

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

Deploy Hitachi Unified Compute Platform Director in Unmanaged Network Mode With Cisco ACI

Reference Architecture Guide

By Hitachi Data Systems

September 2016

Contents

Executive Summary.....	3
Introduction to UCP Networking Components	4
Solution Overview.....	5
UCP 4000 With Cisco Rev 3 Hardware and Software Firmware Versions.....	7
Components	11
UCP Director Network Requirements	11
APIC Policy Configuration Components.....	12
UCP Management Network Configurations in ACI	13
Leaf Static Node Management Node ID and IP Address	13
Leaf Port Layouts	13
Networks and Customer-Defined VLANs	14
UCP Security Domain	15
UCP VLAN Pool	15
UCP Physical Domain	15
Interface Policy Groups	15
Interface Selectors	15
Leaf Profile	15
vPC Domain	15
Tenant UCP	16
Virtual Routing and Forwarding (VRFs)	16
Bridge Domains	16
Application Profile	16
Interface Bridge Domains	16
End Point Group (EPG) Static Bindings	16
VMM Domain Integration	16

APIC Configuration Methods.....17

Caveats17

Conclusion.....17

For More Information.....18

Executive Summary

Hitachi Unified Compute Platform Director (UCP Director) is a single point of management orchestrator designed to speed up an IT administrator's day-to-day deployments on server, storage and network components. It is designed to provide end-to-end visibility of the infrastructure with a unified single pane of glass. Today's evolving network architectures and application requirements are pushing many IT organizations to evaluate and explore software-defined networking (SDN) technology. One such SDN technology is Cisco Application Centric Infrastructure (ACI), which allows application-centric, policy-based network deployments. Network security has also been integrated into deployment via whitelist-based communications and micro-segmentation.

UCP Director is deployed on a validated Hitachi Unified Compute Platform 4000 (UCP 4000) with Rev 3 Cisco appliance architecture consisting of:

- Hitachi Compute Blade 500 (CB 500) chassis with HBx blades.
- Hitachi Virtual Storage Platform (VSP) models: VSP G1000, VSP Gx00, and VSP Fx00.
- Cisco Nexus 3048 Management switches.
- Cisco Nexus 9372PX Ethernet switches.
- Cisco Nexus 9332PQ Ethernet switches.
- Brocade 6510 Fibre Channel switches.
- Brocade 6546 In-Chassis Fibre Channel switches.

This white paper provides the requirements and considerations necessary to deploy UCP Director v4.1.2 or above, in unmanaged network mode, on a UCP 4000 with Rev 3 Cisco Appliance, with Cisco ACI.

Introduction to UCP Networking Components

In a UCP 4000 with Rev 3 Cisco appliance architecture, UCP Director manages the server, storage and networking components. These networking components include:

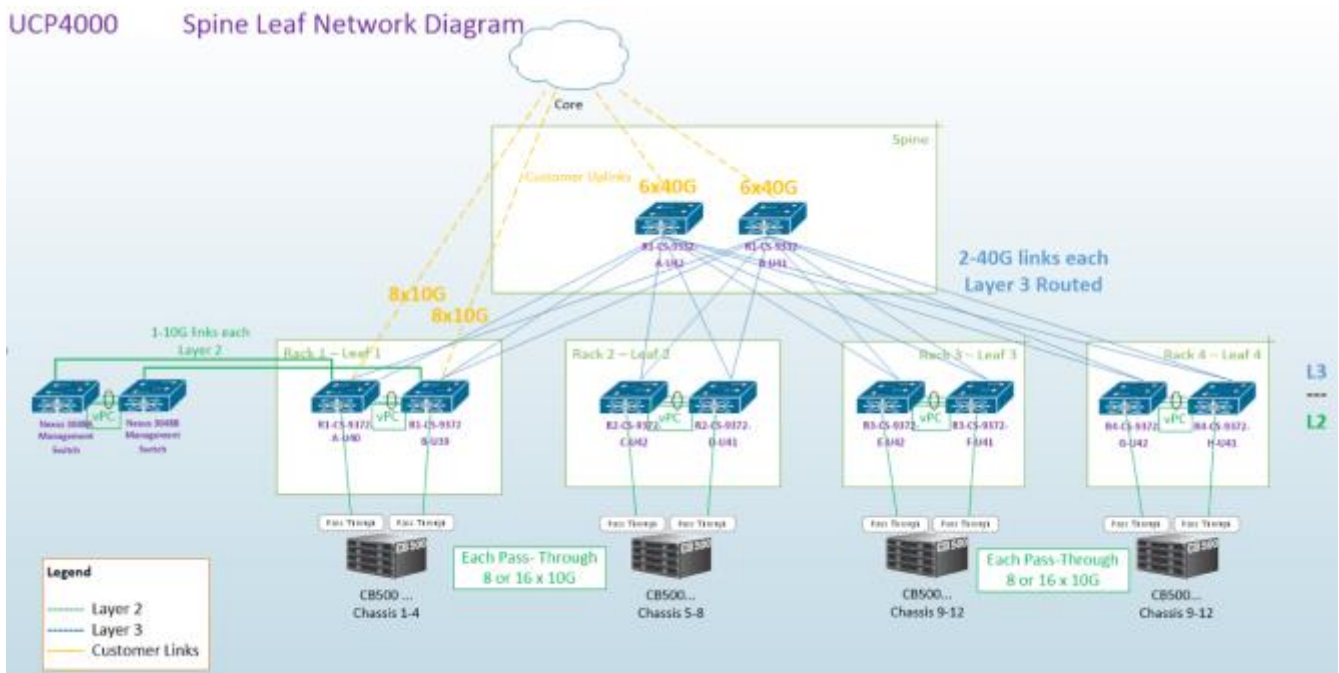
- Cisco Nexus 9332PQ aggregation switches.
- Cisco Nexus 9372PX access switches.
- Cisco Nexus 3048 management switches.

UCP Director manages the Cisco 9372PX and 9332PQ switches. There are two potential UCP Director modes of operation while the switches are running NX-OS: layer 2 mode and layer 3 mode.

In layer 2 mode, vPC-port channels make up the 40G links between the 9372PX and 9332PQ switches. UCP Director manages the server-facing interface-trunk-allowed virtual LANs (VLANs). It propagates VLANs to switch databases in layer 2 mode.

While in layer 3 mode, the individual 40G links are configured as layer 3 links. These have either Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP) and ECMP configured. Hot Standby Router Protocol (HSRP) gateway configurations are deployed on the 9372PX leaf switches (see Figure 1). UCP Director only manages the server-facing interface-trunk-allowed VLANs in layer 3 mode.

Figure 1. UCP 4000 Spine Leaf Network Diagram



UCP4000 = Hitachi Unified Compute Platform 4000, CB500 = Hitachi Compute Blade 500

Solution Overview

The UCP Director currently supports a software mode called unmanaged network mode. This allows organizations with special networking requirements to configure and manage all network configurations while UCP Director manages the server and storage provisioning. By running UCP Director in unmanaged network mode, ACI can be leveraged as network management software (See Figure 2).

UCP Director has specific network requirements that allow autodeploy functionality. With autodeploy functionality, service template deployments are just a click away. This document will cover the UCP network requirements and equivalent configurations needed in ACI.

