

TECHNICAL WHITE PAPER

# Best Practices: FlashArray and Exchange Server

Best practices for Pure Storage® FlashArray with Microsoft Exchange Server 2019.

# Contents

- Introduction .....3**
- Microsoft Exchange on FlashArray .....3**
  - Microsoft Exchange Server 2019 ..... 3
  - Exchange Server Licensing ..... 4
  - Database Availability Groups ..... 4
  - Microsoft Failover Clustering ..... 4
- Sizing Considerations .....4**
  - Storage IOPS ..... 5
  - Database Size ..... 5
  - Pure Data Reduction Rates ..... 5
  - Thin Provisioning ..... 6
  - Filesystem and Allocation Unit Size ..... 6
  - LUN Isolation ..... 6
- Data Protection ..... 7**
  - Full Database Isolation ..... 7
  - VM LUN Consolidation ..... 7
  - Pure Volume Layout Recommendation ..... 7
  - Snapshots and Clones ..... 8
- Disaster Recovery ..... 9**
  - Flash-to-Flash-to-Cloud ..... 9
  - Storage Migration ..... 9
- VMWare Best Practices ..... 10**
  - Compute and Memory ..... 10
  - SCSI Controller ..... 10
  - Network ..... 10
  - Disk Configuration ..... 10
  - Anti-Affinity ..... 10
- Hyper-V Best Practices .....11**
  - Highly Available Virtual Machine ..... 11
  - Anti-affinity ..... 11
- FlashStack and Exchange Server .....11**
- Jetstress Testing ..... 12**
- Microsoft Integration ..... 13**
  - SMI-S ..... 13
  - System Center Virtual Machine Manager (SCVMM) ..... 13
  - System Center Operations Manager (SCOM) ..... 14
- Troubleshooting ..... 14**
- Conclusion ..... 14**
- For More Information ..... 15**



## Introduction

Pure Storage® FlashArray//X is the world's first enterprise-class, all-NVMe flash block storage array. It represents a new class of storage—Shared Accelerated Storage (a term coined by Gartner)—that delivers breakthroughs in performance, simplicity, and consolidation of mixed workloads. With latency as low as 200µs, FlashArray//X's end-to-end NVMe architecture brings new levels of performance to mission-critical business applications as well as cloud-native, web-scale applications. FlashArray's Purity Reduce capabilities include compression and deduplication, offering on average 5:1 data reduction and 10:1 total efficiency. With built-in Purity ActiveCluster™, more applications can now benefit from the simple to deploy, always-on business continuity provided by active/active metro stretch clustering. Pure DirectFlash® with 100% NVMe also enables unprecedented performance density—the kind of density required for mixed-workload consolidation in your cloud. FlashArray//X currently supports ultra-dense 18.3TB DirectFlash modules, for maximum density and full performance. Also, Purity's QoS features mean you can consolidate radically diverse applications without fear of I/O contention.

Organizations have evolved to run a mix of classic business apps and new, modern web-scale apps. Previously these two worlds necessitated radically different architectures, but with FlashArray//X and NVMe, everything can run on a single shared accelerated storage architecture.



---

## Microsoft Exchange on FlashArray

Pure Storage FlashArray™ eliminates IO bottlenecks while significantly reducing latency to 1ms and below. Depending on the data protection architecture, the number of storage objects to manage can be as few as one global datastore or cluster shared volume, to a handful of volumes to manage per Exchange Server. By reducing complexity, amplifying performance, and seamlessly integrating with Microsoft tools, Exchange Administrators are empowered to reduce costs in planning and management while increasing availability.

### Microsoft Exchange Server 2019

Microsoft Exchange Server 2019 is the latest release from Microsoft and the third release following Exchange Server 2013 utilizing managed code in an instanced information store. The new features can be found in Microsoft's [What's New in Exchange Server](#) article. Some highlights that impact storage include:



- A new, faster search engine that is now more scalable, asynchronous, and decentralized.
- A metacache database that can take advantage of solid-state performance.
- Consolidation of the Client Access Server and Mailbox Server roles into the new Mailbox Server Role.
- One Drive for Business & SharePoint Server 2019 integration for linking to files in messages.

