



FlexPod Data Center with VMware vSphere 5.1, and Citrix CloudPlatform 4.2.1

Deployment Guide for FlexPod with VMware vSphere 5.X, and Citrix CloudPlatform 4.2.1 Powered by Apache


Last Updated: February 5, 2014



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors



Deepak Natarajan



Ramesh Kamath



Frank Zhang

Deepak Natarajan, Technical Marketing Engineer, SAVBU, Cisco Systems

Deepak Natarajan is a Technical Marketing Engineer in Cisco Systems Data Center Group (DCG) and specialist on Data Center Solutions. He has over 11 years of extensive experience on Architecture Design, Capacity Planning and Application Performance to build Data Center Cloud. He has worked as a TME and field enablement, customer and field facing, influencing and building Cloud solutions across Compute, Storage, Network, Cloud Management product portfolio in various organizations. He holds a Bachelor degree in Computer Science and is also an Cisco Certified Network Professional (CCNA). Deepak also has strong background in Cisco UCS Systems, Data Center Network, NetApp Storage and Virtualization domain.

Ramesh Kamath, Technical Marketing Engineer, Data Center Solutions Group, NetApp Systems

Ramesh Kamath is a Technical Marketing Engineer for Citrix Solutions at NetApp. He is focused on developing, validating and evangelizing Citrix XenServer, ShareFile and XenDesktop based solutions on NetApp Infrastructure. With over 6 years in the Data Storage Industry, Ramesh holds a Bachelor's Degree from R V College Of Engineering, Bangalore.

Frank Zhang, CloudPlatform Developer, Cloud Platforms Group, Citrix Systems

Frank Zhang is a core CloudPlatform developer in Citrix. He is responsible for developing Baremetal-as-a-Service as well as Cisco UCS integration in CloudPlatform.



Acknowledgment

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

- Mike Brennan
- Siva Sivakumar
- Vadiraja Bhatt
- Roger Barlow
- Manan Shah
- Troy Mangum
- Abhinav Joshi
- Rachel Zhu
- David La Motta
- Kim White
- Animesh Chaturvedi
- Susan Wu
- Parth Jagirdar
- Karthick Radhakrishnan
- Sridhara C



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

Copyright © 2013 Cisco Systems, Inc. All rights reserved.

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.

© 2013 Cisco Systems, Inc. All rights reserved.



FlexPod Data Center with VMware vSphere 5.1, and Citrix CloudPlatform 4.2.1

Executive Summary

This Cisco Validated Design Guide provides design considerations and guidelines for deploying Citrix CloudPlatform 4.2.1 on Cisco Unified Computing System (UCS) and shared NetApp Storage. The document provides information on how to design Compute, Storage, and Network infrastructure layout on Cisco UCS to build Citrix CloudPlatform 4.2.1 over unified protocol (block or file). The deployment scenarios discussed in this document follows the Cisco UCS best practices and recommendations to help ensure that the systems are highly available, scalable, and can be efficiently consolidated and centrally managed in the cloud environment.

Introduction

Companies everywhere are looking for effective strategies to harness the benefits of cloud computing without disrupting current business models. As the industry ventures into the cloud computing era, service providers and entrepreneurs are seeking more efficient and differentiated cloud solutions to reduce the total cost of ownership of IT, acquire the capability to add capacity on demand, charge-back or show-back for services rendered, and attract and retain customers, and increase market share. They need open and flexible cloud solutions that free them from vendor lock-in so they can take full advantage of existing investments and choose the best possible components for their clouds. They need access to source code and open APIs to innovate and build value-added services, all while still having enterprise-class support and services. Their customers want to choose the architecture and hypervisor that's right for them. At the same time, enterprises are looking to the cloud to enable more agile, elastic, on-demand IT services. In both cases, they need the right solutions to build, scale, and manage cloud services.

This joint solution is comprised of several fundamental building blocks, such as, VMware ESXi 5.1 hypervisor, Citrix CloudPlatform 4.2.1, Cisco Unified Computing System, Nexus physical and virtual switches, UCS B-Series Servers, and NetApp FAS 3270 Storage configured with Cluster Mode Data ONTAP 8.1.2. The combination of these building blocks provide a seamless cloud infrastructure management system.



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

Copyright © 2013 Cisco Systems, Inc. All rights reserved.

The reference architecture suggests best practices for configuring and deploying the VMware vSphere ESXi 5.1, Citrix CloudPlatform™ Version 4.2.1 powered by Apache CloudStack on the Cisco Unified Computing System™ (Cisco UCS®) B200 M3 Blade Servers connected to the NetApp FAS 3270 storage array with Cisco Nexus® 5000 Series Switches. Second and third generation Cisco® UCS hardware and software are used in this solution. Citrix CloudPlatform 4.2.1 is a broad solution that includes a commercially certified and packaged Apache CloudStack product. Cisco UCS is an highly scalable, automated, and programmable infrastructure for cloud deployment.

Audience

This document is designed to assist solution architects, sales engineers, field engineers, and consultants with evaluation, planning, design, and deployment of Citrix CloudPlatform Version 4.2.1 on Cisco UCS. The reader should have an architectural understanding of Cisco UCS, Cisco Nexus 5500 Series Switches, Citrix CloudPlatform 4.2.1, NetApp FAS 3270 Storage and NetApp Data ONTAP 8.1.2 Cluster Mode, and related software.

Solution Components

The following components are required to deploy the Citrix CloudPlatform 4.2.1 on the Cisco Unified System design:

- Hardware Components
 - Cisco Unified Computing System
 - Cisco Nexus 5500 Series Switches
 - NetApp FAS Storage System
- Software Components
 - VMware vSphere ESXi 5.1 Hypervisor
 - Citrix CloudPlatform 4.2.1
 - Cisco UCS Manager 2.1(2)

Cisco Unified Computing System

Overview

The Cisco UCS is a next-generation approach to blade and rack server computing. It is an innovative data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Managed as a single system, whether it has one server or 160 servers with thousands of virtual machines, the Cisco UCS decouples scale from complexity. It accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non virtualized systems.

Many of the concepts that make virtualization the ideal platform for delivering cloud apply to Cisco UCS, thus making Cisco UCS also ideal for cloud deployments. Some of the common concepts are:

- Single point of infrastructure management, which enables complete control
- Policy-based infrastructure management
- Using templates to capture desired state and subsequently expedite deployment
- Programmatic control of the entire infrastructure

There are three key areas, inherent in the design, where Cisco UCS excels in comparison to traditional architectures, helping to simplify the design and maintenance of cloud solutions.

The Cisco UCS offers the following features that enable the above concepts in the design solution:

Unified Fabric

Unified fabric can dramatically reduce the number of network adapters, blade-server switches, cables, and management touch points by passing all network traffic to parent Fabric Interconnects, where it can be prioritized, processed, and managed centrally. This approach improves performance, agility, and efficiency and dramatically reduces the number of devices that need to be powered, cooled, secured, and managed.

Embedded Multirole Management

The Cisco UCS Manager is a centralized management application that is embedded on the fabric switch. The Cisco UCS controls all the Cisco UCS elements within a single redundant management domain. These elements include all aspects of system configuration and operation, eliminating the need to use multiple, separate element managers for each system component. Massive reductions in the number of management modules and consoles, and in the proliferation of resident agents on all the hardware (which must be separately managed and updated) are important deliverables of Cisco UCS. Cisco UCS Manager, using role-based access and visibility, helps enable cross-functional communication efficiency and promotes collaboration between data center roles for increased productivity.

Cisco Extended Memory Technology

Significantly enhancing the available memory capacity of some Cisco UCS servers, Cisco Extended Memory Technology helps increase performance for demanding virtualization and large-data-set workloads. Data centers can now deploy very high virtual machine densities on individual servers as well as provide resident memory capacity for databases that need only two processors but can dramatically benefit from more memory. The high-memory dual in-line memory module (DIMM) slot count also lets the users to more cost-effectively scale this capacity using smaller, less costly DIMMs.

Cisco Data Center VM-FEX Virtualization Support and Virtualization Adapter

With Cisco Data Center VM-FEX, virtual machines have virtual links that allow them to be managed in the same way as physical links. Virtual links can be centrally configured and managed without the complexity of traditional systems, which interpose multiple switching layers in virtualized environments. I/O configurations and network profiles move along with virtual machines, helping increase security and efficiency while reducing complexity. Cisco Data Center VM-FEX helps improve performance and reduce network interface card (NIC) infrastructure.

Dynamic Provisioning with Service Profiles

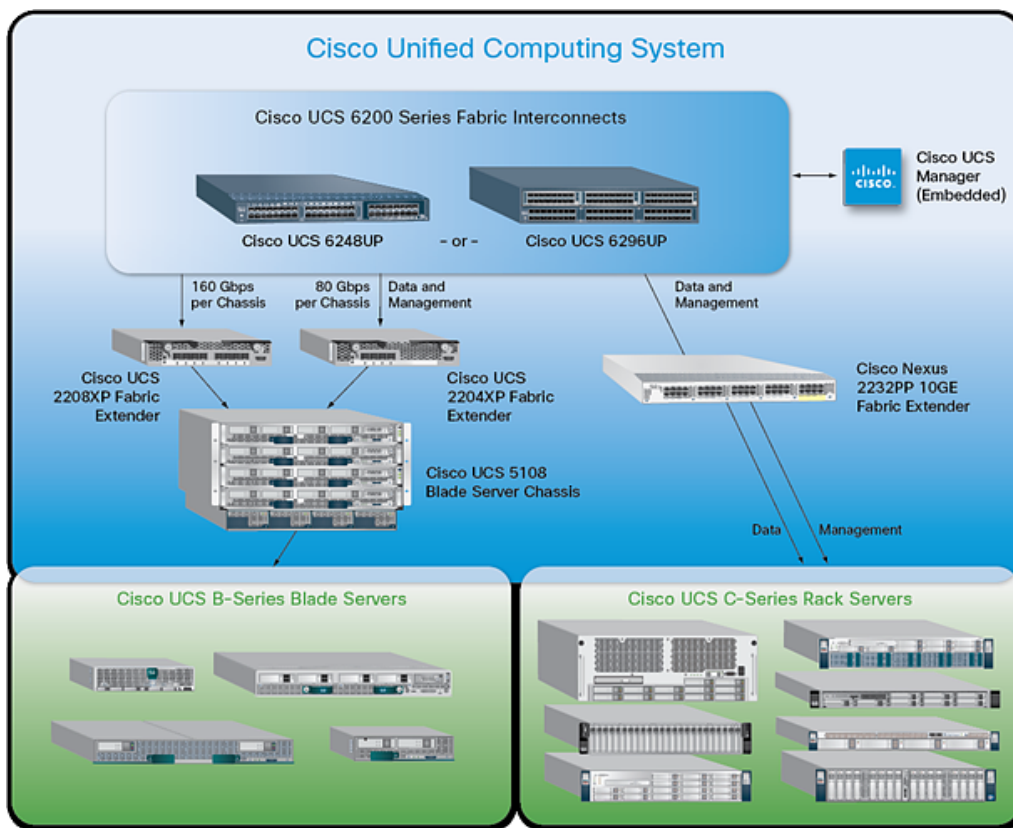
The Cisco UCS Manager delivers service profiles, which contain abstracted server-state information, creating an environment in which everything unique about a server is stored in the fabric, and the physical server is simply another resource to be assigned. Cisco UCS Manager implements role-based and policy-based management focused on service profiles and templates. These mechanisms fully provision one or many servers and their network connectivity in minutes, rather than hours or days.

Cisco UCS Components

This section describes the specific Cisco UCS products used in the Citrix CloudPlatform 4.2.1 on UCS design.

The Cisco Unified Computing System™ with unified fabric dramatically reduces the number of network adapters, blade-server switches, and cables needed as it passes network traffic to parent fabric interconnect. These interconnect process and centrally manage the traffic to improve performance. They also reduce the number of devices that need to be powered, cooled, secured, and managed. The multiple-role management embedded in the Fabric Interconnects to manage configuration and operation, eliminates the need for separate element managers. In addition, Cisco VN-Link Virtualization support gives network links connected to virtual machines the same status as physical links. I/O configurations and network profiles move with the virtual machines, increasing security and efficiency while reducing complexity. This virtual network interface card (vNIC) feature improves performance and reduces NIC infrastructure.