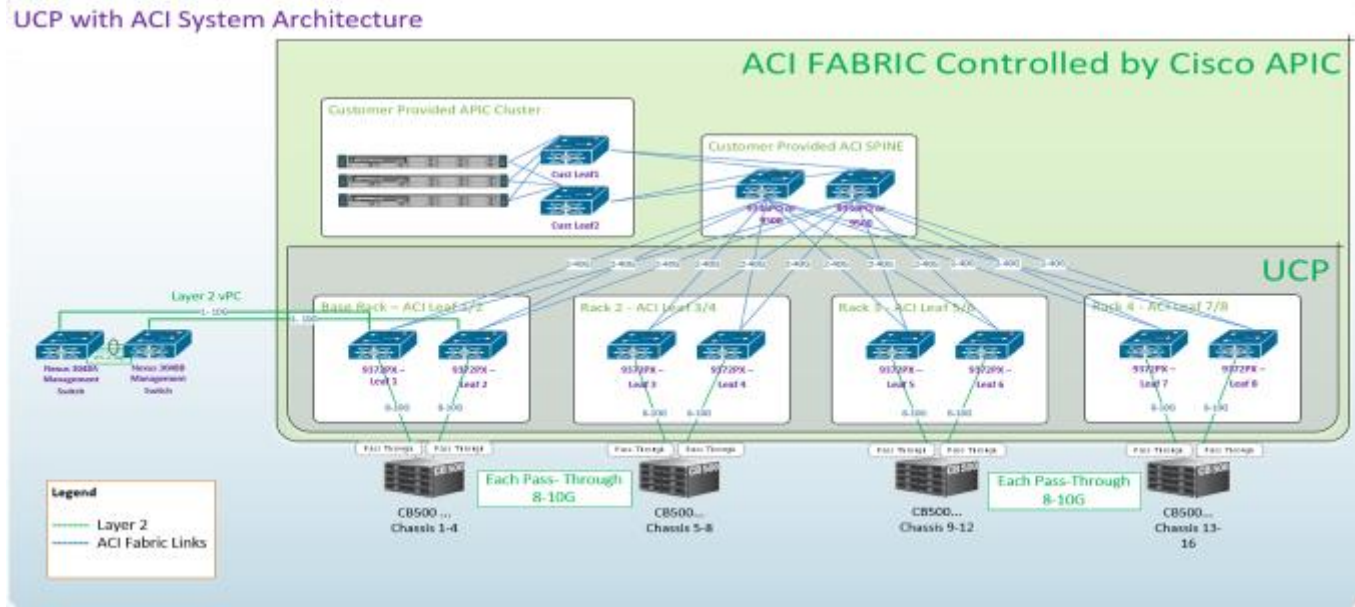
Make Enterprise Apps Agile With UCP 4000

[LEARN MORE](#)

ACI Components

- Spine switches (9336PQ or 9500 series spine switch).
- Leaf switches (9372PX).
- 3 UCS C220 M3 (APIC Cluster)

Figure 2. UCP With ACI System Architecture



UCP4000 = Hitachi Unified Compute Platform 4000, CB500 = Hitachi Compute Blade 500

As shown in the Figure 2, the UCP 4000 with Rev 3 Cisco appliance with an ACI system no longer comes equipped with a set of 9332PQ switches. Instead, the IT administrator provides a set of either 9336PQ or 9500 series spine switches, and has already configured the company’s APIC cluster. See additional prerequisites in Table 1.

Customer Requirements

- Spine switches (9336PQ or 9500 series spine switch).
- 3 UCS C220 M3 (APIC Cluster).
- APIC Cluster software version 1.3 or above.
- Switch software version 11.3 or above.

UCP Components

- Leaf switches (9372PX).

Table 1. Prerequisites

Prerequisites	Verification Check-off
Hitachi Unified Compute Platform 4000 (UCP 4000) with Cisco Rev 3 Appliance	_____
Hitachi Unified Compute Platform Director (UCP Director) v4.1.2 or above deployed in unmanaged network mode	_____

Customer-deployed APIC Cluster with firmware version 1.3 or above	_____
Customer-deployed Spine switches with up to 16 - 40G ports per spine switch	_____

UCP 4000 With Cisco Rev 3 Hardware and Software Firmware Versions

Table 2. Hardware Firmware Versions

Hardware	Version	Quantity
2U4N Management Server		2 nodes
2u4n BIOS	S2S_3A19	
2u4n BMC	3.42.00	
Intel i350	1.5.53	
Intel 82599	2.3.24	
Emulex OCe11102 firmware	10.6.144.21	
Emulex LPe12002 boot (MRC and MRB only)	2.20a4	
Emulex LPe12002 firmware (MRC and MRB only)	2.02a0	
Compute Chassis		
Hitachi Compute Blade 500 (CB 500) SVP Firmware	A0290-B-10707	
CB 500 SVP equipment parameter file	1022	
CB 500 SEL dictionary file	A0151	
CB 520H B3		Up to 8 blades per chassis
CB 520H B3 firmware file	520h_x3_0848 or 520h_x3_0830	
Blade firmware version	08-48 or 08-30	

Hardware	Version	Quantity
BMC version	08-47 or 08-30	
EFI version	08-48 or 08-29	
CB 520H B3 CNA On-board LOM	10.6.144.2704	

