



PURE STORAGE FLASHSTACK CONVERGED INFRASTRUCTURE SOLUTION

Design Guide for Microsoft Exchange Server on FlashStack™

May 2016



TABLE OF CONTENTS

INTRODUCTION	2
GOALS AND OBJECTIVES	2
DESIGN GUIDE PRINCIPLES	2
ABOUT FLASHSTACK	3
WHY FLASHSTACK	3
FLASHSTACK COMPONENTS	4
PURE STORAGE FLASHARRAY	4
CISCO UNIFIED COMPUTING SYSTEM	7
CISCO UCS FABRIC INTERCONNECT	8
CISCO NEXUS 5500UP SWITCHES	8
VMWARE VSPHERE 6.0	9
FLASHSTACK TOPOLOGY	10
DESIGN AND DEPLOYMENT	11
CONFIGURATION	12
NETWORKING AND BACKUP	15
FLASHSTACK VALUE	15
TEST RESULTS SUMMARY	17
INDIVIDUAL SERVER METRICS	18
CONCLUSION	21
REFERENCES	21

INTRODUCTION

FlashStack is a flexible, all-flash converged infrastructure solution from Pure Storage® that brings the flash revolution to your data center. It combines the latest in compute, network, storage hardware and virtualization software into a single, integrated architecture that reduces time to deployment, lowers overall IT costs and reduces deployment risk. Highly efficient components reduce the costs associated with power, cooling and data center space. Based on 100% flash storage, FlashStack provides the performance and reliability business-critical applications demand.

This document describes a reference architecture for deploying 40,000 Exchange Server mail boxes on FlashStack. The test results enumerated in this document explain the performance and headroom as they relate to the maximum achieved IOPs and bandwidth. Pure Storage has validated the reference architecture within its lab – this document presents the hardware and software configuration, the test workload configuration, testing results and further offers implementation and sizing guidance for deploying large number of mail boxes in a relatively smaller space allowing room for other applications.

GOALS AND OBJECTIVES

The goal of this document is to showcase the scalability of FlashStack to accommodate a large Exchange deployment and allow an option to consolidate other applications within the same deployment. This document will demonstrate the following:

- A complete blueprint for 40,000 Exchange Mailbox deployment.
- Information on sizing the infrastructure components.
- Performance characterization of the infrastructure under load.

DESIGN GUIDE PRINCIPLES

The guiding principles for implementing this reference architecture are:


- **Repeatable:** Create a scalable building block that can be easily replicated at any customer site. Publish the version of various firmware under test and weed out any issues in the lab before customers deploy this solution.
- **Available:** Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage, multiple NICs for connectivity, and high availability (HA) clustering.
- **Efficient:** Build a solution that leverages efficient benefits of an all-flash architecture.
- **Simple:** Simplify deployment and ongoing maintenance tasks via automation.
- **Scalable:** Create a design that can start small, but can easily grow to meet the needs of a growing enterprise.

ABOUT FLASHSTACK

FlashStack is a converged infrastructure solution that brings the benefits of an all-flash storage platform to your converged infrastructure deployments. Built on best of breed components from Cisco and Pure Storage, FlashStack provides a converged infrastructure solution that is simple, flexible, efficient, and costs less than legacy converged infrastructure solution based on traditional disk.

FlashStack is available from accredited FlashStack Partners who help provide an excellent converged infrastructure ownership experience. FlashStack Partners have the knowledge and experience necessary to help streamline the sizing, procurement, and delivery of your entire system.

WHY FLASHSTACK



Consistent Performance and Scalability

- Consistent sub-millisecond latency with 100% flash storage.
- Consolidate 100s of enterprise-class applications in a single rack.
- Scale easily, without disruption.
- Repeatable growth through multiple FlashStack deployments.

