

FlashStack for OpenShift 4.3, a User-Provisioned Infrastructure Design Guide

Published: June 2020



In partnership with:



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>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2020 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction.....	6
Audience	6
Purpose of this Document.....	6
What's New in this Release?	7
Solution Summary	7
Technology Overview	9
Cisco Unified Computing System	9
Cisco UCS Differentiators	9
Cisco Nexus 9000 Series Switch	12
Pure Storage FlashArray.....	14
Purity for FlashArray (Purity//FA 5)	14
FlashArray//X Specifications	15
Evergreen™ Storage	17
Pure Service Orchestrator	17
Pure1®	19
Pure1 Manage	19
Pure1 Analyze	19
Pure1 Support.....	19
VMware vSphere	20
VMware vSphere vCenter	20
VMware vSphere Web Client Plugin	20
RedHat OpenShift Container Platform	20
RedHat OpenShift Architecture	20
Solution Design	23
Prerequisites.....	23
Physical Topology.....	23
Compute Connectivity.....	23
Network Connectivity	24
End-to-End Physical Connectivity.....	25
iSCSI End-to-End Data Path.....	25
Compute and Virtualization Layer Design	26
Server Models.....	27
Cisco UCS Service Profiles and Cisco UCS Service Profile Templates	28

Cisco UCS vNIC Templates	28
VMware vSphere.....	28
Virtual Networking Configuration	28
Network Design	29
Cisco Nexus Configuration.....	29
Management Connectivity.....	29
Jumbo Frames	30
Container Infrastructure Virtualization Layer Design.....	30
Container Networking Configuration.....	30
Container Storage Configuration	30
Design Considerations.....	31
Network Considerations	31
Storage Considerations.....	32
Validation.....	35
Test Plan.....	35
Validated Hardware	35
Summary	37
Appendix.....	38
Compute.....	38
Network and Management	38
Storage.....	39
Pure Service Orchestrator	39
Virtualization Layer	39
Compatibility Matrixes.....	39
About the Authors.....	41
Acknowledgements	41



Executive Summary

Red Hat® OpenShift® is an enterprise-ready Kubernetes container platform with full-stack automated operations to manage hybrid cloud and multi-cloud deployments. Red Hat OpenShift is optimized to improve developer productivity and promote innovation.

Red Hat OpenShift includes everything needed for hybrid cloud, enterprise container and Kubernetes development and deployments. It includes an enterprise-grade Linux® operating system, container runtime, networking, monitoring, container registry, authentication, and authorization solutions.

Error! Reference source not found.Error! Reference source not found.Error! Reference source not found.Error! Reference source not found.Combining Red Hat OpenShift with FlashStack, the converged infrastructure from Cisco and Pure Storage, can simplify the deployment and the management of your infrastructure. You can also benefit from improved efficiency, better data protection, lower risk, and the flexibility to scale this highly available enterprise-grade infrastructure stack to accommodate new business requirements and other changes over time.

This combined solution helps organizations achieve the speed, flexibility, security and scale required for all their application modernization and digital transformation initiatives. For many companies, this means adopting micro-services and DevOps practices and moving away from monolithic application development practices as a first step in creating more business value from their applications.

To help customers, business partners, and other interested parties with their digital transformation and to enhance their cloud-native and application modernization practices, Cisco, and Pure Storage partnered to produce this Cisco Validated Design (CVD) for the FlashStack for Red Hat OpenShift Container Platform solution. This document provides a best practice reference architecture that includes design and deployment guidance. With this guide, one can simplify the digital transformation by eliminating or reducing the associated challenges, including cost and risk reduction, scalability, operational flexibility, and security.

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 containerized workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Integration between OpenShift Container Platform and the storage and data management services occur at several levels, all of which are captured in the design aspects of this document. The main storage integration is based on Container Storage Interface (CSI) Driver for Pure Storage block storage systems, which enables container orchestrators such as Kubernetes to manage the life cycle of persistent storage.

After the CSI driver is installed, it becomes part of the Kubernetes framework, which helps accelerates development and complements DevOps practices.

The OCP platform is installed on top of a VMware cluster with the OCP nodes running as Red Hat Enterprise Linux CoreOS (RHCOS) VMs on ESXi hosts which are in turn deployed on the Cisco UCS servers.

Audience

The intended audience for this document includes, but is not limited to, DevOps managers, IT infrastructure managers, application development leaders, business digital transformation leaders, storage and data management managers, sales engineer and architects working with hybrid and private clouds, and other parties that are looking for a tested, market-proven CI solution that offers flexibility and simplicity in support of their cloud native and application modernization needs along with their digital transformation journey.

Purpose of this Document

This document discusses the design for FlashStack implemented with iSCSI, centered around the fourth-generation Cisco UCS Fabric Interconnect, the fifth-generation Cisco UCS Integrated Infrastructure, the Pure Storage FlashArray//X70 R2, and VMware vSphere 6.7. This design provides a Virtual Server Infrastructure on Cisco UCS B200 M5 Blade Servers and Cisco UCS C220 M5 Rack Servers to provide a platform for OpenShift Container Platform deployments.

What's New in this Release?

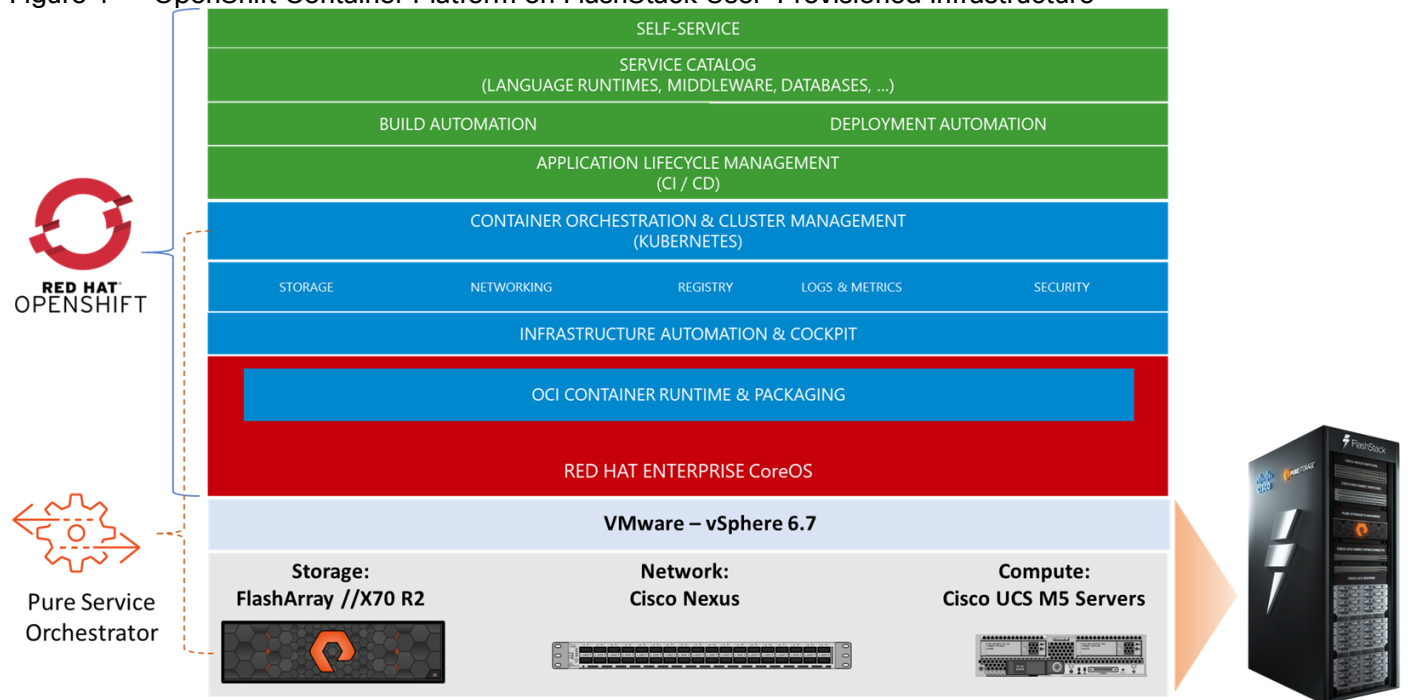
This version of the FlashStack VSI Design introduces the Pure Storage FlashArray//X70 R2 all-flash array along with Cisco UCS B200 M5 Blade Servers and Cisco UCS C220 Rack Servers featuring the Intel Xeon Scalable Family of CPUs as a User Provisioned Infrastructure (UPI) for OpenShift Container Platform (OCP). The design incorporates options for 40Gb iSCSI. Highlights for this design include:

- Pure Service Orchestrator 5.2
- Cisco UCS Manager 4.1
- Cisco Intersight Premier
- RedHat OpenShift Container Platform 4.3

Solution Summary

This solution includes a hardware stack from Cisco and Pure Storage, a hypervisor from VMware, and OCP software platform from Red Hat, and a set of tools for integration and management. These components are integrated so that customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the ground up.

Figure 1 OpenShift Container Platform on FlashStack User-Provisioned Infrastructure



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 container, 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 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 the 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 graphical user interface (GUI), a command-line interface (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 UCS 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 deliver options for both high workload density, as well as high port count, with both supporting either UCS B-Series blade servers, or UCS C-Series rack mount servers. The Fabric Interconnect model featured in this design is the Cisco UCS 6454 Fabric Interconnect; a 54 port 1/10/25/40/100GbE/FCoE switch, supporting 8/16/32Gbps FC ports and up to 3.82Tbps throughput.



[Table 1](#) provides a comparison of the port capabilities of the different Fabric Interconnect models.

Table 1 Cisco UCS 6300 and 6400 Series Fabric Interconnects

Features	6332	6332-16UP	6454	64108
----------	------	-----------	------	-------

Features	6332	6332-16UP	6454	64108
Max 10G ports	96* + 2**	72* + 16	48	96
Max 25G ports	N/A	N/A	48	96
Max 40G ports	32	24	6	12
Max 100G ports	N/A	N/A	6	12
Max unified ports	N/A	16	8	16

* 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, 25Gbps, 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 policies for the servers. FEX technology was developed by Cisco. Up to two FEXs can be deployed in a chassis. For this design we are using the 2408 IOM to provide 25Gbps connective between the blade chassis and the 6400 Series Fabric Interconnect

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/O 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.



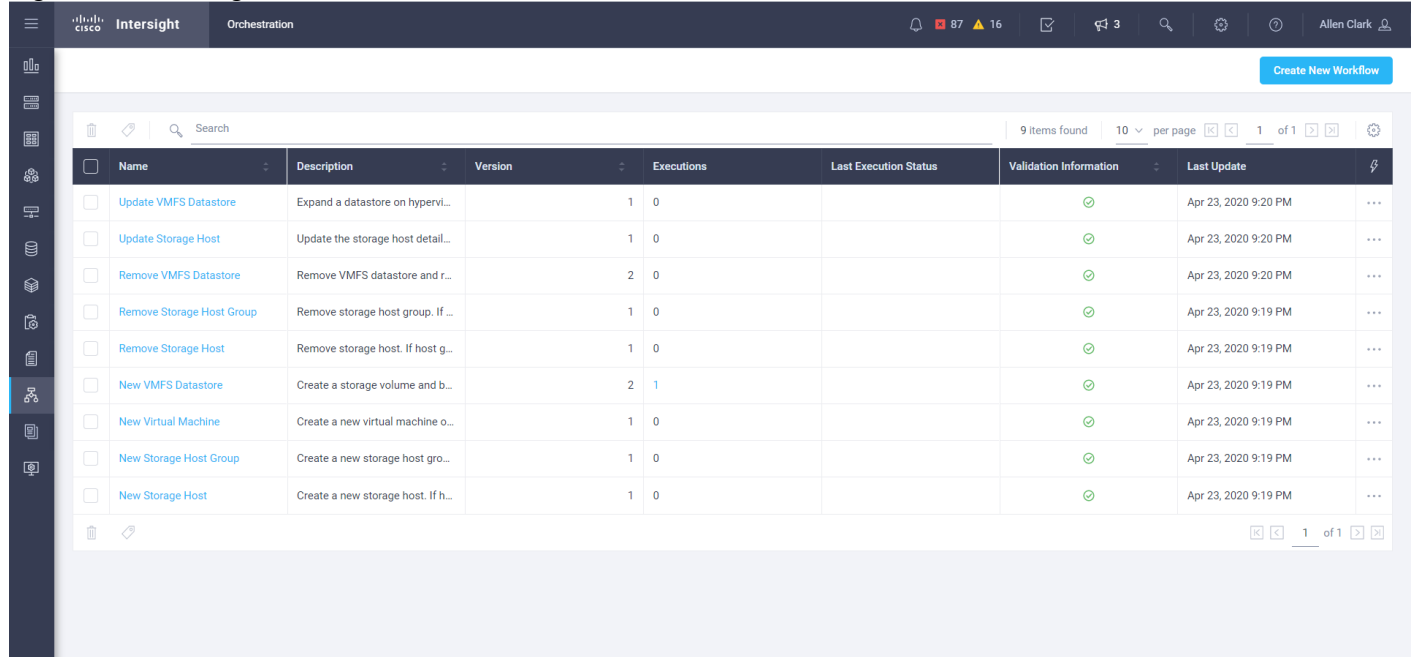
The Cisco Nexus 9336C-FX2 implements Cisco Cloud Scale ASICs, giving flexible, and high port density, intelligent buffering, along with in-built analytics and telemetry.

Cisco Intersight (Optional)

Cisco Intersight™ is a management platform delivered as a service with embedded analytics for your Cisco and 3rd party IT infrastructure. This platform offers an intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in more advanced ways than the prior generations of tools. Cisco Intersight provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with Intersight is quick and easy.

Intersight also provides for the integration, management, and orchestration with the Pure Storage FlashArray//X70 R2. By discovering the FlashArray//X70 R2 and vCenter as part of Intersight, you will be able to automate a variety of task across the converged infrastructure solution. [Figure 2](#) shows the currently available workflow options.