CB 520H B4		Up to 8 blades per chassis
CB 520H B4 firmware file	520h_x4_1002-03	
Blade firmware version	"10-02"	
BMC version	"10-02"	
EFI version	"10-02"	
CB 520H B4 CNA on-board LOM	10.6.144.2704	
Fibre Channel host bus adapter from Hitachi (16Gb/sec)	40-03-0C	

Storage: Hitachi Virtual Storage Platform (VSP)		
VSP G1000	80-03-10-00/00-M054	
VSP G200	83-03-02-20/00	
VSP G400	83-03-02-40/00	
VSP G600	83-03-02-40/00	
VSP G800	83-03-02-60/00	
VSP F400	83-03-02-40/00	
VSP F600	83-03-02-40/00	
VSP F800	83-03-02-60/00	

Cisco		
Cisco Nexus 9372	7.0(3)I1(3a)	Minimum 2 up to 8
Cisco Nexus 3048	6.0(2)U6(5)	2

Hardware	Version	Quantity
ACI		
APIC Cluster	1.3(2h) or 2.0(1n)	Minimum 3 nodes
Cisco Nexus 9372	11.3(2h) or 12.0(1n)	Minimum 2 up to 8
Brocade Fibre Channel		
Brocade 6510	v7.3.1c	Minimum 2 up to 4
Brocade 6546 (in-chassis switch)	v7.3.1c	2 per compute chassis

Table 3. Software Firmware Versions

Software	Version
Software	
Hitachi	
Compute Systems Manager (HCSM)	8.4.0-00
Device Manager (HDvM)	8.4.0-00
Tuning Manager (HTnM)	8.4.0-00
Tuning Manager Agent for RAID (HTnM Agent for RAID)	8.4.0-00
Hitachi Unified Compute Platform (UCP)	
UCP Director	4.1.2.70
UCP Director operations center	4.1.2.70
VMware	
Version 6.0	
vCenter	6.0 U2
ESXi	6.0 U2
Site Recovery Manager (SRM)	6.1
Linux	
CentOS (WDS)	6.6,7.2

Components

2U4N Management Server

The default UCP configuration for the 2U4N management server runs with two half-width nodes. UCP's management stack operates on the two-node configuration as virtual machines that are deployed on VMware vSphere 6.0 U2. The two nodes are configured in a VMware cluster for high availability (HA).

Hitachi Compute Blade 500

Hitachi Compute Blade 500 is a high-end compute platform designed with virtualization and cloud computing in mind. The chassis can scale up to eight servers. It houses pass-through modules for direct 10G connections between blade and leaf switch, as well as Brocade 6546 16Gb/sec in-chassis Fibre Channel switches.

Hitachi CB 520H B4 Blade Servers

The CB 520H B4 blade servers provide world-class performance with Intel Xeon processors and scalable memory. Customizable service templates on the blade servers let you deploy ESXi stateless hosts, Microsoft® Windows®, Linux, and custom-operating systems with the click of a button in UCP Director.

Storage

World-class Hitachi Virtual Storage Platform models, including VSP G1000, VSP Gx00 and VSP Fx00 provide storage for all sorts of business needs. The UCP management stack is deployed on the array of choice.

Networking

The Cisco 9300 series switches provide Cisco's next-generation, platform-supporting NX-OS and ACI modes of operation. The switches are deployed in the access-leaf and aggregation-spine layers of the UCP 4000 with Cisco Rev 3 network. With 10G server-facing interfaces and 40G interswitch connectivity, applications are able to leverage the increased bandwidth and performance. In addition, the Nexus 3048 switches are used as dedicated 1G management switches.

Fibre Channel

The Brocade 6510 switches aggregate Fibre Channel traffic from the in-chassis Brocade 6546 switches.

Application Centric Infrastructure (ACI)

Cisco's ACI is Cisco's next-generation SDN platform, which offers policy-based networking configurations designed around application needs, not the other way around.

Application Policy Infrastructure Controller (APIC)

There are typically 3 UCS nodes running as an APIC cluster that applies policies to the ACI fabric and endpoints. In addition to controlling the infrastructure and pushing policies, the APIC software has built-in-network monitoring, and performs software upgrades.