## Exchange Server Licensing

Microsoft licenses two versions of Exchange Server; Standard, which is limited to 5 databases, and Enterprise, which is limited to 100 databases. In addition to the Exchange Server license, both Windows and a client access license for Exchange and Windows is needed. For details on licensing see [Exchange Server 2019 licensing](#). Larger databases, with more users per server, and larger physical or virtual Exchange Servers can reduce licensing costs.

## Database Availability Groups

Database Availability Groups (DAGs), first implemented in Exchange 2010 and built on the log shipping in Exchange 2007's continuous replication, involves up to 15 passive copies of a database that are behind the active copy by part of a log file to many log files if replication is running behind. These databases can stretch two sites and be configured with an activation preference. Some of the passive copies can be configured in a lagged state, where a delayed log replay time is configured.

Since Exchange is unaware of hypervisor availability features, it is important to configure anti-affinity when virtualizing Exchange. Anti-affinity can be used with careful database placement to prevent all copies of a single database from being collocated on the same physical hypervisor. See the VMware Anti-Affinity or Hyper-V Anti-Affinity section for details on configuring it.

## Microsoft Failover Clustering

A DAG leverages Microsoft Failover Clustering, to decide if nodes are healthy and whether a database failover should occur where a passive copy of a database is mounted, becoming the active copy. The new cloud witness feature in Windows 2016 is not supported by Exchange, and Exchange 2019 itself is only supported on Windows 2019. Both Exchange 2013 and Exchange 2016 are supported on older versions of Windows, and there is a cluster timeout vulnerability during network issues that last longer than five seconds in Windows Server 2012R2 that should be fixed when running that version of the operating system. These recommended changes were made the default in Windows Server 2016. Elden Christensen outlines this in the following blog post that is also referenced in VMware's best practices guide for Exchange Server 2016, [Tuning Failover Cluster Network Thresholds](#). VMware's guidance for Windows Server 2019 is to leave the default thresholds in place.

## Sizing Considerations

Microsoft's [Exchange Server Role Requirements Calculator](#) should be used to size an Exchange 2013, 2016, or 2019 deployment. In Exchange Server 2013-2016, the maximum number of processor cores is 24 and the maximum amount of system RAM is 192GB. Exchange Server 2019 has increased these maximums to 48 processor cores and 256GB of RAM. The Microsoft Exchange Role Requirements Calculator tends to be conservative, choosing a scale-out rather than scale-up strategy which can be costly in Exchange Server Enterprise licenses and more Exchange Servers to manage.



## Storage IOPS

IO for Exchange has been steadily decreasing while storage performance has increased with flash by several orders of magnitude. The IOPS of passive copies are 50% of the active database, and maintenance is 1-2MB per second per database in Exchange 2019.

The IO of Exchange is:

- Transaction Logs—synchronous write—512 bytes to 1MB (Log Buffer Size)
- DB IO—random—32KB pages typically average 40-45KB
- DB Maintenance—256KB sequential read—1 to 2MB per database

Latency requirements:

- Transaction Logs—<10ms
- DB Reads—<20ms
- DB Writes—none but keep DB Page Fault Stall average <1
- DB Maintenance—none

**Best Practice:** Fewer, larger databases will reduce concurrent maintenance IO.

## Database Size

A 2TB database size maximum recommendation in Exchange 2010 is unchanged a decade later. In the Exchange 2010 era, partition sizes in Windows and VMware did not allow larger than 2TB, the first 2TB hard disk started shipping, and 1GbE was the limit. Today's storage sizes are significantly larger with VMware adding support for larger than 2TB in ESXi 5.5, Microsoft started supporting GPT partitions in Windows Server 2003SP1, and most hypervisors are equipped with 10GbE and beyond which reduces database seeding time.