Figure 2 Intersight Orchestration Workflows



Name	Description	Version	Executions	Last Execution Status	Validation Information	Last Update
Update VMFS Datastore	Expand a datastore on hypervi...		1 0		✓	Apr 23, 2020 9:20 PM
Update Storage Host	Update the storage host detail...		1 0		✓	Apr 23, 2020 9:20 PM
Remove VMFS Datastore	Remove VMFS datastore and r...		2 0		✓	Apr 23, 2020 9:20 PM
Remove Storage Host Group	Remove storage host group. If ...		1 0		✓	Apr 23, 2020 9:19 PM
Remove Storage Host	Remove storage host. If host g...		1 0		✓	Apr 23, 2020 9:19 PM
New VMFS Datastore	Create a storage volume and b...		2 1		✓	Apr 23, 2020 9:19 PM
New Virtual Machine	Create a new virtual machine o...		1 0		✓	Apr 23, 2020 9:19 PM
New Storage Host Group	Create a new storage host gro...		1 0		✓	Apr 23, 2020 9:19 PM
New Storage Host	Create a new storage host. If h...		1 0		✓	Apr 23, 2020 9:19 PM

Cisco Workload Optimization Manager (optional)

Instantly scale resources up or down in response to changing demand assuring workload performance. Drive up utilization and workload density. Reduce costs with accurate sizing and forecasting of future capacity.

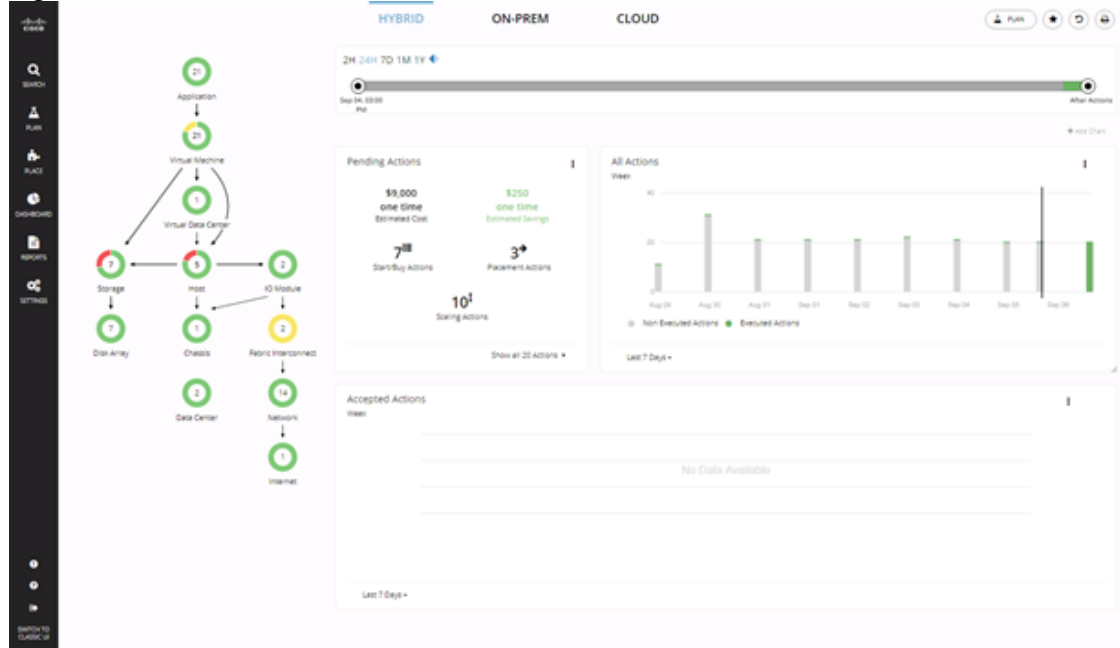
To perform intelligent workload management, Cisco Workload Optimization Manager (CWOM) models your environment as a market of buyers and sellers linked together in a supply chain. This supply chain represents the flow of resources from the datacenter, through the physical tiers of your environment, into the virtual tier and out to the cloud. By managing relationships between these buyers and sellers, CWOM provides closed-loop management of resources, from the datacenter, through to the application.

When you launch CWOM, the Home Page provides the following options:

- Planning
- Placement
- Reports
- Overall Dashboard

The CWOM dashboard provides views specific to On-Prem, the Cloud, or a Hybrid view of infrastructure, applications, and costs.

Figure 3 CWOM Dashboard



For more information about all the capabilities of workload optimization, planning, and reporting, see: <https://www.cisco.com/c/en/us/products/servers-unified-computing/workload-optimization-manager/index.html>

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 percent lower than traditional performance disk arrays.

At the top of the FlashArray line is FlashArray//X – in its 3rd generation is the first mainstream, 100 percent 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 3PB in 6U, 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 an //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 percent 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



FlashArray//X R3 Specifications

	CAPACITY	PHYSICAL
//X10	Up to 73TB / 66.2TiB effective capacity** Up to 22TB / 19.2TiB raw capacity	3U; 640 – 845 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X20	Up to 314TB / 285.4TiB effective capacity** Up to 94TB / 88TiB raw capacity†	3U; 741 – 973 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X50	Up to 663TB / 602.9TiB effective capacity** Up to 185TB / 171TiB raw capacity†	3U; 868 – 1114 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X70	Up to 2286TB / 2078.9TiB effective capacity** Up to 622TB / 544.2TiB raw capacity†	3U; 1084 – 1344 Watts (nominal – peak) 97 lbs (44.0 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X90	Up to 3.3PB / 3003.1TiB effective capacity** Up to 878TB / 768.3TiB raw capacity†	3U – 6U; 1160 – 1446 Watts (nominal – peak) 97 lbs (44 kg) fully loaded; 5.12" x 18.94" x 29.72"
DirectFlash Shelf	Up to 1.9PB effective capacity** Up to 512TB / 448.2TiB raw capacity	3U; 460 - 500 Watts (nominal – peak) 877 lbs (39.8kg) fully loaded; 5.12" x 18.94" x 29.72"

//X Connectivity

ONBOARD PARTS (PER CONTROLLER)	HOST I/O CARDS (3 SLOTS/CONTROLLER)	
<ul style="list-style-type: none"> • 2 x 1/10/25Gb Ethernet • 2 x 1/10/25Gb Ethernet Replication • 2 x 1Gb Management Ports 	<ul style="list-style-type: none"> • 2-port 10GBase-T Ethernet • 2-port 1/10/25Gb Ethernet • 2-port 40Gb Ethernet 	<ul style="list-style-type: none"> • 2-port 25/50Gb NVMe/RoCE • 2-port 16/32Gb Fibre Channel (NVMe-oF Ready) • 4-port 16/32Gb Fibre Channel (NVMe-oF Ready)

[*] Stated //X specifications are applicable to //X R3 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.


[†] Array accepts Pure Storage DirectFlash Shelf and/or Pure.

