

Cisco Desktop as a Service Solution with Deskton



August 6, 2013

Contents

- Contents 1
- About Cisco Validated Design (CVD) Program 4
- 1 Preface 5
 - Introduction 5
 - Target Audience 5
- 2 Cisco DaaS Solution with Deskton Overview 6
 - Virtualized Multiservice Data Center 6
 - Wide Area Network 8
 - Security 8
 - Multi-Tenant Connection Broker 8
 - Cisco DaaS Solution with Deskton Components 8
- 3 Design Considerations 10
 - Tenant Separation 10
 - Compute 10
 - Network 11
 - Network Services 12
 - Storage 12



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2013 Cisco Systems, Inc. All rights reserved

- Desktop 12
- High Availability 13
 - Compute 14
 - Network 16
 - Network Services 16
 - Storage 16
 - Desktop 17
- 4 Solution Validation 18
 - Configuring the Cisco Unified Computing System 18
 - To configure the Cisco Unified Computing System, perform the following steps: 18
 - Installing and Configuring VMware vSphere 32
 - Install VMware ESXi 32
 - Install and Configure VMware vCenter 32
 - Install Licenses 33
 - VMware vSphere Network Configuration for Deskton Appliances 34
 - VMware vSphere Network Configuration for Deskton Tenant Appliances 35
 - Creating the Golden Image 35
 - Steps to Customize the Golden Image Virtual Machine 36
 - Installing the Deskton DaaS Agent. 38
 - Installing the VMware View Agent 39
 - Installing the VMWare View Agent Connect 41
 - Installing and Configuring Deskton 43
 - Installing Deskton Primary Service Provider Appliance 43
 - Configuring the Network on Deskton Primary Service Provider Appliance 47
 - Configuring the Deskton Primary Service Provider Appliance 50
 - Installing the Deskton Secondary Service Provider Appliance 54
 - Installing the Deskton Tenant Resource Manager Appliances 55
 - Adding Desktop Models 57
 - Installing the Deskton Tenant Appliances 57
 - Assigning Hosts to Individual Tenants 60
 - Assigning Quotas for Tenants 63
 - Configuring Tenant Appliance 64
 - Adding the Golden Template into Deskton Tenant Portal 66
 - Creating a Pool of Desktops in the Tenant Portal 68
- 5 Test Setup and Configuration 72
 - Cisco UCS Test Configuration for Single Blade Scalability 72
 - Hardware Components: 72
 - Detailed Windows Configuration 73
 - Testing Methodology and Success Criteria 73
 - Load Generation 73

User Workload Simulation – LoginVSI from Login Consultants	74
Testing Procedure	75
Pre-Test Setup for Single and Multi-Blade Testing	75
Test Run Protocol	75
Success Criteria	76
Login VSI Max	76
Calculating VSIMax	76
VSIMax Dynamic	77
Determining VSIMax	78
6 VDI Test Results	80
Login VSIMax Score	80
Single Blade Maximum Recommended Workload	81
Boot Phase Performance Results	82
Test Phase Performance Results	86
7 Appendix	91
8 Related Documents	167

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit <http://www.cisco.com/go/designzone>.

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

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

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Desktop as a Service Solution with Desktone

© 2013 Cisco Systems, Inc. All rights reserved.

1 Preface

Introduction

The Cisco Desktop as a Service (DaaS) Solution with Desktone enables service providers to offer desktops as a service to their customers. The Cisco DaaS Solution with Desktone is built to be highly scalable and supports multi-tenancy.

Cisco DaaS Solution with Desktone delivers desktops as a service on a common infrastructure. This common infrastructure is hosted in service provider data centers. Microsoft Windows 7 is the operating system of choice for the virtual desktops and is available in persistent and non-persistent flavors. The Cisco DaaS Solution with Desktone is built by combining data center, network, and security technologies from Cisco and desktop virtualization technologies from our ecosystem partners.

Cisco DaaS Solution with Desktone accomplishes complete separation between customers from desktops in the data center, through the data center network, across the wide area network, and down to the customer network. Customer visibility is restricted to only their virtual desktops. However, service providers have insight into how resources as a whole are utilized by different customers.

The pod-based approach adopted for the Cisco DaaS Solution with Desktone enables flexibility in growth based on business needs without requiring a redesign of the data center. This approach ultimately lowers operational costs for the service provider. The pod as a logical unit supports multiple customers at different scale on a common infrastructure.

Network services deliver exceptional user experience for a large number of users without requiring additional resources to be deployed on the customer end. Security policies are enforced between the customer's network and the virtual desktops in the service provider's data center.

Cisco DaaS Solution with Desktone has been tested end-to-end to confirm that all the different components interoperate. With Cisco DaaS Solution with Desktone, service providers are now able to build a multi-tenant desktop as a service solution faster and at reduced risk than ever before.

Target Audience

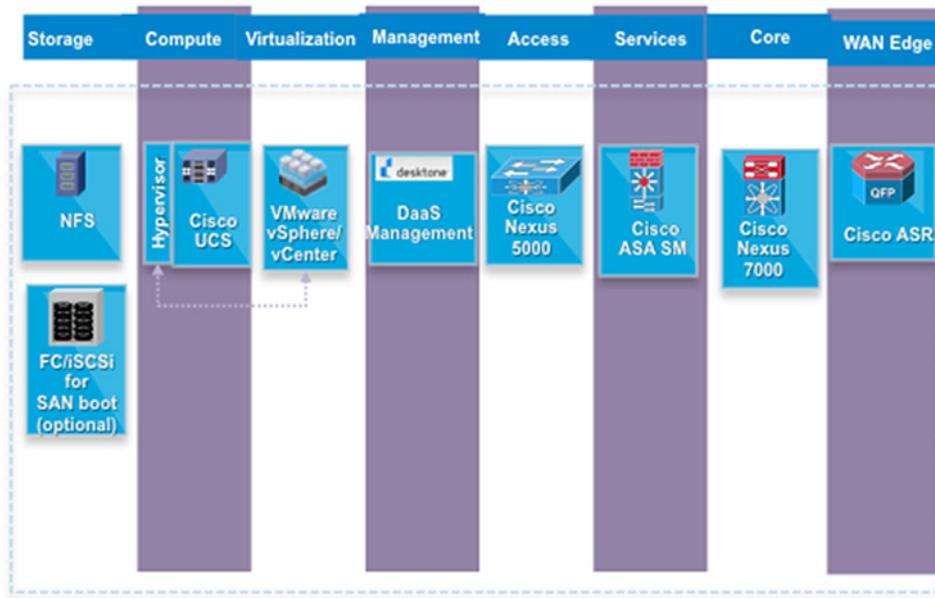
This document is intended for, but not limited to, solution architects and engineers involved in planning and designing Cisco DaaS Solution with Desktone. This document assumes the reader has an understanding of the following:

- Cisco Virtualized Multiservice Data Center
- Cisco Unified Computing System
- Cisco Nexus Switches
- Cisco Catalyst Switches and Service Modules
- Cisco Aggregation Services Router
- Cisco Adaptive Security Appliance
- MPLS VPN
- VMware vSphere
- Desktone
- NFS Storage
- Microsoft Windows 7 Operating System

2 Cisco DaaS Solution with Desktop Overview

This section describes the capabilities of the different products and solutions that are part of the Cisco DaaS Solution with Desktop reference architecture and how these elements are integrated to deliver a scalable, flexible and secure multi-tenant solution for service providers.

Figure 1 System Overview



The building blocks of the Cisco DaaS Solution with Desktop are:

- Virtualized Multiservice Data Center
- Wide Area Network
- Security
- Multi-Tenant Connection Broker

Virtualized Multiservice Data Center

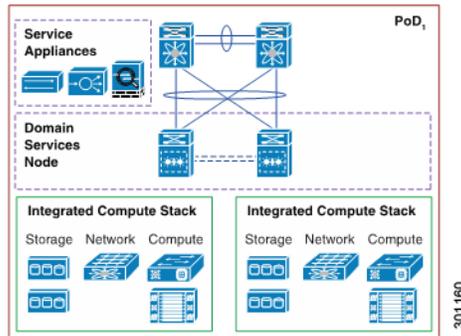
The Cisco Virtualized Multiservice Data Center (VMDC) Version 2.2 is the platform used for Cisco DaaS Solution with Desktop. The platform is ideal for service providers as it is designed for cloud based offerings and provides a multi-tenant infrastructure on which any service can be delivered. The modular architecture is highly scalable and adapts easily to the changing business needs of the service provider. The design has been validated by end-to-end system level testing and offered as part of Cisco Validated Design Program.

The first building block in the Cisco VMDC architecture is an integrated compute stack (ICS) based on Cisco Unified Computing System (UCS).

Compact pod design is the one chosen for Cisco DaaS Solution with Desktop. It uses centralized service node architecture on a collapsed aggregation/core with top-of-rack access.

Multiple ICS blocks are connected through the network infrastructure to form a pod. New pods are built and added to the data center for additional capacity, thereby making the entire solution scalable.

Figure 2 VMDC Pod



VMDC Network is organized into core, aggregation and access layers similar to a campus network design.

The primary function of the core is to provide high performance Layer 3 switching for IP traffic.

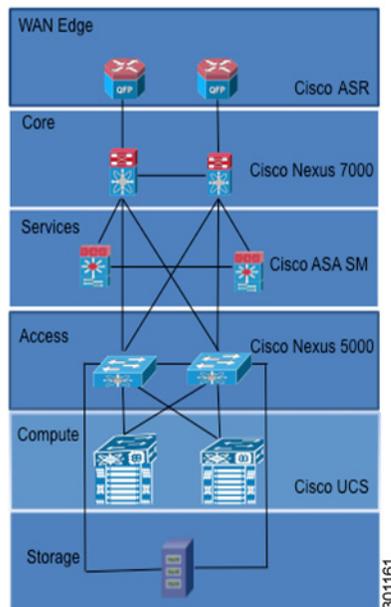
The aggregation layer provides connectivity for the access layer switches in the server farm and aggregates them into a smaller number of interfaces to be connected into the core layer. Services are introduced into data flows at the aggregation layer.

With Cisco DaaS Solution with Deskstone using the compact pod design, the core and aggregation layers are collapsed into one.

The access layer provides connectivity for server farm end nodes. Quality of Service (QoS) markings are done on the display protocol traffic at the access layer of the VMDC network and prioritized in the data center network.

The WAN Edge of the service provider’s data center provides network traffic isolation between different customers as they come into the data center from the WAN.

Figure 3 Cisco VMDC Network Topology



Wide Area Network

The WAN is the primary method of delivery of virtual desktops to customers. A VPN connection extends the customer network into the data center of the service provider enabling seamless access of virtual desktop for users in the customer premises.

MPLS VPN is one of the VPN technologies supported in Cisco DaaS Solution on Deskstone. With MPLS VPN, the traffic separation is built directly into the network without tunneling or encryption. Unique VPN route forwarding (VRFs) at the Provider Edge routers separates the traffic for every customer and allows overlapping IP addresses to exist. A customer's visibility is restricted to their network traffic only.

Based on the Service Level Agreement between the customer and WAN provider, the customer traffic of interest that are marked with appropriate Quality of Service markings is prioritized over the WAN, delivering an enhanced user experience.

Security

Cisco ASA Service Module (SM) is deployed on the Data Service Node, Catalyst 6500 to provide firewall functionality between the customer physical network and the virtual environment hosted by the service provider. The Cisco ASA SM supports multiple contexts, each of which is assigned to a tenant. The context will give every tenant a dedicated firewall to configure specific security policies that they desire.

Multi-Tenant Connection Broker

The Deskstone platform supports connections to virtual desktops over multiple display protocols like Microsoft Remote Desktop Protocol (RDP), VMware PC over IP (PCoIP) from different customers. PCoIP was the one used in the study. The platform authenticates and connects users from different customers to the right desktops in the data center. The IT administrator of customers gets to view only their own desktops. The service provider is the only one who has the visibility on how the resources are used by different customers in the data center.

Cisco DaaS Solution with Deskstone Components

Table 1 *Cisco DaaS Solution on Deskstone Components*

Function	Product
Compute	Cisco UCS 2.1(1a)
Hypervisor	VMware vSphere 5.1
Storage	NFS Storage iSCSI/FC for SAN boot (optional)
Access Layer	Cisco Nexus 5548UP 5.2(1) N1 (2a)
Core/Aggregation Layer	Cisco Nexus 7009 6.1 (2)
WAN Edge	Cisco ASR 9000 4.2.0
Data Center Service Node	Cisco Catalyst 6506 12.2(33) SXJ1

Function	Product
Compute	Cisco UCS 2.1(1a)
Network Services	Cisco ASA SM 8.5(1)
Connection Broker	Deskton 5.3.2

3 Design Considerations

Cisco DaaS Solution with Desktopone has been designed with the following objectives in mind.

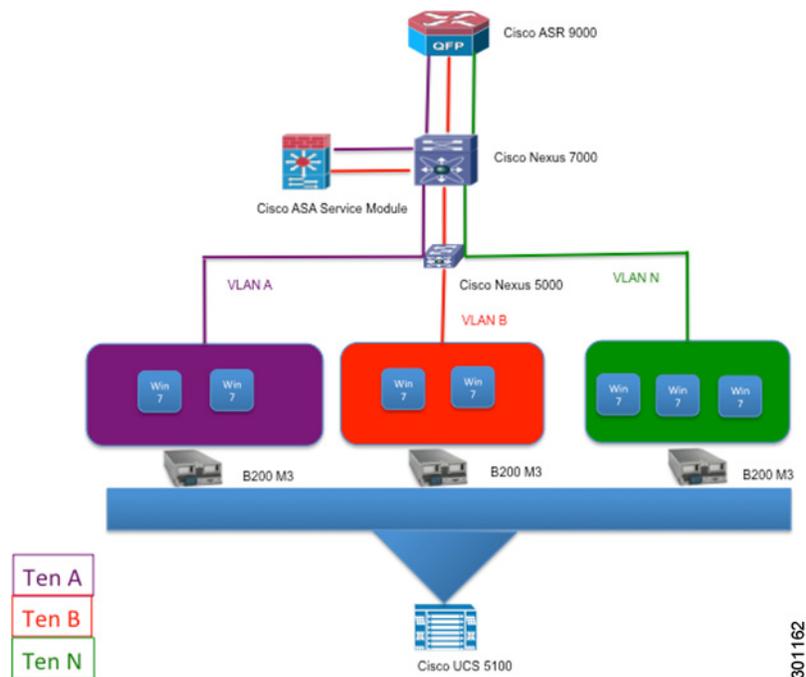
- Tenant Separation
- High Availability

Tenant Separation

Desktops are deployed by enterprises in-house on dedicated infrastructure. To consume desktops on shared infrastructure, customers must be certain that their corporate information is secure and isolated end-to-end. Secure tenant separation is done at the following levels:

- Compute
- Network
- Network Services
- Storage
- Desktop

Figure 4 *Tenant Separation*



Compute

A separate blade server is assigned for the Desktopone Platform management software. This host runs Desktopone service provider appliances and tenant appliances for every customer.

Blade servers are dedicated as a whole to every tenant. This is done to comply with Microsoft licensing agreement. Microsoft Virtual Desktop Access license prohibits running Windows 7 virtual desktops for different customers on the same host. This makes sure that there is complete separation of compute between different customers. Dedicated hosts are maintained even in times of hardware crashes when virtual desktops are moved between hosts.

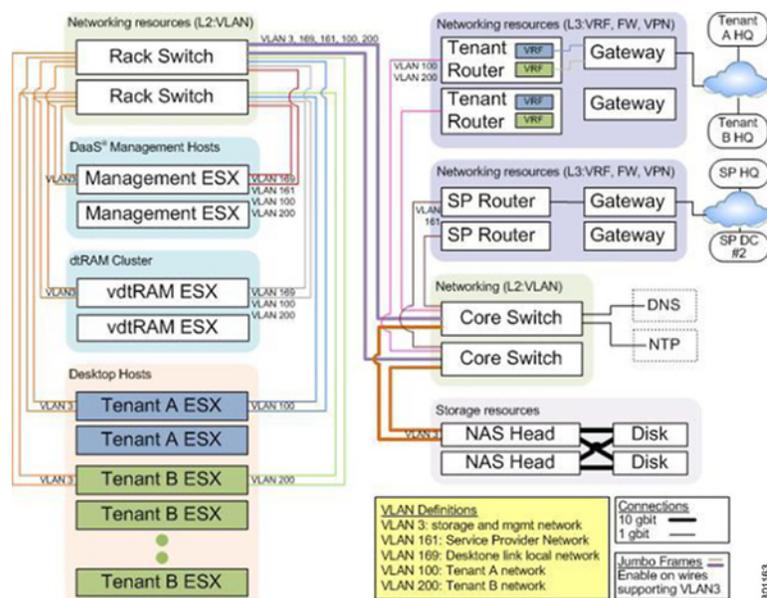
Network

A logically separate network infrastructure is created for every customer. The separation is done at both Layer 2 and Layer 3 using VLAN and VRF respectively.

Layer 3 VPN provides separate links for customers to connect into the service provider data center. VRF at the WAN edge of the data center provides path isolation of customer traffic at Layer 3. A separate routing table is used to make forwarding decisions for every customer's traffic. A separate routing table helps ensure support for overlapping IP address range of customers on a shared network.

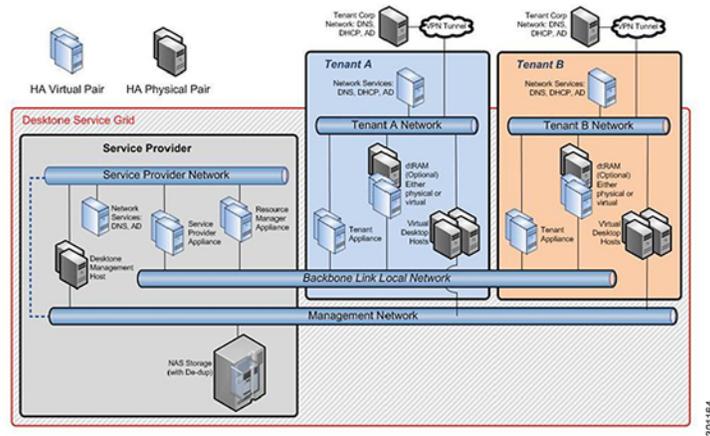
Separation at Layer 2 is done using VLANs. A minimum of one VLAN is assigned to each customer. All virtual desktops and virtual machines that belong to a customer reside on the customer's assigned VLAN. The VLAN ensures that customer's traffic is logically isolated from one another as it traverses different devices in the service provider data center.

Figure 5 Network Separation in Deskstone view



The Deskstone management appliances are located in different networks depending on the application type. The Service Provider Appliance and Resource Manager Appliances reside on the service provider network. The Tenant Appliances reside in the tenant network and cannot be accessed from any other network. There is a backbone link local network which is used for different Deskstone Appliances to talk to each other. The management network is used for the blade servers and storage appliances.

Figure 6 Desktope Service Grid



Network Services

Cisco ASA Service Module deployed on the Data Service Node, Catalyst 6500 is partitioned into multiple virtual firewalls, known as security contexts. The customer's security policies are configured and stored in the tenant configuration file.

Storage

Storage is managed by the VMware vCenter for Desktope. All storage options supported by vCenter are supported storage options for virtual desktops. If NFS is chosen as a storage option, file shares are created for every customer and also one for the Desktope platform. If FC is chosen, one or more Logical Unit Number (LUN) needs to be created for each customer and one for the Desktope platform.

The Desktope platform needs a file share/LUN to host the template used in creation of new tenant appliances on the Desktope platform. This file share/LUN is mounted on the hosts assigned to the Desktope platform.

The virtual desktops for each customer are stored in their corresponding file share/LUN. The file share/LUN of each customer is mounted on the hosts reserved for the customer as VMware datastore. No other host must have access to the datastore of the customer.

Desktop

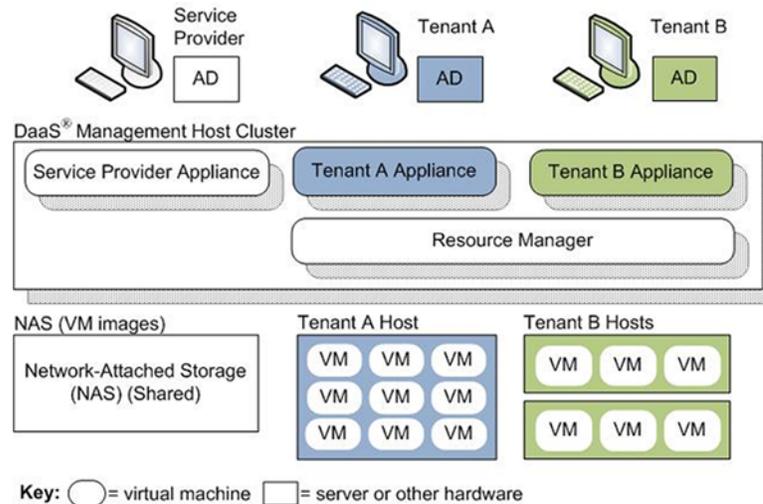
The Desktope Platform manages all incoming connections to the virtual desktops. Based on the user and to which customer the user belongs to, the Desktope Platform connects the user to the right desktop over the preferred display protocol.

The service provider appliance is administered using the Service Center web based Graphical User Interface. The service center portal gives an overview of all the customers sharing the common data center. This portal is also used to manage the system resources for different customers. The service provider appliance should be installed first as it is the foundation for installing rest of the Desktope appliances.

A separate enterprise portal called Enterprise Center is powered by the customer's tenant appliance. The Enterprise Center is only accessible from the corresponding customer's network. This portal restricts visibility to only the desktops the customer owns.

There are two Resource Manager Appliances in a high available environment to manage the resources used for the tenants a whole.

Figure 7 **Desktop Components**



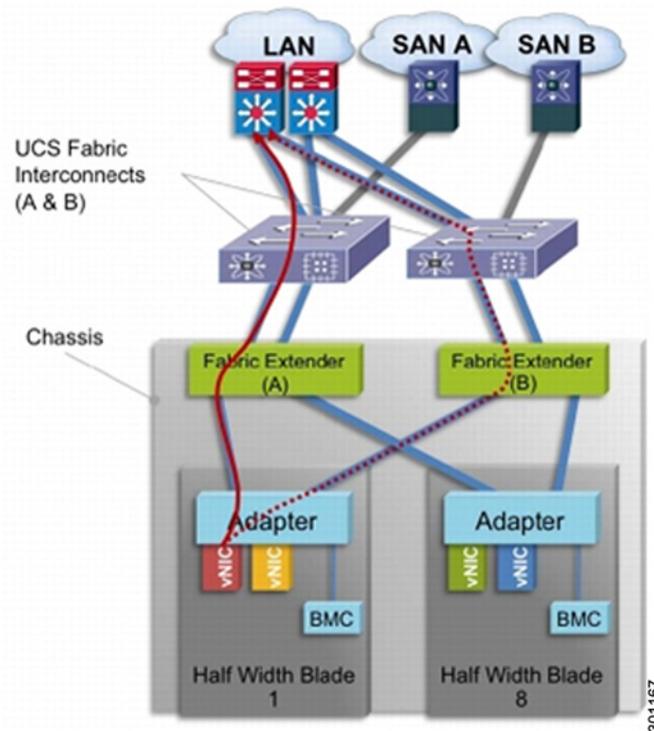
High Availability

High Availability (HA) is a mandatory requirement of any hosted or cloud-based service offerings. For a service to be adopted widely by businesses, service providers need to provide Service Level Assurance (SLA) to their business customers. To meet such SLA, the service provider needs to base their offering on a highly available infrastructure.

VMDC was chosen as platform of choice for Cisco DaaS Solution on Desktop because it can provide redundancy at different levels like Compute, Network and Storage. Redundancy is provided end-to-end all the way from the collapsed core/aggregation, access, compute, and storage layers.

Compute

Figure 8 Redundant paths for Virtual NICs



Cisco Unified Computing System is a fully redundant system. The network adapter is connected to both the fabric planes in active-active node utilizing the full bandwidth. The Fabric Interconnect Uplinks connects to the access layer switches to provide redundant paths for all the virtual machines hosted on the blade server. The Fabric Interconnects should always be deployed in pairs for redundancy and high availability.

As per VMDC recommendation, configure the Cisco UCS Fabric Interconnect uplinks in End Host Mode. In End Host mode, you get the following:

- Uplinks are fully utilized
- Uplinks can be connected to multiple access switches providing redundancy
- Spanning Tree Protocol is not required

Cisco Unified Computing System provides different system class of service to implement quality of service including:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames.



Note Time sensitive applications have to adhere to a strict QoS for optimal performance.

Systems Class is the global operation where entire system interfaces have defined QoS rules.

- By default system has Best Effort Class and FCoE Class.



Note Best effort is equivalent in MQC terminology as “match any”

- FCoE is special Class define for FCoE traffic. In MQC terminology “match cos 3”
- System class allowed with 4 more users define class with following configurable rules.
 - CoS to Class Map
 - Weight: Bandwidth
 - Per class MTU
 - Property of Class (Drop v/s no drop)
- Max MTU per Class allowed is 9216.
- Through Cisco Unified Computing System, we can map one CoS value to particular class.
- Apart from FCoE class there can be only one more class can be configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally system will calculate the bandwidth based on following equation (there will be rounding off the number).

$$\text{➤ \% b/w shared of given Class} = \frac{(\text{Weight of the given priority} * 100)}{\text{Sum of weights of all priority}}$$

Cisco Unified Computing System defines user class names as follows:

- Platinum
- Gold
- Silver
- Bronze

Table 2 *Map between Cisco Unified Computing System and the NXOS*

Cisco UCS Names	NXOS Names
Best effort	Class-default
FC	Class-fc
Platinum	Class-Platinum
Gold	Class-Gold
Silver	Class-Silver
Bronze	Class-Bronze

Table 3 *Class to CoS Map by default in Cisco Unified Computing System*

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Fc	3
Platinum	5

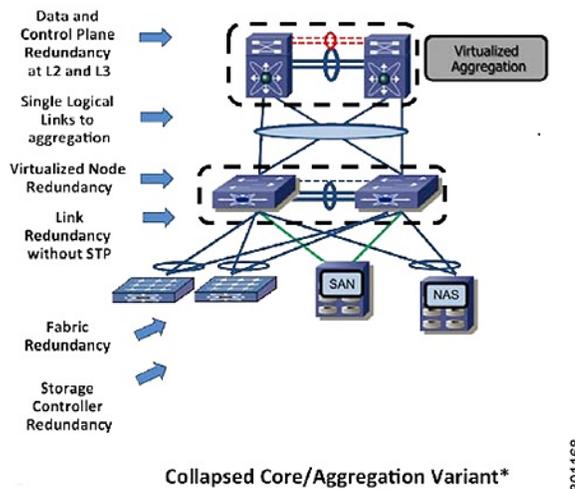
Cisco UCS Class Names	Cisco UCS Default Class Value
Gold	4
Silver	2
Bronze	1

Table 4 Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
Fc	5

Network

Figure 9 Redundant Network Layers



Virtual Port Channel (vPC) helps ensure high availability at Layer 2 in the data center. All available paths are utilized to deliver maximum bandwidth to the hosts. Two upstream Nexus switches act as one logical node delivering both hardware redundancy and load balancing.

Host StandBy Router Protocol (HSRP) provides hardware redundancy and high availability at Layer 3 by having two routers act as a single virtual router with shared IP address and MAC address. The shared address enables hosts to always have an active router to forward packets.

Network Services

Cisco ASA Service Module is deployed on the 2 Catalyst 6500 switches for redundancy. It supports high-speed failover between modules in separate chassis. Both Active-Active and Active-Standby failover is supported.

Storage

As per VMDC recommendation, it is required to have hardware redundancy for storage devices.

Desktop

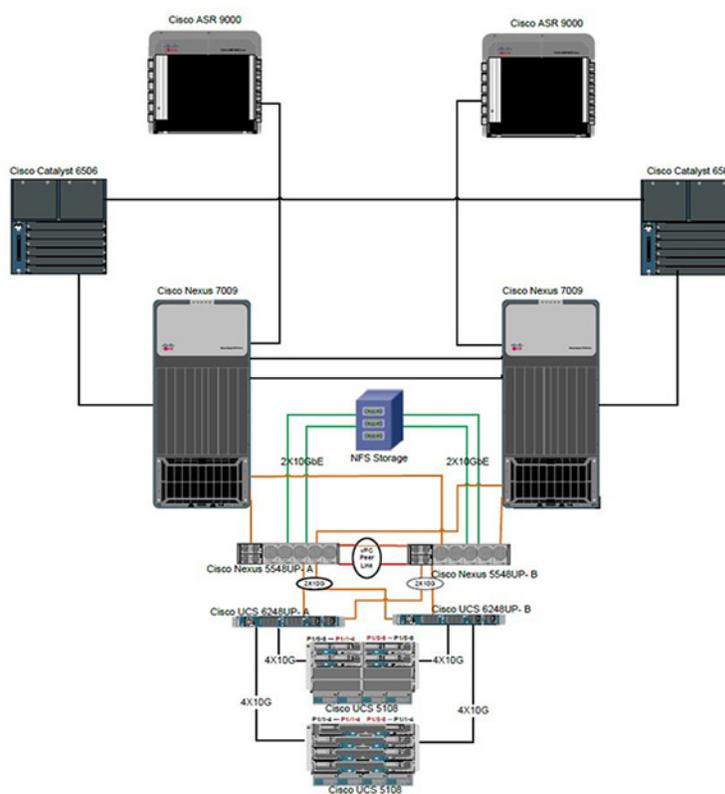
The Desktone platform is installed on two separate blade servers. On each blade server, one instance of Service Provider Appliance, Tenant Appliances and Resource Manager Appliances are installed. One instance is the primary and the other is the secondary. Both primary and secondary instances of the appliances work in active mode. So, either of the appliances can crash with no service interruption. The same holds true if the whole blade server hosting the Desktone appliances were to crash.

When desktop hosts crash, the desktops can be moved over to a new host and the desktops would be available for users. It is recommended to have at least one spare host per tenant to handle blade crashes. This host needs to be on the VLAN assigned to the tenant. It then needs to be assigned to the Tenant Resource Manager to be available for hosting desktops. If and when the crash happens, the desktops need to be migrated to the new host.

4 Solution Validation

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution.

Figure 10 Cisco DaaS Solution on Desktopone Topology



301189

Configuring the Cisco Unified Computing System

To configure the Cisco Unified Computing System, perform the following steps:

- Step 1** Bring up the Cisco 6248UP Fabric Interconnect (FI) and from a serial console connection set the IP address, gateway, and the hostname of the primary FI.
- Step 2** Bring up the second fabric interconnect after connecting the dual cables between them. The second FI automatically recognizes the primary and asks if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname.

**Note**

When Step 2 is completed all access to the FI may be performed remotely. You will also configure the virtual IP address to connect to the FI; you need a total of three IP address to bring it online. You can also wire up the chassis to the FI, using 1, 2, 4 or 8 links per IO Module, depending on your application bandwidth requirement. In this study, four links were connected to each module.

Step 3 Connect using your favorite browser to the Virtual IP address and launch the Cisco UCS Manager. The Java based Cisco UCS Manager enables you to do everything that you could do from the CLI. The GUI methodology is highlighted here.

Step 4 Check the firmware on the system and see if it is current. Visit [http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.0\(4d\)&releind=AVAILABLE&rellifecycle=&reltype=latest](http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.0(4d)&releind=AVAILABLE&rellifecycle=&reltype=latest) to download the most current Cisco UCS Infrastructure and Cisco UCS Manager software. Use the UCS Manager Equipment tab in the left pane, then the Firmware Management tab in the right pane and Packages sub-tab to view the packages on the system. Use the Download Tasks tab to download needed software to the FI.

**Note**

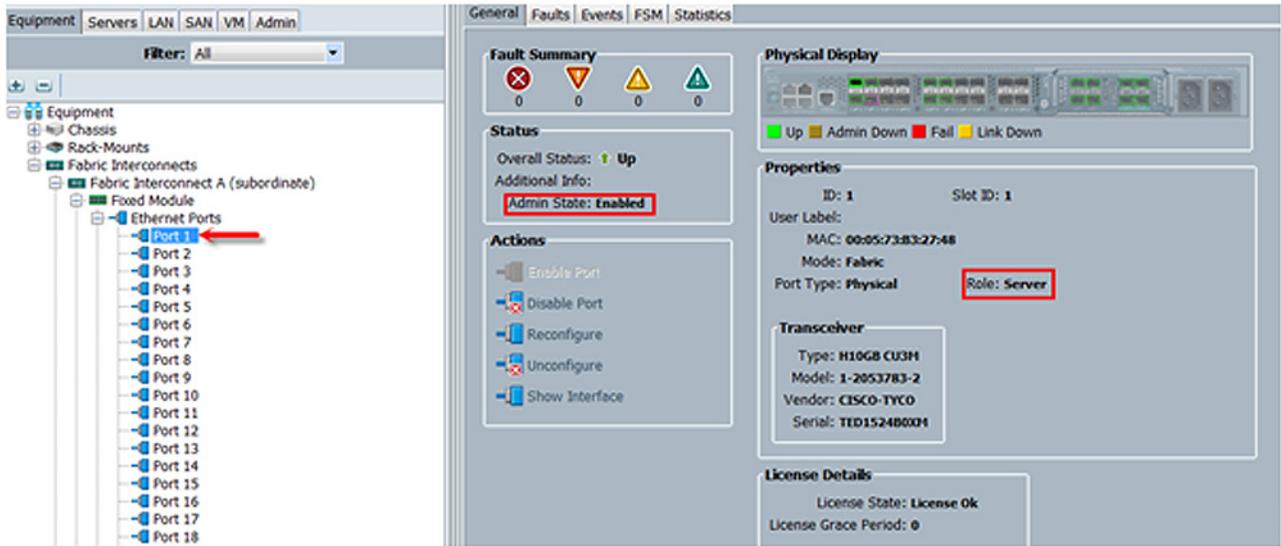
The firmware release used in this paper is 2.1(1a).

The screenshot shows the Cisco UCS Manager interface. The left pane displays a tree view with 'Equipment' selected. The right pane shows the 'Firmware Management' tab with the 'Packages' sub-tab active. A table lists installed packages with columns for Name, Type, State, Vendor, Version, and Deleted on Fabric. Three packages are highlighted with red boxes:

Name	Type	State	Vendor	Version	Deleted on Fabric
ucsi9-bundle-b-series.2.1.1a.B.bin	B Series Bundle	Active		2.1(1a)B	
ucsi9-bundle-c-series.2.1.1a.C.bin	C Series Bundle	Active		2.1(1a)C	
ucsi9-bundle-infra.2.1.1a.A.bin	Infrastructure Bundle	Active		2.1(1a)A	

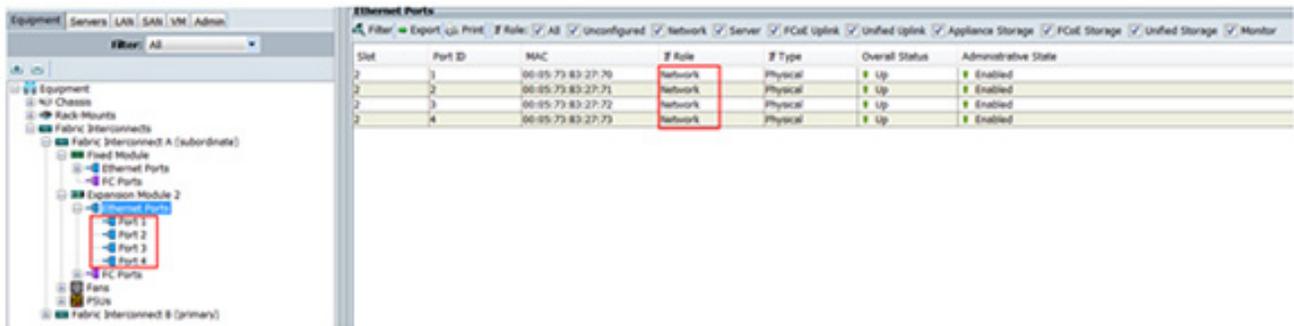
301170

Step 5 Configure and enable the server ports on the FI. These are the ports that will connect the chassis to the FIs.



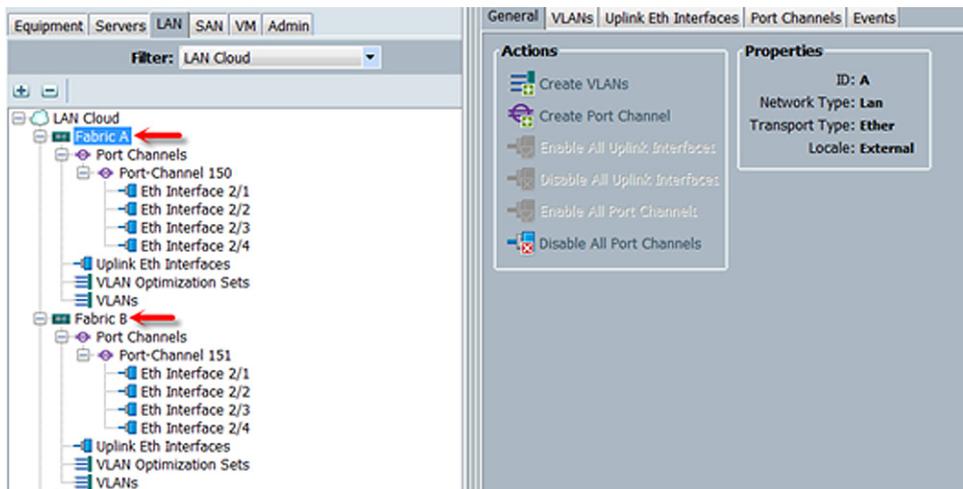
301171

Step 6 Configure and enable uplink Ethernet ports



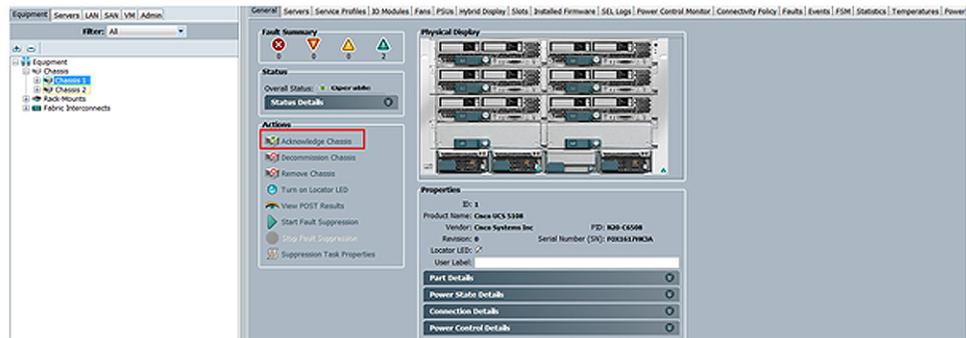
301172

Step 7 On the LAN tab in the Navigator pane, configure the required Port Channels and Uplink Interfaces on both Fabric Interconnects



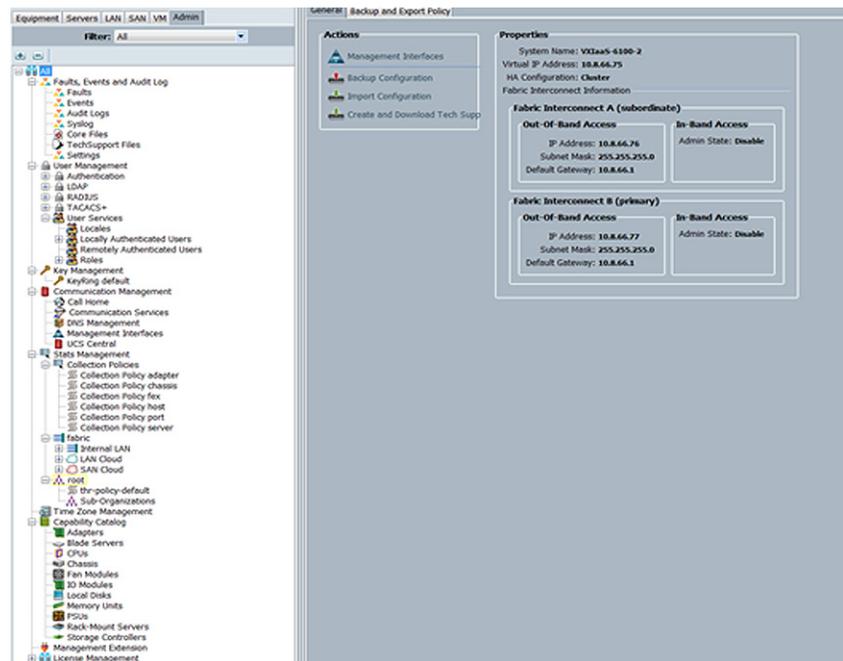
301173

- Step 8** Expand the Chassis node in the left pane and click on each chassis in the left pane.
- Step 9** Click Acknowledge Chassis in the right pane to bring the chassis online and enable blade discovery.