Figure 1 *Cisco UCS Components*



Cisco UCS Fabric Interconnects

Cisco UCS Fabric Interconnects create a unified network fabric throughout Cisco UCS. They provide uniform access to both networks and storage, eliminating the barriers to deployment of a fully virtualized environment based on a flexible, programmable pool of resources. Cisco Fabric Interconnects comprise a family of line-rate, low-latency, lossless 10 Gigabit Ethernet, IEEE Data Center Bridging (DCB), and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus® 5000 Series Switches, Cisco UCS 6200 Series Fabric Interconnects provide additional features and management capabilities that make them the central nervous system of Cisco UCS. The Cisco UCS Manager software runs inside the Cisco UCS Fabric Interconnects. The Cisco UCS 6100 Series Fabric Interconnects expand the Cisco UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS blade server chassis. All chassis and all blades that are attached to interconnect are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6100 Series provides the flexibility to support LAN and SAN connectivity for all blades within its domain at configuration time. Typically deployed in redundant pairs, Cisco UCS Fabric Interconnects provide uniform access to both networks and storage, facilitating a fully virtualized environment.

The Cisco UCS Fabric Interconnects portfolio currently consists of the Cisco 6100 and 6200 Series Fabric Interconnects. In this design we have used the Cisco 6248UP FI:

- Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnects is a one-rack-unit (1RU) 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect providing more than 1-terabit-per-second (Tbps) throughput with low latency. It has 32 fixed ports of Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE Enhanced Small Form-Factor Pluggable (SFP+) ports.

One expansion module slot can provide up to 16 additional Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS Fabric Extenders

The Cisco UCS 2100 and 2200 Series Fabric Extenders multiplex, and forward all traffic from blade servers in a chassis to a parent Cisco UCS Fabric Interconnects over 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis or virtual machines on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the fabric interconnect. At the core of the Cisco UCS Fabric Extenders are application-specific integrated circuit (ASIC) processors developed by Cisco that multiplex all traffic.

Up to two fabric extenders can be placed in a blade chassis. In this design the Cisco UCS 2208XP Fabric Extender

- The Cisco UCS 2208XP Fabric Extender

It has eight 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.

Cisco UCS Virtual Interface Card

The Cisco UCS VIC 1240 and 1280 enable a policy-based, stateless, agile server infrastructure that can present up to 256 PCI Express (PCIe) standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs. In addition, the Cisco UCS VIC 1280 supports Cisco Data Center VM-FEX technology, which extends the Cisco UCS Fabric Interconnects ports to virtual machines, simplifying server virtualization deployment.

- Cisco UCS VIC 1240

A Cisco innovation, the Cisco UCS VIC 1240 is a four-port 10 Gigabit Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

- Cisco UCS VIC 1280

A Cisco innovation, the Cisco UCS VIC 1280 is an eight-port 10 Gigabit Ethernet, FCoE-capable mezzanine card designed exclusively for Cisco UCS B-Series Blade Servers.

Cisco UCS 5100 Series Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a 6RU blade chassis that accepts up to eight half-width Cisco UCS B-Series Blade Servers or up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The Cisco UCS 5108 can accept four redundant power supplies with automatic load sharing and failover and two Cisco UCS 2100 or 2200 Series Fabric Extenders. The chassis is managed by Cisco UCS chassis management controllers, which are mounted in the Cisco UCS fabric extenders and work in conjunction with Cisco UCS Manager to control the chassis and its components.

A single Cisco UCS managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time, Cisco UCS supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows Cisco UCS to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two Cisco UCS Fabric Extenders slots accessible from the back panel
- Eight Cisco UCS Blade Server slots accessible from the front panel

Cisco UCS B200 M3 Blade Servers

The Cisco UCS B200 M3 Blade Server delivers performance, versatility, and density without compromise. It addresses the broadest set of workloads, from IT and web infrastructure to distributed databases. Building on the success of the Cisco UCS B200 M2 Blade Server, the enterprise-class Cisco UCS B200 M3 Blade Server further extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M3 harnesses the power of the latest Intel Xeon processor E5-2600 product family, with up to 384 GB of RAM (using 16-GB DIMMs), two disk drives, and up to dual 4x 10 Gigabit Ethernet throughput. In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches in each blade chassis. With a larger power budget per blade server, Cisco can design uncompromised expandability and capabilities in its blade servers, as evidenced by the new Cisco UCS B200 M3, with its leading memory slot and drive capacity.

Cisco Nexus 5500 Series Switch

The Cisco Nexus 5000 Series Switch is designed for data center environments with cut-through technology that enables consistent, low-latency Ethernet solutions with front-to-back or back-to-front cooling and data ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data-center-class Cisco NX-OS Software for high reliability and ease of management.

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

- Cisco Nexus 5548UP

The Cisco Nexus 5548UP is a 1RU 10 Gigabit Ethernet (10 GE), Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE) switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports and one expansion slot supporting modules with 10 Gigabit Ethernet and FCoE ports or connectivity to Fibre Channel SANs with 8/4/2/1 Gbps Fibre Channel switch ports, or both.

NetApp Storage

The NetApp FAS3270 can handle today's diverse, virtualized workloads and easily respond to future expansion. The NetApp FAS3270 series meets the storage needs of business applications in both virtual and traditional environments in a cost-effective manner. It is ideal for demanding business applications in virtualized environments and dramatically reduces the consumption of raw storage, power, cooling, and space with NetApp FAS3270 highly efficient storage utilization.

NetApp FAS3270 supports the FC, FCoE, IP SAN (iSCSI), NFS, CIFS, HTTP, FTP storage networking. It provides high availability features such as Alternate Control Path (ACP), Ethernet-based service processor and Data ONTAP management interface; redundant hot-swappable controllers, cooling fans, power supplies, and optics. Additionally it supports the highly available controller configurations, such as, active-active controller with controller failover and multipath HA storage, active-active controller with stretch (non switch) and fabric-attached MetroCluster V-Series Storage Acceleration Appliance SA320.

VMware vSphere ESXi 5.1

Virtualization is a proven technology that enables multiple virtual machines to run on a single physical server. Each virtual machine is completely isolated from other machines and is decoupled from the underlying host by a thin layer of software known as a hypervisor. This allows each virtual machine to run different operating systems and applications. Because the machines have been decoupled from the underlying host, the guest can also be moved from one physical server host to another in production or development environment, this is known as live migration. These attributes have transformed the organization's approach to virtual computing. The range of products offered by VMware meet the business needs of an ever evolving IT Infrastructure.

Built on the powerful vSphere hypervisor, ESXi 5.1 is a complete, managed server virtualization platform. ESXi technology is extensively acknowledged as the fastest and most secure virtualization software in the industry. ESXi efficiently manages virtual servers and provides an economical server consolidation ensuring business continuity.

Citrix CloudPlatform 4.2.1

Citrix CloudPlatform, powered by Apache Cloudstack, is the only future-proofed, application-centric cloud solution proven to reliably orchestrate both existing scale-up enterprise workloads and scale-out cloud-native application workloads within a single unified cloud management platform. Citrix CloudPlatform combines the best private cloud foundation for enterprise workloads like CRM and ERP with true Amazon-style scale, elasticity and operational efficiency for cloud-native application workloads like social applications, Big Data and HPC.

Enabling Scale-up and Scale-out Workloads with CloudPlatform 4.2.1

Citrix CloudPlatform enables organizations to create workload specific availability zones or regions which can support high performance, massively parallel workloads on distributed, low cost infrastructure. Enterprises can run complex enterprise workloads with best in class virtualization and networking to deliver true private cloud isolation, ensure SLA, compliance, security and availability for running mission critical workloads. For organizations that are delivering cloud-native application workloads, Citrix CloudPlatform provides customers with availability, storage and network with proven scalability for public and private clouds with over 40,000 hosts per region.

Dedicated Private Cloud Isolation

Cloud infrastructure layers—zones, pods, clusters and hosts—and virtual machine resources—CPU, memory, storage and network—can be granularly grouped into different isolated logical partitions for true multi-tenant private cloud deployments. Dedicated resources can be applied to any of the infrastructure layers to support virtual private cloud use cases to meet requirements for compliance, security and performance.

Logical Isolation

With CloudPlatform, enterprises can organize their private cloud into multiple discrete logical abstraction layers including availability zones or regions, each comprising resources in one or more physical data centers. CloudPlatform uses a logical user isolation hierarchy that includes Domain, Sub-domain, Account and Users. Usage quotas can be applied to any layer in the hierarchy. This logical isolation hierarchy can be used to model the organizational structure in typical enterprises in a domain to represent a business unit; users can also be grouped into Accounts which could represent a team. A sub-domain could represent a division in a larger business unit. The same isolation hierarchy can be applied to a service provider context in which an Account could represent an individual customer. CloudPlatform ensures that all memory, CPU, network, and storage resources are both available and isolated from one user account to another.

Regions

Cloud Administrators can define availability regions consisting of multiple zones and/or data centers. The benefits to implementing regions are increased scalability and availability, geographic availability, lower latency and ensuring corporate compliance. Regions enable application workloads to be launched and deployed across multiple availability zones from the same template. Application availability would not be impacted if a given zone were to fail because the VMs supporting the service will be running in multiple zones.

Hypervisor Agnostic

CloudPlatform supports the leading commercially supported hypervisors including Citrix® XenServer™, VMware® vSphere® Oracle® VM (OVM), and KVM. Bare metal infrastructure without a hypervisor is also supported. Customers have complete freedom to choose the right hypervisor or hypervisors for their workload instead of being locked into technology from one single vendor.

Traditional application availability zones typically begin with bare metal or a supported hypervisor, such as VMware® vSphere or Citrix XenServer™ which supports live migration of VMs. CloudPlatform has a two-tier storage hierarchy. Root and data volumes are stored on a primary storage tier that typically resides on host local storage, FC, iSCSI or NFS for performance. VM templates and volume snapshots are stored on a secondary storage tier, typically NFS or object storage.

Virtual Machine Operations and Management

CloudPlatform provides efficient lifecycle management of virtual machines, including creation and maintenance under a single platform. Root and data volumes reside on the primary storage tier and are created at VM creation. When a VM is destroyed, the root volume is also destroyed. In the case of data volumes that are attached to the VM, they do not get destroyed when the VM is destroyed which means

that administrators can take scheduled or adhoc snapshots of the volume to preserve configuration states for backup or data recovery. OS and ISO templates can be imported, created, and stored across zones or regions on the secondary storage tier for optimized storage efficiency.

Advanced VMware Integration

Virtualization has been broadly adopted by enterprises and service providers alike over the past decade. With the VMware® vCenter® integration, CloudPlatform can help organizations move their enterprise workloads seamlessly from the data center to the cloud and still leverage their existing VMware investments, configurations and expertise. vCenter® features such as virtual hardware hot add, vMotion®, VMware HA & DRS, Storage vMotion, and CPU and RAM overcommit are all enabled in CloudPlatform. CloudPlatform leverages VMware's dvSwitch and PVLAN features to provide further network segmentation and VM isolation on the same network. Additional storage integrations allow for VM level snapshots as well as volume level snapshots.

Guest Isolation using Security Groups

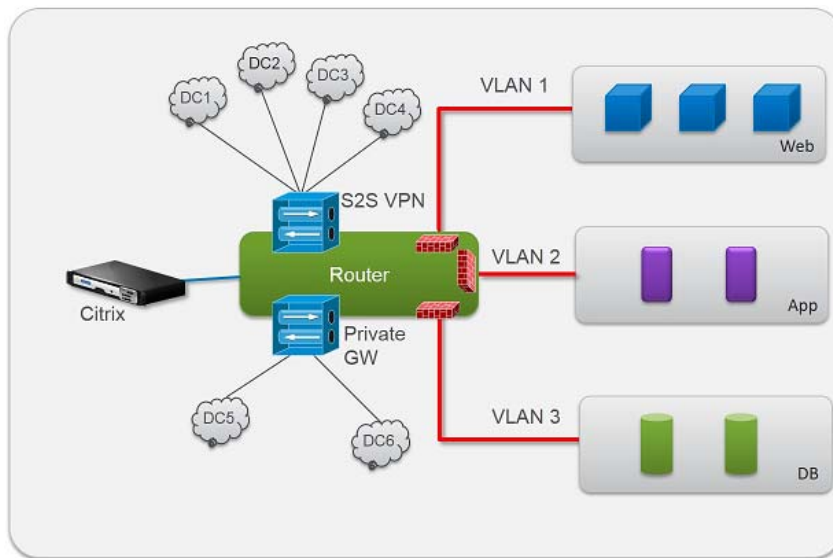
Guest Instances require isolation from other instances running within the same Zone. When guest VMs need to communicate with each other over a network, CloudPlatform provides two isolation methods, Security Groups and VLANs. Although Security Groups can be implemented in both the Basic and Advanced Networking models, VLANs are only available in advanced networking. Advanced Networking allows the cloud administrator to provide custom IP ranges for different accounts. In a Basic Networking setup, the guest instances use the same IP range as the underlying CloudPlatform and Hypervisor architecture.

