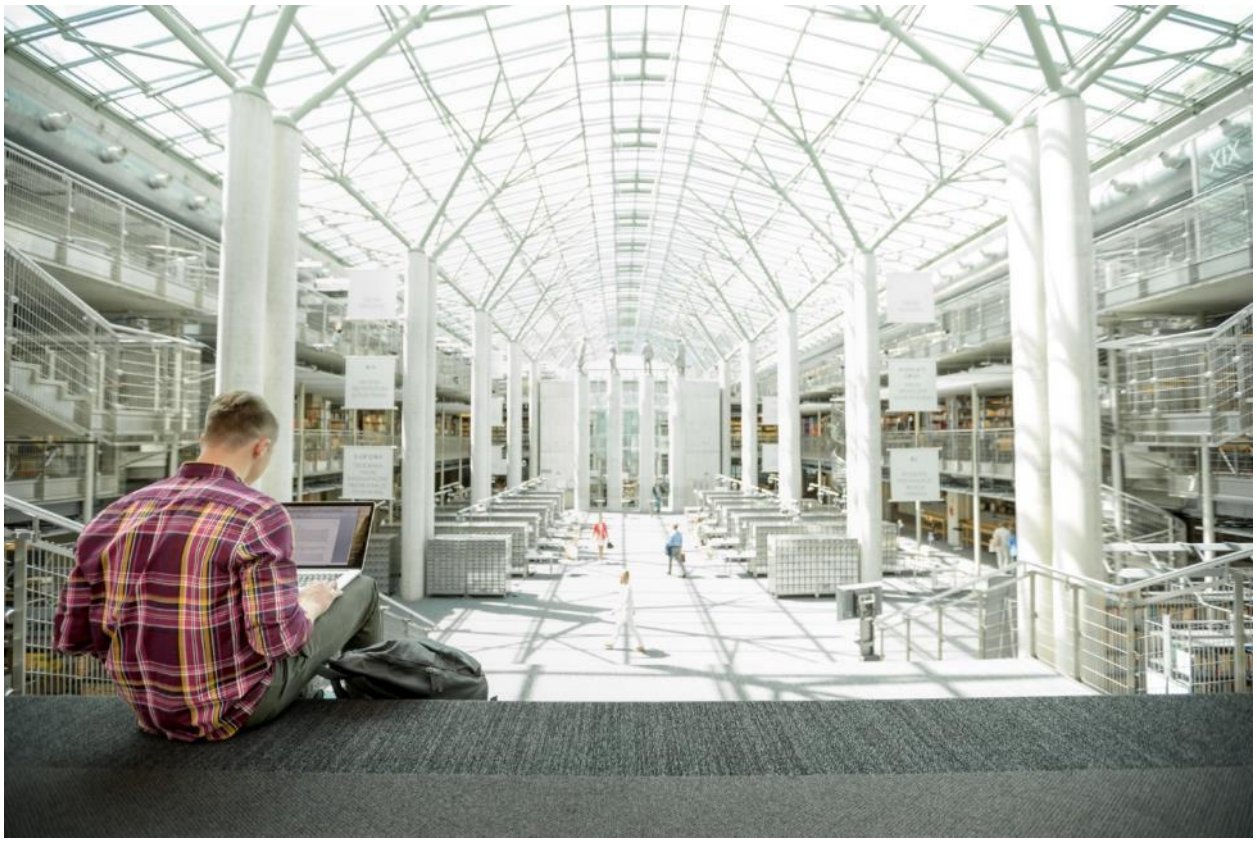

FlashStack Virtual Server Infrastructure for VMware vSphere 6.7 Update 1 Design Guide

Last Updated: July 7, 2020

CISCO VALIDATED DESIGN

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

http://www.cisco.com/go/designzone.

Table of Contents

Executive Summary

Cisco Validated Designs consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

This document discusses the design principles that go into the FlashStack™ solution, which is a validated Converged Infrastructure (CI) jointly developed by Cisco and Pure Storage. The solution is a predesigned, best-practice data center architecture with VMware vSphere built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and Pure Storage FlashArray//X R2 all flash array supporting either iSCSI or Fibre Channel storage access.

The solution architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Server Infrastructure (VSI).

Solution Overview

Introduction

In the current industry there is a trend for pre-engineered solutions which standardize the data center infrastructure, offering the business operational efficiencies, agility and scale to address cloud, bi-modal IT and their business. Their challenge is complexity, diverse application support, efficiency and risk; all these are met by FlashStack with:

- Reduced complexity, automatable infrastructure and easily deployed resources
- Robust components capable of supporting high performance and high bandwidth virtualized applications
- Efficiency through optimization of network bandwidth and in-line storage compression with de-duplication
- Risk reduction at each level of the design with resiliency built into each touch point

Cisco and Pure Storage have partnered to deliver this Cisco Validated Design, which uses best of breed storage, server and network components to serve as the foundation for virtualized workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

In this document we describe a reference architecture detailing a Virtual Server Infrastructure composed of Cisco Nexus switches, Cisco UCS Compute, Cisco MDS Multilayer Fabric Switches, and a Pure Storage FlashArray//X70 R2 delivering a VMware vSphere 6.7 Update1 hypervisor environment.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to modernize their infrastructure to meet their business needs.

Purpose of this Document

This document discusses the design for FlashStack, implemented with either FC or iSCSI, centered around the Cisco UCS 6454 Fabric Interconnect and the Pure Storage FlashArray//X70 R2, delivering a Virtual Server Infrastructure on Cisco UCS B200 M5 Blade Servers and Cisco UCS C220 M5 Rack Servers running VMware vSphere 6.7 U1.

What's New in this Release?

This version of the FlashStack VSI Design introduces the Pure Storage FlashArray//X70 R2 all-flash array along with Cisco UCS B200 M5 Blade Servers featuring the Intel Xeon Scalable Family of CPUs. The design incorporates options for 40Gb iSCSi as well as 32Gb Fibre Channel, both delivered with new design options and features. Highlights for this design include:

- Pure Storage FlashArray//X70 R2
- Pure1 Analyze - VM Analytics
- VMware vSphere 6.7 U1

- VMware Virtual Volumes
- Cisco UCS Manager 4.0
- Cisco UCS B200 M5
- Cisco UCS C220 M5
- Cisco UCS 6454 Fabric Interconnect
- Cisco UCS 1400 Series Virtual Interface Cards
- Cisco Intersight
- Cascade Lake CPU Support (UCS version 4.0(4a) and later)

Solution Summary

FlashStack provides a jointly supported solution by Cisco and Pure Storage. Bringing 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: It 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 resiliency: 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.

Technology Overview

This section provides a technical overview of the compute, network, storage and management components in this solution.

Cisco Unified Computing System

Cisco Unified Computing System™ (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10-100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- Compute - The compute piece of the system incorporates servers based on latest Intel's x86 processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.
- Network - The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lowers costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- Storage access - Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. The use of Policies, Pools, and Profiles allows for simplified storage connectivity management
- Management - The system uniquely integrates compute, network and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired server configurations

Cisco UCS Differentiators

Cisco Unified Computing System has revolutionized the way servers are managed in data center and provides several unique differentiators that are outlined below:

- Embedded Management – Servers in Cisco UCS are managed by embedded software in the Fabric Interconnects.
- Unified Fabric – Cisco UCS uses a wire-once architecture, where a single Ethernet cable is used from the FI from the server chassis for LAN, SAN and management traffic. Adding compute capacity does not require additional connections. This converged I/O reduces overall capital and operational expenses.
- Auto Discovery – The server discovery process is automatic when a blade is inserted into a chassis or a rack server is connected to the Fabric Interconnect
- Policy Based Resource Classification – Once a compute resource is discovered, it can be automatically classified to a resource pool based on configured policies.
- Combined Rack and Blade Server Management – Cisco UCS Manager is hardware form factor agnostic and can manage both blade and rack servers under the same management domain.

- Model based Management Architecture – Cisco UCS Manager architecture and management database is model based. An open XML API is provided to operate on the management model which enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, and Templates – The management approach in Cisco UCS Manager is based on defining policies, pools and templates. This enables a straight forward approach in managing compute, network and storage resources while decreasing the opportunities for misconfigurations
- Policy Resolution – In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy.
- Service Profiles and Stateless Computing – A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource that meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in Multi-Tenancy Support – The combination of a profiles-based approach using policies, pools and templates and policy resolution with organizational hierarchy to manage compute resources makes Cisco UCS Manager suitable for multi-tenant environments.

Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, integrated management for all software and hardware components in Cisco UCS. Using Cisco Single Connect technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive GUI, a CLI, or a through a robust application programming interface (API).

Cisco UCS Manager is embedded into the Cisco UCS Fabric Interconnects and provides a unified management interface that integrates server, network, and storage. Cisco UCS Manger performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers a comprehensive set of XML API for third party integration, exposes thousands of integration points and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Cisco UCS Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point for connectivity and management for the entire Cisco Unified Computing System. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by the Cisco UCS Manager. Cisco UCS FIs provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, SAN and management traffic using a single set of cables

The 4th generation (6400) Fabric delivers options for both high workload density, as well as high port count, with both supporting either Cisco UCS B-Series blade servers, or Cisco UCS C-Series rack mount servers.

Figure 1 Cisco UCS 6454 Fabric Interconnect Model Featured in this Design

The Cisco UCS 6454 Fabric Interconnect is a 54 port 1/10/25/40/100GbE/FCoE switch, supporting 8/16/32Gbps FC ports and up to 3.82Tbps throughput. This model is aimed at higher port count environments that can be configured with 32Gbps FC connectivity to Cisco MDS switches or FC direct attached storage.

Table 1 provides a comparison of the port capabilities of the different Fabric Interconnect models.

Table 1 Cisco UCS 6200 and 6300 Series Fabric Interconnects

Features	6248	6296	6332	6332-16UP	6454
Max 10G ports	48	96	96* + 2**	72* + 16	48
Max 25G ports	N/A	N/A	N/A	N/A	48
Max 40G ports	N/A	N/A	32	24	6
Max 100G ports	N/A	N/A	N/A	N/A	6
Max unified ports	48	96	N/A	16	8
Max FC ports	48x 2/4/8G	96x 2/4/8G	N/A	16x 4/8/16G	8x 8/16/32G

** Using 40G to 4x10G breakout cables* *** Requires QSA module*

Cisco UCS Fabric Extenders

The Cisco UCS Fabric extenders (FEX) or I/O Modules (IOMs) multiplex and forwards all traffic from servers in a blade server chassis to a pair of Cisco UCS Fabric Interconnects over a 10Gbps or 40Gbps unified fabric links. All traffic, including traffic between servers on the same chassis, or between virtual machines on the same server, is forwarded to the parent fabric interconnect where Cisco UCS Manager runs, managing the profiles and polices for the servers. FEX technology was developed by Cisco. Up to two FEXs can be deployed in a chassis.