The database size limit should be predicated on a data protection strategy and the business service level agreement (SLA). For details see the Data Protection section.

## Pure Data Reduction Rates

Most Exchange data is in the database itself, or the EDB file. Pure Storage's Purity Storage Operating System invokes many levels of data reduction, is always on, with no configuration required. Most Exchange deployments on FlashArray see between 2.5:1 and 3:1 data reduction where more than one copy of the databases are hosted on the same FlashArray. Designs that house a single copy of a database on a FlashArray will see less data reduction. This might be required for no single point of failure at the rack, row, or building layer. Designs that exclude highly reducible files will see less data reduction. These files include the operating system, Exchange binaries, Exchange Service logs, transaction logs, restore LUNs, and the transport database and its transaction logs. A database that is seeded will have a very high data reduction rate, that over time will lower due to the unique header information written to each database page. Many customers have had great success maintaining great data reduction rates by periodically (once or twice a year), refreshing database copies from a recent [Volume Shadow-Copy Service](#) (VSS) snapshot. VSS is a framework that consists of an application writer, a storage provider, and a requestor enabling application-consistent snapshots without impacting availability.



## Thin Provisioning

The FlashArray architecture is thinly provisioned at the Pure Volume level and it is trivial and virtually instantaneous to increase the size of a Pure Volume even when more flash is added to an array. From a VMware perspective, consider **eagerzeroedthick** when provisioning disks for virtual machines.

## Filesystem and Allocation Unit Size

[NTFS](#) is required for the system partition, Exchange binaries, and the Transport Database. Only a mailbox database and its corresponding transaction logs can use the Resilient File System ([ReFS](#)) with Exchange 2013+. ReFS is Microsoft's newest filesystem that adds features designed to detect and fix corruption.

**Best Practice:** Use NTFS with a 64KB allocation unit size with support for large file record segments, even though ReFS is supported. NTFS has an over 26-year history of reliability and compatibility.

```
Format-Volume -DriveLetter D -FileSystem NTFS -AllocationUnitSize 64KB -UseLargeFRS
```

If electing to use ReFS, ensure the following:

- ReFS must be declared on the Database Availability Group (DAG) with the following command:  

```
Set-DatabaseAvailabilityGroup -Identity DAGName -Filesystem ReFS
```
- Alternatively, it can be specified during DAG creation:  

```
New-DatabaseAvailabilityGroup -Name QDAG -Filesystem ReFS
```
- ReFS data integrity features should be disabled for the EDB files or volume. It is disabled by default and can be disabled if accidentally enabled:  

```
Get-Item 'p:\db\pdb.edb' | Get-FileIntegrity Set-FileIntegrity P:\ -Enable $False
```
- Allocation Unit size should be set to 64KB regardless of partition type for volumes that contain databases and transaction logs. To format a volume for ReFS with a 64KB AUS:  

```
Format-Volume -DriveLetter P -FileSystem ReFS -AllocationUnitSize 65536  
-SetIntegrityStreams $false
```

## LUN Isolation

Most legacy concerns with isolating everything came down to performance. A spinning disk was more performant if you separated sequential from random IO, and isolating the operating system and paging file increased performance and made server recovery faster. There isn't a performance consideration for isolation on FlashArray, because Pure FlashArray is significantly faster than hard disk and hybrid solutions. Everything revolves around data protection. Most deployments consolidate the operating system, Exchange binaries, paging file, and Exchange transport database to the C: drive, and move the production user mailbox database to other volumes. Some data protection strategies recommend isolating the paging file so that it will not be part of the backup or replication set.



## Data Protection

Data Protection, and not performance, is the primary factor to consider when architecting storage for Exchange. While fewer storage objects are preferred, weigh whether restoring a database LUN directly to the original location is required, or if exposing and copying files that leverage [offloaded data transfer](#) (ODX) is fast enough.