When using Security Groups, each account has a default Security Group that is automatically created. When new instances are created, they are assigned to one or more Security Groups. Users can create additional Security Groups at any time. Communication between guest instances can happen only if they are assigned to the same security group. The use of Ingress and Egress rules on the Security Group control the flow of traffic, both in and out of the group.

Virtual Private Clouds

Virtual Private Cloud (VPC) is a private, isolated grouping of resources in CloudPlatform. Enterprise applications in traditional data centers have network tiers connected to them for connection to databases, load balancing, and firewalls. A VPC can have its own virtual network topology that resembles a traditional physical network.

Using the VPC feature, cloud administrators can launch VMs in a virtual network containing private addresses to recreate the network architecture of their traditional enterprise applications, including IP ranges and VLANs. The ability to group similar kinds of instances based on IP address range or network tier definition makes the data center transformation to the cloud computing easier. Through the recreation of traditional enterprise network topologies including granular network segmentation for traditional applications, the transition to cloud computing can be simplified.

Figure 2 *n-Tier Applications*

In traditional multi-tier application, VLANs have been traditionally used to produce network segregation and correct traffic flow. In the above multi-tier example, the application is composed of a web front end tier, an application tier and a database tier. Each tier is isolated by an individual VLAN. For large scale deployments, administrators using this n-tier application architecture in Citrix CloudPlatform to create granular networks, routing, firewall and load balancing policies, or affinity/anti-affinity rules for individual tiers.

Advanced Cloud Networking

Cloud operators can create advanced cloud networking configurations and Network-as-a-Service offerings including Portable IP capabilities, global server load balancing (GSLB), and AWS-like Health Checks to ensure application availability. In addition, CloudPlatform has an in-built virtual router that provides granular control of network services like DHCP, Network Address Translation (NAT), load balancing, firewall, and port forwarding.

CloudPlatform integrates with enterprise class Application Delivery Controllers (ADCs) to provide server load balancing. Advanced features such as compression, connection multiplexing, caching, and SSL offload found in ADCs are becoming increasingly helpful off load overburdened networks and servers and increasing application availability. CloudPlatform supports a broad ecosystem of ADCs and networking devices including Citrix Netscaler®, F5 load balancers, Cisco hardware and software (UCS, Nexus 1000v, ASA1000v), Juniper firewalls and VMware Distributed Virtual Switch. Advanced networking requirements for architecture, scale, SLAs, load balancing can be met using best of breed networking solution for the most demanding cloud-native application workloads and availability requirements.

In particular, the Citrix NetScaler Global Server Load Balancing (GSLB) feature is enabled through CloudPlatform which enables distribution of traffic across multiple sites and helps to manage disaster recovery. GSLB works by controlling how the system routes incoming client requests by directing DNS requests to the best-performing GSLB site in a distributed Internet environment. Pre-defined NetScaler policies and configurations can be orchestrated by CloudPlatform to send traffic. For example, a policy could direct traffic to the closest availability zone, a region with the lowest latency or the least amount of load, or to a secondary data center in case of an outage.

Intelligent load balancers such as Citrix Netscaler can be configured to perform AWS-style Health Checks on backend services through CloudPlatform. NetScaler will perform periodic checks on backend services based on a set of service level parameters to be monitored. When a VM fails the Health Check, Citrix Netscaler will automatically remove the VM from the load balancer pool and route the incoming requests only onto healthy VMs. Once the VM successfully passes the health checks again, the load balancer will add the VMs back into the resource pool.

CloudPlatform with Citrix NetScaler offers AutoScale technology that automatically expands and contracts the cloud according to business demands. Citrix Netscaler has the ability to monitor CPU usage, server health or application responsiveness. Working in unison with CloudPlatform, changes to application load can prompt Citrix Netscaler to scale up or scale down the corresponding backend services or guest VMs.

Citrix NetScaler comes with a choice of configurations, as a physical appliance or as a Virtual Machine that runs on Citrix XenServer. CloudPlatform treats Citrix NetScaler just like any other infrastructure resource, for which it can be added into the resource pool.

Object Storage

Cloud-native application workloads that make use of object storage will have transparent access to storage objects across geographic and logically defined locations. The ability to access object storage in a region or across multiple zones increases workload availability and operations efficiency. Object Storage can be used to store persistent data given a zone failure. Also the same object storage can provide secondary storage for Infrastructure-as-a Service and Storage-as-a-Service.

Cisco and Citrix - Better Together for production clouds

With over 200+ clouds in production, Citrix CloudPlatform is the trusted cloud management platform to orchestrate the world's most demanding workloads. Leading telecommunications companies and higher education institutions like Bell Canada, BT, COLT, and University of Sao Paulo to web-centric companies like Edmunds.com have chosen Cisco UCS and Citrix CloudPlatform to run their clouds in production and at scale.

Cisco UCS Manager

Cisco UCS Manager is an embedded, unified manager that provides a single point of management for Cisco UCS. Cisco UCS Manager can be accessed through an intuitive GUI, a command-line interface (CLI), or the comprehensive open XML API. It manages the physical assets of the server and storage and LAN connectivity, and it is designed to simplify the management of virtual network connections through integration with several major hypervisor vendors. It provides IT departments with the flexibility to allow people to manage the system as a whole, or to assign specific management functions to individuals based on their roles as managers of server, storage, or network hardware assets. It simplifies operations by automatically discovering all the components available on the system and enabling a stateless model for resource use.

The elements managed by Cisco UCS Manager include:

1. CloudPlatform allows to provision bare metal host as a service, in order to enable easy expansion of the cloud by leveraging the programmability of the UCS converged infrastructure. CloudPlatform provides UCS Plugin which can automatically understand the UCS environment, server profiles, etc. so CloudPlatform administrators can deploy a bare metal OS on a Cisco UCS System. An overview of the steps involved in using UCS with CloudPlatform:

2. Set up your Cisco UCS Blade Servers, profiles, and Cisco UCS Manager according to Cisco documentation
3. Registering the Cisco UCS Manager with CloudPlatform
4. Associate a Service Profile with a UCS blade
5. Provision the blade as a bare metal host as described

Reference Architecture

This section presents physical and logical high-level design considerations for Cisco UCS networking and computing with VMware ESXi 5.1 virtualization on NetApp storage with Citrix CloudPlatform to build private cloud deployments.

Using Cisco's Unified Fabric for the networking component of the cloud infrastructure helps to ensure that once cables are connected, they need not be rerouted if there are changes in the cloud workload requirements. The Cisco UCS helps to ensure that the servers can be delivered quickly and in an automated fashion.

Citrix CloudPlatform 4.2.1 on FlexPod Overview

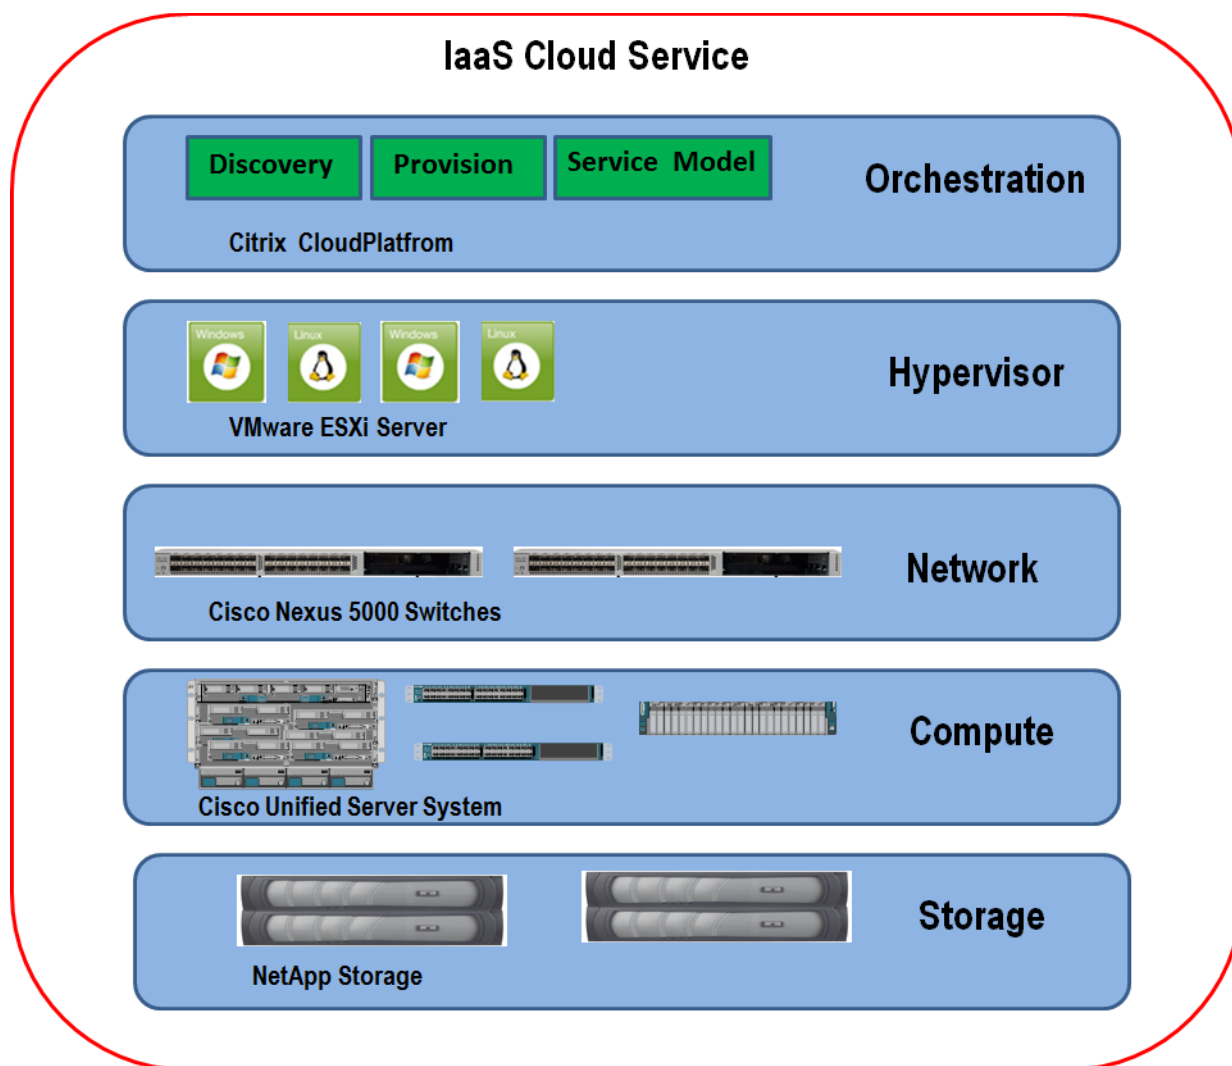
Cisco Flexpod and Citrix CloudPlatform powered by Apache CloudStack combines converged converged compute, networking, and storage solution developed by Cisco and NetApp with a proven, application centric, turn-key orchestration solution from Citrix offers true secure and scalable IT infrastructure with on-demand access to virtual server and storage resources.

Well suited for Enterprise private clouds, CloudPlatform provides a system that allows for delivering capacity on demand, allowing users to provision their own resources based on policy, and charge-back or show back to internal customers, departments or business units.

This document describes design methodology followed to deployment of IaaS (Infrastructure as a Service) based on the components below to build scalable, dynamic and multi-tenant cloud solution deployment models:

- Cisco Unified Computing System
- Cisco Nexus physical and virtual Switches
- NetApp FAS Storage
- VMware vSphere ESXi Hypervisor
- Citrix CloudPlatform 4.2.1

Figure 3 *IaaS Cloud Service Model*



Cisco Unified Compute System

Cisco UCS is designed from the start to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards (VICs), even the number and type of I/O interfaces is programmed dynamically, making all the servers ready to power up and take any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration and associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation accelerates provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Nexus Unified Network System

Cisco Unified Fabric delivers reliable, scalable, agile, and cost-effective network services to servers, storage, and applications while improving the user experience. It facilitates better support of virtualization and cloud services with improved staff utilization, more efficient resource utilization (more load on servers and storage), low-latency options, lower TCO, and better resiliency and uptime. It offers high-performance, low-latency, and highly available networks. These networks serve diverse data center needs, including the lossless requirements for block-level storage traffic. A Cisco Unified Fabric network carries multiprotocol traffic to connect storage (Fibre Channel, FCoE, Small Computer System Interface over IP [iSCSI], and network-attached storage [NAS]) as well as general data traffic. Fibre Channel traffic can be on its own fabric or part of a converged fabric with FCoE. Offering the best of both LAN and SAN environments, Cisco Unified Fabric enables storage network users to take advantage of the economy of scale, robust vendor community, and aggressive performance roadmap of Ethernet while providing the high-performance, lossless characteristics of a Fibre Channel network.

