

CASE STUDY: DATA CENTER REFRESH PROJECT

Jason Gates
Data Center Architect
Presidio
jgates@presidio.com

Randall Borge
Senior Network Architect
Presidio
rborge@presidio.com

Table of Contents

1.0	Abstract	5
2.0	Scope	5
3.0	Solution Overview	6
3.1	Project Overview.....	6
3.2	Overview of Network Devices.....	6
3.3	Network Diagrams	8
3.4	Host Table	9
3.5	Chassis-Based Hardware Platforms	10
3.5.1	Nexus 7000 BC1-CR-1.....	10
3.5.2	Nexus 7000 BC1-CR-2.....	11
3.5.3	Nexus 7000 BC2-CR-1.....	11
3.5.4	Nexus 5596 BC1-DS-1.....	11
3.5.5	Nexus 5596 BC1-DS-2.....	12
3.5.6	Nexus 5596 BC2-DS-1.....	12
3.6	Software Versions	13
3.6.1	Nexus 7000.....	13
3.6.2	Nexus 5000.....	13
3.6.3	Nexus 2000.....	13
3.7	Licensing.....	14
3.7.1	Nexus 7000.....	14
3.7.2	Nexus 5000.....	16
4.0	Physical Layer	17
4.1	Management Interfaces.....	17
4.1.1	Nexus 7000 Supervisor Management Interface (Mgmt 0)	17
4.1.2	Nexus 5000 Switch Management Interface (Mgmt 0).....	18
4.1.3	Nexus Fabric Interconnect Switch Management Interface (Mgmt 0)	18
5.0	Layer 2 Features and Design	18
5.1	Virtual Device Context (VDC) Configuration	18
5.1.1	VDC Information	19
5.2	VLAN Database.....	19
5.2.1	BC1-CR-1	19
5.2.2	BC1-CR-2	20
5.3	VLAN Trunk Protocol (VTP) Configuration	20
5.3.1	Atlanta VTP Configuration	21
5.3.2	Santa Clara VTP Configuration	21
5.4	Nexus 7000 Spanning Tree Configuration	21
5.4.1	BC1-A Spanning Tree Configuration.....	21
5.4.2	BC1-B Spanning Tree Configuration.....	21
5.5	Spanning Tree Extensions	21
5.6	Virtual Port Channel Configuration.....	22
5.6.1	vPC Domain Configuration Details	22

5.6.2	BC1-CR-1	22
5.6.3	BC1-CR-2	23
5.6.4	Port Channel Configuration	23
5.6.5	BC1-CR-1 Port Channel Configuration Table.....	24
5.6.6	BC1-CR-2 Port Channel Configuration Table.....	24
6.0	Layer 3 Configuration	25
6.1	Layer 3 Interface Configurations.....	25
6.1.1	BC1-CR7K1.....	25
6.1.2	BC1-CR7K2.....	25
6.1.3	BC2-CR7K1.....	26
6.2	Routing Configuration.....	26
6.2.1	EIGRP Configuration.....	26
6.2.2	BGP Configuration.....	26
6.2.3	Static Routing Configuration.....	27
7.0	Multicast Routing Configuration.....	27
8.0	Device Management and Security Configuration.....	28
8.1	Nexus 7000 Control Plane Policing (CoPP).....	28
9.0	Quality of Service (QoS) Configuration.....	29
9.1	QoS Configuration.....	30
10.0	Advanced Features	30
10.1	OTV Configuration.....	30
10.1.1	Multihomed Sites and Load Balancing.....	30
10.1.2	OTV BCS1.....	32
10.1.3	OTV BCS-2	32
10.1.4	OTV BCS2-1	33
11.0	Device Management and Security Configuration.....	33
11.1	System Management Configuration	33
11.1.1	Message of the Day (MOTD) Banner	33
11.1.2	Network Time Protocol (NTP) Server.....	33
11.1.3	Timezone Settings.....	34
11.1.4	DNS Settings.....	34
11.1.5	Syslog Configuration	34
11.2	Smart Call Home Configuration	34
11.2.1	Example Call Home Configuration	35
11.3	Simple Network Management Protocol V2 (SNMP) Configuration	35
11.3.1	SNMPv2 Configuration.....	35
11.4	NetFlow Configuration.....	35
12.0	Compute & Storage Architecture	36
12.1	Overview	36
12.2	Data Center Diagrams.....	37
12.2.1	World Pay Architecture Design Diagram	37
12.3	Overview of Platform Features.....	37
12.3.1	Cisco Unified Computing System	37

12.4	Overview of Cisco UCS	37
12.4.1	EMC VNX Series Storage Arrays	39
	Chapter 1.....	41
	Chapter 2.....	41
	Chapter 3.....	41
	Chapter 4.....	41
	Chapter 5.....	41
	Chapter 6.....	41
	Chapter 7.....	41
	Chapter 8.....	41
	Chapter 9.....	41
	Chapter 10.....	41
10.5.1	Virtual Blended Storage Pool (FAST VP).....	41
10.5.2	FAST Cache.....	44
3.3.1	Pool-based Thin Provisioning	46
12.4.2	Storage Configuration	47
12.5	EMC VNX 5400 Configuration	47
12.5.1	Basic Configuration Information	47
12.5.2	Storage Processor Summary	48
12.5.3	Disk Summary	48
12.5.4	Storage Pools / RAID Groups.....	48
12.5.5	Initial Proposed RAID Group Layout	49
13.0	Cisco Unified Computing System (ucs) Production Environment	49
13.1	Cisco UCS Configuration.....	49
13.1.1	Fabric Interconnects	49
13.1.2	Server Configuration	50
13.1.3	Admin Configuration.....	51
14.0	Conclusion	51

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell EMC's views, processes or methodologies.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.

1.0 ABSTRACT

Technology refreshes can be challenging for companies due to overwhelming choices of technology platforms, vendor choice, and complexity. The main drivers for tech refreshes are cost savings, consolidation, replacing aging equipment, maintenance renewal, increasing application HA/performance, and other factors. "World of Art" is a company selling high end art, collectibles, and 24/7 auction events with locations around the globe. World of Art has decided to refresh the technology in their main Atlanta data center and co-location Santa Clara; which includes compute, networking, storage, and backups. The objective of this case study is to detail the design considerations for the project such as advanced concepts, load balancers, iSCSI, virtual port channels, etc. This design document will detail the architecture, implementation strategy, and expected results of following best practices combined with real world examples. It considers steps during the high and low level design phases.

This design document is intended for EMC Proven Professionals, network engineers, SAN engineers, IT managers, and IT directors.

2.0 SCOPE

The EMC Proven Professional Firm has developed a solution that addresses the requirements from World of Art personnel obtained from the RFP documentation and detailed design meetings.

This document contains the detailed design elements associated with the deployed solution:

- Build-out of two new data centers including Nexus 7000, 5500, and 2000 series switches.
- Storage Area Network build-out with VNX[®] 5400.
- Physical network design, which includes physical connections between all of the components.
- Logical design of the network, which takes into account World of Art requirements listed in the RFP.
- VNX array information and detailed feature analysis.
- Cisco Unified Computing System information and detailed feature analysis.

3.0 SOLUTION OVERVIEW

3.1 Project Overview

The Datacenter Design Project provides the following features and benefits:

- Separation of traffic between different networks.
- High Availability and Redundancy.
- Quality of Service for Voice and Business Applications (identified by World of Art).
- Increased throughput through the use of up to 40 Gigabit interfaces.

3.2 Overview of Network Devices

Nexus 7000 Series Switch

The modular Cisco Nexus 7000 Series Switches deliver the highest-density 10 Gigabit Ethernet ports in the market, with up to 768 10-Gbps ports and more than 17 terabits per second (Tbps) of switching capacity, support for 40 and 100 Gigabit Ethernet interfaces, and a comprehensive Cisco NX-OS Software feature set. The Cisco Nexus 7000 Series Switches are designed to meet the requirements of next-generation data centers, Cisco Data Center Interconnect (DCI) solutions, data center aggregation, spine and leaf architecture, and data center end-of-row (EoR) access designs.

Built on a zero-service-loss hardware and software architecture, the Cisco Nexus 7000 Series offers the kind of high availability needed in a next-generation data center, in which virtualization increases the scope of downtime and unified fabric demands Fibre Channel-like availability to properly support storage services. The Cisco Nexus 7000 Series was built with manageability in mind and incorporates a number of unique features, including integrated lights-out management and integrated packet capture and decoding. The Cisco Nexus 7000 also offers innovative switch virtualization capabilities, which, in combination with the switch's density, allows customers to greatly simplify their switching infrastructure, reducing costs, power and cooling load, and management complexity.

Cisco Nexus 5000 Series Switch

Even with its svelte rack switch form factor, the Cisco Nexus 5000 Series of switches offers numerous innovations. The low-latency, low-cost, 10 Gigabit Ethernet switch was first to market with support for Data Center Ethernet, which improves the reliability and scalability of Ethernet for data center purposes. The switch was also the first to deliver Fibre Channel over Ethernet (FCoE), which allows storage traffic to be reliably carried over an Ethernet infrastructure. These features are prime examples of Cisco's holistic approach, with Cisco working with companies such as Emulex, Intel, QLogic, and VMware to bring a complete implementable solution to market. While the Cisco Nexus 5000 Series is designed for most data center environments, its low-latency characteristics also make it an ideal candidate for high-performance computing applications. The Cisco Nexus 5000 Series is available in two variations. The Cisco Nexus 5020 Switch offers 40 ports of lossless 10 Gigabit Ethernet and two uplink slots that can support a combination of 10 Gigabit Ethernet and Fibre Channel ports. The Cisco Nexus 5010 Switch

offers the same features and capabilities as the Cisco Nexus 5020, but in a smaller form factor: 20 fixed, lossless, 10 Gigabit Ethernet ports and one uplink slot.

The Cisco Nexus 5500 platform extends the industry-leading versatility of the Cisco Nexus 5000 Series purpose-built 10 Gigabit Ethernet data center-class switches and provides innovative advances toward higher density, lower latency, and multilayer services. The Cisco Nexus 5500 platform is well suited for enterprise-class data center server access-layer deployments across a diverse set of physical, virtual, storage-access, and high-performance computing (HPC) data center environments.

The Cisco Nexus 5596UP Switch is a 2RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 1920 Gbps of throughput and up to 96 ports. The switch has 48 unified ports and three expansion slots.