For more information about the benefits of FEX, see: http://www.cisco.com/c/en/us/solutions/data-center-virtualization/fabric-extender-technology-fex-technology/index.html

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a fundamental building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server architecture. The Cisco UCS blade server chassis uses an innovative unified fabric with fabric-extender technology to lower TCO by reducing the number of network interface cards (NICs), host bus adapters (HBAs), switches, and cables that need to be managed, cooled, and powered. It is a 6-RU chassis that can house up to 8 x half-width or 4 x full-width Cisco UCS B-series blade servers. A passive mid-plane provides up to 80Gbps of I/O bandwidth per server slot and up to 160Gbps for two slots (full-width). The rear of the chassis contains two I/O bays to house a pair of Cisco UCS 2000 Series Fabric Extenders to enable uplink connectivity to FIs for both redundancy and bandwidth aggregation.

Front View

Back View

Cisco UCS VIC

The Cisco UCS Virtual Interface Card (VIC) 1400 Series extends the network fabric directly to both servers and virtual machines so that a single connectivity mechanism can be used to connect both physical and virtual servers with the same level of visibility and control. Cisco VICs provide complete programmability of the Cisco UCS I/O infrastructure, with the number and type of I/O interfaces configurable on demand with a zero-touch model. The VIC presents virtual NICs (vNICs) as well as virtual HBAs (vHBAs) from the same adapter, according to their provisioning specifications within UCSM.

Cisco UCS B200 M5 Servers

The enterprise-class Cisco UCS B200 M5 Blade Server extends the Cisco UCS portfolio in a half-width blade form-factor. This M5 server uses the latest Intel® Xeon® Scalable processors with up to 28 cores per processor, 3TB of RAM (using 24 x128GB DIMMs), 2 drives (SSD, HDD or NVMe), 2 GPUs and 80Gbps of total I/O to each server. It supports the Cisco VIC 1440 adapter to provide 40Gb FCoE connectivity to the unified fabric.

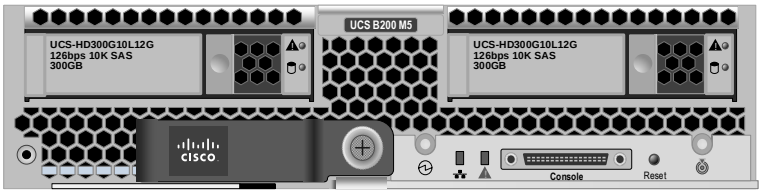

Cisco UCS C220 M5 Servers

The enterprise-class Cisco UCS C220 M5 Rack Server extends the Cisco UCS portfolio in a 1RU rack server. This M5 server uses the latest Intel® Xeon® Scalable processors with up to 28 cores per processor, 3TB of RAM (using 24 x128GB DIMMs), 10 drives (SSD, HDD or NVMe), 2 GPUs and up to 200Gbps of unified I/0 to the server. It supports a variety of other PCIe options detailed in the spec sheet. https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m5-sff-specsheet.pdf

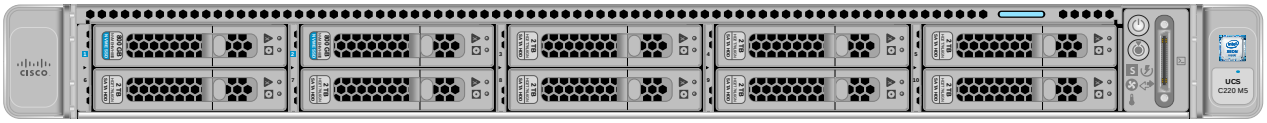

Cisco Nexus 9000 Series Switch

The Cisco Nexus 9000 Series Switches offer both modular and fixed 10/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of non-blocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

Supporting either Cisco ACI or NX-OS, the Nexus delivers a powerful 40/100Gbps platform offering up to 7.2 TBps of bandwidth in a compact 1RU TOR switch. The Nexus featured in this design is the Nexus 9336C-FX2 implemented in NX-OS standalone mode.

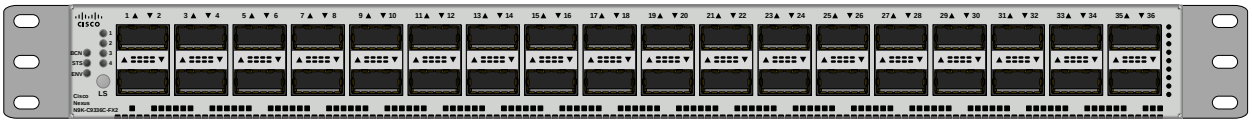

The Nexus 9336C-FX2 implements Cisco Cloud Scale ASICs, giving flexible, and high port density, intelligent buffering, along with in-built analytics and telemetry.

Cisco MDS 9000 Series Switch

The Cisco MDS 9000 family of multilayer switches give a diverse range of storage networking platforms, allowing you to build a highly scalable storage network with multiple layers of network and storage management intelligence. Fixed and modular models implement 4-32Gbps FC, 10-40Gbps FCoE/FCIP, and up to 48 TBps of switching bandwidth.

The MDS 9132T Fibre Channel Switch is featured in this design as the option for the Fibre Channel network

This 1RU switch offers a variable port count, support for NVME over FC, and SAN analytics.

Cisco Intersight

Cisco Intersight gives IT operations management to claimed devices across differing sites, presenting these devices within a unified dashboard. The adaptive management of Intersight provides visibility and alerts to firmware management, showing compliance across managed UCS domains, as well as proactive alerts for upgrade recommendations. Integration with Cisco TAC allows the automated generation and upload of tech support files from the customer.

Pure Storage FlashArray

The Pure Storage FlashArray family delivers purpose-built, software-defined all-flash power and reliability for businesses of every size. FlashArray is all-flash enterprise storage that is up to 10X faster, more space and power efficient, more reliable, and far simpler than other available solutions. Critically, FlashArray also costs less, with a TCO that's typically 50% lower than traditional performance disk arrays.

At the top of the FlashArray line is FlashArray//X – the first mainstream, 100% NVMe, enterprise-class all-flash array. //X represents a higher performance tier for mission-critical databases, top-of-rack flash deployments, and Tier 1 application consolidation. //X, at 1PB in 3U, with hundred-microsecond range latency and GBs of bandwidth, delivers an unprecedented level of performance density that makes possible previously unattainable levels of consolidation.

FlashArray//X is ideal for cost-effective consolidation of everything on flash. Whether accelerating a single database, scaling virtual desktop environments, or powering an all-flash cloud, there is a //X model that fits your needs.

Purity for FlashArray (Purity//FA 5)

At the heart of every FlashArray is Purity Operating Environment software. Purity//FA5 implements advanced data reduction, storage management, and flash management features, enabling organizations to enjoy Tier 1 data services for all workloads, proven 99.9999% availability over two years (inclusive of maintenance and generational upgrades), completely non-disruptive operations, 2X better data reduction versus alternative all-flash solutions, and – with FlashArray//X – the power and efficiency of DirectFlash™. Moreover, Purity includes enterprise-grade data security, comprehensive data protection options, and complete business continuity via ActiveCluster multi-site stretch cluster. All these features are included with every array.

FlashArray//X Specifications

	CAPACITY	PHYSICAL
//X10	Up to 55 TB / 53.5 TiB effective capacity** Up to 20 TB / 18.6 TiB raw capacity	3U 490 – 600 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
//X20	Up to 275 TB / 251.8 TiB effective capacity** Up to 87 TB / 80.3 TiB raw capacity††	3U 620 – 688 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
//X50	Up to 650 TB / 602.8 TiB effective capacity** Up to 183 TB / 171 TiB raw capacity†	3U 620 – 760 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
//X70	Up to 1.3 PB / 1238.5 TiB effective capacity** Up to 366 TB / 320.1 TiB raw capacity†	3U 915 – 1345 Watts (nominal – peak) 97 lbs (44.0 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
//X90	Up to 3 PB / 3003.1 TiB effective capacity** Up to 878 TB / 768.3 TiB raw capacity†	3U – 6U 1100 – 1570 Watts (nominal – peak) 97 lbs (44 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
DIRECT FLASH SHELF	Up to 1.9 PB effective capacity** Up to 512 TB / 448.2 TiB raw capacity	3U 460 - 500 Watts (nominal – peak) 87.7 lbs (39.8kg) fully loaded 5.12" x 18.94" x 29.72" chassis

//X CONNECTIVITY

Onboard Ports (per controller)
- 2 x 1/10/25 Gb Ethernet
- 2 x 1/10/25 Gb Ethernet Replication
- 2 x 1Gb Management Ports

Host I/O Cards (3 slots/controller)
- 2-port 10GBase-T Ethernet
- 2-port 1/10/25 Gb Ethernet
- 2-port 40 Gb Ethernet
- 2 Port 50Gb Ethernet (NVMe-oF Ready)***
- 2-port 16/32 Gb Fibre Channel (NVMe-oF Ready)
- 4-port 16/32 Gb Fibre Channel (NVMe-oF Ready)

* Stated //X specifications are applicable to //X R2 versions.

** Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1 and does not include thin provisioning or snapshots.

*** FlashArray //X currently supports NVMe-oF via RoCEv2 with a roadmap for FC-NVMe and TCP-NVMe.

Evergreen Storage

Customers can deploy storage once and enjoy a subscription to continuous innovation via Pure's Evergreen™ Storage ownership model: expand and improve performance, capacity, density, and/or features for 10 years or more – all without downtime, performance impact, or data migrations. Pure has disrupted the industry's 3-5-year rip-and-replace cycle by engineering compatibility for future technologies right into its products, notably with the NVMe-Ready Guarantee for //M and online upgrade from any //M to //X.