Operational Simplicity

- Fully tested, validated, and documented for rapid deployment
- Reduced management complexity
- Auto-aligned 512b architecture eliminates storage alignment headaches
- No storage tuning or tiers necessary

Lowest TCO

- Dramatic savings in power, cooling and space with 100% Flash.
- Industry leading data reduction
- FlashArray controller upgrades included every three years with ForeverFlash™

Enterprise Grade Resiliency

- Highly available architecture and redundant components
- Non-disruptive operations
- Upgrade and expand without downtime or performance loss
- Native data protection: snapshots and replication

FLASHSTACK COMPONENTS


PURE STORAGE FLASHARRAY



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

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 FlashArray//m.

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.



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>

WHY FLASHARRAY

The FlashArray//m expands upon the FlashArray's modular, stateless architecture, designed to enable expandability and upgradability for generations. The FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your business' needs today and tomorrow.

The Pure Storage FlashArray is ideal for:

Accelerating Databases and Applications Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.

Virtualizing and Consolidating Workloads Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.

Delivering the Ultimate Virtual Desktop Experience Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance for under \$50/desktop.

Protecting and Recovering Vital Data Assets Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

Consistent Performance FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

Less Cost than Disk Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure. With VDI workloads data reduction is typically > 10:1.

Mission-Critical Resiliency FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

Disaster Recovery Built-In FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

Simplicity Built-In FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, Windows PowerShell, Python, REST API, or OpenStack.

PURITY OPERATING ENVIRONMENT

Purity operating system runs on a FlashArray that implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.

Storage Software Built for Flash The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.

Granular and Adaptive Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.


Best Data Reduction Available FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

Highly Available and Resilient FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.

Backup and Disaster Recovery Built-In FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.

PURE1

Pure1™ is the management platform of FlashArray that provides the following features:



Pure1 Manage By combining local web-based management with cloud-based monitoring, Pure1Manage allows you to manage your FlashArray wherever you are – with just a web browser.

Pure1 Connect A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//m to all your data center and cloud monitoring, management, and orchestration tools.

Pure1 Support FlashArray//m is constantly cloud-connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.


Pure1 Collaborate Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.



EVERGREEN STORAGE

Tired of the 3-5 year array replacement merry-go-round? The move to FlashArray//m can be your last data migration. Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. Pure Storage's vision for Evergreen Storage is delivered by a combination of the FlashArray's stateless, modular architecture and the ForeverFlash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

CISCO UNIFIED COMPUTING SYSTEM



The Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that unites compute, network, storage access, and virtualization into an organized structure aimed to reduce total cost of ownership and introduce vastly improved infrastructure deployment mechanisms at scale. UCS incorporates a unified network fabric with scalable, modular and powerful x86-architecture servers. With an innovative and proven design, Cisco UCS delivers an architecture that increases cost efficiency, agility, and flexibility beyond what traditional blade and rack-mount servers provide. Cisco makes organizations more effective by addressing the real problems that IT managers and executives face and solves them on a systemic level.

GREATER TIME-ON-TASK EFFICIENCY

Automated configuration can change an IT organization's approach from reactive to proactive. The result is more time for innovation, less time spent on maintenance, and faster response times. These efficiencies allow IT staff more time to address strategic business initiatives. They also enable better quality of life for IT staff, which means higher morale and better staff retention—both critical elements for long-term efficiency.

Cisco UCS Manager is an embedded, model-based management system that allows IT administrators to set a vast range of server configuration policies, from firmware and BIOS settings to network and storage connectivity. Individual servers can be deployed in less time and with fewer steps than in traditional environments. Automation frees staff from tedious, repetitive, time-consuming chores that are often the source of errors that cause downtime, making the entire data center more cost effective.

EASIER SCALING

Automation means rapid deployment, reduced opportunity cost, and better capital resource utilization. With Cisco UCS, rack-mount and blade servers can move from the loading dock and into production in a "plug-and-play" operation. Automatically configure blade servers using predefined policies simply by inserting the devices into an open blade chassis slot. Integrate rack-mount servers by connecting them to top-of-rack Cisco Nexus fabric extenders. Since policies make configuration automated and repeatable, configuring 100 new servers is as straightforward as configuring one server, delivering agile, cost-effective scaling.