Cisco Nexus 2000 Series Fabric Extenders

The Cisco Nexus 2000 Series Fabric Extenders offer a unique approach designed specifically to give customers a means of granularly transitioning from Gigabit Ethernet to 10 Gigabit Ethernet and to a unified fabric. The Cisco Nexus 2000 Series sits on top of a server rack and essentially acts as a remote line card for an upstream switch and becomes an extension of the switch, so software, configuration, and policy are all inherited from the upstream switch; even advanced features such as FCoE and Cisco VN-Link support are inherited. This approach offers two primary benefits to the customer. First, total cost of ownership (TCO) is reduced because of simpler cabling requirements (primarily intrarack) and because there are fewer switches to manage. Second, the Cisco Nexus 2000 Series allows customers to support their existing Gigabit Ethernet attached servers while providing access to advanced features and maintaining a consistent management and operations environment across the data center. The initial model of the fabric extender is the Cisco Nexus 2148T Fabric Extender, which supports 48 Gigabit Ethernet downlinks and 4 10 Gigabit Ethernet uplinks. The combination of the Cisco Nexus 2248PQ and upstream Cisco Nexus 5000 Series Switches and Cisco Nexus 6000 Series Switches provides a cost-effective access layer and scalable strategy for 10 GE and FCoE at the server access layer. It also helps to enable a smooth migration to a 40 GE network fabric.

3.3 Network Diagrams




Figure 1 Primary Atlanta Data Center




Figure 2 Santa Clara and Atlanta Connectivity

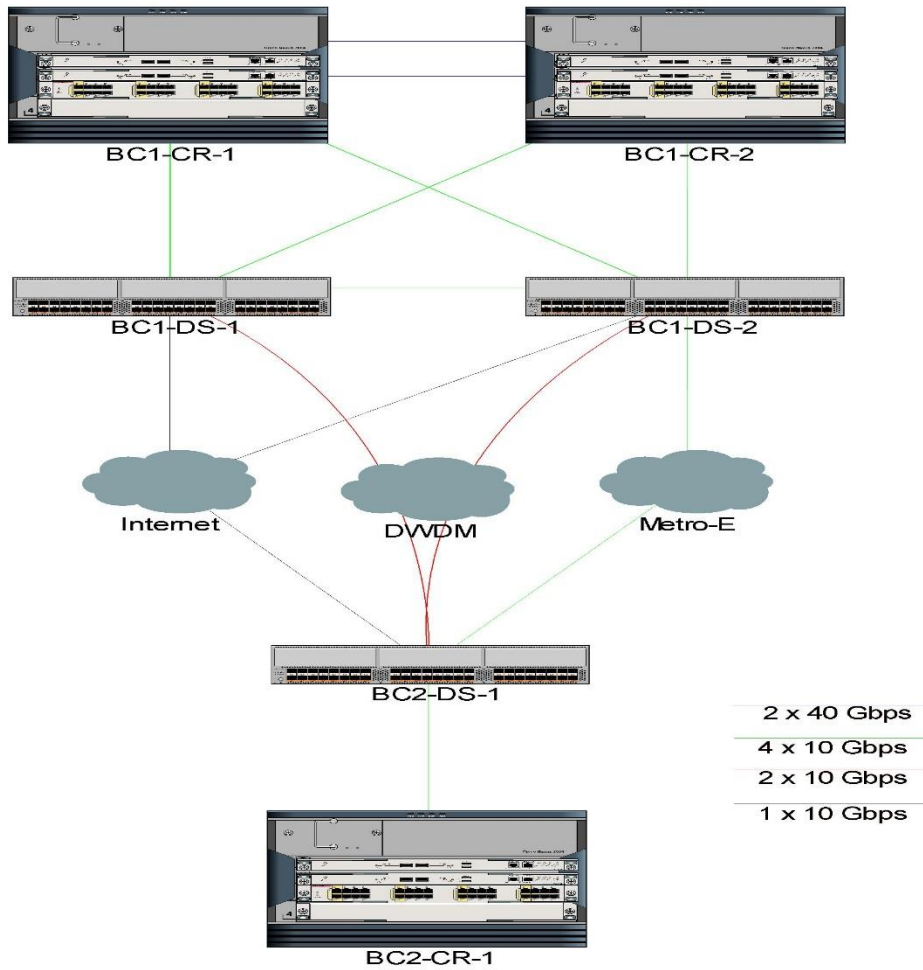


Figure 3 WAN Connection

3.4 Host Table

<i>Hostname</i>	<i>Make/Model</i>	<i>Description</i>	<i>Data Center Location</i>
BC1-CR-1	Cisco Nexus 7004	Core Switch	Atlanta
BC1-CR-2	Cisco Nexus 7004	Core Switch	Atlanta
BC2-CR-1	Cisco Nexus 7004	Core Switch	Santa Clara
BC1-DS-1	Cisco Nexus 5596	Distribution/Access switch	Atlanta
BC1-DS-2	Cisco Nexus 5596	Distribution/Access switch	Atlanta
BC2-DS-1	Cisco Nexus 5596	Distribution/Access switch	Santa Clara

BC1-FI-1	Cisco FI 6248UP	Fabric Interconnect for UCS
BC1-FI-2	Cisco FI 6248UP	Fabric Interconnect for UCS
BC2-FI-1	Cisco FI 6248UP	Fabric Interconnect for UCS
BC2-FI-2	Cisco FI 6248UP	Fabric Interconnect for UCS
BC1-LTM-1	F5 BIG-LTM-4000S	Local Traffic Manager
BC1-LTM-2	F5 BIG-LTM-4000S	Local Traffic Manager
BC2-LTM-1	F5 BIG-LTM-4000S	Local Traffic Manager
BC1-GTM-1	F5 BIG-DNS-2200S	Global Traffic Manager
BC2-GTM-1	F5 BIG-DNS-2200S	Global Traffic Manager
BC1-FW-1	Checkpoint 13500	Next GEN threat prevention
BC1-FW-2	Checkpoint 13500	Next GEN threat prevention
BC2-FW-1	Checkpoint 13500	Next GEN threat prevention

3.5 Chassis-Based Hardware Platforms

The following section details the hardware configuration of the chassis-based LAN hardware being installed as part of this project. This includes which model of linecards are installed in each slot of the switch.

3.5.1 Nexus 7000 BC1-CR-1

<i>Slot</i>	<i>Module</i>	<i>Description</i>
1	N7K-SUP2	Supervisor Module 2
2	N7K-SUP2	Supervisor Module 2
3	N7K-F312FQ-25	40 Gbps Ethernet Module
4		

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>	<i>Serial #</i>
191	Nexus 2248PQF		10	48	FOC1808R1FK

3.5.2 Nexus 7000 BC1-CR-2

<i>Slot</i>	<i>Module</i>	<i>Description</i>
1	N7K-SUP2	Supervisor Module 2
2	N7K-SUP2	Supervisor Module 2
3	N7K-F312FQ-25	40 Gbps Ethernet Module
4		

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>
192	Nexus 2248PQF		10	48

3.5.3 Nexus 7000 BC2-CR-1

<i>Slot</i>	<i>Module</i>	<i>Description</i>
1	N7K-SUP2	Supervisor Module 1X
2	N7K-SUP2	Supervisor Module 1X
3	N7K-F312FQ-25	40 Gbps Ethernet Module
4		

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>
191	Nexus 2248PQF		10	48

3.5.4 Nexus 5596 BC1-DS-1

<i>XBAR Slot</i>	<i>Module</i>	<i>Description</i>
1	Nexus 5596	Fixed Ports (48)

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>
111	Nexus 2248TPE		1G	48
131	Nexus 2248TPE		1G	48
161	Nexus 2248TPE		1G	48
181	Nexus 2248TPE		1G	48
199	Nexus 2248TPE		1G	48

3.5.5 Nexus 5596 BC1-DS-2

<i>XBAR Slot</i>	<i>Module</i>	<i>Description</i>
1	Nexus 5596	Fixed Ports (48)

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>
111	Nexus 2248TPE		1G	48
131	Nexus 2248TPE		1G	48
161	Nexus 2248TPE		1G	48
181	Nexus 2248TPE		1G	48
199	Nexus 2248TPE		1G	48

3.5.6 Nexus 5596 BC2-DS-1

<i>XBAR Slot</i>	<i>Module</i>	<i>Description</i>
1	Nexus 5596	Fixed Ports (48)
2		

<i>FEX #</i>	<i>FEX Model</i>	<i>Location</i>	<i>Speed</i>	<i>Ports</i>
199	Nexus 2248TPE		1G	48

3.6 Software Versions

3.6.1 Nexus 7000

There are 3 different types of software used on the Nexus 7000 series switch these are the kickstart image, system software image, and EPLD image.

n7000-s2-dk9.6.2.8a.bin

n7000-s2-epld.6.2.8.img

n7000-s2-kickstart.6.2.8a.bin

3.6.2 Nexus 5000

There are 2 different types of software used on the Nexus 5000 series switch these are the kickstart image, and system software image.

n5000-uk9-kickstart.7.0.3.N1.1.bin

n5000-uk9.7.0.3.N1.1.bin

3.6.3 Nexus 2000

The Nexus 2000 Fabric Extender runs software that is automatically downloaded from the Nexus 5000 series switch, and is not upgradable outside of this process.

3.7 Licensing

3.7.1 Nexus 7000

In addition to the Base License included with the Nexus 7000, there are multiple feature-based licenses available that enable additional functionality on the device. The following chart has an overview of the feature-based licenses, and the additional functionality they enable.

<i>Feature License</i>	<i>Features</i>
Enterprise Services Package (LAN_ENTERPRISE_SERVICES_PKG)	<ul style="list-style-type: none">• Open Shortest Path First (OSPF) Protocol• Border Gateway Protocol (BGP)• Intermediate System-to-Intermediate System (IS-IS) Protocol (Layer 3 only)• Protocol Independent Multicast (PIM) (sparse mode, bidirectional mode, and source-specific mode (SSM))• Multicast Source Discovery Protocol (MSDP)• Policy-Based Routing• Generic routing encapsulation (GRE) tunnel• Enhanced Interior Gateway Routing Protocol (EIGRP)
Advanced Services Package (LAN_ADVANCED_SERVICES_PKG)	<ul style="list-style-type: none">• Cisco TrustSec• 4 Virtual Device Contexts (VDCs)
VDC Licenses (VDC_PKG)	<ul style="list-style-type: none">• Increments four VDC licenses that allow the Cisco Nexus 7000 Series Supervisor 2 Enhanced module to support eight VDCs
Transport Services Package (LAN_TRANSPORT_SERVICES_PKG)	<ul style="list-style-type: none">• Overlay Transport Virtualization (OTV)• Locator/ID Separation Protocol (LISP)
Scalable Feature Package	<ul style="list-style-type: none">• A single license per system enables all XL-capable

(LAN_SCALABLE_FEATURE_PKG)	<ul style="list-style-type: none"> I/O modules to operate in XL mode.
Enhanced Layer 2 Package (ENHANCED_LAYER2_PKG)	<ul style="list-style-type: none"> FabricPath support on the F Series module
MPLS Services Package (MPLS_PKG)	<ul style="list-style-type: none"> Multiprotocol Label Switching (MPLS)
Storage Enterprise Package (ENTERPRISE_PKG)	<ul style="list-style-type: none"> Inter-VSAN routing (IVR) over Fibre Channel and FCoE IVR Network Address Translation (NAT) over Fibre Channel VSAN-based Access Control Fabric Binding for open systems

World of Art has purchased the following licenses:

Enterprise Services Package:	Yes
Advanced Services Package:	Yes
VDC Licenses:	Yes
Transport Services Package:	Yes
Scalable Feature Package:	No
Enhanced Layer 2 Package:	Yes
MPLS Services Package:	No
Storage Enterprise Package:	No

Licenses are first copied to the bootflash, then installed using the command:

install license bootflash:license_file.lic

3.7.2 Nexus 5000

In addition to the Base License included with the Nexus 5000, there are multiple feature-based licenses available that enable additional functionality on the Nexus.