[1] Gartner, Magic Quadrant for Primary Storage, 17 September 2019. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

Evergreen™ Storage

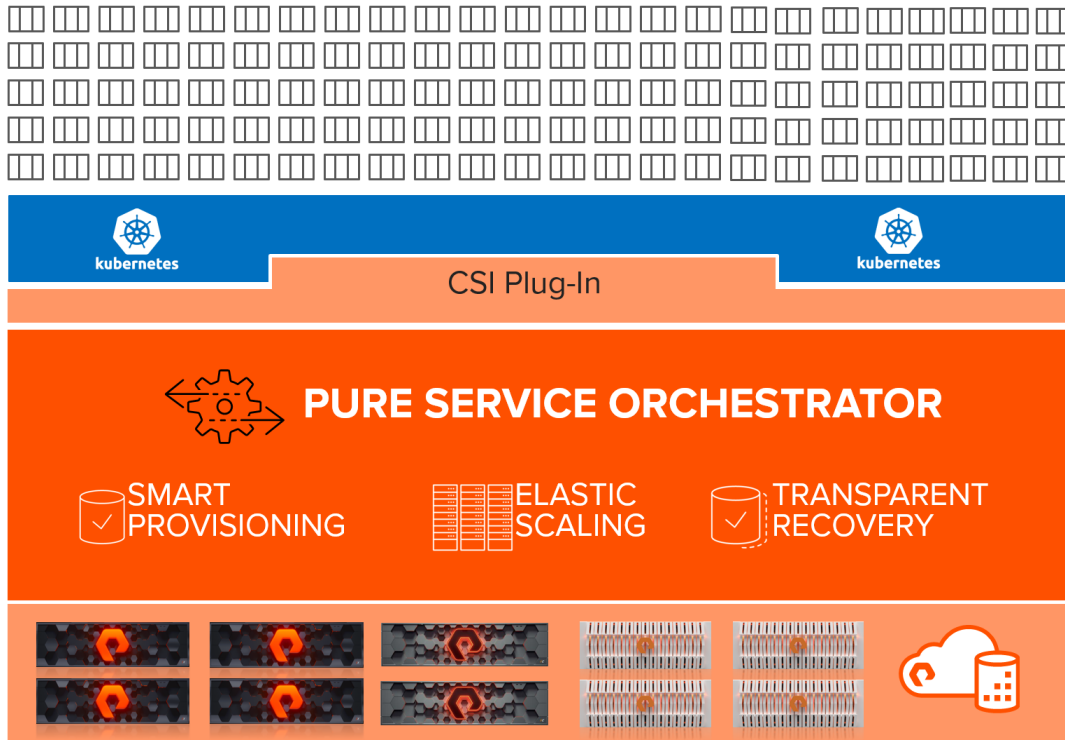
Customers can deploy storage once and enjoy a subscription to continuous innovation via Pure’s Evergreen Storage ownership model: expand and improve performance, capacity, density, and/or features for 10 years or more – all without downtime, performance impact, or data migrations. Pure has disrupted 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.

7 YEARS OF NON-DISRUPTIVE EVERGREEN IMPROVEMENTS



Pure Service Orchestrator

Pure Service Orchestrator is designed to provide a similar experience as the public cloud that your developers expect.

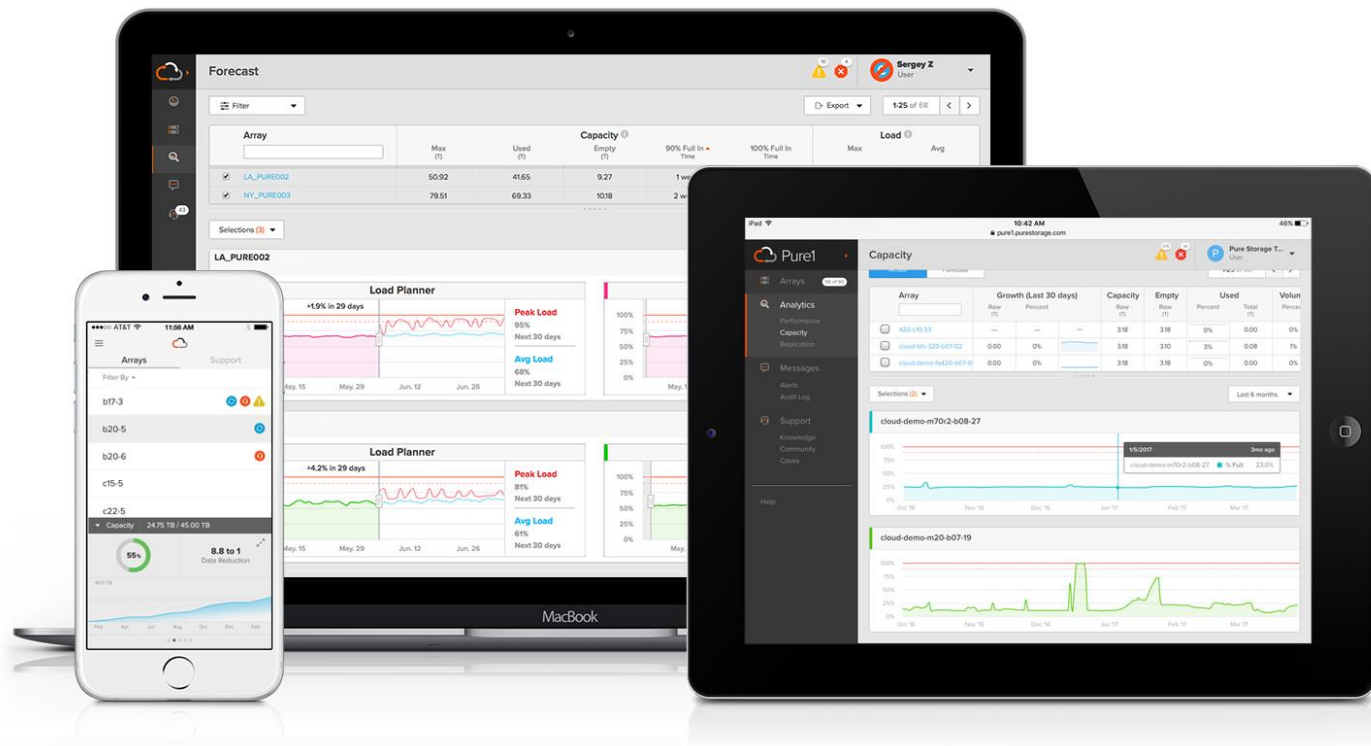


Pure Service Orchestrator integrates seamlessly with the container orchestration frameworks (Kubernetes) and functions as the control plane virtualization layer that enables your containerized environment to move away from consuming storage-as-a-device to consuming storage-as-a-service:

- **Smart Policy Provisioning:** Delivers storage on demand via policy. Pure Service Orchestrator makes the best provisioning decision for each storage request by assessing multiple factors such as performance load, capacity utilization, and health of a given storage system in real-time. Building upon the effortless foundation of Pure Storage, Pure Service Orchestrator delivers an effortless experience and frees the Container Admin from writing complex pod definitions. For example, the default storage requests to Pure Service Orchestrator are as simple as specifying the storage size. No size specified, no problem! Pure Service Orchestrator will service those requests too.
- **Elastic Scaling:** Scales container storage across multiple FlashArray, FlashBlade™ or mix of both systems, supports mix of file and block, and delivers the flexibility to have varied configurations for each storage system that is part of the service. With the astonishing ease of expanding your storage service with a single command, you can start small and scale storage seamlessly and quickly as the needs of your containerized environments grow.
- **Transparent Recovery:** Self-heals to ensure your services stay robust. For example, Pure Service Orchestrator prevents accidental data corruption by ensuring a storage volume is bound to a single persistent volume claim at any given time. If Kubernetes master sends a request to attach the same volume to another Kubernetes node, Pure Service Orchestrator will disconnect the volume from the original Kubernetes node first before attaching it to a new Kubernetes node. This behavior ensures that if a Kubernetes cluster “split-brain” condition occurs (where the Kubernetes master and node become out of sync due to loss of communication), simultaneous I/O to the same storage volume which can corrupt data is prevented.