### Full Database Isolation

The most complex scenario involves isolating each database and each corresponding transaction log set to a separate LUN in windows, meaning that for every database you will add two LUNs. This scenario is most common when utilizing the Pure Storage Hardware VSS Provider, where the administrator wants the ability to perform a VSS restore of the database LUN to the original location, and then control roll-forward recovery of the transaction logs. In the case of a database copy DAG with 200 unique databases, there would be 800 database copies, and 1600 distinct LUNs at the storage layer or Pure Volumes spread between two data centers.

### VM LUN Consolidation

The opposite extreme would be to have only one LUN per physical Exchange Server, or one LUN exposed to the hypervisor, and utilize virtual hard disk files for the Exchange Server virtual machines. The same 800 database copies could be placed on just two LUNs, one in each datacenter. This is more common with small Exchange deployments with less than five Exchange Servers. The good news is that Pure supports either extreme, yet in the next section a case is made for moderation. It is not recommended to consolidate everything onto a single LUN because there should not be a single point of failure. If a thin-provisioned Pure Volume is created with a fixed size, and runs out of space due to poor monitoring, no further writes will be accepted, causing any databases residing on that volume to go offline.

**Best Practice:** Ensure copies of the same database are placed on separate Pure Volumes.

### Pure Volume Layout Recommendation

It is recommended to have a database layout that can best be leveraged by the data protection strategy which can be influenced by the following:

- Some backup applications can only backup the entire server or virtual machine. Ensure the database copy spread is such that every server does not need to be fully backed up, since you can only restore one copy of the database and reseed—either with Exchange through the network, or with Pure and cloning.
- Backing up every single Exchange Server will cause a complex restore scenario where each database copy will need to be checked for which backup was the most recent to prevent data loss (recovery point objective, or RPO) and can substantially add to the recovery time objective (RTO).
- Further, some backup applications will only back up passive database copies. If this is true and the entire machine must be backed up, consider having active and passive Exchange Servers. Remember that Microsoft does not support hypervisor virtual machine snapshots of Exchange servers, so backup applications that must do this should only restore individual database and transaction logs, and not attempt to restore the VM itself.
- Some network links between sites cannot support a stretched DAG. Some replication applications replicate at the Windows volume layer. Placing too many databases on a single volume can cause the replication window to be too large and can cause failures.



- A database layout where servers mirror each other enables features such as Pure Clones to rapidly seed databases regardless of size. Place several databases so that the total size on the LUN meets the recovery SLA, should all those databases need to be restored from a backup repository.

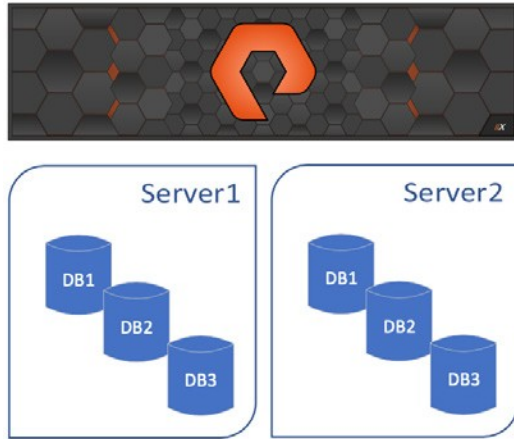


FIGURE 1. DB mirror layout

### Snapshots and Clones

Exchange requires VSS backup or native data protection (circular logging). Native data protection does not protect against malware, accidents, or longer-term (>two weeks) recovery. Up to two-week recovery requires additional copies of the database, typically one in each datacenter configured as a lagged database copy.

Exchange VSS can utilize a third-party backup application, or requestor, to take a software VSS snapshot. For virtual machines, this typically involves a copy-on-write snapshot that creates child files for each virtual hard disk file for each snapshot. This is usually limited to a handful of snapshots as each additional snapshot creates a child virtual hard disk file where all new writes are written. The more snapshots that exist increases the performance cost and increases disk latency. When deleting a software snapshot, all the data must be merged into the master virtual hard disk file which can stun the VM, impact the amount of IO, increase latency, and send thousands to hundreds of thousands of Ios to the storage subsystem. With higher IO VMs it can take hours to delete a snapshot while reducing total IOPS and increasing latency. Some backup applications attempt to minimize this impact by scheduling backup to off hours, and by creating a single snapshot, copying the changed blocks, and then immediately removing the software snapshot.