The following chart has an overview of the feature-based licenses, and the additional functionality they enable.

Feature License	Features
<i>FabricPath Services Package</i> (<i>ENHANCED_LAYER2_PKG</i>)	<ul style="list-style-type: none">• FabricPath
<i>Layer 3 Base Services Package</i> (<i>LAN_BASE_SERVICES_PKG</i>)	<ul style="list-style-type: none">• Static routing• RIPv2• OSPFv2• EIGRP stub• HSRP• VRRP• IGMP v2/v3• routed ACL• uRPF <p>NOTE: OSPF scalability is limited to 256 dynamically learned routes</p>
<i>Layer 3 Enterprise Services Package</i> (<i>LAN_ENTERPRISE_SERVICES_PKG</i>)	<ul style="list-style-type: none">• Full EIGRP• OSPF with scalability up to 8000 routes• BGP and VRF-lite (IP-VPN)• maximum routes supported by L3 hardware 8000 entries
<i>Storage Protocols Services Package</i>	<ul style="list-style-type: none">• Native Fibre Channel

(FC_FEATURES_PKG)

- FCoE
- NPV
- FC Port Security
- Fabric Binding

FCoE NPV Package

(FCOE_NPV_PKG)

- FCoE NPV

World of Art has purchased the following licenses:

FabricPath Services Package	Yes
Layer 3 Base Services Package:	No
Layer 3 Enterprise Services Package:	No
Storage Protocol Services Package:	Yes
FCoE NPV Package:	No

Licenses are first copied to the bootflash, then installed using the command:

install license bootflash:license_file.lic

4.0 PHYSICAL LAYER

4.1 Management Interfaces

4.1.1 *Nexus 7000 Supervisor Management Interface (Mgmt 0)*

Nexus 7000 Management Interfaces

VDC	Network Address	Netmask	Default Gateway
BC1-CR-1	10.200.1.10	255.255.255.0	10.200.1.1
BC1-CR-2	10.200.1.11	255.255.255.0	10.200.1.1
BC2-CR-1	10.201.1.10	255.255.255.0	10.201.1.1

4.1.2 Nexus 5000 Switch Management Interface (Mgmt 0)

Nexus 5000 Management Interface Configuration Information

<i>Device Hostname</i>	<i>Network Address</i>	<i>Netmask</i>	<i>Default Gateway</i>
BC1-DS-1	10.200.1.12	255.255.255.0	10.200.1.1
BC1-DS-2	10.200.1.13	255.255.255.0	10.200.1.1
BC2-DS-1	10.201.1.12	255.255.255.0	10.201.1.1

4.1.3 Nexus Fabric Interconnect Switch Management Interface (Mgmt 0)

Nexus Fabric Interconnect Management Interface Configuration Information

<i>Device Hostname</i>	<i>Network Address</i>	<i>Netmask</i>	<i>Default Gateway</i>
BC1-FI	10.200.1.60	255.255.255.0	10.200.1.1
BC1-FI-1	10.200.1.61	255.255.255.0	10.200.1.1
BC1-FI-2	10.200.1.62	255.255.255.0	10.200.1.1
BC2-FI	10.201.1.60	255.255.255.0	10.201.1.1
BC2-FI-1	10.201.1.61	255.255.255.0	10.201.1.1
BC2-FI-2	10.201.1.62	255.255.255.0	10.201.1.1

5.0 LAYER 2 FEATURES AND DESIGN

5.1 Virtual Device Context (VDC) Configuration

Cisco NX-OS software supports VDCs on the Nexus 7000 series switches, which partition a single physical device into multiple logical devices that provide fault isolation, management isolation, address allocation isolation, service differentiation domains, and adaptive resource management. You can manage a VDC instance within a physical device independently. Each VDC appears as a unique device to the connected users. A VDC runs as a separate logical entity within the physical device, maintains its own unique set of running software processes, has its own configuration, and can be managed by a separate administrator. Each physical Ethernet interface can belong to only one VDC at any given time.

NOTE: By default all Interfaces are in the default VDC. Interfaces not listed in the table, below, remain in the default VDC.

5.1.1 VDC Information

The following VDCs will be configured, with interfaces allocated as indicated:

<i>Switch Hostname</i>	<i>VDC Name</i>	<i>Interfaces Allocated to VDC</i>	<i>Interface Type</i>
BC1-CR-1	BC1-CR-1	-	-
BC1-CR-1	BC1-CR-1-VDC-1	Ethernet 3/1-12	40 Gig

<i>Switch Hostname</i>	<i>VDC Name</i>	<i>Interfaces Allocated to VDC</i>	<i>Interface Type</i>
BC1-CR-2	BC1-CR-2	-	-
BC1-CR-2	BC1-CR-2-VDC-1	Ethernet 3/1-12	40 Gig

<i>Switch Hostname</i>	<i>VDC Name</i>	<i>Interfaces Allocated to VDC</i>	<i>Interface Type</i>
BC2-CR-1	BC2-CR-1	-	-
BC2-CR-1	BC2-CR-1-VDC-1	Ethernet 3/1-12	40 Gig

5.2 VLAN Database

As part of the data center implementation, the following VLANs will be created to support connectivity in the switches/VDCs listed.

5.2.1 BC1-CR-1

<i>VLAN #</i>	<i>VLAN NAME</i>	<i>Mode</i>
2	Oracle RAC	
6	iSCSI	
20	Production	
48	Safari	
300	Management	

310	Fail-over Sync
400	MetroE
410	DWDM
900	Internet-Public
910	Internet-DMz
920	Internet-Inside-1
921	Internet-Inside-2

5.2.2 BC1-CR-2

<i>VLAN #</i>	<i>VLAN NAME</i>	<i>Mode</i>
2	iSCSI	
6	vMotion	
20	Production	
48	Safari	
300	Management	
400	MetroE	
410	DWDM	
900	Internet-Public	
910	Internet-DMz	
920	Internet-Inside-1	
921	Internet-Inside-2	

5.3 VLAN Trunk Protocol (VTP) Configuration

VLAN Trunk Protocol (VTP) reduces administration in a switched network. When you configure a new VLAN on one VTP server, the VLAN is distributed through all switches in the domain. This reduces the

need to configure the same VLAN everywhere. VTP is a Cisco-proprietary protocol that is available on most of the Cisco Catalyst series products.

VTP has four different modes that it can operate in:

- Server – Makes VLAN changes and pushes information to client devices.
- Client – Receives VTP updates from a server. Cannot make VLAN changes.
- Transparent – Does not participate in VTP, but will pass VTP through. Has a standalone VLAN database.
- Off – Does not run VTP at all. VTP messages will not pass through. This option is only available on some switches.

5.3.1 Atlanta VTP Configuration

```
ntp domain DTC  
ntp mode server
```

5.3.2 Santa Clara VTP Configuration

```
ntp domain CDR  
ntp mode server
```

World of Art has opted to run their switches in server mode.

5.4 Nexus 7000 Spanning Tree Configuration

For increased spanning tree convergence speed, the Rapid per VLAN Spanning Tree will be used in all VDCs on the Nexus 7000s, and on the Nexus 5000 switches. To support Virtual Port Channel (VPC) functionality, the Nexus 7000 Switch A will be the spanning tree root for all VLANs shown above.

5.4.1 BC1-A Spanning Tree Configuration

```
spanning-tree vlan 1-999 priority 4096 (LAN)
```

5.4.2 BC1-B Spanning Tree Configuration

```
spanning-tree vlan 1-999 priority 4096 (LAN)
```

5.5 Spanning Tree Extensions

For vPC Loop-Free Topologies, the following spanning-tree configurations are recommended:

vPC Peer Link Ports:

- spanning-tree port type network (automatic)

Port Channel Interfaces connected to Uplinked Switches:

- spanning-tree port type normal
- spanning-tree guard root

Interfaces connected to Host Devices

- spanning-tree port type edge
- spanning-tree bpdguard enabled

5.6 Virtual Port Channel Configuration

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus switches to appear as a single port channel by a third device. You can use only Layer 2 port channels in the vPC. A vPC domain is associated to a single VDC, so all vPC interfaces belonging to a given vPC domain must be defined in the same VDC. You must have a separate vPC peer-link and peer-keepalive link infrastructure for each VDC deployed. Consolidating a vPC pair (two vPC peer devices of the same domain) in two VDCs of the same physical device is not supported. The vPC peer link must use 10-Gigabit Ethernet ports for both ends of the link or the link will not form.

If the vPC peer link is a Classic Ethernet trunk port, then the system will be running standard vPC. If the peer link is running Fabric Path, then the system will be running vPC+. In the case of vPC+ a virtual switch ID needs to be defined for each vPC+ pair.

5.6.1 vPC Domain Configuration Details

<i>Device Name</i>	<i>VPC Domain ID</i>	<i>Keepalive Dst IP</i>	<i>Keepalive VRF</i>
BC1-CR-1	999	10.200.1.14	Management
BC1-CR-2	999	10.200.1.13	Management

5.6.2 BC1-CR-1

```
vpc domain 999
role priority 1000
peer-keepalive destination 10.200.1.13 source 10.200.1.14
peer-gateway
auto-recover
```

5.6.3 BC1-CR-2

```
vpc domain 999
role priority 100
peer-keepalive destination 10.200.1.14 source 10.1.0.13
peer-gateway
auto-recover
```

5.6.4 Port Channel Configuration

LACP allows for up to 16 interfaces to be configured in a port-channel. A maximum of 8 interfaces can be active, with up to 8 interfaces in standby mode.

The device load balances traffic across all operational interfaces in a port channel by hashing the addresses in the frame to a numerical value that selects one of the links in the channel. Port channels provide load balancing by default. Port-channel load-balancing uses MAC addresses, IP addresses, or Layer 4 port numbers to select the link. Port-channel load balancing uses either source or destination addresses or ports, or both source and destination addresses or ports.