Pure1®

Pure1, a 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 VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, 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

RedHat OpenShift Container Platform

RedHat OpenShift Architecture

The RedHat OpenShift Container Platform is a container application platform that brings together CRI-O and Kubernetes and provides an API and web interface to manage these services. OpenShift Container Platform allows you to create and manage containers. Containers are standalone processes that run within their own environment,

independent of operating system and the underlying infrastructure. OpenShift helps developing, deploying, and managing container-based applications. It provides a self-service platform to create, modify, and deploy applications on demand, thus enabling faster development and release life cycles. OpenShift Container Platform has a microservices-based architecture of smaller, decoupled units that work together. It runs on top of a Kubernetes cluster, with data about the objects stored in etcd, a reliable clustered key-value store.

Kubernetes Infrastructure

Within OpenShift Container Platform, Kubernetes manages containerized applications across a set of CRI-O runtime hosts and provides mechanisms for deployment, maintenance, and application-scaling. The CRI- service packages, instantiates, and runs containerized applications.

A Kubernetes cluster consists of one or more masters and a set of nodes. This solution design includes HA functionality at the hardware as well as the software stack. Kubernetes cluster is designed to run in HA mode with three (3) master nodes and a minimum of two (2) worker nodes to help ensure that the cluster has no single point of failure

Deployment Node

The Deployment node is a Linux server used to run the OpenShift Container Platform deployment program from. This node can also be used to manage the deployment

Bootstrap Node

The OpenShift Container Platform Bootstrap node is a temporary server used during cluster provisioning to with the initial configuration of the Master nodes to create the permanent cluster control plane. It boots by using an Ignition configuration file that describes how to create the cluster.

OpenShift Master Node

In a Kubernetes cluster, the master nodes run services that are required to control the Kubernetes cluster. In OpenShift Container Platform, the master machines are the control plane. They contain more than just the Kubernetes services for managing the OpenShift Container Platform cluster.

Multiple master nodes are required in a high availability (HA) environment to allow for failover if the leading master host fails. There are three Master nodes in our test environment to provide HA cluster, each is deployed on a separate host (ESXi) for redundancy.

OpenShift Worker Node

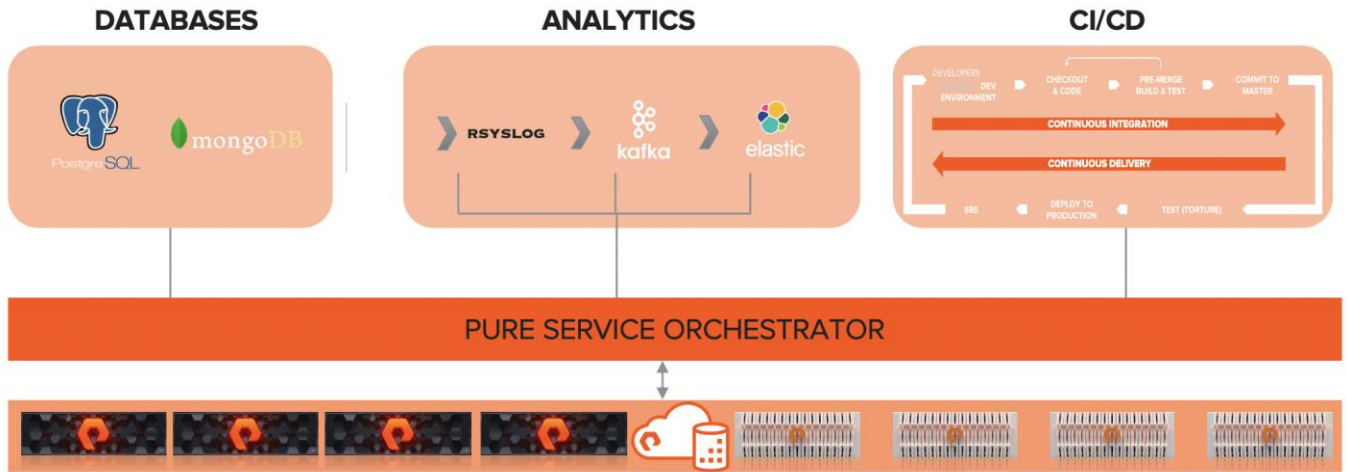
In a Kubernetes cluster, the worker nodes are where the actual workloads requested by Kubernetes users run and are managed. The worker nodes advertise their capacity and the scheduler, which is part of the master services, determines on which nodes to start containers and Pods. Important services run on each worker node, including CRI-O, which is the container engine, Kubelet, which is the service that accepts and fulfills requests for running and stopping container workloads, and a service proxy, which manages communication for pods across workers.

There are two (2) worker nodes in our initial test environment, additional worker nodes were added during the testing procedures to verify the scalability of the solution.

Pure Service Orchestrator

PSO creates an abstracted orchestration layer that runs within your Kubernetes cluster and utilizes native tools in Linux to provide persistent storage for your applications. PSO installs with a single line of CLI and can scale with a simple upgrade command. This provides developers, platform teams, or Site Reliability Engineers (SRE) with complete self-service for persistent storage, with no knowledge of physical infrastructure required. PSO supports

Pure FlashArray™, FlashBlade™, and Pure Cloud Block Store™. This allows you to use the same APIs for private, hybrid, or all public cloud deployments of Kubernetes.



Solution Design

The FlashStack Solutions Design implements a Container 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 datacenter 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, DHCP, and NTP services
- Two Load Balancers for use by OpenShift
- One HTTP server to host Ignition configuration files
- On Linux machine to run the initial installation from


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 external networks and the FlashArray are converged within the fabric interconnect to the Cisco UCS servers through the IOM. These IOM connections from the Cisco UCS Fabric Interconnects to the IOMs are automatically configured as port channels by specifying a Chassis/FEX Discovery Policy within UCSM.

These connections from the Cisco UCS 6454 Fabric Interconnect to the 2408XP IOM hosted within the chassis are shown in [Figure 4](#).

Figure 4 Compute Connectivity



The 2408XP IOM are shown with 4x25Gbps ports to deliver an aggregate of 200Gbps to the chassis, full population of the 2408XP IOM can support 8x25Gbps ports, allowing for an aggregate of 200Gbps per fabric and 400Gbps 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 Cisco UCS servers and the FlashArray//X R2.