NetApp Unified Storage System

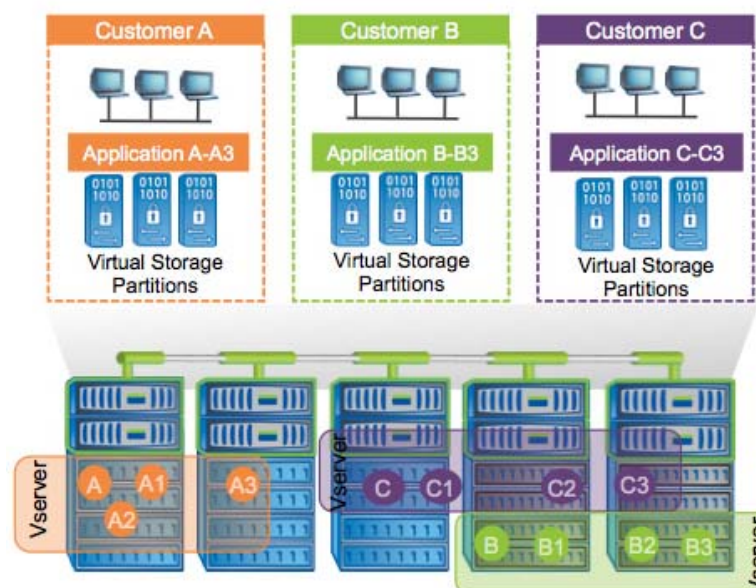
NetApp solutions begin with Clustered Data ONTAP 8.2.1, the fundamental software platform that runs on all NetApp storage systems. Data ONTAP 8.2.1 is a highly optimized, scalable operating system that supports mixed NAS and SAN environments and a range of protocols, including Fibre Channel, iSCSI, FCoE, NFS, and CIFS.

Clustered Data ONTAP offers compelling value to both the enterprise private cloud and public cloud service providers. The software's unified architecture services a range of workloads—from the most demanding to the most unique across virtualized environments, scale-out NAS, and enterprise applications. During the past year companies across all industries have adopted clustered Data ONTAP with a 4x increase in clustered nodes resulting in 100% quarter over quarter growth of deployments. Cloud providers, whose business reputation depends on continuously running operations, are bringing centralized management to the data center with FlexPod® from NetApp and Cisco with clustered Data ONTAP. The new software gives organizations and cloud service providers the capability to rapidly and cost effectively deliver new services and capacity with maximum application uptime. Clustered Data ONTAP 8.2 cuts through the performance, availability, and efficiency limits of traditional hardware silos, empowering IT to nondisruptively align the storage infrastructure with changing business and application demands.

Industry leading capabilities in a clustered data ONTAP are:

- Nondisruptive operations with greater than five-9s reliability (99.999%), providing continuous data access during scheduled downtime and dynamic load balancing without disruptive data migrations
- Seamless scalability with up to 69PB of storage and 24 controller nodes, 49,000 LUNs, 12,000 NAS volumes supporting over 100,000 clients and single container up to 20 PB
- Proven storage and operational efficiency within a unified and multi-tenant architecture that meets the needs of enterprises, small, and midsized businesses

Figure 4 NetApp Clustered ONTAP Storage Architecture



An array of NetApp tools and enhancements are available to augment the storage platform. These tools assist in deployment, backup, recovery, replication, management, and data protection. This solution makes use of a subset of these tools and enhancements. Through the use of secure multitenancy, cloud deployments can successfully provide shared storage infrastructure, a core component of cloud by providing secure tenant data separation, address space, authentication, high availability and data protection.

VMware ESXi Server Cloud Hypervisor System

VMware ESXi hypervisor is a complete, managed server virtualization platform built on a powerful ESXi hypervisor. The ESXi technology is widely accredited as the fastest and most secure virtualization software in the industry. ESXi server is designed for efficient management of virtual servers and delivers cost-effective server consolidation and business continuity.

With VMware ESXi Server, Cloud providers can automate key IT processes to improve service delivery and business continuity for virtual environments resulting in both time and money savings while providing more responsive IT services. Some of the important features in a cloud environment are:

- **High Availability** - automatically restarts virtual machines if a failure occurs at the VM, hypervisor, or server level. The auto restart capability allows users to protect all virtualized applications and bring higher levels of availability to the business. Memory Optimization reduces costs and improves application performance and protection by sharing unused server memory between VMs on the host server.
- **Dynamic Workload - Balancing** improves system utilization and increases application performance by automatically balancing two virtual machines within a resource pool. Workload balancing intelligently places VMs on a most suitable host in the resource pool by matching application requirements to the available hardware resources.

- Automated VM- Protection and Recovery cloud administrators - can create snapshot and archival policies. Regularly scheduled snapshots help to protect against data loss in case of a VM failure. The policies established are based on snapshot type, frequency, amount of historical data that is retained, and an archive location. Recovering a VM is completed by simply choosing the last good known archive.
- Storage ESXiMotion moves - live running virtual machines and their associated virtual disk image within and across resource pools leveraging local and shared storage. This enables users to move a VM between tiers of storage when a VM is limited by storage capacity, and perform maintenance and upgrades with zero downtime.
- Site Recovery- provides site-to-site disaster recovery planning and services for virtual environments. Site recovery is easy to set up, fast to recover, and has the ability to frequently test to ensure disaster recovery plans remain valid.
- Host Power management - takes advantage of embedded hardware features to lower data center electricity consumption by dynamically consolidating VMs on fewer systems and then powering off under utilized servers as the demand for services fluctuates.

Citrix CloudPlatform Cloud Orchestrator System

Citrix CloudPlatform allows virtually unlimited computing power in public, private or hybrid deployments, can be accessed on-demand, and provides real-time usage and metering actual data used. It helps to automate the distribution of compute, network, and storage while adhering to defined policies on load balancing, data security, and compliance. It simplifies and accelerates service delivery by combining self-service provisioning with a catalog of custom-built and pre-defined machine images and gain real-time visibility and reporting within your cloud environment to ensure compliance, security and comprehensive customer usage metering.

Build and deploy CloudPlatform includes a management server and the extensions to run industry-standard hypervisor software such as Citrix Xen Server, VMWare vSphere™, and KVM in your cloud. Deployed on a server farm, the management server can manage resources such as hosts, storage devices, and IP addresses.

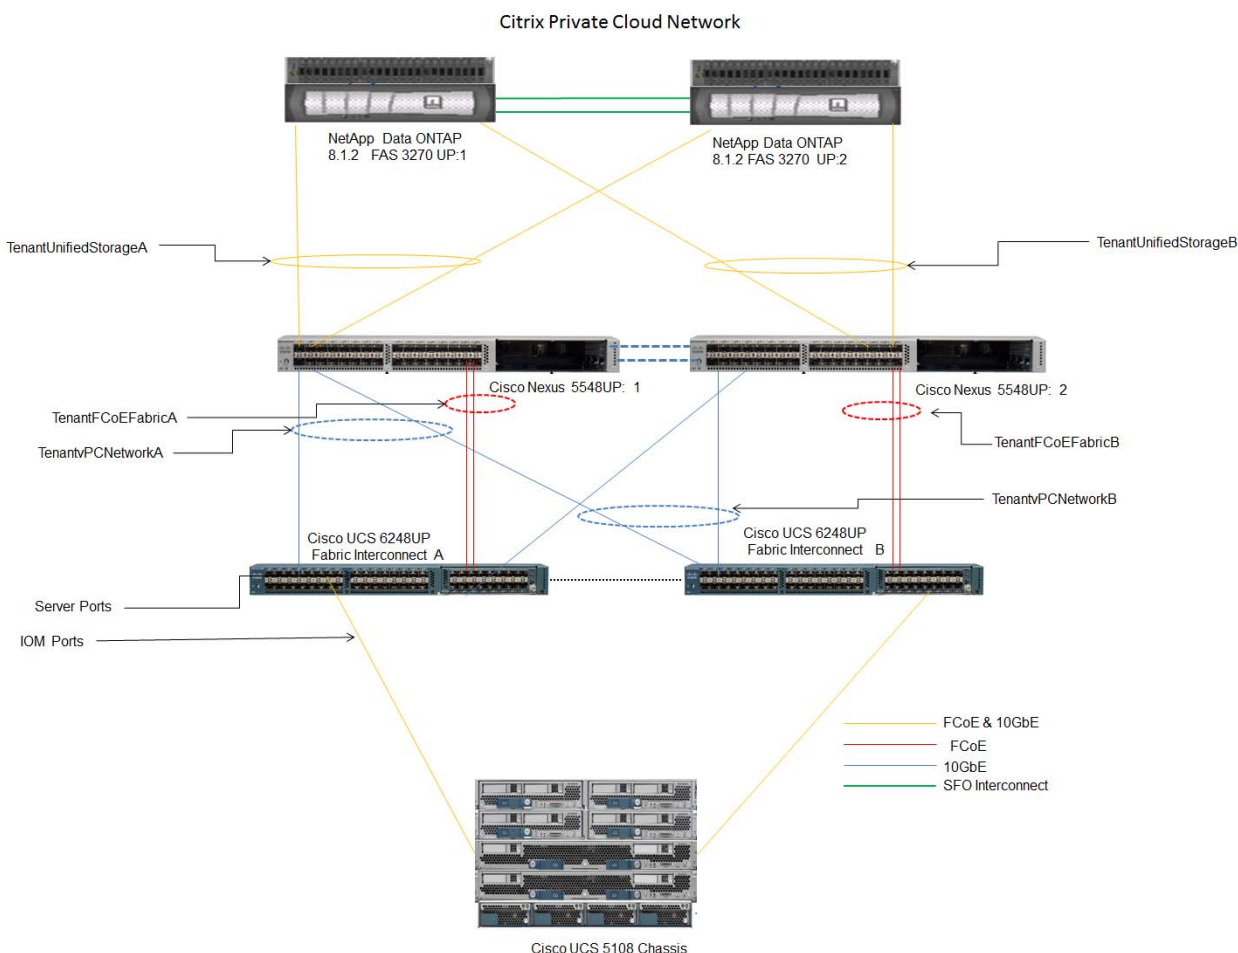
CloudPlatform provides all that is needed to sell on-demand, self-service VM instances, storage volumes and networking configurations over the Internet, or set up an on-premise private cloud. It manages by offering visibility into the aggregate storage, IP pools, CPU, memory and other resources used by the cloud, as well as a chronological view of events within it. Once the cloud is built, additional accounts can be set up, VMs can be imported, and controls can be developed and monitored.

Cisco UCS and Citrix CloudPlatform Network Design

This section explains Cisco UCS networking design considerations when deploying Citrix CloudPlatform 4.2.1 in a multi-tenant environment. (Multi-tenancy refers to multiple internal customers for Enterprise Private Clouds or to multiple external customers for Service Provider Public Clouds). In this design multi-tenant Management, guest and corresponding storage network traffic is isolated using the same Cisco UCS infrastructure by physically defining ethernet uplink ports and logical VLAN networks to provide data security in cloud environment. This design also reduces OpEx and CapEx compared to a topology in which a separate dedicated physical servers and switches are deployed to handle multi-tenant network traffic.

Figure 5 presents a detailed view of physical topology, identifying the various levels of the architecture and some of the main components of Cisco UCS in a cloud network design.

Figure 5 *Citrix CloudPlatform 4.2.1 on Cisco UCS Network Topology*



As shown in [Figure 5](#), a pair of Cisco UCS 6248UP Fabric Interconnects carry cloud multi-tenant network traffic from the Cisco UCS Blade Servers with the help Cisco Nexus 5548UP Switches. Both Cisco UCS 6248UP Fabric Interconnects and Cisco Nexus 5548UP Switches are clustered with the peer link between them to provide high availability. There are two virtual Port Channels (vPCs) that are configured to provide network access paths for the Cisco UCS Blade Servers to establish northbound connections with the Cisco Nexus Switches. Each vPC has VLANs created for cloud tenants management, guest, storage network data paths. Two other virtual Port Channels (vPCs) are configured to provide storage network paths from NetApp Controllers on northbound Nexus Switches.

The same cloud network design can be duplicated to provide multiple multi-tenants, network isolation on cloud on a single Cisco UCS, by provisioning separate vPCs with dedicated physical Ethernet uplink ports and logical VLANs to handle network data traffic.

To handle network scalability in cloud, and large multi-tenant network data traffic bandwidth at any given time, Cisco UCS vPCs allow to dynamically add more Ethernet uplink ports on the fly to accommodate cloud network bandwidth on demand.