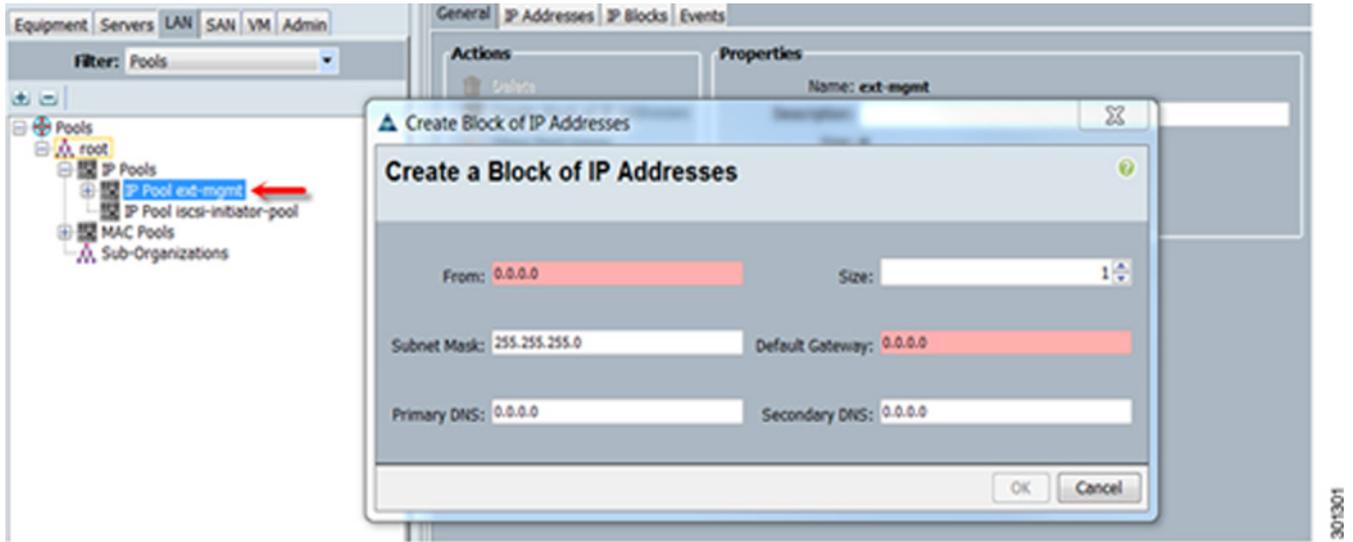
301174

- Step 10** From the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and licensing. Time zone Management (including NTP time source(s)) and uploading your license files are critical steps in the process.

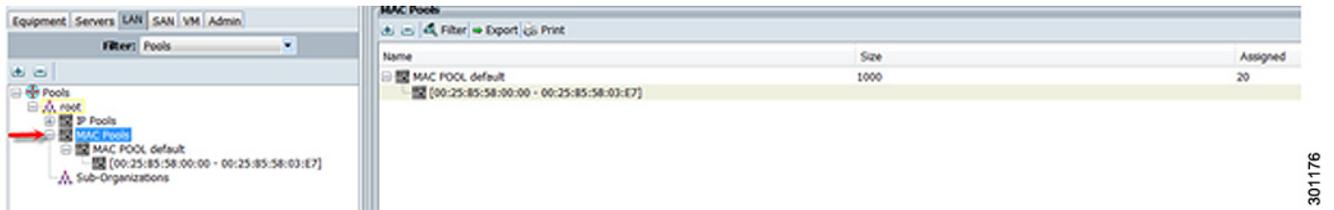


301175

Step 11 On the LAN tab, expand Pools, IP Pools. Right click on IP Pool ext-mgmt. Configuring your Management IP Pool (which provides IP based access to the KVM of each UCS Blade Server).



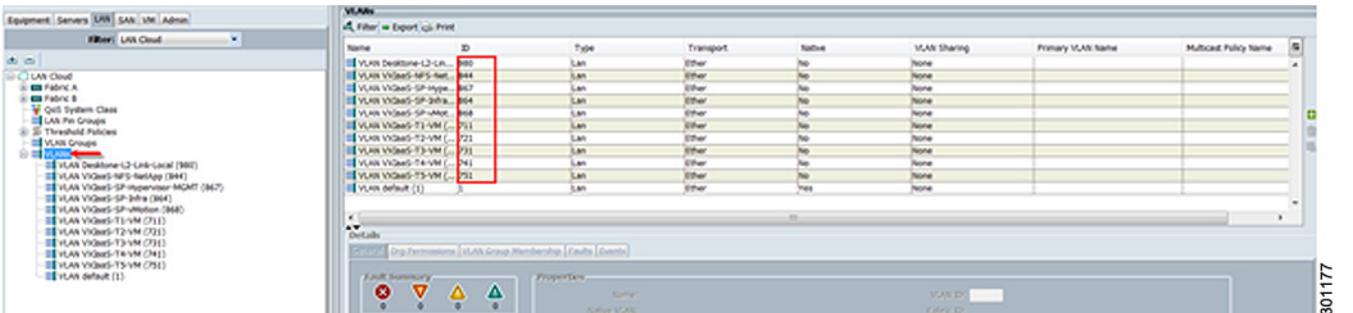
Step 12 From the LAN tab in the navigator, under the Pools node, create a MAC address pool of sufficient size for the environment.



Step 13 From the LAN tab in the navigator pane, configure the VLANs for the environment. You need one VLAN for Service Provider network, one for Link-Local Network and one for each tenant.



Note In this study, a separate VLAN is used for storage.



Step 14 Enable the different priorities, Platinum, Gold and Silver for QoS System Class

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input checked="" type="checkbox"/>	5	<input checked="" type="checkbox"/>	10	22	9000	<input type="checkbox"/>
Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	20	normal	<input type="checkbox"/>
Silver	<input checked="" type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	18	normal	<input type="checkbox"/>
Bronze	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	15	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	11	normal	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	14	fc	N/A

301178

Step 15 Add the QoS policy with its corresponding priority

Actions

- Delete
- Show Policy Usage

Properties

Name: **Platinum**

Egress

Priority: **Platinum**

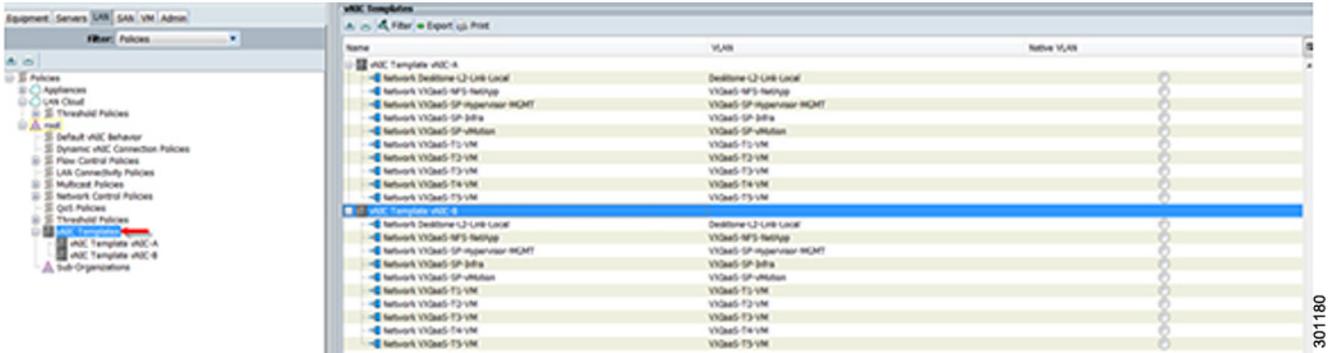
Burst(Bytes): 10240

Rate(kbps): line-rate

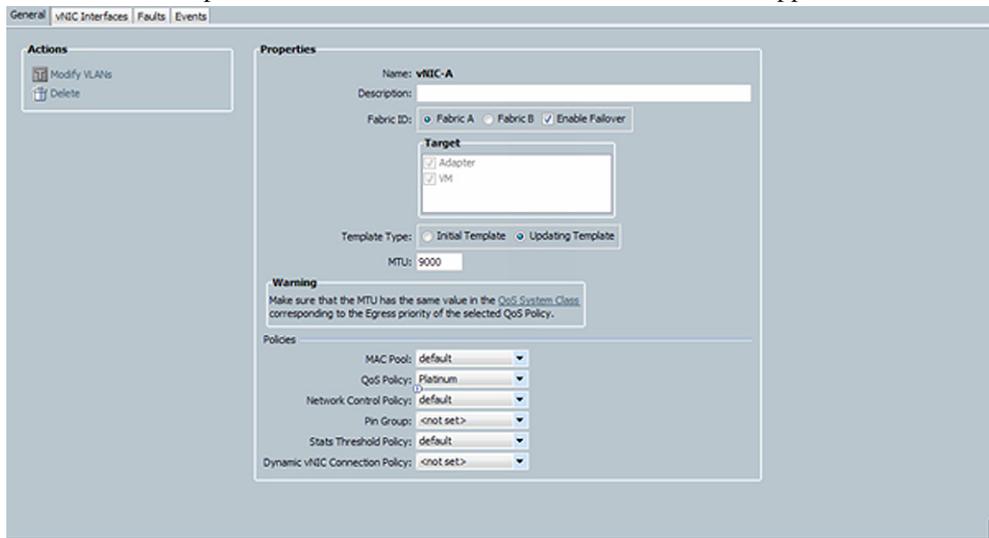
Host Control: None Full

301179

Step 16 From the LAN tab in the navigator pane, under the policies node configure the vNIC templates to be used in the Service Profiles.



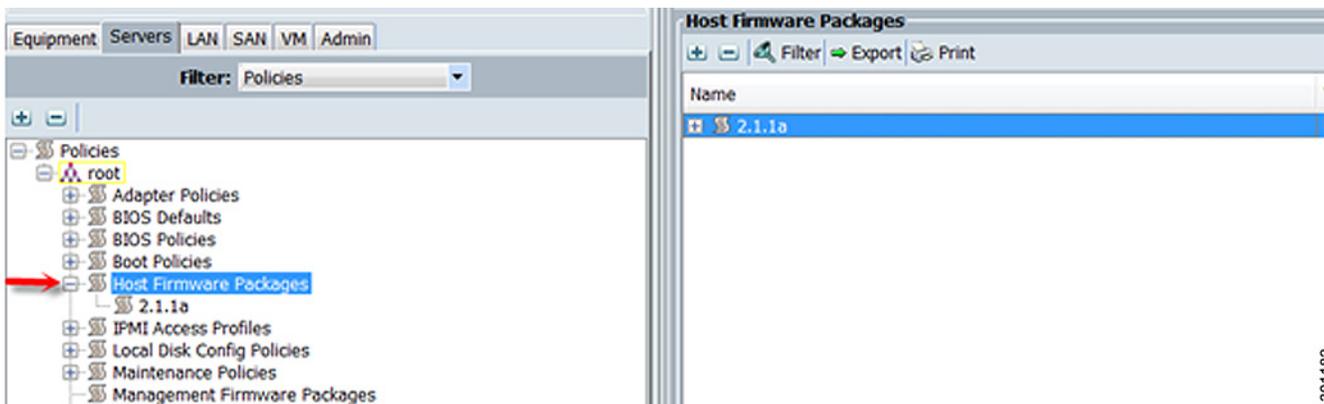
Step 17 Create vNIC templates for both fabrics and select the VLANs to be supported, MTU size and QoS policy.



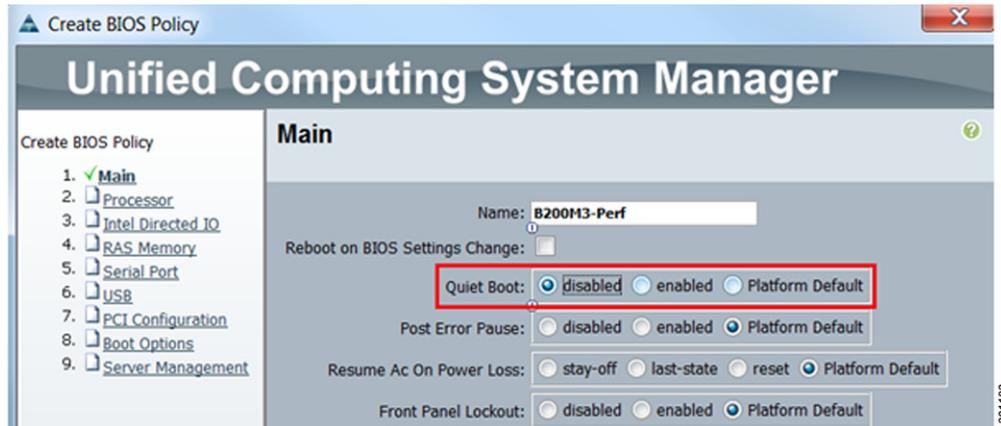
Step 18 New in Cisco UCS Manager 2.1(1a) is a method to set Host Firmware Package polices that can be set by package version across the Cisco UCS domain rather than by server model.



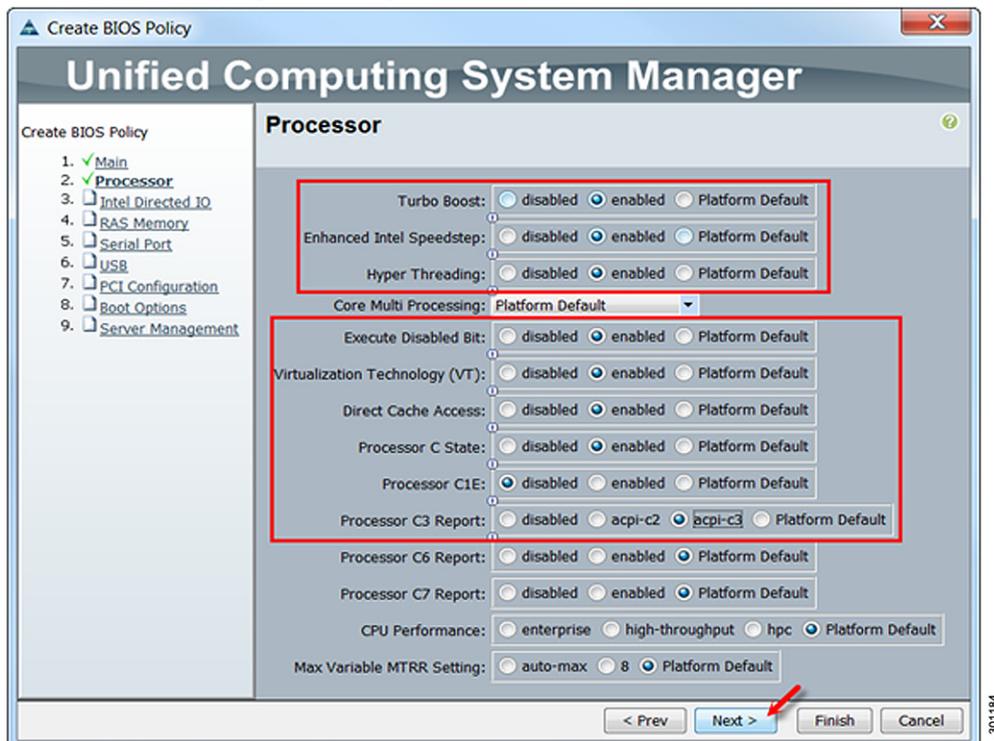
Note You can create specific packages for different models or for specific purposes.



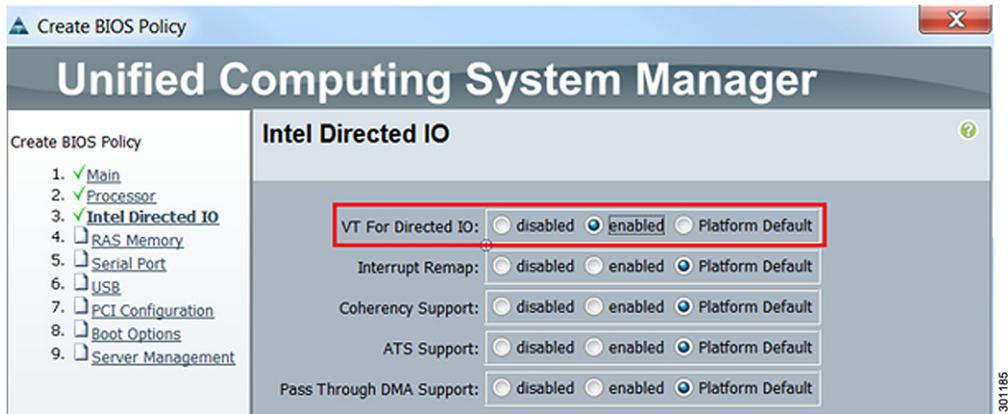
- Step 19** Create a BIOS policy under Servers -> Policies -> BIOS Policies for the Cisco UCS B200 M3 blade servers.



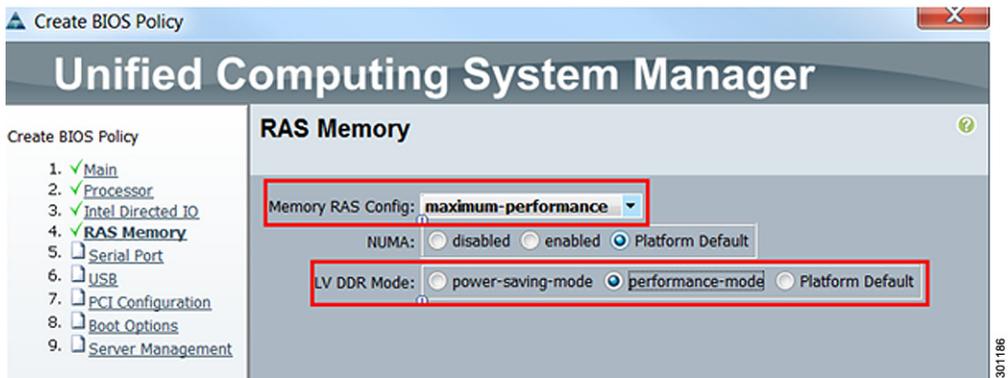
- Step 20** Choose the highlighted settings for Processor.



Step 21 Enable VT for Directed IO.



Step 22 Choose the following Memory settings:



Step 23 Keep the default settings for the other options and click Finish.

Step 24 Create a service profile template using the pools, templates, and policies configured above.

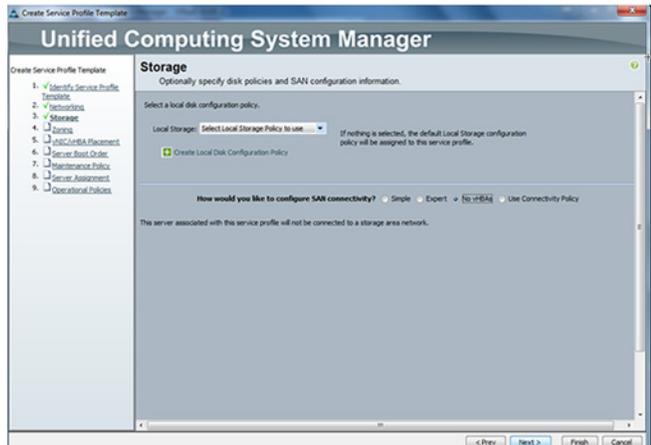


- Step 25** For the Identify Service Profile Template, enter a unique name, select the type Updating Template, and click Next.

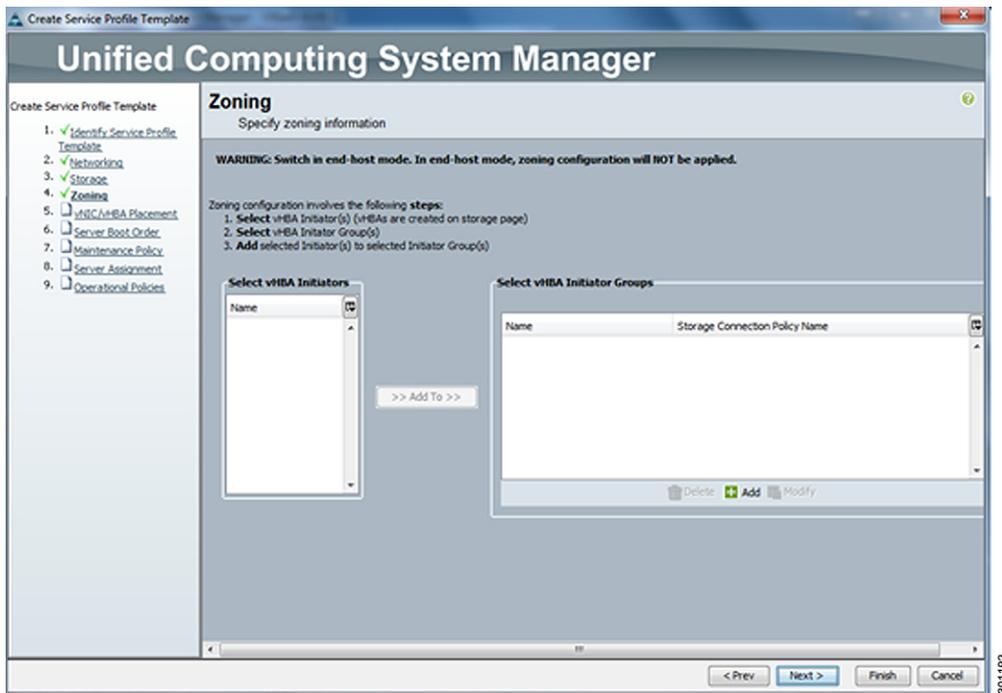
- Step 26** For the Networking option, select Expert and click Add in the adapters window.

- Step 27** Create vNIC and check Use vNIC Template.

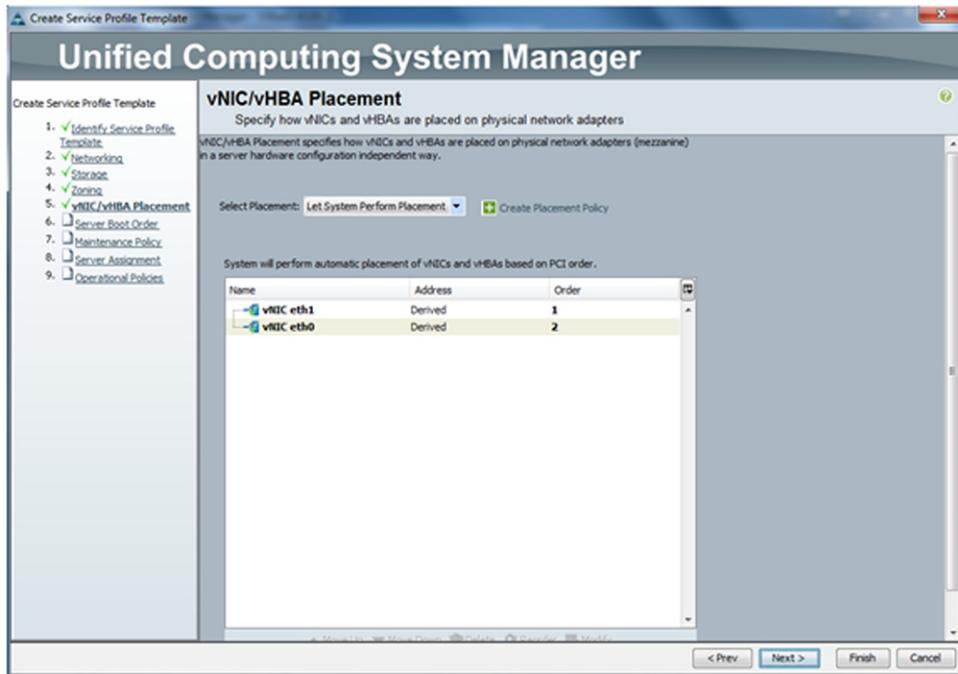
Step 28 Select No vHBAs option since SAN is not used in this study.



Step 29 Click Next since no Zoning is required.

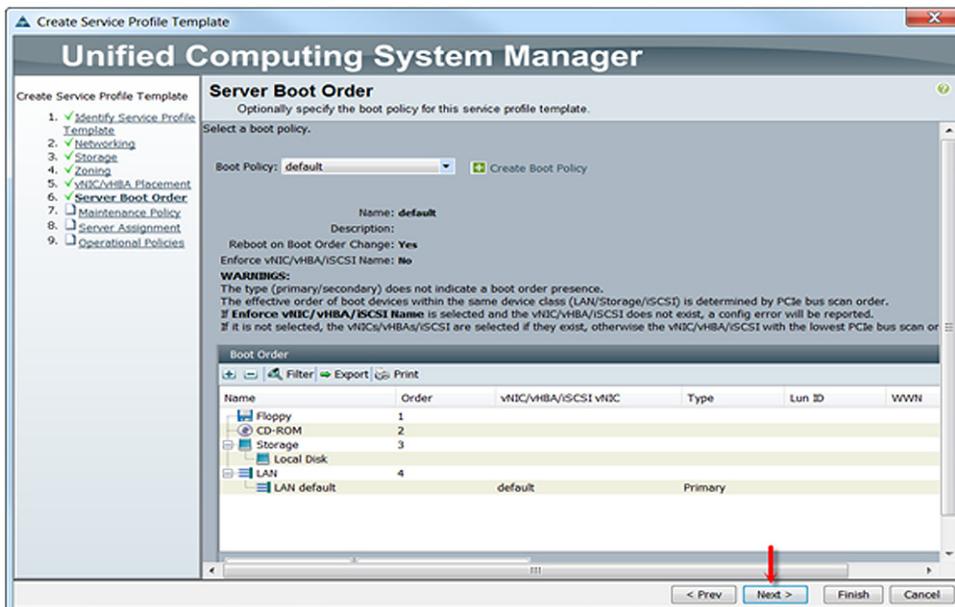


Step 30 Accept Default placement and click Next.



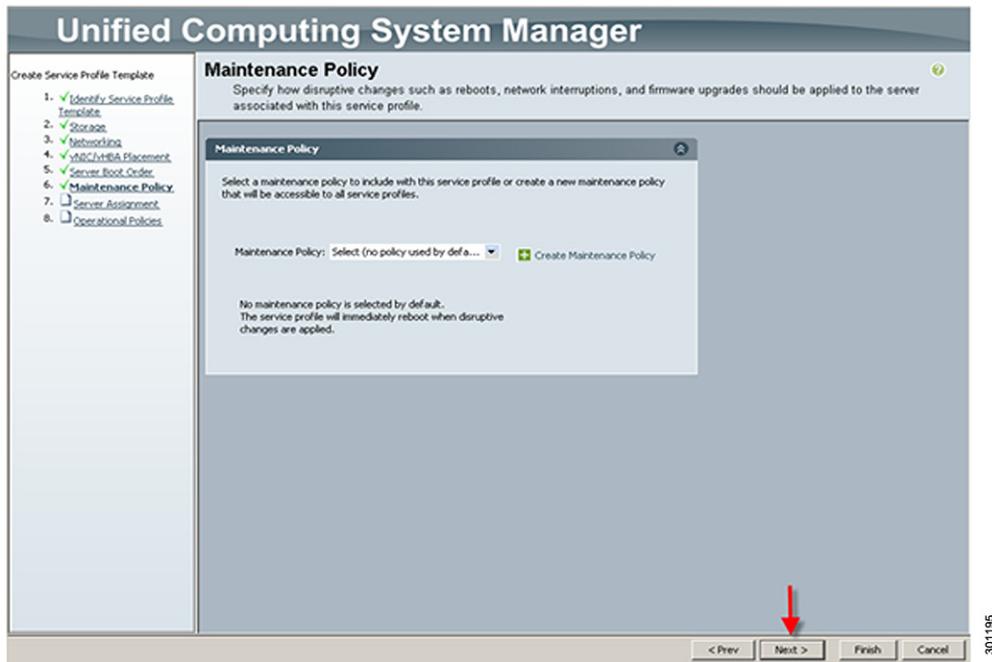
301193

Step 31 Select default boot policy.

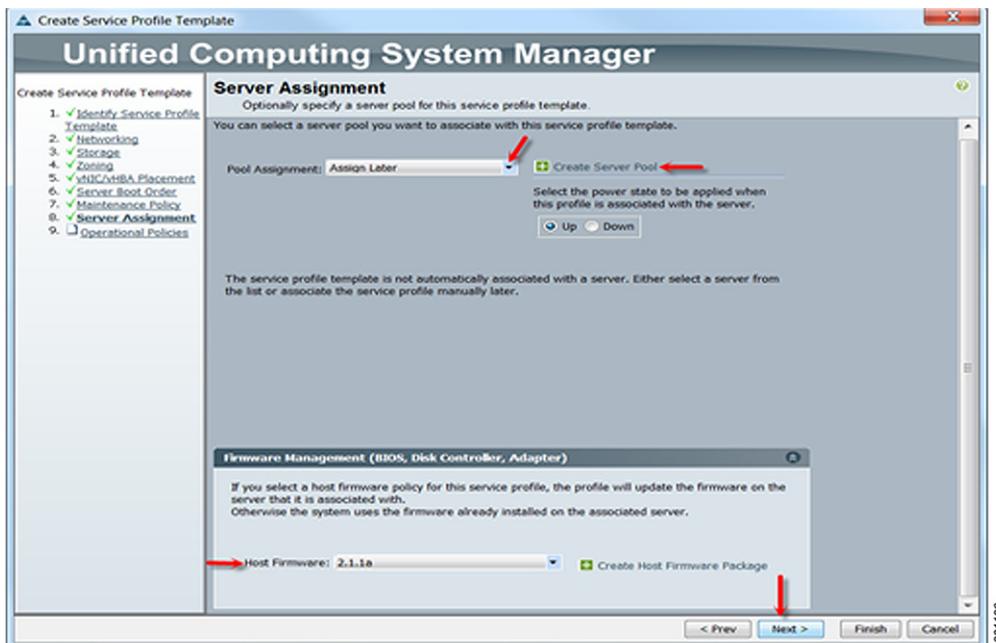


301194

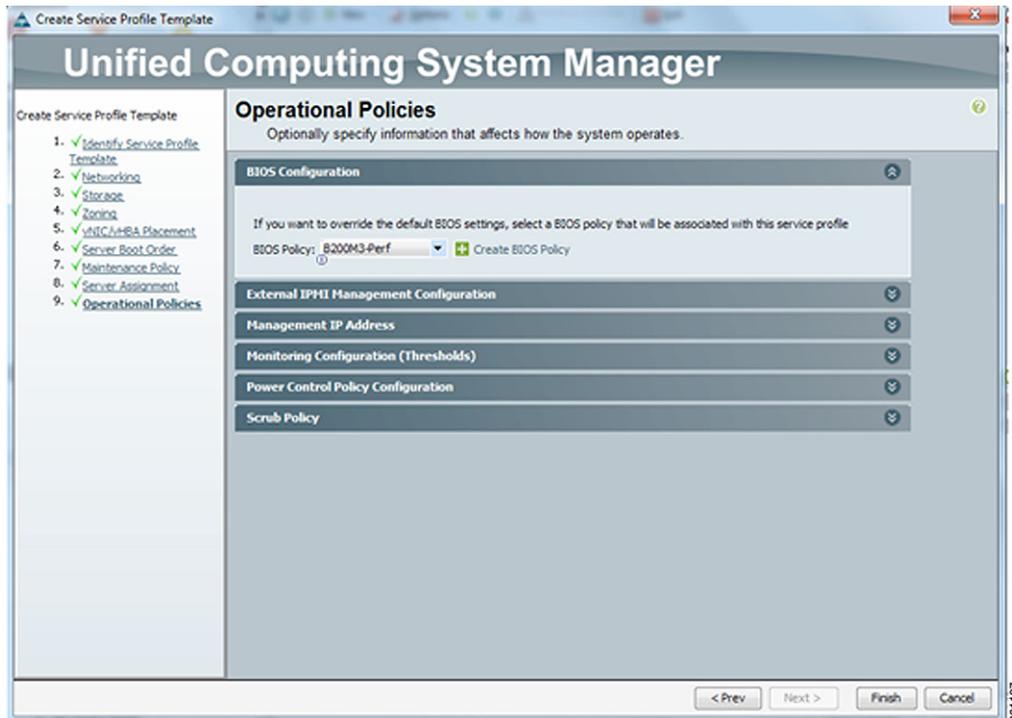
Step 32 A Maintenance policy was not used in this study. Click Next to continue.



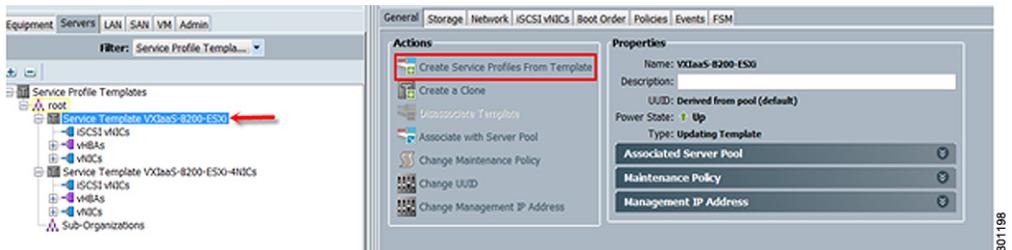
Step 33 For the Server Assignment option select default pool created or create a new server pool. Select Host Firmware package to be applied and click Next.



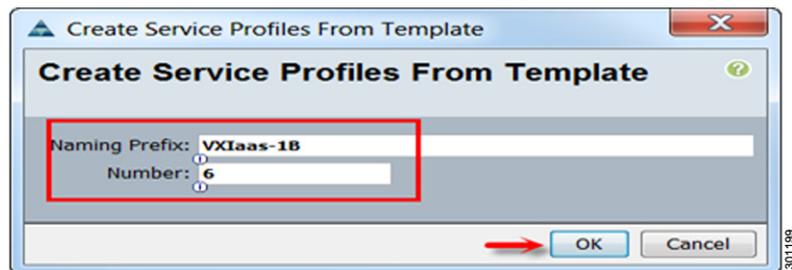
Step 34 Choose the BIOS policy created and click Finish to complete the Service Profile Template.

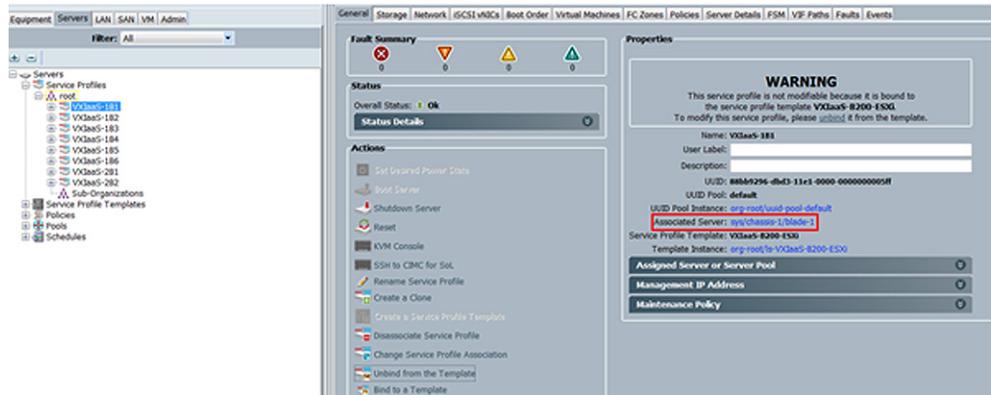


Step 35 Create a Service Profile from Template.



Step 36 Provide a Naming Prefix and the number of service profiles to be created from the template.



Step 37 Service profiles are assigned to the servers.

Installing and Configuring VMware vSphere

Install VMware ESXi

ESXi was installed from the Cisco UCS Manager (UCSM) KVM console using a ESXi 5.1 ISO image downloaded from the VMware site.

The IP address, hostname, and NTP server were configured using Direct Console ESXi Interface accessed from Cisco UCS Manager KVM console.

See the following VMware documentation for details about configuring network settings:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc_50/GUID-26F3BC88-DAD8-43E7-9EA0-160054954506.html.

Install and Configure VMware vCenter

A Cisco UCS B200 M3 blade server was used to host the virtual machines used for SP management including Microsoft Active Directory, DNS, DHCP server, etc., and also for hosting the vCenters used for Desktop environment.

To manage hypervisors and virtual machines on hosts reserved for Desktop Appliances, a dedicated vCenter server instance was installed on a Windows 2008R2 virtual machine.

An identical vCenter server was dedicated for managing the hypervisors and desktops for all tenants.

Table 5 vCenter Server Configuration

VMware vCenter Server			
OS:	Windows 2008 R2	Service Pack:	
CPU:	4vCPUs	RAM:	16GB
Disk:	80GB	Network:	1x10Gbps

To support vCenter instance, one Microsoft SQL Server 2008 R2 was created to host vCenter database.

It is recommended to utilize fault tolerance at the SQL Server level, refer to Microsoft documentation about configuring SQL Server clusters:

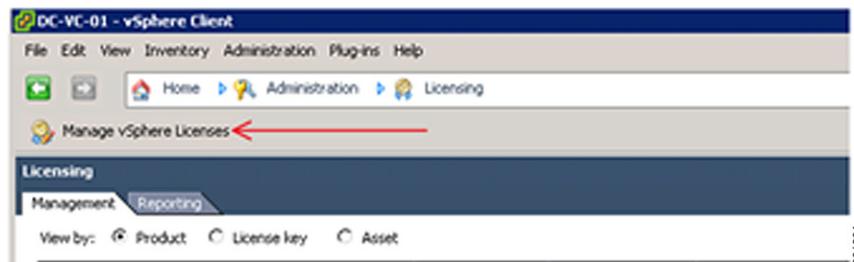
[http://msdn.microsoft.com/en-us/library/ms189134\(v=sql.105\).aspx](http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx)

To install and configure vCenter, perform the following steps:

-
- Step 1** Install the Microsoft® SQL Server® 2008 R2 Native Client for ODBC connections
(<http://www.microsoft.com/en-us/download/details.aspx?id=16978> look for Native Client for your architecture)
- Step 2** Create a System DSN (control panel, administrative tools, Data Sources ODBC) and connect to your vCenter-SQL server.
-  **Note** Make sure to use FQDNs throughout this installation.
-
- Step 3** Install vCenter server package and connect to the database.
- Step 4** Connect your vSphere client to vCenter and create a datacenter.
- Step 5** Create a self-signed certificate.
(http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514).

Install Licenses

-
- Step 1** Connect to vCenter using vSphere client.
- Step 2** Go to Home --> Administration --> Licensing and click Manage vSphere Licenses.



Step 3 Add License keys for vCenter and Hosts and click Next.

Manage vSphere Licenses

Add License Keys
If you have recently purchased, upgraded, combined, or split keys, add them to your inventory now.

Add License Keys
Assign Licenses
Remove License Key
Confirm Changes

Enter new vSphere license keys (one per line):

Enter optional label for new license keys:

Add License Keys

30/202

Step 4 Enter the license key and add an optional label.

Step 5 Repeat the above-mentioned procedure to create another instance of vCenter; one for Desktop appliances and another one for tenant desktops.

VMware vSphere Network Configuration for Desktop Appliances

The VLANs for service provider network is added in vSwitch of the vCenter reserved for Desktop appliances.

A separate VLAN is used for Desktop Platform management traffic. It is a non-routable subnet with a link-local address. A link-local address is an IP address used only for communications within a link (segment of a local network) or a point-to-point connection to which a host is connected. Routers do not forward packets with link-local addresses. The address block 169.254.1.0 through 169.254.254.255 is reserved for link-local addressing in Internet Protocol Version 4. You cannot choose addresses outside this range.



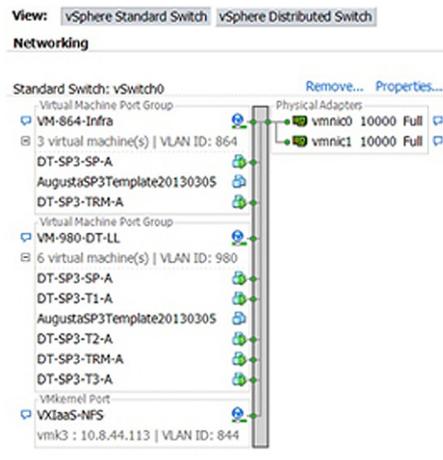
Note

If you have more than one data center, the link-local address space must be unique (non-overlapping) across data centers.