VIRTUAL BLADE CHASSIS

With a separate network and separate management for each chassis, traditional blade systems are functionally an accidental architecture based on an approach that compresses all the components of a rack into each and every chassis. Such traditional blade systems are managed with multiple management tools that are combined to give the illusion of convergence for what is ultimately a more labor-intensive, error-prone and costly delivery methodology. Rack-mount servers are not integrated and must be managed separately or through additional tool sets, adding complexity, overhead, and the burden of more time.

Architecturally, Cisco UCS blade and rack-mount servers are joined into a single virtual blade chassis that is centrally managed yet physically distributed across multiple blade chassis, rack-mount servers, and even racks and rows. This capability is delivered through Cisco fabric interconnects that provide redundant connectivity, a common management and networking interface, and enhanced flexibility. This larger virtual chassis, with a single redundant point of management, results in lower infrastructure cost per server, with fewer management touch points, and lower administration, capital, and operational costs.

CISCO UCS FABRIC INTERCONNECT

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a core part of Cisco Unified Computing System. It is usually deployed in redundant pairs, the Cisco UCS 6248UP Fabric Interconnects provide uniform access to both networks and storage. Here we are configuring it for the End-Host Mode.

Left: Cisco UCS
6248UP

Right: Cisco Nexus
5500UP switch



CISCO NEXUS 5500UP SWITCHES

Cisco Nexus 5500UP switches, using cut-through architecture, supports line-rate 10 Gigabit Ethernet on all ports while maintaining consistently low latency independent of packet size and services enabled. It supports a set of network technologies known collectively as Data Center Bridging (DCB) that increases the reliability, efficiency, and scalability of Ethernet networks. These features allow the switches to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN and cluster environments. Its ability to connect Fibre Channel over Ethernet (FCoE) to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling. The Nexus 5500UP series switch we used was a Nexus 5548UP (see Figure 3 below), which is a 1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports and one expansion slot.

VMWARE VSPHERE 6.0

VMware vSphere is the industry-leading virtualization platform for building cloud infrastructures. It enables IT to meet SLAs (service-level agreements) for the most demanding business critical applications, at the lowest TCO (total cost of ownership). vSphere accelerates the shift to cloud computing for existing data centers and also underpins compatible public cloud offerings, forming the foundation for the industry's only hybrid cloud model. With the support of more than 3,000 applications from more than 2,000 ISV partners, vSphere is the trusted platform for any application.

- VMware vSphere Hypervisor Architecture provides a robust, production-proven, high-performance virtualization layer. It enables multiple virtual machines to share hardware resources with performance that can match (and in some cases exceed) native throughput.
- Each vSphere Hypervisor 6.0 instance can support as many as 480 logical CPUs, 12TB of RAM, and 1024 virtual machines. By leveraging the newest hardware advances, ESXi 6.0 enables the virtualization of applications that were once thought to be non-virtualizable.
- VMware ESXi 6.0 has dramatically increased the scalability of the platform. With vSphere Hypervisor 6.0, clusters can scale to as many as 64 hosts, up from 32 in previous releases. With 64 hosts in a cluster, vSphere 6.0 can support 8,000 virtual machines in a single cluster. This enables greater consolidation ratios, more efficient use of VMware vSphere Distributed Resource Scheduler (vSphere DRS), and fewer clusters that must be separately managed.
- VMware vSphere Virtual Machine File System (VMFS) allows virtual machines to access shared storage devices (Fibre Channel, iSCSI, etc.) and is a key enabling technology for other vSphere components such as VMware vSphere Storage vMotion.
- VMware vSphere Storage APIs provide integration with supported third-party data protection, multipathing and storage array solutions.
- Optimization with the Software-Defined Data Center – Allocates resources dynamically with virtual storage, compute, and networking to manage and deliver desktop services on demand.
- Central Image Management – Central image management for physical, virtual, and BYO devices.
- Hybrid-cloud flexibility – Provides an architecture built for onsite and cloud-based deployment

FLASHSTACK TOPOLOGY

In this section, we will learn more about the high-level topology of the FlashStack reference architecture. FlashStack comprises the following hardware and software components:


HARDWARE

- **Network:** Cisco Nexus 5500UP and Cisco UCS Fabric Interconnect 6248UP for external and internal connectivity of IP and FC network.
- **Storage:** Pure Storage FlashArray//m 20 with Fibre Channel connectivity
- **Compute:** Cisco UCS B440 M2 Blade Server

SOFTWARE

- Cisco UCS Manager
- Microsoft Exchange Server 2013
- Pure Storage PowerShell SDK
- Pure Storage Management interface

The following diagram (Figure 4) shows a detailed topology of the reference architecture configuration. A major goal of the architecture is to build out a highly redundant and resilient infrastructure. We used powerful servers with dual Fibre Channel ports connected redundantly to two SAN switches that were connected to redundant FC target ports on the FlashArray//m. The servers also have redundant network connectivity. It is configured in End-Host Mode.



FlashStack connectivity diagram

DESIGN AND DEPLOYMENT


In this section, we explain the design and configuration of individual components in the test setup that was deployed to validate the performance of Microsoft Exchange solution.

The logical architecture of the Microsoft Exchange Server is explained in Figure 5, below. In the test setup, 40,000 Exchange mailboxes were simulated—with each blade server having 4 databases, 1250 mailboxes were deployed in each database. A total of eight servers were deployed per Database Availability Group (DAG). To recreate a high-availability environment, another DAG with a similar setup was set up. Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it works with the other servers in the DAG to provide automatic recovery from failures that affect mailbox databases, such as a disk, server, or network failure. The blade servers with the mail box VMs as a cluster is connected to an 8GB/s Fibre channel SAN switch. The SAN switch establishes a connection between the blade server cluster and FlashArray//m.

The simulated IOPS/second/mailbox was 1. At the end of the performance testing, we observed that a headroom of around 20% was still available.

Other attributes are as below:

Attribute	Value
Database LUN size	3 TB
Log LUN size	200 GB
Total database size for performance testing	78.99 TB



Logical architecture of the Microsoft Exchange Server deployment

The following are the recommendations to achieve optimal performance for the individual mailbox server virtualized instances:

- Windows Server NTFS volumes should use a 64KB Allocation Unit Size (AUS) for Microsoft Exchange databases and logs. This applies to Windows Server 2008 R2, Windows Server 2012 and Windows Server 2012 R2.
- All Service Pack updates or Operating System Hotfixes should be applied for all components in the infrastructure.
- Configure individual virtual instances of the mailbox servers with adequate CPU and memory resources. For the Jetstress tests performed we used a Cisco UCS B200-M4 dedicating 4 vCPUs and 40GB memory per virtual machine instance.

CONFIGURATION

This section explains the configuration of each FlashStack component with respect to Exchange Mailbox in order to produce the desired results.

CISCO UCS SERVER IN EXCHANGE MAILBOX CONFIGURATION

A pair of Cisco UCS Fabric Interconnects 6248UP, and 3 Cisco UCS B-series B200-M4 blade servers were deployed for hosting Microsoft Exchange Server 2013. The UCS manager, UCS Fabric Interconnects and the components in the chassis were upgraded to 2.2.3f firmware level.

The server has Cisco VIC 1280 cards and they were connected through four ports from each Cisco Fabric extender of the Cisco UCS chassis to the Cisco Fabric Interconnect, they were in turn connected to Cisco Nexus 5548UP Switch for upstream connectivity to access the Pure Storage FlashArray//m LUNs. The server configuration is described in the table on page 19.

CISCO UCS SERVER CONFIGURATION	
Component	Description
Processor	4X Intel Xeon E74870 @ 2.4GHz, 40 Cores
Memory	256 GB @ 1600 MHz, regular voltage
HBA	4 X 8G ports on Cisco UCS VIC 1280 (UCSB-MLOM-40G-03)
NIC	4 X 10G ports on Cisco UCS VIC 1280 (UCSB-MLOM-40G-03)

CISCO UCS SERVICE PROFILE EXCHANGE MAILBOX CONFIGURATION

In order to facilitate rapid deployment of UCS servers, a service profile template was created with the following characteristics.