Cabling Information for Network Devices


The following cabling information is provided as a reference for cabling the physical equipment in a Citrix CloudPlatform 4.2.1. The tables in this section include both local and remote device, and port locations in order to simplify cabling requirements.

This document assumes that the out-of-band management ports are plugged into an existing management infrastructure at the deployment site.

Be sure to follow the cabling instructions provided in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

[Figure 6](#) shows Citrix CloudPlatform 4.2.1 network cabling diagram. The labels indicate connections to the end points rather than the port numbers on the physical device. For example, connection A is a 10 Gb vPC port connected from Cisco Nexus 5548UP switch A to Cisco Fabric Interconnects A and connection P is a 10 Gb vPC peer links connected from Cisco Nexus 5548UP Switch A to Cisco Nexus 5548UP Switch B.

Figure 6 Citrix Cloud Network Cabling Diagram



The following table depicts the major cabling sections in the architecture:

1. Inter switch links
2. Virtual Port Channel connectivity
3. Infrastructure connectivity.

Table 1 shows the Ethernet cabling information layout on Cisco Nexus 5548UP Switch A, Cisco UCS Fabric Interconnects A and B, local device and NetApp Storage Controller A and B, local port channels and remote device port.

Table 1 *Ethernet Cabling on Cisco Nexus Switch A*

Cable ID on both end	Ethernet Interface	VLAN	Mode	Speed	Port Channel	Remote Device Port
A	Eth1/15	1,602-603,192-193,200,300	Trunk	10G	101	Fabric Interconnects A 1/15
C	Eth1/16	1,602-603,192-193,200,300	Trunk	10G	102	Fabric Interconnects B 1/15
G	Eth 1/11	1,602-603,1000-1001	Trunk	1G	None	Nexus 1100-A LOM A
H	Eth 1/12	1,602-603,1000-1001	Trunk	1G	None	Nexus 1100-B LOM A
P	Eth1/ 13	100	Trunk	10G	100	VPC peer link
E	Eth1/17	192	Access	10G	103	NetApp Controller A E1a
F	Eth 1/18	192	Access	10G	104	NetApp Controller B E1a
Not shown	Eth1/14	1,602-603,192-193,200,300	trunk	10G	None	Uplink to Infrastructure n/w

[Table 2](#) displays the Ethernet cabling information layout on Cisco Nexus 5548UP Switch B, UCS Fabric Interconnects A and B, local device and NetApp Storage Controller A and B, local port channels remote device port.

Table 2 *Ethernet Cabling on Cisco Nexus Switch B*

Cable ID on both end	Ethernet Interface	VLAN	Mode	Speed	Port Channel	Remote Device Port
B	Eth1/15	1,602-603,192-193,200,300	Trunk	10G	101	Fabric Interconnects A 1/16
D	Eth1/16	1,602-603,192-193,200,300	Trunk	10G	102	Fabric Interconnects B 1/16

Table 2 Ethernet Cabling on Cisco Nexus Switch B

I	Eth 1/11	1,602-603,1000-1001	Trunk	1G	None	Nexus 1100-A LOM B
J	Eth 1/12	1,602-603,1000-1001	Trunk	1G	None	Nexus 1100-B LOM B
P	Eth1/ 13	100	Trunk	10G	100	VPC peer link
G	Eth 1/17	192	Access	10G	103	NetApp Controller B E1b
H	Eth 1/18	192	Access	10G	104	NetApp Controller B E1b
Not shown	Eth1/14	1,602-603,192-193,200, 300	Trunk	10G	None	Uplink to Infrastructure n/w

Connect all the cables as outlined in the tables and the corresponding images.

Ethernet Cabling on Cisco Nexus Switch A and Ethernet Cabling on Cisco Nexus Switch B, shows the Cisco Nexus 5548UP vPC configurations with the vPC domains and corresponding vPC names and IDs in the Citrix cloud. To provide Layer 2 and 3 switching, a pair of Cisco Nexus 5548UP Switches with upstream switching are deployed, providing high availability in the event of failure to Cisco UCS to handle management, guest and storage data traffic. In the Cisco Nexus 5548UP topology, a single vPC feature is enabled to provide high availability, faster convergence in the event of a failure, and greater throughput.

Cisco UCS and Cloud Storage Design

This section explains Cisco UCS and cloud storage design considerations when deploying Citrix CloudPlatform 4.2.1 in a multi-tenant environment. In cloud each tenant's storage network traffic is completely isolated physically and logically by creating Virtual Logical SAN Network (VSAN). This allows creation of completely isolated fabric topologies, each with its own set of fabric services, on top of a scalable common physical infrastructure. Since each VSAN possesses its own zoning service, zoning is configured within each VSAN independently and has no effect on other VSANs and zoning services. This design also reduces OpEx and CapEx compared to the topology in which a separate dedicated physical switch is deployed to handle cloud multi-tenant storage traffic.

To support scalability, the Cisco UCS and Cisco Nexus fabric allows aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy by creating Port Channels.

In this deployment model, we will create VSANs for isolating multi-tenant storage data traffic and create FC Port Channels for link aggregation.

Cabling Information for Storage Devices

The following information is provided as a reference for cabling the physical equipment in a Citrix CloudPlatform 4.2.1 environment. The tables in this section provide both local and remote device and port locations in order to simplify cabling requirements.

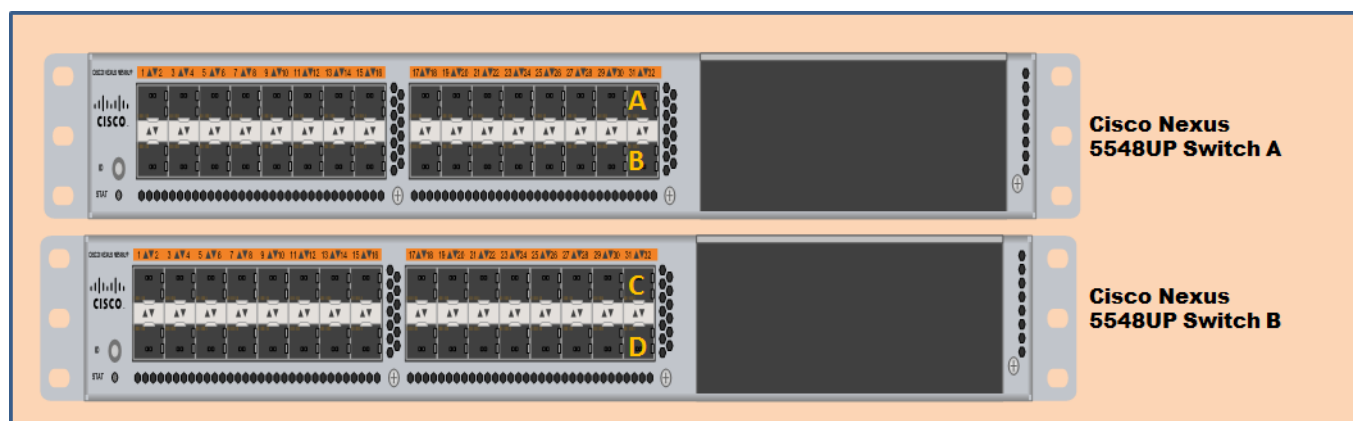
This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site.

The deployment procedures described use specific port locations. Be sure to follow the cabling directions provided in this section. Any change in the port location will result in a consequent change in the deployment procedure.

Before starting, be sure the configuration matches what is described in the tables and figures in this section.

Figure 7 shows Citrix CloudPlatform 4.2.1 network cabling diagram. The labels indicate connections to the end points rather than the port numbers on the physical device. For example, connection A is an FCoE 10-gb unified FCoE port connected from Cisco Nexus 5548UP switch A to Cisco Fabric Interconnects A.

Figure 7 Uplink Network Cabling for the Citrix CloudPlatform 4.2.1



Cisco UCS Quality-of-Service System and Policy

In a multi-tenant cloud environment service level availability plays a major role in defining various cloud services which are built on cloud infrastructure. One of the infrastructures is the cloud network for handling data traffic internally and externally. In this paper we will focus on how to define network service level agreements based on Ethernet network bandwidth offered which can later be provisioned to handle cloud multi-tenants data traffic.

Cisco UCS uses IEEE Data Center Bridging (DCB) to handle all traffic within Cisco UCS. This industry-standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. System classes determine how the DCB bandwidth in these virtual lanes is allocated across the entire Cisco UCS platform.

Each system class reserves a specific segment of bandwidth for a specific type of traffic, providing an assured level of traffic management even in an over-subscribed system. For example, the Fibre Channel priority system class can be configured to determine the percentage of DCB bandwidth allocated to FCoE traffic.

Table 3 **System Classes defined as per the Quality of Service**

System Class	Description
<ul style="list-style-type: none"> Platinum Priority Gold Priority Silver Priority Bronze Priority 	These classes set the Quality of Service (QoS) for all the servers. QoS system classes are defined in the service profile associated to the server. Each of these system classes manages one lane of traffic. All properties of these system classes are available to assign custom settings and policies.
Best-Effort Priority	This class sets the QoS for the lane that is reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy to allow it to drop data packets if required (can we write why the data packets will be dropped? In what circumstances?).
Fibre Channel Priority	This class sets the QoS for the lane that is reserved for FCoE traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy to help ensure that it never drops data packets.

To provide network based Service Level Agreement on Citrix CloudPlatform 4.2.1 environment for multi-tenant network requirements, the following system class SLAs have been defined. These SLAs are directly mapped to the cloud network offering which are provided to cloud tenants for choosing network infrastructure based on cost.

- Platinum-Network-SLA

To meet multi-tenant network requirements, to provide higher bandwidth, and lower latency, Cisco UCS QoS system class offers Platinum class with the highest weight (bandwidth) and a maximum transmission unit (MTU) of 9000 defined for handling large Cloud I/O network.

For example, to handle virtual machine migration across server pools in a cluster, define Platinum-Network-SLA policy for management traffic that is applied on Cisco UCS Static vNICs, part of management network bond interfaces.

- Gold-Network-SLA

To provide slightly lower bandwidth, and lower latency network for multi-tenant network requirements compared to Platinum-Network-SLA, Cisco UCS QoS system class offers Gold-Network-SLA with the second highest weight (bandwidth) and a maximum transmission unit (MTU) of 9000 for handling large Cloud I/O network traffic.

- Silver-Network-SLA

On the same cloud network SLA requirements to provide lower bandwidth and lower latency network for multi-tenant network requirements, Silver-Network-SLA is defined. The Cisco UCS QoS system class offers Silver-Network-SLA with the third highest weight (bandwidth) and a Maximum Transmission Unit (MTU) of 9000 for handling large Cloud I/O network traffic.


Note

In this paper we will define Cisco UCS QoS policies based on the described Cloud Network SLA Offerings which are related to specific customer needs and can be altered depending on your cloud network infrastructure offering.

To handle VM migration across server pools in a cluster that requires higher network bandwidth and lower latency, Platinum-Network-SLA policy for management traffic is defined and is later applied on Cisco UCS static vNICs that are part of management network bond interfaces. To handle NFS storage traffic, for handling operation such as cloud VM backup/ restore, templates, ISO or primary storage for storing VMs and so on, which require better bandwidth, Gold-Network-SLA policy for handling all storage related traffic is defined and is later applied on Cisco UCS static vNICs that are part of storage network bond interfaces. To handle VMs' public or private traffic, Silver-Network-SLA policy for such guest traffic is defined and is later applied on Cisco UCS static vNICs that are part of guest network bond interfaces.

After defining cloud network SLAs, Cisco UCS allows to define QoS policies that are assigned to the outgoing traffic for a vNIC. It is required to include a QoS policy with in a vNIC policy, which is then included in a service profile to configure the vNIC.

Table 4 lists the QoS policy name with the corresponding priority, weight, and MTU value. These values are applied to static and dynamic vNICs in the Microsoft SQL Server deployment environment.

Table 4 *Cisco UCS QoS Policies and Network SLAs*

Policy Name	Priority	Weight (Percentage)	MTU
Platinum-Net-SLA	Platinum	10	9000
Gold-Net-SLA	Gold	9	9000
Silver-Net-SLA	Silver	8	9000

Configuring Devices

This section describes the steps and procedure to configure the hardware devices in the unified design.