The configured VLANs must be the same across all management hosts.

- VLAN 864 for Service Provider Network
- VLAN 980 for Link Local
- VLAN 844 for Storage

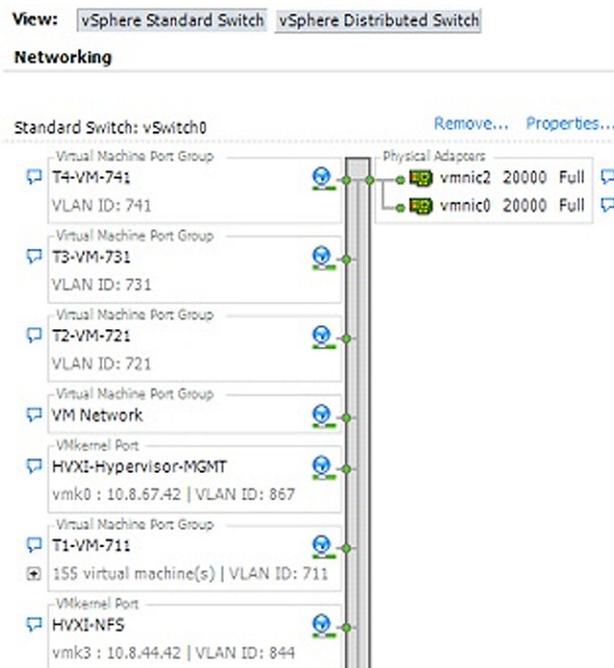
Figure 11 vSwitch for Deskstone Appliances



VMware vSphere Network Configuration for Deskstone Tenant Appliances

The tenant VLANs and the VLAN for storage are added to the vSwitch of the vCenter reserved for tenant desktops. Tenant 1 was used for this study. The VLAN assigned for Tenant 1 is 711.

Figure 12 vSwitch for Tenant hosts



Creating the Golden Image

Before defining a virtual machine as your gold template you need to create your template. A new OS installation is strongly recommended, which should be customized to VDI best practices.

For details about Windows 7 VDI best practices, see <http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf>

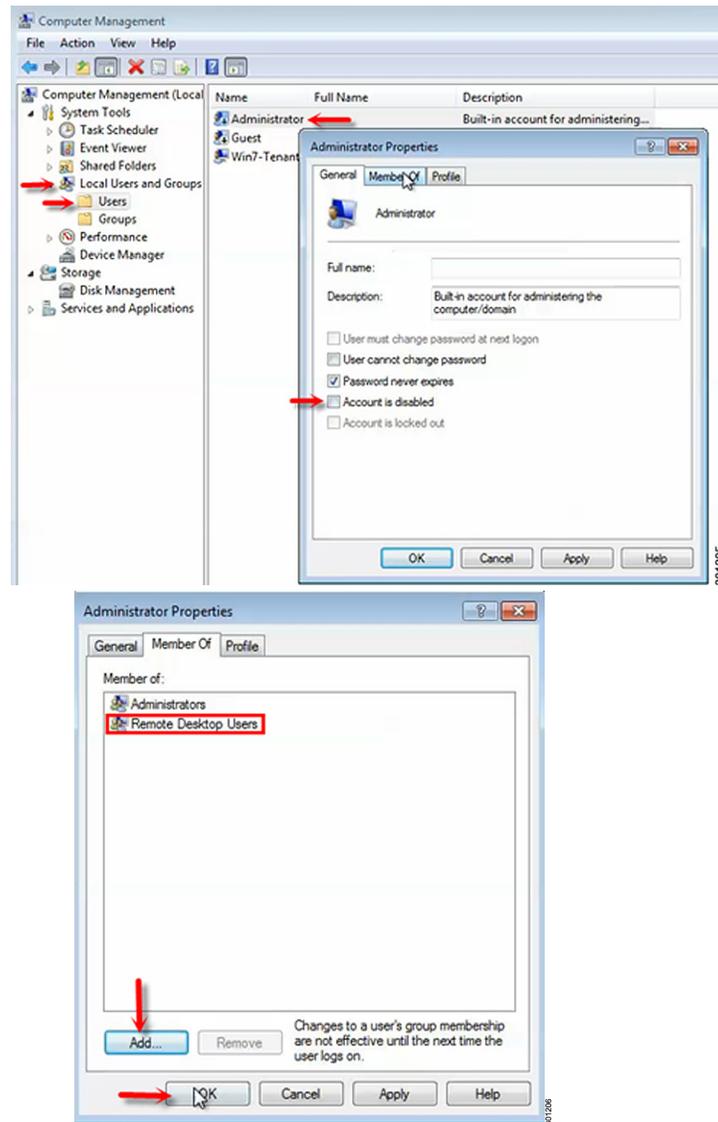
**Note**

32bit Windows 7 OS with 1 vCPU, 1.5GB RAM, 24GB HDD, and one 10GB vNIC were used for testing.

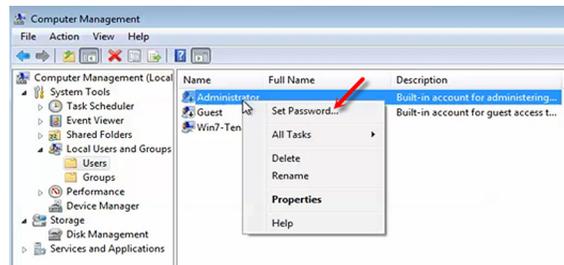
Steps to Customize the Golden Image Virtual Machine

- Step 1** Install VMware tools on Windows 7 VM.
- Step 2** Verify NIC settings. Confirm the adapter type is VMXNET3.
- Step 3** Set the power option to High Performance.
- Step 4** Enable Administrator account and RDP access.
- Step 5** Right-click Computer and select Manage.
- Step 6** Select Local Users and Groups --> Users.
- Step 7** Right click the Administrator user and select Properties.
- Step 8** From the General tab, uncheck Account is disabled.

Step 9 From the Member Of tab, confirm Administrator is a member of “Remote Desktop Users”



Step 10 Set the Administrator password.



Step 11 Apply PCoIP GPO (.adm file) and configure protocol settings.

Step 12 Confirm the Windows firewall is disabled.

Step 13 Confirm the Windows Updates are current and then disable Windows Updates.

Step 14 Log in as Administrator and remove all other accounts on the virtual machine.

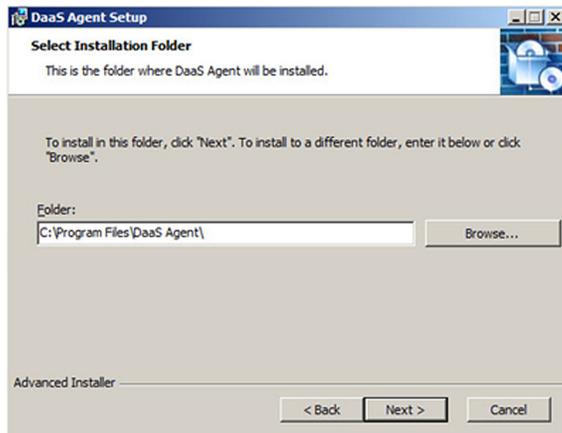
Installing the Desktop DaaS Agent.

Step 1 Copy the DaaSAgent_5.3.1.msi installer file onto the Windows 7 VM and double-click to start the Install.

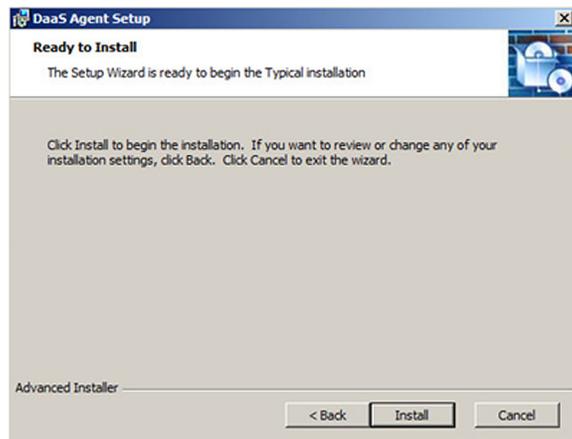
Step 2 Click Next.



Step 3 Select a specific path for installation or leave it as default. Click Next.



Step 4 Click Next.



Step 5 Click Finish.



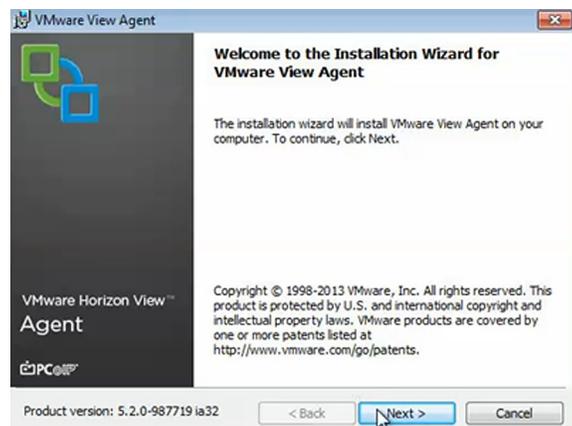
Note

The DHCP option code 74 needs to be configured with the tenant appliance IP addresses. This will be utilized by the DaaS Agent to register the virtual desktop to the tenant appliance.

Installing the VMware View Agent

Step 1 Copy the VMware-viewagent-5.2.0-987719.exe file onto the Windows 7 virtual machine.

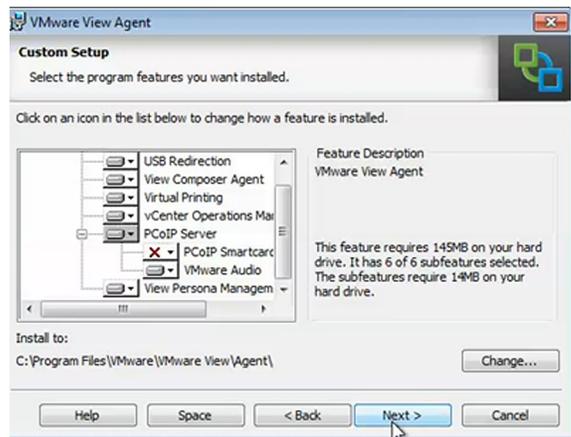
Step 2 Click Next.



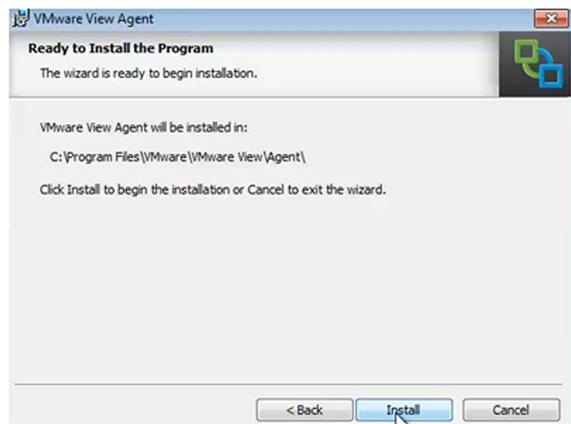
Step 3 Accept the User License Agreement. Click Next.



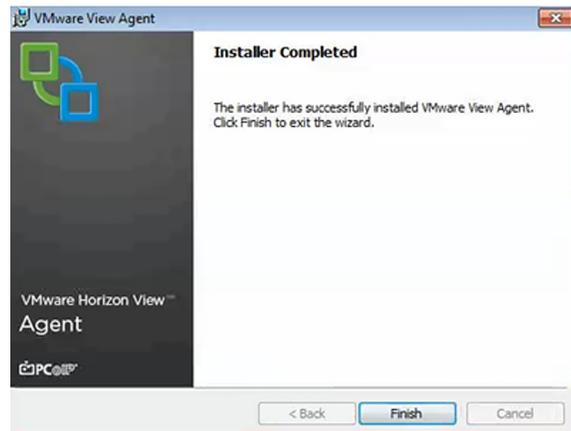
Step 4 Click Next.



Step 5 Click Install.



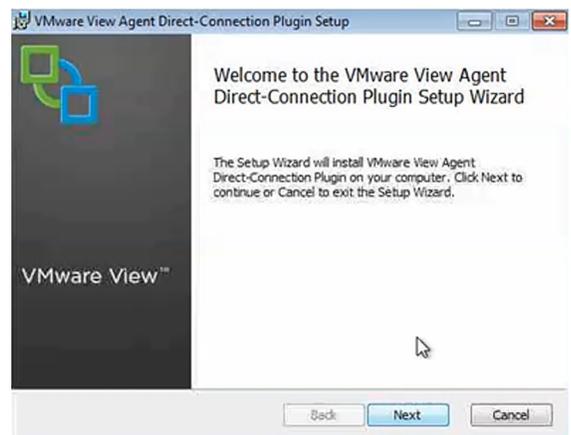
Step 6 Click Finish. A pop-up will then display and prompt a reboot.



Installing the VMWare View Agent Connect

Step 1 Copy the VMware-viewagent-direct-connection-5.2.0-987719.exe installer file to Windows 7 Golden virtual machine.

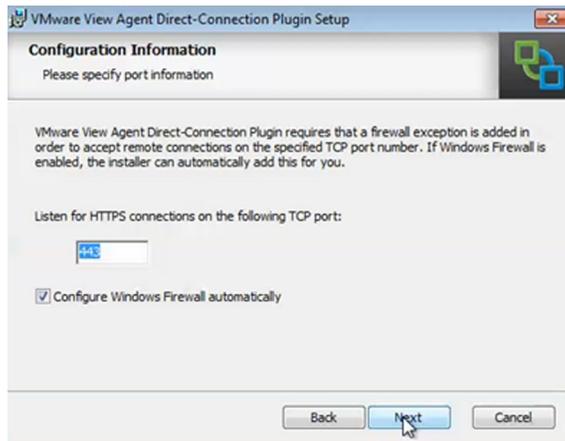
Step 2 Click Next.



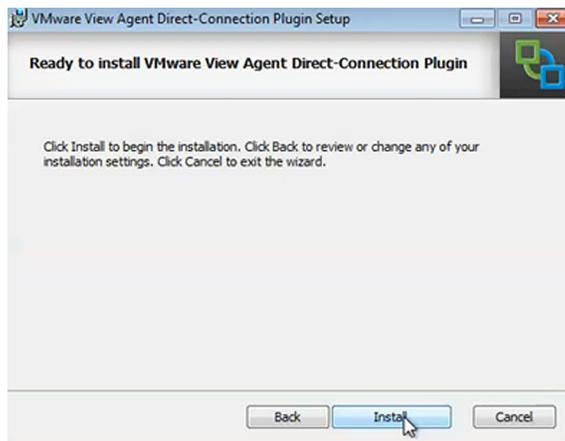
Step 3 Accept the End User License Agreement. Click Next.



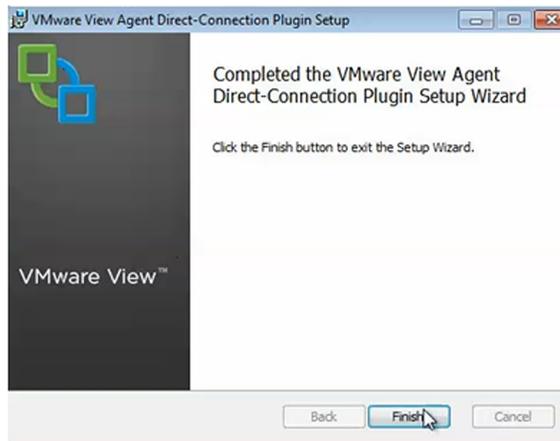
Step 4 Click Next.



Step 5 Click Install.



Step 6 Click Finish.



Note

Golden virtual machine must exist on the same datastore and VLAN where the desktops would reside.

Repeat the steps outlined in section [Creating the Golden Image](#) for preparing the Windows 7 golden image for other tenants.

Installing and Configuring Desktone

Building a Desktone environment requires the installation of the following components:

- Desktone Service Provider Appliances
- Desktone Resource Manager Appliances
- Desktone Tenant Appliances

Installing Desktone Primary Service Provider Appliance

Download the following required files for installation.



Note

Use the worksheets from Desktone to make sure you capture all the required information before you begin the installation.

Table 6

File Contents	File Names
Appliance template	AugustaSP3Template20130305.ova
5.3 debians	dt-platform-5_3_0.deb and dt-aux-1_1_0.deb
Patch	Dt-platform-5_3_0_patch_1.tgz,
Dt-platform-5_3_0_patch_2.tgz	



Note

On both management hosts, add the Service Provider datastore. Be sure to use the same name on each host.

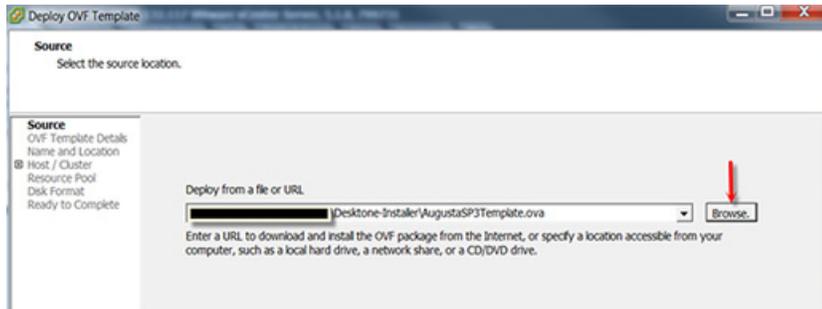
Step 1 From the vSphere Client for Service Provider, click on File, Select File and select Deploy OVF Template to deploy the first copy of the ova file.



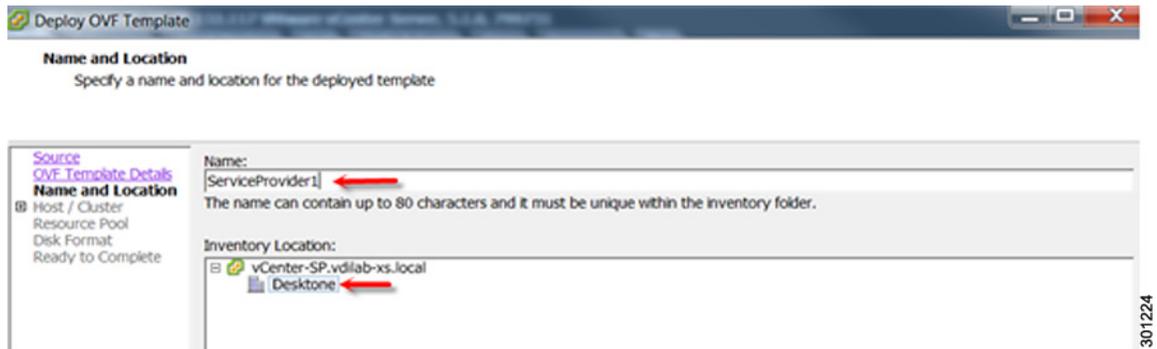
Note

This ova file becomes the first Service Provider appliance.

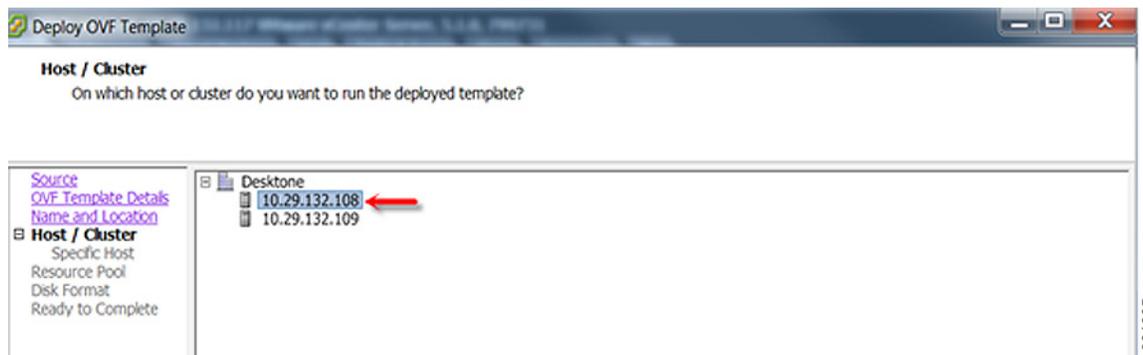
Step 2 Click Browse and select AugustaSP3Template.ova file

Step 3 Click Next.**Step 4** Review the OVF template details and click Next.**Step 5** Select a Name and the data center for Service Provider Appliance.

Note Do not use any special characters for the name.



Step 6 Select the host where the virtual machine will reside on and click Next.

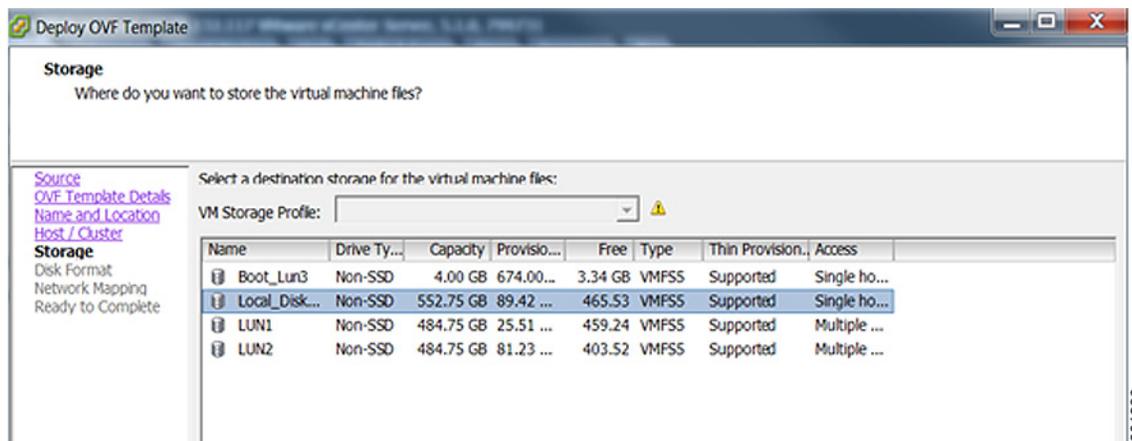


Step 7 Deploy the Deskton Service Provider Appliance to local storage on one of the two management hosts.



Note

- By default, the Deskton Platform will clone out management appliances on the local disk (through a localdatastore). This is considered a best practice. However, if desired, it is possible to use shared storage for management appliances.
- If using shared storage for management appliances on vCenter, there are a few guidelines:
- Any shared storage (NFS, iSCSI, or FC) can be used.
- Datastores must be manually created on each of the management hosts. The datastore name must be identical (case sensitive) on each management host.



Step 8 Select Thin Provision.

Step 9 Click Next.

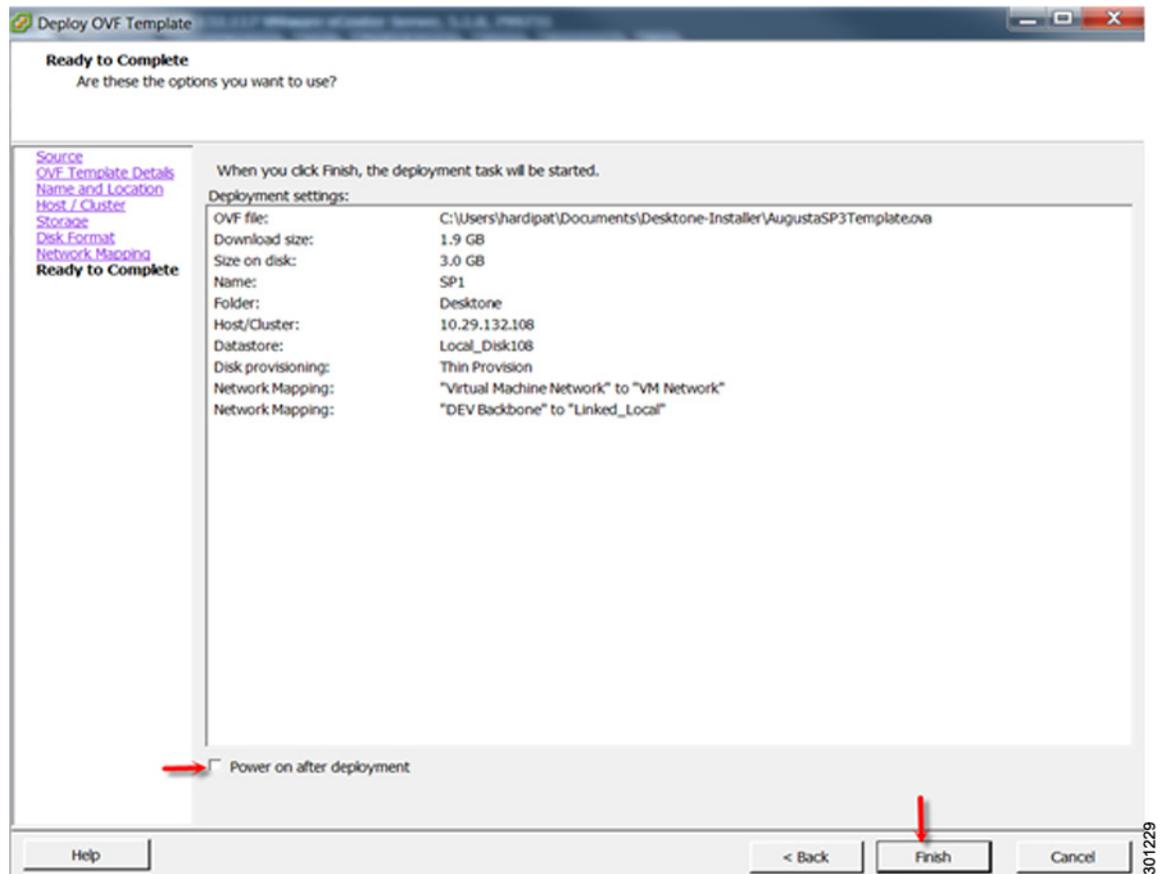
The screenshot shows the 'Deploy OVF Template' wizard at the 'Disk Format' step. The title bar reads 'Deploy OVF Template'. Below the title bar, the text asks 'In which format do you want to store the virtual disks?'. The 'Datastore' is set to 'Local_Disk108' and 'Available space (GB)' is '465.5'. There are three radio button options: 'Thick Provision Lazy Zeroed', 'Thick Provision Eager Zeroed', and 'Thin Provision'. A red arrow points to the 'Thin Provision' option. On the left side, there is a navigation pane with links for 'Source OVF Template Details', 'Name and Location Host / Cluster', 'Storage', 'Disk Format', 'Network Mapping', and 'Ready to Complete'. A vertical ID '301227' is on the right.

Step 10 Select the service provider network as the destination for the first source network; Virtual Machine Network and Link Local for the second source network, Dev Network.

The screenshot shows the 'Deploy OVF Template' wizard at the 'Network Mapping' step. The title bar reads 'Deploy OVF Template'. Below the title bar, the text asks 'What networks should the deployed template use?'. The left navigation pane shows 'Network Mapping' as the active step. The main area is titled 'Map the networks used in this OVF template to networks in your inventory'. It contains two tables: 'Source Networks' and 'Destination Networks'. The 'Source Networks' table has two rows: 'Virtual Machine Network' and 'DEV Backbone'. The 'Destination Networks' table has two rows: 'VM Network' and 'Linked_Local'. Red arrows point from 'Virtual Machine Network' to 'VM Network' and from 'DEV Backbone' to 'Linked_Local'. Below the tables is a 'Description:' field containing the text 'The Virtual Machine Network network'. A vertical ID '301228' is on the right.

Step 11 Review the deployment settings and select the check box to power on the virtual machine after deployment

Step 12 Click Finish.



Step 13 Repeat the steps 1-12 to deploy the second copy of the ova file, which becomes the template for all subsequent Desktope management appliances. For the template, do not check the box to power on after deployment.



Note The name of the template must be unique across all datacenters.

Configuring the Network on Desktope Primary Service Provider Appliance

Step 1 From the vSphere Client open a console for Desktope Primary Service Provider Appliance.

Step 2 Logon using the default credentials:

- Username: deskstone
- Password: Deskstone1

Step 3 Begin the bootstrap process by executing the following command:

```
sudo /usr/local/desktopone/scripts/bootstrap.sh
Enter password for Desktope: Deskstone1
```

```
desktopone@template:~$ sudo /usr/local/desktopone/scripts/bootstrap.sh
[sudo] password for desktopone:
```

Step 4 The bootstrap script prompts you to enter the network information.



Note The host reboots after entering the network information. Because the node is not configured until the reboot completes, disregard any error messages displayed on the console.

Field	Sample Value	Notes
Existing multi datacenter setup	No	Select “No” to install the first datacenter
Datacenter Name	Deskstone	Use Location Name if you have multiple DCs
IP for Eth1 (backbone)	169.254.1.10	For the Backbone network (must be a link-local address)
Enter netmask CIDR format 0-32	24	Use /22 for large number of tenants
Enter Ip for eth0 (SP)	10.29.132.119	For the SP network
Enter netmask CIDR format 0-32	24	For the SP network
Enter Gateway	10.29.132.1	For the SP network
Enter hostname	Sp1.vdilab-xs.local	
Enter nameserver	10.29.132.30	
Enter NTP servers	Ntp.ubuntu.com	
Is this an HA Service Provider appliance setup?	Yes	
Enter the floating IP address	10.29.132.120	
Enter psql password	XYZabc123	This alters the psql passwords for admin, master, slave and slony user.
Appliance password	Password123	The user-defined password for Service Provider appliances in this datacenter. Any Service Provider appliances accessible by ssh requires this custom password.
Does this configuration Look correct?	Yes or No	Review the information applied and select Yes or No.

```

Is this Service Provider appliance part of an existing multi Data Center setup?
If this is the first Data Center, answer no.

Answer yes or no: no
Enter Data Center Name:
Deskstone
Enter IP for eth1 (backbone):
169.254.1.10
Enter netmask CIDR format 0-32:
24
Enter IP for eth0 (SP):
10.29.132.119
Enter netmask CIDR format 0-32:
24
Enter Gateway:
10.29.132.1
Enter hostname:
Sp1.vdilab-xs.local
Enter nameserver:
10.29.132.30
Enter NTP servers (if no more NTP servers, press <RETURN>)
If you press enter for the first NTP server, [ntp.ubuntu.com] will be configured
.

Is this an HA Service Provider appliance setup?:
Answer yes or no: yes
Enter the floating IP address:
10.29.132.120

```

301231

```

After reboot put the latest dt-platform package in the directory: /data/repo/

Does this configuration look correct?

Data Center Name: Deskstone
Hostname: Sp1.vdilab-xs.local
eth1 IP: 169.254.1.10
eth1 Netmask: 255.255.255.0
Backbone IP Block: 169.254.1.0/24
eth0 IP: 10.29.132.119
eth0 Netmask: 255.255.255.0
eth0 CIDR: 24
Gateway: 10.29.132.1
HA Transit Server IP: 169.254.1.11
Floating IP: 10.29.132.120
Data Center UID: 256a98db-ed4e-4604-9dfe-b388a5913e31
Nameserver: 10.29.132.30
Data Center Master: True
Multi Data Center: False
NTP Server 1: ntp.ubuntu.com
UMGR UID: 09b83774-cabf-4250-aa01-d55102950da0
Answer yes or no: yes

```

301232

- Step 5** After rebooting, login with the configured username/password: deskstone/Password123
- Step 6** Copy the following files to the /tmp directory on the service provider appliance.

```
dt-platform-5_3_0.deb
dt-aux-1_1_0.deb
```

- Step 7** ssh into the service provider appliance using login credentials. Move the files copied to /tmp directory into the /data/repo directory on the appliance using the commands:

```
sudo mv /tmp/dt-platform-5_3_0.deb /data/repo
sudo mv /tmp/dt-aux-1_1_0.deb /data/repo
```

```
desktone@Sp1:~$ sudo mv /tmp/dt-platform-5_3_0.deb /data/repo/
[sudo] password for desktone:
desktone@Sp1:~$
desktone@Sp1:~$
desktone@Sp1:~$ sudo mv /tmp/dt-aux-1_1_0.deb /data/repo/
```

301234

- Step 8** Run the bootstrap script a second time to install the Desktone software:

```
sudo /usr/local/desktone/scripts/bootstrap.sh
```



Note It might take up to five minutes for the appliance to start after the reboot. Because the node is not configured until the reboot cycle completes, you can disregard any error messages displayed on the console.

```
desktone@Sp1:~$ sudo /usr/local/desktone/scripts/bootstrap.sh
Selecting previously deselected package dt-aux.
(Reading database ... 37442 files and directories currently installed.)
Unpacking dt-aux (from /data/repo/dt-aux-1_1_0.deb) ...
This is Pre Install script
Setting up dt-aux (1.1.0) ...
This is Post Install script

Processing triggers for ureadahead ...
Selecting previously deselected package dt-platform-5-3-0.
(Reading database ... 37480 files and directories currently installed.)
Unpacking dt-platform-5-3-0 (from ../repo/dt-platform-5_3_0.deb) ...
This is Pre Install script
Setting up dt-platform-5-3-0 (5.3.0) ...
This is Post Install script

Shutting down sfcf: done.
```

301235

Configuring the Desktone Primary Service Provider Appliance

- Step 1** Browse to the Desktone service provider portal by entering the URL or IP address:
<https://<IP address of service provider appliance>/service>
- Step 2** Enter the information from the service provider Active Directory and click Save.



Note The first time, browsing the service center, you will be prompted for Domain Bind and Group Info.

301236

301237

Step 3 Enter a username, password and domain based on the Active Directory information on the service center login page.

- Step 4** Enter the IP address or FQDN for vCenter hosting primary Service Provider Appliance, username and password and click Discover Server.

Discover Management Server

Use this form to discover the host management server which holds the virtual machine template for installing Deskstone appliances.

Once the management server has been discovered and the capacity metrics have been saved, the virtual machines on that server will be displayed. Select the virtual machine to be used as the appliance template.

* IP Address / Hostname

* Username

* Password

Discovering host and calculating capacity...

- Step 5** Enter the memory over allocation ratio for each discovered host.



Note If the server is too small to accommodate the ratios, you may be prompted to re-configure them. Click Save to set the ratios.

10.29.132.108

The host has been discovered successfully. Please review the host capacity derived from the ratios you entered, and ensure there is sufficient capacity compared to the consumed resources determined from existing VMs on the host.

Memory Installed	Memory Overallocation Ratio	Total VM Memory	Current Used VM Memory
255 GB	* <input type="text" value="1.0"/>	= 255 GB	4 GB
CPUs Installed	Virtual to Physical CPU Ratio	Total VM CPU Count	Current Used VM CPU Count
16	* <input type="text" value="10.0"/>	= 160	2

Adjust the resource ratios above as needed, then click the Save button.

- Step 6** Repeat step 5 for the other ESXi host.

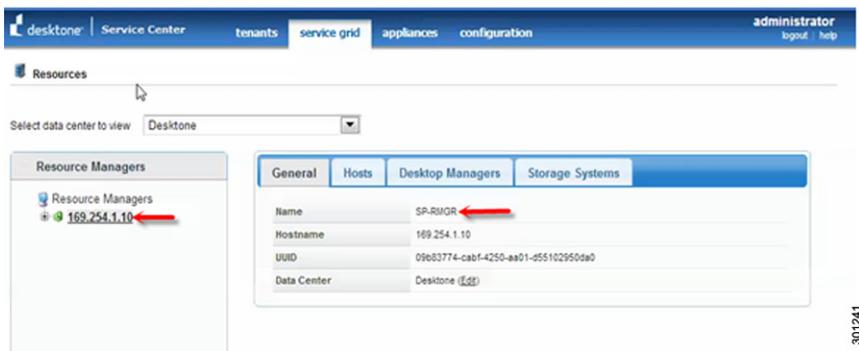
- Step 7** After setting the ratios, the Service Center displays a pop-up listing the virtual machines on your host or in your vCenter Data Center.

Step 8 Select the Deskstone appliance template to be used for cloning additional management appliances.



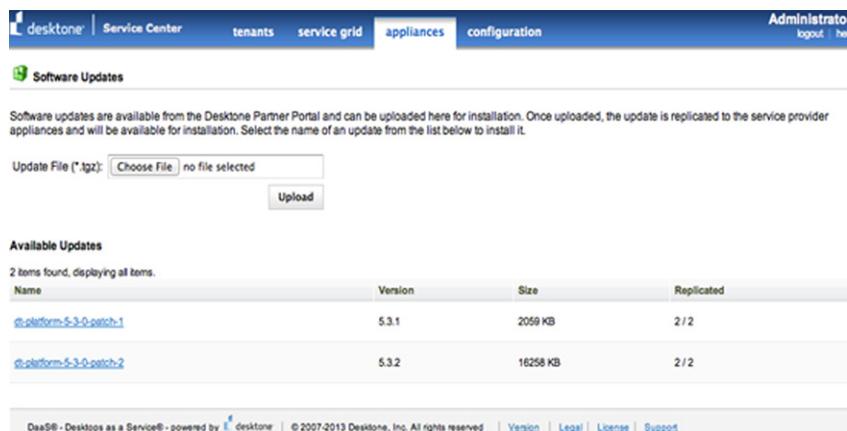
Step 9 Go to service grid --> resources. In the Resource Managers panel on the left, click the IP address of the resource manager.

Step 10 From the General tab, in the Name field, double-click on the IP address of the resource manager. Change the name to a user-friendly name and click OK.



Step 11 Apply the patches for the base image of service provider appliance by clicking appliances --> software update.

Step 12 Browse to the patches to be applied and click on Upload.



Step 13 SSH into the service provider appliance using login credentials and execute the following commands:

```
cd /data/repo
sudo dpkg -i dt-platform-5_3_0_patch_1.deb
sudo dpkg -i dt-platform-5_3_0_patch_2.deb
sudo service dtService restart
```

Installing the Desktop Secondary Service Provider Appliance

Step 1 Go to Service grid --> data center and click Edit.

The screenshot shows the Desktopone Service Center interface. The top navigation bar includes 'Service Center', 'tenants', 'service grid', 'appliances', and 'configuration'. The user is logged in as 'administrator'. The 'Data Centers' section is active, showing a table with one entry: 'Deskstone'. The 'Edit' button for this entry is highlighted with a red arrow.

Name	Friendly Name	Description	VMs	Hosts	Resource Managers	Desktop Managers	
Deskstone	Deskstone	Deskstone		7	2	2	2 Edit

Step 2 Verify the displayed information and click Add Appliances.

The screenshot shows the 'Edit Data Center' dialog box. The fields are as follows:

* Name	Deskstone
Friendly Name	Deskstone
Description	Deskstone
Auto assigned RMgr	169.254.1.12
NTP Server	[ntp.ubuntu.com]
Subnet Mask	255.255.255.0
IP Address Block	169.254.1.0/24
Network ID Type	vlan
Network ID	132
Backbone Network Type	vlan
Backbone Network ID	55
DNS Server	10.29.132.30
Gateway	10.29.132.1

Buttons: Save, Add Appliances, Cancel

Step 3 Select Service Provider Appliance from the Appliances and enter values for the fields to create the Appliance.

The screenshot shows the 'Add Appliances' dialog box. The fields are as follows:

* Name	SP2	* Friendly Name	SP2
* IP Address	10.29.132.121	* Start Date	05/20/2013
		Start Time	00:00

Buttons: Add Appliances

**Note**

Reservation will create a virtual machine for the second service provider appliance for High Availability, customize and install. This screen is accessible by navigating to Appliances --> Reservations and clicking the details on the appropriate reservation. From here, you can check the status of the SP2 appliance.

SP2	cloneVM	100			completed	successful
SP2	customize	100			completed	successful
SP2	install	100			completed	successful

Friendly Name	Name	Service Provider	Tenant	Template	RMgr	Desktop Mgr	Primary	Organization Primary	Data Center	Tenant	State	Version	Actions
ServiceProvider1	ServiceProvider1	✓	-	-	✓	-	✓	✓	Deskstone	Service Provider	Installed	5.3.1	Actions
AugustaSP2Template20130305	AugustaSP2Template20130305	✗	-	✓	-	-	✓	-	Deskstone	Service Provider	Reserved	5.3.0	
SP2	SP2	✓	-	-	✓	-	-	-	Deskstone	Service Provider	Installed	5.3.1	Actions

301245

Installing the Deskstone Tenant Resource Manager Appliances

Step 1 In the service center console, select service grid --> data centers.

Name	Friendly Name	Description	VMs	Hosts	Resource Managers	Desktop Managers	Actions
Deskstone	Deskstone	Deskstone	0	1	1	0	Edit

301246

Step 2 Click Add Appliances.