Exchange VSS can leverage the Pure Storage Hardware VSS Provider. This provider can be called by some thirdparty backup applications and DiskShadow, which ships in Windows. The benefit to the hardware provider is that the snapshots are created at the Pure FlashArray Volume layer. These Pure Volume snapshots are created and deleted in the millisecond range, have no performance impact, and do not cause tens of thousands of Ios on deletion. These snapshots can be copied (clones) and presented to other servers, which enables use cases such as:

- **Test/Dev:** Will that patch, update, or Exchange Cumulative Update break something?
- **Forensics/Discovery:** Instantly present a copy of a database as a [recovery database](#) to search through mailboxes or package message items into a PST file.
- **Rapid seed:** Create three additional copies of a 4TB database and instantly present it to other Exchange Servers in a DAG. Add a database copy with seeding postponed and be up in seconds, instead of hours or days.





- **Refresh a Pure volume with a recent snapshot:** Over time, data reduction rates will lessen due to the nature of independent database copies writing unique header information. In seconds periodically refresh additional database copies to sustain high data reduction rates.

## Disaster Recovery

Exchange VSS backups can be used to recover from a disaster. Ensure the backup repository is not stored on the same storage array as the production Exchange Servers, or else that storage array is a single point of failure.

Most Exchange administrators no longer backup Exchange Servers, only the databases and transaction logs. This is because Microsoft does not support [virtual machine snapshots](#) of Exchange servers, and because it is trivial to recover an [Exchange Server with setup switches](#). For example: `Setup /m:recoverserver`

A [recovery database](#) can be used to mount a database file, and a mailbox restore can be issued against data in that database where it is restored to the user's mailbox. This is often used after mounting an empty database, or [dial tone recovery](#), to get mail flow working immediately and, later, recovering messages.

[Database portability](#) is another option if waiting for an unsupported VM restoration, or recovering an Exchange Server will take a long time. Simply present the database to a still-functioning Exchange Server and then rehome the users to that database. This can be accomplished in a minute or two after migrating the database file(s) to a surviving server, which for Pure using Volume Snapshots, would take seconds.

## Flash-to-Flash-to-Cloud

Flash offers orders of magnitude more performance in restoration speed, and many customers are [moving backup repositories to Pure FlashBlade®](#) for incredible restore performance measured in minutes, not hours or days. Older backups that must be retained are moved to the cloud after [Pure ObjectEngine™](#) reduces the size of the backup—reducing cloud storage and bandwidth costs. [Veeam](#), [Rubrik](#), and [Commvault](#) directly integrate with FlashArray and FlashBlade to help recover Exchange Servers or databases quickly.

## Storage Migration

Storage Migration can be difficult when the source and target storage array is not a Pure Storage FlashArray. Limiting downtime and reducing end-to-end time to completion are two common goals. Taking a database or a database copy offline, copying the files, and then playing drive letter or mount point swap is prone to error, and takes the most administrative time. Recommended methods with which customers have reported the most success follow.

- **Move Mailbox:** New Databases can be created on the new storage, and users can be moved to these new databases. This has the smallest administrative impact, but it can take a long time, and causes a lot of network traffic. Since all mail is moved to new databases, it will all be written to the new database transaction log first, and not properly sizing the transaction log LUN, or ensuring additional backup intervals, can cause an out of space condition and an outage. Existing backup jobs would have to be changed to account for the new and removed databases.
- **Hypervisor Storage Migration:** Both VMware's Storage vMotion and Microsoft's Storage Live Migration can be used to move Exchange Server VMs to new storage without impacting availability. Existing backup jobs would not be impacted after migrating.