This section offers a detail of how to configure and deploy the UCS to make it ready to deploy the storage and CloudPlatform 4.2.1 on UCS:

- Cisco Nexus Switch Configuration
- Cisco UCS Deployment
- Uplink Port Channel Creation
- Storage Configurations

Cisco Nexus Switch Configurations

In this section vPC design is explained as per which the switches and other components will be configured and set up. The following section provides step-by-step procedure to configure the ports, features, VLANs, and Ethernet interfaces on the Nexus switches. The following tasks are detailed in this section:

- vPC Mapping in the Switch
- Configuring the Cisco Nexus Switch
- Enabling feature and Global configuration
- Configuring VLANs
- Configuring Port channel

- Configuring Virtual port channel
- Configuring the Port Channel
- Configuring the Ethernet Interfaces

vPC Mapping in the Switch

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 5000 Series devices to appear as a single port channel by a third device. The third device can be a switch, server, or any other networking device that supports port channels. A vPC can provide Layer 2 multipathing, which allows you to create redundancy and increase bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.

Table 5 shows the Cisco Nexus 5548UP vPC configurations with the vPC domains and corresponding vPC names and IDs in cloud. To provide Layer 2 and 3 switching, a pair of Cisco Nexus 5548UP Switches with upstream switching are deployed, providing high availability in the event of failure to Cisco UCS to handle management, guest and storage data traffic. In the Cisco Nexus 5548UP topology, a single vPC feature is enabled to provide high availability, faster convergence in the event of a failure, and greater throughput.

Table 5 **vPC Mapping**

vPC Domain	vPC Name	vPC ID
100	TenantvPCNetworkA	101
100	TenantvPCNetworkB	102
100	TenantUnifiedStorageA	103
100	TenantUnifiedStorageB	104
100	TenantFCoEFabric	10

In the vPC design table, a single vPC domain, Domain 100, is created across Cisco Nexus 5548UP member switches to define vPCs to carry specific network traffic. This topology defines four vPCs with IDs 101 through 104. vPC IDs 101 and 102 are defined for traffic from Cisco UCS Fabric Interconnects, and vPC IDs 103 and 104 are defined for iSCSI traffic, and vPC 105 and 106 for NFS traffic on NetApp storage. These vPCs are managed within the Cisco Nexus 5548UP, which connects Cisco UCS Fabric Interconnects and the NetApp storage system.

Configuring the Cisco Nexus Switch

The first time that you access Cisco Nexus 5000 Series, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the Ethernet interface. This information is required to configure and manage the switch.

Preparing to Configure the Switch

Before you configure Cisco Nexus 5000 Series switch for the first time, you need the following information:

- Administrator password



Note

If a password is weak (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password.

- If you are using an IPv4 address for the management interface, you need the following information:
 - IPv4 subnet mask for the switch's management interface
 - IPv4 address of the default gateway (optional)
- SSH service on the switch (optional)

To enable this service, select the type of SSH key (dsa/rsa/rsa1) and number of SSH key bits (768 to 2048).
- NTP server IPv4 address (optional)
- SNMP community string (optional)
- Switch name (optional)

This is your switch prompt
- An additional login account and password (optional)



Note

If you are using IPv4, be sure to configure the IPv4 route, the IPv4 default network address, and the IPv4 default gateway address to enable SNMP access.

Default Login

The switch has the network administrator as a default user (admin). You cannot change the default user at any time.

There is no default password so you must explicitly configure a strong password. If you configure and subsequently forget the password, you have the option to recover the password.



Note

If you enter the write erase command and reload the switch, you must reconfigure the default user (admin) password using the setup procedure.

To set the initial configurations on the Cisco Nexus 5548 switches, follow these steps:

For Cisco Nexus A and Cisco Nexus B

The NX-OS setup should automatically start on initial boot and connection to the serial or console port of the switch. Enter the following commands to configure the Cisco Nexus Switch:

1. Enter yes to enforce secure password standards: yes
2. Enter the password for the administrator (adminuser): <xxxxx>
3. Enter the password a second time to commit the password; <xxxxx>
4. Enter yes to enter the basic configuration dialog: yes
5. Create another login account (yes/no) [n]: Enter
6. Configure read-only SNMP community string (yes/no) [n]: Enter
7. Configure read-write SNMP community string (yes/no) [n]: Enter
8. Enter the switch name: <Nexus A Switch name> Enter

9. Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
10. Mgmt0 IPv4 address: <Nexus A mgmt0 IP> Enter
11. Mgmt0 IPv4 netmask: <Nexus A mgmt0 netmask> Enter
12. Configure the default gateway? (yes/no) [y]: Enter
13. IPv4 address of the default gateway: <Nexus A mgmt0 gateway> Enter
14. Enable the telnet service? (yes/no) [n]: Enter
15. Enable the ssh service? (yes/no) [y]: Enter
16. Type of ssh key you would like to generate (dsa/rsa):rsa
17. Number of key bits <768–2048>:1024 Enter
18. Configure the ntp server? (yes/no) [y]: n Enter
19. NTP server IPv4 address: <NTP Server IP> Enter
20. Enter basic FC configurations (yes/no) [n]: Enter
21. Would you like to edit the configuration? (yes/no) [n]: Enter

**Note**

Be sure to review the configuration summary before enabling it.

22. Use this configuration and save it? (yes/no) [y]: Enter

**Note**

Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus A or B.

**Note**

Log in as user admin with the password previously entered.

Enabling Features and Global Configuration

Enabling Features

To enable the vPC on the devices and define the communication between the linked devices, Link Aggregation Control Protocol (LACP), virtual port channel (vPC), and interface-vlan features must be activated on both the Cisco Nexus switches. To set these features on the two switches follow these steps:

For Cisco Nexus A and Cisco Nexus B:

1. Type configure t to enter the global configuration mode.
2. Type feature lacp.
3. Type feature interface-vlan.
4. Type feature vpc.

Setting Global Configurations

vPC eliminates the Spanning Tree Protocol (STP) blocked ports, and it is enabled as global configurations, which ensures that all the ports on the Cisco Nexus Switches appear as network port. To set the STP global configurations, follow these steps:

For Cisco Nexus A and Cisco Nexus B:

1. From the global configuration mode, type spanning-tree port type network default to make sure that, by default, the ports are considered as network ports with regard to spanning-tree.
2. Type spanning-tree port type edge bpduguard default to enable bpduguard on all edge ports by default.
3. Type spanning-tree port type edge bpdufilter default to enable bpdufilter on all edge ports by default.

Configuring VLANs

A VLAN is a group of end stations in a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but the end stations can be grouped even if they are not physically located on the same LAN segment.

This section details how the VLANS can be configured on both the Nexus switches.

In this design the VLANS have been defined as listed in [Table 6](#).

Table 6 *VLANs for Citrix CloudPlatform 4.2.1 setup*

VLAN Name	VLAN Purpose	ID used in this document	Network Address
Management	For Mgmt and VM Migration traffic	602	10.65.121.0/24
BareMetal-VLAN20	For Baremetal host PXE Boot traffic	20	20.1.1.0/24
Guest-VLAN	For guest VM public traffic	603	10.65.122.0/24
Guest-VLAN-1000	For guest VM Private traffic	1000	10.1.0.0/24
Guest-VLAN-1001	For guest VM Private traffic	1001	10.2.0.0/24
NFS	For NFS Storage traffic	193	193.191.1.0/24
vMotion	For vMotion traffic	192	192.191.1.0/24

For Cisco Nexus A and Cisco Nexus B

To configure the VLANS on both the Cisco Nexus Switches, follow these steps:

1. Type config-t.
2. Type vlan <Management VLAN ID>.
3. Type name Management.
4. Type exit.
5. Type vlan <BareMetal-VLAN20 VLAN ID>.
6. Type name BareMetal-VLAN.
7. Type exit.

8. Type `vlan <guest VLAN ID>`.
9. Type `name guest-VLAN`.
10. Type `exit`.
11. Type `vlan <guest VLAN ID>`.
12. Type `name guest-VLAN-1000`.
13. Type `exit`.
14. Type `vlan <guest VM VLAN ID>`.
15. Type `name guest-VLAN-1001`.
16. Type `exit`.
17. Type `vlan <NFS VLAN ID>`.
18. Type `name NFS`.
19. Type `exit`.
20. Type `vlan <vMotion VLAN ID>`.
21. Type `name vMotion`.
22. Type `exit`.
23. Type `vlan <FCoE VLAN ID>`.
24. Type `name VLAN0200`.
25. Type `fcoe vsan <ID 200>`.
26. Type `state <active>`.
27. Type `no <shutdown>`.
28. Type `exit`.
29. Type `vlan <FCoE VLAN ID>`.
30. Type `name VLAN0300`.
31. Type `fcoe vsan <ID 300>`.
32. Type `state <active>`.
33. Type `no <shutdown>`.
34. Type `exit`.

Configuring Port Channels

This section provides details to create and configure individual port descriptions for troubleshooting steps and verification. To create the port channels to the Cisco Nexus switches A and B, follow these steps:

Cisco Nexus 5548 A

1. From the global configuration mode, type `interface Eth1/13`.
2. Type `description <VPC Peer Link>`.
3. Type `exit`.
4. Type `interface Eth1/15`.

5. Type description <Cisco UCS Manager A:Eth1/5>.
6. Type exit.
7. Type interface Eth1/16.
8. Type description <Cisco UCS Manager B: Eth1/5>.
9. Type exit.
10. Type interface Eth1/17.
11. Type description <NetApp Controller Flexpod A:E1a>.
12. Type exit.
13. Type interface Eth1/18.
14. Type description <NetApp Controller Flexpod B:E1a>.
15. Type exit.
16. Type interface Eth1/31.
17. Type description <FCoE Uplink To Nexus 5548 B>.
18. Type exit.
19. Type interface Eth1/32.
20. Type description <FCoE Uplink To Nexus 5548 B>.
21. Type exit.

Cisco Nexus 5548 B

1. From the global configuration mode, type interface Eth1/13.
2. Type description <VPC Peer Link>.
3. Type exit.
4. Type interface Eth1/15.
5. Type description <Cisco UCS Manager A:Eth1/5>.
6. Type exit.
7. Type interface Eth1/16.
8. Type description <Cisco UCS Manager B:Eth1/5>.
9. Type exit.
10. Type interface Eth1/17.
11. Type description <NetApp Controller Flexpod A:E1b>.
12. Type exit.
13. Type interface Eth1/18.
14. Type description <NetApp Controller Flexpod B:E1b>.
15. Type exit.
16. Type interface Eth1/31.
17. Type description <FCoE Uplink To Nexus 5548 B>.
18. Type exit.
19. Type interface Eth1/32.

20. Type description <FCoE Uplink To Nexus 5548 B>.
21. Type exit.

Configuring Virtual Port Channels

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 5000 Series devices to appear as a single port channel by a third device. To configure the virtual Port Channels (vPCs) on Cisco Nexus switches A and B, follow these steps:

Cisco Nexus 5548 A

1. From the global configuration mode, type vpc domain <Nexus vPC domain ID>.
2. Type role priority 10.
3. Type peer-keepalive destination <Nexus B mgmt0 IP> source <Nexus A mgmt0IP>.
4. Type exit.
5. Type interface Port-Channel 100.
6. Type vpc peer-link.
7. Type exit.
8. Type interface Port-Channel 101.
9. Type vpc 101.
10. Type exit.
11. Type interface Port-Channel 102.
12. Type vpc 102.
13. Type exit.
14. Type interface Port-Channel 103.
15. Type vpc 103.
16. Type exit.
17. Type interface Port-Channel 104.
18. Type vpc 104.
19. Type exit.
20. Type interface Port-Channel 10.
21. Type exit.
22. Type copy run start.

Cisco Nexus 5548 B

1. From the global configuration mode, type vpc domain <Nexus vPC domain ID>.
2. Type role priority 20.
3. Type peer-keepalive destination <Nexus A mgmt0 IP> source <Nexus B mgmt0IP>.
4. Type exit.
5. Type interface Port-Channel 100.

6. Type vpc peer-link.
7. Type exit.
8. Type interface Port-Channel 101.
9. Type vpc 101.
10. Type exit.
11. Type interface Port-Channel 102.
12. Type vpc 102.
13. Type exit.
14. Type interface Port-Channel 103.
15. Type vpc 103.
16. Type exit.
17. Type interface Port-Channel 104.
18. Type vpc 104.
19. Type exit.
20. Type interface Port-Channel 10.
21. Type exit.
22. Type copy run start.