1. We configured boot from SAN policy so that the server booted from a Pure Storage boot LUN.
2. We kept every other setting to the default, we didn't tweak any parameters.
3. The Ethernet and FC adapter policy was set to Microsoft Windows Server 2012 R2.
4. The BIOS defaults were used for the B200-M2 blade servers.

5. We configured two vHBA FC adapters and four vNIC Eth adapters on the Cisco VIC cards to avoid any single point of failure.
6. We deployed three service profiles from the template and associated it with the blade servers in.

FLASHARRAY STORAGE IN EXCHANGE MAILBOX CONFIGURATION

As was mentioned earlier Pure Storage strives to reduce the complexity and management associated with the deployment of applications. All Pure Storage FlashArrays are configured for optimal multi-workload performance and do not require any specialized configurations or tuning. The only detail to consider is for volume design as follows:

- Separate the databases from the logs on individual volumes. For example, ExDb-1 has four 3TB volumes created to accommodate four databases and two copies.
- Four ExLog-1 volumes of 200GB are created for logs.

These are not hard set rules and can be modified based on your Exchange administration and recovery requirements.

HYPERVISOR CONFIGURATION

Microsoft Exchange Server does not make a requirement for a particular hypervisor. The test results reported in this report VMware vSphere 6.0 was used to host the virtualized mailbox servers. Pure Storage provides our best practices for VMware vSphere available here:

www.purestorage.com/resources/type-a/WP-PureStorageandVMwarevSphereBestPracticesGuide_Request.html

The following tables provides details on the configuration of the individual components.

STORAGE HARDWARE	
Component	Description
Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)	Fibre Channel (FC), SAS
Storage model and OS/firmware revision	Pure Storage FlashArray//m20 running Purity 4.5.8 Windows Hardware Compatibility List: https://www.windowsservercatalog.com/item.aspx?itemId=6e9fe3dd-23ba-d504-b29d-a203ed3d6b70&bCatID=1282
Storage cache	N/A
Number of storage controllers	2
Number of storage ports	4 (2 per controller)
Maximum bandwidth of storage connectivity to host	5 GB/s
Switch type/model/firmware revision	Cisco MDS9148S (BIOS 2.1.17, System 6.2(9)), Cisco Nexus 9K (BIOS 07.17, NXOS 6.1(2)I3(3a)), Cisco Fabric Interconnect 6248UP (2.2(3f))
HBA model and firmware	Cisco UCS VIC 1340
Number of HBAs/host	2
Host server type	Cisco C-220 (2 CPUs / 65,536MB), Cisco B200-M4 (2 CPUs / 131,072MB)
Total number of disks tested in solution	20
Maximum number of spindles can be hosted in the storage	N/A

STORAGE SOFTWARE	
Component	Description
HBA driver	8 GB Fibre Channel SAN
Hypervisor	VMware vSphere 6.0.U1 (Vmware-ESXi-6.0.0-3073146-Custom-Cisco-6.0.1.1)
HBA QueueTarget Setting	32
HBA QueueDepth Setting	32
Multi-Pathing	Microsoft Windows Server 2012 R2 Multipath-IO (Device Specific Module)
Host OS	Windows Server 2012 R2 Datacenter (6.3.9600)
ESE.dll file version	15.00.0516.026
Replication solution name/version	N/A

STORAGE DISK CONFIGURATION (MAILBOX STORE DISKS)	
Attribute	Description
Disk type, speed and firmware revision	N/A (all-flash storage solution)
Raw capacity per disk (GB)	1,024 GB
Number of physical disks in test	20
Total raw storage capacity (GB)	19,046.4 GB (18.6 TB)
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	N/A
RAID level	RAID-3D
Total formatted capacity	11.17 TB
Storage capacity utilization	5.8% (108.71TB / 18.6 TB)
Database capacity utilization	4.2% (78.99TB / 18.6 TB)