Field	Value
* Name	Deskstone
Friendly Name	Deskstone
Description	Deskstone
Auto assigned RMGr	
NTP Server	[ntp.ubuntu.com]
Subnet Mask	255.255.255.0
IP Address Block	169.254.1.0/24
Network ID Type	vlan
Network ID	132
Backbone Network Type	vlan
Backbone Network ID	55
DNS Server	10.29.132.30
Gateway	10.29.132.1

Buttons: Save, Add Appliances, Cancel

Step 3 Select the Appliance Type as Resource Manager and fill in the appropriate information in the value field for Primary and Secondary appliance and new reservation fields.

Appliance Install

Appliance Type: Resource Manager

Primary Appliance

- Name: RMGR1
- IP Address: 10.29.132.108

Secondary Appliance

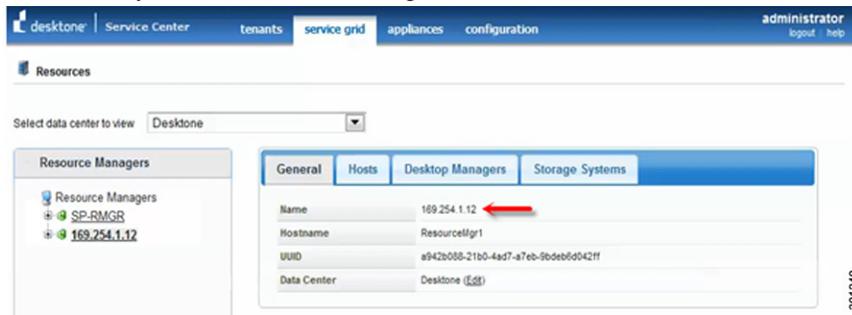
- Name: RMGR2
- IP Address: 10.29.132.109

New Reservation

- Friendly Name: RSMGR
- Start Date: 05/20/2013
- Start Time: 00:00

Buttons: Create Appliance, Cancel

- Step 4** From the general tab, in the name field double-click on the IP address and change the name for the Resource Manager. For example, TenantRMGR and click OK. You can do this by going to the Service Center console, select Service grid --> resources. In the Resource Managers panel select the IP address for the newly created resource manager.



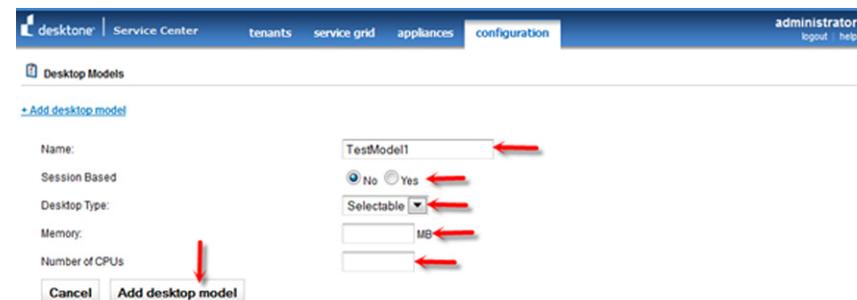
- Step 5** Select the Tenant Resource Manager as the default for the Datacenter.



Note It is very important to verify that the default Resource Manager for the Datacenter is the Tenant Resource Manager, not the Service Provider Resource Manager.

Adding Desktop Models

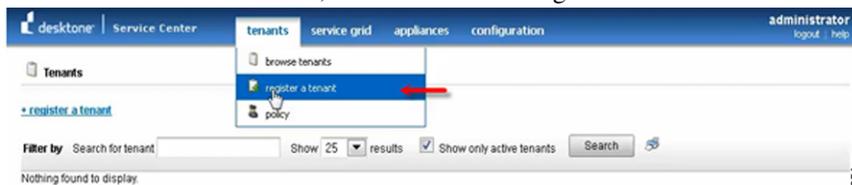
- Step 1** Desktop Model is used to create multiple desktops of a certain configuration and priced accordingly. Tenants can then choose from the available desktop models.



- Step 2** Add Desktop Model by going into the Service Center console, Select Configuration --> Desktop Models and click Add Desktop Model.

Installing the Deskstone Tenant Appliances

- Step 1** In the service center console, select tenants --> register a tenant.



Step 2 From the General Info tab, enter the required fields of Tenant Name, Administrator Name, and Database Password.

deskstone | Service Center | tenants | service grid | appliances | configuration | administrator | logout | help

Register a tenant

General Info

* Tenant Name: TenantA

* Administrator Name: [Redacted]

* Database Password: [Redacted] Password must be at least 4 characters long

Administrator Phone: [Empty]

Administrator Email: [Empty]

Technical contact name: [Empty]

Technical contact phone: [Empty]

Technical contact email: [Empty]

CRM System URL: [Empty]

Technical Support URL: [Empty]

Style Sheet URL: [Empty]

License File URL: [Empty]

Business Support URL: [Empty]

Save and Create Appliances | Save and Finish Later | Cancel

301252

Step 3 Enter the Network Information.



Note Network ID is an install time decision and cannot be changed after the tenant has been installed. Additional networks can be added to a tenant at any time. The first network added would be the default for the tenant and also the network that the tenant appliances will reside on.

deskstone | Service Center | tenants | service grid | appliances | configuration | administrator | logout | help

Register a tenant

General Info

Networks

Custom Fields

Networks

Add the tenant information for the required data centers below so appliances and desktops can be added. All fields are required if you want to create appliances now.

Data Center: Deskstone Add

Deskstone

* Network ID: 53

* Network ID Type: VLAN

* Gateway: 10.10.53.1

* DNS Server: 10.10.53.30

* Subnet Mask: 255.255.255.0

Default: [Green checkmark]

Save and Create Appliances | Save and Finish Later | Cancel

301253

- Step 4** From the Custom Fields tab, enter any site-specific information you want to maintain. These are freemform text fields with no data validation; the content is optional.

desktop | Service Center | tenants | service grid | appliances | configuration | administrator

Register a tenant

General Info
Networks
Custom Fields

Custom Fields

Custom Field 1
Custom Field 2
Custom Field 3
Custom Field 4
Custom Field 5
Custom Field 6

Save and Create Appliances | Save and Finish Later | Cancel

301254

- Step 5** Click Save and Create Appliances.

- Step 6** From the Tenant Install page, enter the values for the tenant Appliances and click Create Appliances.

desktop | Service Center | tenants | service grid | appliances | configuration | administrator

Tenant Install

Data Center: Desktop

Primary Appliance

* Name: TenantA1
* IP Address: 10.10.53.31

Secondary Appliance

* Name: TenantA2
* IP Address: 10.10.53.32

Floating IP Address: 10.10.53.33

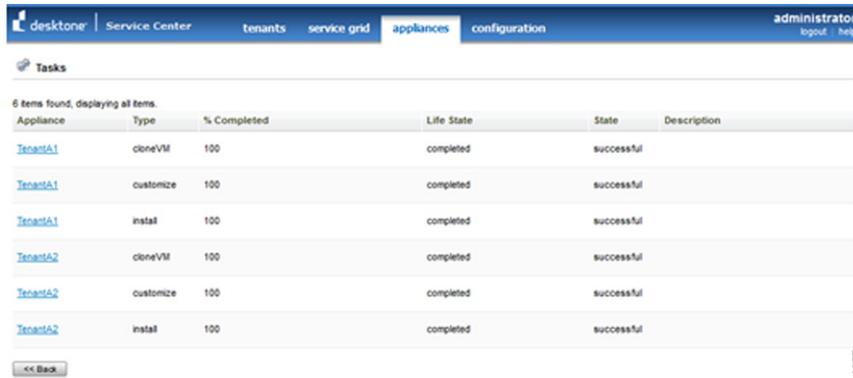
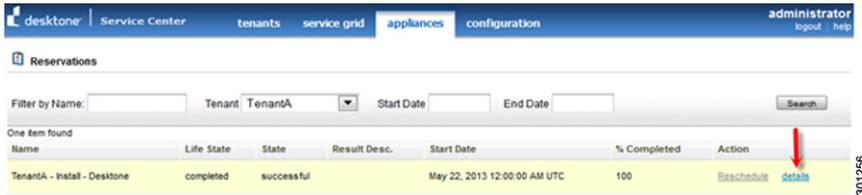
Schedule Install

* Start Date: 05/22/2013
Start Time: 00:00
format:(HMM) 00:00-23:59 GMT —use 00:00 for "now"

Create Appliances | Cancel

301255

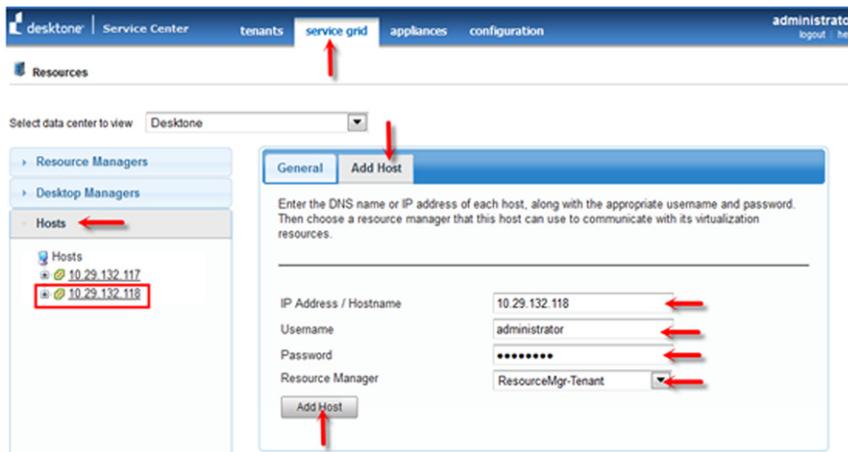
Step 7 To monitor the status of the tenant appliances being created, go to the service center console, Select Appliances --> Reservations and click on details to view the reservation of interest.



Step 8 Repeat the steps 1-7 to create Tenant Appliances for every new tenant.

Assigning Hosts to Individual Tenants

- Step 1** Assign a host by clicking Add host. Do this by going to the service center console, Select service grid --> resources.
- Step 2** From the left side of the screen select Hosts.
- Step 3** Click Add Host.
- Step 4** Enter the IP address or FQDN for vCenter created for Tenant Desktop Hosts and credentials.



Step 5 Edit the host by selecting the vCenter for Tenant desktop hosts.

The screenshot shows the Cisco Desktop Manager interface. The top navigation bar includes 'desktop', 'Service Center', 'tenants', 'service grid', 'appliances', and 'configuration'. The user is logged in as 'administrator'. The main content area is titled 'Resources' and shows a tree view of hosts under 'Desktop Managers'. The host '10.29.132.107' is selected. The configuration panel on the right shows the following details:

Property	Value
Host role	Active
Host purpose	Desktop
LUID	2a7259ad-e722-e211-0210-ad030000000a
Address	10.29.132.107
Type	Virtual Center
Status	connected
CPU Count	16
Total Memory (GB)	255 GB
Shared	No
Virtual to Physical CPU Ratio	
Memory Overallocation Ratio	
Total VM CPU Count	
Current Used VM CPU Count	
Total VM Memory (GB)	0.0 GB
Current Used VM Memory (GB)	0.0 GB
Version	5.1.0 VMware ESX 5.1.0 build-030463

Step 6 Change the memory overallocation ratio if required and click Save.

The screenshot shows the resource allocation ratios for host 10.29.132.107. The page displays a table with the following data:

Memory Installed	Memory Overallocation Ratio	Total VM Memory	Current Used VM Memory
255 GB	1.5	383 GB	24 GB
CPUs Installed	Virtual to Physical CPU Ratio	Total VM CPU Count	Current Used VM CPU Count
16	10.0	160	14

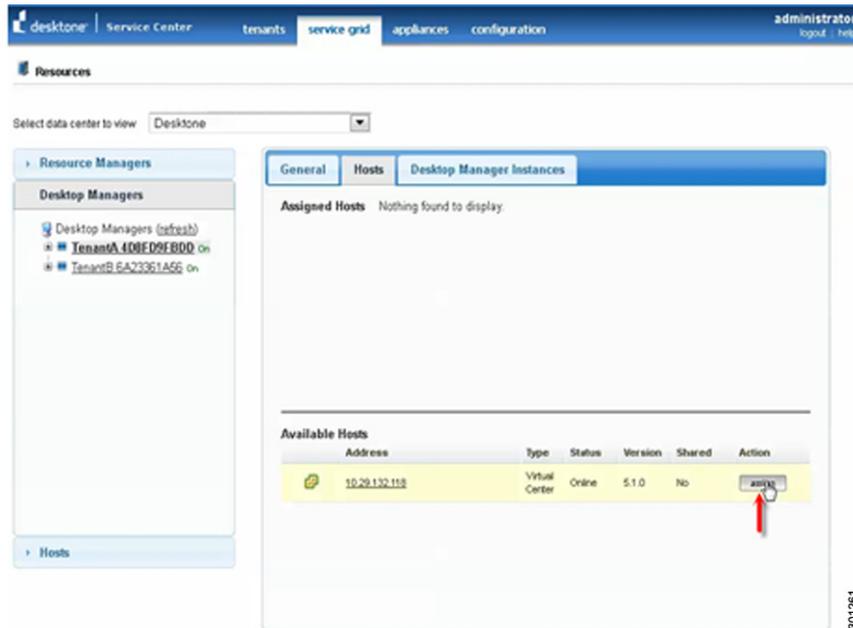
Adjust the resource ratios above as needed, then click the Save button.



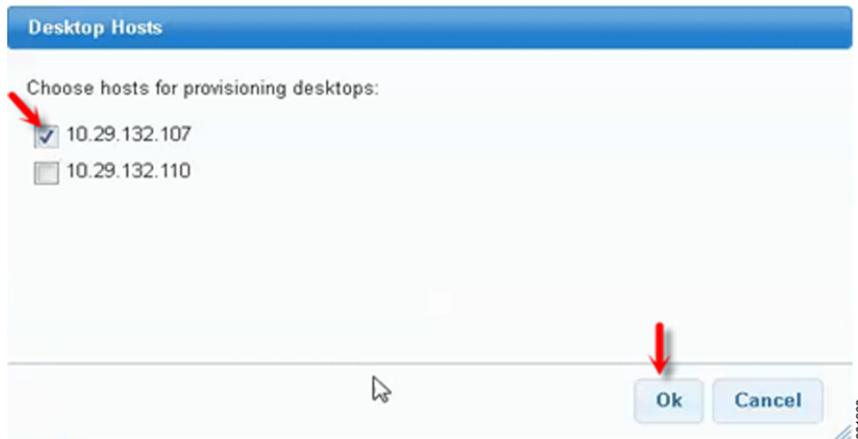
Note This option only allows for ratios to be increased and may not be decreased after the initial setup.

Step 7 Repeat the above steps 1-6 for all ESXi hosts reserved for the tenant.

Step 8 Assign a host to the Desktop Manager by clicking on Assign under Available Hosts from the Desktop Managers tab.



Step 9 Select the hosts in the vCenter for this tenant.



Step 10 Repeat steps 9 and 10 for the other hosts reserved for the tenant.

Assigning Quotas for Tenants

Step 1 To assign a quota, go in the Service Center console -> tenants and edit the tenant.

The screenshot shows the Service Center console with the 'tenants' tab selected. A table lists two tenants:

Id	Tenant Name	Admin Contact	VM Quota	Action
1001	TenantA	hardpat	0	Edit
1002	TenantB	amulhayan	0	Edit

Step 2 Enter the values in the VM Quota column for respective Desktop Models and click Update.



Note If the quota you entered in the VM Quota field is too large, the system will automatically reduce the quota to the largest possible quota.

The screenshot shows the 'Editing TenantA' page with the 'Quotas' tab selected. A table titled 'Desktop Model Quota' is displayed:

Name	Type	Memory	CPUs	Session Model	VM Quota	In Use
Normal	selectable	2048	1	No	5	0
Utility	selectable	2048	1	No	2	0
VDI-Model1	dynamic	1024	1	Yes	2	0
VDI-Model2	selectable	1536	1	No	2	0

An 'Update' button is located at the bottom of the table.

Step 3 Check the desired protocols for the tenant and click Update.

Protocol Quota

Set the protocol quota to a number or unlimited. Unlimited allows unrestricted use of the selected protocol.

Protocol	Unlimited	Quota	In Use
RDP	<input checked="" type="checkbox"/>	0	3
RGS	<input type="checkbox"/>	0	0
HDX	<input type="checkbox"/>	0	0
VNC	<input type="checkbox"/>	0	0
NX	<input type="checkbox"/>	0	0
PCoIP	<input checked="" type="checkbox"/>	0	3

301265

Configuring Tenant Appliance

Step 1 Browse to the Desktop Enterprise Center by entering the URL or IP address in a browser:
<https://<IP address of Tenant Appliance>/admin>

Step 2 Enter the values under Domain Bind, Group Info and Domain Join Info when prompted and click Save.

desktop | Enterprise Center

Register a domain

Save successful. Proceed with the other tabs or save this URL and enter it in a new browser if you wish to edit the other tabs later.
<https://10.10.53.33/admin/config/editDomainFlow.action?ticket=1QdGcUEdckhUegZHfA==>

Domain Bind | Group Info | Domain Join Info

* Name: tenanta
NETBIOS domain name. For example, SALES.

* Domain Suffix: tenanta.local
For example, mycompany.com

* Protocol: ldap

* Port: 389

Domain Controller IPs: 10.10.53.30
Preferred domain controller IP list with comma separation

* Context: dc=tenanta,dc=local
AD context of domain. For example, dc=mycompany,dc=com

* Domain Bind Account DN: cn=admin,dc=tenanta,dc=local
Distinguished name of admin user for this domain. For example, cn=admin,dc=tenanta,dc=local

* Password: *****
Password for user named above.

* Password verify: *****

301266

Step 3 Enter the values under Group Info and click Save.

deskstone Enterprise Center

Register a domain

Start typing the first several characters of the group distinguished name and then wait for suggestions to appear.

Domain Bind Group Info Domain Join Info

* Admin Groups =administrators,cn=builtin,dc=tenanta,dc=local Super Admin - Enterprise Admin
Add Admin Group
For example, CI=myadmins,ou=groups

* User Groups cn=users,cn=builtin,dc=tenanta,dc=local
Add User Group
For example, CI=myusers,ou=groups

Notice: Only users assigned directly to this domain will be available. Child/trusted domain users are not currently supported.

Save Clear

301267

Step 4 Enter the values under Domain Join Info and click Save.

deskstone Enterprise Center

Register a domain

In the Domain Join User field, enter the login name of a user having unlimited privileges to join computers to this domain.

Domain Bind Group Info Domain Join Info

Domain Join User administrator
AD login used to join VM to this domain. For example, john.doe

Domain Join Password
Password for domain user.

Domain Join Password Verify

* Primary DNS server IP 10.10.53.30
IP address of primary DNS server.

Secondary DNS server IP 127.0.0.1
IP address of secondary DNS server. Leave this blank if you do not have one.

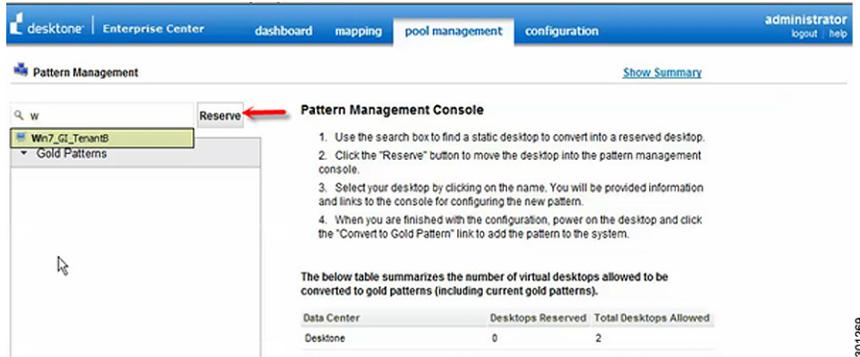
Save Clear

301268

Step 5 Login into the Desktopone Enterprise Center login screen again to continue the configuration.

Adding the Golden Template into Desktopone Tenant Portal

- Step 1** Click on Pool Management -> Pattern Management and type in the Golden Windows 7 VM and click Reserve.



Pattern Management Console

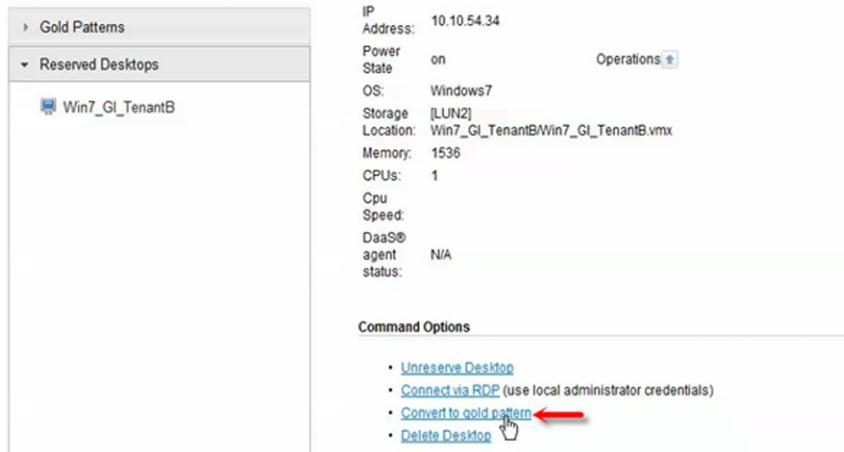
1. Use the search box to find a static desktop to convert into a reserved desktop.
2. Click the "Reserve" button to move the desktop into the pattern management console.
3. Select your desktop by clicking on the name. You will be provided information and links to the console for configuring the new pattern.
4. When you are finished with the configuration, power on the desktop and click the "Convert to Gold Pattern" link to add the pattern to the system.

The below table summarizes the number of virtual desktops allowed to be converted to gold patterns (including current gold patterns).

Data Center	Desktops Reserved	Total Desktops Allowed
Desktopone	0	2

301269

- Step 2** Click Convert to Gold Pattern.



Gold Patterns

- Reserved Desktops
 - Win7_Gl_TenantB

IP Address: 10.10.54.34

Power State: on Operations

OS: Windows7

Storage: [LUN2]

Location: Win7_Gl_TenantB/Win7_Gl_TenantB.vmx

Memory: 1536

CPUs: 1

Cpu Speed:

DaaS agent status: N/A

Command Options

- Unreserve Desktop
- Connect via RDP (use local administrator credentials)
- Convert to gold pattern
- Delete Desktop

301270

- Step 3** Enter the values for the fields shown below and click Convert to Pattern.



Note Do not check the Override Licensing option if you want to use KMS.

Gold Pattern Data

Override Licensing Override the default licensing scheme (KMS or MAK) by using a licensing key

License key: [REDACTED]

* Time Zone: (UTC-08:00) Pacific Time(US & Canada)

* Company Name: TENANTB

Password for local administrator: Administrator

* Password: [REDACTED]

* Password verify: [REDACTED]

Convert to Pattern Cancel

Step 4 Verify the progress under Configuration -> Task and Events.

Enterprise Center dashboard mapping pool management configuration administrator

Tasks and Events

Filter by Status: - Select - Pools: ---

Date	State	Description	% Completed
May 24, 2013 6:13:13 PM UTC	prepared		0
May 24, 2013 6:13:14 PM UTC	queued		0
May 24, 2013 6:13:14 PM UTC	running		0
May 24, 2013 6:13:16 PM UTC	running	Called Resource Manager to convert pattern to template	10
May 24, 2013 6:13:16 PM UTC	running	Sending Element inventory to Fabric	15
May 24, 2013 6:13:20 PM UTC	running	Seal message sent to VM, waiting for machine to shut down	25
May 24, 2013 6:14:20 PM UTC	success	Successfully converted VM to gold pattern	100

Step 5 Enable the Gold pattern by clicking Enable.

Step 6 Click Update.

**Note**

Reseal option is required only when the Golden Pattern virtual machine is powered on.

301273

Creating a Pool of Desktops in the Tenant Portal

Step 1 Click Pool Management -> Create Pool and select Individual Desktop Based.

301274

Step 2 Enter the values and click Customize Pool.

**Note**

There are two desktop types: Static and Dynamic. Static desktops are assigned to individual users and the same desktop is available to users at all times. Dynamic desktops are available to any user for the duration of their session.

deskstone | Enterprise Center | dashboard | mapping | pool management | configuration

Create New Pool

Please fill in the Pool specifications below.

Data Center: Deskstone

Name: TenantB-DesktopPool

Desktop Model: VDI-Model2

Protocols: RDP VNC RGS HDX NX PCoIP

Desktop Type: Static

Gold Pattern: Win7_GL_TenantB

Pool Size: 2 (2 Remaining)

Clear Customize Pool

301275

Step 3 Enter the values under Policies as applicable for each tenant and click Review Pool.

deskstone | Enterprise Center | dashboard | mapping | pool management | configuration | administrator | logout | help

Create New Pool > Deskstone > TenantB-DesktopPool > Policies

The provisioning policy defines all the attributes required for provisioning new virtual desktops in this pool.

VM name composition rule: TenantB-VM (incremental number, 1-12 character base name)

Computer OU: ou=Deskstone/VMs. This is an optional field.

Domain Join: No Yes (Allow the desktop to join the selected domain)

Domain: TENANTB (NETBIOS domain this pool belongs to)

Assigned Groups: Add User Group (AD groups to assign to this pool. For example, cn=deskstoneusers,ou=groups)

Next >>

<< Back Review Pool

301276

Step 4 Click Create.

deskone | Enterprise Center | dashboard | mapping | pool management | configuration | administrator

Create New Pool > TenantB-DesktopPool > Confirm

Please review the pool information below. [Click the Create button to create the pool](#)

POOL CONFIGURATION

Name	TenantB-DesktopPool
Pool Type	Fixed size
Pool Size	2
Based On	
Gold Pattern	Win7_GI_TenantB
Desktop Type	Static
Desktop Model	VDI-Model2
Protocols	PCoP RDP
Default Protocol	None
Computer OU	
Session Timeout for VM (in ms)	3600000

WINDOWS CONFIGURATION

VM name composition rule	TenantB-VM
Assigned Groups	
Domain Join	Yes
Domain	TENANTB
Run once script	
Remote Applications	0

<< Back **Create**

Step 5 You can monitor the task of new desktops being created under Pool Management --> Tasks and Events.



Note

- Desktops will go through restart several times for customization steps like obtaining DHCP IP address, joining Domain, etc.,
- Two virtual displays with 128MB of video memory configuration are used by default when Desktopone creates desktops for any desktop pool using PCoIP.

deskone | Enterprise Center | dashboard | mapping | pool management | configuration | administrator

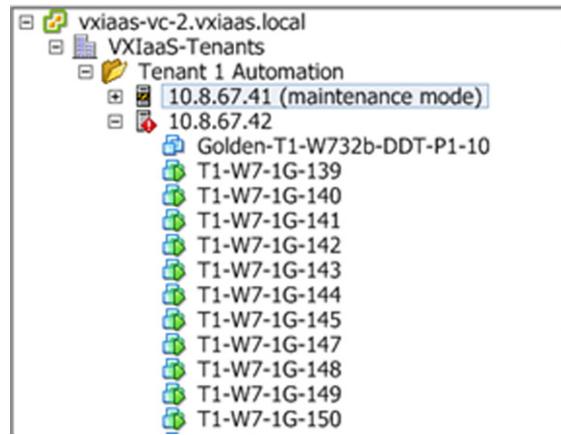
Tasks and Events

Filter by Status: - Select - Pools: ---

Expanding pool TenantB-DesktopPool				36% Complete
Date	State	Description	% Completed	
May 29, 2013 6:57:23 PM UTC	prepared		0	
May 29, 2013 6:57:24 PM UTC	queued		0	
May 29, 2013 6:57:24 PM UTC	running		0	
May 29, 2013 6:57:33 PM UTC	running	Cloning out VM 'TenantB-VM101'	0	
May 29, 2013 6:57:45 PM UTC	running		36	

Expanding pool TenantB-DesktopPool				46% Complete
Date	State	Description	% Completed	
May 29, 2013 6:56:55 PM UTC	prepared		0	
May 29, 2013 6:57:04 PM UTC	queued		0	
May 29, 2013 6:57:14 PM UTC	running		0	
May 29, 2013 6:57:23 PM UTC	running	Cloning out VM 'TenantB-VM100'	0	
May 29, 2013 6:57:45 PM UTC	running		46	
May 29, 2013 6:58:02 PM UTC	running	Finished post clone customization, attempting to join the domain	90	

Figure 13 *Desktops Created on a Tenant*



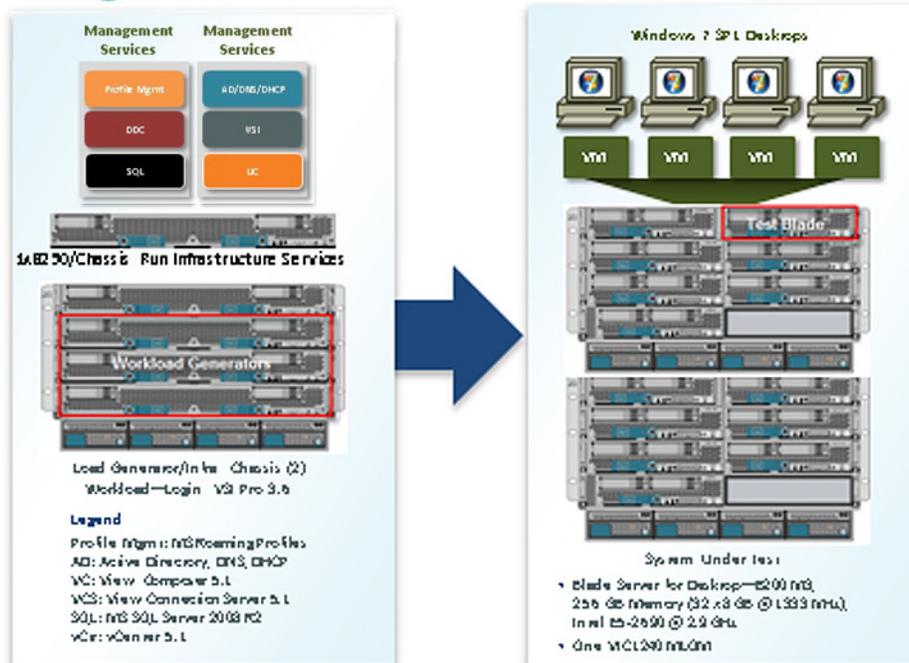
5 Test Setup and Configuration

For this project, a single Cisco UCS B200 M3 blade server in a single chassis for performance was used in testing.

Cisco UCS Test Configuration for Single Blade Scalability

Figure 14 Cisco UCS B200 M3 Blade Server for Single Server Scalability

Cisco UCS B200 M3 Blade Server Single Blade Test Result— 203 Users



301280

Hardware Components:

- 2 X Cisco Nexus 7009
- 2 X Cisco Catalyst 6506 with ASA-SM
- 2 X Cisco ASR 9006
- 2 X Cisco Nexus 5548UP Access Switches
- 2 X Cisco UCS Fabric Interconnect 6248UPs
- 1 X Cisco UCS B200-M3 - Virtual Desktop host:
 - 2 x E5-2690 @ 2.9 GHz CPUs
 - 256GB of memory (16 GB X 16 DIMMS @ 1666 MHz)

- 1 x VIC-1240 Converged Network Adapter/Blade
- 2 X Cisco UCS B series - Desktone Management hosts with minimum of:
 - 2 x Intel Xeon 5680 @ 3.333 GHz CPUs
 - 96 GB of memory
 - 1 X Converged Network Adapter/Blade
- 1 X NFS System storage array:
 - 2 x Service Controllers
 - 2 x 10 GBe Ports per Controller
 - SSD/SAS drives to support atleast 25 IOPS per desktop
 - Boot – The ratio of Read/Write is 9:1 (approx)
 - Run – The ratio of Read/Write is 4:3 (approx)

Detailed Windows Configuration

- 215 Desktop Configuration: Windows 7 SP1 32 bit, 1vCPU, 1 GB of memory (1GB Reserved), 24 GB/VM, 2 virtual displays with 128MB of Video Memory.
- 155 Desktop Configuration: Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory (1GB Reserved), 24 GB/VM, 2 virtual displays with 128MB of Video Memory.

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco RTP labs with joint support from Desktone. The test results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the Hosted VDI model under test.

Test metrics were gathered from the hypervisor, virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for this hardware configuration and the results were found to be relatively consistent from one test to the next.

Load Generation

Within the test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the Desktone 5.3.2 environment and executing a typical end-user workflow. To generate load within the environment, an auxiliary software application was required to generate the end user connection to the Desktone environment, to provide unique user credentials, to initiate the workload, and to evaluate the end user experience.

In the Hosted VDI test environment, session launchers were used to simulate multiple users making a direct connection to the Desktone 5.3.2 connection server through a VMware PCoIP protocol connection.

User Workload Simulation – LoginVSI from Login Consultants

One of the most critical factors of validating a Deskton 5.3.2 deployment is identifying a real-world user workload that is easy for customers to replicate and standardized across platforms to allow customers to realistically test the impact of a variety of worker tasks. To accurately represent a real-world user workload, a third-party tool from Login Consultants was used throughout the Hosted VDI testing.

The tool has the benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for individual desktop throughout large scale testing, including login storms.

The Virtual Session Indexer (Login Consultants' Login VSI 3.6) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and hence allows customers to easily replicate the testing results in their environment.



Note

In this testing, the tool was used to benchmark in a VDI environment only.

Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.

Login VSI simulates a medium workload user (also known as knowledge worker) running generic applications such as: Microsoft Office 2007 or 2010, Internet Explorer 8 including a Flash video applet and Adobe Acrobat Reader (Note: For the purposes of this test, applications were installed locally).

Like real users, the scripted Login VSI session will leave multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. This workload emulated a medium knowledge working using Office, IE, printing and PDF viewing.

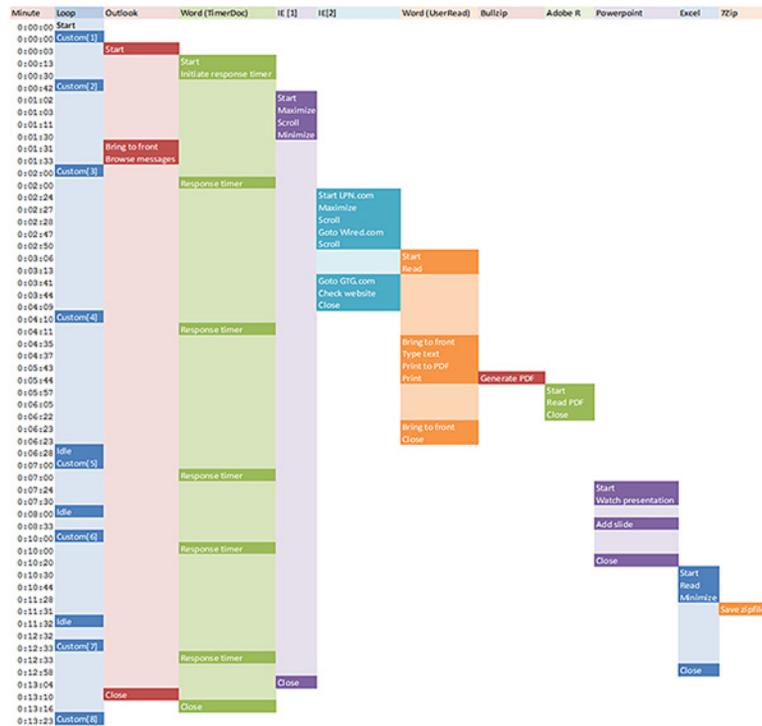
- Once a session has been started the medium workload will repeat every 12 minutes.
- During each loop the response time is measured every 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The type rate is 160ms for each character.
- Approximately two minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007/2010, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and gettheglass.com.
- 480 p Flash application.
- Word 2007/2010, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007/2010, a very large randomized sheet is opened.
- PowerPoint 2007/2010, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

A graphical representation of the medium workload is shown below.

Figure 15 Graphical overview



You can obtain additional information on Login VSI from <http://www.loginvsi.com>.

Testing Procedure

The following protocol was used for each test cycle in this study to insure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

All virtual machines were shut down utilizing VMware vCenter. All Launchers for the test were also shut down from VMware vCenter. Then they were started until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in thirty minutes. Additionally, Cisco requires all sessions started to become active within two minutes after the session is launched.

For each of the three consecutive runs on single blade (155 or 215 User) the process detailed below was followed:

1. Time 0:00:00 Start ESXtop Logging on VDI host blade used in test run.
2. Time 0:05 Power-on desktops with the Deskton Tenant Appliance.
3. Time 0:35 All desktops started on a single blade.
4. Time 0:50 Start Login VSI 3.6 Test utilizing 11 Launchers.

5. Time 1:20 Launch desktop sessions.
6. Time 1:22 Desktop sessions become active.
7. Time 1:35 Login VSI Test Ends.
8. Time 1:50 Logoff desktop sessions.
9. Time 2:00 Terminate all logging.

Success Criteria

There were multiple metrics that were captured during each test run, but the success criteria for determining the maximum workload per blade maintaining acceptable end-user experience for a single test run was based on the key metric, VSI Max. The Login VSI Max evaluates the user response time during increasing user load and assesses the point at which end-user experience becomes unacceptable.

Login VSI Max

VSI Max represents the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users as indicated during the workload execution. The user response time has a threshold of 4000ms and all users response times are expected to be less than 4000ms in order to assume that the user interaction with the virtual desktop is at a functional level. VSI Max is reached when the response times reaches or exceeds 4000ms for 6 consecutive occurrences. If VSI Max is reached, that indicates the point at which the user experience has significantly degraded. The response time is generally an indicator of the host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.



Note

In the prior version of Login VSI, the threshold for response time was 2000ms. The workloads and the analysis have been upgraded in Login VSI 3 to make the testing more aligned to real-world use. In the medium workload in Login VSI 3.0, a CPU intensive 480p flash movie is incorporated in each test loop. In general, the redesigned workload would result in an approximate 20% decrease in the number of users passing the test versus Login VSI 2.0 on the same server and storage hardware.

Calculating VSIMax

Typically the desktop workload is scripted in a 12-14 minute loop when a simulated Login VSI user is logged on. After the loop is finished it will restart automatically. Within each loop the response times of seven specific operations is measured in a regular interval: six times in within each loop. The response times if these seven operations are used to establish VSIMax. The seven operations from which the response times are measured are:

- Copy new document from the document pool in the home drive
 - This operation will refresh a new document to be used for measuring the response time. This activity is mostly a file-system operation.
- Starting Microsoft Word with a document.
 - This operation will measure the responsiveness of the Operating System and the file system. Microsoft Word is started and loaded into memory; also the new document is automatically loaded into Microsoft Word. When the disk I/O is extensive or even saturated, this will impact the file open dialogue considerably.

- Starting the “File Open” dialogue.
 - This operation is handled for small part by Word and a large part by the operating system. The file open dialogue uses generic subsystems and interface components of the OS. The OS provides the contents of this dialogue.
- Starting Notepad.
 - This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe itself through execution. This operation seems instant from an end-user’s point of view.
- Starting the “Print” dialogue.
 - This operation is handled for a large part by the OS subsystems, as the print dialogue is provided by the OS. This dialogue loads the print-subsystem and the drivers of the selected printer. As a result, this dialogue is also dependent on disk performance.
- Starting the “Search and Replace” dialogue.
 - This operation is handled within the application completely; the presentation of the dialogue is almost instant. Serious bottlenecks on application level will impact the speed of this dialogue.
- Compress the document into a zip file with 7-zip command line.
 - This operation is handled by the command line version of 7-zip. The compression will very briefly spike CPU and disk I/O.

These measured operations with Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations are consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

With Login VSI 3.0 and later it is now possible to choose between ‘VSImax Classic’ and ‘VSImax Dynamic’ results analysis. For these tests, VSImax Dynamic analysis was utilized.

VSIMax Dynamic

VSImax Dynamic is calculated when the response times are consistently above a certain threshold. However, this threshold is now dynamically calculated on the baseline response time of the test.

Five individual measurements are weighted to better support this approach:

- Copy new doc from the document pool in the home drive: 100%
- Microsoft Word with a document: 33.3%
- Starting the “File Open” dialogue: 100%
- Starting “Notepad”: 300%
- Starting the “Print” dialogue: 200%
- Starting the “Search and Replace” dialogue: 400%
- Compress the document into a zip file with 7-zip command line 200%
- A sample of the VSImax Dynamic response time calculation is displayed below:

Figure 16 VSIMax Dynamic Response Time

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	33.3%	467
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	300%	150
Print Dialogue (PRINT)	220	200%	440
Replace Dialogue (FIND)	10	400%	40
Zip documents (ZIP)	130	200%	230
VSIMax Dynamic Response Time			1837

301282

Then the average VSIMax response time is calculated based on the amount of active Login VSI users logged on to the system. For this the average VSIMax response times need to consistently higher than a dynamically calculated threshold.