You can configure the load-balancing mode to apply to all port channels that are configured on the entire device or on specified modules. The per-module configuration takes precedence over the load-balancing configuration for the entire device. You can configure one load-balancing mode for the entire device, a different mode for specified modules, and another mode for the other specified modules. You cannot configure the load-balancing method per port channel

The Nexus supports the following load-balancing methods:

- Destination MAC address
- Source MAC address
- Source and destination MAC address
- Destination IP address
- Source IP address
- Source and destination IP address
- Source TCP/UDP port number
- Destination TCP/UDP port number
- Source and destination TCP/UDP port number

For a Layer 2 frame, the default load-balancing method is source and destination MAC address. For a Layer 3 frame, the default load-balancing method is source and destination MAC address and the source and destination IP address. For a Layer 4 frame, the default load-balancing method is the source and destination MAC address, the source and destination IP address, and the source and destination port number.

Port Channel Load Balancing Mode Configuration: Default Methods

5.6.5 BC1-CR-1 Port Channel Configuration Table

<i>PC #</i>	<i>Member Ports</i>	<i>Device Connected</i>	<i>vPC Number</i>	<i>Port Type</i>
101	1/1-4	2248	NA	FEX
201	1/5-8	UCS-A	201	Edge
202	1/9-10	UCS-A-ISCSI	202	Edge
301	1/11-14	UCS-B	301	Edge
302	1/15-16	UCS-B-ISCSI	302	Edge
401	1/17-18	VNX-A	401	Edge
402	1/19-20	VNX-B	402	Edge
999	1/47-48	Peer-Keepalive 7K	999	Normal

5.6.6 BC1-CR-2 Port Channel Configuration Table

<i>PC #</i>	<i>Member Ports</i>	<i>Device Connected</i>	<i>vPC Number</i>	<i>Port Type</i>
102	1/1-4	2248	NA	FEX
201	1/5-8	UCS-A	201	Edge
202	1/9-10	UCS-A-ISCSI	202	Edge
301	1/11-14	UCS-B	301	Edge
302	1/15-16	UCS-B-ISCSI	302	Edge
401	1/17-18	VNX-A	401	Edge
402	1/19-20	VNX-B	402	Edge
999	1/47-48	Peer-Keepalive 7K	999	Normal

6.0 LAYER 3 CONFIGURATION

6.1 Layer 3 Interface Configurations

6.1.1 BC1-CR7K1

<i>Interface</i>	<i>Description</i>	<i>IP Address</i>	<i>Netmask</i>
VLAN2	iSCSI	10.200.2.254	/24
VLAN6	vMotion	10.200.6.254	/24
VLAN20	Production	10.200.20.254	/24
VLAN48	Safari	10.200.48.254	/24
VLAN300	Management	10.200.1.254	/24
VLAN400	MetroE	10.150.1.250	/24
VLAN410	DWDM	10.200.0.241	/29
VLAN921	Internet-Inside	10.200.0.13	/28
Lo0	Loopback	10.200.0.254	/32

6.1.2 BC1-CR7K2

<i>Interface</i>	<i>Description</i>	<i>IP Address</i>	<i>Netmask</i>
VLAN2	iSCSI	10.200.2.253	/24
VLAN6	vMotion	10.200.6.253	/24
VLAN20	Production	10.200.20.253	/24
VLAN48	Safari	10.200.48.253	/24
VLAN300	Management	10.200.1.253	/24
VLAN400	MetroE	10.150.1.251	/24
VLAN410	DWDM	10.200.0.242	/29
VLAN921	Internet-Inside	10.200.0.12	/28
Lo0	Loopback	10.200.0.253	/32

6.1.3 BC2-CR7K1

<i>Interface</i>	<i>Description</i>	<i>IP Address</i>	<i>Netmask</i>
VLAN2	iSCSI	10.201.2.1	/24
VLAN6	vMotion	10.201.6.1	/24
VLAN20	Production	10.201.20.252	/24
VLAN48	Safari	10.201.48.1	/24
VLAN300	Management	10.201.1.1	/24
VLAN400	MetroE	10.150.1.252	/24
VLAN410	DWDM	10.200.0.243	/29
VLAN921	Internet-Inside	10.201.0.14	/28
Lo0	Loopback	10.201.0.254	/32

6.2 Routing Configuration

6.2.1 EIGRP Configuration

World of Art will be using EIGRP as their enterprise wide routing protocol. EIGRP will be configured in the following Virtual Device Contexts: LAN and WAN

EIGRP configuration template:

```
router eigrp 1
  redistribute direct route-map <NAME>
  redistribute static route-map <NAME>

interface Vlan####
  ip router eigrp 1

(Examples of prefix-list and route map)
ip prefix-list <NAME> seq 5 permit 0.0.0.0/0 le 32

route-map <NAME> permit 10
  match ip address prefix-list <NAME>
```

6.2.2 BGP Configuration

```
Router bgp 650XX
neighbor 192.168.98.X remote-as 65000
neighbor 192.168.98.X soft-reconfiguration inbound
network 10.X.X.0 mask 255.255.X.X
network 10.X.X.0 mask 255.255.X.X
network 10.X.X.0 mask 255.255.X.X
```

6.2.3 Static Routing Configuration

Static Routing will be configured between the Nexus 7000 and the Firewall for connectivity to the DMZ servers.

Static Routing will be configured as follows:

```
ip route 10.200.1.0 255.255.255.0 10.200.1.21 (Firewall IP address)
```

7.0 MULTICAST ROUTING CONFIGURATION

IP multicast is a method of forwarding the same set of IP packets to a number of hosts within a network. You can use multicast in both IPv4 and IPv6 networks to provide efficient delivery of data to multiple destinations. Multicast involves both a method of delivery and discovery of senders and receivers of multicast data, which is transmitted on IP multicast addresses called groups. The routers in the network listen for receivers to advertise their interest in receiving multicast data from selected groups. The routers then replicate and forward the data from sources to the interested receivers. Multicast data for a group is transmitted only to those LAN segments with receivers that requested it.

The Internet Group Management Protocol (IGMP) snooping software examines Layer 2 IP multicast traffic within a VLAN to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help the routers forward IGMP membership reports. The IGMP snooping software responds to topology change notifications. By default, IGMP snooping is enabled on the device. Cisco NX-OS supports multicasting with Protocol Independent Multicast (PIM) sparse mode. PIM is IP routing protocol-independent and can leverage whichever unicast routing protocols are used to populate the unicast routing table. In PIM sparse mode, multicast traffic is sent only to locations of the network that specifically request it. PIM dense mode is not supported by Cisco NX-OS.

All the Layer 3 interfaces and SVI for the core devices (BC1-CR7K1/BC1-CR7K2/BC2-CR7K) are going to be configured to support PIM/SSM and IGMPv3 for the group range 232.0.0.0/5.

8.0 DEVICE MANAGEMENT AND SECURITY CONFIGURATION

8.1 Nexus 7000 Control Plane Policing (CoPP)

The NX-OS device provides control plane policing to prevent denial-of-service (DoS) attacks from impacting performance.

The supervisor module divides the traffic that it manages into three functional components or planes:

- Data plane—Handles all the data traffic. The basic functionality of a NX-OS device is to forward packets from one interface to another. The packets that are not meant for the switch itself are called the transit packets. These packets are handled by the data plane.
- Control plane—Handles all routing protocol control traffic. These protocols, such as the Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) Protocol, and Protocol Independent Multicast (PIM) Protocol, send control packets between devices. These packets are destined to router addresses and are called control plane packets.
- Management plane—Runs the components meant for NX-OS device management purposes such as the command-line interface (CLI) and Simple Network Management Protocol (SNMP).

The supervisor module has both the management plane and control plane and is critical to the operation of the network. Any disruption or attacks to the supervisor module will result in serious network outages. For example, excessive traffic to the supervisor module could overload and slow down the performance of the entire NX-OS device. Attacks on the supervisor module can be of various types such as DoS that generates IP traffic streams to the control plane at a very high rate. These attacks force the control plane to spend a large amount of time in handling these packets and prevents the control plane from processing genuine traffic.

Default Policing Policies

The NX-OS software installs the default copp-system-policy policy to protect the supervisor module from DoS attacks. Set the level of protection by choosing one of the following CoPP policy options:

- **Strict**—This policy is 1 rate and 2 color and has BC value of 250 ms, except for the important class, which has a value of 1000 ms.
- **Moderate**—This policy is 1 rate and 2 color and has a BC value of 310 ms, except for the important class, which has a value of 1250 ms. These values are 25 percent greater than the Strict policy.
- **Lenient**—This policy is 1 rate and 2 color and has a BC value of 375 ms, except for the important class, which has a value of 1500 ms. These values are 50 percent greater than the Strict policy.

Control Plane Policing is configured in the default VDC, only, but applies to all other VDCs.

<i>Feature Name</i>	<i>Status</i>
Control Plane Policing	Enabled
Policy Option Selected	Lenient

Policing option Lenient was selected. Strict and Moderate options are also available should World of Art personnel wish to change to these in the future.

9.0 QUALITY OF SERVICE (QoS) CONFIGURATION

QoS features are used to provide the most desirable flow of traffic through a network. QoS allows you to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance. The control of traffic is based on the fields in the packets that flow through the system.

You use classification to partition traffic into classes. You classify the traffic based on the port characteristics (CoS field) or the packet header fields that include IP precedence, Differentiated Services Code Point (DSCP), Layer 2 to Layer 4 parameters, and the packet length. The values used to classify traffic are called match criteria. When you define a traffic class, you can specify multiple match criteria, you can choose to not match on a particular criterion, or you can determine traffic class by matching any or all criteria. Traffic that fails to match any class is assigned to a default class of traffic called class-default.

Marking is the setting of QoS information that is related to a packet. You can set the value of standard QoS fields IP precedence, DSCP and Class of Service (CoS), and internal labels that can be used in subsequent actions. Marking is used to identify the traffic type for use in policing, queuing, and scheduling traffic (only CoS is used in scheduling). Policing is the monitoring of data rates for a particular class of traffic. The device can also monitor associated burst sizes. Three "colors," or conditions, are determined by the policer depending on the data rate parameters supplied: conform (green), exceed (yellow), or violate (red). You can configure only one action for each condition. When the data rate exceeds the user-supplied values, packets are either marked down or dropped. You can define single-rate, dual-rate, and color-aware policers. Single-rate policers monitor the specified committed information rate (CIR) of traffic. Dual-rate policers monitor both CIR and peak information rate (PIR) of traffic. Color-aware policers assume that traffic has been previously marked with a color. The queuing and scheduling process allows you to control the bandwidth allocated to traffic classes, so you achieve the desired trade-off between throughput and latency. You can apply weighted random early detection (WRED) to a class of traffic, which allows packets to be dropped based on the Class of Service field. The WRED algorithm allows you to perform proactive queue management to avoid traffic congestion. You can schedule traffic by imposing a maximum data rate on a class of traffic so that excess packets are retained in a queue to smooth (constrain) the output rate. The QoS queuing features are enabled by default. Specific QoS-type features, policing and marking, are enabled only when a policy is attached to an interface. Specific policies are enabled when that policy is attached to an interface. By default, the

device always enables a system default queuing policy, or system-defined queuing policy map, on each port and port channel. When you configure a queuing policy and apply the new queuing policy to specified interfaces, the new queuing policy replaces the default queuing policy and those rules now apply.