STORAGE DISK CONFIGURATION (TRANSACTIONAL LOG DISKS)	
Attribute	Description
Disk type, speed and firmware revision	N/A (all-flash storage solution)
Raw capacity per disk (GB)	1,024 GB
Number of Spindles in test	N/A (all-flash storage solution)
Total raw storage capacity (GB)	19,046.4 GB (18.6 TB)
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	N/A
RAID level	RAID-3D
Total formatted capacity	800 GB (4 x 200 GB LUNs)

NETWORKING AND BACKUP

All Exchange Server deployments are concerned with high-availability. This is achieved with Exchange by using DAG that automates recovery at the database-level after a database, server or network failure. We recommend multiple physical or virtual NICs that are connected to different networks to segregate user traffic from database replication traffic.

As mentioned earlier, the use of RDMs (Raw Device Mapping) should be considered for Exchange to leverage application-consistent VSS snapshots for data protection. Pure Storage provides a VSS hardware provider that can be downloaded directly from a Web Management Interface and installed to work with supported backup software.

FLASHSTACK VALUE

The value of deploying the Microsoft Exchange server solution on FlashStack Converged Infrastructure is to have additional space for other applications. In this deployment, more than 5000 mailboxes have been deployed on a given server which exceeds the threshold value thus having enough headroom for the other applications to run seamlessly on the same server platform. The scalability of FlashArray//m20 is 150K IOPS; however, in our test setup, we have seen that 40,000 IOPS are utilized for deploying the 40,000 mailboxes (@ 1 IOPS/mailbox), thus leaving 110K IOPS bandwidth available for additional applications to be run. This estimated number of IOPS and bandwidth availability can strengthen the multi-workload scenario in multiple ways. The high availability feature of each FlashStack component also ensures zero downtime.


The main values for deploying Microsoft Exchange Solution on FlashStack are:

1. **Simplified deployment and workload consolidation:** Exchange administrators can consolidate all of the Exchange infrastructure needs on FlashStack and also use the setup for consolidation of additional workloads.
2. **Simplified day to day management, infrastructure availability and elasticity:** Manage the infrastructure with industry standards tools such as UCS Director and VMware vCenter, which reduce the burden of day to day management. Highly available infrastructure that can scale up as the needs of the business grow.
3. **Securing and protecting MS Exchange Server:** At rest data encryption for all data. Built in snapshots and data replication allow easy protection of the data on FlashStack.
4. **Ultimate performance:** Consistent performance for Exchange deployments with sub-ms latency.

The FlashArray//m20 used in this testing demonstrates the following:

- 40,000 mail boxes contained in just 3U of rack space of Storage

On average, during the solution testing, we found the average number of mailboxes to be 44945.708; in the above representation, we can see that each server instance exceeded the threshold of 5000 mailboxes quite easily. This demonstrates that FlashArray//m20 supports 40,000 users while leaving enough room to support additional applications on the same array. Exchange administrators can use this available bandwidth to deploy additional applications and users.



Actual performance vs required thresholds per Mailbox Server

- Up to 4:1 data reduction ratios for Exchange data even when multiple workloads are co-located

FlashArray//m's FlashRecover Snapshot technology enables Exchange administrators to rapidly seed other Database Available Groups (DAGs) by leveraging multiple data reduction techniques, such as storage level compression, deduplication and thin provisioning. These features deliver an average of 4:1 data reduction ratio even when Exchange is deployed along with other production workloads on the same FlashArray//m. without compromising performance and resiliency. In most cases, customers co-locate Microsoft Exchange with other applications on FlashArray//m.

The result is substantial reduction in operating expenses with low power, space and cooling requirements in the data center. In addition, the elimination of storage silos and complexity helps data center managers to focus on value added business projects.

In our test setup, when we replaced legacy disk storage with Pure Storage FlashArrays, power and cooling savings totaled \$104,327 over three years and assumes a cost per KWH for power of \$0.14 and a cost or KWH for cooling of \$0.10 (see table below). We have risk-adjusted the savings downward by 7% to reflect regional KWH rate differentials.