UCP Director Network Requirements

In network-unmanaged mode, UCP Director provides management of server and storage components while giving the customer the flexibility to configure the network. Components and features of unmanaged network mode include:

- CB 520H B4 servers provisioning.
- VSP G1000, VSP Gx00 and VSPFx00 provisioning.

- Provisioning ESXi Standalone, ESXi Cluster, Linux, Windows and Custom OS.
- Server, storage and network monitoring.

UCP Director has specific network requirements to allow components to communicate with each other. Many of the out-of-band management interfaces are reachable via the Nexus 3048 1G management network, but a vPC port-channel is required between the Nexus 3048 switches and first rack 9372PX leaf switches. The management stack that resides on the 2U4N management nodes run in an ESXi cluster. These nodes are connected to the first rack of the 9372PX leaf switches, via 10G connections, and 3048 management switch, via 1G connections.

All CB 520H B4 server-facing interfaces are 10G and connected to 9372PX leaf switches.

- Two-port mode allows the number of blades to scale up to 128 servers.
- Four-port mode allows the number of servers to scale to 88 servers.

Each blade server is connected to a Hitachi pass-through module, which then directly connects to the 9372PX leaf switches at 10G speeds per port.

UCP Director provides autodeploy functionality for ESXi, Windows and Linux OS. Autodeploy requires that each CB 520H B4 blade server switch port is configured as a trunk interface, with a native VLAN configured. This native VLAN is the default management VLAN specified by the customer. It is required, because when the blade is initially powered-on, PXE-boot has no VLAN tagging mechanism. Thus, these untagged packets will reside on the native VLAN and the DHCP server will be sitting on the same VLAN. It will be ready to respond to the initial DHCP discover packets and ready to proceed with the operating system deployment. Either ESXi autodeploy or WDS mechanisms are used to bring up the respective servers with their operating systems.

In a typical UCP 4000 with Cisco Rev 3 appliance, deployed with UCP Director, the network is managed. The network is preconfigured and the switches are running in NX-OS mode upon arrival at the customer site. With ACI integration, this is not possible; switches must be running in ACI mode, and network configurations must be applied at the APIC Cluster level. The UCP network settings need to be configured in a way that satisfies the blade provisioning network requirements on the ACI Fabric.

APIC Policy Configuration Components

Leaf Static Node Management Addresses

When the 9372PX leaf switches are converted to ACI mode, the out-of-band management IP address should be configured. In ACI, this is known as the static node management address, and each spine and leaf switch should be assigned an IP for management reachability.

Security Domain

A security domain is a means for tagging components. Users in the same domain can access each other's security domain if they have access rights.

VLAN Pool

A VLAN pool must be defined with all of the LANs that will be used in a domain.

Physical Domain

A physical domain consists of all of the physical resources, including ports, port-channels and VLANs, that are used for endpoint groups in the domain.

Interface Access Policy

The interface access policies define the settings on an interface, such as whether or not to enable Link Layer Discovery Protocol (LLDP), Cisco Discovery Protocol (CDP) or Spanning Tree Protocol (STP). These policies are then applied on the switch level in specific interface groupings.

vPC Domain and Port-Channel

A vPC domain can be configured between two leaf switches in an ACI fabric, and vPC port-channels can be configured facing traditional non-ACI switches. UCP requires this port-channel setup between the first-rack leafs and the Nexus 3048 management switches for management traffic.

Tenant

A tenant in ACI is a container for application policies that can be reused for specific customers, users or organizations.

Bridge Domain

A bridge domain houses specific layer 2 broadcast domains. It is comparable to a VLAN; however, multiple VLANs can be part of a bridge domain. Additionally, a bridge domain should be associated to a VRF.

Application Profile

An application profile can be used to group specific application endpoint needs. For example, a server may need access to multiple network segments, and the application profile may have many endpoint groups created that can be applied to the server's interfaces.

Endpoint Groups

An endpoint group (EPGs) is a grouping of application endpoints. The endpoint may be a bare-metal server, virtual machine (VM) or external device.

802.1P Tagging

802.1P tagging is analogous to the native VLAN in ACI.