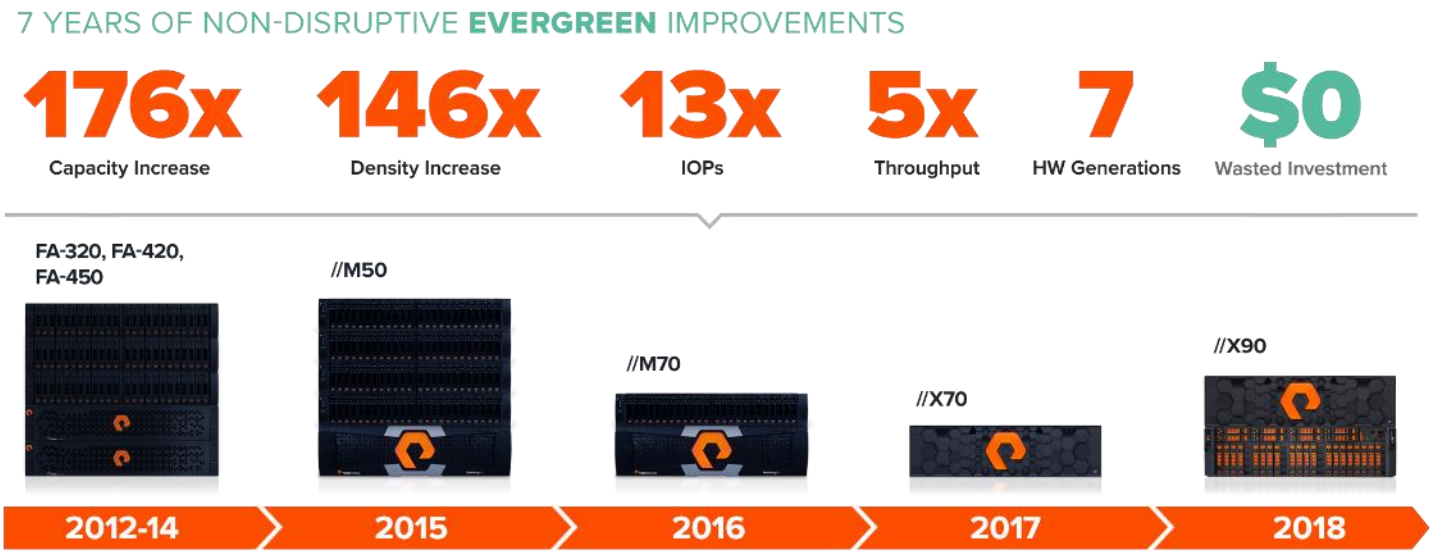

Pure1

Pure1®, our cloud-based management, analytics, and support platform, expands the self-managing, plug-n-play design of Pure all-flash arrays with the machine learning predictive analytics and continuous scanning of Pure1 Meta™ to enable an effortless, worry-free data platform.

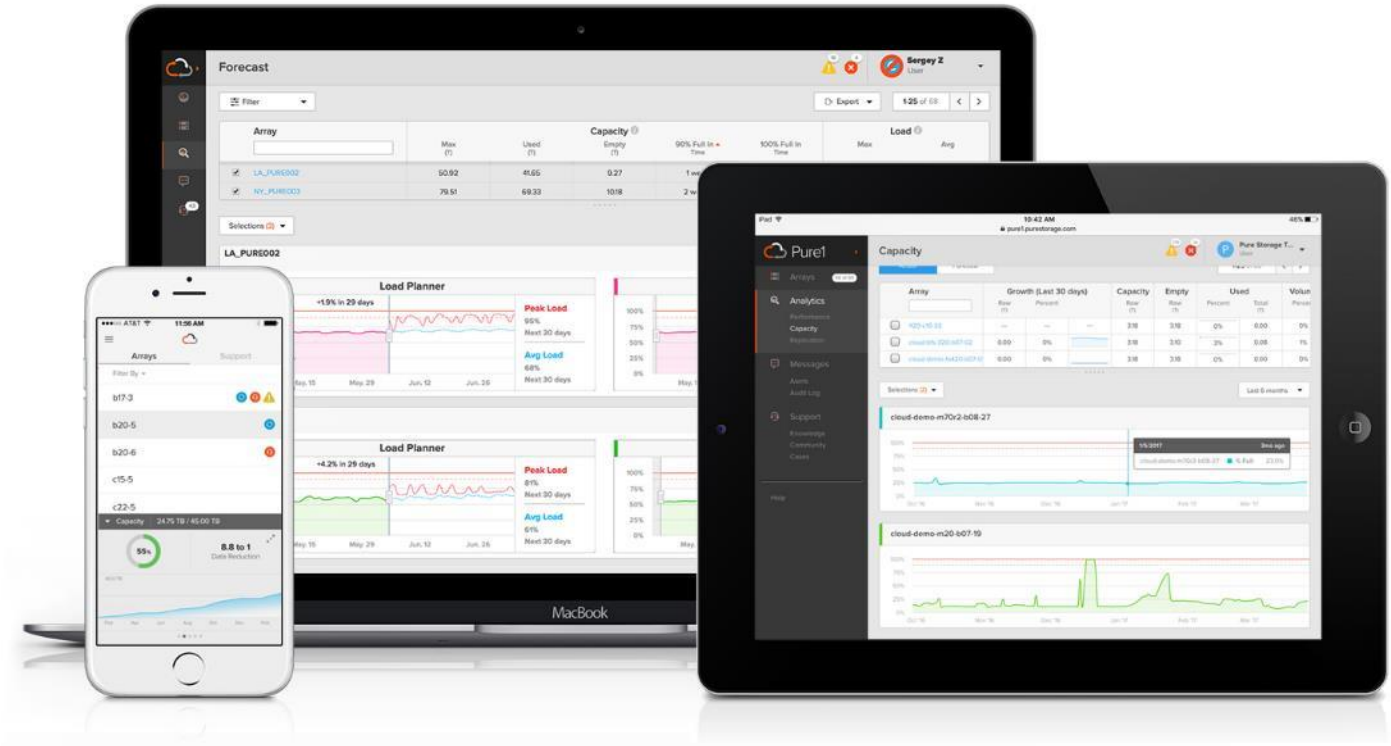

Pure1 Manage

In the Cloud IT operating model, installing and deploying management software is an oxymoron: you simply login. Pure1 Manage is SaaS-based, allowing you to manage your array from any browser or from the Pure1 Mobile App – with nothing extra to purchase, deploy, or maintain. From a single dashboard you can manage all your arrays, with full visibility on the health and performance of your storage.

Pure1 Analyze

Pure1 Analyze delivers true performance forecasting – giving customers complete visibility into the performance and capacity needs of their arrays – now and in the future. Performance forecasting enables intelligent consolidation and unprecedented workload optimization.

Pure1 Support

Pure combines an ultra-proactive support team with the predictive intelligence of Pure1 Meta to deliver unrivaled support that's a key component in our proven FlashArray 99.9999% availability. Customers are often surprised and delighted when we fix issues they did not even know existed.

Pure1 META

The foundation of Pure1 services, Pure1 Meta is global intelligence built from a massive collection of storage array health and performance data. By continuously scanning call-home telemetry from Pure's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues and optimize workloads. The result is both a white glove customer support experience and breakthrough capabilities like accurate performance forecasting.

Meta is always expanding and refining what it knows about array performance and health, moving the Data Platform toward a future of self-driving storage.

VMware vSphere

VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and virtual machines from a single console and aggregates performance monitoring of clusters, hosts, and virtual machines. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, virtual machines, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

VMware vSphere Web Client Plugin

The VMware vSphere Web Client is a cross-platform web application used to interact with the VMware vCenter server – for managing clusters, hosts, and virtual machines in data centers. The Web Client provides full vSphere client functionality including capabilities for configuring hosts, clusters, networks, datastores, or datastore clusters.

The Pure Storage vSphere Web Client Plugin extends the vSphere Web Client with capability for managing Pure Storage FlashArray volumes and snapshots in a vCenter context. The Plugin makes it possible to create, monitor, and manage VMware datastores based on FlashArray volumes (and related snapshots) completely from within the vSphere Web Client.

The Plugin also supports these actions:

- Opening a Purity GUI session in the vSphere Web Client
- Setting a datastore's Path Selection Policy to Round Robin, for multipathing
- Assigning datastores membership in Pure Storage protection groups
- Creating a FlashArray host group for a vSphere cluster
- Initiating or scheduling a space reclamation run on a FlashArray
- Configuring role-based access control for vSphere and FlashArray users
- Working with FlashArray virtual volumes (VVols).

Pure1 Analyze – VM Analytics

With VM Analytics, customers get full-stack visibility from storage all the way up to a virtual machine, there are four basic functions of VM Analytics – you can change which metric is displayed on the node, you can highlight everything that's connected to the node, you can filter what you're seeing, and you can graph the metrics from the node. These are fairly simple actions, but they make a highly versatile product.

At the very core of the product is the topology map which shows Disks, Virtual Machines, Hosts, Datastores, Volumes, and Arrays in the cluster that's selected. By clicking any of the nodes, you can easily visualize what's connected to that object. For example, you can select a virtual machine to see which volume or array it's on. Or you could go the other way and see which virtual machines are living on a particular volume by selecting the volume.

With visibility throughout the stack, you have insight into the latency, bandwidth, and IOPs of your workflows – and the data points you need to quickly resolve issues and pinpoint latency problems or other bottlenecks.