To determine this dynamic threshold, first the average baseline response time is calculated. This is done by averaging the baseline response time of the first 15 Login VSI users on the system.

The formula for the dynamic threshold is: Avg. Baseline Response Time x 125% + 3000. As a result, when the baseline response time is 1800, the VSIMax threshold will now be $1800 \times 125\% + 3000 = 5250\text{ms}$.

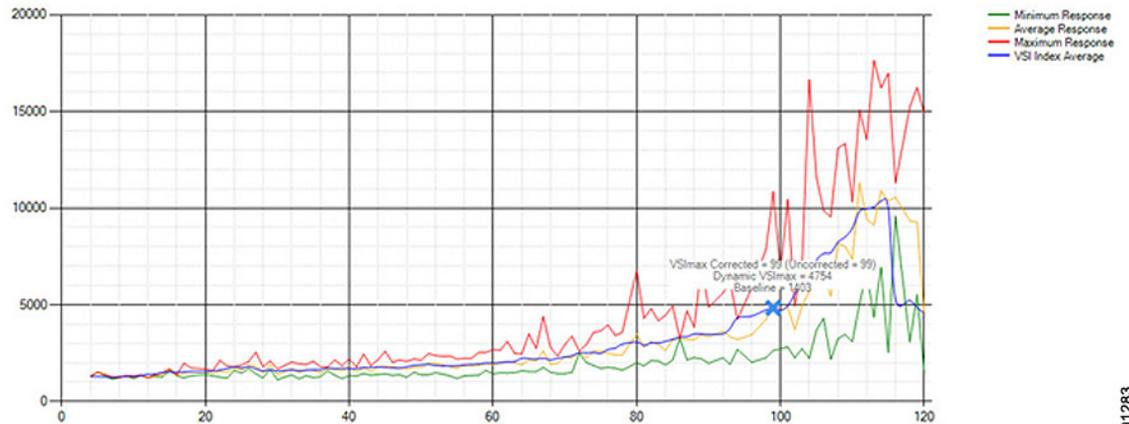
Especially when application virtualization is used, the baseline response time can wildly vary per vendor and streaming strategy. Therefore it is recommend to use VSIMax Dynamic when comparisons are made with application virtualization or anti-virus agents. The resulting VSIMax Dynamic scores are aligned again with saturation on a CPU, Memory or Disk level, also when the baseline response time are relatively high.

Determining VSIMax

The Login VSI analyzer will automatically identify the VSIMax. In the example below the VSIMax is 98. The analyzer will automatically determine “stuck sessions” and correct the final VSIMax score.

- Vertical axis: Response Time in milliseconds
- Horizontal axis: Total Active Sessions

Figure 17 Sample Login VSI Analyzer Graphic Output



301283

- Red line: Maximum Response (worst response time of an individual measurement within a single session)
- Orange line: Average Response Time within for each level of active sessions
- Blue line: the VSImax average.
- Green line: Minimum Response (best response time of an individual measurement within a single session)

In testing, the total number of users in the test run had to login, become active and run at least one test loop and log out automatically without reaching the VSI Max to be considered a success.



Note

A technical issue was discovered with the VSIMax dynamic calculation in testing on Cisco UCS B200 M3 blade servers where the VSIMax Dynamic was not reached during extreme conditions. Working with Login Consultants, a methodology was devised to validate the testing without reaching VSIMax Dynamic until such time as a new calculation is available.

The Login VSI “pass” criteria, accepted by Login Consultants for this testing, is as follows:

- Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, Memory utilization, Storage utilization and Network utilization.
- Will utilize Login VSI to launch version 3.6 medium workloads, including flash.
- Number of Launched Sessions must equal Active Sessions within two minutes of the last session launched in a test.
- The Deskton Enterprise Center will be monitored throughout the steady state to insure that:
 - All running sessions report in use throughout the steady state
 - No sessions move to unregistered or agent not available state at any time during Steady State
- Within 20 minutes of the end of the test, all sessions on all Launchers must have logged out automatically and the Login VSI Agent must have shut down.
- The Cisco Validated Design will be published with the process recommended above and will note that VSIMax dynamic was not reached during testing due to a technical issue with the analyzer formula that calculates VSIMax.

6 VDI Test Results

The purpose of this testing is to provide the data needed to validate Deskstone 5.3.2 infrastructure in an end-to-end Cisco environment.

Server scalability based on desktops supported in a single server is very useful to determine the total number of servers needed for a deployment. Different CPU, memory, storage and network metrics can be used to estimate sizing guidelines for necessary data center components and overall deployment. The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of Deskstone 5.3.2 with VMware vSphere 5.1.

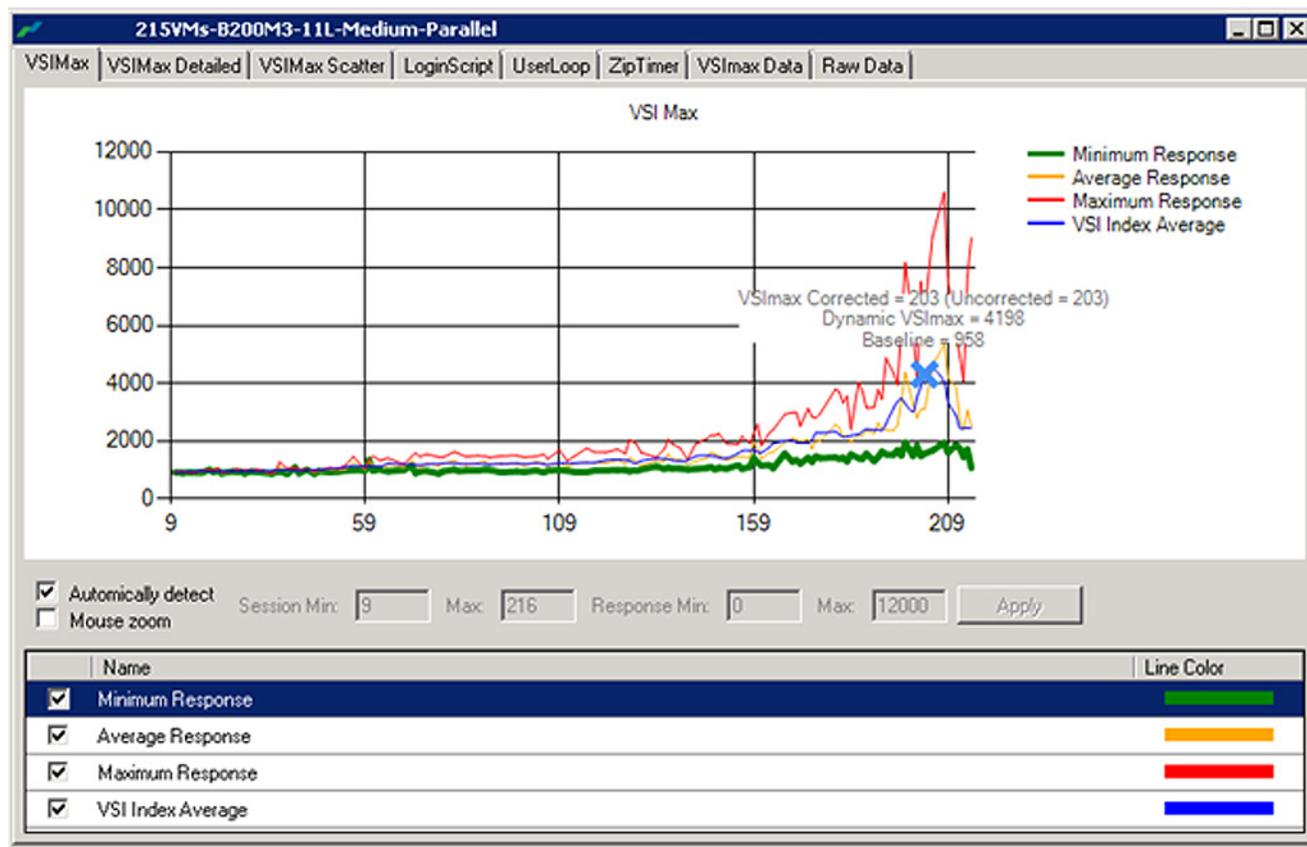
The results provided in this section are based on the testing done on a Cisco UCS B200 M3 blade server running ESXi 5.1 hypervisor managed by vCenter 5.1 to host Deskstone automated pools with static assignment and provision Windows 7 SP1 full virtual desktops using a NFS storage system.

Two test sequences, each containing three consecutive test runs generating the same result, were performed to establish single server performance.

Login VSIMax Score

One of the stress tests on a single blade server was conducted to establish the official Login VSI Max Score. Test results show that 215 Medium Workload (with flash) Windows 7 SP1 sessions on a single server achieved a Login VSI Max score of 203. The Login VSI score was achieved on three consecutive runs and is shown in [Figure 18](#) below.

Figure 18 Login VSIMax Reached: 203 Users



301284

Single Blade Maximum Recommended Workload

This section details the results from the Deskstone 5.3.2 Hosted VDI single blade server validation testing. The primary success criteria used to validate the overall success of the test cycle is an output chart from Login Consultants' VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)



Note

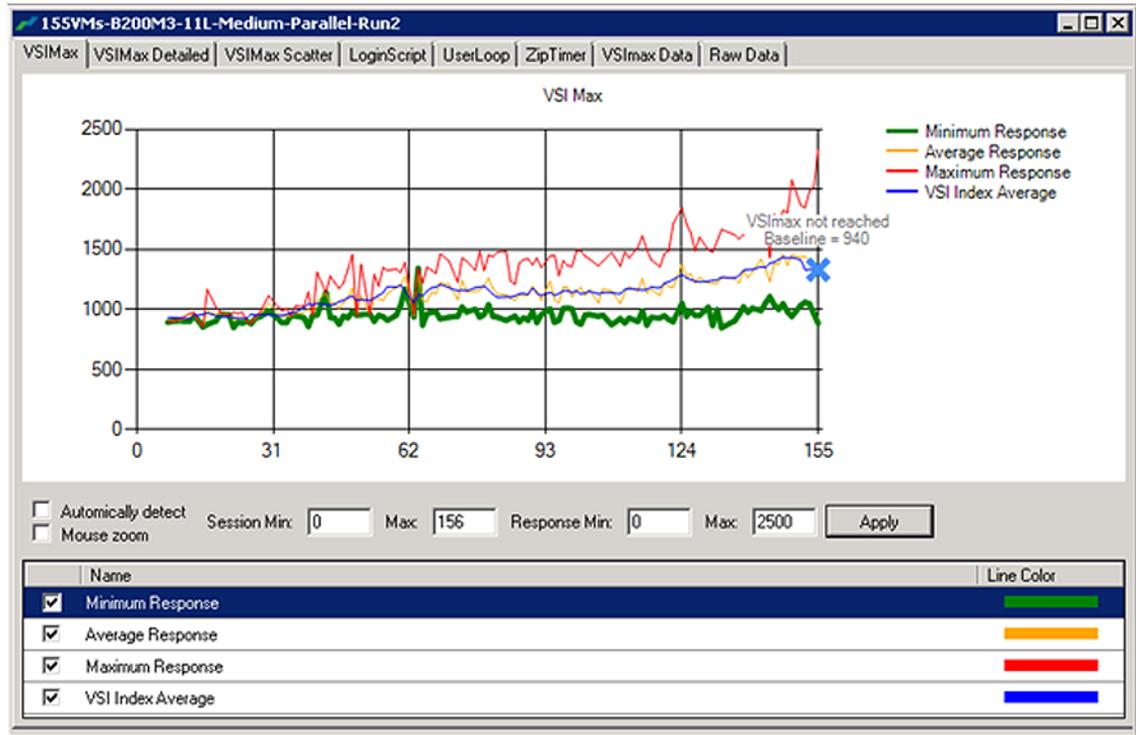
VSIMax Dynamic in testing was not reached due to a technical issue with the analyzer formula that calculates VSIMax. See section [Determining VSIMax](#) for a discussion of this issue.

A single server test using approximately 24% lower user density than prescribed by the Login VSI Max was executed to achieve a successful pass of the test with server hardware performance in a realistic range. The recommended maximum load for Deskstone 5.3.2 Hosted Virtual Desktops on a Cisco UCS B200 M3 blade server running the Login VSI Medium workload is 155 desktops given adequate storage capability and enough CPU resources. CPU utilization limit of 90% was a consideration to determine the maximum virtual machine density per blade.

Additionally, graphs detailing the CPU, Memory and network utilization during boot phase and peak session loads are also presented.

The charts below present the recommended maximum Login VSI Medium workload loading on a single blade server. The maximum recommended workload for Desktope 5.3.2 Windows 7 32-bit virtual desktops is 155 per B200 M3 blade server. This chart shows that the average and maximum application response times are below 1.5secs and 2.5secs respectively.

Figure 19 155 Desktope 5.3.2 Desktop Sessions on VMware ESXi 5.1 below 2500 ms



Boot Phase Performance Results

The following graphs detail CPU, Memory, Disk and Network performance on a single Cisco UCS B200-M3 blade server during the boot phase collected by esxtop polling data every 10 secs.

The first two charts show the CPU and Core utilizations well below 35% during the boot phase.

The third chart is the memory utilization chart showing the memory allocated to the virtual desktops as they boot up. Because enough memory is available, 1.5GB of RAM was assigned to each VM. A 1GB memory reservation for each VM was configured to minimize or prevent the hypervisor from swapping.

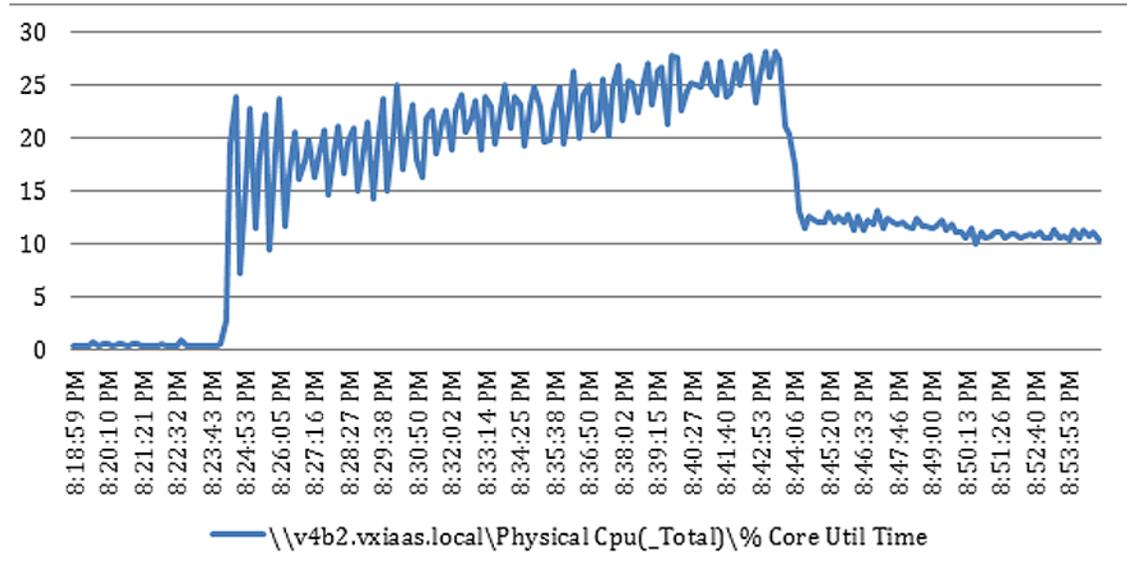
Subsequent charts show the storage performance, in terms of read and write I/O load on the storage system generated by 155 VMs on a single server during the boot phase. The I/O Bandwidth data chart shows the network bandwidth utilization associated with the storage traffic.



Note

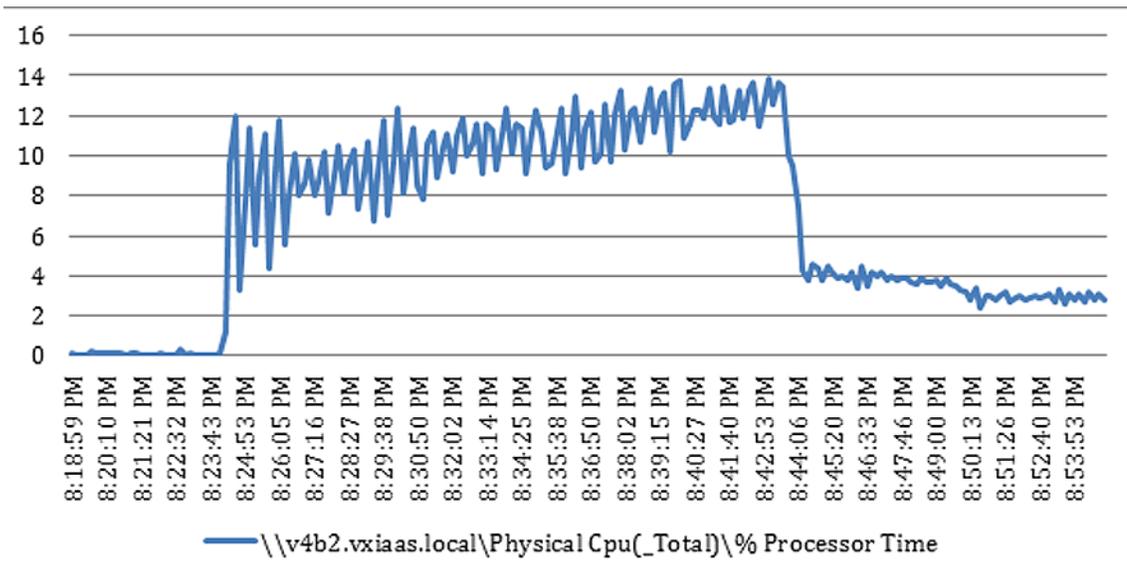
Some of the charts are in Mbytes/sec instead of Mbits/sec. and can be a starting point for estimating the bandwidth needs in the data center.

Figure 20 155 Users Single Cisco UCS B200 M3 CPU Core Utilization - Boot Phase



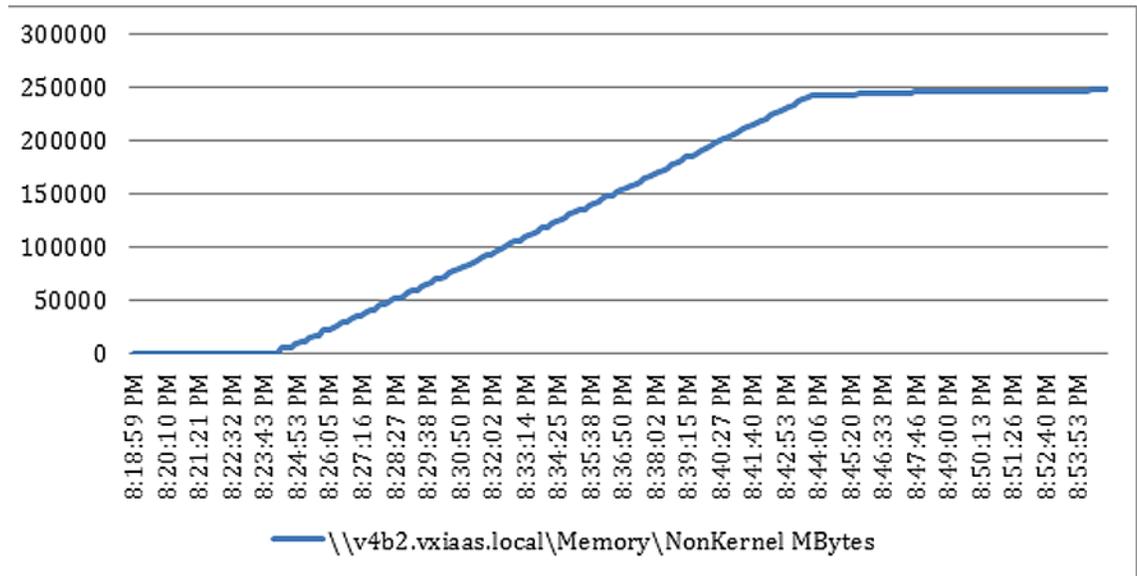
301286

Figure 21 155 Users Single Cisco UCS B200 M3 Processor Time - Boot Phase



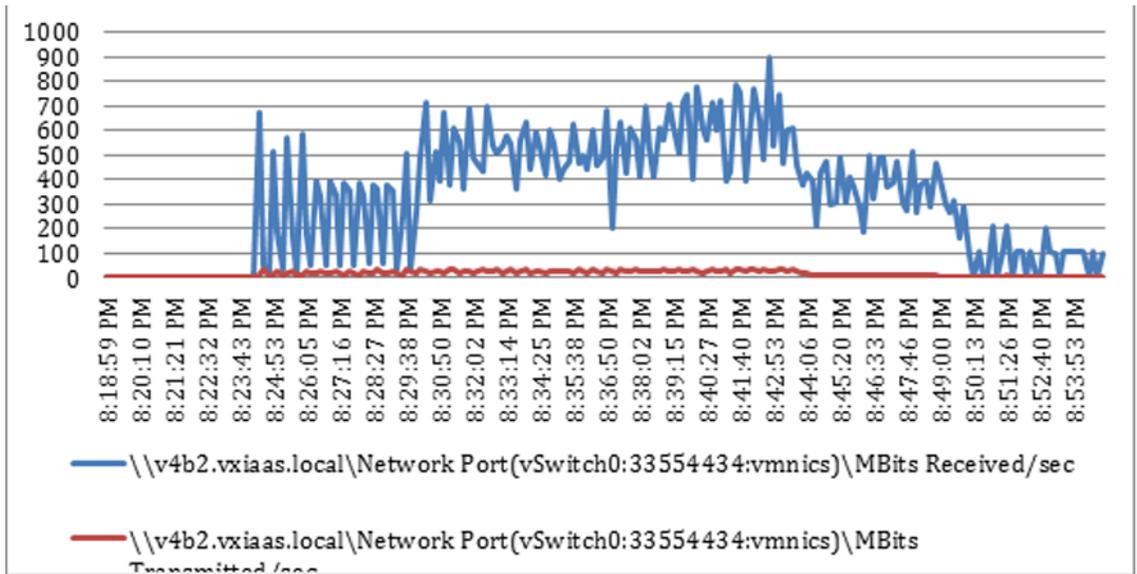
301287

Figure 22 155 Users Single Cisco UCS B200 M3 NonKernel Memory – Boot Phase



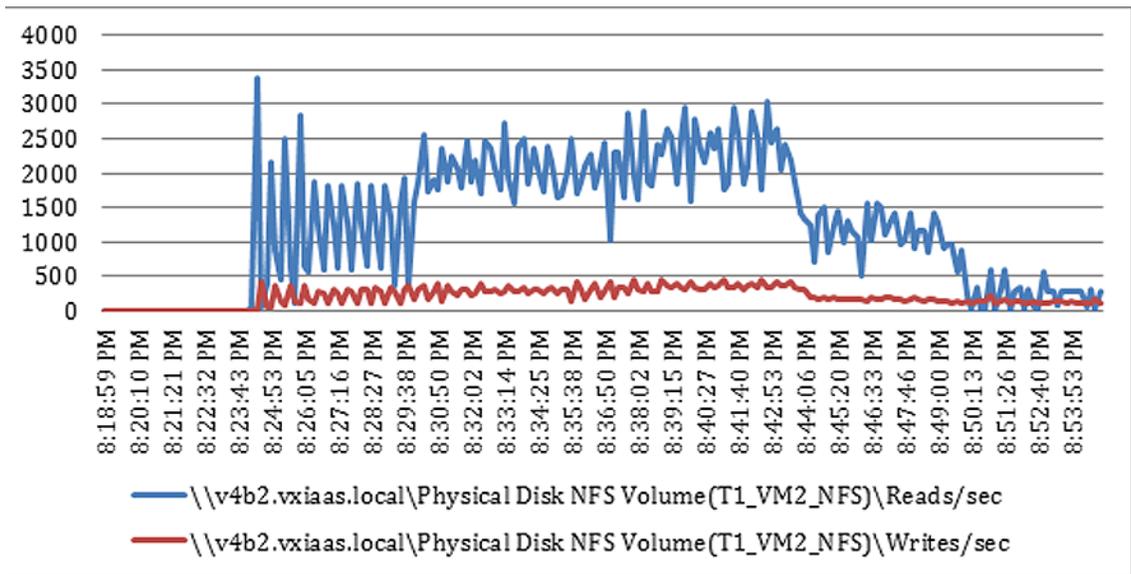
301288

Figure 23 155 Users Single Cisco UCS B200 M3 Cisco VIC1240 MLOM Network Adapter Mbps Receive/Transmit - Boot Phase



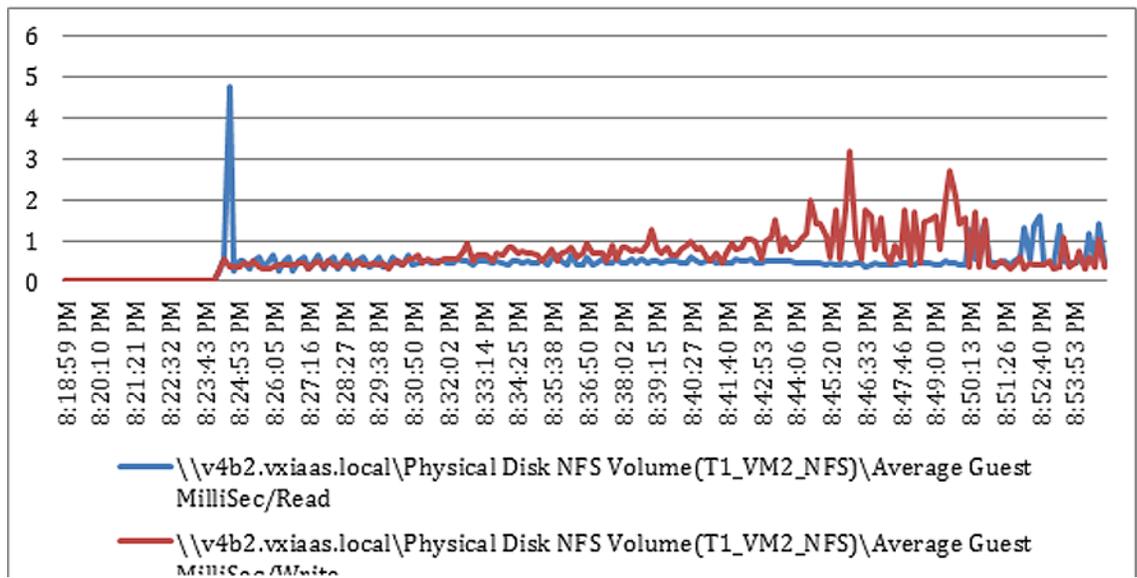
301289

Figure 24 155 Users Single Cisco UCS B200 M3 NFS Read/Write IOPS - Boot Phase



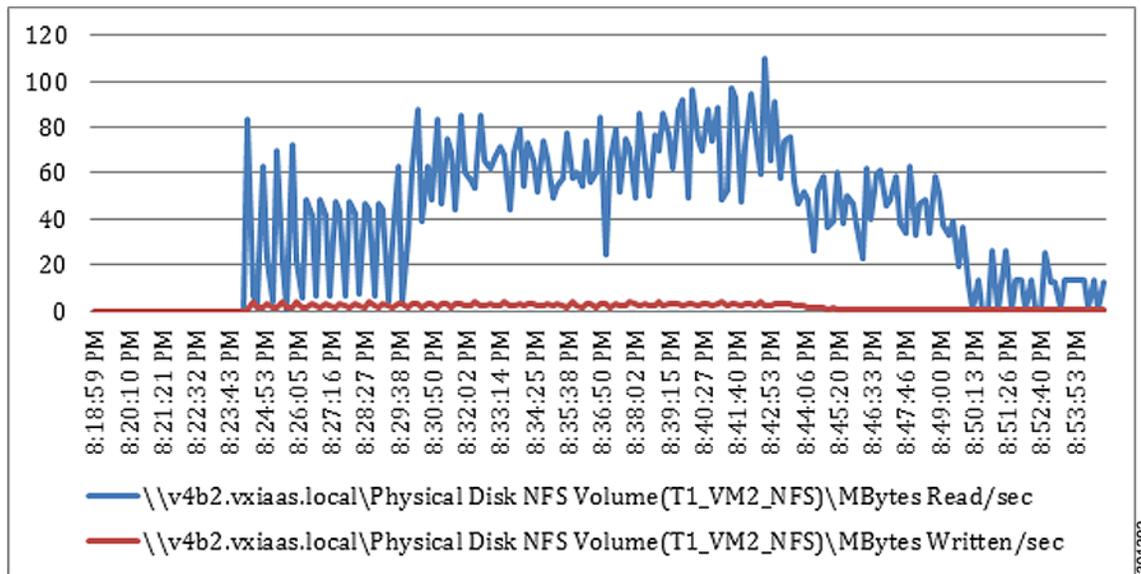
301290

Figure 25 155 Users Single Cisco UCS B200 M3 NFS Read/Write Latency - Boot Phase



301291

Figure 26 155 Users Single Cisco UCS B200 M3 NFS MBps Read/Write - Boot Phase



Test Phase Performance Results

The following graphs detail CPU, Memory, Disk and Network performance on the Single Cisco UCS B200-M3 blade servers during the test phase collected by esxtop polling data every 10 secs at the recommended maximum load of 155 Windows 7 virtual desktops.

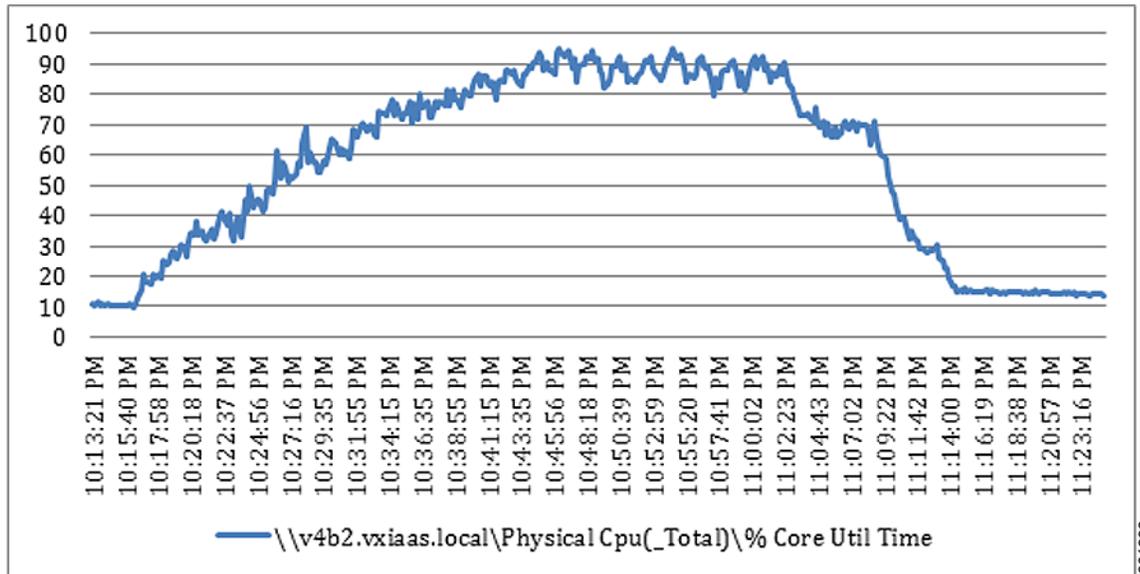
The first charts show the Core utilizations of ~90% during the steady workload phase. The second chart shows the CPU utilization time of ~55% which indicates that the CPU resources were not overcommitted.

The third chart is the memory utilization chart showing the memory allocated to the virtual desktops during the entire workload. Because enough memory is available, 1.5GB of RAM was assigned to each virtual machine. A 1GB memory reservation for each virtual machine was configured to minimize or prevent the hypervisor from swapping.

The next few charts show storage and network performance, in terms of read and write I/O load on the network and the storage system generated 155 VMs on a single server during the entire workload. The I/O Bandwidth data chart shows the network bandwidth utilization associated with the storage traffic.

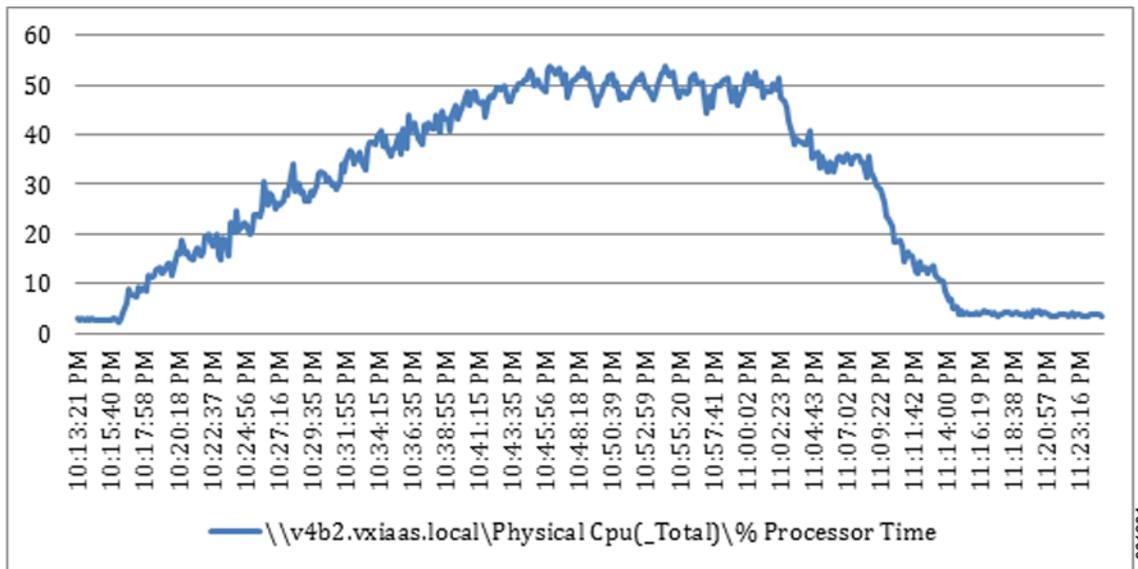
NOTE: Some of the charts are in Mbytes/sec instead of Mbits/sec. and can be a starting point for estimating the bandwidth needs in the data center.