UCP Management Network Configurations in ACI

Leaf Static Node Management Node ID and IP Address

When the UCP 9372PX leaf nodes join the customer ACI fabric, an out-of-band management IP address should be specified. In ACI, nodes need to be assigned a node ID. Typically, the spine nodes are assigned a 2xx numbering scheme, while the leaf nodes are assigned a 1xx numbering scheme. If certain leaf nodes are already deployed, the 101-108 default UCP node ID scheme may not be usable. In this case, the customer will have to specify the desired node ID and management IP address

Leaf Port Layouts

Table 4. Port Allocation (2 Port Server)

Usage 2 Port Mode	Ethernet Ports	Speed
9372 A,B,C,D,E,F,G,H	E1/1 - 32	10Gb/sec
9372 A,B to 2U4N Mgmt Server	E1/45 - 47	10Gb/sec
9372 A,B ISL to 3048	E1/48	10Gb/sec

Figure 3. Switch Layout (2 Port Server)





Table 5. Port Allocation (4 Port Server)

Usage 4 Port Mode	Ethernet Ports	Speed
9372 A,B	E1/1 - 32	10Gb/sec
9372 A,B to 2U4N Mgmt Server	E1/45 - 47	10Gb/sec
9372 A,B ISL to 3048	E1/48	10Gb/sec
9372 C,D,E,F,G,H	E1/1 - 48	10Gb/sec

Figure 4. Switch Layout (4 Port Server)



Networks and Customer-Defined VLANs

When purchasing the UCP, the organization should define the VLANs listed in Table 6. The customer-defined values are used to identify the VLAN for each of the network segments.

Table 6. VLAN Allocation

Network	Customer-Defined VLAN
ucpmanagement	
ucpvmotion	
ucpcompute	

UCP Security Domain

The UCP Security Domain will be defined as “UCP.” All relevant UCP components, such as the UCP Physical Domain and UCP Tenant, will be tagged accordingly.

- UCP.

UCP VLAN Pool

The VLAN pool will be defined as “UCP-VLAN-Pool.” It will be configured in the UCP security domain with the customer-defined VLANs set to dynamic mode.

- UCP-VLAN-Pool.

UCP Physical Domain

The Physical Domain will be defined as “UCP-Physical-Domain” and the UCP-VLAN-Pool will be defined under this “UCP-Physical-Domain.”

- UCP-Physical-Domain.

Interface Policy Groups

The following two policy groups, which are part of the “UCP-Physical-Domain,” should be defined:

- Server facing policy group, “UCPServer-Access-Port-Policy”
- vPC port-channel policy group, “UCP-vPC-Port-Channel-Interface”

Cisco Discovery Protocol (CDP) is enabled on both policy groups. The vPC Port-Channel Policy Group specifies that Link Aggregation Control Protocol (LACP) mode active is used.

Interface Selectors

Before interface policy groups can be applied, interface selectors must be created to define which leaf interfaces are grouped together. A vPC port-channel interface selector with port E1/48 needs to be created for the vPC port-channel interface connection to the Nexus 3048 management switches. The server-facing, interface selector will differ depending on whether two- or four-port mode is selected. First rack leaves will also have ports 45-47 defined. The interface policy groups will be applied depending on the port’s respective use case. See [Leaf Port Layouts](#) for the two- and four-port allocation diagrams.

Leaf Profile

The leaf profile consists of the interface selectors and leaf groups. For example, the vPC port-channel interface selector can be applied to the leaf group of nodes 101 and 102. The same will be configured for the server-facing interfaces. See [Leaf Port Layouts](#) for the two- and four-port allocation diagrams.

vPC Domain

A vPC domain 901 will be defined for the first rack 9372PX leaf switches.

Tenant UCP

The tenant “UCP” is defined as the policy container for ucpsmanagement, ucpsmotion, and ucpscompute. It is also part of the UCP security domain.

- UCP.

Virtual Routing and Forwarding (VRFs)

The following three VRFs are defined:

- ucpsmanagement – Used for ucpsmanagement network.
- ucpsmotion – Used for ucpsmotion network.
- ucpscompute – Used for all other compute network traffic.

Bridge Domains

Three bridge domains are defined below. For ucpsmanagement bridge domain, the L2 Unknown Unicast setting should be configured as “Flood,” due to the fact that the Nexus 3048 management switch is an external device.