Adding Port Channel Configurations

After the port channels have been created, it is required to configure them individually to meet the design requirements. The ports are configured as per the vPC Mapping defined earlier in [Table 6](#). To add the port channel configurations on the Cisco Nexus Switches A and B, follow these steps:

Cisco Nexus 5548 A

1. From the global configuration mode, type interface Port-Channel 100.
2. Type switchport mode trunk.
3. Type switchport trunk native vlan <Native VLAN ID>.
4. Type switchport trunk allowed vlan <MGMT VLAN ID, BareMetal-VLAN VLAN ID GUESTVM VLAN ID, guest-VLAN-1000 VLANID, guest-VLAN-1001 VLANID, NFS VLAN ID, vMotion VLAN ID>.
5. Type spanning-tree port type network.
6. Type no shutdown.
7. Type exit.
8. Type interface Port-Channel 101.
9. Type switchport mode trunk.
10. Type switchport trunk native vlan <Default VLAN ID>.
11. Type switchport trunk allowed vlan <MGMT VLAN ID, BareMetal-VLAN VLANID, guest-VLAN VLAN ID, guest-VLAN-1000 VLAN ID, guest-VLAN-1001 VLAN ID, NFS VLAN ID, vMotion VLAN ID>.
12. Type spanning-tree port type edge trunk.

13. Type no shut.
14. Type exit.
15. Type interface Port-Channel 102.
16. Type switchport mode trunk.
17. Type switchport trunk native vlan <Native VLAN ID>.
18. Type switchport trunk allowed vlan <NFS VLAN ID>.
19. Type spanning-tree port type edge trunk.
20. Type no shut.
21. Type exit.
22. Type interface Port-Channel 103.
23. Type switchport mode trunk.
24. Type switchport trunk native vlan <Native VLAN ID>.
25. Type switchport trunk allowed vlan <NFS VLAN ID>.
26. Type no shut.
27. Type exit.
28. Type interface Port-Channel 104.
29. Type switchport mode trunk.
30. Type switchport trunk native vlan <Native VLAN ID>.
31. Type no shut.
32. Type exit.
33. Type interface Port-Channel 10.
34. Type switchport mode trunk.
35. Type switchport trunk native vlan <Native VLAN ID>.
36. Type switchport trunk allowed vlan <200>.
37. Type no shut.
38. Type exit.
39. Type copy run start.

Cisco Nexus 5548 B

1. From the global configuration mode, type interface Port-Channel 100.
2. Type switchport mode trunk.
3. Type switchport trunk native vlan <Native VLAN ID>.
4. Type switchport trunk allowed vlan <MGMT VLAN ID, BareMetalVlan VLANID, guestVM VLAN ID, guest-VLAN-1000 VLANID, guest-VLAN-1001 VLANID,NFS VLAN ID, vMotion VLAN ID>.
5. Type spanning-tree port type network.
6. Type no shutdown.
7. Type exit.

8. Type interface Port-Channel 101.
9. Type switchport mode trunk.
10. Type switchport trunk native vlan <Default VLAN ID>
11. Type switchport trunk allowed vlan <MGMT VLAN ID, BareMetalVlan VLANID, guest-VLAN VLAN ID, guest-VLAN-1000 VLANID, guest-VLAN-1001 VLANID, NFS VLAN ID, vMotion VLAN ID>
12. Type spanning-tree port type edge trunk.
13. Type no shut.
14. Type exit.
15. Type interface Port-Channel 102.
16. Type switchport mode trunk.
17. Type switchport trunk native vlan <Native VLAN ID>
18. Type switchport trunk allowed vlan <MGMT VLAN ID, BareMetalVlan VLANID, guest-Vlan VLAN ID, guest-VLAN-1000 VLANID, guest-VLAN-1001 VLANID, NFS VLAN ID, vMotion VLAN ID>.
19. Type spanning-tree port type edge trunk.
20. Type no shut.
21. Type exit.
22. Type interface Port-Channel 103.
23. Type switchport mode trunk.
24. Type switchport trunk native vlan <Native VLAN ID>.
25. Type no shut.
26. Type exit.
27. Type interface Port-Channel 104.
28. Type switchport mode trunk.
29. Type switchport trunk native vlan <Native VLAN ID>.
30. Type no shut.
31. Type exit.
32. Type interface Port-Channel 10.
33. Type switchport mode trunk.
34. Type switchport trunk native vlan <Native VLAN ID>.
35. Type switchport trunk allowed vlan <300>.
36. Type no shut.
37. Type exit.
38. Type copy run start.

Configuring Ethernet Interfaces

Every port channel created must have their interfaces configured to join with the relevant VLANs. To configure the Ethernet interfaces on the port channels between the devices, follow these steps:

Cisco Nexus 5548 A

1. From the global configuration mode, type interface Eth1/13.
2. Type channel-group 100 mode active.
3. Type no shutdown.
4. Type exit.
5. Type interface Eth1/15.
6. Type channel-group 101 mode active.
7. Type no shutdown.
8. Type exit.
9. Type interface Eth1/16.
10. Type channel-group 102 mode active.
11. Type no shutdown.
12. Type exit.
13. Type interface Eth1/17.
14. Type channel-group 103 mode active.
15. Type no shutdown.
16. Type exit.
17. Type interface Eth1/18.
18. Type channel-group 104 mode active.
19. Type no shutdown.
20. Type exit.
21. Type interface Eth1/17.
22. Type channel-group 103 mode active.
23. Type no shutdown.
24. Type exit.
25. Type interface vfc 17.
26. Type bind interface Ethernet 1/17.
27. Type switch port trunk allowed vsan 200.
28. Type exit.
29. Type interface vfc 18.
30. Type bind interface Ethernet 1/18.
31. Type switch port trunk allowed vsan 200.
32. Type exit.
33. Type copy run start.

Cisco Nexus 5548 B

1. From the global configuration mode, type interface Eth1/13.

2. Type channel-group 100 mode active.
3. Type no shutdown.
4. Type exit.
5. Type interface Eth1/15.
6. Type channel-group 101 mode active.
7. Type no shutdown.
8. Type exit.
9. Type interface Eth1/16.
10. Type channel-group 102 mode active.
11. Type no shutdown.
12. Type exit.
13. Type interface Eth1/17.
14. Type channel-group 103 mode active.
15. Type no shutdown.
16. Type exit.
17. Type interface Eth1/18.
18. Type channel-group 104 mode active.
19. Type no shutdown.
20. Type exit.
21. Type interface vfc 17.
22. Type bind interface Ethernet 1/17.
23. Type switch port trunk allowed vsan 200.
24. Type exit.
25. Type interface vfc 18.
26. Type bind interface Ethernet 1/18.
27. Type switch port trunk allowed vsan 200.
28. Type exit.
29. Type copy run start.

Configuring Nexus 1100 Ethernet Ports

This section provides details to create and configure individual port descriptions for troubleshooting steps and verification. To create the ethernet ports to the Cisco Nexus switches A and B, follow these steps:

Cisco Nexus 5548 A

1. From the global configuration mode, type interface Eth1/11.
2. Type description <Nexus 1100 – A LOM A>.
3. Type switchport mode trunk.

4. Type switchport trunk allowed vlan 602,603,1000,1001,50.
5. Type exit.
6. Type interface Eth1/12.
7. Type description <Nexus 1100 – B LOM A>.
8. Type switchport mode trunk.
9. Type switchport trunk allowed vlan 602,603,1000,1001,50.
10. Type exit.
11. Type copy run start.

Cisco Nexus 5548 B

1. From the global configuration mode, type interface Eth1/11.
2. Type description <Nexus 1100 – A LOM B>.
3. Type switchport mode trunk.
4. Type switchport trunk allowed vlan 602,603,1000,1001,50.
5. Type exit.
6. Type interface Eth1/12.
7. Type description <Nexus 1100 – B LOM B>.
8. Type switchport mode trunk.
9. Type switchport trunk allowed vlan 602,603,1000,1001,50.
10. Type exit.
11. Type copy run start.

FlexPod Cisco Nexus 1000V vSphere

The following sections provide detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1000V software DVS Switch in a FlexPod configuration. Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the separate ESX hosts. By the end of this section, a Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned.

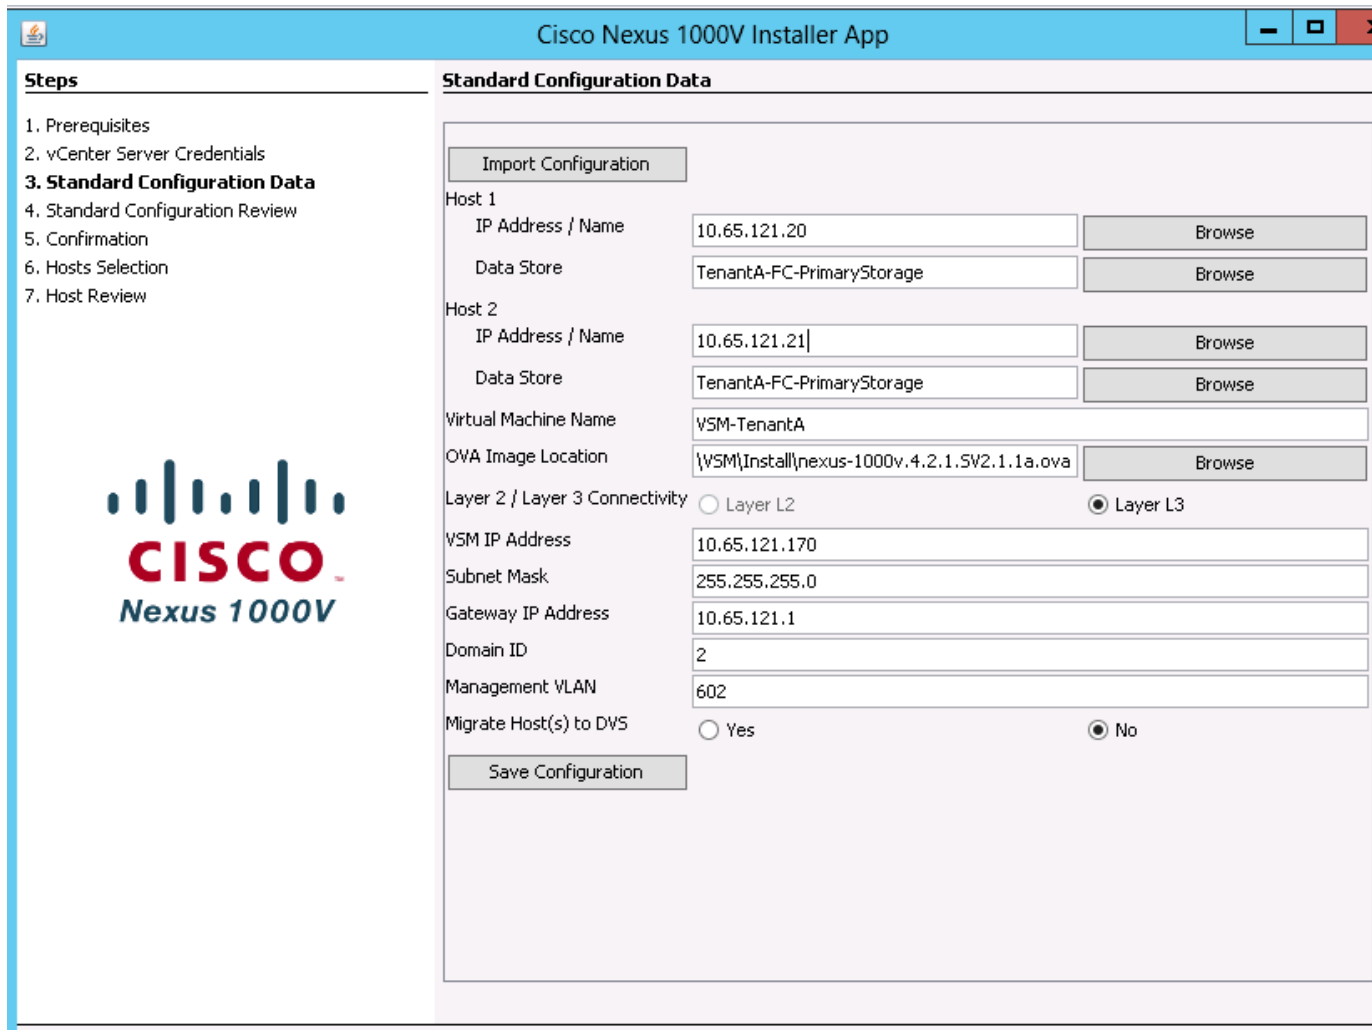
This procedure assumes that the Cisco Nexus 1000V software version 4.2.1(1)SV2(1.1a) has been downloaded from www.cisco.com and expanded. It also assumes that two ESX host is been installed with VMware vSphere 5.1 with Enterprise Plus licensing and VMware vCenter 5.1 application is installed which is part of Infrastructure Pod.

Installation of Primary and Secondary Cisco Nexus 1000V VSMs

1. Launch Nexus1000V-install_CNx.jar file from command prompt on Infrastructure Windows client
2. Click the **Cisco Nexus 1000V Complete Installation** radio button.
3. Select Standard Type.
4. Click **Next**
5. Enter vCenter IP Address <xx.xx.xx.xx>, Enter User ID <administrator> and Password <XXXXXX>.