POWER AND COOLING SAVINGS						
Ref.	Metric	Calc./Source	Year 1	Year 2	Year 3	Total
E1	Power and cooling costs— legacy disk	Interviews	\$37,563	\$43,197	\$49,677	\$130,437
E2	Power and cooling costs— Pure Storage	Interviews	\$5,258	\$6,046	\$6,953	\$18,257
E3	Power and cooling savings with Pure Storage	E1-E2	\$32,305	\$37,151	\$42,724	\$112,180
Et	Total power and cooling savings	E3	\$32,305	\$37,151	\$42,724	\$112,180
	Risk Adjustment	↓7%				
Elr	Power and cooling savings (risk-adjusted)	Et-7%	\$30,044	\$34,550	\$39,733	\$104,327

Source: Forrester Research, Inc.

1 IOPS per mail box with enough head room for growth and for running other workloads.

TEST RESULTS SUMMARY

This section provides a high-level summary of the test data for Microsoft Exchange mailbox.

TEST GOALS

Reliability

A number of tests in the framework are to check Reliability tests runs for 24 hours. The goal is to verify the storage can handle high IO load for a long period of time. Both log and database files will be analyzed for integrity after the stress test to ensure no database/log corruption.

Storage Performance Results

The Primary Storage performance testing is designed to exercise the storage with maximum sustainable Exchange type of IO for 2 hours. The test is to show how long it takes for the storage to respond to an IO request under load. The data below is the sum of all of the logical disk I/O's and average of all the logical disks I/O latency in the 2 hours' test duration. Each server is listed separately and the aggregate numbers across all servers are listed as well.

Customers should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

INDIVIDUAL SERVER METRICS

Upon the completion of a two-hour test, the following results were observed. They are listed as the sum of I/O's across Storage Groups and the average latency across all Storage Groups on a per server basis. During the tests,

- No errors were reported in the saved event log file.
- No errors were reported for the database and log checksum processes.

WIN-ME-0001	
Database I/O	Value
Database Disks Transfers/sec	5439.899
Database Disks Reads/sec	3633.815
Database Disks Writes/sec	1806.084
Average Database Disk Read Latency (ms)	1.05275
Average Database Disk Write Latency (ms)	1.55275
Transaction Log I/O	Value
Log Disks Writes/sec	811.039
Average Log Disk Write Latency (ms)	0.412

WIN-ME-0004	
Database I/O	Value
Database Disks Transfers/sec	6409.672
Database Disks Reads/sec	4295.139
Database Disks Writes/sec	2114.533
Average Database Disk Read Latency (ms)	1.162
Average Database Disk Write Latency (ms)	2.00925
Transaction Log I/O	Value
Log Disks Writes/sec	932.149
Average Log Disk Write Latency (ms)	0.4405

WIN-ME-0007	
Database I/O	Value
Database Disks Transfers/sec	5539.138
Database Disks Reads/sec	3698.984
Database Disks Writes/sec	1840.154
Average Database Disk Read Latency (ms)	1.033
Average Database Disk Write Latency (ms)	1.51375
Transaction Log I/O	Value
Log Disks Writes/sec	825.697
Average Log Disk Write Latency (ms)	0.41125

WIN-ME-0010	
Database I/O	Value
Database Disks Transfers/sec	5351.198
Database Disks Reads/sec	3576.194
Database Disks Writes/sec	1775.004
Average Database Disk Read Latency (ms)	1.07025
Average Database Disk Write Latency (ms)	1.7
Transaction Log I/O	Value
Log Disks Writes/sec	798.777
Average Log Disk Write Latency (ms)	0.415

WIN-ME-0013	
Database I/O	Value
Database Disks Transfers/sec	5993.923
Database Disks Reads/sec	4007.353
Database Disks Writes/sec	1986.57
Average Database Disk Read Latency (ms)	0.9545
Average Database Disk Write Latency (ms)	1.41325
Transaction Log I/O	Value
Log Disks Writes/sec	890.468
Average Log Disk Write Latency (ms)	0.37