- **Additional Database Copies:** Add an additional database copy for each database and place them on the new storage. Exchange will seed the database and then these can be cloned until the correct number of database copies exist on the new storage and copies that exist on the legacy storage can be removed. Existing backup jobs may have to be changed to account for the new and removed database copies.

## VMWare Best Practices

VMware has very good Exchange Best Practices for both [Exchange Server 2016](#) and [Exchange Server 2019](#). The focus here is on the most common misconfigurations.

### Compute and Memory

VMware supports the full 48-core maximum for Exchange Server 2019, however only with two sockets. Be sure to configure the VM with the correct number of cores per socket to equal two sockets with up to 48 total cores. Microsoft recommends a 1:1 physical to virtual CPU mapping but supports up to 1:2.

### SCSI Controller

There are many storage queues, and customers transitioning from hard disk or hybrid storage to FlashArray may not be aware of them, as the hard disks were often the bottleneck.

At the virtual machine layer, VMware recommends using up to three Paravirtual SCSI adapters for increased performance, which requires VMware Tools be installed and kept up to date.

### Network

Use VMXNET3 network adapters. This provides better throughput while reducing hypervisor CPU, and the network is critical for a healthy DAG.

Enable jumbo frames on the virtual switches where vMotion traffic is enabled. Ensure jumbo frames is also enabled on the physical network infrastructure.

If using iSCSI from the FlashArray to ESXi or directly to the guest, ensure jumbo frames is enabled end to end.

### Disk Configuration

The [MetaCacheDatabase](#) (MCDB) accelerates Exchange when placed on flash storage and VMDKs can be created and presented to VMs when marked as [MediaType SSD](#) within the guest operating system.

### Anti-Affinity

Anti-affinity rules enforce VM separation and should be utilized to prevent all copies of the same database from existing on the same hypervisor host. ESXi 5.5 requires the `das.respectVmVmAntiAffinityRules` configuration option be set to true. For newer versions of ESXi, this is automatic.



## Hyper-V Best Practices

### Highly Available Virtual Machine

When a VM is created with Hyper-V manager it is not highly available. PowerShell includes a switch, and System Center Virtual Machine Manager has a checkbox. If the VM is not a failover cluster resource, and the hypervisor goes down, the VM is not failed over.

```
Get-VM <VMName> | Add-ClusterVirtualMachineRole
```

### Anti-affinity

Anti-affinity rules enforce VM separation and should be utilized to prevent all copies of the same database from existing on the same hypervisor host. [Anti-affinity groups](#) can be created in PowerShell.

First, create an anti-affinity group and name it.

```
$ExchServers1 = New-Object System.Collections.Specialized.StringCollection
$ExchServers1.Add("DAG Group")
```

Next add the Exchange VMs to the new group.

```
(Get-ClusterGroup -Name exch1).AntiAffinityClassNames = $ExchServers1
(Get-ClusterGroup -Name exch2).AntiAffinityClassNames = $ExchServers1
```

View the VMs and ensure they are assigned to the correct Anti-Affinity group.

```
Get-ClusterGroup | Select-Object -Property Name,AntiAffinityClassNames
```

Name	AntiAffinityClassNames
Available Storage	{}
Cluster Group	{}
exch1	{DAG Group}
exch2	{DAG Group}

## FlashStack and Exchange Server

FlashStack® provides a jointly supported solution by Cisco and Pure Storage, delivering a carefully validated architecture built on superior compute, world-class networking, and the leading innovations in all-flash storage. The portfolio of validated offerings from FlashStack includes but is not limited to the following:

- **Consistent performance:** FlashStack provides higher, more consistent performance than disk-based solutions and delivers a converged infrastructure based on all-flash that provides non-disruptive upgrades and scalability.
- **Cost savings:** FlashStack uses less power, cooling, and data center space when compared to legacy disk/hybrid storage. It provides industry-leading storage data reduction and exceptional storage density.