- ucpsmanagement – Used for ucpsmanagement network.
- ucpsmotion – Used for ucpsmotion network.
- ucpscompute – Used for all other compute network traffic.

Application Profile

The following endpoint groups are defined under the application profile “UCP”:

- ucpsmanagement – Part of ucpsmanagement bridge domain.
- ucpsmotion – Part of ucpsmotion bridge domain.
- ucpscompute – Part of ucpscompute bridge domain.

Interface Bridge Domains

Interface bridge domains are comparable to switched virtual interfaces (SVIs) and can be defined under each bridge domain.

End Point Group (EPG) Static Bindings

The static bindings apply the specific endpoint configurations to specific leaf interfaces. In UCP, the ucpsmanagement EPG needs to be statically bound to all server-facing interfaces in 802.1P mode, with the management VLAN configured. The vPC port-channel interface should be set as trunk interface, with the management VLAN configured. If vmotion is required, the ucpsmotion EPG should also be deployed on all server-facing interfaces. The ucpscompute EPG, as well as any additional EPGs, should be deployed as necessary by the customer.

- ucpsmanagement – Access (802.1P).
- ucpsmotion/ucpscompute – Trunk.

VMM Domain Integration

ACI leverages the benefits of microsegmentation by supporting integration into vCenter with a distributed virtual switch (DVS). Currently, only DVS has been tested to work with UCP Director.

APIC Configuration Methods

The UCP APIC configuration guide supports the following two methods:

- CLI.
- GUI.

The CLI and GUI both call the same APIs and are equivalent methods of configuring the APIC Cluster. Please reference the UCP APIC configuration guide for step-by-step CLI or GUI, UCP network configuration steps.

Caveats

Cisco has outlined the following hardware limitations for the 9372PX switches with regards to multiple EPGs.

- On older hardware operating on APIC version 1.3(1) and NX-OS software 11.3(1), packets configured with multiple EPGs (802.1P mode and trunk mode) packets will exit the port tagged as “VLAN zero.”
- On newer hardware operating on APIC version 1.3(1) and NX-OS software 11.3(1), packets configured with multiple EPGs (802.1P mode and trunk mode) packets will exit the port tagged as “VLAN untagged.”

Cisco defines the following as new hardware:

- N9K-C93180YC-EX 48-port 10/25 Gigabit and 6-port 40/100 Gigabit QSFP28.
- N9K-C9504-FM-E supporting 100 Gigabit I/O modules.
- N9K-C9508-FM-E supporting 100 Gigabit I/O modules.
- N9K-X9732C-EX 32-port 40/100 Gigabit QSFP28.

Cisco defines the following as older hardware:

- N9K-C9372PX.

In CentOS, packets tagged as “VLAN zero” are dropped by the operating system. As a workaround, the *modprobe 8021q* module must be loaded in the kernel. Once loaded, the incoming packets can be correctly identified by the CentOS.

During the UCP service template deployment of CentOS 6.6 or 7.2, it is required that the server-facing port only have the *ucpmanagement* EPG configured. Configuring multiple EPGs with the *ucpmanagement* EPG will cause the EPG CentOS deployment to fail. Once the operating system is successfully deployed, the *modprobe 8021q* module should be loaded into the kernel. Then additional EPGs can be configured on the leaf server-facing interface.

Bring-up of vPC port-channel between the Nexus 3048 management switches and Rack1 Nexus 9372PX Leaf switches, ports may not aggregate. If this occurs, it may be necessary to flap the individual ports and port-channels on both sides to aggregate the vPC port-channel correctly.

Conclusion

UCP Director in unmanaged network mode can seamlessly integrate with Cisco ACI to leverage the power of both UCP's server and storage provisioning and Cisco's software-defined networking. For many corporations, the ability to quickly

scale and deploy is becoming a requirement with today's ever-changing data-centers, hybrid-cloud and cloud environments.

For More Information

Additional information about the Cisco Nexus 9000 Series ACI-mode switches can be found on [Cisco's website](#). Find more information about the Hitachi Unified Computer Platform products and technologies at <https://www.hds.com/en-us/products-solutions/converged-infrastructure.html>. Pricing information is available through authorized resellers. For a demonstration of the capabilities of UCP Director, please contact serversales@HDS.com.