9.1 QoS Configuration

- QoS Enabled: Yes
- QoS Classification required: Yes
- QoS Marking required: Yes
- QoS Policing required: No
- QoS Queuing required: Yes

World of Art provided the following QoS Classification / Marking Requirements:

According to the RFP documents, World of Art requires the following:

- 200Mb/sec minimum bandwidth for Real-time (Voice) traffic
- 200Mb/sec minimum bandwidth for Video traffic
- 200Mb/sec minimum bandwidth for Critical application traffic (Hybrid Exchange 2013 deployment, IBM R2 database, Oracle RAC solution)

10.0 ADVANCED FEATURES

10.1 OTV Configuration

World of Art has chosen to use OTV to extend certain VLANs between two data centers. The Atlanta data center will have two OTV edge devices for redundancy purposes.

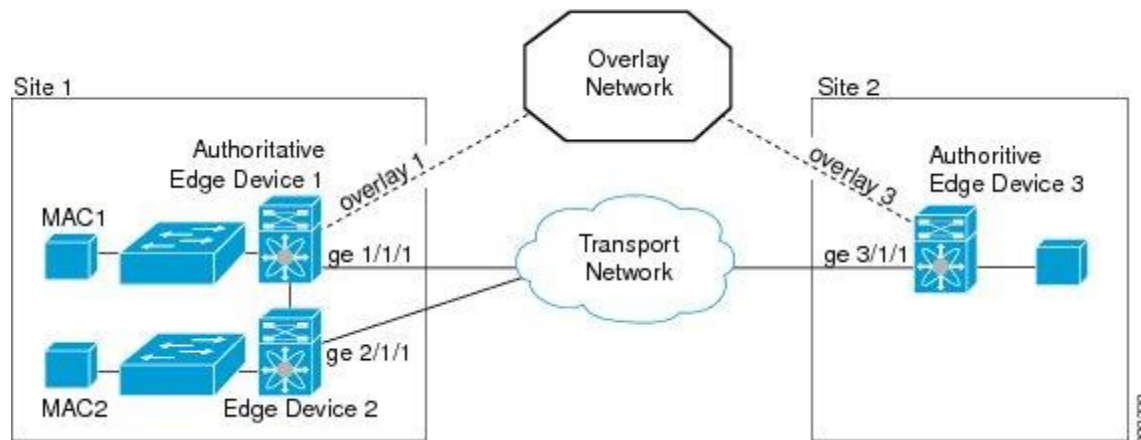
10.1.1 *Multihomed Sites and Load Balancing*

For resiliency and load balancing, a site can have multiple edge devices.

When more than one edge device exists in a site and both participate in the same overlay network, the site is considered multihomed. For the VLANs that are extended using OTV, one edge device is elected as an AED on a per-VLAN basis. OTV leverages a local VLAN to establish an adjacency between edge devices on their internal interfaces. The local VLAN that is shared by the internal interfaces is the site VLAN. The adjacency establishment over the site VLAN determines whether the other edge device is still present and which edge device is authoritative for what VLANs.

Load balancing is achieved because each edge device is authoritative for a subset of all VLANs that are transported over the overlay. Link utilization to and from the transport is optimized.

The figure below shows the AED that is selected for a multihomed site in an overlay network.



The following components need to be defined for each OTV Edge Device:

Internal Interface – This interface is a Layer 2 trunk that will carry all VLANs that will be extended into the OTV edge device

Join Interface – This interface is a Layer 3 interface that connects to the data network. This interface is used to communicate with the OTV edge devices at other sites and is used as the source address for all encapsulated traffic from this particular OTV edge device.

Overlay Interface – This interface holds the majority of the OTV configuration. It defines which VLANs will be extended and how OTV neighbors will be discovered.

The overlay interface is a logical interface that connects to remote edge devices in an overlay network through an associated physical interface on the transport network. From the perspective of MAC-based forwarding in a site, an overlay interface is simply another bridged interface. As a bridged interface, unicast MAC addresses are associated with an overlay interface. No STP packets are forwarded over an overlay interface. Unknown unicast packets are also not flooded on an overlay interface. From the perspective of IP transport, an overlay interface is not visible.

OTV encapsulates Layer 2 frames in IP packets and transmits them to an overlay interface.

Site VLAN – This VLAN is used for communication between OTV edge devices at the same site. This allows for Authoritative Edge Device role to be negotiated. This VLAN should not be extended on the overlay.

Site Identifier – This ID is used to help the OTV edge devices determine which peers are considered in the same location vs. a different one. This is used to help prevent loops between redundant OTV edge devices.

Neighbor Discovery – There are two modes of neighbor discovery for OTV. The first is multicast mode. This mode requires an Any-Source Multicast (ASM) address for neighbor discovery (Control Group) and a range of Source-Specific Multicast (SSM) addresses to carry multicast and broadcast traffic (Data Group). Multicast mode is more efficient because it lets the network replicate packets as needed. The second

mode is unicast mode. This mode requires at least one adjacency server so that all OTV edge devices can discover each other. This mode is simpler since it does not require a multicast enabled network, but can be less efficient due to head-end replication of packets that need to go to multiple sites.

Extended VLANs – This is the range of VLANs that will be extended between all sites participating in the configured overlay network.

The configuration parameters for each OTV edge device in this design are listed below:

10.1.2 OTV BCS1

<i>OTV Configuration Details</i>	
Internal Interface	interface e4/12
Join Interface	interface e4/11
Overlay Interface	int overlay1
Site VLAN	1
Site Identifier	0000.0000.0001
Neighbor Discovery (Multicast/Unicast)	Unicast
Primary Adjacency Server (Unicast)	192.168.0.2
Secondary Adjacency Server (Unicast)	
Extended VLANs	1010,1090

10.1.3 OTV BCS-2

<i>OTV Configuration Details</i>	
Internal Interface	interface e4/12
Join Interface	interface e4/11
Overlay Interface	int overlay1
Site VLAN	1
Site Identifier	0000.0000.0001
Neighbor Discovery (Multicast/Unicast)	Unicast
Primary Adjacency Server (Unicast)	192.168.0.2
Secondary Adjacency Server (Unicast)	
Extended VLANs	1010,1090

10.1.4 OTV BCS2-1

OTV Configuration Details	
Internal Interface	interface Gig 0/0/3 interface Gig 0/0/2
Join Interface	
Overlay Interface	int overlay1
Site VLAN	2
Site Identifier	0000.0000.0002
Neighbor Discovery (Multicast/Unicast)	Unicast
Primary Adjacency Server (Unicast)	192.168.0.1
Secondary Adjacency Server (Unicast)	192.168.0.2
Extended VLANs	1010,1090

11.0 DEVICE MANAGEMENT AND SECURITY CONFIGURATION

11.1 System Management Configuration

Each Cisco switch and router will be configured with the following System Management configuration:

11.1.1 Message of the Day (MOTD) Banner

```
banner motd #
```

```
*****
```

```
    This Device is the corporate property of  
    World of Art
```

```
    WARNING: Unauthorized access to this system is prohibited.  
    Violators are subject to criminal and civil penalties.
```

```
*****#
```

11.1.2 Network Time Protocol (NTP) Server

```
ntp server 10.200.1.72  
ntp server 10.1.0.2  
ntp server 10.1.0.3  
ntp server 10.10.4.6  
ntp master 5
```

11.1.3 Timezone Settings

clock timezone EST -5 0
clock summer-time EDT recurring

11.1.4 DNS Settings

ip domain-name worldofart.com
ip name-server 10.10.4.10
ip name-server 10.10.4.11

11.1.5 Syslog Configuration

By default the Nexus sends all log messages with timestamp enabled.

logging buffered 81920
logging enable
logging host 10.0.10.106

11.2 Smart Call Home Configuration

Smart Call Home analyzes Call Home messages and provides background information and recommendations. For known issues, particularly online diagnostics failures, Automatic Service Requests are generated with the Cisco TAC.

Smart Call Home offers the following features:

- Continuous device health monitoring and real-time diagnostic alerts.
- Analysis of Call Home messages and, if needed, Automatic Service Request generation, routed to the correct TAC team, including detailed diagnostic information to speed problem resolution.
- Secure message transport directly from your device or through a downloadable Transport Gateway (TG) aggregation point. You can use a TG aggregation point to support multiple devices or in cases where security dictates that your devices may not be connected directly to the Internet.
- Web-based access to Call Home messages and recommendations, inventory, and configuration information for all Call Home devices. Provides access to associated field notices, security advisories, and end-of-life information.

<i>Feature Name</i>	<i>Enabled/Disabled</i>
Smart Call Home	disabled

11.2.1 Example Call Home Configuration

```
callhome
contract-id 99887766
site-id 99
email-contact sysadmin@worldofart.com
phone-contact 678-503-3090
streetaddress 3625 Cumberland Blvd Atlanta Ga 30339
distribute
transport email reply-to sysadmin@worldofart.com
transport email mail-server 10.0.10.999 port 25 priority 50 use-vrf management
periodic-inventory notification interval 30
```

11.3 Simple Network Management Protocol V2 (SNMP) Configuration

SNMP Access is controlled through Control Plane Policing

<i>Feature Name</i>	<i>Enabled/Disabled</i>
SNMPv2	Enabled

11.3.1 SNMPv2 Configuration

```
snmp-server community XXXXXX RO
snmp-server community XXXXXX RW
snmp-server contact Sys Admin
snmp-server location Atlanta or Santa Clara
snmp-server enable traps
snmp-server host 10.0.10.106 version 2c XXXXXXXX
```

11.4 NetFlow Configuration

<i>Feature Name</i>	<i>Enabled/Disabled</i>
NetFlow	Enabled

NetFlow Version: 9

Example NetFlow Configuration:

```
flow exporter <NAME>
version 9
destination X.X.X.X (World of Art to provide once Netflow Server is online)
```

```
flow record <NAME>
  match ipv4 source address
  match ipv4 destination address
  collect transport tcp flags
  collect counter bytes
  collect counter packets
  collect flow sampler id
  collect ip version
flow monitor <NAME>
  record <RECORD NAME>
  exporter <EXPORTER NAME>
```

Interfaces where NetFlow is enabled:

```
ip flow monitor <NAME> input
ip flow monitor <NAME> output
```

Netflow Santa Clara

```
flow exporter FlowExporter1
  destination 10.0.10.106
  source GigabitEthernet0/0/0
  transport udp 2055
  export-protocol netflow-v5
```

12.0 COMPUTE & STORAGE ARCHITECTURE

12.1 Overview


The converged infrastructure data center design utilizing Cisco Unified Computing System (UCS) B-Series, EMC VNX Series storage arrays, and VMware vSphere provides the following features and benefits:

- High availability and redundancy through fabric connectivity and physical architecture.
- Increased throughput with the use of 10 Gigabit interfaces.
- Horizontal scalability within the UCS Fabric Interconnect for additional chassis.
- Optimized architecture for a virtualized environment.