- **Simplicity:** FlashStack requires low ongoing maintenance and reduces operational overhead. It also scales simply and smoothly, in step with business requirements.
- **Deployment choices:** FlashStack is available as a custom-built single unit from FlashStack partners, but organizations can also deploy using equipment from multiple sources, including equipment they already own.
- **Unique business model:** The Pure Storage Evergreen Storage™ model enables companies to keep their storage investments forever, which means no more forklift upgrades and no more downtime.
- **Mission-critical resilience:** FlashStack offers best-in-class performance by providing active-active resiliency, no single point of failure, and non-disruptive operations—enabling organizations to maximize productivity.
- **Support choices:** Focused, high-quality, single-number reach for FlashStack support is available from FlashStack Authorized Support Partners. Single-number support is also available directly from Cisco Systems as part of the Cisco Solution Support for Data Center offering. Support for FlashStack components is also available from Cisco, VMware, and Pure Storage individually and leverages TSANet for resolution of support queries between vendors.

The [Cisco FlashStack Design Guide](#) for Exchange Server is a blueprint for deploying 40,000 Exchange mailboxes on Pure FlashArray, Cisco UCS Servers, Cisco Fabric Interconnects, and Cisco Nexus Switches.

## Jetstress Testing

[Jetstress](#) is a Microsoft tool that calls directly into the jet engine (ese.dll) simulating the Exchange Information Store disk IO. It was created when 100 IOPS per hard disk device was the standard, where slightly higher could be achieved with 15k rpm spindles, and significantly lower with SATA or nearline SAS. Jetstress does not do a good job of testing flash storage, as it is not a benchmark and wasn't designed for storage where each flash device can sustain tens of thousands of IO. It is a good quick sanity test for the network and MPIO configuration where achieving ~20k IOPS at low latency proves the solution is not misconfigured. If the entire physical or virtual Exchange infrastructure will be built out and then tested with Jetstress, remember that a single LUN will hold one to four databases and with 0.05 IOPS averages per user, 4000 users only requires 200 transactional IOPS plus 1-2MB/sec of maintenance reads.

**Test an Exchange mailbox profile:** Running a burn-in test to shake out the system isn't a bad idea, but it should be done with care, and a specific thread count. If you let JetStress "auto-tune" on any all-flash array, or even a RAM drive, the tool itself will break down. You'll see high latencies not from the storage, but JetStress itself not being able to write fast enough. This is evidenced as the threads increase the higher the latency (>30ms) without a significant increase in the achieved IOPS. At the same time, look at the storage console or perfmon and it will be sitting at less than 1-2ms of latency. JetStress thinks the storage is adequate and further increases the thread count in a vicious loop. If you choose to run the 'test subsystem throughput' test, manually select a saner thread count and increment it yourself.

The Jetstress tool breaks down when there are more than two to three threads per database. Since the thread count is global, this is very apparent with greater than 16 threads when deploying with less than five databases. As the database count is increased, the total thread count can be increased. Evidence of too many threads will be that the write latency difference between what Jetstress reports, and what is logged in perfmon will widen. It will stay within 15 to 20ms and sharply increase until other failures are experienced. Both -1018 corruption errors and Database Fault Stalls per second are indications that there are too many threads.





Figure 2: Setting Jetstress thread count..

**Best Practices:**

- Select Test an Exchange mailbox profile, do not test disk subsystem throughput.
- Select Suppress tuning and use thread count (global). Sixteen threads is a good start, just be aware that it will provide orders of magnitude more IOPS than necessary if building out the entire Exchange Server deployment. If testing many databases, try to keep the thread count to within two or three times the number of databases.
- When defining the database configuration, be sure to select the correct number of copies per database so that replication IOPS or reads against the transaction logs occur.

**Microsoft Integration**

**SMI-S**

Pure FlashArray supports [SMI-S natively in Purity](#), and the SMI-S provider can be enabled simply in the Pure FlashArray Management GUI.