WIN-ME-0016	
Database I/O	Value
Database Disks Transfers/sec	5373.341
Database Disks Reads/sec	3589.651
Database Disks Writes/sec	1783.69
Average Database Disk Read Latency (ms)	1.08025
Average Database Disk Write Latency (ms)	1.72975
Transaction Log I/O	Value
Log Disks Writes/sec	803.038
Average Log Disk Write Latency (ms)	0.4155

WIN-ME-0019	
Database I/O	Value
Database Disks Transfers/sec	5497.137
Database Disks Reads/sec	3671.74
Database Disks Writes/sec	1825.397
Average Database Disk Read Latency (ms)	1.0465
Average Database Disk Write Latency (ms)	1.502
Transaction Log I/O	Value
Log Disks Writes/sec	818.388
Average Log Disk Write Latency (ms)	0.42025

WIN-ME-0022

Database I/O	Value
Database Disks Transfers/sec	5341.4
Database Disks Reads/sec	3568.865
Database Disks Writes/sec	1772.535
Average Database Disk Read Latency (ms)	1.06675
Average Database Disk Write Latency (ms)	1.51625
Transaction Log I/O	Value
Log Disks Writes/sec	799.088
Average Log Disk Write Latency (ms)	0.406

AGGREGATE PERFORMANCE ACROSS ALL SERVER METRICS

The sum of I/O's across servers in solution and the average latency across all servers in solution.

Database I/O	Value
Database Disks Transfers/sec	44945.708
Database Disks Reads/sec	30041.741
Database Disks Writes/sec	14903.967
Average Database Disk Read Latency (ms)	1.05825
Average Database Disk Write Latency (ms)	1.617125
Transaction Log I/O	Value
Log Disks Writes/sec	6678.644
Average Log Disk Write Latency (ms)	0.4113125

SOFTWARE MATRIX

Component	Software
Server Application	Microsoft Exchange Server 2013
Operating System	Windows Server 2012 R2 Update
Performance Testing	Microsoft Exchange Jetstress 2013
File System	NTFS (64kb Allocation Unit Size)

CONCLUSION

We set out to prove that FlashStack Converged Infrastructure with Pure Storage FlashArray provides an ultimate infrastructure platform for Microsoft Exchange 2013 and we have achieved consistently impressive results. The test results/data presented in this document is based on the tests introduced in Microsoft's ESRP test framework.

The test results achieved shows that we can deploy 40000 mailboxes easily and we have enough headroom in the infrastructure to run other applications effortlessly with relatively low latency.

REFERENCES

Below are the references to various documents pertaining to Microsoft Exchange 2013:

- **ESRP** <https://technet.microsoft.com/en-us/office/dn756396.aspx>
- **Exchange server 2013** [https://technet.microsoft.com/en-us/library/bb124558\(v=exchg.150\).aspx](https://technet.microsoft.com/en-us/library/bb124558(v=exchg.150).aspx)
- **MSFT best practices for Exchange** http://download.microsoft.com/download/4/A/C/4AC32FD3-220E45DCAA97DBDBE19C15B2/Best_Practices_for_Virtualizing_and_Managing_Exchange_2013.pdf
- **Pure1 Community best practices for Windows** <http://community.purestorage.com/t5/Interoperability-Best-Practice/Windows-Server-Best-Practices/ta-p/166>

© 2016 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, ForeverFlash and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries.

Microsoft Windows Server and Microsoft Exchange are registered trademarks of Microsoft in the U.S. and other countries. Cisco®, Cisco UCS®, Cisco Nexus®, Cisco Unified Computing System™ are trademarks of Cisco in the U.S. and other countries. VMware vSphere®, VMware vMotion®, VMware ESXi™, and VMware vSphere Distributed Resource Scheduler™ are registered trademarks of VMware in the U.S. and other countries. Intel® and Xeon® are registered trademarks of Intel 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, Inc.
Twitter: [@purestorage](#)
[www.purestorage.com](#)

650 Castro Street, Suite #260
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: [sales@purestorage.com](#)
Support: [support@purestorage.com](#)
Media: [pr@purestorage.com](#)