Figure 27 155 Users Single Cisco UCS B200 M3 CPU Core Utilization - Test Phase



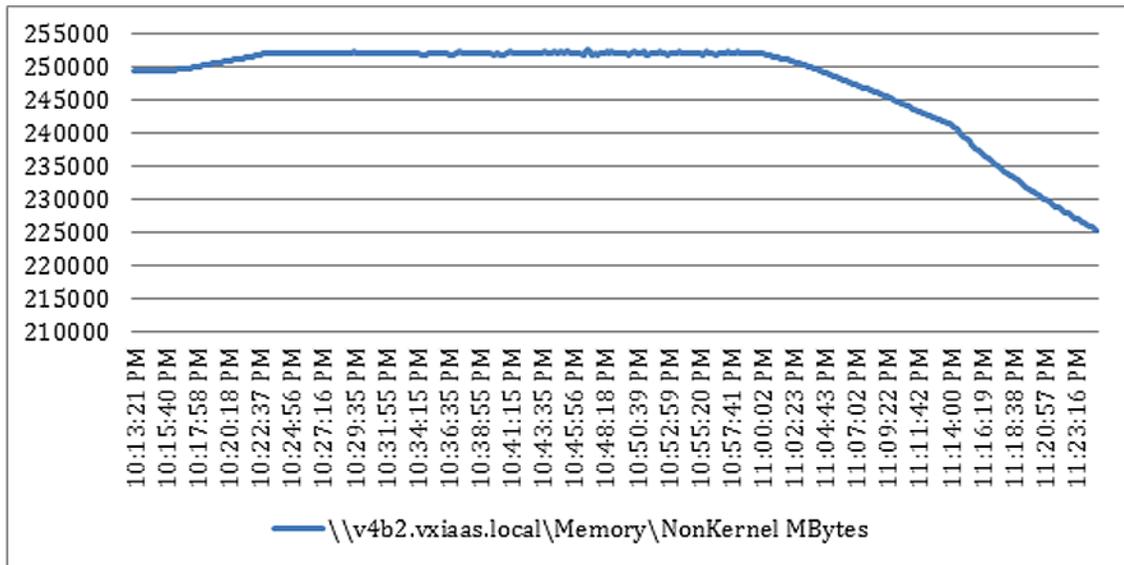
301293

Figure 28 155 Users Single Cisco UCS B200 M3 CPU Processor Time - Test Phase



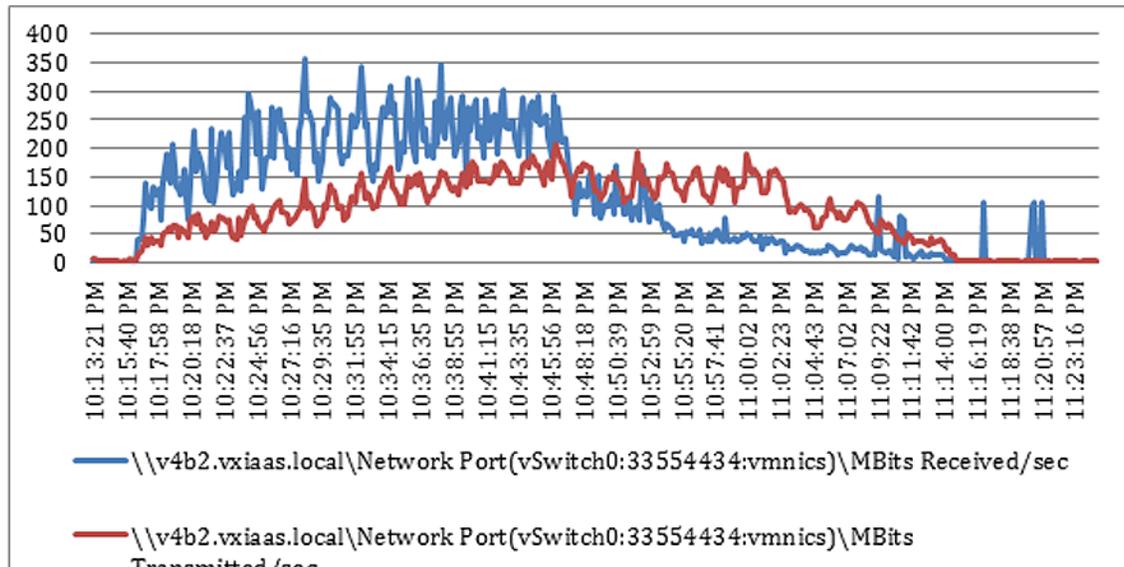
301294

Figure 29 155 Users Single Cisco UCS B200 M3 NonKernel Memory - Test Phase



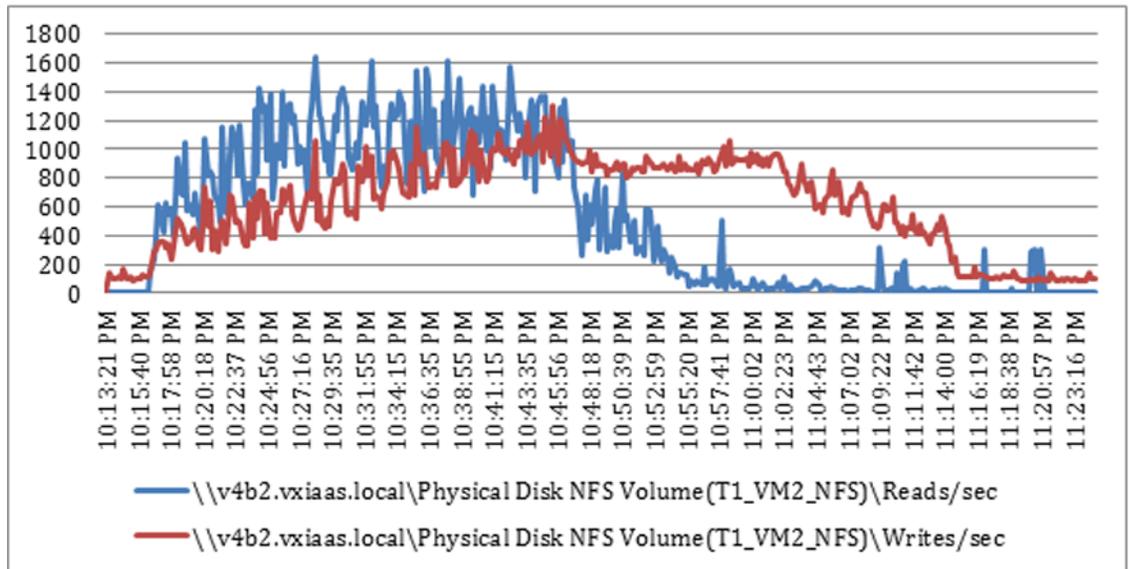
301296

Figure 30 155 Users Single Cisco UCS B200 M3 Cisco VIC1240 MLOM Network Adapter Mbps Receive/Transmit - Test Phase



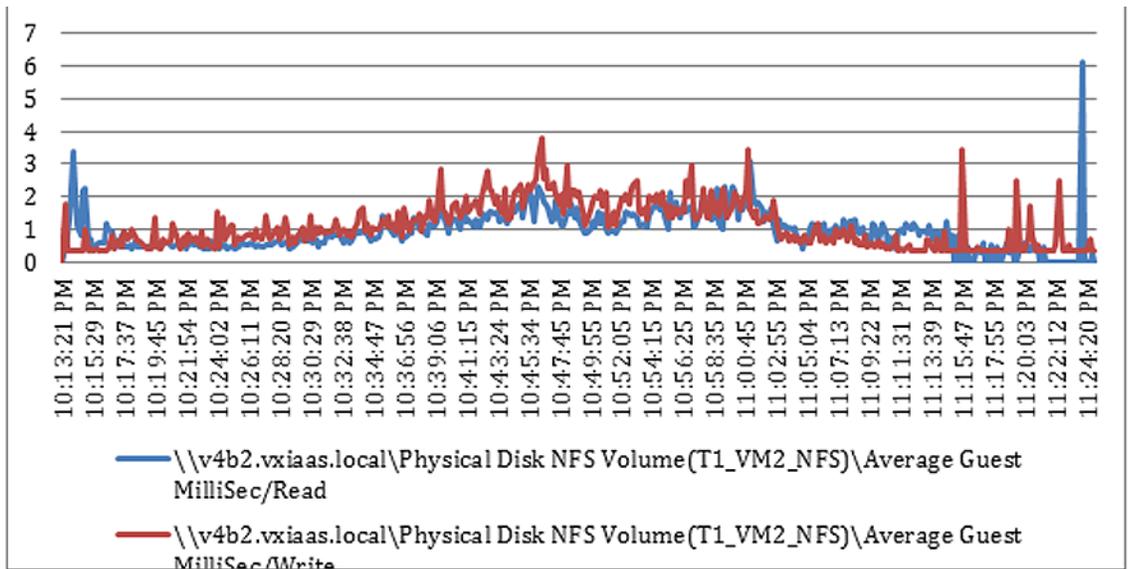
301296

Figure 31 155 Users Single Cisco UCS B200 M3 NFS Read/Write IOPS - Test Phase



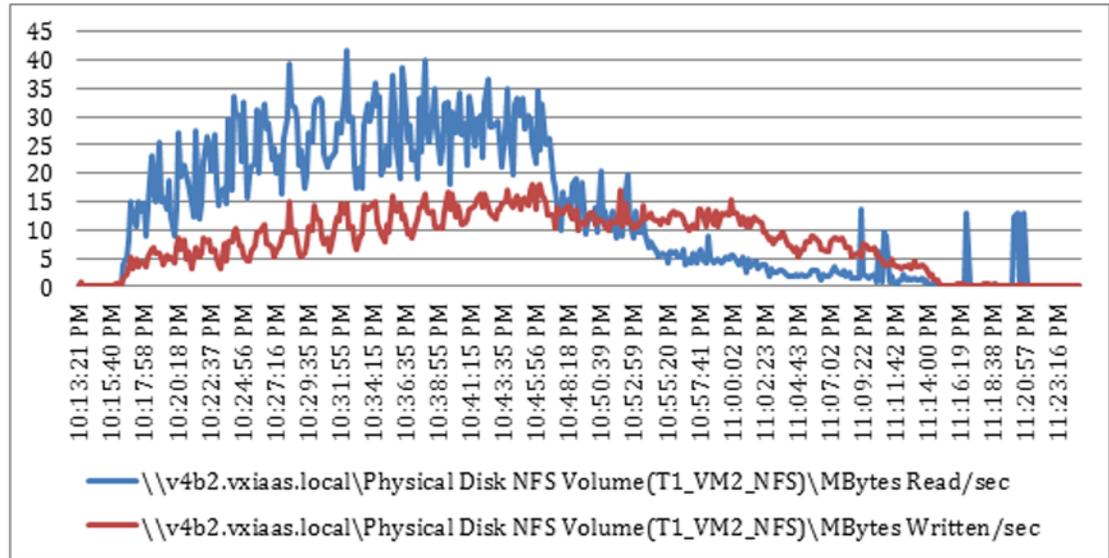
301297

Figure 32 155 Users Single Cisco UCS B200 M3 NFS Read/Write Latency - Test Phase



301298

Figure 33 155 Users Single Cisco UCS B200 M3 NFS MBps Read/Write - Test Phase



7 Appendix

Cisco Nexus 5548 configurations

```

!Time: Sat Jun  1 11:53:31 2013
version 5.2(1)N1(2a)
hostname VXIaaS-N5K-1
no feature telnet
cfs eth distribute
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature fex
username admin password 5 $1$yZiJ2hxG$u0ooTIZcFlu4hbEkhFKlz.  role network-admin
no password strength-check
banner motd #Nexus 5000 Switch
ip domain-lookup
ip access-list BULK-DATA
  10 permit tcp any eq 32111 any
ip access-list CALL-SIGNALING
  10 permit tcp any any eq 2748
  20 permit tcp any any eq 5060
  30 permit tcp any any eq 2000
ip access-list MULTIMEDIA-STREAMING
  10 permit tcp any eq 9427 any
ip access-list TRANSACTIONAL-DATA
  10 permit tcp any eq 3389 any
  20 permit udp any eq 50002 any
  30 permit tcp any eq 50002 any
  40 permit udp any eq 4172 any
  50 permit tcp any eq 4172 any
  60 permit tcp any eq 1494 any
class-map type qos class-fcoe
class-map type qos match-any BULK-DATA
  match access-group name BULK-DATA
class-map type qos match-any CALL-SIGNALING
  match access-group name CALL-SIGNALING
class-map type qos match-any TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-DATA
class-map type qos match-any MULTIMEDIA-STREAMING
  match access-group name MULTIMEDIA-STREAMING
class-map type queuing BULK
class-map type queuing BULK-DATA
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos HVD-ACCESS-PORT
  class CALL-SIGNALING
    set dscp 24
  class MULTIMEDIA-STREAMING
    set dscp 26
  class TRANSACTIONAL-DATA
    set dscp 18
  class BULK-DATA
    set dscp 10
  class class-default
policy-map type queuing UPSTREAM-PORT
  class type queuing class-default

```

```

class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos
  service-policy type network-qos jumbo
policy-map type control-plane copp-system-policy-customized
  class copp-system-class-default
    police cir 2048 kbps bc 6400000 bytes
hardware profile tcam feature interface-qos limit 50
fex 110
  pinning max-links 1
  description "FEX0110"
  diagnostic bootup level complete
snmp-server user admin network-admin auth md5 0x5b310d304479611ca2ebf839f5c267d7 priv
0x5b310d304479611ca2ebf839f5c267d7 localizedkey
ntp server 10.8.79.254 prefer
vrf context management
  ip route 0.0.0.0/0 10.8.66.1
vlan configuration 731
  service-policy type qos input HVD-ACCESS-PORT
vlan 1
vlan 505
  name VXIaaS-DH-Machines
vlan 711
  name Tenant-1-VM-A
vlan 721
  name Tenant-2-VM-A
vlan 731
  name Tenant-3-VM-A
vlan 741
  name Tenant-4-VM-A
vlan 751
  name Tenant-5-VM-A
vlan 761
  name Tenant-6-VM-A
vlan 844
  name HVXI -NFS

vlan 864
  name VXIaaS-SP-Infra
vlan 865
  name VXIaaS-SP-DH
vlan 867
  name VXIaaS-SP-Hypervisor-MGMT
vlan 868
  name VXIaaS-SP-vMotion
vlan 941
  name Tenant-4-vWAAS
vlan 980
  name Desktone-L2-Link-Local
vpc domain 500
  peer-keepalive destination 10.8.66.9 source 10.8.66.8
port-profile default max-ports 512
interface Vlan1

```

```

interface port-channel1
  description VXIaaS-N5K VPC Peer
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  spanning-tree port type network
  speed 10000
  vpc peer-link
interface port-channel55
  description VXIaaS-N7K Pair
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
  spanning-tree port type normal
  speed 10000
  vpc 55
interface port-channel60
  description vxiaas-n1010-a
  switchport mode trunk
  switchport trunk native vlan 867
  switchport trunk allowed vlan 711,867
  spanning-tree port type edge
interface port-channel65
  description vxiaas-n1010-b
  switchport mode trunk
  switchport trunk native vlan 867
  switchport trunk allowed vlan 711,867
  spanning-tree port type edge
interface port-channel100
  description VXIaaS-6100-1-a
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000
  vpc 100
interface port-channel110
  switchport mode fex-fabric
  fex associate 110
  spanning-tree port type edge trunk
  vpc 110
interface port-channel150
  description VXIaaS-6100-2-a
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000
  vpc 150
interface port-channel151
  description VXIaaS-6100-2-b
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000
  vpc 151
interface port-channel200
  description VXIaaS-6100-1-b
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000
  vpc 200

```

```

interface Ethernet1/1
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/2
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/3
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/4
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
  description VXIaaS-N7K
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 55 mode active
interface Ethernet1/14
  description VXIaaS-N7K
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 55 mode active
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
  description VXIaaS-6100-2-a
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
  channel-group 150 mode active
interface Ethernet1/26
  description VXIaaS-6100-2-a
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
  channel-group 150 mode active
interface Ethernet1/27
  description VXIaaS-6100-2-b

```

```

switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
channel-group 151 mode active
interface Ethernet1/28
description VXIaaS-6100-2-b
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
channel-group 151 mode active
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
description VXIaaS-6100-1-a
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
channel-group 100 mode active
interface Ethernet1/32
description VXIaaS-6100-1-b
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
channel-group 200 mode active
interface Ethernet2/1
interface Ethernet2/2
interface Ethernet2/3
interface Ethernet2/4
interface Ethernet2/5
interface Ethernet2/6
interface Ethernet2/7
interface Ethernet2/8
interface Ethernet2/9
interface Ethernet2/10
interface Ethernet2/11
interface Ethernet2/12
interface Ethernet2/13
interface Ethernet2/14
interface Ethernet2/15
switchport mode fex-fabric
fex associate 110
channel-group 110
interface Ethernet2/16
switchport mode fex-fabric
fex associate 110
channel-group 110
interface mgmt0
ip address 10.8.66.8/24
interface Ethernet110/1/1
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/2
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/3
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/4
description vxiaas-n1010-a

```

```
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/5
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/6
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/7
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/8
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/9
interface Ethernet110/1/10
interface Ethernet110/1/11
interface Ethernet110/1/12
interface Ethernet110/1/13
interface Ethernet110/1/14
interface Ethernet110/1/15
interface Ethernet110/1/16
interface Ethernet110/1/17
interface Ethernet110/1/18
interface Ethernet110/1/19
interface Ethernet110/1/20
interface Ethernet110/1/21
interface Ethernet110/1/22
interface Ethernet110/1/23
interface Ethernet110/1/24
interface Ethernet110/1/25
interface Ethernet110/1/26
interface Ethernet110/1/27
interface Ethernet110/1/28
interface Ethernet110/1/29
interface Ethernet110/1/30
interface Ethernet110/1/31
interface Ethernet110/1/32
interface Ethernet110/1/33
interface Ethernet110/1/34
interface Ethernet110/1/35
interface Ethernet110/1/36
interface Ethernet110/1/37
interface Ethernet110/1/38
interface Ethernet110/1/39
interface Ethernet110/1/40
interface Ethernet110/1/41
interface Ethernet110/1/42
interface Ethernet110/1/43
interface Ethernet110/1/44
```

```

interface Ethernet110/1/45
interface Ethernet110/1/46
interface Ethernet110/1/47
interface Ethernet110/1/48
clock timezone EST -5 0
clock summer-time EDT 2 Sun Mar 2:00 1 Sun Nov 2:00 60
line console
    exec-timeout 0
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.2a.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.2a.bin

!Time: Sat Jun  1 12:19:44 2013
version 5.2(1)N1(2a)
hostname VXIaaS-N5K-2
no feature telnet
cfs eth distribute
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature fex
username admin password 5 $1$vb55EyW$H21ygbp9047zkiTIjFkx60  role network-admin
no password strength-check
banner motd #Nexus 5000 Switch
ssh key rsa 2048
ip domain-lookup
ip access-list BULK-DATA
    10 permit tcp any any eq 2748
ip access-list CALL-SIGNALING
    10 permit tcp any any eq 2748
    20 permit tcp any any eq 5060
    30 permit tcp any any eq 2000
ip access-list MULTIMEDIA-STREAMING
    10 permit tcp any eq 9427 any
ip access-list TRANSACTIONAL-DATA
    10 permit tcp any eq 3389 any
    20 permit udp any eq 50002 any
    30 permit tcp any eq 50002 any
    40 permit udp any eq 4172 any
    50 permit tcp any eq 4172 any
    60 permit tcp any eq 1494 any
class-map type qos class-fcoe
class-map type qos match-any BULK-DATA
    match access-group name BULK-DATA
class-map type qos match-any CALL-SIGNALING
    match access-group name CALL-SIGNALING
class-map type qos match-any TRANSACTIONAL-DATA
    match access-group name TRANSACTIONAL-DATA
class-map type qos match-any MULTIMEDIA-STREAMING
    match access-group name MULTIMEDIA-STREAMING
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
policy-map type qos HVD-ACCESS-PORT
    class CALL-SIGNALING
        set dscp 24
    class MULTIMEDIA-STREAMING

```

```

        set dscp 26
    class TRANSACTIONAL-DATA
        set dscp 18
    class BULK-DATA
        set dscp 10
    class class-default
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
policy-map type network-qos jumbo
    class type network-qos class-fcoe
        pause no-drop
        mtu 2158
    class type network-qos class-default
        mtu 9216
        multicast-optimize
system qos
    service-policy type network-qos jumbo
policy-map type control-plane copp-system-policy-customized
    class copp-system-class-default
        police cir 2048 kbps bc 6400000 bytes
hardware profile tcam feature interface-qos limit 50
fex 110
    pinning max-links 1
    description "FEX0110"
    diagnostic bootup level complete
snmp-server user admin network-admin auth md5 0x5fad43c003bbc5cc41071a4c7ca92e5e priv
0x5fad43c003bbc5cc41071a4c7ca92e5e localizedkey
ntp server 10.8.79.254
vrf context management
    ip route 0.0.0.0/0 10.8.66.1
vlan configuration 731
    service-policy type qos input HVD-ACCESS-PORT
vlan 1
vlan 505
    name VXIaaS-DH-Machines
vlan 711
    name Tenant-1-VM-A
vlan 721
    name Tenant-2-VM-A
vlan 731
    name Tenant-3-VM-A
vlan 741
    name Tenant-4-VM-A
vlan 751
    name Tenant-5-VM-A
vlan 761
    name Tenant-6-VM-A
vlan 844
    name HVXI- NFS
vlan 864
    name VXIaaS-SP-Infra
vlan 865
    name VXIaaS-SP-DH
vlan 867
    name VXIaaS-SP-Hypervisor-MGMT
vlan 868
    name VXIaaS-SP-vMotion
vlan 941
    name Tenant-4-vWAAS
vlan 980

```

```

    name Desktone-L2-Link-Local
vpc domain 500
    peer-keepalive destination 10.8.66.8 source 10.8.66.9
    auto-recovery
port-profile default max-ports 512
interface Vlan1
interface port-channel1
    description VXIaaS-N5K VPC Peer
    switchport mode trunk
    switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
    spanning-tree port type network
    speed 10000
    vpc peer-link
interface port-channel55
    description VXIaaS-N7K Pair
    switchport mode trunk
    switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
    spanning-tree port type normal
    speed 10000
    vpc 55
interface port-channel60
    description vxiaas-n1010-a
    switchport mode trunk
    switchport trunk native vlan 867
    switchport trunk allowed vlan 711,867
    spanning-tree port type edge
interface port-channel65
    description vxiaas-n1010-b
    switchport mode trunk
    switchport trunk native vlan 867
    switchport trunk allowed vlan 711,867
    spanning-tree port type edge
interface port-channel100
    description VXIaaS-6100-1-a
    switchport mode trunk
    switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
    spanning-tree port type edge trunk
    spanning-tree bpdufilter enable
    speed 10000
    vpc 100
interface port-channel110
    switchport mode fex-fabric
    fex associate 110
    spanning-tree port type edge trunk
    vpc 110
interface port-channel150
    description VXIaaS-6100-2-a
    switchport mode trunk
    switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
    spanning-tree port type edge trunk
    spanning-tree bpdufilter enable
    speed 10000
    vpc 150
interface port-channel151
    description VXIaaS-6100-2-b
    switchport mode trunk
    switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
    spanning-tree port type edge trunk
    spanning-tree bpdufilter enable
    speed 10000
    vpc 151
interface port-channel200
    description VXIaaS-6100-1-b

```

```

switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
spanning-tree port type edge trunk
spanning-tree bpdupfilter enable
speed 10000
vpc 200
interface Ethernet1/1
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/2
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/3
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/4
  description VXIaaS-N5K-1
  switchport mode trunk
  switchport trunk allowed vlan
505,711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 1 mode active
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
  description VXIaaS-N7K
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 55 mode active
interface Ethernet1/14
  description VXIaaS-N7K
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865,867-868,941,980
  channel-group 55 mode active
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
  description VXIaaS-6100-2-a
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
  channel-group 150 mode active
interface Ethernet1/26

```

```

description VXIaaS-6100-2-a
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
channel-group 150 mode active
interface Ethernet1/27
description VXIaaS-6100-2-b
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
channel-group 151 mode active
interface Ethernet1/28
description VXIaaS-6100-2-b
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,980
channel-group 151 mode active
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
description VXIaaS-6100-1-a
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
channel-group 100 mode active
interface Ethernet1/32
description VXIaaS-6100-1-b
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864,867-868,941,980
channel-group 200 mode active
interface Ethernet2/1
interface Ethernet2/2
interface Ethernet2/3
interface Ethernet2/4
interface Ethernet2/5
interface Ethernet2/6
interface Ethernet2/7
interface Ethernet2/8
interface Ethernet2/9
interface Ethernet2/10
interface Ethernet2/11
interface Ethernet2/12
interface Ethernet2/13
interface Ethernet2/14
interface Ethernet2/15
switchport mode fex-fabric
fex associate 110
channel-group 110
interface Ethernet2/16
switchport mode fex-fabric
fex associate 110
channel-group 110
interface mgmt0
ip address 10.8.66.9/24
interface Ethernet110/1/1
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/2
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/3
description vxiaas-n1010-a

```

```
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/4
description vxiaas-n1010-a
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 60 mode active
interface Ethernet110/1/5
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/6
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/7
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/8
description vxiaas-n1010-b
switchport mode trunk
switchport trunk native vlan 867
switchport trunk allowed vlan 711,867
channel-group 65 mode active
interface Ethernet110/1/9
interface Ethernet110/1/10
interface Ethernet110/1/11
interface Ethernet110/1/12
interface Ethernet110/1/13
interface Ethernet110/1/14
interface Ethernet110/1/15
interface Ethernet110/1/16
interface Ethernet110/1/17
interface Ethernet110/1/18
interface Ethernet110/1/19
interface Ethernet110/1/20
interface Ethernet110/1/21
interface Ethernet110/1/22
interface Ethernet110/1/23
interface Ethernet110/1/24
interface Ethernet110/1/25
interface Ethernet110/1/26
interface Ethernet110/1/27
interface Ethernet110/1/28
interface Ethernet110/1/29
interface Ethernet110/1/30
interface Ethernet110/1/31
interface Ethernet110/1/32
interface Ethernet110/1/33
interface Ethernet110/1/34
interface Ethernet110/1/35
interface Ethernet110/1/36
interface Ethernet110/1/37
interface Ethernet110/1/38
```

```

interface Ethernet110/1/39
interface Ethernet110/1/40
interface Ethernet110/1/41
interface Ethernet110/1/42
interface Ethernet110/1/43
interface Ethernet110/1/44
interface Ethernet110/1/45
interface Ethernet110/1/46
interface Ethernet110/1/47
interface Ethernet110/1/48
clock timezone EST -5 0
clock summer-time EDT 2 Sun Mar 2:00 1 Sun Nov 2:00 60
line console
  exec-timeout 0
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.2a.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.2a.bin

```

Cisco Nexus 7009 configurations:

```

!Time: Sat Jun  1 12:22:38 2013
version 6.1(2)
hostname VXIaaS-DC-N7K-1
vdc VXIaaS-DC-N7K-1 id 1
  limit-resource module-type f2
  allocate interface Ethernet4/1-48
  allocate interface Ethernet5/1-48
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
feature telnet
cfs eth distribute
feature ospf
feature bgp
feature eigrp
feature interface-vlan
feature dotlx
feature hsrp
feature lacp
feature dhcp
feature vpc
feature wccp
logging level pixm 2
username admin password 5 $1$QAUkrln2$EgIAciPkN0CLqHbf7SmQf/  role network-admin
no password strength-check
ip domain-lookup
ip domain-name vxiaas.local
ip domain-list cisco.com
ip domain-list vxiaas.local
ip name-server 10.8.64.100
class-map type qos match-any VOICE
  match dscp 46
class-map type qos match-all BULK-DATA
  match dscp 10,12,14
class-map type qos match-any SCAVENGER
  match dscp 8
class-map type qos match-any CALL-SIGNALING

```

```

    match dscp 24
class-map type qos match-any NETWORK-CONTROL
    match dscp 48
class-map type qos match-any TRANSACTIONAL-DATA
    match dscp 18,20,22
class-map type qos match-any MULTIMEDIA-STREAMING
    match dscp 26,28,30
class-map type qos match-any MULTIMEDIA-CONFERENCING
    match dscp 34,36,38
class-map type queuing match-any BROADCAST-VIDEO
policy-map type queuing VPN-EDGE
copp profile strict
snmp-server user admin network-admin auth md5 0xde4b54725e78b89933e0602c1a882917 priv
0xde4b54725e78b89933e0602c1a882917 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.8.79.254
vlan 1,711,721,731,741,751,761,844,864-868,941,980
vrf context GT-Out
vrf context T1-In
vrf context T1-Out
vrf context T2-In
vrf context T2-Out
vrf context T3-In
vrf context T3-Out
vrf context T4-In
vrf context T4-Out
    ip wccp 61
    ip wccp 62
vrf context T5-In
vrf context T5-Out
vrf context management
    ip route 0.0.0.0/0 10.8.66.1
vlan 711
    name Tenant-1-VM-A
vlan 721
    name Tenant-2-VM-A
vlan 731
    name Tenant-3-VM-A
vlan 741
    name Tenant-4-VM-A
vlan 751
    name Tenant-5-VM-A
vlan 761
    name Tenant-6-VM-A
vlan 844
    name HVXI-NFS
vlan 864
    name VXIaaS-SP-Infra
vlan 865
    name VXIaaS-SP-DH
vlan 866
    name VXIaaS-OOB-MGMT
vlan 867
    name VXIaaS-SP-Hypervisor-MGMT
vlan 868
    name VXIaaS-SP-vMotion
vlan 941
    name Tenant-4-vWAAS
vlan 980
    name Desktone-L2-Link-Local

```

```

spanning-tree vlan 711,721,731,741,751,761,864-868,980 priority 4096
ip prefix-list BGP-DIRECT seq 5 permit 10.8.0.0/16 le 30

ip prefix-list COMMON-SUM seq 5 permit 10.10.0.0/17
ip prefix-list EIGRP-BGP seq 5 permit 10.10.0.0/17
route-map BGP-DIRECT permit 10
  match ip address prefix-list BGP-DIRECT
route-map COMMON-SUM permit 10
  match ip address prefix-list COMMON-SUM
route-map EIGRP-BGP permit 10
  match ip address prefix-list EIGRP-BGP
service dhcp
ip dhcp relay
vpc domain 42
  peer-switch
  role priority 100
  peer-keepalive destination 10.8.66.22 source 10.8.66.21
  peer-gateway
  auto-recovery
  ip arp synchronize
interface Vlan1
interface Vlan711
  vrf member T1-In
  no ip redirects
  ip address 10.9.104.2/21
  ip ospf passive-interface
  ip router ospf VXIaaS area 0.0.0.20
  hsrp version 2
  hsrp 711
    priority 110
    ip 10.9.104.1
  ip dhcp relay address 10.9.10.150
  ip dhcp relay address 10.9.10.151
  description Tenant 1 VMs
  no shutdown
  mtu 9216
interface Vlan721
  vrf member T2-In
  no ip redirects
  ip address 10.9.112.2/21
  ip ospf passive-interface
  ip router ospf VXIaaS area 0.0.0.20
  hsrp version 2
  hsrp 721
    preempt
    priority 110
    ip 10.9.112.1
  ip dhcp relay address 10.9.20.150
  ip dhcp relay address 10.9.20.151
  description Tenant 2 VMs
  no shutdown
  mtu 9216
interface Vlan731
  vrf member T3-In
  no ip redirects
  ip address 10.9.120.2/21
  ip ospf passive-interface
  ip router ospf VXIaaS area 0.0.0.20
  hsrp version 2
  hsrp 731
    preempt
    priority 110
    ip 10.9.120.1
  ip dhcp relay address 10.9.30.150

```

```
    ip dhcp relay address 10.9.30.151
    description Tenant 3 VMs
    no shutdown
    mtu 9216
interface Vlan741
    vrf member T4-In
    no ip redirects
    ip address 10.9.128.2/21
    ip ospf passive-interface
    ip router ospf VXIaaS area 0.0.0.20
    hsrp version 2
    hsrp 741
        preempt
        priority 110
        ip 10.9.128.1
    ip dhcp relay address 10.9.40.150
    ip dhcp relay address 10.9.40.151
    description Tenant 4 VMs
    no shutdown
    mtu 9216
interface Vlan751
    vrf member T5-In
    no ip redirects
    ip address 10.9.136.2/21
    ip ospf passive-interface
    ip router ospf VXIaaS area 0.0.0.20
    hsrp version 2
    hsrp 751
        preempt
        priority 110
        ip 10.9.136.1
    ip dhcp relay address 10.9.50.150
    ip dhcp relay address 10.9.50.151
    description Tenant 5 VMs
    no shutdown
    mtu 9216
interface Vlan864
    no ip redirects
    ip address 10.8.64.2/24
    ip router eigrp 42
    ip passive-interface eigrp 42
    hsrp version 2
    hsrp 864
        priority 110
        ip 10.8.64.1
    description VXIaaS-SP-Infra
    no shutdown
    mtu 9216
interface Vlan865
    no ip redirects
    ip address 10.8.65.2/24
    ip router eigrp 42
    ip passive-interface eigrp 42
    hsrp version 2
    hsrp 865
        priority 110
        ip 10.8.65.1
    description VXIaaS-SP-DH
    no shutdown
    mtu 9216
interface Vlan866
    no ip redirects
    ip address 10.8.66.2/24
    ip router eigrp 42
```

```

ip passive-interface eigrp 42
hsrp version 2
hsrp 866
  priority 110
  ip 10.8.66.1
description VXIaaS-OOB-MGMT
no shutdown
mtu 9216
interface Vlan867
no ip redirects
ip address 10.8.67.2/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 867
  priority 110
  ip 10.8.67.1
description VXIaaS-SP-Hypervisor-MGMT
no shutdown
mtu 9216
interface Vlan868
no ip redirects
ip address 10.8.68.2/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 868
  priority 110
  ip 10.8.68.1
description VXIaaS-SP-vMotion
no shutdown
mtu 9216
interface Vlan941
vrf member T4-Out
no ip redirects
ip address 10.9.192.26/29
ip ospf passive-interface
ip router ospf VXIaaS area 0.0.0.10
hsrp version 2
hsrp 941
  preempt
  priority 110
  ip 10.9.192.25
ip dhcp relay address 10.9.40.150
ip dhcp relay address 10.9.40.151
description Tenant 4 vWAAS
no shutdown
mtu 9216
interface port-channel1
description vPC Peer Link Between N7Ks
switchport
switchport mode trunk
switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
switchport trunk allowed vlan add 864-868,941,980
spanning-tree port type network
mtu 9216
vpc peer-link
interface port-channel2
description NON-vPC Peer Link Between N7Ks
mtu 9216
interface port-channel2.718
description Tenant 1 N7K-N7K Outside
mtu 9216

```

```
encapsulation dot1q 718
vrf member T1-Out
no ip redirects
ip address 10.8.78.25/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.10
no shutdown
interface port-channel2.719
description Tenant 1 N7K-N7K Inside
mtu 9216
encapsulation dot1q 719
vrf member T1-In
no ip redirects
ip address 10.8.78.29/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.20
no shutdown
interface port-channel2.728
description Tenant 2 N7K-N7K Outside
mtu 9216
encapsulation dot1q 728
vrf member T2-Out
no ip redirects
ip address 10.8.78.57/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.10
no shutdown
interface port-channel2.729
description Tenant 1 N7K-N7K Inside
mtu 9216
encapsulation dot1q 729
vrf member T2-In
no ip redirects
ip address 10.8.78.61/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.20
no shutdown
interface port-channel2.738
description Tenant 3 N7K-N7K Outside
mtu 9216
encapsulation dot1q 738
vrf member T3-Out
no ip redirects
ip address 10.8.78.89/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.10
no shutdown
interface port-channel2.739
description Tenant 1 N7K-N7K Inside
mtu 9216
encapsulation dot1q 739
vrf member T3-In
no ip redirects
ip address 10.8.78.93/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.20
no shutdown
interface port-channel2.748
description Tenant 4 N7K-N7K Outside
mtu 9216
encapsulation dot1q 748
vrf member T4-Out
no ip redirects
ip address 10.8.78.121/30
```

```

    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface port-channel2.749
    description Tenant 1 N7K-N7K Inside
    mtu 9216
    encapsulation dot1q 749
    vrf member T4-In
    no ip redirects
    ip address 10.8.78.125/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.20
    no shutdown
interface port-channel2.758
    description Tenant 5 N7K-N7K Outside
    mtu 9216
    encapsulation dot1q 758
    vrf member T5-Out
    no ip redirects
    ip address 10.8.78.153/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface port-channel2.759
    description Tenant 1 N7K-N7K Inside
    mtu 9216
    encapsulation dot1q 759
    vrf member T5-In
    no ip redirects
    ip address 10.8.78.157/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.20
    no shutdown
interface port-channel2.788
    description Global Tenant N7K-N7K Outside
    mtu 9216
    encapsulation dot1q 788
    vrf member GT-Out
    no ip redirects
    ip address 10.8.78.249/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface port-channel3
    description VXIaaS-OOB-MGMT-SW
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 866
    vpc 3
interface port-channel4
    description VXIaaS-OOB-Tenant-MGMT-SW
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 866
    vpc 4

interface port-channel5
    description Common N5K Pair
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 844,864-865
    spanning-tree port type edge trunk
    mtu 9216

```

```

vpc 5
interface port-channel6
  description VXIaaS-OOB-SP-MGMT-SW
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 866
vpc 6
interface port-channel55
  description VXIaaS-SP-5K-Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
  switchport trunk allowed vlan add 867-868,941,980
  spanning-tree port type normal
  mtu 9216
vpc 55
interface port-channel103
  description L3 Link To Service 6500 VSS
  no lacp graceful-convergence
interface port-channel103.611
  description T1 Outside Service 6500
  encapsulation dot1q 611
  vrf member T1-Out
  no ip redirects
  ip address 10.8.76.1/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel103.613
  description T1 Inside Service 6500
  encapsulation dot1q 613
  vrf member T1-In
  no ip redirects
  ip address 10.8.76.9/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel103.621
  description T2 Outside Service 6500
  encapsulation dot1q 621
  vrf member T2-Out
  no ip redirects
  ip address 10.8.76.65/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel103.623
  description T2 Inside Service 6500
  encapsulation dot1q 623
  vrf member T2-In
  no ip redirects
  ip address 10.8.76.73/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel103.631
  description T3 Outside Service 6500
  encapsulation dot1q 631
  vrf member T3-Out
  no ip redirects
  ip address 10.8.76.129/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown

```

```

interface port-channel103.633
  description T3 Inside Service 6500
  encapsulation dot1q 633
  vrf member T3-In
  no ip redirects
  ip address 10.8.76.137/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel103.641
  description T4 Outside Service 6500
  encapsulation dot1q 641
  vrf member T4-Out
  no ip redirects
  ip address 10.8.76.193/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
  ip wccp 61 redirect in
interface port-channel103.643
  description T4 Inside Service 6500
  encapsulation dot1q 643
  vrf member T4-In
  no ip redirects
  ip address 10.8.76.201/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel103.651
  description T5 Outside Service 6500
  encapsulation dot1q 651
  vrf member T5-Out
  no ip redirects
  ip address 10.8.77.1/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel103.653
  description T5 Inside Service 6500
  encapsulation dot1q 653
  vrf member T5-In
  no ip redirects
  ip address 10.8.77.9/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel103.681
  description Global Tenant Outside Service 6500
  encapsulation dot1q 681
  vrf member GT-Out
  no ip redirects
  ip address 10.8.77.193/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface Ethernet4/1
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active

```

```
no shutdown
interface Ethernet4/2
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active
  no shutdown
interface Ethernet4/3
interface Ethernet4/4
interface Ethernet4/5
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet4/6
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet4/7
interface Ethernet4/8
interface Ethernet4/9
interface Ethernet4/10
interface Ethernet4/11
interface Ethernet4/12
interface Ethernet4/13
interface Ethernet4/14
interface Ethernet4/15
interface Ethernet4/16
interface Ethernet4/17
  description Common N5K Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 844,864-865
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 5 mode active
  no shutdown
interface Ethernet4/18
interface Ethernet4/19
interface Ethernet4/20
interface Ethernet4/21
interface Ethernet4/22
interface Ethernet4/23
interface Ethernet4/24
interface Ethernet4/25
interface Ethernet4/26
interface Ethernet4/27
interface Ethernet4/28
interface Ethernet4/29
interface Ethernet4/30
interface Ethernet4/31
interface Ethernet4/32
interface Ethernet4/33
interface Ethernet4/34
interface Ethernet4/35
interface Ethernet4/36
interface Ethernet4/37
interface Ethernet4/38
interface Ethernet4/39
```

```

interface Ethernet4/40
interface Ethernet4/41
interface Ethernet4/42
interface Ethernet4/43
interface Ethernet4/44
interface Ethernet4/45
interface Ethernet4/46
interface Ethernet4/47
interface Ethernet4/48
interface Ethernet5/1
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active
  no shutdown
interface Ethernet5/2
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active
  no shutdown
interface Ethernet5/3
  switchport
interface Ethernet5/4
  switchport
interface Ethernet5/5
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet5/6
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet5/7
interface Ethernet5/8
interface Ethernet5/9
interface Ethernet5/10
interface Ethernet5/11
interface Ethernet5/12
interface Ethernet5/13
interface Ethernet5/14
interface Ethernet5/15
interface Ethernet5/16
interface Ethernet5/17
  description Common N5K Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 844,864-865
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 5 mode active
  no shutdown
interface Ethernet5/18
interface Ethernet5/19

```

```
interface Ethernet5/20
interface Ethernet5/21
interface Ethernet5/22
interface Ethernet5/23
interface Ethernet5/24
interface Ethernet5/25
  description VXIaaS-OOB-MGMT-SW
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 866
  channel-group 3 mode active
  no shutdown
interface Ethernet5/26
  description VXIaaS-OOB-Tenant-MGMT-SW
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 866
  channel-group 4 mode active
  no shutdown
interface Ethernet5/27
  description VXIaaS-GW (10.81 net)
  ip address 10.8.79.242/30
  ip router eigrp 42
  ip summary-address eigrp 42 10.8.0.0/16
  no shutdown

interface Ethernet5/28
  description VXIaaS-OOB-SP-MGMT-SW
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 866
  channel-group 6 mode active
  no shutdown
interface Ethernet5/29
interface Ethernet5/30
interface Ethernet5/31
interface Ethernet5/32
interface Ethernet5/33
  description Common N5K-1 L3
  mtu 9216
  no ip redirects
  ip address 10.10.0.26/30
  ip router eigrp 42
  ip passive-interface eigrp 42
  no shutdown
interface Ethernet5/34
  description Common N5K-2 L3
  mtu 9216
  no ip redirects
  ip address 10.10.0.34/30
  ip router eigrp 42
  ip passive-interface eigrp 42
  no shutdown
interface Ethernet5/35
interface Ethernet5/36
interface Ethernet5/37
  description VXIaaS-DC-ASR9K-1
  no shutdown
interface Ethernet5/37.712
  description T1 Link to ASR9K-1
  encapsulation dot1q 712
  vrf member T1-Out
  no ip redirects
  ip address 10.8.78.1/30
```

```

    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/37.722
    description T2 Link to ASR9K-1
    encapsulation dot1q 722
    vrf member T2-Out
    no ip redirects
    ip address 10.8.78.33/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/37.732
    description T3 Link to ASR9K-1
    encapsulation dot1q 732
    vrf member T3-Out
    no ip redirects
    ip address 10.8.78.65/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/37.742
    description T4 Link to ASR9K-1
    encapsulation dot1q 742
    vrf member T4-Out
    no ip redirects
    ip address 10.8.78.97/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    ip wccp 62 redirect in
    no shutdown
interface Ethernet5/37.752
    description T5 Link to ASR9K-1
    encapsulation dot1q 752
    vrf member T5-Out
    no ip redirects
    ip address 10.8.78.129/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/37.782
    description Global Tenant Link to ASR9K-1
    encapsulation dot1q 782
    vrf member GT-Out
    no ip redirects
    ip address 10.8.78.225/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/38
    description VXIaaS-DC-ASR9K-2
    no shutdown
interface Ethernet5/38.713
    description T1 Link to ASR9K-2
    encapsulation dot1q 713
    vrf member T1-Out
    no ip redirects
    ip address 10.8.78.5/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.0
    no shutdown
interface Ethernet5/38.723
    description T2 Link to ASR9K-2
    encapsulation dot1q 723

```

```

vrf member T2-Out
no ip redirects
ip address 10.8.78.37/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/38.733
description T3 Link to ASR9K-2
encapsulation dot1q 733
vrf member T3-Out
no ip redirects
ip address 10.8.78.69/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/38.743
description T4 Link to ASR9K-2
encapsulation dot1q 743
vrf member T4-Out
no ip redirects
ip address 10.8.78.101/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
ip wccp 62 redirect in
no shutdown
interface Ethernet5/38.753
description T5 Link to ASR9K-2
encapsulation dot1q 753
vrf member T5-Out
no ip redirects
ip address 10.8.78.133/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/38.783
description Global Tenant Link to ASR9K-2
encapsulation dot1q 783
vrf member GT-Out
no ip redirects
ip address 10.8.78.229/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown

interface Ethernet5/39
interface Ethernet5/40
interface Ethernet5/41
description L3 Link To Service 6500 VSS
channel-group 103 mode active
no shutdown
interface Ethernet5/42
description L3 Link To Service 6500 VSS
channel-group 103 mode active
no shutdown
interface Ethernet5/43
interface Ethernet5/44
interface Ethernet5/45
description VXIaaS-SP-5K
switchport
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
switchport trunk allowed vlan add 867-868,941,980
mtu 9216
channel-group 55 mode active