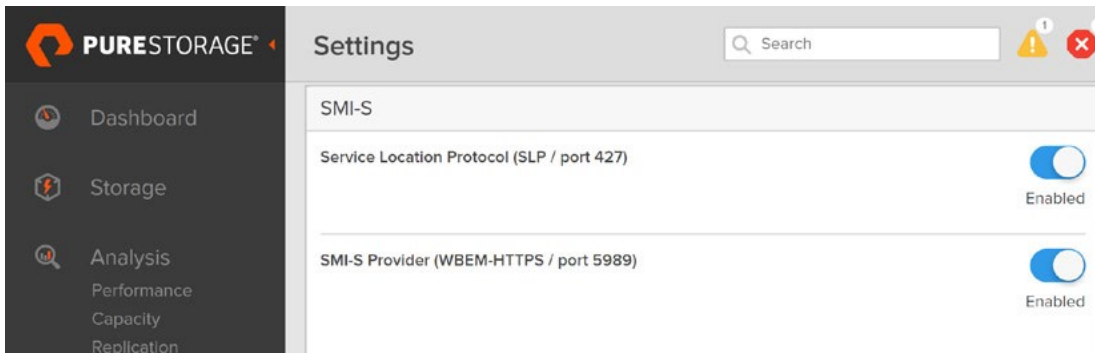


FIGURE 3. SMI-S provider is enabled on Pure FlashArray

The two most common clients for SMI-S are System Center Virtual Machine Manager (SCVMM), and the [Windows Standards-Based Storage Management](#) feature in Windows.

**System Center Virtual Machine Manager (SCVMM)**

Pure FlashArray can be [added to SCVMM](#) to show consumed and available capacity. LUNs can be created, and virtual machines can be deployed on the FlashArray directly in SCVMM.






Figure 4. FlashArray properties in SCVMM 2019

## System Center Operations Manager (SCOM)

The Pure Storage [FlashArray Management Pack](#) for Microsoft System Center Operations Manager enables monitoring of health, capacity, and usage to diagnose and troubleshoot issues across applications, workloads, and infrastructure.

## Troubleshooting

When utilizing iSCSI, the three most common misconfiguration issues are:

- Jumbo frames are not correctly configured end to end. This is the storage, the initiator, the switch, and if applicable the vSwitch layers. From windows ensure you can ping with 8000 from initiator to the FlashArray.  

```
ping -l 8000 <flasharray-iscsi-IP>
```
- High latency because a device in the network path is dropping packets. Get a Wireshark of the behavior and look at the basic analysis functionality for dropped packets, retransmissions, and out of order errors.
  - Select the Analyze menu, then Expert Information
- High latency because the throughput is not uniform to and from each device in the network path. See [the Pure Microsoft Platform Guide](#) for details.

## Conclusion

Exchange Server has evolved and become the mission-critical messaging platform of choice for many organizations. As your data platform evolves, so should your storage, backed by the FlashArray//X product offerings. Utilizing all-flash technology and our Purity Operating Environment can help meet the demands of running a variety of workloads associated with Exchange Server maintenance tasks to optimize environment resources. Pure Storage reduces the complexity of dealing with mixed workloads by providing consistent performance across mixed workload environments.



## For More Information

- [Pure Storage FlashArray//X Data Sheet](#)
- [Pure Storage Microsoft Platform Guide](#)
- [Pure Storage VMware Platform Guide](#)
- [FlashStack Converged Infrastructure Solutions](#)
- [Veeam and Pure Solution Brief](#)
- [Highly Available Rapid Restore with Veeam Backup and Replication and FlashBlade](#)
- **Best Practices**
  - [Pure Storage Best Practices Guide for Hyper-V 2016](#)
  - [Microsoft Exchange Server 2016 on vSphere Best Practices](#)
  - [Microsoft Exchange Server 2019 on VMware Best Practices](#)
- **Design Guides**
  - [Design Guide for Microsoft Applications with Hyper-V on FlashStack Mini CI](#)
  - [Design Guide for Microsoft Exchange Server on FlashStack™](#)

©2020 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. 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. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS 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, #400  
Mountain View, CA 94041

[purestorage.com](https://www.purestorage.com)

800.379.PURE