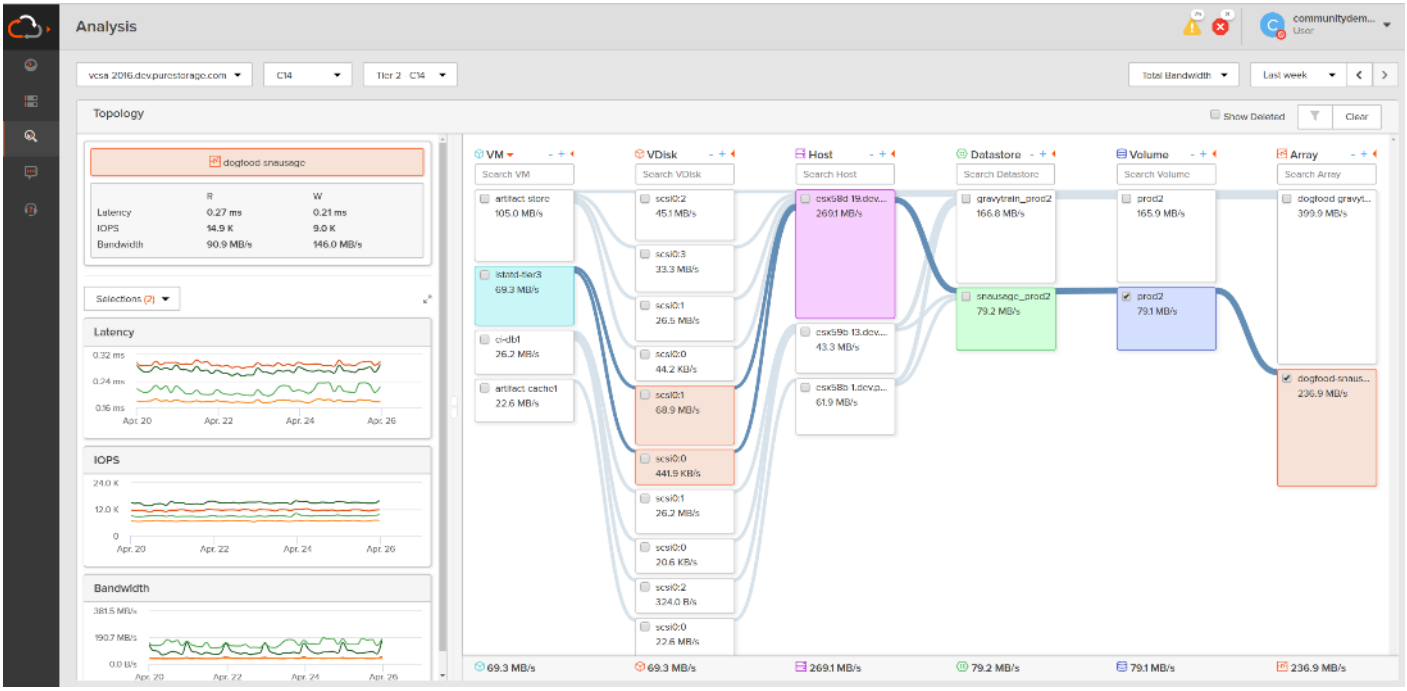
Analysis
communitydem...
User
vcsa-2016.dev.purestorage.com
C14
Tier 2 C14
Total Bandwidth
Last week
Topology
Show Deleted
Clear
dogfood sneusage
R
W
Latency
0.27 ms
0.21 ms
IOPS
14.9 K
9.0 K
Bandwidth
90.9 MB/s
146.0 MB/s
Selections (2)
Latency
0.32 ms
0.24 ms
0.16 ms
Apr. 20
Apr. 22
Apr. 24
Apr. 26
IOPS
24.0 K
12.0 K
0
Bandwidth
381.5 MB/s
190.7 MB/s
0.0 B/s
VM
Search VM
artifact store
105.0 MB/s
istatd-tier3
69.3 MB/s
ci-db1
26.2 MB/s
artifact cache1
22.6 MB/s
VDisk
Search VDisk
scsi0:2
45.1 MB/s
scsi0:3
33.3 MB/s
scsi0:1
26.5 MB/s
scsi0:0
44.2 KB/s
scsi0:1
68.9 MB/s
scsi0:0
441.9 KB/s
scsi0:1
26.2 MB/s
scsi0:0
20.6 KB/s
scsi0:2
324.0 B/s
scsi0:0
22.6 MB/s
Host
Search Host
csk98d 19.dev...
269.1 MB/s
csk99b 13.dev...
43.3 MB/s
csk98b 1.dev.p...
61.9 MB/s
Datastore
Search Datastore
gravytrain_prod2
166.8 MB/s
sneusage_prod2
79.2 MB/s
Volume
Search Volume
prod2
165.9 MB/s
prod2
79.1 MB/s
Array
Search Array
dogfood gravyt...
399.9 MB/s
dogfood-snaus...
236.9 MB/s
69.3 MB/s
69.3 MB/s
269.1 MB/s
79.2 MB/s
79.1 MB/s
236.9 MB/s

Solution Design

The FlashStack Solutions Design implements a Virtual Server Infrastructure that is scalable, reliable, and redundant, using the best practices of Cisco and Pure Storage. This section explains the architecture used in the solution.

Prerequisites

The FlashStack data center is intended to provide a Virtual Server Infrastructure that addresses the primary needs of hosting virtual machines. This design assumes an existing management, security, and routing infrastructure to provide necessary connectivity. This includes but is not limited to the following:

- Out-of-Band management network
- A terminal server for console access
- Layer 3 connectivity to adjacent enterprise network and/or the Internet
- DNS and NTP services

Physical Topology

Compute Connectivity

Each compute chassis in the design is redundantly connected to the managing fabric interconnects with at least two ports per IOM. Ethernet traffic from the upstream network and Fibre Channel frames coming from the FlashArray are converged within the fabric interconnect to be both Ethernet and Fibre Channel over Ethernet, and transmitted to the UCS servers through the IOM. These IOM connections from the UCS Fabric Interconnects to the IOMs are automatically configured as port channels by specifying a Chassis/FEX Discovery Policy within UCSM.

Each rack server in the design is redundantly connected to the managing fabric interconnects with at least one port to each FI. Ethernet traffic from the upstream network and Fibre Channel frames coming from the FlashArray are converged within the fabric interconnect to be both Ethernet and Fibre Channel over Ethernet, and transmitted to the UCS server.

These connections from the UCS 6454 Fabric Interconnect to the 2208XP IOM hosted within the chassis are shown in Figure 2.

Figure 2 Compute Connectivity

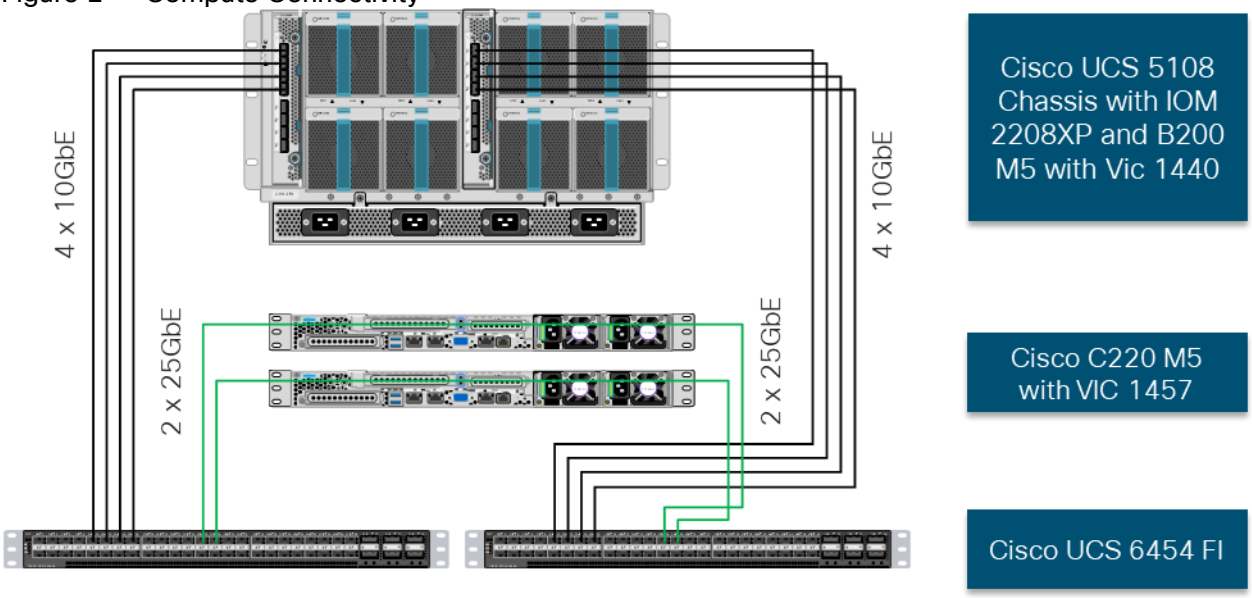

The 2208XP IOM are shown with 4x10Gbps ports to deliver an aggregate of 80Gbps to the chassis, full population of the 2208XP IOM can support 8x10Gbps ports, allowing for an aggregate of 160Gbps to the chassis.

Network Connectivity

The Layer 2 network connection to each Fabric Interconnect is implemented as Virtual Port Channels (vPC) from the upstream Nexus Switches. In the switching environment, the vPC provides the following benefits:

- Allows a single device to use a Port Channel across two upstream devices
- Eliminates Spanning Tree Protocol blocked ports and use all available uplink bandwidth
- Provides a loop-free topology
- Provides fast convergence if either one of the physical links or a device fails
- Helps ensure high availability of the network

The upstream network switches can connect to the Cisco UCS 6454 Fabric Interconnects using 10G, 25G, 40G, or 100G port speeds. In this design, the 100G ports from the 40/100G ports on the 6454 (1/49-54) were used for the virtual port channels. In the iSCSI design, this would also transport the storage traffic between the UCS servers and the FlashArray//X R2.

Figure 3 Network Connectivity

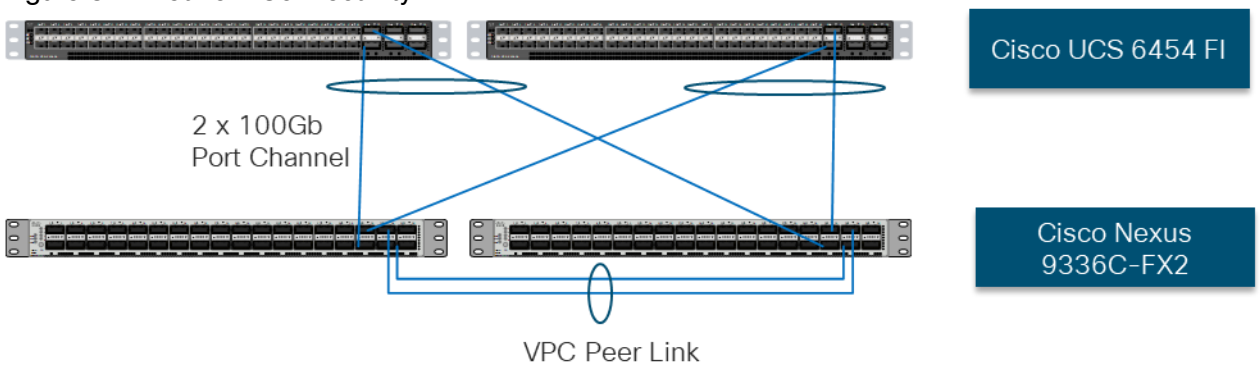

Fibre Channel Storage Connectivity

The //X70 R2 platform is connected through both MDS 9132Ts to their respective Fabric Interconnects in a traditional air-gapped A/B fabric design. The Fabric Interconnects are configured in N-Port Virtualization (NPV) mode, known as FC end host mode in UCSM. The MDS has N-Port ID Virtualization (NPIV) enabled. This allows F-port channels to be used between the Fabric Interconnect and the MDS, providing the following benefits:

- Increased aggregate bandwidth between the fabric interconnect and the MDS
- Load balancing across the FC uplinks
- High availability in the event of a failure of one or more uplinks

Figure 4 Fibre Channel Storage Connectivity

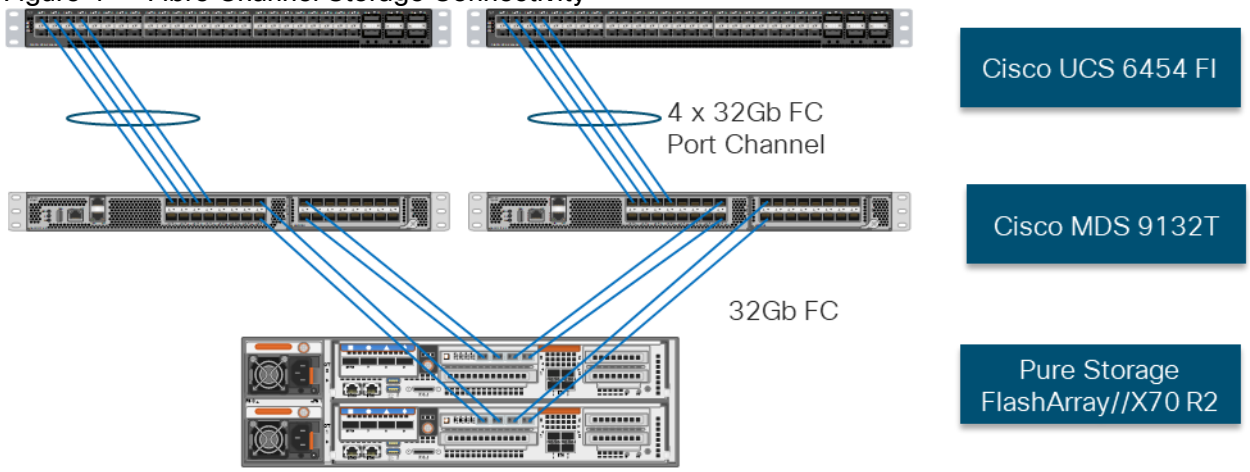

End-to-End Physical Connectivity

FC End-to-End Data Path

The FC end to end path in the design is a traditional air-gapped fabric with identical data path through each fabric as detailed below:

- Each Cisco UCS Server is equipped with a VIC 1400 Series adapter.
- In the Cisco B200 M5 server, a VIC 1440 provides 2x10Gbps to IOM A and 2x10Gbps to IOM B via the Cisco UCS Chassis 5108 chassis backplane

- In the Cisco C220 M5 server, a VIC 1457 is used with 2x 25Gbps connections to FI-A and 2x5Gbps FI-B for a total of 4x25Gbps of uplink bandwidth.
- Each IOM is connected to its respective Cisco UCS 6454 Fabric Interconnect using a port-channel for 4-8 links.
- Each Cisco UCS 6454 FI connects to the MDS 9132T for the respective SAN fabric using an F-Port channel.
- The Pure FlashArray//X70 R2 is connected to both MDS 9132T switches to provide redundant paths through both fabrics.

Figure 5 FC End-to-End Data Path

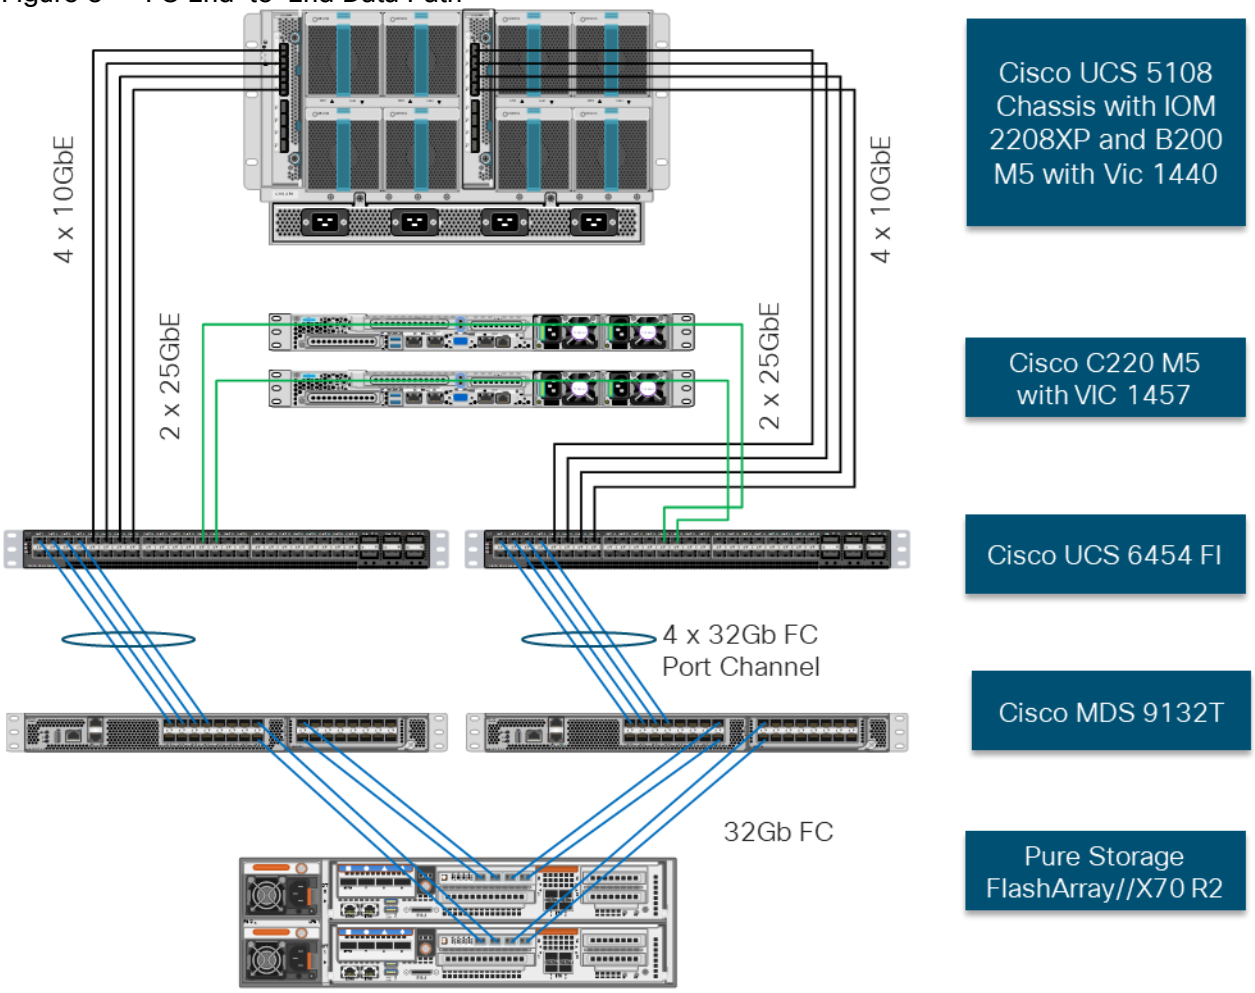

The components of this integrated architecture shown in Figure 5 are:

- Cisco Nexus 9336C-FX2 – 10/25/40/100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6454 Fabric Interconnect – Unified management of UCS compute, and the compute's access to storage and networks.
- Cisco UCS B200 M5 – High powered blade server, optimized for virtual computing.
- Cisco UCS C220 M5 – High powered rack server, optimized for virtual computing

- Cisco MDS 9132T – 32Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- Pure Storage //X70 R2

Direct Attach storage in FC-Switch mode is validated for this solution with UCS Version 4.0(4a) and later

iSCSI End-to-End Data Path

The iSCSI end to path present in the design is leverages the Nexus 9336C-FX2 networking switches to carry storage traffic.

- Each Cisco UCS Server is equipped with a VIC 1400 Series adapter
- In the Cisco B200 M5 server this is the VIC 1440 with 20Gb to IOM A and 20Gbps to IOM B through the Cisco UCS Chassis 5108 chassis backplane
- In the Cisco UCS C220 M5 server is this the VIC 1457 with 4x 10/25Gbps connection
- Each IOM is connected to its respective Cisco UCS 6454 Fabric Interconnect using a port-channel for 4-8 links
- Each Cisco UCS C-Series Server is attached through a 25Gb port to each Cisco UCS 6454 FI
- Each Cisco UCS 6454 FI connects to the Nexus 9336C-FX2 through 2x 100Gbps virtual port channels
- Each controller on the Pure FlashArray//X70 R2 is connected to each Nexus 9336C-FX2 switch through 2x 40Gbps connections to provide redundant paths