12.2 Data Center Diagrams

The following diagram represent the physical architecture design of the Atlanta data center. The equipment is identical in the Santa Clara data center.

12.2.1 World Pay Architecture Design Diagram



12.3 Overview of Platform Features

12.3.1 Cisco Unified Computing System

12.4 Overview of Cisco UCS

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

The UCS architecture is a cohesive architecture that includes the following hardware components (shown in Figure 1 below):

- Cisco UCS 6200 Series Fabric Interconnects
- Cisco UCS 2200 Series Fabric Extenders (or IO Modules)
- Cisco UCS 5100 Series Blade Server Chassis
- Cisco UCS B-Series Blade Servers

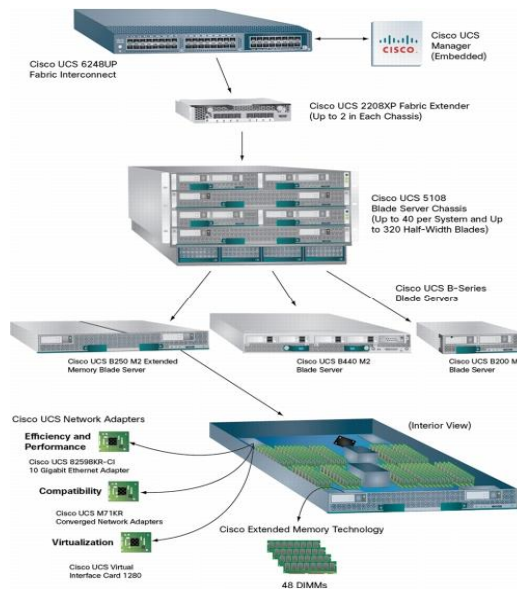
The Cisco UCS 6200 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet and Fiber Channel over Ethernet (FCoE) functions.

The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both the LAN and SAN connectivity for all blades within its domain

Cisco UCS 2200 Series Fabric Extenders bring the unified fabric into the blade server enclosure, providing 10 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The Cisco UCS 2200 Series extends the I/O fabric between the Cisco UCS 6200 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fiber Channel over Ethernet (FCoE) fabric or iSCSI deployment to connect all blades and chassis together. Since the fabric extender is similar to a distributed line card, it does not do any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling the Cisco Unified Computing System to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

Figure 1 – Cisco Unified Computing System Architecture



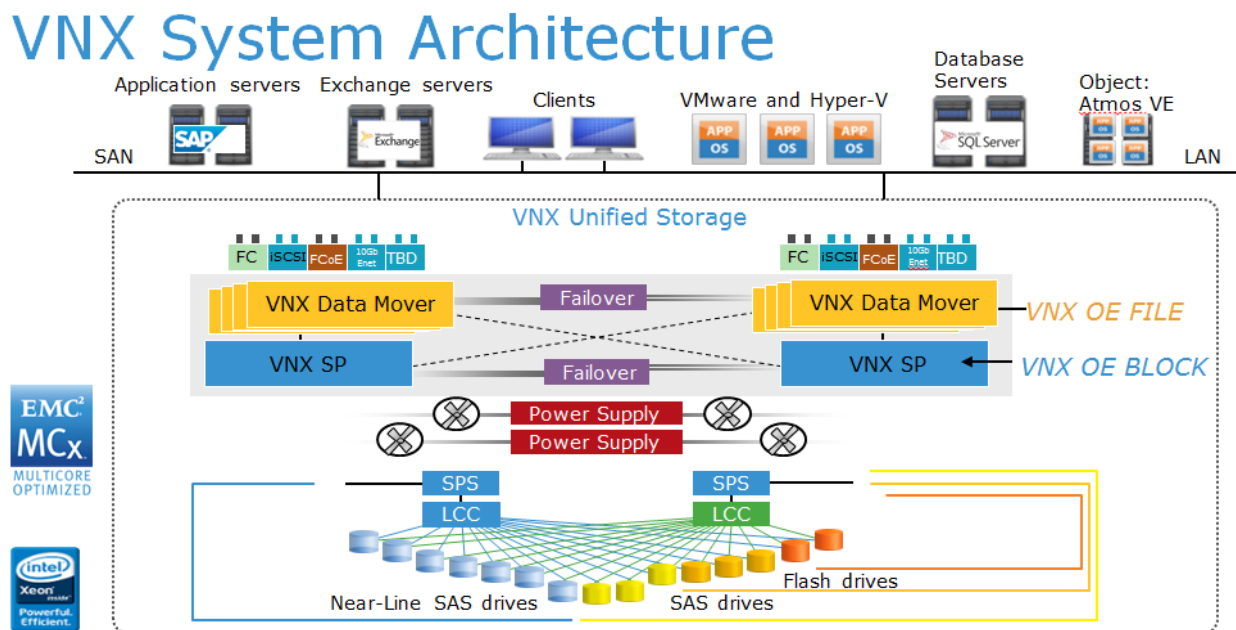
The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis for today's and tomorrow's data center while helping reduce TCO. The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A chassis can house up to eight half-width Cisco UCS B-

Series Blade Servers and can accommodate both half- and full-width blade form factors. The Cisco UCS B-Series Blade Servers are designed to increase performance, energy efficiency, and flexibility for demanding virtualized and non-virtualized applications. Based on Intel Xeon Series Processors, Cisco UCS B-Series Blade Servers adapt processor performance to application demands and intelligently scale energy use based on utilization. Each Cisco UCS B-Series Blade Server uses converged network adapters (CNAs) for access to the unified fabric. This design reduces the number of adapters, cables, and access-layer switches while still allowing traditional LAN and SAN connectivity.

12.4.1 EMC VNX Series Storage Arrays

EMC VNX Series provides high-performing unified storage with unsurpassed simplicity and efficiency, optimized for virtual applications. The VNX Series achieves new levels of performance, protection, compliance, and ease of management. The VNX series also leverages a single platform for file and block data services, centralized management makes administration simple, and data efficiency services reduce your capacity requirements up to 50 percent. The array is optimized for virtual applications with VMware and Hyper-V integration.

The following figure shows an overview of the VNX architecture:

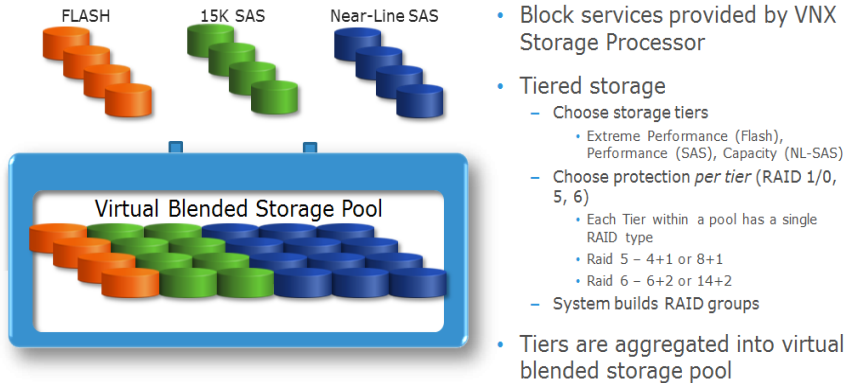


VNX System Architecture

This picture illustrates a unified storage product with scalable controllers. The VNX includes a fully integrated block processing component, the Storage Processor. In the VNX, the CPU and memory have been physically separated from the I/O complex.


The VNX disk technology used is either 2.5" or 3.5" technology and includes Flash drives, 7.2k, 10k, or 15k rpm SAS drive types, connecting natively to the SAS interface. The 7.2K high capacity drives are also referred to as Near-Line SAS (NL-SAS).

VNX Modular Architecture—Pools



VNX Modular Architecture – Pools. Physical devices are assigned to a pool and each specific drive type (SSD, SAS, or NL-SAS) is configured with a specific RAID type (RAID 1/0, RAID 5, RAID 6). All devices in the corresponding tier within a pool is configured with this RAID type and the system is built with the required RAID sets to ensure physical resiliency for the pool. Each tier must have the same RAID level ex. RAID 5 4+1 or RAID 5 8+1

The FAST Suite consists of two main components, FAST Virtual Pool (VP) and FAST Cache. The following picture shows an overview of these two technologies:



As shown in the above picture, FAST VP and Fast Cache use different data granularity (slices) in moving the data:

- **FAST VP:** Moves dynamically between the SSDs, HDD, and NL-HDD tiers in a 256 MB granularity.
- **FAST Cache:** Moves data dynamically between the FAST Cache and the disks (SSD, HDD, and NL-HDD) in 64 kB granularity.

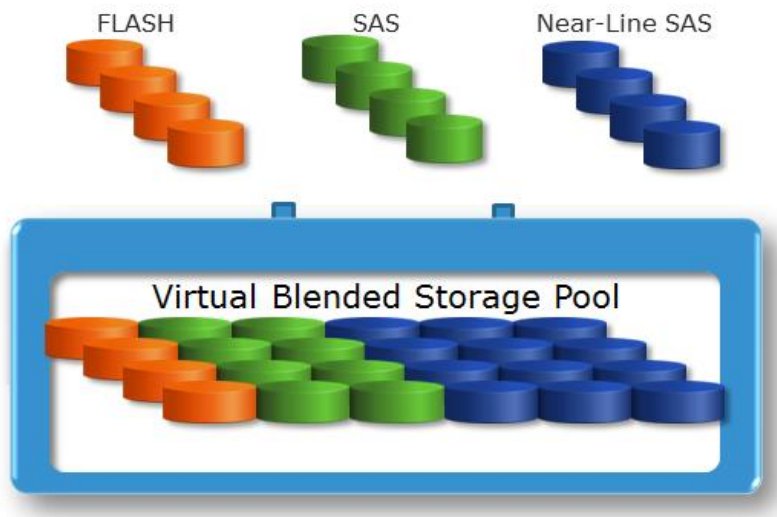
10.5.1 Virtual Blended Storage Pool (FAST VP)

FAST Virtual Pool (VP) assigns automatically different categories of data to different types of storage media within a tiered pool. Data categories can be based on performance requirements, frequency of use, cost, and other considerations. This is done by retaining the most frequently accessed or important data on fast, high performance (more expensive) drives, and moving the less frequently accessed and less important data to lower performance (less expensive) drives.