Figure 5 Network Connectivity



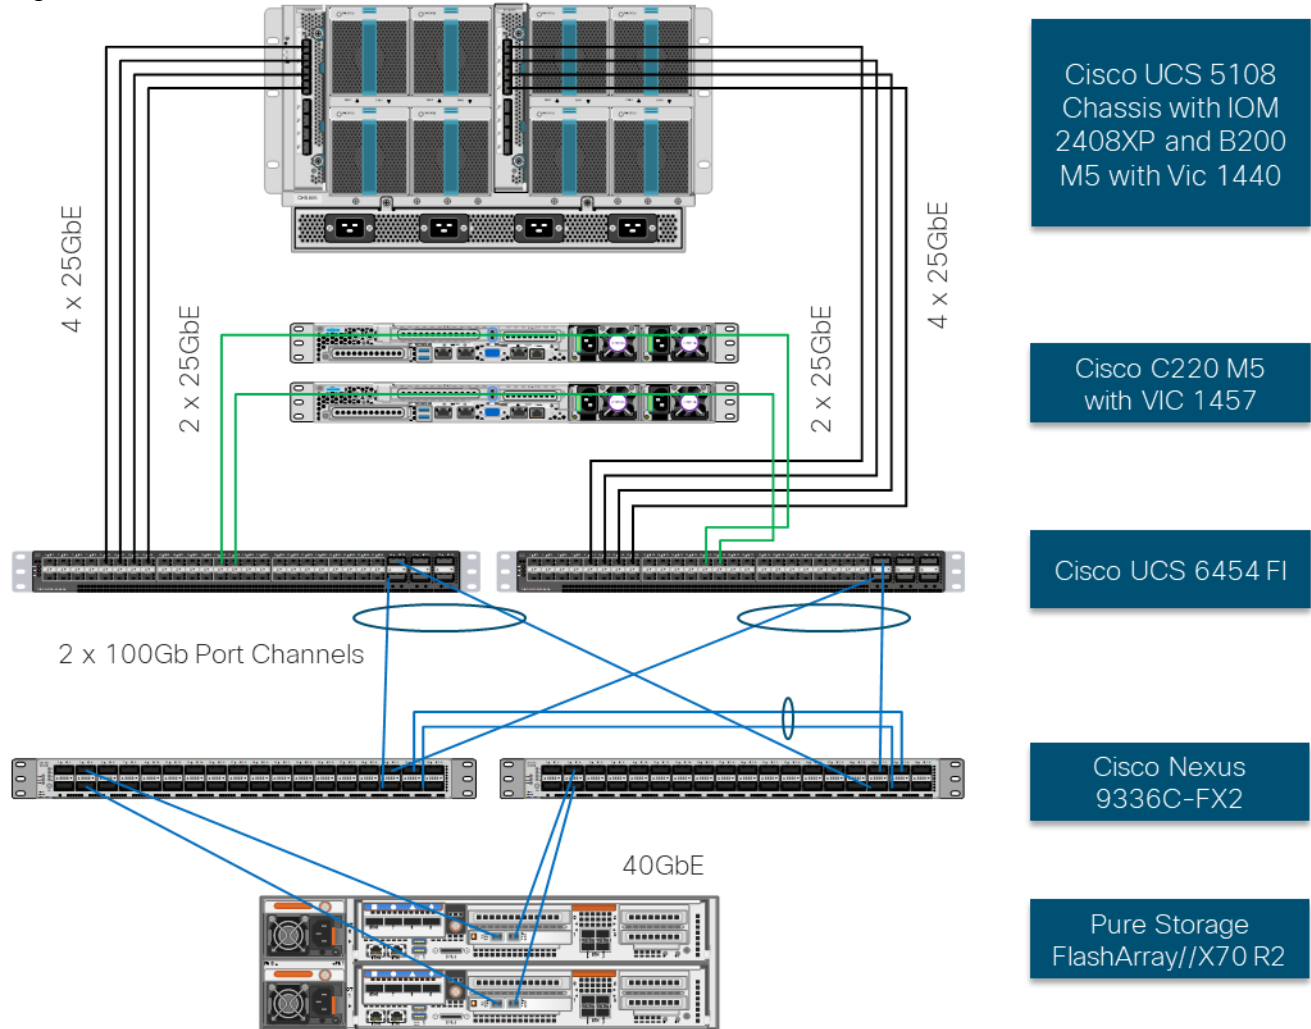
End-to-End Physical Connectivity

iSCSI End-to-End Data Path

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

- Each Cisco UCS server is equipped with a Cisco UCS VIC 1400 Series adapter.
- In the Cisco UCS B200 M5 server this is the Cisco UCS VIC 1440 with 20Gb to IOM A and 20Gbps to IOM B via the Cisco UCS Chassis 5108 chassis backplane.
- In the Cisco UCS C220 M5 server is this the Cisco USC 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 via a 25Gb port to each Cisco UCS 6454 FI.
- Each Cisco UCS 6454 FI connects to the Cisco Nexus 9336C-FX2 via 2x 100Gbps virtual port channels.
- Each controller on the Pure Storage FlashArray//X70 R2 is connected to each Cisco Nexus 9336C-FX2 switch via 2x 40Gbps connections to provide redundant paths.

Figure 6 iSCSI End-to-End Data Path



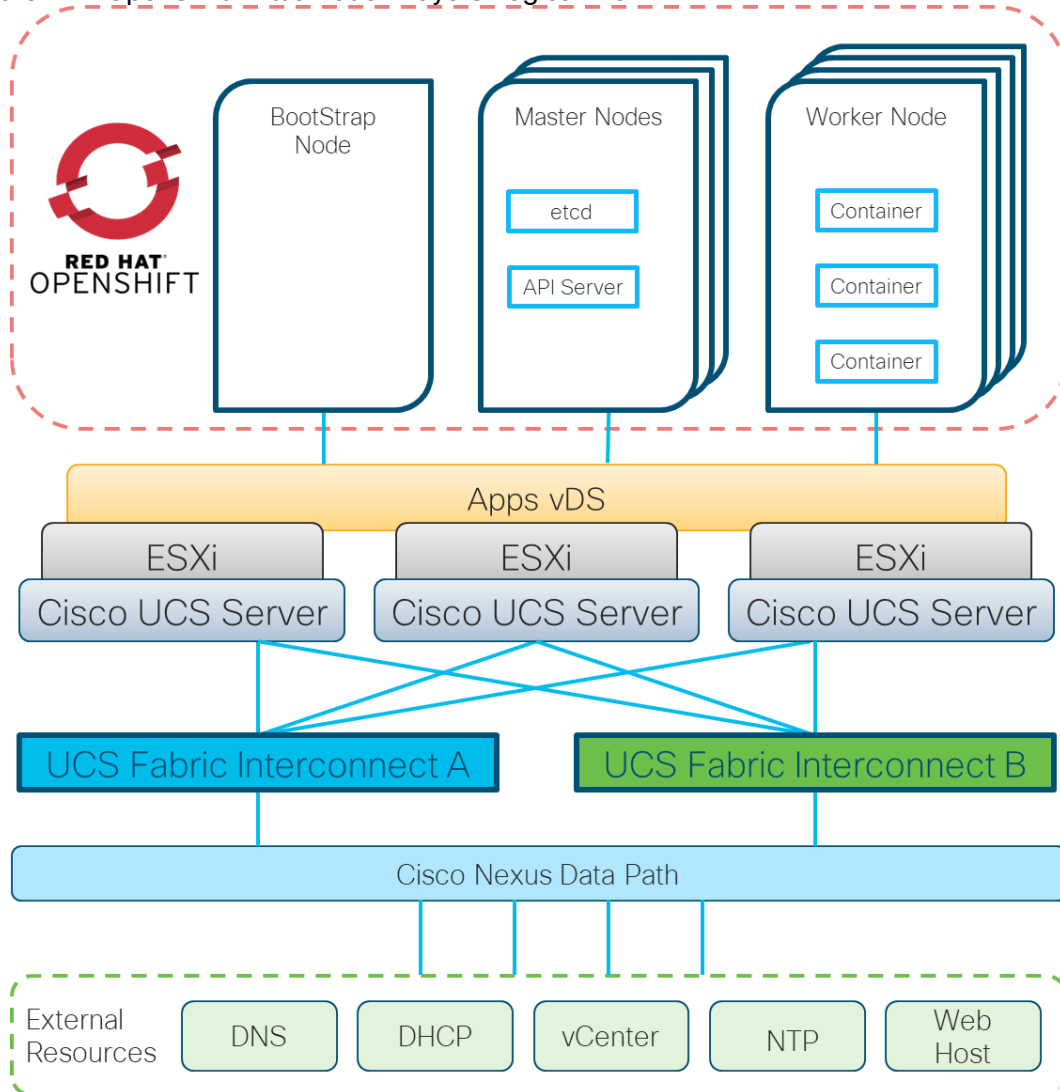
The components of this integrated architecture shown in [Figure 6](#) are:

- Cisco Nexus 9336C-FX2 – 10/25/40/100Gb capable, LAN connectivity to the Cisco UCS compute resources and Pure Storage resource.
- Cisco UCS 6454 Fabric Interconnect – Unified management of Cisco 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 FlashArray//X70 R2

Compute and Virtualization Layer Design

The RedHat OpenShift Container Platform cluster is deployed as virtual machines on top of the VMware cluster. The Containers are connected to the Apps vDS through the first network interface on the Worker nodes to provide access to the services on other Worker Nodes and to outside resources.

Figure 7 OpenShift Virtualization Layers Logical View



The connectivity of details for the ESXi and Cisco UCS layers are detailed in the following section.

Server Models

The following 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 Cisco 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.
- Cisco UCS 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 Cisco 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.
- Cisco UCS 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 Cisco UCS Service Profile Templates that allow rapid creation, as well as guaranteed consistency of the hosts at the Cisco UCS hardware layer.

Cisco UCS vNIC Templates

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

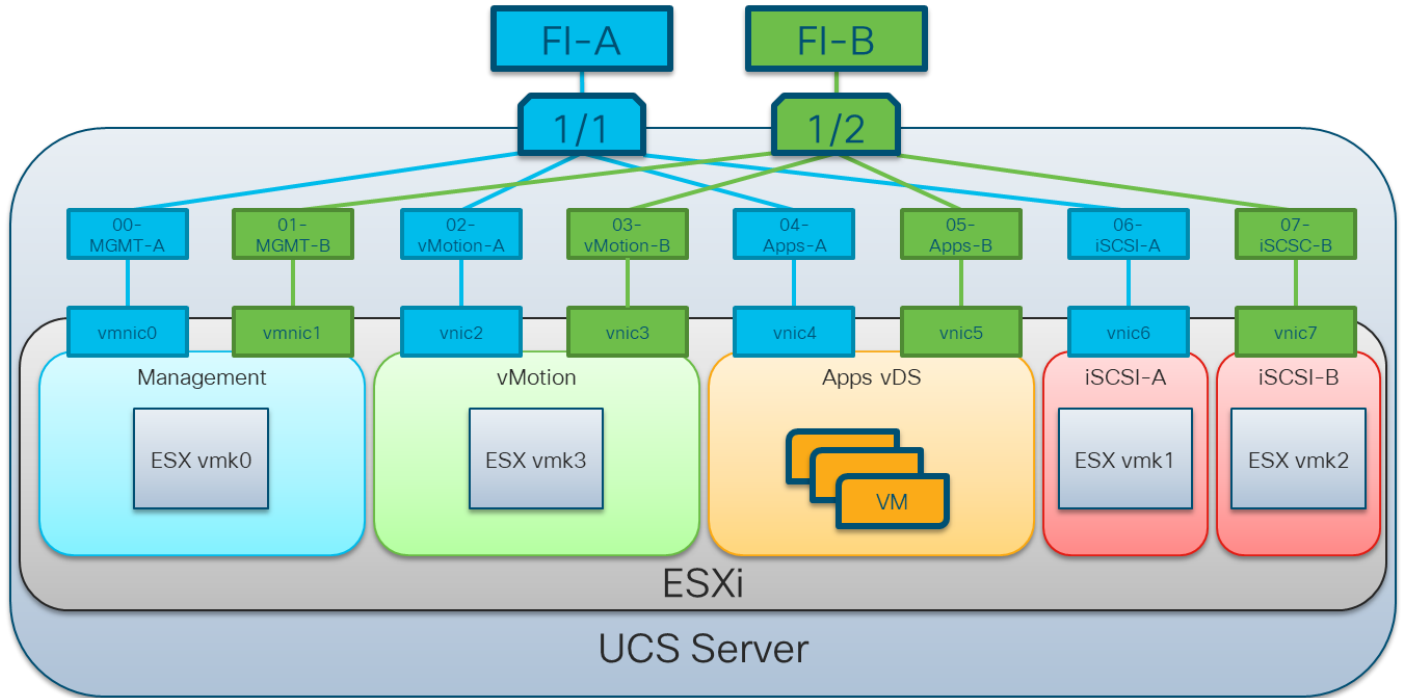
VMware vSphere

This design is implemented using vSphere 6.7 for the vCenter and ESXi version.

Virtual Networking Configuration

The virtual networking configuration utilizes the Cisco UCS VIC adapter to present multiple adapters to ESXi as shown below:

Figure 8 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 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.

Network Design

Cisco Nexus Configuration

The Cisco Nexus configuration covers 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.

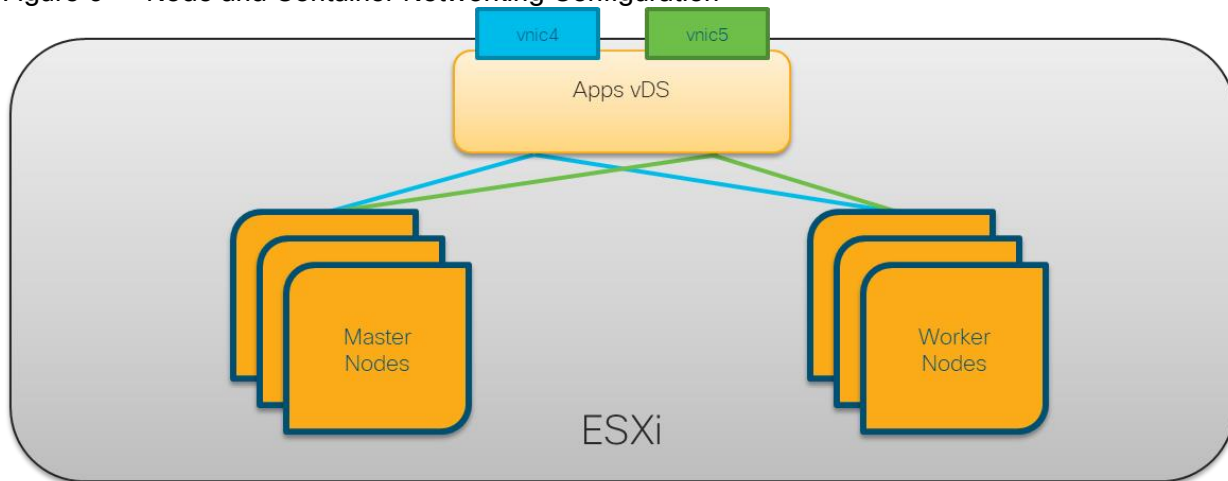
Container Infrastructure Virtualization Layer Design