Figure 6 iSCSI End-to-End Data Path

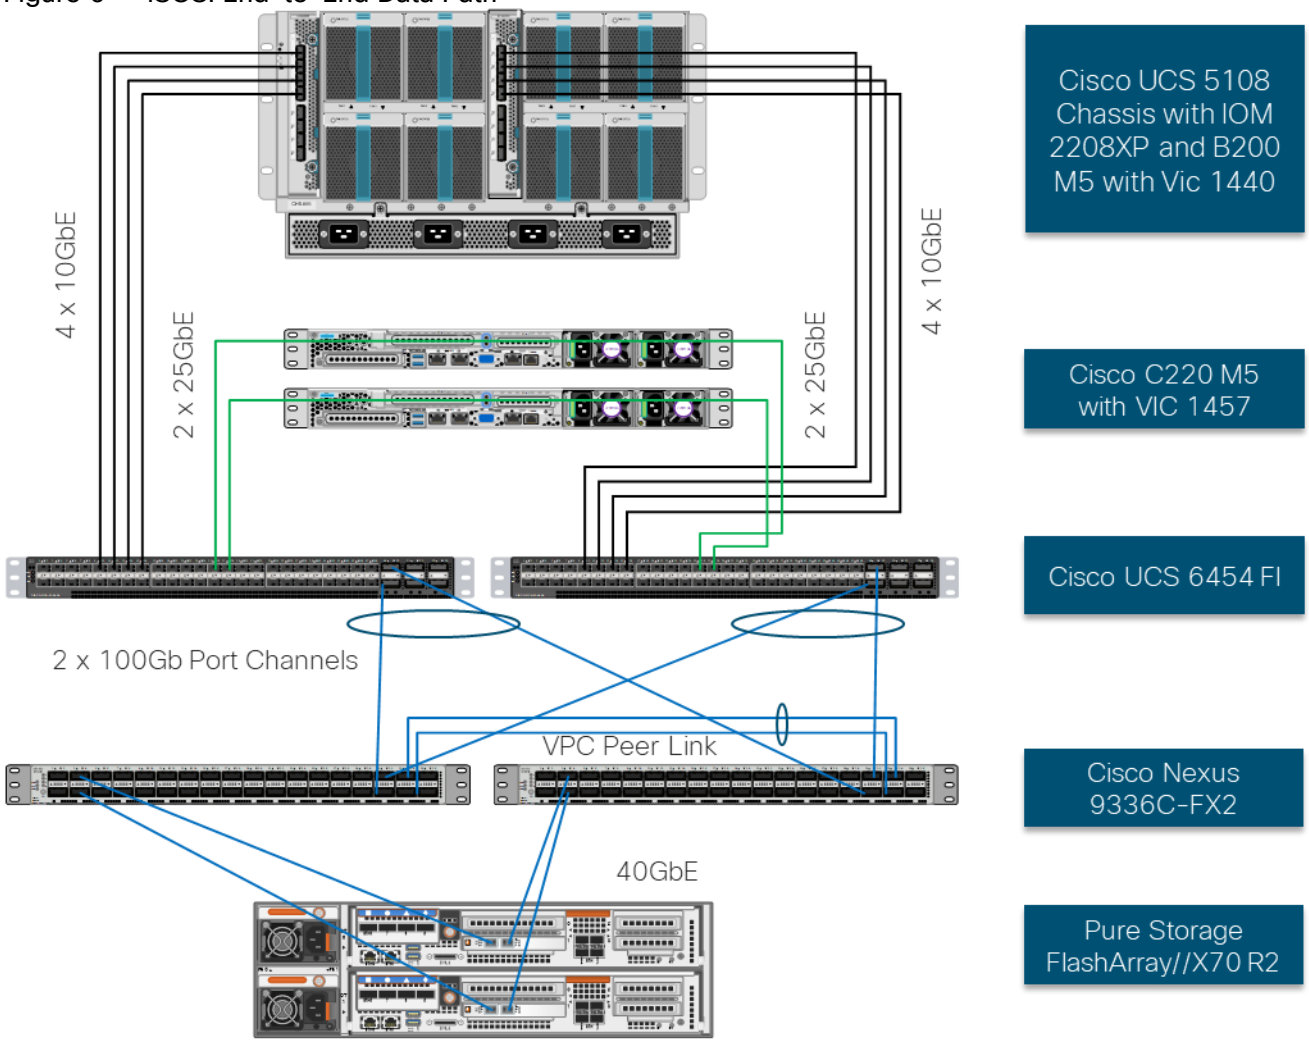

The components of this integrated architecture shown in Figure 6 are:

- Cisco Nexus 9336C-FX2 – 10/25/40/100Gb capable, LAN connectivity to the UCS compute resources and Pure Storage resource
- Cisco UCS 6454 Fabric Interconnect – Unified management of UCS compute, and the compute's access to storage and networks
- Cisco UCS B200 M5 – High powered blade server, optimized for virtual computing
- Cisco UCS C220 M5 – High powered rack server, optimized for virtual computing
- Pure Storage //X70 R2

Direct Attach Storage is validated for this solution with UCS Version 4.0(4a) and later

Compute and Virtualization Layer Design

Server Models

The following topics are the compute options used in this design.

Cisco UCS B-Series

The Cisco UCS B200 M5 servers were selected for this converged infrastructure. Supporting up to 3TB of memory in a half width blade format, these UCS servers are ideal virtualization hosts. These servers are configured in the design with:

- Diskless SAN boot – Persistent operating system installation, independent of the physical blade for true stateless computing
- VIC 1440 – Dual-port 40Gbps capable of up to 256 Express (PCIe) virtual adapters

Cisco UCS C-Series

The Cisco UCS C220 M5 servers were selected for this converged infrastructure. Supporting up to 3TB of memory in a 1RU rack mount format, these UCS servers are ideal virtualization hosts. These servers are configured in the design with:

- Diskless SAN boot – Persistent operating system installation, independent of the physical blade for true stateless computing
- VIC 1457 – Quad-port 10/25Gpbs capable of up to 256 Express (PCIe) virtual adapters

Cisco UCS Service Profiles and Cisco UCS Service Profile Templates

Cisco UCS Service Profiles (SP) were configured with identity information pulled from pools (WWPN, MAC, UUID, etc) as well as policies covering firmware to power control options. These SP are provisioned from UCS Service Profile Templates that allow rapid creation, as well as guaranteed consistency of the hosts at the UCS hardware layer.

Cisco UCS vNIC Templates

Virtual Network Interface cards are created as virtual adapters from the UCS VICs to present the installed operating system. vNIC Templates provide a repeatable, reusable, and adjustable sub-component of the SP template for handling these vNICs. These vNIC templates were adjusted with the options for:

- Fabric association or failover between fabrics
- VLANs that should be carried
- Native VLAN specification
- VLAN and setting consistency with another vNIC template
- vNIC MTU
- MAC Pool specifications

VMware vSphere

This design is implemented using vSphere 6.7 for the vCenter and ESXi version. This design leverages VMware Virtual Volumes (VVols) for storage integration.

Virtual Networking Configuration

The virtual networking configuration utilizes the UCS VIC adapter to present multiple adapters too ESXi as shown in Figure 7.

Figure 7 Virtual Networking Configuration

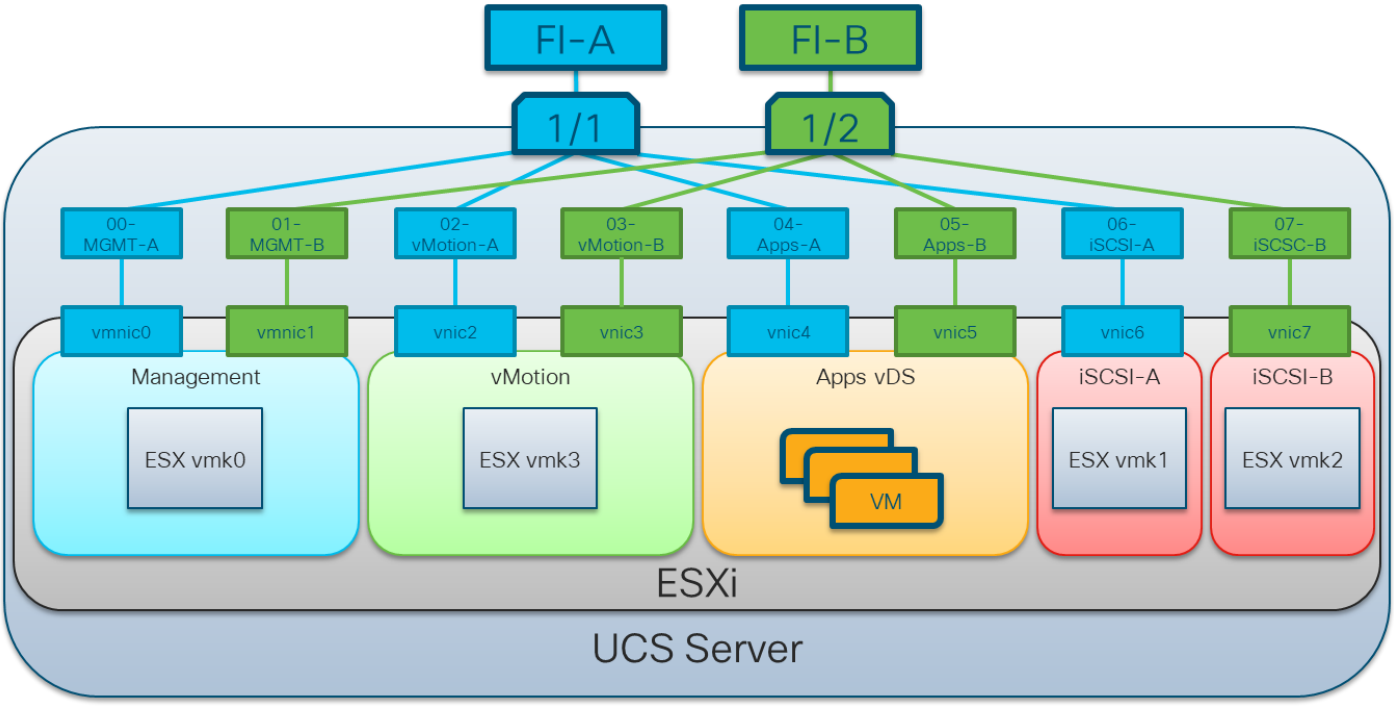

Adapters are created in pairs to provide redundancy to the vSwitches. One pair is created for management, one pair for vmotion, one pair for application data. In the Fibre Channel design, a pair of vHBAs will be created. In the iSCSI design an additional pair of vNICs will be created. VMware vDS (virtual distributed switches) are configured for Infrastructure and for Application traffic, with iSCSI traffic carried within standard vSwitches.

Virtual Volumes Configuration

VMware Virtual Volumes (VVols), first introduced in VMware vSphere version 6.0, is a storage technology that provides policy-based, granular storage configuration and control of virtual machines. Through API-based interaction with underlying storage, VMware administrators can maintain storage configuration compliance using only native VMware interfaces.

Version 5.x of Purity//FA software introduced support for FlashArray-based vSphere Virtual Volumes (VVols). The accompanying FlashArray Plugin for the vSphere Web Client (the Plugin) makes it possible to create, manage, and use VVols that are based on FlashArray volumes from within the Web Client.

