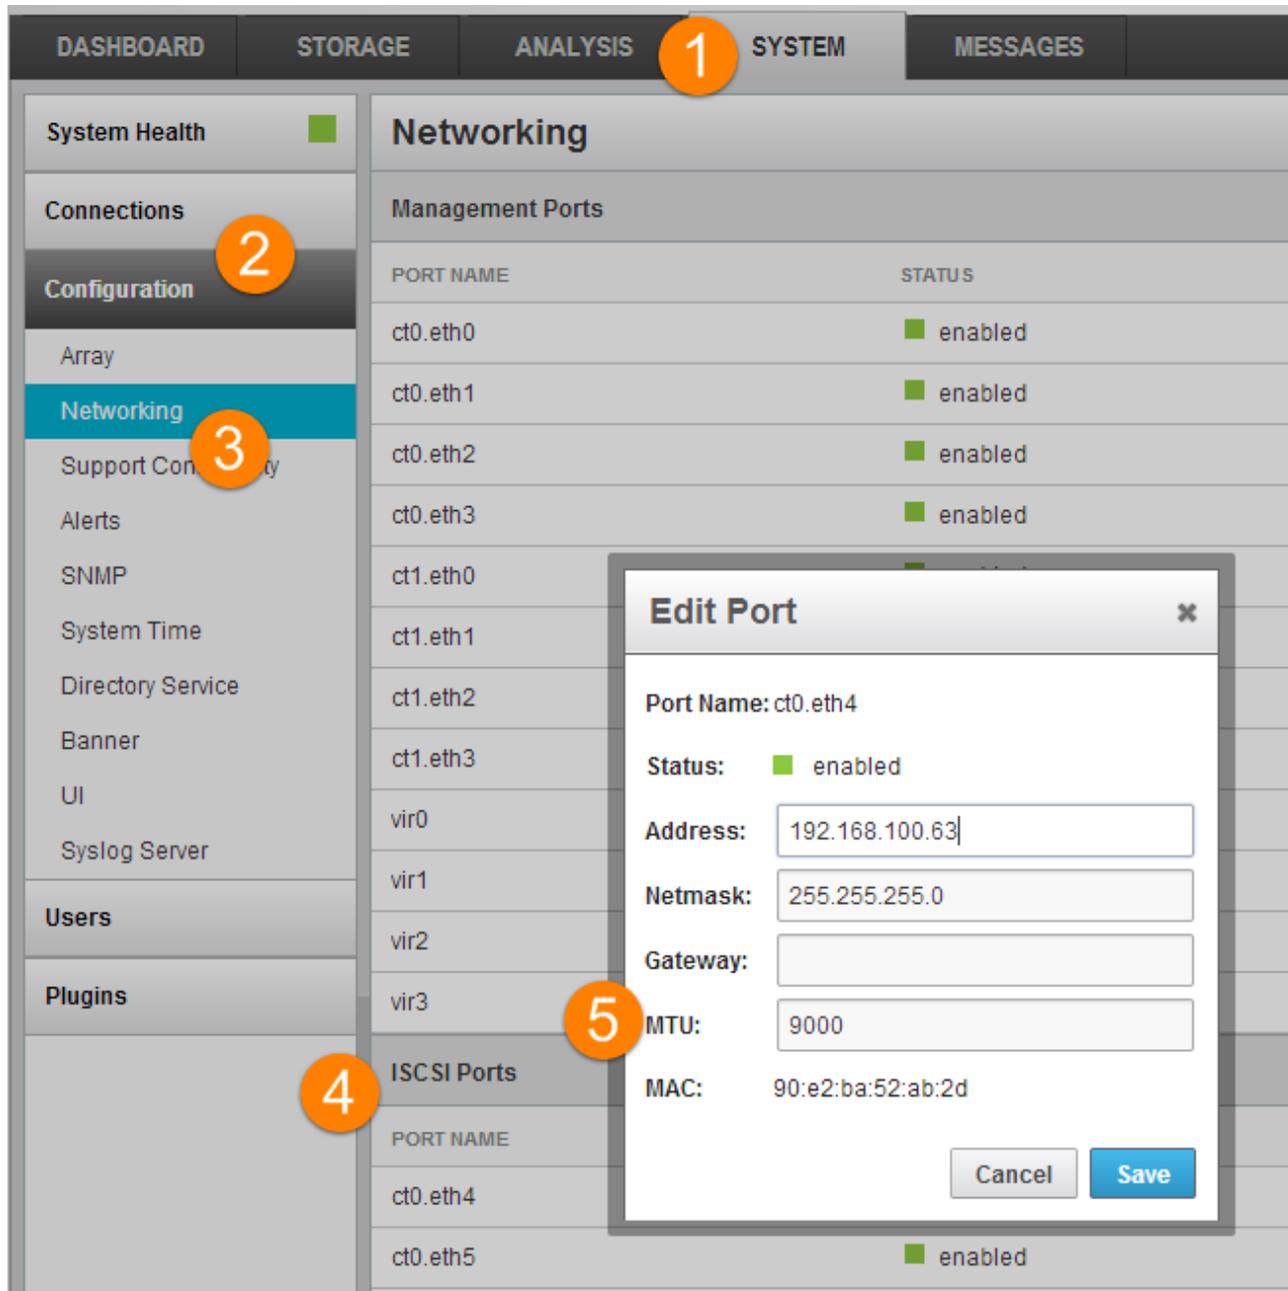


Figure: Configuring Jumbo Frames on the Pure Storage FlashArray

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.3efb9d4xxxxxxxxx97
```

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:

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	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

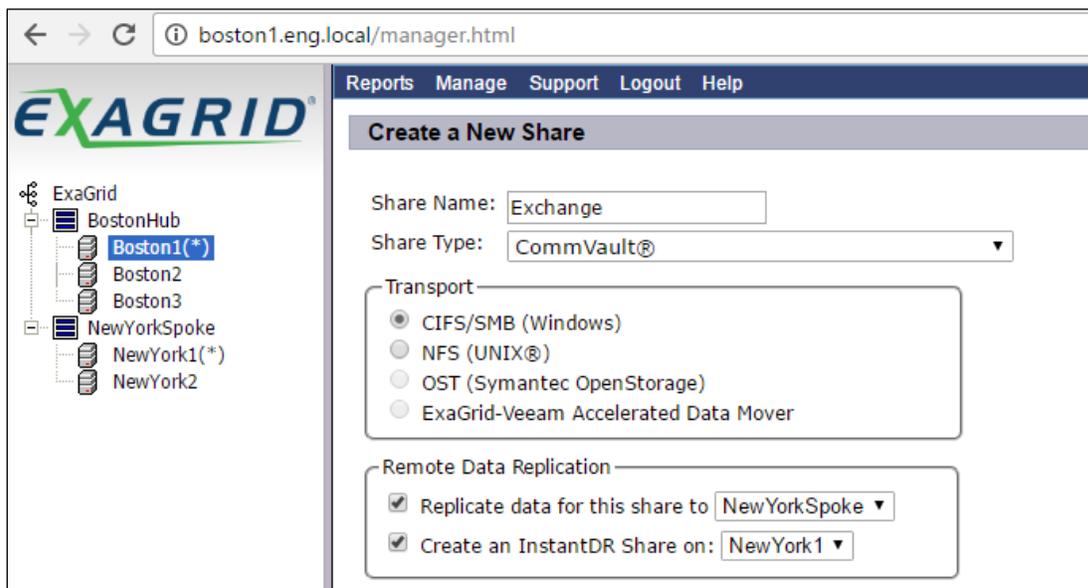
Server	Available	Consumed	Shared	Retention Space		Total Capacity	Over Capacity	Replication Status
				% Available				
TC176S1	86,619.29 GB	5,776.33 GB	5%	94%		92,395.62 GB	0.00 GB	Complete
TC176SP1OB1	29,081.24 GB	1,717.30 GB	3%	94%		30,798.54 GB	0.00 GB	Complete
TC176SP1OB2	28,251.80 GB	2,546.74 GB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

Details on Commvault and ExaGrid Configuration

The screen shots below show more details on the Commvault and ExaGrid configuration. ExaGrid provides best practices for this configuration, and Commvault has several knowledgebase articles covering best practices.