```

```

no shutdown
interface Ethernet5/46
description VXIaaS-SP-5K
switchport
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
switchport trunk allowed vlan add 867-868,941,980
mtu 9216
channel-group 55 mode active
no shutdown
interface Ethernet5/47
interface Ethernet5/48
interface mgmt0
vrf member management
ip address 10.8.66.21/24
interface loopback0
description MGMT Loopback
ip address 10.8.79.250/32
ip router eigrp 42
ip passive-interface eigrp 42
interface loopback1
description T1 Outside Loopback
vrf member T1-Out
ip address 192.168.10.5/32
ip router ospf VXIaaS area 0.0.0.0

interface loopback2
description T2 Outside Loopback
vrf member T2-Out
ip address 192.168.20.5/32
ip router ospf VXIaaS area 0.0.0.0
interface loopback3
description T3 Outside Loopback
vrf member T3-Out
ip address 192.168.30.5/32
ip router ospf VXIaaS area 0.0.0.0
interface loopback4
description T4 Outside Loopback
vrf member T4-Out
ip address 192.168.40.5/32
ip router ospf VXIaaS area 0.0.0.0
interface loopback5
description T5 Outside Loopback
vrf member T5-Out
ip address 192.168.50.5/32
ip router ospf VXIaaS area 0.0.0.0
interface loopback8
description Global Tenant Outside Loopback
vrf member GT-Out
ip address 192.168.80.5/32
ip router ospf VXIaaS area 0.0.0.0
interface loopback10
description T1 Inside Loopback
vrf member T1-In
ip address 192.168.10.9/32
ip router ospf VXIaaS area 0.0.0.20
interface loopback20
description T2 Inside Loopback
vrf member T2-In
ip address 192.168.20.9/32
ip router ospf VXIaaS area 0.0.0.20
interface loopback30
description T3 Inside Loopback
vrf member T3-In

```

```

    ip address 192.168.30.9/32
    ip router ospf VXIaaS area 0.0.0.20
interface loopback40
    description T4 Inside Loopback
    vrf member T4-In
    ip address 192.168.40.9/32
    ip router ospf VXIaaS area 0.0.0.20
interface loopback50
    description T5 Inside Loopback
    vrf member T5-In
    ip address 192.168.50.9/32
    ip router ospf VXIaaS area 0.0.0.20
clock timezone EST -5 0
clock summer-time EDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
line console
    exec-timeout 0
line vty
    exec-timeout 420
boot kickstart bootflash:/n7000-s1-kickstart.6.1.2.bin sup-1
boot system bootflash:/n7000-s1-dk9.6.1.2.bin sup-1
boot kickstart bootflash:/n7000-s1-kickstart.6.1.2.bin sup-2
boot system bootflash:/n7000-s1-dk9.6.1.2.bin sup-2
router eigrp 42
    default-metric 10000 100 255 1 1500
    redistribute bgp 42 route-map EIGRP-BGP
router ospf VXIaaS
    vrf GT-Out
        router-id 8.8.8.15
        area 0.0.0.10 nssa default-information-originate
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T1-In
        router-id 1.1.1.15
        area 0.0.0.20 nssa
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T1-Out
        router-id 1.1.1.10
        area 0.0.0.10 nssa default-information-originate
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T2-In
        router-id 2.2.2.15
        area 0.0.0.20 nssa
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T2-Out
        router-id 2.2.2.10
        area 0.0.0.10 nssa default-information-originate
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T3-In
        router-id 3.3.3.15
        area 0.0.0.20 nssa
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T3-Out
        router-id 3.3.3.10
        area 0.0.0.10 nssa default-information-originate
        log-adjacency-changes detail
        auto-cost reference-bandwidth 100 Gbps
    vrf T4-In
        router-id 4.4.4.15
        area 0.0.0.20 nssa

```

```

    log-adjacency-changes detail
    auto-cost reference-bandwidth 100 Gbps
vrf T4-Out
  router-id 4.4.4.10
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T5-In
  router-id 5.5.5.15
  area 0.0.0.20 nssa
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T5-Out
  router-id 5.5.5.10
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
router bgp 42
  address-family ipv4 unicast
    redistribute direct route-map BGP-DIRECT
    aggregate-address 10.8.0.0/16 summary-only
    maximum-paths 4
  neighbor 10.10.0.25 remote-as 1000
    address-family ipv4 unicast
      route-map COMMON-SUM in
  neighbor 10.10.0.33 remote-as 1000
    address-family ipv4 unicast
      route-map COMMON-SUM in

Time: Sat Jun  1 12:21:48 2013
version 6.1(2)
hostname VXIaaS-DC-N7K-2
vdc VXIaaS-DC-N7K-2 id 1
  limit-resource module-type f2
  allocate interface Ethernet4/1-48
  allocate interface Ethernet5/1-48
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
feature telnet
cfs eth distribute
feature ospf
feature bgp
feature eigrp
feature interface-vlan
feature dot1x
feature hsrp
feature lacp
feature dhcp
feature vpc
feature wccp
logging level pixm 2
username admin password 5 $1$EBvmxlZW$V1JNH1Nyrvy6kVcEtaTbI1 role network-admin
no password strength-check
ip domain-lookup
ip domain-name vxiaas.local

```

```

ip domain-list cisco.com
ip domain-list vxiaas.local
ip name-server 10.8.64.100
class-map type qos match-any VOICE
  match dscp 46
class-map type qos match-all BULK-DATA
  match dscp 10,12,14
class-map type qos match-any SCAVENGER
  match dscp 8
class-map type qos match-any CALL-SIGNALING
  match dscp 24
class-map type qos match-any NETWORK-CONTROL
  match dscp 48
class-map type qos match-any TRANSACTIONAL-DATA
  match dscp 18,20,22
class-map type qos match-any MULTIMEDIA-STREAMING
  match dscp 26,28,30
class-map type qos match-any MULTIMEDIA-CONFERENCING
  match dscp 34,36,38
class-map type queuing match-any BROADCAST-VIDEO
policy-map type queuing VPN-EDGE
copp profile strict
snmp-server user admin network-admin auth md5 0x7b83230d09d201df578c1f817cb425da priv
0x7b83230d09d201df578c1f817cb425da localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.8.79.254
vlan 1,711,721,731,741,751,761,844,864-868,941,980
vrf context GT-Out
vrf context T1-In
vrf context T1-Out
vrf context T2-In
vrf context T2-Out
vrf context T3-In
vrf context T3-Out
vrf context T4-In
vrf context T4-Out
  ip wccp 61
  ip wccp 62
vrf context T5-In
vrf context T5-Out
vrf context management
  ip route 0.0.0.0/0 10.8.66.1
vlan 711
  name Tenant-1-VM-A
vlan 721
  name Tenant-2-VM-A
vlan 731
  name Tenant-3-VM-A
vlan 741
  name Tenant-4-VM-A
vlan 751
  name Tenant-5-VM-A
vlan 761
  name Tenant-6-VM-A
vlan 844
  name HVXI-NetApp-NFS
vlan 864
  name VXIaaS-SP-Infra
vlan 865
  name VXIaaS-SP-DH

```

```

vlan 866
  name VXIaaS-OOB-MGMT
vlan 867
  name VXIaaS-SP-Hypervisor-MGMT
vlan 868
  name VXIaaS-SP-vMotion
vlan 941
  name Tenant-4-vWAAS
vlan 980
  name Desktone-L2-Link-Local

spanning-tree vlan 711,721,731,741,751,761,864-868,980 priority 4096
ip prefix-list BGP-DIRECT seq 5 permit 10.8.0.0/16 le 30
ip prefix-list COMMON-SUM seq 5 permit 10.10.0.0/17
ip prefix-list EIGRP-BGP seq 5 permit 10.10.0.0/17
route-map BGP-DIRECT permit 10
  match ip address prefix-list BGP-DIRECT
route-map COMMON-SUM permit 10
  match ip address prefix-list COMMON-SUM
route-map EIGRP-BGP permit 10
  match ip address prefix-list EIGRP-BGP
service dhcp
ip dhcp relay
vpc domain 42
  peer-switch
  peer-keepalive destination 10.8.66.21 source 10.8.66.22
  peer-gateway
  auto-recovery
  ip arp synchronize
interface Vlan1
interface Vlan711
  vrf member T1-In
  no ip redirects
  ip address 10.9.104.3/21
  ip ospf passive-interface
  ip router ospf VXIaaS area 0.0.0.20
  hsrp version 2
  hsrp 711
    ip 10.9.104.1
  ip dhcp relay address 10.9.10.150
  ip dhcp relay address 10.9.10.151
  description Tenant 1 VMs
  no shutdown
  mtu 9216
interface Vlan721
  vrf member T2-In
  no ip redirects
  ip address 10.9.112.3/21
  ip ospf passive-interface
  ip router ospf VXIaaS area 0.0.0.20
  hsrp version 2
  hsrp 721
    ip 10.9.112.1
  ip dhcp relay address 10.9.20.150
  ip dhcp relay address 10.9.20.151
  description Tenant 2 VMs
  no shutdown
  mtu 9216

interface Vlan731
  vrf member T3-In
  no ip redirects
  ip address 10.9.120.3/21
  ip ospf passive-interface

```

```
ip router ospf VXIaaS area 0.0.0.20
hsrp version 2
hsrp 731
  ip 10.9.120.1
ip dhcp relay address 10.9.30.150
ip dhcp relay address 10.9.30.151
description Tenant 3 VMs
no shutdown
mtu 9216
interface Vlan741
vrf member T4-In
no ip redirects
ip address 10.9.128.3/21
ip ospf passive-interface
ip router ospf VXIaaS area 0.0.0.20
hsrp version 2
hsrp 741
  ip 10.9.128.1
ip dhcp relay address 10.9.40.150
ip dhcp relay address 10.9.40.151
description Tenant 4 VMs
no shutdown
mtu 9216
interface Vlan751
vrf member T5-In
no ip redirects
ip address 10.9.136.3/21
ip ospf passive-interface
ip router ospf VXIaaS area 0.0.0.20
hsrp version 2
hsrp 751
  ip 10.9.136.1
ip dhcp relay address 10.9.50.150
ip dhcp relay address 10.9.50.151
description Tenant 5 VMs
no shutdown
mtu 9216
interface Vlan864
no ip redirects
ip address 10.8.64.3/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 864
  ip 10.8.64.1
description VXIaaS-SP-Infra
no shutdown
mtu 9216

interface Vlan865
no ip redirects
ip address 10.8.65.3/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 865
  ip 10.8.65.1
description VXIaaS-SP-DH
no shutdown
mtu 9216
interface Vlan866
no ip redirects
ip address 10.8.66.3/24
ip router eigrp 42
```

```

ip passive-interface eigrp 42
hsrp version 2
hsrp 866
  ip 10.8.66.1
description VXIaaS-OOB-MGMT
no shutdown
mtu 9216
interface Vlan867
no ip redirects
ip address 10.8.67.3/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 867
  ip 10.8.67.1
description VXIaaS-SP-Hypervisor-MGMT
no shutdown
mtu 9216
interface Vlan868
no ip redirects
ip address 10.8.68.3/24
ip router eigrp 42
ip passive-interface eigrp 42
hsrp version 2
hsrp 868
  ip 10.8.68.1
description VXIaaS-SP-vMotion
no shutdown
mtu 9216
interface Vlan941
vrf member T4-Out
no ip redirects
ip address 10.9.192.27/29
ip ospf passive-interface
ip router ospf VXIaaS area 0.0.0.10
hsrp version 2
hsrp 941
  ip 10.9.192.25
ip dhcp relay address 10.9.40.150
ip dhcp relay address 10.9.40.151
description Tenant 4 vWAAS
no shutdown
mtu 9216
interface port-channel1
description vPC Peer Link Between N7Ks
switchport
switchport mode trunk
switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
switchport trunk allowed vlan add 864-868,941,980
spanning-tree port type network
mtu 9216
vpc peer-link
interface port-channel2
description NON-vPC Peer Link Between N7Ks
mtu 9216
interface port-channel2.718
description Tenant 1 N7K-N7K Outside
mtu 9216
encapsulation dot1q 718
vrf member T1-Out
no ip redirects
ip address 10.8.78.26/30
ip ospf network point-to-point

```

```
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface port-channel2.719
  description Tenant 1 N7K-N7K Inside
  mtu 9216
  encapsulation dot1q 719
  vrf member T1-In
  no ip redirects
  ip address 10.8.78.30/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel2.728
  description Tenant 2 N7K-N7K Outside
  mtu 9216
  encapsulation dot1q 728
  vrf member T2-Out
  no ip redirects
  ip address 10.8.78.58/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel2.729
  description Tenant 1 N7K-N7K Inside
  mtu 9216
  encapsulation dot1q 729
  vrf member T2-In
  no ip redirects
  ip address 10.8.78.62/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel2.738
  description Tenant 3 N7K-N7K Outside
  mtu 9216
  encapsulation dot1q 738
  vrf member T3-Out
  no ip redirects
  ip address 10.8.78.90/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel2.739
  description Tenant 1 N7K-N7K Inside
  mtu 9216
  encapsulation dot1q 739
  vrf member T3-In
  no ip redirects
  ip address 10.8.78.94/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel2.748
  description Tenant 4 N7K-N7K Outside
  mtu 9216
  encapsulation dot1q 748
  vrf member T4-Out
  no ip redirects
  ip address 10.8.78.122/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel2.749
  description Tenant 1 N7K-N7K Inside
```

```

mtu 9216
encapsulation dot1q 749
vrf member T4-In
no ip redirects
ip address 10.8.78.126/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.20
no shutdown
interface port-channel2.758
description Tenant 5 N7K-N7K Outside
mtu 9216
encapsulation dot1q 758
vrf member T5-Out
no ip redirects
ip address 10.8.78.154/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.10
no shutdown
interface port-channel2.759
description Tenant 1 N7K-N7K Inside
mtu 9216
encapsulation dot1q 759
vrf member T5-In
no ip redirects
ip address 10.8.78.158/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.20
no shutdown
interface port-channel2.788
description Global Tenant N7K-N7K Outside
mtu 9216
encapsulation dot1q 788
vrf member GT-Out
no ip redirects
ip address 10.8.78.250/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.10
no shutdown
interface port-channel3
description VXIaaS-OOB-MGMT-SW
switchport
switchport mode trunk
switchport trunk allowed vlan 866
vpc 3
interface port-channel4
description VXIaaS-OOB-Tenant-MGMT-SW
switchport
switchport mode trunk
switchport trunk allowed vlan 866
vpc 4
interface port-channel5
description Common N5K Pair
switchport
switchport mode trunk
switchport trunk allowed vlan 844,864-865
spanning-tree port type edge trunk
mtu 9216
vpc 5
interface port-channel6
description VXIaaS-OOB-SP-MGMT-SW
switchport
switchport mode trunk
switchport trunk allowed vlan 866
vpc 6

```

```
interface port-channel55
  description VXIaaS-SP-5K-Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
  switchport trunk allowed vlan add 867-868,941,980
  spanning-tree port type normal
  mtu 9216
  vpc 55
interface port-channel104
  description L3 Link To Service 6500 VSS
  no lACP graceful-convergence
interface port-channel104.612
  description T1 Outside Service 6500
  encapsulation dot1q 612
  vrf member T1-Out
  no ip redirects
  ip address 10.8.76.5/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel104.614
  description T1 Inside Service 6500
  encapsulation dot1q 614
  vrf member T1-In
  no ip redirects
  ip address 10.8.76.13/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel104.622
  description T2 Outside Service 6500
  encapsulation dot1q 622
  vrf member T2-Out
  no ip redirects
  ip address 10.8.76.69/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel104.624
  description T2 Inside Service 6500
  encapsulation dot1q 624
  vrf member T2-In
  no ip redirects
  ip address 10.8.76.77/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.20
  no shutdown
interface port-channel104.632
  description T3 Outside Service 6500
  encapsulation dot1q 632
  vrf member T3-Out
  no ip redirects
  ip address 10.8.76.133/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.10
  no shutdown
interface port-channel104.634
  description T3 Inside Service 6500
  encapsulation dot1q 634
  vrf member T3-In
  no ip redirects
  ip address 10.8.76.141/30
  ip ospf network point-to-point
```

```

    ip router ospf VXIaaS area 0.0.0.20
    no shutdown
interface port-channel104.642
    description T4 Outside Service 6500
    encapsulation dot1q 642
    vrf member T4-Out
    no ip redirects
    ip address 10.8.76.197/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
    ip wccp 61 redirect in
interface port-channel104.644
    description T4 Inside Service 6500
    encapsulation dot1q 644
    vrf member T4-In
    no ip redirects
    ip address 10.8.76.205/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.20
    no shutdown
interface port-channel104.652
    description T5 Outside Service 6500
    encapsulation dot1q 652
    vrf member T5-Out
    no ip redirects
    ip address 10.8.77.5/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface port-channel104.654
    description T5 Inside Service 6500
    encapsulation dot1q 654
    vrf member T5-In
    no ip redirects
    ip address 10.8.77.13/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.20
    no shutdown
interface port-channel104.682
    description Global Tenant Outside Service 6500
    encapsulation dot1q 682
    vrf member GT-Out
    no ip redirects
    ip address 10.8.77.197/30
    ip ospf network point-to-point
    ip router ospf VXIaaS area 0.0.0.10
    no shutdown
interface Ethernet4/1
    description vPC Peer Link Between N7Ks
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
    switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
    switchport trunk allowed vlan add 864-868,941,980
    mtu 9216
    channel-group 1 mode active
    no shutdown

interface Ethernet4/2
    description vPC Peer Link Between N7Ks
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735

```

```
switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
switchport trunk allowed vlan add 864-868,941,980
mtu 9216
channel-group 1 mode active
no shutdown
interface Ethernet4/3
interface Ethernet4/4
interface Ethernet4/5
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet4/6
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet4/7
interface Ethernet4/8
interface Ethernet4/9
interface Ethernet4/10
interface Ethernet4/11
interface Ethernet4/12
interface Ethernet4/13
interface Ethernet4/14
interface Ethernet4/15
interface Ethernet4/16
interface Ethernet4/17
  description Common N5K Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 844,864-865
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 5 mode active
  no shutdown
interface Ethernet4/18
interface Ethernet4/19
interface Ethernet4/20
interface Ethernet4/21
interface Ethernet4/22
interface Ethernet4/23
interface Ethernet4/24
interface Ethernet4/25
interface Ethernet4/26
interface Ethernet4/27
interface Ethernet4/28
interface Ethernet4/29
interface Ethernet4/30
interface Ethernet4/31
interface Ethernet4/32
interface Ethernet4/33
interface Ethernet4/34
interface Ethernet4/35
interface Ethernet4/36
interface Ethernet4/37
interface Ethernet4/38
interface Ethernet4/39
interface Ethernet4/40
interface Ethernet4/41
interface Ethernet4/42
interface Ethernet4/43
interface Ethernet4/44
interface Ethernet4/45
```

```

interface Ethernet4/46
interface Ethernet4/47
interface Ethernet4/48
interface Ethernet5/1
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active
  no shutdown
interface Ethernet5/2
  description vPC Peer Link Between N7Ks
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 711,714-715,721,724-725,731,734-735
  switchport trunk allowed vlan add 741,744-745,751,754-755,761,844
  switchport trunk allowed vlan add 864-868,941,980
  mtu 9216
  channel-group 1 mode active
  no shutdown
interface Ethernet5/3
interface Ethernet5/4
interface Ethernet5/5
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet5/6
  description NON-vPC Peer Link Between N7Ks
  mtu 9216
  channel-group 2 mode active
  no shutdown
interface Ethernet5/7
interface Ethernet5/8
interface Ethernet5/9
interface Ethernet5/10
interface Ethernet5/11
interface Ethernet5/12
interface Ethernet5/13
interface Ethernet5/14
interface Ethernet5/15
interface Ethernet5/16
interface Ethernet5/17
  description Common N5K Pair
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 844,864-865
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 5 mode active
  no shutdown
interface Ethernet5/18
interface Ethernet5/19
interface Ethernet5/20
interface Ethernet5/21
interface Ethernet5/22
interface Ethernet5/23
interface Ethernet5/24
interface Ethernet5/25
  description VXIaaS-OOB-MGMT-SW
  switchport

```

```

switchport mode trunk
switchport trunk allowed vlan 866
channel-group 3 mode active
no shutdown
interface Ethernet5/26
description VXIaaS-OOB-Tenant-MGMT-SW
switchport
switchport mode trunk
switchport trunk allowed vlan 866
channel-group 4 mode active
no shutdown
interface Ethernet5/27
description VXIaaS-GW (10.81 net)
ip address 10.8.79.246/30
ip router eigrp 42
ip summary-address eigrp 42 10.8.0.0/16
no shutdown
interface Ethernet5/28
description VXIaaS-OOB-SP-MGMT-SW
switchport
switchport mode trunk
switchport trunk allowed vlan 866
channel-group 6 mode active
no shutdown
interface Ethernet5/29
interface Ethernet5/30
interface Ethernet5/31
interface Ethernet5/32
interface Ethernet5/33
description Common N5K-1 L3
mtu 9216
no ip redirects
ip address 10.10.0.30/30
ip router eigrp 42
ip passive-interface eigrp 42
no shutdown
interface Ethernet5/34
description Common N5K-2 L3
mtu 9216
no ip redirects
ip address 10.10.0.38/30
ip router eigrp 42
ip passive-interface eigrp 42
no shutdown
interface Ethernet5/35
interface Ethernet5/36
interface Ethernet5/37
description desc VXIaaS-DC-ASR9K-1
no shutdown
interface Ethernet5/37.714
description T1 Link to ASR9K-1
encapsulation dot1q 714
vrf member T1-Out
no ip redirects
ip address 10.8.78.9/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown

interface Ethernet5/37.724
description T2 Link to ASR9K-1
encapsulation dot1q 724
vrf member T2-Out
no ip redirects

```

```

ip address 10.8.78.41/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/37.734
  description T3 Link to ASR9K-1
  encapsulation dot1q 734
  vrf member T3-Out
  no ip redirects
  ip address 10.8.78.73/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  no shutdown
interface Ethernet5/37.744
  description T4 Link to ASR9K-1
  encapsulation dot1q 744
  vrf member T4-Out
  no ip redirects
  ip address 10.8.78.105/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  ip wccp 62 redirect in
  no shutdown
interface Ethernet5/37.754
  description T5 Link to ASR9K-1
  encapsulation dot1q 754
  vrf member T5-Out
  no ip redirects
  ip address 10.8.78.137/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  no shutdown
interface Ethernet5/37.784
  description Global Tenant Link to ASR9K-1
  encapsulation dot1q 784
  vrf member GT-Out
  no ip redirects
  ip address 10.8.78.233/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  no shutdown

interface Ethernet5/38
  description VXIaaS-DC-ASR9K-2
  no shutdown
interface Ethernet5/38.715
  description T1 Link to ASR9K-2
  encapsulation dot1q 715
  vrf member T1-Out
  no ip redirects
  ip address 10.8.78.13/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  no shutdown
interface Ethernet5/38.725
  description T2 Link to ASR9K-2
  encapsulation dot1q 725
  vrf member T2-Out
  no ip redirects
  ip address 10.8.78.45/30
  ip ospf network point-to-point
  ip router ospf VXIaaS area 0.0.0.0
  no shutdown
interface Ethernet5/38.735

```

```

description T3 Link to ASR9K-2
encapsulation dot1q 735
vrf member T3-Out
no ip redirects
ip address 10.8.78.77/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/38.745
description T4 Link to ASR9K-2
encapsulation dot1q 745
vrf member T4-Out
no ip redirects
ip address 10.8.78.109/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
ip wccp 62 redirect in
no shutdown
interface Ethernet5/38.755
description T5 Link to ASR9K-2
encapsulation dot1q 755
vrf member T5-Out
no ip redirects
ip address 10.8.78.141/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/38.785
description Global Tenant Link to ASR9K-2
encapsulation dot1q 785
vrf member GT-Out
no ip redirects
ip address 10.8.78.237/30
ip ospf network point-to-point
ip router ospf VXIaaS area 0.0.0.0
no shutdown
interface Ethernet5/39
interface Ethernet5/40
interface Ethernet5/41
description L3 Link To Service 6500 VSS
channel-group 104 mode active
no shutdown
interface Ethernet5/42
description L3 Link To Service 6500 VSS
channel-group 104 mode active
no shutdown
interface Ethernet5/43
interface Ethernet5/44
interface Ethernet5/45
description VXIaaS-SP-5K
switchport
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
switchport trunk allowed vlan add 867-868,941,980
mtu 9216
channel-group 55 mode active
no shutdown
interface Ethernet5/46
description VXIaaS-SP-5K
switchport
switchport mode trunk
switchport trunk allowed vlan 711,721,731,741,751,761,844,864-865
switchport trunk allowed vlan add 867-868,941,980
mtu 9216

```

```

    channel-group 55 mode active
    no shutdown
interface Ethernet5/47
interface Ethernet5/48
interface mgmt0
    vrf member management
    ip address 10.8.66.22/24
interface loopback0
    description MGMT Loopback
    ip address 10.8.79.251/32
    ip router eigrp 42
    ip passive-interface eigrp 42
interface loopback1
    description T1 Outside Loopback
    vrf member T1-Out
    ip address 192.168.10.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback2
    description T2 Outside Loopback
    vrf member T2-Out
    ip address 192.168.20.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback3
    description T3 Outside Loopback
    vrf member T3-Out
    ip address 192.168.30.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback4
    description T4 Outside Loopback
    vrf member T4-Out
    ip address 192.168.40.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback5
    description T5 Outside Loopback
    vrf member T5-Out
    ip address 192.168.50.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback8
    description Global Tenant Outside Loopback
    vrf member GT-Out
    ip address 192.168.80.6/32
    ip router ospf VXIaaS area 0.0.0.0
interface loopback10
    description T1 Inside Loopback
    vrf member T1-In
    ip address 192.168.10.10/32
    ip router ospf VXIaaS area 0.0.0.20
interface loopback20
    description T2 Inside Loopback
    vrf member T2-In
    ip address 192.168.20.10/32
    ip router ospf VXIaaS area 0.0.0.20
interface loopback30
    description T3 Inside Loopback
    vrf member T3-In
    ip address 192.168.30.10/32
    ip router ospf VXIaaS area 0.0.0.20

interface loopback40
    description T4 Inside Loopback
    vrf member T4-In
    ip address 192.168.40.10/32
    ip router ospf VXIaaS area 0.0.0.20
interface loopback50

```

```

description T5 Inside Loopback
vrf member T5-In
ip address 192.168.50.10/32
ip router ospf VXIaaS area 0.0.0.20
clock timezone EST -5 0
clock summer-time EDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
line console
  exec-timeout 0
line vty
  exec-timeout 420
boot kickstart bootflash:/n7000-s1-kickstart.6.1.2.bin sup-1
boot system bootflash:/n7000-s1-dk9.6.1.2.bin sup-1
boot kickstart bootflash:/n7000-s1-kickstart.6.1.2.bin sup-2
boot system bootflash:/n7000-s1-dk9.6.1.2.bin sup-2
router eigrp 42
  default-metric 10000 100 255 1 1500
  redistribute bgp 42 route-map EIGRP-BGP
router ospf VXIaaS
vrf GT-Out
  router-id 8.8.8.11
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T1-In
  router-id 1.1.1.16
  area 0.0.0.20 nssa
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T1-Out
  router-id 1.1.1.11
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T2-In
  router-id 2.2.2.16
  area 0.0.0.20 nssa
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T2-Out
  router-id 2.2.2.11
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T3-In
  router-id 3.3.3.16
  area 0.0.0.20 nssa
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T3-Out
  router-id 3.3.3.11
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T4-In
  router-id 4.4.4.16
  area 0.0.0.20 nssa
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T4-Out
  router-id 4.4.4.11
  area 0.0.0.10 nssa default-information-originate
  log-adjacency-changes detail
  auto-cost reference-bandwidth 100 Gbps
vrf T5-In

```

```

router-id 5.5.5.16
area 0.0.0.20 nssa
log-adjacency-changes detail
auto-cost reference-bandwidth 100 Gbps
vrf T5-Out
router-id 5.5.5.11
area 0.0.0.10 nssa default-information-originate
log-adjacency-changes detail
auto-cost reference-bandwidth 100 Gbps
router bgp 42
address-family ipv4 unicast
redistribute direct route-map BGP-DIRECT
aggregate-address 10.8.0.0/16 summary-only
maximum-paths 4
neighbor 10.10.0.29 remote-as 1000
address-family ipv4 unicast
route-map COMMON-SUM in
neighbor 10.10.0.37 remote-as 1000
address-family ipv4 unicast
route-map COMMON-SUM in

```

Cisco Catalyst 6506 configuration:

```

!
version 12.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service counters max age 5
hostname VXIaaS-DC-SVC-6500
boot-start-marker
boot system flash bootflash:s72033-ipservicesk9_wan-mz.122-33.SXJ1.bin
boot-end-marker
security passwords min-length 1
no logging console
enable password cisco_123
username admin password 0 cisco_123
aaa new-model
aaa authentication login default local
aaa session-id common
clock timezone EST -5
clock summer-time EDT recurring
firewall autostate
firewall multiple-vlan-interfaces
firewall switch 1 module 3 vlan-group 1,3
firewall switch 2 module 3 vlan-group 1,3
firewall vlan-group 1 910,915
firewall vlan-group 3 616,617,626,627,636,637,646,647,656,657
ip vrf GT-In
ip vrf GT-Out
ip vrf T1-In
ip vrf T1-Out
ip vrf T2-In
ip vrf T2-Out
ip vrf T3-In
ip vrf T3-Out
ip vrf T4-In
ip vrf T4-Out
ip vrf T5-In
ip vrf T5-Out
ip vrf T6-In
ip vrf T6-Out
ip vrf T7-In
ip vrf T7-Out

```

```
ip ssh source-interface Vlan866
no ip domain-lookup
ip domain-name vxiaas.local
vtp mode transparent
switch virtual domain 42
  switch mode virtual
mls netflow interface
spanning-tree mode rapid-pvst
spanning-tree extend system-id
diagnostic bootup level complete
redundancy
  main-cpu
  auto-sync running-config
mode sso
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
vlan 616
  name T1-ASA-Out
vlan 617
  name T1-ASA-In
vlan 626
  name T2-ASA-Out
vlan 627
  name T2-ASA-In
vlan 636
  name T3-ASA-Out
vlan 637
  name T3-ASA-In
vlan 646
  name T4-ASA-Out
vlan 647
  name T4-ASA-In
vlan 656
  name T5-ASA-Out
vlan 657
  name T5-ASA-In
vlan 686
  name GT-ASA-Out
vlan 687
  name GT-ASA-In
vlan 866
  name VXIaaS-OOB-MGMT
vlan 910
  name ASA-FW-LAN-FO
vlan 915
  name ASA-FW-STATE-FO
interface Loopback1
  description T1 Outside Loopback
  ip vrf forwarding T1-Out
  ip address 192.168.10.7 255.255.255.255
interface Loopback2
  description T2 Outside Loopback
  ip vrf forwarding T2-Out
  ip address 192.168.20.7 255.255.255.255
interface Loopback3
  description T3 Outside Loopback
  ip vrf forwarding T3-Out
  ip address 192.168.30.7 255.255.255.255
interface Loopback4
  description T4 Outside Loopback
  ip vrf forwarding T4-Out
  ip address 192.168.40.7 255.255.255.255
interface Loopback5
  description T5 Outside Loopback
```

```

ip vrf forwarding T5-Out
ip address 192.168.50.7 255.255.255.255
interface Loopback8
description GT Outside Loopback
ip vrf forwarding GT-Out
ip address 192.168.80.7 255.255.255.255
interface Loopback10
description T1 Inside Loopback
ip vrf forwarding T1-In
ip address 192.168.10.8 255.255.255.255
interface Loopback20
description T2 Inside Loopback
ip vrf forwarding T2-In
ip address 192.168.20.8 255.255.255.255
interface Loopback30
description T3 Inside Loopback
ip vrf forwarding T3-In
ip address 192.168.30.8 255.255.255.255
interface Loopback40
description T4 Inside Loopback
ip vrf forwarding T4-In
ip address 192.168.40.8 255.255.255.255
interface Loopback50
description T5 Inside Loopback
ip vrf forwarding T5-In
ip address 192.168.50.8 255.255.255.255
interface Loopback80
description GT Inside Loopback
ip vrf forwarding GT-In
ip address 192.168.80.8 255.255.255.255
interface Port-channel10
no switchport
no ip address
switch virtual link 1
mls qos trust cos
no mls qos channel-consistency
interface Port-channel15
description VPN ASA Outside
no switchport
ip vrf forwarding GT-Out
ip address 10.8.77.209 255.255.255.248
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel16
description VPN ASA Inside
no switchport
no ip address
no ip redirects
no ip unreachable
no ip proxy-arp
interface Port-channel16.635
encapsulation dot1Q 635
ip vrf forwarding T3-Out
ip address 10.8.76.153 255.255.255.248
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel20
no switchport
no ip address
switch virtual link 2

```

```
mls qos trust cos
no mls qos channel-consistency
interface Port-channel103
description N7K-1 Uplink
no switchport
no ip address
interface Port-channel103.611
description T1 Outside N7K-1
encapsulation dot1Q 611
ip vrf forwarding T1-Out
ip address 10.8.76.2 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.613
description T1 Inside N7K-1
encapsulation dot1Q 613
ip vrf forwarding T1-In
ip address 10.8.76.10 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.621
description T2 Outside N7K-1
encapsulation dot1Q 621
ip vrf forwarding T2-Out
ip address 10.8.76.66 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.623
description T2 Inside N7K-1
encapsulation dot1Q 623
ip vrf forwarding T2-In
ip address 10.8.76.74 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.631
description T3 Outside N7K-1
encapsulation dot1Q 631
ip vrf forwarding T3-Out
ip address 10.8.76.130 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
shutdown
interface Port-channel103.633
description T3 Inside N7K-1
encapsulation dot1Q 633
ip vrf forwarding T3-In
ip address 10.8.76.138 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.641
description T4 Outside N7K-1
encapsulation dot1Q 641
```

```

ip vrf forwarding T4-Out
ip address 10.8.76.194 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.643
description T4 Inside N7K-1
encapsulation dot1Q 643
ip vrf forwarding T4-In
ip address 10.8.76.202 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.651
description T5 Outside N7K-1
encapsulation dot1Q 651
ip vrf forwarding T5-Out
ip address 10.8.77.2 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.653
description T5 Inside N7K-1
encapsulation dot1Q 653
ip vrf forwarding T5-In
ip address 10.8.77.10 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel103.681
description GT Outside N7K-1
encapsulation dot1Q 681
ip vrf forwarding GT-Out
ip address 10.8.77.194 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104
description N7K-2 Uplink
no switchport
no ip address
interface Port-channel104.612
description T1 Outside N7K-2
encapsulation dot1Q 612
ip vrf forwarding T1-Out
ip address 10.8.76.6 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.614
description T1 Inside N7K-2
encapsulation dot1Q 614
ip vrf forwarding T1-In
ip address 10.8.76.14 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip ospf network point-to-point

```

```
interface Port-channel104.622
description T2 Outside N7K-2
encapsulation dot1Q 622
ip vrf forwarding T2-Out
ip address 10.8.76.70 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.624
description T2 Inside N7K-2
encapsulation dot1Q 624
ip vrf forwarding T2-In
ip address 10.8.76.78 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.632
description T3 Outside N7K-2
encapsulation dot1Q 632
ip vrf forwarding T3-Out
ip address 10.8.76.134 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
shutdown
interface Port-channel104.634
description T3 Inside N7K-2
encapsulation dot1Q 634
ip vrf forwarding T3-In
ip address 10.8.76.142 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.642
description T4 Outside N7K-2
encapsulation dot1Q 642
ip vrf forwarding T4-Out
ip address 10.8.76.198 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.644
description T4 Inside N7K-2
encapsulation dot1Q 644
ip vrf forwarding T4-In
ip address 10.8.76.206 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
interface Port-channel104.652
description T5 Outside N7K-2
encapsulation dot1Q 652
ip vrf forwarding T5-Out
ip address 10.8.77.6 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip ospf network point-to-point
```

```

interface Port-channel104.654
  description T5 Inside N7K-2
  encapsulation dot1Q 654
  ip vrf forwarding T5-In
  ip address 10.8.77.14 255.255.255.252
  no ip redirects
  no ip unreachableables
  no ip proxy-arp
  ip ospf network point-to-point
interface Port-channel104.682
  description GT Outside N7K-2
  encapsulation dot1Q 682
  ip vrf forwarding GT-Out
  ip address 10.8.77.198 255.255.255.252
  no ip redirects
  no ip unreachableables
  no ip proxy-arp
  ip ospf network point-to-point
interface TenGigabitEthernet1/1/1
  description N7K-1 Uplink
  no switchport
  no ip address
  channel-group 103 mode active
interface TenGigabitEthernet1/1/2
  description N7K-2 Uplink
  no switchport
  no ip address
  channel-group 104 mode active
interface TenGigabitEthernet1/1/3
  description VPN ASA Outside
  no switchport
  no ip address
  channel-group 15 mode active
interface TenGigabitEthernet1/1/4
  description VPN ASA Inside
  no switchport
  no ip address
  channel-group 16 mode active
interface GigabitEthernet1/5/1
  no switchport
  no ip address
  shutdown
interface GigabitEthernet1/5/2
  no switchport
  no ip address
  shutdown
interface GigabitEthernet1/5/3
  description VXIaaS-OOB-MGMT SW1
  switchport
  switchport access vlan 866
  switchport mode access
  arp timeout 1500
interface TenGigabitEthernet1/5/4
  no switchport
  no ip address
  mls qos trust cos
  channel-group 10 mode on
interface TenGigabitEthernet1/5/5
  no switchport
  no ip address
  mls qos trust cos
  channel-group 10 mode on
interface TenGigabitEthernet2/1/1
  description N7K-1 Uplink

```

```
no switchport
no ip address
channel-group 103 mode active
interface TenGigabitEthernet2/1/2
description N7K-2 Uplink
no switchport
no ip address
channel-group 104 mode active
interface TenGigabitEthernet2/1/3
description VPN ASA Inside
no switchport
no ip address
channel-group 16 mode active
interface TenGigabitEthernet2/1/4
description VPN ASA Outside
no switchport
no ip address
channel-group 15 mode active
interface GigabitEthernet2/5/1
no switchport
no ip address
shutdown
interface GigabitEthernet2/5/2
no switchport
no ip address
shutdown
interface GigabitEthernet2/5/3
description VXIaaS-OOB-MGMT SW2
switchport
switchport access vlan 866
switchport mode access
arp timeout 1500
interface TenGigabitEthernet2/5/4
no switchport
no ip address
mls qos trust cos
channel-group 20 mode on
interface TenGigabitEthernet2/5/5
no switchport
no ip address
mls qos trust cos
channel-group 20 mode on
interface Vlan1
no ip address
shutdown
interface Vlan616
description T1 ASA Outside
ip vrf forwarding T1-Out
ip address 10.8.76.33 255.255.255.248
interface Vlan617
description T1 ASA Inside
ip vrf forwarding T1-In
ip address 10.8.76.41 255.255.255.248
interface Vlan626
description T2 ASA Outside
ip vrf forwarding T2-Out
ip address 10.8.76.97 255.255.255.248
interface Vlan627
description T2 ASA Inside
ip vrf forwarding T2-In
ip address 10.8.76.105 255.255.255.248
interface Vlan636
description T3 ASA Outside
ip vrf forwarding T3-Out
```

```

    ip address 10.8.76.161 255.255.255.248
interface Vlan637
    description T3 ASA Inside
    ip vrf forwarding T3-In
    ip address 10.8.76.169 255.255.255.248
interface Vlan646
    description T4 ASA Outside
    ip vrf forwarding T4-Out
    ip address 10.8.76.225 255.255.255.248
interface Vlan647
    description T4 ASA Inside
    ip vrf forwarding T4-In
    ip address 10.8.76.233 255.255.255.248
interface Vlan656
    description T5 ASA Outside
    ip vrf forwarding T5-Out
    ip address 10.8.77.33 255.255.255.248
interface Vlan657
    description T5 ASA Inside
    ip vrf forwarding T5-In
    ip address 10.8.77.41 255.255.255.248
interface Vlan866
    description VXIaaS-OOB-MGMT
    ip address 10.8.66.13 255.255.255.0
    no ip redirects
    no ip unreachable
    no ip proxy-arp
    arp timeout 1500
router ospf 10 vrf T1-Out
    router-id 1.1.1.12
    log-adjacency-changes detail
    auto-cost reference-bandwidth 100000
    nsf ietf
    capability vrf-lite
    area 10 nssa
    redistribute static subnets
    passive-interface Loopback1
    passive-interface Vlan616
    network 10.8.76.2 0.0.0.0 area 10
    network 10.8.76.6 0.0.0.0 area 10
    network 10.8.76.33 0.0.0.0 area 10
    network 192.168.10.7 0.0.0.0 area 10
router ospf 20 vrf T2-Out
    router-id 2.2.2.12
    log-adjacency-changes detail
    auto-cost reference-bandwidth 100000
    nsf ietf
    capability vrf-lite
    area 10 nssa
    redistribute static subnets
    passive-interface Loopback2
    passive-interface Vlan626
    network 10.8.76.66 0.0.0.0 area 10
    network 10.8.76.70 0.0.0.0 area 10
    network 10.8.76.97 0.0.0.0 area 10
    network 192.168.20.7 0.0.0.0 area 10
router ospf 30 vrf T3-Out
    router-id 3.3.3.12
    log-adjacency-changes detail
    auto-cost reference-bandwidth 100000
    nsf ietf
    capability vrf-lite
    area 10 nssa
    redistribute static subnets