Figure 8 Virtual Volumes Configuration

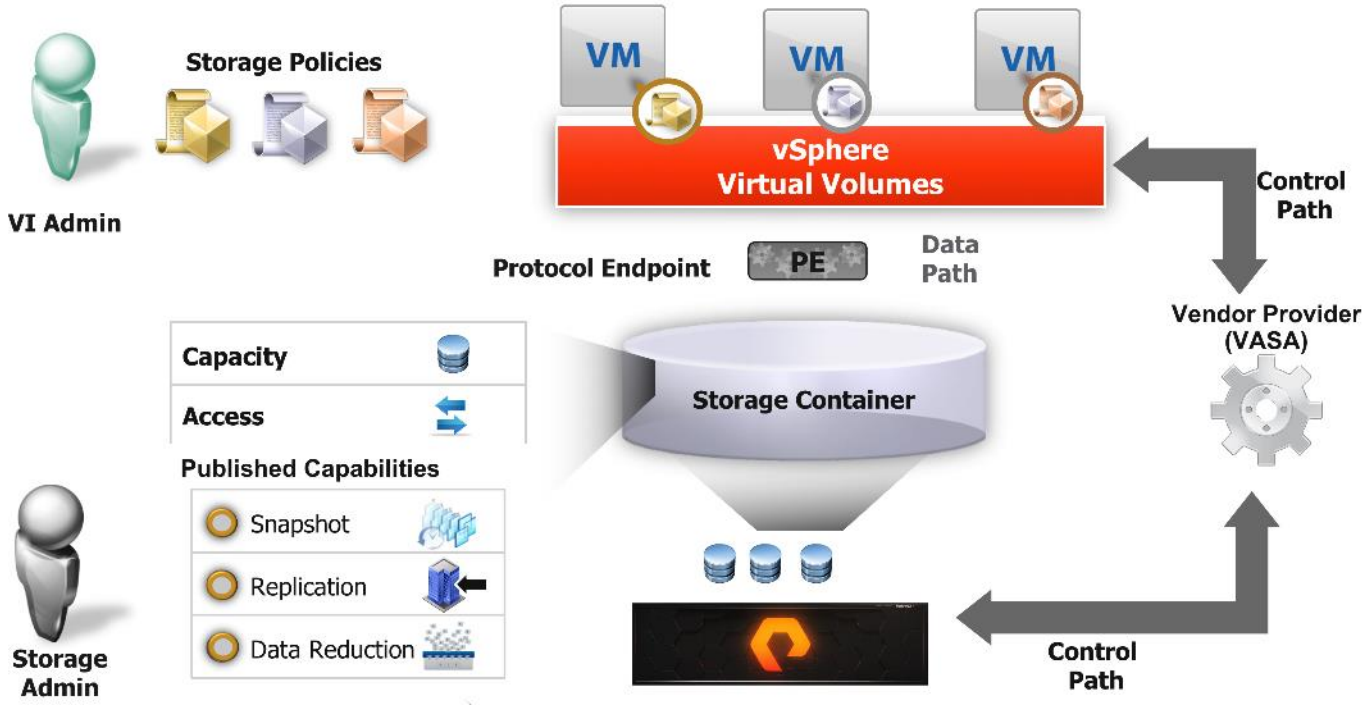

To start using VVols with the Pure Storage FlashArray, the Array's Storage Providers must be registered in vCenter. The Protocol Endpoint is then connected to the Host Group, and finally the VVol Datastore is created.

- Register the FlashArray's Storage Providers

  The two main methods to register the Storage Providers is to manually add them in vCenter or to use the Pure Storage vSphere Web Client Plugin to register the storage provider for the given FlashArray. Pure Storage recommends using the Plugin to register the FlashArray's Storage Providers. CT0.ETH0 and CT1.ETH0

- Connect Protocol Endpoint to Host Group/s

  Users will need to connect the protocol endpoint when manually creating the VVol Datastore in vCenter. This must be done via the FlashArray CLI. Should the Pure Storage Plugin be used, the plugin will automatically connect the PE to the host group that correlates to the ESXi Cluster that is used to create the VVol Datastore with the Plugin.

- Create Virtual Volume Storage Container (Datastore)

- The VVol Datastore can be created manually or created with the Pure Storage Plugin. When creating the VVol Datastore mount the VVol Datastore to all hosts in the ESXi Cluster.

- Storage Policy Based Management (SPMB)

  Storage Policies can be applied to VVols virtual machines that can automatically place the volumes for the VM into FlashArray protection groups. Policies can be imported from existing FlashArray protection group policies or can be created in vCenter to match the requirements of existing Protection Groups.

For more information about the implementation of VMware Virtual Volumes with Pure Storage FlashArray, refer to the Web Guide: Implementing vSphere Virtual Volumes with FlashArray.

Network Design

Nexus Configuration

The Nexus configuration explains the basic networking needs within the stack for Layer 2 to Layer 3 communication.

The following NX-OS features are implemented within the Nexus switches for the design:

- feature lacp – Allows for the utilization of Link Aggregation Control Protocol (802.3ad) by the port channels configured on the switch. Port channeling is a link aggregation technique offering link fault tolerance and traffic distribution (load balancing) for improved aggregate bandwidth across member ports.
- feature vpc – Allows for two Nexus switches to provide a cooperative port channel called a virtual Port Channel (vPC). vPCs present the Nexus switches as a single "logical" port channel to the connecting upstream or downstream device. Configured vPCs present the same fault tolerance and traffic distribution as a port-channel but allow these benefits to occur across links distributed between the two switches. Enablement of vPCs will require a connecting vPC Peer-Link between the two switches, as well as an out of band vPC keep alive to handle switch isolation scenarios.

Management Connectivity

Out-of-band management is handled by an independent switch that could be one currently in place in the customer's environment. Each physical device had its management interface carried through this Out-of-band switch, with in-band management carried as a differing VLAN within the solution for ESXi, vCenter and other virtual management components.

Jumbo Frames

An MTU of 9216 was configured at all network levels to allow for the usage of jumbo frames as needed by the guest OS and application layer.

Storage Design

LUN Multipathing per HBA and Different Pathing Options

This design implements Single Initiator-Multi Target (SI-MT) zoning in conjunction with single vHBAs per fabric on the Cisco UCS infrastructure. This means that each vHBA within Cisco UCS will see multiple paths on their respective fabric to each LUN. Using this design requires the use of Cisco Smart Zoning within the MDS switches.

Different pathing options including Single Initiator-Single Target (SI-ST) are supported. This option would increase the administrative work and exposure to user error.

Balance your bandwidth and application needs, vSphere cluster utilization, and availability requirements when evaluating alternative pathing options to deploy this solution.

Smart Zoning

Zoning is set for Single Initiator/Multiple Target (SI-MT) to optimize traffic intended to be specific to the initiator (UCS host vHBA) and the targets (VSP controller ports). Using SI-MT zoning provides reduced administrative overhead versus configuring single initiator/single target zoning, and results in the same SAN switching efficiency when configured with Smart Zoning.

For more information about smart zoning, go to section Configuring and Managing Zones of the Cisco MDS 9000 Series Fabric Configuration Guide, Release 8.x.

Design Considerations

Cisco Nexus 9000 Series vPC Best Practices

The following Cisco Nexus 9000 design best practices and recommendations were used in this design.

vPC Peer Keepalive Link Considerations

- It is recommended to have a dedicated layer 3 link for vPC peer keepalive, followed by out-of-band management interface (mgmt0) and lastly, routing the peer keepalive link over an existing Layer3 infrastructure between the existing vPC peers.
- vPC peer keepalive link should not be routed over a vPC peer-link.
- The out-of-band management network is used as the vPC peer keepalive link in this design.

vPC Peer Link Considerations

- Only vPC VLANs are allowed on the vPC peer-links. For deployments that require non-vPC VLAN traffic to be exchanged between vPC peer switches, deploy a separate Layer 2 link for this traffic.
- Only required VLANs are allowed on the vPC peer links and member ports – prune all others to minimize internal resource consumption.

Cisco UCS Fabric Interconnect (FI) Best Practices

The following Cisco UCS Fabric Interconnect design best practices and recommendations were used in this design.

Ethernet End-Host Mode

- This is the default switch mode for the UCS Fabric Interconnect.
- In this mode, the FI will only learn MAC addresses from devices connected on Server and Appliance ports
- In this mode, the FI does not run spanning-tree and handles loop avoidance using a combination of Deja-Vu check and Reverse Path Forwarding (RFP). Storage Considerations

Boot from SAN

When utilizing UCS Server technology it is recommended to configure Boot from SAN and store the boot partitions on remote storage, this enabled architects and administrators to take full advantage of the stateless nature of service profiles for hardware flexibility across lifecycle management of server hardware generational changes, Operating Systems/Hypervisors and overall portability of server identity. Boot from SAN also removes the need to populate local server storage creating more administrative overhead.

Pure Storage FlashArray Considerations

The following are some of the Pure Storage FlashArray considerations:

- Make sure Each FlashArray Controller is connected to BOTH storage fabrics (A/B).
- Within Purity it's best practice to map Hosts to Host Groups and then Host Groups to Volumes, this ensures the Volume is presented on the same LUN ID to all hosts and allows for simplified management of ESXi Clusters across multiple nodes.

- How big should a Volume be? When managing storage, you create a volume based on the size required, availability and performance are taken care of via RAID-HD and DirectFlash Software. As an administrator, you can create 1 10TB volume or 10 1TB Volumes and their performance/availability will be the same, so instead of creating volumes for availability or performance you can think about recoverability, manageability and administrative considerations. Like what data do I want to present to this application or what data do I want to store together so I can replicate it to another site/system/cloud.

Port Connectivity

10/25/40GbE connectivity support – while both 10 and 25 Gbps is provided via 2 onboard NICs on each FlashArray controller if any more interfaces are required or if 40GbE connectivity is also required then make sure there is provision for additional NICs have been included in the original FlashArray BOM.

16/32Gb Fiber Channel support (N-2 support) – Pure Storage offer up to 32Gb FC support on the latest FlashArray//X series arrays. Always make sure the correct number of HBAs and the speed of SFPs are included in the original FlashArray BOM.

Oversubscription

To reduce the impact of an outage or maintenance scheduled downtime it is good practice when designing fabrics to provide oversubscription of bandwidth, this enables a similar performance profile during component failure and protects workloads from being impacted by a reduced number of paths during a component failure or maintenance event. Oversubscription can be achieved by increasing the number of physically cabled connections between storage and compute. These connections can then be utilized to deliver performance and reduced latency to the underlying workloads running on the solution.

Topology

When configuring your SAN, it's important to remember that the more hops you have, the more latency you will see. For best performance, the ideal topology is a "Flat Fabric" where the FlashArray is only one hop away from any applications being hosted on it. For iSCSI, we recommend that you do not add routing to your storage LAN.

VMware Virtual Volumes Considerations

When configuring a Pure Storage FlashArray with Virtual Volumes, the FlashArray will only be able to provide the VASA Service to an individual vCenter at this time. vCenters that are in Enhanced Linked Mode will each be able to communicate with the same FlashArray, however vCenters that are not in Enhanced Linked Mode cannot both use VVols on the same FlashArray. Should multiple vCenters need to use the same FlashArray for VVols, they should be configured in Enhanced Linked Mode.

Ensure that the Config VVol is either part of an existing FlashArray Protection Group, Storage Policy that includes snapshots or manual snapshots of the Config VVol are taken. This will help with the virtual machine recovery process if the virtual machine is deleted.