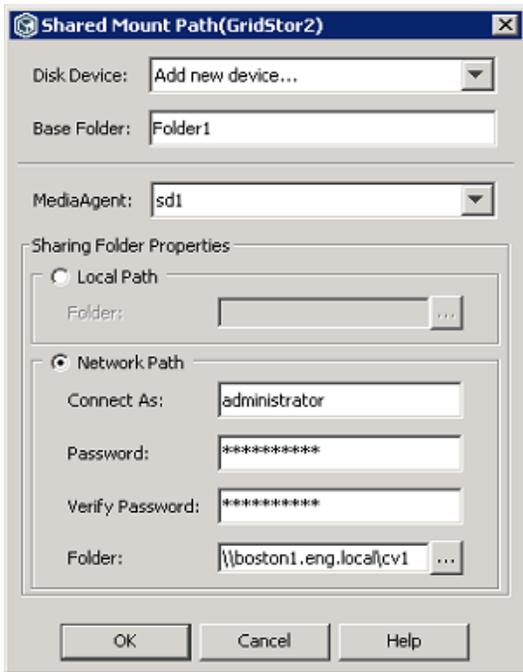
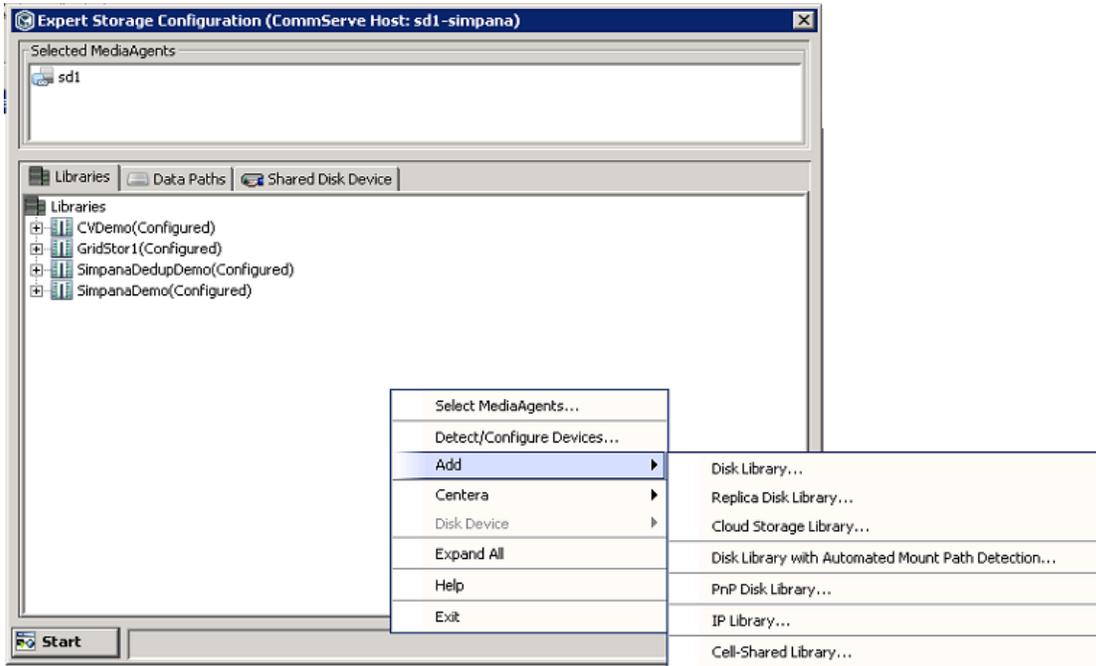
Creating Commvault Shares

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



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

Each ExaGrid Commvault share is added to Commvault as a Disk Library.



Summary

The powerful combination of Pure Storage all-flash storage, Commvault backup 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 Commvault'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/>

<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 Authors

Tom Gillispie

Director, Application Interop & Product Management

ExaGrid

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)

Solutions Architect (VMware)

PureStorage

Sunny has extensive experience in defining and developing technical product solutions and 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 (a 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 that can be given to the product marketing team. Sunny was also developing demos with real-world applications like MySQL and Aerospike, and showing the performance of PCIe Cards and answering any questions about the product at various conferences like VMworld and FlashMemory Summit. Apart from work, Sunny likes to swim and do Crossfit.