```

```
passive-interface Loopback3
passive-interface Vlan636
network 10.8.76.130 0.0.0.0 area 10
network 10.8.76.134 0.0.0.0 area 10
network 10.8.76.153 0.0.0.0 area 10
network 10.8.76.161 0.0.0.0 area 10
network 192.168.30.7 0.0.0.0 area 10
router ospf 40 vrf T4-Out
router-id 4.4.4.12
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 10 nssa
redistribute static subnets
passive-interface Loopback4
passive-interface Vlan646
network 10.8.76.194 0.0.0.0 area 10
network 10.8.76.198 0.0.0.0 area 10
network 10.8.76.225 0.0.0.0 area 10
network 192.168.40.7 0.0.0.0 area 10
router ospf 50 vrf T5-Out
router-id 5.5.5.12
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 10 nssa
redistribute static subnets
passive-interface Loopback5
passive-interface Vlan656
network 10.8.77.2 0.0.0.0 area 10
network 10.8.77.6 0.0.0.0 area 10
network 10.8.77.33 0.0.0.0 area 10
network 192.168.50.7 0.0.0.0 area 10
router ospf 15 vrf T1-In
router-id 1.1.1.13
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 20 nssa default-information-originate
passive-interface Loopback10
passive-interface Vlan617
network 10.8.76.10 0.0.0.0 area 20
network 10.8.76.14 0.0.0.0 area 20
network 10.8.76.41 0.0.0.0 area 20
network 192.168.10.8 0.0.0.0 area 20
router ospf 25 vrf T2-In
router-id 2.2.2.13
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 20 nssa default-information-originate
passive-interface Vlan627
network 10.8.76.74 0.0.0.0 area 20
network 10.8.76.78 0.0.0.0 area 20
network 10.8.76.105 0.0.0.0 area 20
network 192.168.20.8 0.0.0.0 area 20
router ospf 35 vrf T3-In
router-id 3.3.3.13
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
```

```

nsf ietf
capability vrf-lite
area 20 nssa default-information-originate
passive-interface Loopback30
passive-interface Vlan637
network 10.8.76.138 0.0.0.0 area 20
network 10.8.76.142 0.0.0.0 area 20
network 10.8.76.169 0.0.0.0 area 20
network 192.168.30.8 0.0.0.0 area 20
router ospf 45 vrf T4-In
router-id 4.4.4.13
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 20 nssa default-information-originate
passive-interface Loopback40
passive-interface Vlan647
network 10.8.76.202 0.0.0.0 area 20
network 10.8.76.206 0.0.0.0 area 20
network 10.8.76.233 0.0.0.0 area 20
network 192.168.40.8 0.0.0.0 area 20
router ospf 55 vrf T5-In
router-id 5.5.5.13
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 20 nssa default-information-originate
passive-interface Loopback50
passive-interface Vlan657
network 10.8.77.10 0.0.0.0 area 20
network 10.8.77.14 0.0.0.0 area 20
network 10.8.77.41 0.0.0.0 area 20
network 192.168.50.8 0.0.0.0 area 20
router ospf 80 vrf GT-Out
router-id 8.8.8.12
log-adjacency-changes detail
auto-cost reference-bandwidth 100000
nsf ietf
capability vrf-lite
area 10 nssa
redistribute static subnets
passive-interface Loopback8
network 10.8.77.194 0.0.0.0 area 10
network 10.8.77.198 0.0.0.0 area 10
network 10.8.77.209 0.0.0.0 area 10
network 192.168.80.7 0.0.0.0 area 10
ip classless
ip forward-protocol nd
ip route 0.0.0.0 0.0.0.0 10.8.66.1
ip route vrf T1-In 0.0.0.0 0.0.0.0 10.8.76.43
ip route vrf T1-Out 10.9.104.0 255.255.248.0 10.8.76.35
ip route vrf T2-In 0.0.0.0 0.0.0.0 10.8.76.107
ip route vrf T2-Out 10.9.112.0 255.255.248.0 10.8.76.99
ip route vrf T3-In 0.0.0.0 0.0.0.0 10.8.76.171
ip route vrf T3-Out 10.9.120.0 255.255.248.0 10.8.76.163
ip route vrf T4-In 0.0.0.0 0.0.0.0 10.8.76.235
ip route vrf T4-Out 10.9.128.0 255.255.248.0 10.8.76.227
ip route vrf T5-In 0.0.0.0 0.0.0.0 10.8.77.43
ip route vrf T5-Out 10.9.136.0 255.255.248.0 10.8.77.35
no ip http server
no ip http secure-server
control-plane

```

```

dial-peer cor custom
line con 0
  exec-timeout 0 0
  logging synchronous
line vty 0 4
  exec-timeout 300 0
  password cisco_123
  logging synchronous
  transport input telnet ssh
mac-address-table aging-time 1800
module provision switch 1
  slot 1 slot-type 148 port-type 60 number 4 virtual-slot 17
  slot 2 slot-type 207 port-type 106 number 1 virtual-slot 18
  slot 3 slot-type 330 port-type 111 number 3 virtual-slot 19
  slot 5 slot-type 254 port-type 31 number 2 port-type 61 number 1 port-type 60 number
2 virtual-slot 21
module provision switch 2
  slot 1 slot-type 148 port-type 60 number 4 virtual-slot 33
  slot 2 slot-type 207 port-type 106 number 1 virtual-slot 34
  slot 3 slot-type 330 port-type 111 number 3 virtual-slot 35
  slot 5 slot-type 254 port-type 31 number 2 port-type 61 number 1 port-type 60 number
2 virtual-slot 37
end

```

Cisco ASA Service Module Tenant 1 configuration:

```

ASA Version 8.5(1) <context>
hostname T1
enable password MkHB0dBX0zv6EXRs encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
names
interface Vlan616
  nameif outside
  security-level 10
  ip address 10.8.76.35 255.255.255.248 standby 10.8.76.36
interface Vlan617
  nameif inside
  security-level 100
  ip address 10.8.76.43 255.255.255.248 standby 10.8.76.44
access-list inside-acl extended permit ip any any
access-list inside-acl extended permit udp any any
access-list inside-acl extended permit tcp any any
access-list inside-acl extended permit icmp any any
access-list outside-acl extended permit ip any any
access-list outside-acl extended permit udp any any
access-list outside-acl extended permit tcp any any
access-list outside-acl extended permit icmp any any
access-list mgmt-acl extended permit icmp any any
access-list mgmt-acl extended permit tcp any any
access-list mgmt-acl extended permit udp any any
access-list mgmt-acl extended permit ip any any
pager lines 24
logging buffered debugging
mtu outside 1500
mtu inside 1500
icmp unreachable rate-limit 1 burst-size 1
icmp permit any outside
icmp permit any inside
no asdm history enable
arp timeout 1500
access-group outside-acl in interface outside per-user-override
access-group inside-acl in interface inside per-user-override
route outside 0.0.0.0 0.0.0.0 10.8.76.33 1
route inside 10.9.104.0 255.255.248.0 10.8.76.41 1

```

```

timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:01:00
timeout floating-conn 0:00:00
no snmp-server location
no snmp-server contact
telnet timeout 5
ssh timeout 5
no threat-detection statistics tcp-intercept
class-map inspection_default
  match default-inspection-traffic
policy-map type inspect dns preset_dns_map
  parameters
    message-length maximum client auto
    message-length maximum 512
policy-map global_policy
  class inspection_default
    inspect dns preset_dns_map
    inspect ftp
    inspect h323 h225
    inspect h323 ras
    inspect ip-options
    inspect netbios
    inspect rsh
    inspect rtsp
    inspect skinny
    inspect esmtp
    inspect sqlnet
    inspect sunrpc
    inspect tftp
    inspect sip
    inspect xdmcp
service-policy global_policy global
Cryptochecksum:bc9bd98b298c334af9d76007204a043b
: end

```

Cisco ASR 9006 configurations:

```

!! IOS XR Configuration 4.2.0
!! Last configuration change at Mon Jan 14 18:43:43 2013 by admin
hostname VXIaaS-DC-9K-1
logging disable
domain name vxiaas.local
cdp
vrf GT
vrf T1
vrf T2
vrf T3
vrf T4
vrf T5
vrf management
line console
  exec-timeout 0 0
  length 32
interface Loopback1
  description Tenant 1 VRF Loopback
  vrf T1
  ipv4 address 192.168.10.1 255.255.255.255
interface Loopback2
  description Tenant 2 VRF Loopback
  vrf T2

```

```
    ipv4 address 192.168.20.1 255.255.255.255
interface Loopback3
  description Tenant 3 VRF Loopback
  vrf T3
  ipv4 address 192.168.30.1 255.255.255.255
interface Loopback4
  description Tenant 4 VRF Loopback
  vrf T4
  ipv4 address 192.168.40.1 255.255.255.255
interface Loopback5
  description Tenant 5 VRF Loopback
  vrf T5
  ipv4 address 192.168.50.1 255.255.255.255
interface Loopback8
  description Global Tenant VRF Loopback
  vrf GT
  ipv4 address 192.168.80.1 255.255.255.255
interface MgmtEth0/RSP0/CPU0/0
  description MGMT Interface
  vrf management
  ipv4 address 10.8.66.15 255.255.255.0
interface MgmtEth0/RSP0/CPU0/1
  shutdown
interface GigabitEthernet0/1/0/0
  description Global Tenant Internet
  vrf GT
  ipv4 point-to-point
  ipv4 address 192.168.200.2 255.255.255.252
interface GigabitEthernet0/1/0/1
  shutdown
interface GigabitEthernet0/1/0/2
  shutdown
interface GigabitEthernet0/1/0/3
  shutdown
interface GigabitEthernet0/1/0/4
  shutdown
interface GigabitEthernet0/1/0/5
  shutdown
interface GigabitEthernet0/1/0/6
  shutdown
interface GigabitEthernet0/1/0/7
  shutdown
interface GigabitEthernet0/1/0/8
  shutdown
interface GigabitEthernet0/1/0/9
  shutdown
interface GigabitEthernet0/1/0/10
  shutdown
interface GigabitEthernet0/1/0/11
  shutdown
interface GigabitEthernet0/1/0/12
  shutdown
interface GigabitEthernet0/1/0/13
  shutdown
interface GigabitEthernet0/1/0/14
  shutdown
interface GigabitEthernet0/1/0/15
  shutdown
interface GigabitEthernet0/1/0/16
  shutdown
interface GigabitEthernet0/1/0/17
  shutdown
interface GigabitEthernet0/1/0/18
  shutdown
```

```
interface GigabitEthernet0/1/0/19
 shutdown
interface GigabitEthernet0/1/0/20
 shutdown
interface GigabitEthernet0/1/0/21
 shutdown
interface GigabitEthernet0/1/0/22
 shutdown
interface GigabitEthernet0/1/0/23
 shutdown
interface GigabitEthernet0/1/0/24
 shutdown
interface GigabitEthernet0/1/0/25
 shutdown
interface GigabitEthernet0/1/0/26
 shutdown
interface GigabitEthernet0/1/0/27
 shutdown
interface GigabitEthernet0/1/0/28
 shutdown
interface GigabitEthernet0/1/0/29
 shutdown
interface GigabitEthernet0/1/0/30
 shutdown
interface GigabitEthernet0/1/0/31
 shutdown
interface GigabitEthernet0/1/0/32
 shutdown
interface GigabitEthernet0/1/0/33
 shutdown
interface GigabitEthernet0/1/0/34
 shutdown
interface GigabitEthernet0/1/0/35
 shutdown
interface GigabitEthernet0/1/0/36
 shutdown
interface GigabitEthernet0/1/0/37
 shutdown
interface GigabitEthernet0/1/0/38
 shutdown
interface GigabitEthernet0/1/0/39
 shutdown
interface TenGigE0/0/0/0
 description VXIaaS-N7K-1
 cdp
interface TenGigE0/0/0/0.712
 description T1 N7K-1
 vrf T1
 ipv4 point-to-point
 ipv4 address 10.8.78.2 255.255.255.252
 encapsulation dot1q 712
interface TenGigE0/0/0/0.722
 description T2 N7K-1
 vrf T2
 ipv4 point-to-point
 ipv4 address 10.8.78.34 255.255.255.252
 encapsulation dot1q 722
interface TenGigE0/0/0/0.732
 description T3 N7K-1
 vrf T3
 ipv4 point-to-point
 ipv4 address 10.8.78.66 255.255.255.252
 encapsulation dot1q 732
interface TenGigE0/0/0/0.742
```

```
description T4 N7K-1
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.98 255.255.255.252
encapsulation dot1q 742
interface TenGigE0/0/0/0.752
description T5 N7K-1
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.130 255.255.255.252
encapsulation dot1q 752
interface TenGigE0/0/0/0.782
description Global Tenant N7K-1
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.226 255.255.255.252
encapsulation dot1q 782
interface TenGigE0/0/0/1
description VXIaaS-N1K-2
cdp
interface TenGigE0/0/0/1.714
description T1-N7K-2
vrf T1
ipv4 point-to-point
ipv4 address 10.8.78.10 255.255.255.252
encapsulation dot1q 714
interface TenGigE0/0/0/1.724
description T2-N7K-2
vrf T2
ipv4 point-to-point
ipv4 address 10.8.78.42 255.255.255.252
encapsulation dot1q 724
interface TenGigE0/0/0/1.734
description T3-N7K-2
vrf T3
ipv4 point-to-point
ipv4 address 10.8.78.74 255.255.255.252
encapsulation dot1q 734
interface TenGigE0/0/0/1.744
description T4-N7K-2
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.106 255.255.255.252
encapsulation dot1q 744
interface TenGigE0/0/0/1.754
description T5-N7K-2
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.138 255.255.255.252
encapsulation dot1q 754
interface TenGigE0/0/0/1.784
description GT-N7K-2
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.234 255.255.255.252
encapsulation dot1q 784
interface TenGigE0/0/0/2
cdp
shutdown
interface TenGigE0/0/0/3
cdp
shutdown
interface TenGigE0/0/0/4
description SP-ASR9K-1
```

```

cdp
interface TenGigE0/0/0/4.511
  vrf T1
  ipv4 point-to-point
  ipv4 address 10.8.2.2 255.255.255.252
  encapsulation dot1q 511
interface TenGigE0/0/0/4.521
  vrf T2
  ipv4 point-to-point
  ipv4 address 10.8.2.10 255.255.255.252
  encapsulation dot1q 521
interface TenGigE0/0/0/4.531
  vrf T3
  ipv4 point-to-point
  ipv4 address 10.8.2.18 255.255.255.252
  encapsulation dot1q 531
interface TenGigE0/0/0/4.541
  vrf T4
  ipv4 point-to-point
  ipv4 address 10.8.2.26 255.255.255.252
  encapsulation dot1q 541
interface TenGigE0/0/0/4.551
  vrf T5
  ipv4 point-to-point
  ipv4 address 10.8.2.34 255.255.255.252
  encapsulation dot1q 551
interface TenGigE0/0/0/5
  description SP-ASR9K-2
  cdp
interface TenGigE0/0/0/5.513
  vrf T1
  ipv4 point-to-point
  ipv4 address 10.8.2.66 255.255.255.252
  encapsulation dot1q 513
interface TenGigE0/0/0/5.523
  vrf T2
  ipv4 point-to-point
  ipv4 address 10.8.2.74 255.255.255.252
  encapsulation dot1q 523
interface TenGigE0/0/0/5.533
  vrf T3
  ipv4 point-to-point
  ipv4 address 10.8.2.82 255.255.255.252
  encapsulation dot1q 533
interface TenGigE0/0/0/5.543
  vrf T4
  ipv4 point-to-point
  ipv4 address 10.8.2.90 255.255.255.252
  encapsulation dot1q 543
interface TenGigE0/0/0/5.553
  vrf T5
  ipv4 point-to-point
  ipv4 address 10.8.2.98 255.255.255.252
  encapsulation dot1q 553
interface TenGigE0/0/0/6
  description VXIaaS-ASR9K-2
  cdp
interface TenGigE0/0/0/6.716
  description T1 ASR9K-2
  vrf T1
  ipv4 point-to-point
  ipv4 address 10.8.78.17 255.255.255.252
  encapsulation dot1q 716
interface TenGigE0/0/0/6.726

```

```
description T2 ASR9K-2
vrf T2
ipv4 point-to-point
ipv4 address 10.8.78.49 255.255.255.252
encapsulation dot1q 726
interface TenGigE0/0/0/6.736
description T3 ASR9K-2
vrf T3
ipv4 point-to-point
ipv4 address 10.8.78.81 255.255.255.252
encapsulation dot1q 736
interface TenGigE0/0/0/6.746
description T4 ASR9K-2
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.113 255.255.255.252
encapsulation dot1q 746
interface TenGigE0/0/0/6.756
description T5 ASR9K-2
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.145 255.255.255.252
encapsulation dot1q 756
interface TenGigE0/0/0/6.786
description GT ASR9K-2
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.241 255.255.255.252
encapsulation dot1q 786
interface TenGigE0/0/0/7
description VXIaaS-ASR9K-2
cdp
interface TenGigE0/0/0/7.717
description T1 ASR9K-2
vrf T1
ipv4 point-to-point
ipv4 address 10.8.78.21 255.255.255.252
encapsulation dot1q 716
interface TenGigE0/0/0/7.727
description T2 ASR9K-2
vrf T2
ipv4 point-to-point
ipv4 address 10.8.78.53 255.255.255.252
encapsulation dot1q 726
interface TenGigE0/0/0/7.737
description T3 ASR9K-2
vrf T3
ipv4 point-to-point
ipv4 address 10.8.78.85 255.255.255.252
encapsulation dot1q 736
interface TenGigE0/0/0/7.747
description T4 ASR9K-2
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.117 255.255.255.252
encapsulation dot1q 746
interface TenGigE0/0/0/7.757
description T5 ASR9K-2
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.149 255.255.255.252
encapsulation dot1q 756
interface TenGigE0/0/0/7.787
description GT ASR9K-2
```

```
vrf GT
  ipv4 point-to-point
  ipv4 address 10.8.78.245 255.255.255.252
  encapsulation dot1q 786
rd-set 100:110
  end-set
rd-set 200:220
  end-set
rd-set 300:330
  end-set
rd-set 400:440
  end-set
rd-set 500:550
  end-set
rd-set 800:880
  end-set
route-policy all
  pass
  end-policy
route-policy GT-BGP
  pass
  end-policy
route-policy T1-BGP
  pass
  end-policy
route-policy T2-BGP
  pass
  end-policy
route-policy T3-BGP
  pass
  end-policy
route-policy T4-BGP
  pass
  end-policy
route-policy T5-BGP
  pass
  end-policy
route-policy GT-OSPF
  pass
  end-policy
route-policy T1-OSPF
  pass
  end-policy
route-policy T2-OSPF
  pass
  end-policy
route-policy T3-OSPF
  pass
  end-policy
route-policy T4-OSPF
  pass
  end-policy
route-policy T5-OSPF
  pass
  end-policy
router static
  vrf management
  address-family ipv4 unicast
    0.0.0.0/0 10.8.66.1
router ospf VXIaaS
  vrf GT
  router-id 8.8.8.1
  auto-cost reference-bandwidth 100000
  default-information originate
```

```
redistribute bgp 42 route-policy GT-BGP
area 0
  interface Loopback8
    network point-to-point
    passive enable
  interface TenGigE0/0/0/0.782
    network point-to-point
    neighbor 10.8.78.225
  interface TenGigE0/0/0/1.784
    network point-to-point

neighbor 10.8.78.233
interface TenGigE0/0/0/6.786
  network point-to-point
  neighbor 10.8.78.242
interface TenGigE0/0/0/7.787
  network point-to-point
  neighbor 10.8.78.246

vrf T1
  router-id 1.1.1.1
  auto-cost reference-bandwidth 100000
  default-information originate
  redistribute bgp 42 route-policy T1-BGP
  area 0
    interface Loopback1
      network point-to-point
      passive enable
  interface TenGigE0/0/0/0.712
    network point-to-point
    neighbor 10.8.78.1
  interface TenGigE0/0/0/1.714
    network point-to-point
    neighbor 10.8.78.9
  interface TenGigE0/0/0/6.716
    network point-to-point
    neighbor 10.8.78.18
  interface TenGigE0/0/0/7.717
    network point-to-point
    neighbor 10.8.78.22

vrf T2
  router-id 2.2.2.1
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback2
      network point-to-point
      passive enable
  interface TenGigE0/0/0/0.722
    network point-to-point
    neighbor 10.8.78.33
  interface TenGigE0/0/0/1.724
    network point-to-point
    neighbor 10.8.78.41
  interface TenGigE0/0/0/6.726
    network point-to-point
    neighbor 10.8.78.50
  interface TenGigE0/0/0/7.727
    network point-to-point
    neighbor 10.8.78.54

vrf T3
  router-id 3.3.3.1
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
```

```
interface Loopback3
  network point-to-point
  passive enable
interface TenGigE0/0/0/0.732
  network point-to-point
  neighbor 10.8.78.65
interface TenGigE0/0/0/1.734
  network point-to-point
  neighbor 10.8.78.73
interface TenGigE0/0/0/6.736
  network point-to-point
  neighbor 10.8.78.82
interface TenGigE0/0/0/7.737
  network point-to-point
  neighbor 10.8.78.86
vrf T4
  router-id 4.4.4.1
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback4
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.742
      network point-to-point
      neighbor 10.8.78.97
    interface TenGigE0/0/0/1.744
      network point-to-point
      neighbor 10.8.78.105
    interface TenGigE0/0/0/6.746
      network point-to-point
      neighbor 10.8.78.114
    interface TenGigE0/0/0/7.747
      network point-to-point
      neighbor 10.8.78.118
vrf T5
  router-id 5.5.5.1
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback5
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.752
      network point-to-point
      neighbor 10.8.78.129
    interface TenGigE0/0/0/1.754
      network point-to-point
      neighbor 10.8.78.137
    interface TenGigE0/0/0/6.756
      network point-to-point
      neighbor 10.8.78.146
    interface TenGigE0/0/0/7.757
      network point-to-point
      neighbor 10.8.78.150
router bgp 42
  address-family ipv4 unicast
  address-family vpnv4 unicast
vrf GT
  rd 800:880
  bgp router-id 8.8.8.1
  address-family ipv4 unicast
  redistribute ospf VXiaaS route-policy GT-OSPF
  neighbor 192.168.200.1
```

```
remote-as 40
address-family ipv4 unicast
  route-policy all in
  route-policy all out
vrf T1
  rd 100:110
  bgp router-id 1.1.1.1
  address-family ipv4 unicast
    redistribute ospf VXIaaS route-policy T1-OSPF
neighbor 10.8.2.1
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
neighbor 10.8.2.65
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
vrf T2
  rd 200:220
  bgp router-id 2.2.2.1
  address-family ipv4 unicast
    redistribute ospf VXIaaS route-policy T2-OSPF
neighbor 10.8.2.9
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
neighbor 10.8.2.73
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
vrf T3
  rd 300:330
  bgp router-id 3.3.3.1
  address-family ipv4 unicast
    redistribute ospf VXIaaS route-policy T3-OSPF
neighbor 10.8.2.17
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
neighbor 10.8.2.81
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
vrf T4
  rd 400:440
  bgp router-id 4.4.4.1
  address-family ipv4 unicast
    redistribute ospf VXIaaS route-policy T4-OSPF
neighbor 10.8.2.25
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
neighbor 10.8.2.89
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
```

```

vrf T5
  rd 500:550
  bgp router-id 5.5.5.1
  address-family ipv4 unicast
    redistribute ospf VXIaaS route-policy T5-OSPF
neighbor 10.8.2.33
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
neighbor 10.8.2.97
  remote-as 20
  address-family ipv4 unicast
    route-policy all in
    route-policy all out
ssh server vrf management
end

!! IOS XR Configuration 4.2.0
!! Last configuration change at Mon Jan 14 18:48:08 2013 by admin
hostname VXIaaS-DC-9K-2
logging disable
domain name vxiaas.local
cdp
vrf GT
vrf T1
vrf T2
vrf T3
vrf T4
vrf T5
vrf management
line console
  exec-timeout 0 0
  length 23
interface Loopback1
  description Tenant 1 VRF Loopback
  vrf T1
  ipv4 address 192.168.10.2 255.255.255.255
interface Loopback2
  description Tenant 2 VRF Loopback
  vrf T2
  ipv4 address 192.168.20.2 255.255.255.255
interface Loopback3
  description Tenant 3 VRF Loopback
  vrf T3
  ipv4 address 192.168.30.2 255.255.255.255
interface Loopback4
  description Tenant 4 VRF Loopback
  vrf T4
  ipv4 address 192.168.40.2 255.255.255.255
interface Loopback5
  description Tenant 5 VRF Loopback
  vrf T5
  ipv4 address 192.168.50.2 255.255.255.255
interface Loopback8
  description Global Tenant VRF Loopback
  vrf GT
  ipv4 address 192.168.80.2 255.255.255.255
interface MgmtEth0/RSP0/CPU0/0
  description MGMT Interface
  vrf management
  ipv4 address 10.8.66.16 255.255.255.0
interface MgmtEth0/RSP0/CPU0/1
  shutdown

```

```
interface GigabitEthernet0/1/0/0
description Global Tenant Internet
vrf GT
ipv4 point-to-point
ipv4 address 192.168.200.6 255.255.255.252
interface GigabitEthernet0/1/0/1
shutdown
interface GigabitEthernet0/1/0/2
shutdown
interface GigabitEthernet0/1/0/3
shutdown
interface GigabitEthernet0/1/0/4
shutdown
interface GigabitEthernet0/1/0/5
shutdown
interface GigabitEthernet0/1/0/6
shutdown
interface GigabitEthernet0/1/0/7
shutdown
interface GigabitEthernet0/1/0/8
shutdown
interface GigabitEthernet0/1/0/9
shutdown
interface GigabitEthernet0/1/0/10
shutdown
interface GigabitEthernet0/1/0/11
shutdown
interface GigabitEthernet0/1/0/12
shutdown
interface GigabitEthernet0/1/0/13
shutdown
interface GigabitEthernet0/1/0/14
shutdown
interface GigabitEthernet0/1/0/15
shutdown
interface GigabitEthernet0/1/0/16
shutdown
interface GigabitEthernet0/1/0/17
shutdown
interface GigabitEthernet0/1/0/18
shutdown
interface GigabitEthernet0/1/0/19
shutdown
interface GigabitEthernet0/1/0/20
shutdown
interface GigabitEthernet0/1/0/21
shutdown
interface GigabitEthernet0/1/0/22
shutdown
interface GigabitEthernet0/1/0/23
shutdown
interface GigabitEthernet0/1/0/24
shutdown
interface GigabitEthernet0/1/0/25
shutdown
interface GigabitEthernet0/1/0/26
shutdown
interface GigabitEthernet0/1/0/27
shutdown
interface GigabitEthernet0/1/0/28
shutdown
interface GigabitEthernet0/1/0/29
shutdown
interface GigabitEthernet0/1/0/30
```

```

shutdown
interface GigabitEthernet0/1/0/31
shutdown
interface GigabitEthernet0/1/0/32
shutdown
interface GigabitEthernet0/1/0/33
shutdown
interface GigabitEthernet0/1/0/34
shutdown
interface GigabitEthernet0/1/0/35
shutdown
interface GigabitEthernet0/1/0/36
shutdown
interface GigabitEthernet0/1/0/37
shutdown
interface GigabitEthernet0/1/0/38
shutdown
interface GigabitEthernet0/1/0/39
shutdown
interface TenGigE0/0/0/0
description VXIaaS-N1K-1
cdp
interface TenGigE0/0/0/0.713
description T1 N7K-1
vrf T1
ipv4 point-to-point
ipv4 address 10.8.78.6 255.255.255.252
encapsulation dot1q 713
interface TenGigE0/0/0/0.723
description T2 N7K-1
vrf T2
ipv4 point-to-point
ipv4 address 10.8.78.38 255.255.255.252
encapsulation dot1q 723
interface TenGigE0/0/0/0.733
description T3 N7K-1
vrf T3
ipv4 point-to-point
ipv4 address 10.8.78.70 255.255.255.252
encapsulation dot1q 733
interface TenGigE0/0/0/0.743
description T4 N7K-1
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.102 255.255.255.252
encapsulation dot1q 743
interface TenGigE0/0/0/0.753
description T5 N7K-1
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.134 255.255.255.252
encapsulation dot1q 753
interface TenGigE0/0/0/0.783
description Global Tenant N7K-1
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.230 255.255.255.252
encapsulation dot1q 783
interface TenGigE0/0/0/1
description VXIaaS-N1K-2
cdp
interface TenGigE0/0/0/1.715
description T1 N7K-2
vrf T1

```

```
    ipv4 point-to-point
    ipv4 address 10.8.78.14 255.255.255.252
    encapsulation dot1q 715
interface TenGigE0/0/0/1.725
    description T2 N7K-2
    vrf T2
    ipv4 point-to-point
    ipv4 address 10.8.78.46 255.255.255.252
    encapsulation dot1q 725
interface TenGigE0/0/0/1.735
    description T3 N7K-2
    vrf T3
    ipv4 point-to-point
    ipv4 address 10.8.78.78 255.255.255.252
    encapsulation dot1q 735
interface TenGigE0/0/0/1.745
    description T4 N7K-2
    vrf T4
    ipv4 point-to-point
    ipv4 address 10.8.78.110 255.255.255.252
    encapsulation dot1q 745
interface TenGigE0/0/0/1.755
    description T5 N7K-2
    vrf T5
    ipv4 point-to-point
    ipv4 address 10.8.78.142 255.255.255.252
    encapsulation dot1q 755
interface TenGigE0/0/0/1.785
    description GT-N7K-2
    vrf GT
    ipv4 point-to-point
    ipv4 address 10.8.78.238 255.255.255.252
    encapsulation dot1q 785
interface TenGigE0/0/0/2
    cdp
    shutdown
interface TenGigE0/0/0/3
    cdp
    shutdown
interface TenGigE0/0/0/4
    description SP-ASR9K-1
    cdp
interface TenGigE0/0/0/4.512
    vrf T1
    ipv4 point-to-point
    ipv4 address 10.8.2.6 255.255.255.252
    encapsulation dot1q 512
interface TenGigE0/0/0/4.522
    vrf T2
    ipv4 point-to-point
    ipv4 address 10.8.2.14 255.255.255.252
    encapsulation dot1q 522
interface TenGigE0/0/0/4.532
    vrf T3
    ipv4 point-to-point
    ipv4 address 10.8.2.22 255.255.255.252
    encapsulation dot1q 532
interface TenGigE0/0/0/4.542
    vrf T4
    ipv4 point-to-point
    ipv4 address 10.8.2.30 255.255.255.252
    encapsulation dot1q 542
interface TenGigE0/0/0/4.552
    vrf T5
```

```

    ipv4 point-to-point
    ipv4 address 10.8.2.38 255.255.255.252
    encapsulation dot1q 552
interface TenGigE0/0/0/5
    description SP-ASR9K-2
    cdp
interface TenGigE0/0/0/5.514
    vrf T1
    ipv4 point-to-point
    ipv4 address 10.8.2.70 255.255.255.252
    encapsulation dot1q 514
interface TenGigE0/0/0/5.524
    vrf T2
    ipv4 point-to-point
    ipv4 address 10.8.2.78 255.255.255.252
    encapsulation dot1q 524
interface TenGigE0/0/0/5.534
    vrf T3
    ipv4 point-to-point
    ipv4 address 10.8.2.86 255.255.255.252
    encapsulation dot1q 534
interface TenGigE0/0/0/5.544
    vrf T4
    ipv4 point-to-point
    ipv4 address 10.8.2.94 255.255.255.252
    encapsulation dot1q 544
interface TenGigE0/0/0/5.554
    vrf T5
    ipv4 point-to-point
    ipv4 address 10.8.2.102 255.255.255.252
    encapsulation dot1q 554
interface TenGigE0/0/0/6
    description VXIaaS-ASR9K-1
    cdp
interface TenGigE0/0/0/6.716
    description T1 ASR9K-1
    vrf T1
    ipv4 point-to-point
    ipv4 address 10.8.78.18 255.255.255.252
    encapsulation dot1q 716
interface TenGigE0/0/0/6.726
    description T2 ASR9K-1
    vrf T2
    ipv4 point-to-point
    ipv4 address 10.8.78.50 255.255.255.252
    encapsulation dot1q 726
interface TenGigE0/0/0/6.736
    description T3 ASR9K-1
    vrf T3
    ipv4 point-to-point
    ipv4 address 10.8.78.82 255.255.255.252
    encapsulation dot1q 736
interface TenGigE0/0/0/6.746
    description T4 ASR9K-1
    vrf T4
    ipv4 point-to-point
    ipv4 address 10.8.78.114 255.255.255.252
    encapsulation dot1q 746
interface TenGigE0/0/0/6.756
    description T5 ASR9K-1
    vrf T5
    ipv4 point-to-point
    ipv4 address 10.8.78.146 255.255.255.252
    encapsulation dot1q 756

```

```
interface TenGigE0/0/0/6.786
description GT ASR9K-1
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.242 255.255.255.252
encapsulation dot1q 786
interface TenGigE0/0/0/7
description VXIaaS-ASR9K-1
cdp
interface TenGigE0/0/0/7.717
description T1 ASR9K-1
vrf T1
ipv4 point-to-point
ipv4 address 10.8.78.22 255.255.255.252
encapsulation dot1q 716
interface TenGigE0/0/0/7.727
description T2 ASR9K-1
vrf T2
ipv4 point-to-point
ipv4 address 10.8.78.54 255.255.255.252
encapsulation dot1q 726
interface TenGigE0/0/0/7.737
description T3 ASR9K-1
vrf T3
ipv4 point-to-point
ipv4 address 10.8.78.86 255.255.255.252
encapsulation dot1q 736
interface TenGigE0/0/0/7.747
description T4 ASR9K-1
vrf T4
ipv4 point-to-point
ipv4 address 10.8.78.118 255.255.255.252
encapsulation dot1q 746
interface TenGigE0/0/0/7.757
description T5 ASR9K-1
vrf T5
ipv4 point-to-point
ipv4 address 10.8.78.150 255.255.255.252
encapsulation dot1q 756
interface TenGigE0/0/0/7.787
description GT ASR9K-1
vrf GT
ipv4 point-to-point
ipv4 address 10.8.78.246 255.255.255.252
encapsulation dot1q 786
rd-set 100:110
end-set
rd-set 200:220
end-set
rd-set 300:330
end-set
rd-set 400:440
end-set
rd-set 500:550
end-set
rd-set 800:880
end-set
route-policy all
pass
end-policy
route-policy GT-BGP
pass
end-policy
route-policy T1-BGP
```

```

    pass
  end-policy
route-policy T2-BGP
  pass
  end-policy
route-policy T3-BGP
  pass
  end-policy
route-policy T4-BGP
  pass
  end-policy
route-policy T5-BGP
  pass
  end-policy
route-policy GT-OSPF
  pass
  end-policy
route-policy T1-OSPF
  pass
  end-policy
route-policy T2-OSPF
  pass
  end-policy
route-policy T3-OSPF
  pass
  end-policy
route-policy T4-OSPF
  pass
  end-policy
route-policy T5-OSPF
  pass
  end-policy
router static
  vrf management
    address-family ipv4 unicast
      0.0.0.0/0 10.8.66.1
router ospf VXiaaS
  redistribute bgp 42 route-policy GT-BGP
  vrf GT
  router-id 8.8.8.2
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback8
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.783
      network point-to-point
      neighbor 10.8.78.229
    interface TenGigE0/0/0/1.785
      network point-to-point
      neighbor 10.8.78.237
    interface TenGigE0/0/0/6.786
      network point-to-point
      neighbor 10.8.78.241
    interface TenGigE0/0/0/7.787
      network point-to-point
      neighbor 10.8.78.245
  vrf T1
  router-id 1.1.1.2
  auto-cost reference-bandwidth 100000
  default-information originate
  redistribute bgp 42 route-policy T1-BGP
  area 0

```

```
interface Loopback1
  network point-to-point
  passive enable
interface TenGigE0/0/0/0.713
  network point-to-point
  neighbor 10.8.78.5
interface TenGigE0/0/0/1.715
  network point-to-point
  neighbor 10.8.78.13
interface TenGigE0/0/0/6.716
  network point-to-point
  neighbor 10.8.78.17
interface TenGigE0/0/0/7.717
  network point-to-point
  neighbor 10.8.78.21
vrf T2
  router-id 2.2.2.2
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback2
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.723
      network point-to-point
    interface TenGigE0/0/0/1.725
      network point-to-point
      neighbor 10.8.78.45
    interface TenGigE0/0/0/6.726
      network point-to-point
      neighbor 10.8.78.49
    interface TenGigE0/0/0/7.727
      network point-to-point
      neighbor 10.8.78.53
vrf T3
  router-id 3.3.3.2
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback3
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.733
      network point-to-point
      neighbor 10.8.78.69
    interface TenGigE0/0/0/1.735
      network point-to-point
      neighbor 10.8.78.77
    interface TenGigE0/0/0/6.736
      network point-to-point
      neighbor 10.8.78.81
    interface TenGigE0/0/0/7.737
      network point-to-point
      neighbor 10.8.78.85
vrf T4
  router-id 4.4.4.2
  auto-cost reference-bandwidth 100000
  default-information originate
  area 0
    interface Loopback4
      network point-to-point
      passive enable
    interface TenGigE0/0/0/0.743
      network point-to-point
```

```

neighbor 10.8.78.101
interface TenGigE0/0/0/1.745
network point-to-point
neighbor 10.8.78.109
interface TenGigE0/0/0/6.746
network point-to-point
neighbor 10.8.78.113
interface TenGigE0/0/0/7.747
network point-to-point
neighbor 10.8.78.117
vrf T5
router-id 5.5.5.2
auto-cost reference-bandwidth 100000
default-information originate
area 0
interface Loopback5
network point-to-point
passive enable
interface TenGigE0/0/0/0.753
network point-to-point
neighbor 10.8.78.133
interface TenGigE0/0/0/1.755
network point-to-point
neighbor 10.8.78.141
interface TenGigE0/0/0/6.756
network point-to-point
neighbor 10.8.78.145
interface TenGigE0/0/0/7.757
network point-to-point
neighbor 10.8.78.149
router bgp 42
address-family ipv4 unicast
address-family vpnv4 unicast
vrf GT
rd 800:880
bgp router-id 8.8.8.2
address-family ipv4 unicast
redistribute ospf VXIaaS route-policy GT-OSPF
neighbor 192.168.200.5
remote-as 40
address-family ipv4 unicast
route-policy all in
route-policy all out
vrf T1
rd 100:110
bgp router-id 1.1.1.2
address-family ipv4 unicast
redistribute ospf VXIaaS route-policy T1-OSPF
neighbor 10.8.2.5
remote-as 20
address-family ipv4 unicast
route-policy all in
route-policy all out
neighbor 10.8.2.69
remote-as 20
address-family ipv4 unicast
route-policy all in
route-policy all out
vrf T2
rd 200:220
bgp router-id 2.2.2.2
address-family ipv4 unicast
redistribute ospf VXIaaS route-policy T2-OSPF
neighbor 10.8.2.13

```

```
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
neighbor 10.8.2.77
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
vrf T3
rd 300:330
bgp router-id 3.3.3.2
address-family ipv4 unicast
  redistribute ospf VXIaaS route-policy T3-OSPF
neighbor 10.8.2.21
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
neighbor 10.8.2.85
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
vrf T4
rd 400:440
bgp router-id 4.4.4.2
address-family ipv4 unicast
  redistribute ospf VXIaaS route-policy T4-OSPF
neighbor 10.8.2.29
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
neighbor 10.8.2.93
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
vrf T5
rd 500:550
bgp router-id 5.5.5.2
address-family ipv4 unicast
  redistribute ospf VXIaaS route-policy T5-OSPF
neighbor 10.8.2.37
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
neighbor 10.8.2.101
remote-as 20
address-family ipv4 unicast
  route-policy all in
  route-policy all out
ssh server vrf management
end
```

8 Related Documents

The following links provide more detailed information regarding:

- [Cisco VMDC](#)
- [Cisco UCS](#)
- [Cisco Nexus Switch](#)
- [Cisco Catalyst 6500 Switch](#)
- [Cisco ASR](#)
- [Cisco ASA](#)
- [Desktone](#)

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