Keep in mind that there are some FlashArray limits on Volume Connections per Host, Volume Count and Snapshot Count. For more details about FlashArray limits, review the following: https://support.purestorage.com/FlashArray/PurityFA/General_Troubleshooting/Pure_Storage_FlashArray_Limits

When a Storage Policy is applied to a VVol virtual machine, the Volumes associated with that virtual machine are added to the designated protection group when applying the policy to the virtual machine. Should replication be part of the policy, be mindful of the amount of virtual machines using that storage policy and replication group. A large amount of virtual machines with a high change rate could cause replication to miss its schedule due to increases replication bandwidth and time needed to complete the scheduled snapshot. Pure Storage recommends VVol virtual machines that have Storage Policies applied be balanced between protection groups. Currently Pure Storage recommends 20 to 30 virtual machines per Storage Policy Replication Group.

Pure Storage FlashArray Best Practices for VMware vSphere

The following Pure Storage best practices for VMware vSphere should be followed as part of a design:

- For hosts earlier than 6.0 Patch 5 or 6.5 Update 1, Configure Round Robin and an I/O Operations Limit of 1 for every FlashArray device. This is no longer needed for later versions of ESXi. The best way to do this is to create an ESXi SATP Rule on every host (below). This will make sure all devices are set automatically.

  ```
  esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "PURE" -M "FlashArray" -P "VMW_PSP_RR" -O "iops=1"
  ```

- For iSCSI, disable DelayedAck and set the Login Timeout to 30 seconds. Jumbo Frames are optional.
- In vSphere 6.x, if hosts have any VMFS-5 volumes, change EnableBlockDelete to enabled. If it is all VMFS-6, this change is not needed.
- For VMFS-5, Run UNMAP frequently.
- For VMFS-6, keep automatic UNMAP enabled.
- When using vSphere Replication and/or when you have ESXi hosts running EFI-enabled VMs set the ESXi parameter Disk.DiskMaxIOSize to 4 MB.
- DataMover.HardwareAcceleratedMove, DataMover.HardwareAcceleratedInit, and VMFS3.HardwareAcceleratedLocking should all be enabled.
- Ensure all ESXi hosts are connected to both FlashArray controllers. Ideally at least two paths to each. Aim for total redundancy.
- Install VMware tools whenever possible.
- Queue depths should be left at the default. Changing queue depths on the ESXi host is considered to be a tweak and should only be examined if a performance problem (high latency) is observed.
- When mounting snapshots, use the ESXi resignature option and avoid force-mounting.
- Configure Host Groups on the FlashArray identically to clusters in vSphere. For example, if a cluster has four hosts in it, create a corresponding Host Group on the relevant FlashArray with exactly those four hosts—no more, no less.
- Use Paravirtual SCSI adapters for virtual machines whenever possible.
- Atomic Test and Set (ATS) is required on all Pure Storage volumes. This is a default configuration and no changes should normally be needed.
- UseATSForHBOnVMFS5 should be enabled. This was introduced in vSphere 5.5 U2 and is enabled by default. It is NOT required though.

For more information about the VMware vSphere Pure Storage FlashArray Best Practices, refer to the Web Guide: FlashArray® VMware Best Practices.

Pure Storage FlashArray Best Practices for VMware Virtual Volumes (VVols)

Along with the above Pure Storage Best Practices for VMware vSphere the following should be considered as part of a design that includes the implementation of VVols as part of the Solution:

- Create a Local FlashArray Array Admin user to register the storage provider with vs using the local pureuser account, vvols-admin for example.
- Use the Round Robin pathing policy (default) for the Protocol Endpoint.
- Use the Pure Storage vSphere Web Client Plugin to register the FlashArray storage provider and create the VVols Datastore if possible.
- If manually registering the storage providers, Register both controllers' storage providers with CT0.ETH0 and CT1.ETH0
- If manually registering the storage providers, you will need to connect the protocol endpoint to Host Groups vs individual hosts if manually connecting the PE to hosts from the FlashArray CLI.
- A single PE should be sufficient for the design utilizing the default device queue depth for the PE.
- Keep VM Templates on VVols when deploying new VVol virtual machines from a template.
- When resizing a virtual machine's VMDK that resides on a VVol is required, complete the task from vSphere vCenter GUI and not the FlashArray GUI.
- Have all ESXi Hosts, vCenter Server and FlashArray synced to the same NTP Server.

For more information about VVols Best practices, refer to the following KB article:
https://support.purestorage.com/Solutions/VMware_Platform_Guide/003Virtual_Volumes_-_VVols/VVols_Best_Practices

Validation

Test Plan

The solution was validated by deploying virtual machines running the vdbench tool. The system was validated for resiliency by failing various aspects of the system under load. Examples of the types of tests executed include:

- Failure and recovery of fibre channel booted ESXi hosts in a cluster
- Rebooting of fibre channel booted hosts
- Failure and recovery of redundant links to Flash Array controllers from MDS switches for Fibre Channel
- Service Profile migration between blades
- Failure of partial and complete IOM links to Fabric Interconnects
- Failure and recovery of iSCSI booted ESXi hosts in a cluster
- Rebooting of iSCSI channel booted hosts
- Failure and recovery of redundant links to Flash Array controllers from Nexus switches for iSCSI
- Failure and recovery of a Nexus switch
- Removal and recovery of FlashArray//X R2 Controller

Validated Hardware

Table x describes the hardware and software versions used during solution validation. It is important to note that Cisco, Pure Storage, and VMware have compatibility matrixes that should be referenced to determine support and are available in the AppendixAppendix.

Table 2 Validated Hardware

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	7.0(3)I7(5)
Compute	Cisco UCS Fabric Interconnect 6454	4.0(2b)
	Cisco UCS 2208XP IOM	4.0(2b)
	Cisco UCS B200 M5	4.0(2b)
	VMware vSphere	6.7 U1 VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
	ESXi 6.7 U1 nenic	1.0.26.0

Component		Software Version/Firmware Version
	ESXi 6.7 U1 nfnic	4.0.0.20
	Virtual Machine Virtual Hardware Version	14
Storage	Pure Flash Array //X70 R2	5.1.9
	Cisco MDS 9132T	8.3(2)

Summary

FlashStack delivers a platform for Enterprise and cloud datacenters using Cisco UCS Blade Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches and Fibre Channel, or iSCSI attached Pure Storage FlashArray//X70 R2. FlashStack is designed and validated using compute, network and storage best practices for high performance, high availability, and simplicity in implementation and management.

This CVD validates the design, performance, management, scalability, and resilience that FlashStack provides to customers.

Appendix

Compute

Cisco Unified Computing System:

http://www.cisco.com/en/US/products/ps10265/index.html

Cisco UCS 6300 Series Fabric Interconnects:

http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html

Cisco UCS 6400 Series Fabric Interconnects:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html

Cisco UCS 5100 Series Blade Server Chassis:

http://www.cisco.com/en/US/products/ps10279/index.html

Cisco UCS 2300 Series Fabric Extenders:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-675243.html

Cisco UCS 2200 Series Fabric Extenders:
https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/data_sheet_c78-675243.html

Cisco UCS B-Series Blade Servers:

http://www.cisco.com/en/US/partner/products/ps10280/index.html

Cisco UCS VIC Adapters:

http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

http://www.cisco.com/en/US/products/ps10281/index.html

Network and Management

Cisco Nexus 9000 Series Switches:

http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html

Cisco Nexus 9000 vPC Configuration Guide:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/interfaces/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Interfaces_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Interfaces_Configuration_Guide_chapter_0111.html

Cisco Data Center Network Manager:

https://www.cisco.com/c/en/us/products/cloud-systems-management/prime-data-center-network-manager/datasheet-listing.html

Storage

Pure Storage FlashArray//X:

https://www.purestorage.com/products/flasharray-x.html

Pure FlashStack Compatibility Matrix:

https://support.purestorage.com/FlashStack/Product_Information/FlashStack_Compatibility_Matrix

https://support.purestorage.com/FlashArray/Getting_Started/Compatibility_Matrix

https://support.purestorage.com/FlashArray/FlashArray_Hardware/99_General_FA_HW_Troubleshooting/FlashArray_Transceiver_and_Cable_Support

Cisco MDS 9000 Series Multilayer Switches:

http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html

Pure Storage VMware Integration:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/001VMwareBestPractices/hhhWeb_Guide%3A_FlashArray_VMware_Best_Practices

vSphere Virtual Volumes

https://support.purestorage.com/Solutions/VMware_Platform_Guide/003Virtual_Volumes_-_VVols/Guides_and_How_To's/Web_Guide%3A_Implementing_vSphere_Virtual_Volumes_with_FlashArray

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/virtualvolumes/vmw-vsphere-virtual-volumes-getting-started-guide.pdf

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/virtualvolumes/vmware-virtual-volumes-faq.pdf

Virtualization Layer

VMware vCenter Server:

http://www.vmware.com/products/vcenter-server/overview.html

VMware vSphere:

https://www.vmware.com/products/vsphere

Compatibility Matrixes

Cisco UCS Hardware Compatibility Matrix:

https://ucshcltool.cloudapps.cisco.com/public/

Cisco Nexus Recommended Releases for Nexus 9K:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html

Cisco MDS Recommended Releases:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html

Cisco Nexus and MDS Interoperability Matrix:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx/Matrix1.html

VMware and Cisco Unified Computing System:

http://www.vmware.com/resources/compatibility

About the Authors

Allen Clark, Technical Marketing Engineer, Cisco Systems, Inc.

Allen Clark has over 15 years of experience working with enterprise storage and data center technologies. As a member of various organizations within Cisco, Allen has worked with hundreds of customers on implementation and support of compute and storage products. Allen holds a bachelor's degree in Computer Science from North Carolina State University and is a dual Cisco Certified Internetwork Expert (CCIE 39519, Storage Networking and Data Center)

Craig Waters, Solutions Architecture / Product Management, Pure Storage Inc.

Craig Waters has over 20 years of experience in the ICT industry working in the Customer, Integrator and Vendor spaces. Craig has specialized in data center technologies across Compute, Networking and Storage Infrastructure including Virtualization and Container platforms supporting Business Critical Applications. Craig has a wide breadth of experience building architectural designs taking Availability, Manageability, Performance, Recoverability and Security requirements into consideration. More recently Craig has been involved in Product Management bringing key Pure Storage Cisco integrations to market leveraging his customer experience background.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.
- Alex Carver, Solutions Architect, Pure Storage Inc.