Container Networking Configuration

OpenShift Management Connectivity

Master nodes and Worker nodes will have management connectivity through the Applications vDS and Application VLAN. This virtual distributed switch will also provide network connectivity for any Containers requiring outside resources. This logical connectivity is shown in [Figure 9](#).

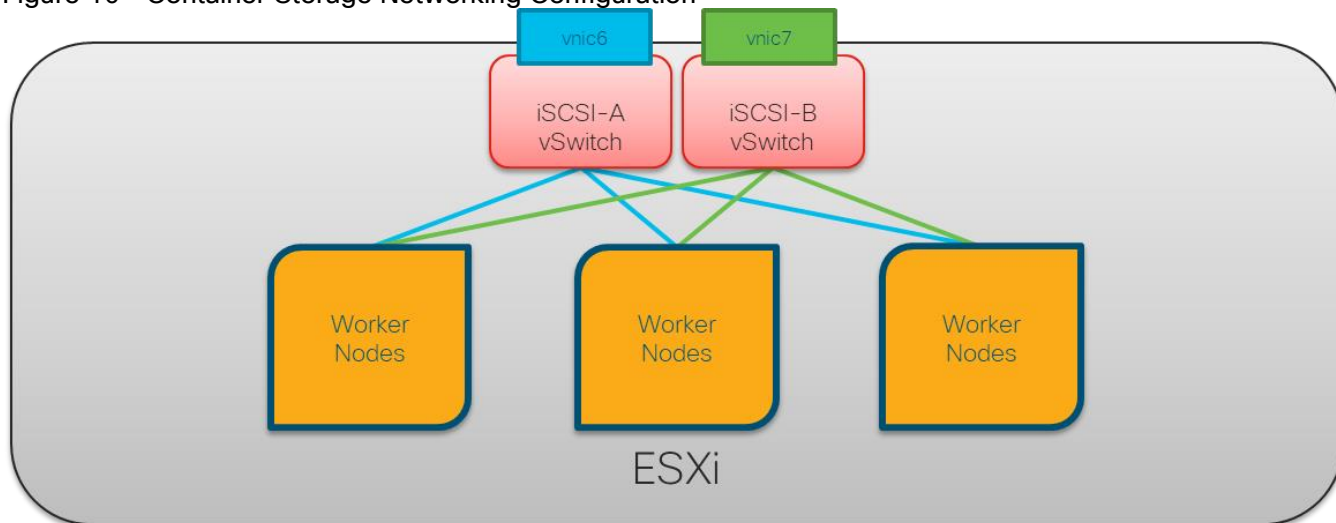
Figure 9 Node and Container Networking Configuration



Container Storage Configuration

This solution leverages the Pure Service Orchestrator (PSO) CSI plugin to provide direct access from the worker nodes to the Pure Storage FlashArray//X70 R2 over the iSCSI vSwitches and associated VLANs. The logical connectivity for this is shown in [Figure 10](#).

Figure 10 Container Storage Networking Configuration



This connectivity is leveraged when Persistent Volumes (PV) and Persistent Volume Claims (PVC) are created within the OpenShift platform. Using Pure Service Orchestrator, the PV, PVC, and backing LUN are all created and mapped to provide persistent storage to a Container or Pod.

Design Considerations

Network 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 Cisco 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 Cisco UCS Server technology with shared storage 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

- 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? With Purity we remove the complexities of aggregates, RAID groups etc. When managing storage, you just 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.

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, it's recommended that you do not add routing to your storage LAN.

Pure Service Orchestrator

PSO can be deployed via an Operator or from the Helm chart:

- PSO has Operator-based install available for the CSI plugin. This install method does not need Helm installation.
- The Pure Flex Operator is supported for OpenShift and Kubernetes (Pure Flex Operator is the preferred installation method for FlexVolume on OpenShift version 3.11). The CSI Operator should be used for OpenShift 4.1 and 4.2.
- Note Use the CSI Helm install method for OpenShift 4.3 and higher with the adoption of Helm3 in OpenShift.

For installation, see the Flex Operator Documentation or the CSI Operator Documentation:

<https://github.com/purestorage/pure-csi-driver/tree/master/operator-csi-plugin>

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 recommended.
- 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.
- UseATSFoRHBOonVMFS5 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, please refer to:

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

Validation

Test Plan

The solution was validated by deploying and running the Weathervane 2.0 tool. The system was validated for resiliency by failing various aspects of the system under load. Examples of the types of tests executed include:

- 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 Cisco Nexus switches for iSCSI
- Failure and recovery of a Cisco Nexus switch
- Removal and recovery of FlashArray//X R2 Controller
- Adding of an OpenShift Worker node
- Restarting OpenShift Worker node
- Restarting OpenShift Master node

Validated Hardware

[Table 2](#) lists 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 [Appendix](#).

Table 2 Validated Hardware

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	7.0(3)I7(8)
Compute	Cisco UCS Fabric Interconnect 6454	4.1(1c)
	Cisco UCS 2408XP IOM	4.1(1c)
	Cisco UCS B200 M5	4.1(1c)
	VMware vSphere	6.7 U3
	ESXi 6.7 U3 nenic	1.0.31.0-1OEM nenic

Component		Software Version/Firmware Version
	VM Virtual Hardware Version	13
	Pure Storage FlashArray //X70 R2	5.3.2
Storage	Pure Storage FlashArray //X70 R2	5.3.2
Container	RedHat OpenShift	4.3

Summary

FlashStack delivers a platform for Enterprise and cloud datacenters using Cisco UCS Blade Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, and 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 2400 Series Fabric Extenders:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-742624.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

Pure Storage VMware Integration:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/vSphere_Web_Client_Plugin/Guide%3A_vSphere_Web_Client_Plugin_-_All_Versions

Pure Service Orchestrator

CSI Operator:

<https://github.com/purestorage/pure-csi-driver/tree/master/operator-csi-plugin>

Helm Chart:

<https://github.com/purestorage/helm-charts/tree/master/pure-k8s-plugin>

<https://github.com/purestorage/helm-charts/tree/master/pure-csi>

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 Centre 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:

- Sreeni Edula, Technical Marketing Engineer, Cisco Systems, Inc.
- Brian Everitt, Technical Marketing Engineer, Cisco Systems, Inc.
- Simon Dodsley, Principle Field Solutions Architect, Pure Storage, Inc.