FAST VP functionality is available for both block data and file data. FAST VP optimizes storage utilization by automatically moving data between and within storage tiers.

VNX uses FAST VP-optimized SSD drives. When comparing these drives to other SSD's, FAST VP-optimized drives are more cost effective and appropriate when data change rates are more moderate. FAST VP-optimized drives cannot be used for FAST Cache. The VNX series used for DCP moves the data in 256 MB slices (in the previous series it was 1 GB) based on the access patterns of the I/O.

The following picture shows the different tiers based on the disk technology used for FAST VP:

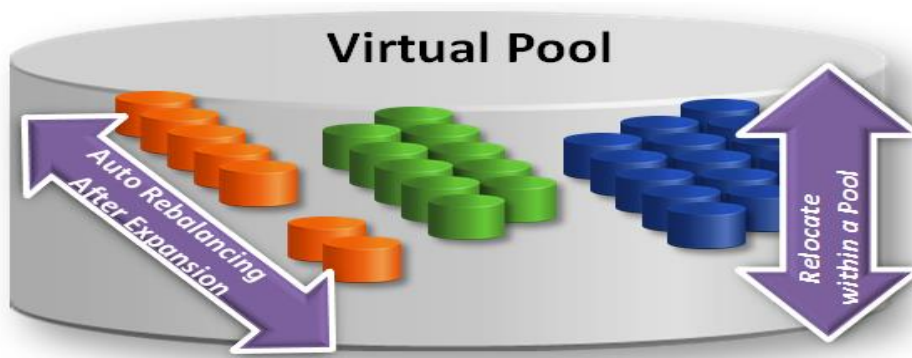


Fast VP has the following key facts:

- Fast VP tiers are based on disk technology:
 - Extreme Performance (Flash)
 - Performance (SAS)

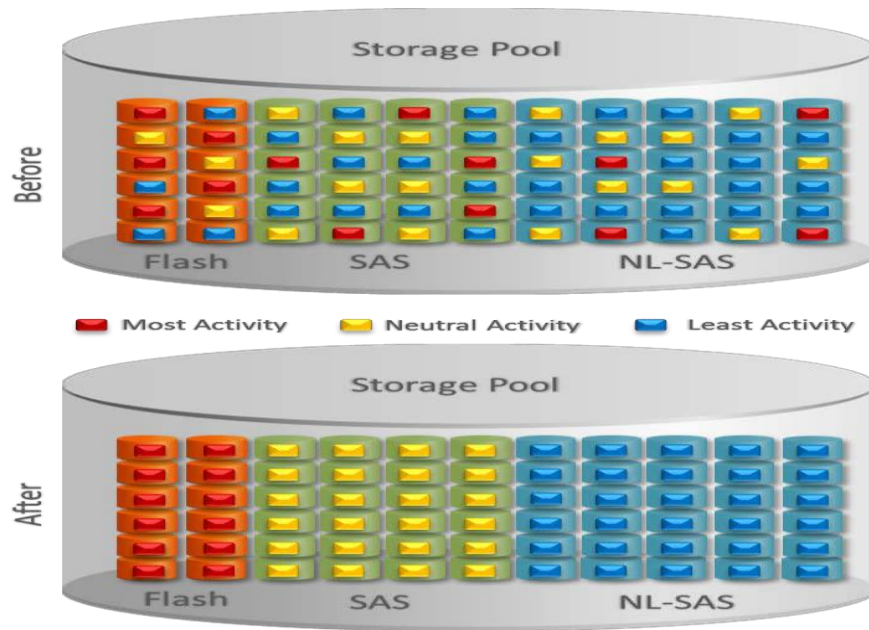
- Capacity (NL-SAS)
- Protection *per tier* (RAID5 and RAID6):
 - Each tier within a pool has a single RAID type
 - Used RAID type per disk technology:
 - SSD RAID 5 4+1
 - SAS RAID 5 8+1
 - NLSAS RAID 6 14+2
 - System builds RAID groups
- Tiers are aggregated into a virtual blended storage pool
- Virtual blended storage pool is provisioned for the file pool
- FAST is also known as “fully automated storage tiering”
- VP is also known as “virtual provisioning”

As a high-level overview, the following picture shows how FAST VP works:



Virtual Pool

VNX leverages FAST VP to migrate data to high-performance drives or high-capacity drives, depending on end-user access. As a result, customers require fewer drives and receive the best ROI from those drives that are configured.



Comparison of storage pools without FAST VP and with FAST VP

The storage pools at the top of the above figure show the initial storage configuration. After implementing FAST VP (at the bottom of the above figure) the system proactively optimizes the storage pool by moving the 256 MB slices of sub-LUN data to the most effective drive. This ensures that the appropriate data is housed on the right tier at the right time, which significantly increases efficiency and performance.

FAST VP moves the data between the different tiers based on a tiering policy. A tiering policy specifies where the initial placement of the data will be done and how this data will be relocated within the pool during the scheduled and/or manually invoked relocation periods. FAST VP bases decisions for how data relocation occurs on performance statistics collected every hour.

The following FAST VP tiering policies are available:

- Highest Available Tier
- Auto-Tier
- Start High the Auto-Tier
- Lowest Available Tier
- No Data Movement

User can set all FAST VP tiering policies except the “No Data Movement” policy both during and after LUN creation. The “No Data Movement” policy is only available after LUN creation.

In the World Pay environment, the FAST VP tiering policy “Auto-Tier” is used. The tiering (movement of the data/slices) is done between 21:00 and 05:00 local time.

The “Auto-Tier” policy works in following way:

A small portion of a large set of data may be responsible for most of the I/O activity in a system. FAST VP allows for moving a small percentage of the “hot” data to a higher tier while maintaining the rest of the data in the lower tiers. The “Auto-Tier” policy automatically relocates data to the most appropriate tier based on the activity level of each data slice. Slices provisioned to a LUN are relocated based on the highest performance disk drives available and the LUN’s slice temperature. Although this setting relocates data based on the LUN’s performance statistics, LUNs set with “Highest available Tier” take precedence.

The initial placement of slices within the Auto-Tier policy is based on the available capacity of a disk pool. For example, if 70% of a pool’s free capacity resides in the lowest available disks, then 70% of the new slices are placed on them.

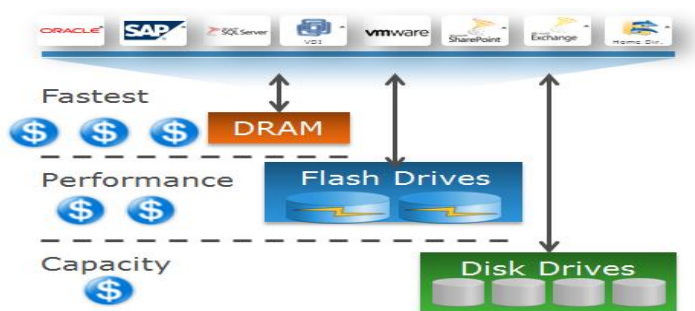
10.5.2 FAST Cache

FAST Cache is a large capacity secondary cache that uses enterprise Flash drives positioned between the Storage Processor’s DRAM-based primary cache and hard disk drives (HDD). The VNX series uses this feature to extend the array’s read-write cache and ensure that unpredictable I/O spikes are serviced at flash speeds, which benefits all applications.

FAST Cache perfectly complements FAST VP, as it works at a more granular level by copying 64 KB slices onto flash drives reserved for FAST Cache, depending upon the I/O characteristics. Multicore FAST Cache allows for faster initial cache warm-up, which results in immediate FAST Cache benefits.

Repeated access of the same 64 KB slice of data causes the policy engine to promote that data to FAST Cache. FAST Cache also works in quicker cycles than FAST VP; it has a response time on the order of microseconds to milliseconds. FAST Cache reacts to I/O spikes and maintains access levels by acting as an extension to onboard memory.

FAST Cache is most appropriate for workloads with a high locality of reference, for example, applications that access a small area of storage with very high frequency, such as database indices and reference tables. FAST Cache was not designed for very large I/O streams that are sequential, such as backups, because each 64 KB slice is accessed only once. Multicore Cache, however, uses various algorithms to optimize sequential I/O and handle these workloads. FAST Cache is most useful for handling spikes in I/O, and it benefits applications that may have unpredictable I/O profiles.

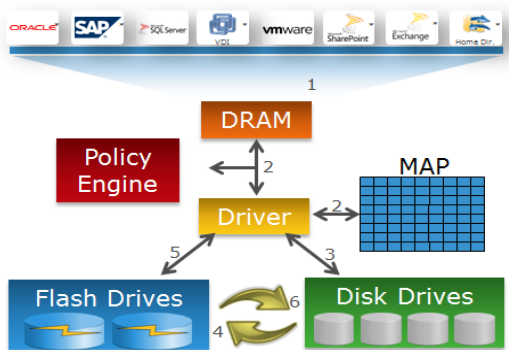


FAST Cache overview

FAST Cache has the following key facts:

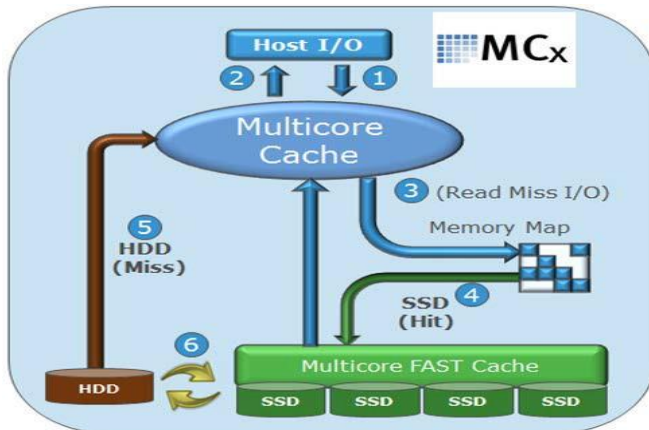
- Support for file
- Extends the system cache with Flash drives (1.2 TB of cache)
- Hot data automatically ends up in FAST Cache
- RAID 1 for Read/Write protection
- Transparent to SP failure; no need to warm up the cache

The following picture shows the steps of the FAST Cache process:



FAST Cache process

1. Page requests satisfied from DRAM if available
2. If not, FAST Cache driver checks map to determine where page is located
3. Page request satisfied from disk drive if not in FAST Cache
4. Policy Engine copies the page to FAST Cache if it is being used frequently
5. Subsequent requests for this page satisfied from FAST Cache
6. Dirty pages are copied back to disk drives as background activity



Virtual Provisioning is a strategy for efficiently managing space in a storage area network (SAN) by allocating physical storage on an "as needed" basis and is the underlying technology for Thin Provisioning. Thin Provisioning is a technology to present more storage to an application than is physically available.

Thin Provisioning enables organizations to reduce storage costs by increasing capacity utilization and simplifying storage management. Thin Provisioning also helps to reduce power and cooling requirements and reduce capital expenditures.

3.3.1 Pool-based Thin Provisioning


Thin Provisioning utilizes storage pool technology. A pool is somewhat analogous to a RAID Group, which is a physical collection of drives on LUNs.

Storage Pools allow you to take advantage of advanced data services like FAST VP, compression, and deduplication. Multiple drive types can be mixed into a pool to create multiple tiers with each tier having its own RAID configuration. They can contain a few drives or hundreds of drives whereas RAID Groups are limited to 16 drives. Because of the large number of drives supported in a pool, pool-based provisioning spreads workloads over many resources requiring minimal planning and management effort.

Storage pools can be used to create thick and thin LUNs. With thin provisioning, the user capacity (storage perceived by the host) can be larger than the available capacity on the storage system. Thin LUNs can be sized to accommodate growth without regard for currently available assets. Physical storage is assigned to the server in a capacity-on-demand fashion from the shared pool. The primary difference between thin LUNs compared to thick LUNs is that thin LUNs have the ability to present more storage to an application than what is physically allocated. Presenting storage that is not physically available avoids underutilizing the storage system's capacity. Data and LUN metadata is written to thin LUNs in 8 KB slices. Thin LUNs consume storage on an as-needed basis from the underlying pool. As new writes come into a thin LUN, more physical space is allocated in 256 MB slices. Unlike a thin LUN, a thick LUN's capacity is fully reserved and allocated on creation so it will never run out of capacity.

Virtual Provisioning technology also supports existing VNX features such as hot sparing, proactive sparing, and the ability to migrate data between thin LUNs, thick LUNs, or classic LUNs without incurring application downtime. The ability to non-disruptively migrate data to different LUN and disk types provides the best solution for meeting your changing application and business requirements without incurring downtime.

Thin Provisioning allows storage administrators to allocate storage on demand. It presents a host with the total amount of storage that has been requested; however, it only allocates storage on the array that is actually being used.



Example of an oversubscribed pool

The above used terms have the following meaning:

- **Total Capacity** is the amount of physical capacity available to all LUNs in the pool.
- **Total Allocation** is the amount of physical capacity that is currently assigned to LUNs.
- **Total Subscription** is the total capacity reported to the host.
- **Oversubscribed Capacity** is the amount of capacity that exceeds the capacity in a

12.4.2 Storage Configuration

12.4.2.1 VMware Datastores

Name	Size	VMFS Version	Cluster Assigned
EMC_LUN 6_DataStore	3TB	5.58	QTS Prod
EMC_LUN 7_DataStore	3TB	5.58	QTS Prod
EMC_LUN 8_DataStore	3TB	5.58	QTS Prod
EMC_LUN 9_DataStore	3TB	5.58	QTS Prod
EMC_LUN 10_DataStore	3TB	5.58	QTS Prod
EMC_LUN 11_DataStore_UC	1TB	5.58	UC Prod
EMC_LUN 12_DataStore_UC	1TB	5.58	UC-Prod
EMC_LUN 13_DataStore_UC	1TB	5.58	UC-Prod

12.5 EMC VNX 5400 Configuration

12.5.1 Basic Configuration Information

EMC Site ID:	
Model Number	5400
Hardware Type	VNX
System Type	Block/File(Unified)
Primary Hostname (cs0)	VNX5400_CS_1
Secondary Hostname (cs1)	VNX5400_CS_2
Connect Home S/N	TBD
Control Station 0 IP	10.200.1.51
Control Station 1 IP	10.200.1.52
SPA IP	10.200.1.53
SPB IP	10.200.1.54
Subnet Mask	255.255.255.0
Gateway	10.200.1.1
DNS Servers	10.10.4.5
NTP Server	64.90.182.55

12.5.2 Storage Processor Summary

	Storage Processor A	Storage Processor B
Storage System Serial Number	TBD	TBD
Model	VNX5400	VNX5400
Revision	05.33.000.5.015	05.33.000.5.015
PROM Revision	7.0.0	7.0.0
Memory	8192 MB	8192 MB
Statistics Logging	Enabled	Enabled

12.5.3 Disk Summary

Size (GB)	Type	Total
200	SSD	35
600	SAS 10k	47
2000	SAS 7.2k	87

12.5.4 Storage Pools / RAID Groups

Based on the evaluation of the storage available and the performance needs of the applications, EMC Proven Professional firm will configure one general pool and leave the SSD non-fast cache drives unallocated for future VDI needs.

Total Disks: 169

- General Pool Pool 0 (SQL/Exchange) w/FAST Cache
 - 900GB 10k SAS drives RAID5
 - 20 Disks, 600GB 10k SAS(4+1)
 - 3TB 7.2k NL-SAS drives RAID6
 - 16 Disks, 3TB 7.2k NL-SAS(14+2)

12.5.5 Initial Proposed RAID Group Layout

Enclosure	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
DAE3	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-0 3TB NL-SAS RAID6(14+2)	HS 3TB NL- SAS	HS 3TB NL- SAS	Pool-1 3TB NL-SAS RAID6(4+2)	Pool-1 3TB NL-SAS RAID6(4+2)	Pool-1 3TB NL-SAS RAID6(4+2)	Pool-1 3TB NL-SAS RAID6(4+2)	Pool-1 3TB NL-SAS RAID6(4+2)	Pool-1 3TB NL-SAS RAID6(4+2)																
DAE2	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)	Pool-2 3TB NL-SAS RAID6(14+2)											
DAE1	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)	Pool-1 3TB NL-SAS RAID6(14+2)											
BPE 0	Pool-1 Vault ROOGB SAS	Pool-1 Vault ROOGB SAS	Pool-1 Vault ROOGB SAS	Pool-1 Vault ROOGB SAS	Pool-1 iSSD SAS	Pool-1 ROOGB SAS RAID5(4+1)	Pool-1 ROOGB SAS RAID5(4+1)	Pool-1 ROOGB SAS RAID5(4+1)	Pool-1 ROOGB SAS RAID5(4+1)	Pool-1 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Pool-0 ROOGB SAS RAID5(4+1)	Cache 100GB SSD	Cache 100GB SSD	Cache 100GB SSD	Cache 100GB SSD	HS 1TB SSD	

13.0 CISCO UNIFIED COMPUTING SYSTEM (UCS) PRODUCTION ENVIRONMENT

13.1 Cisco UCS Configuration

The following sections detail the configuration of the Cisco UCS environment. Not all features may be used in every implementation.

13.1.1 Fabric Interconnects

The Cisco UCS 6296UP 96-Port Fabric Interconnect is a core part of the Cisco Unified Computing System. Typically deployed in redundant pairs, the Cisco UCS 6296UP Fabric Interconnects provide uniform access to both networks and storage. The fabric interconnects provide the management and communication backbone for the Cisco UCS B-Series Blades and UCS 5100 Series Blade Server Chassis.

Port Configuration

Fabric Interconnect A

Port	Device Connected	Function	Port	Device Connected	Function
1/1	Chassis-1, IOM1/1	Server	1/25		
1/2	Chassis-1, IOM1/2	Server	1/26		
1/3			1/27	Nexus 7004	Network
1/4			1/28	Nexus 7004	Network
1/5			1/29	Nexus 7004	Network
1/6			1/30	Nexus 7004	Network
1/7			1/31	Nexus 7004	iSCSI
1/8			1/32	Nexus 7004	iSCSI
1/9	Chassis-2, IOM1/1	Server			
1/10	Chassis-2, IOM1/2	Server	Mgmt0	Nexus 2248TP	Management
1/11					

Fabric Interconnect B

Port	Device Connected	Function	Port	Device Connected	Function
1/1	Chassis-1, IOM2/1	Server	1/25		
1/2	Chassis-1, IOM2/2	Server	1/26		
1/3			1/27	Nexus 7004	Network
1/4			1/28	Nexus 7004	Network
1/5			1/29	Nexus 7004	Network
1/6			1/30	Nexus 7004	Network
1/7			1/31	Nexus 7004	iSCSI
1/8			1/32	Nexus 7004	iSCSI
1/9	Chassis-2, IOM2/1	Server	Mgmt0	Nexus 2248TP	Management
1/10	Chassis-2, IOM2/2	Server			
1/11					

13.1.1.1 Addressing and Name Information

	Name	IP Address	Subnet Mask	Gateway
Fabric Interconnect A	UCS-FI-A	10.200.1.18	255.255.255.0	10.200.1.1
Fabric Interconnect B	UCS-FI-B	10.200.1.19	255.255.255.0	10.200.1.1

13.1.1.2 vNIC Templates

Name	Fabric ID	Type	VLAN's	MAC Pool
vNIC01	Fabric A	Updating	2,6,20,48,300,310,910	MAC-POOL-FIC-A
vNIC02	Fabric B	Updating	2,6,20,48,300,310,910	MAC-POOL-FIC-B
vNIC03	Fabric A	Updating	2,6,20,48,300,310,910	MAC-POOL-FIC-A
vNIC04	Fabric B	Updating	2,6,20,48,300,310,910	MAC-POOL-FIC-B

13.1.2 Server Configuration

The following outlines the configuration information for Server settings within UCS Manager.

13.1.2.1 UUID Suffix Pools

Name	Size
UUID-POOL-01	256

13.1.2.2 Maintenance Policies

Name	Reboot Policy
USER-ACK	User Ack

13.1.3 Admin Configuration

The following outlines the configuration information for Admin settings within UCS Manager.

13.1.3.1 Sub-Organizations

Function	Value
Sub-Organization Name	World Pay
Locales	QTS
Server Pool	Chassis-1/Blade-1, Chassis-1/Blade-2, Chassis-1/Blade-3, Chassis-2/Blade-1, Chassis-2/Blade-2

14.0 CONCLUSION

This design document educates the Client (World Pay) on all configuration options that are implemented and impacting existing business or future needs.

World Pay had to increase their ITs organization's ability to respond to new business needs while continuing to cut costs and improve overall reliability. The overall objective of any data center hardware/software refresh should include the following benefits:

- Reduced total cost of ownership.
- Improved infrastructure redundancy and high availability.
- Improved network and application performance, that can be measured and validated.
- Flexibility to implement private cloud focused around service management, applications, and organization.
- Greater systems compliance.
- Improved effectiveness of security including physical and logical endpoints.

EMC VNX series combined with Cisco data center technologies provides excellent hardware and software solutions to meet the requirements listed above.

Dell EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL EMC MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell EMC software described in this publication requires an applicable software license.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.