6. Click **Next**
7. Enter or Browse ESX Host 1 IP Address <xx.xx.xx.xx> and shared Data Store <xxxxxxxx>.
8. Enter or Browse ESX Host 2 IP Address <xx.xx.xx.xx> and shared Data Store <xxxxxxxx>.
9. Enter Virtual Machine Name <VSM-TenantA>.
10. Browse the path to locate Nexus 1000V software <nexus-1000v.4.2.1.1.SV2.1.1a.ova> file.
11. By default **Layer 3** radio button will be selected for Layer 2 / Layer 3 Connectivity.

Figure 8 Cisco Nexus 1000V Standard Configuration Data details



Cisco Nexus 1000V Installer App

Steps

1. Prerequisites
2. vCenter Server Credentials
- 3. Standard Configuration Data**
4. Standard Configuration Review
5. Confirmation
6. Hosts Selection
7. Host Review

Standard Configuration Data

Import Configuration

Host 1

IP Address / Name: 10.65.121.20 Browse

Data Store: TenantA-FC-PrimaryStorage Browse

Host 2

IP Address / Name: 10.65.121.21 Browse

Data Store: TenantA-FC-PrimaryStorage Browse

Virtual Machine Name: VSM-TenantA

OVA Image Location: \\VSM\Install\nexus-1000v.4.2.1.1.SV2.1.1a.ova Browse

Layer 2 / Layer 3 Connectivity: ☐ Layer L2 ☒ Layer L3

VSM IP Address: 10.65.121.170

Subnet Mask: 255.255.255.0

Gateway IP Address: 10.65.121.1

Domain ID: 2

Management VLAN: 602

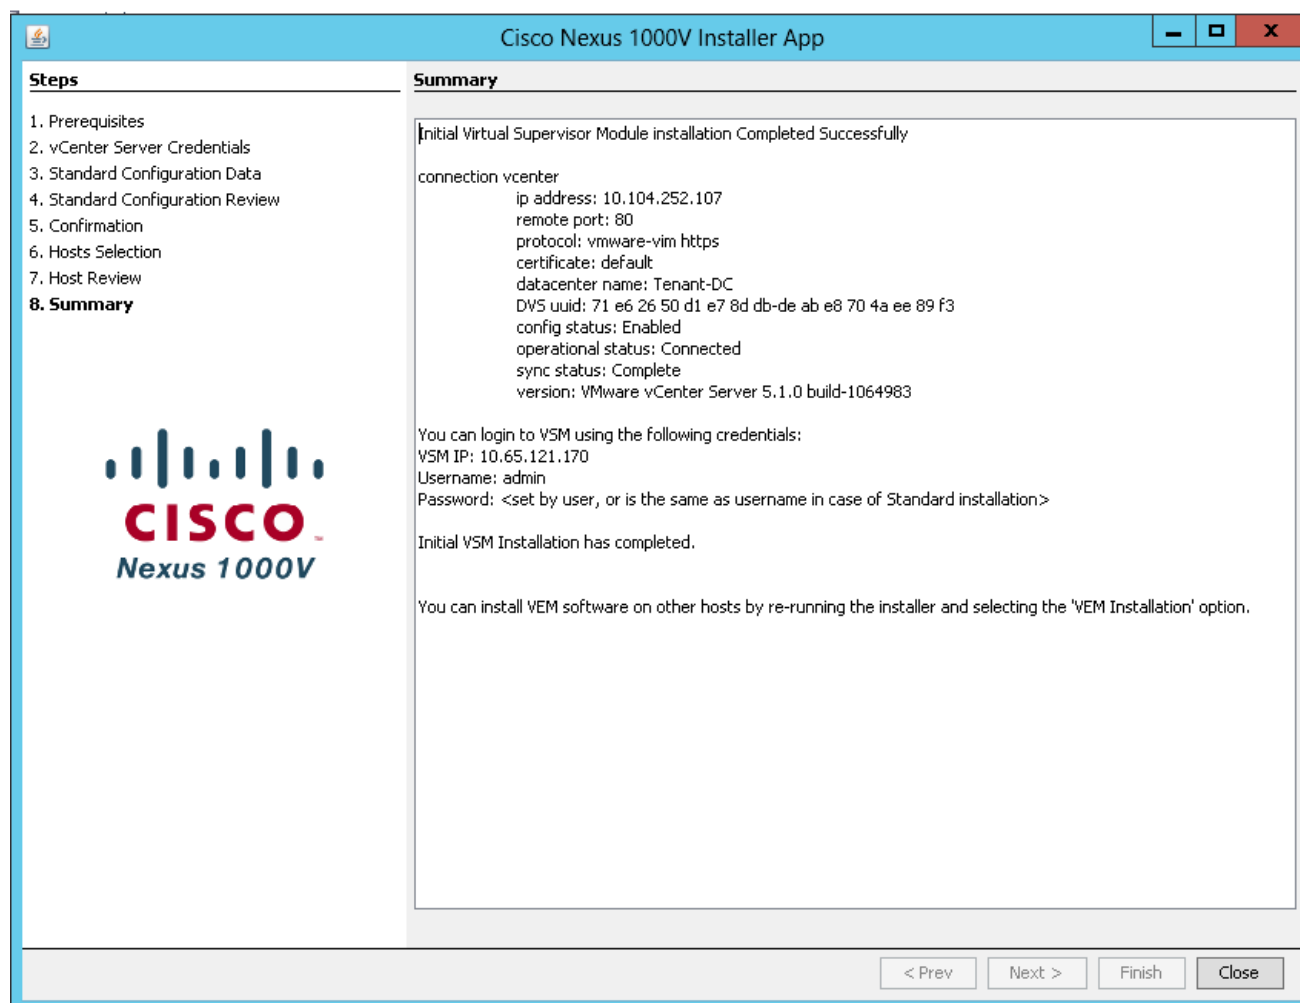
Migrate Host(s) to DVS: ☐ Yes ☒ No

Save Configuration

CISCO
Nexus 1000V

12. Click **Next**
13. Check Standard Configuration Review data
14. Click **Next**
15. Click **No** radio button Do you want to add more modules
16. Click **Next**
17. Click **Next** in Host Selection
18. Click **Finish** in Host review

Figure 9 Cisco Nexus 1000V Standard Installation Summary details



Configure Cisco Nexus 1000V Virtual Appliance

Cisco Nexus 1000V

To configure Cisco Nexus 1000-V VSM-TenantA-1, follow these steps:

1. Login to VSM-TenantA-1 VM on ESX Host 10.65.121.20 using VMware VI Client

```
Enter the User Name <<admin>>
Enter Password <<admin>>
Enter <<Conf Terminal>>
Enter <<feature ssh>> <<feature telnet>> <<feature http-server>>
Enter <<Copy run start>>
```

2. Perform Base Configuration of the Primary VSM

To perform the base configuration of the primary VSM, follow these steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t
```

```

ntp server <<var_global_ntp_server_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id 602>>
name IB-MGMT-VLAN
vlan <<var_guest_vlan_id 603>>
name guestVM-VLAN
vlan <<var_guest_vlan_id 1000>>
name guestVM-VLAN1000
vlan <<var_guest_vlan_id 1001>>
name guestVM-VLAN1001
exit
port-profile type ethernet TenantA-guest-Uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<default 1>>,<<var_ib-mgmt_vlan_id 602>>, <<
var_guest_vlan_id 603>>,<<var_guest_vlan_id 1000>>, <<var_N1KV_Ctrl_vlan_id 50>>, <<
var_guest_vlan_id 1001>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id 602>>, << var_guest_vlan_id 603>>,
<<var_guest_vlan_id 1000>>, << var_guest_vlan_id 1001>>
system mtu 9000
state enabled
port-profile type vethernet TenantA-VSM-VEM-Network
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id 602>>
no shutdown
capability l3control
system vlan <<var_ib-mgmt_vlan_id 602>>
state enable
exit
copy run start

```

Cisco UCS Deployment

The following section provides a detailed procedure for configuring the Cisco Unified Computing System for use in the FlexPod environment. These steps should be followed precisely because a failure to do so could result in an improper configuration.

This section describes how to perform the following tasks that helps in setting up Cisco UCS:

- Configuring Cisco UCS 6248 Fabric Interconnects
- Modifying the Chassis Discovery Policy
- Creating Uplink Port-Channel
- Enabling Jumbo frames
- Enabling QoS Policy in Cisco UCS 6248 Fabric Interconnects
- Enabling QoS Qualifications on Cisco Nexus 5548UP Switch
- Configuring VLAN on Cisco UCS 6248 Fabric Interconnects

Configuring Cisco UCS 6248 Fabric Interconnects

To perform the initial setup of the Cisco UCS 6248 Fabric Interconnects A and B, follow these steps:

Fabric Interconnects A

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect.
2. At the prompt to enter the configuration method, enter console to continue.
3. If asked to either do a new setup or restore from backup, enter setup to continue.
4. Enter y to continue to set up a new fabric interconnect.
5. Enter y to enforce strong passwords.
6. Enter the password for the admin user.
7. Enter the same password again to confirm the password for the admin user.
8. When asked if this Fabric Interconnects is part of a cluster, answer y to continue.
9. Enter A for the switch fabric.
10. Enter the cluster name for the system name.
11. Enter the Mgmt0 IPv4 address.
12. Enter the Mgmt0 IPv4 netmask.
13. Enter the IPv4 address of the default gateway.
14. Enter the cluster IPv4 address.
15. To configure DNS, answer y.
16. Enter the DNS IPv4 address.
17. Answer y to set up the default domain name.
18. Enter the default domain name.
19. Review the settings that were printed to the console, and if they are correct, answer yes to save the configuration.
20. Wait for the login prompt to make sure the configuration has been saved.

Fabric Interconnects B

1. Connect to the console port on the second Cisco UCS 6248 fabric interconnect.
2. When prompted to enter the configuration method, enter console to continue.
3. The installer detects the presence of the partner Fabric Interconnects and adds this Fabric Interconnects to the cluster. Enter y to continue the installation.
4. Enter the admin password for the first Fabric Interconnects.
5. Enter the Mgmt0 IPv4 address.
6. Answer yes to save the configuration.
7. Wait for the login prompt to confirm that the configuration has been saved.

Cisco UCS Manager-Login

After the Cisco UCS Fabric Interconnects has been configured, the next step would be to download and start using the Cisco UCS Manager software. To login to the Cisco UCS Manager follow these steps:

1. Open a Web browser and navigate to the Cisco UCS 6248 Fabric Interconnects cluster address.
2. Select the Launch link to access Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.

4. When prompted, enter admin for the username and enter the administrative password and click **Login** to log into the Cisco UCS Manager.
5. Click the **Equipment** tab in the left pane.
6. Select Fabric Interconnects A.
7. In the right pane, click the **General** tab.
8. Select Configure Unified Ports.
9. Select Yes to launch the wizard.
10. Use the slider tool and move one position to the left to configure the last two ports (31 and 32) as FC uplink ports.
11. Ports 31 and 32 now have the “B” indicator indicating their reconfiguration as FC uplink ports.
12. Click **Finish**.
13. Click **OK**.
14. The Cisco UCS Manager GUI will close as the primary Fabric Interconnects reboot.
15. Upon successful reboot, open a Web browser and navigate to the Cisco UCS 6248 Fabric Interconnects cluster address.
16. When prompted, enter admin as username and enter the administrative password.

Modifying Chassis Discovery Policy

This section provides details for modifying the chassis discovery policy, as the base architecture includes two uplinks from each Fabric Extender installed in the Cisco UCS chassis.

If the workload on the cloud is expected to be network centric, you may want to provide more bandwidth for the B200 M3 Blade Servers. Increasing the number of links between Fabric Extenders (IOM) on UCS Chassis and Fabric Interconnects would reduce bandwidth oversubscription. Change the number of links in the “Chassis Discovery Policy” under the equipment to match physical number of links. Ideally, the chassis discovery policy should be modified before configuring server ports. If you add or remove links between Fabric Interconnects and Fabric Extender after a chassis is discovered, you need to “Acknowledge chassis”.



Note

Acknowledging chassis will not reboot the servers in a chassis.

To modify the chassis discovery policy, login to the Cisco UCS Manager and follow these steps:

1. Click the **Equipment** tab in the left pane.
2. In the right pane, click the **Policies** tab.
3. Under Global Policies, change the Chassis Discovery Policy to 2-link.
4. Click **Save Changes**.

Configuring Port Channels

To configure the port channels follow these steps:

1. Click the **Equipment** tab in the left pane.
2. Select Fabric Interconnects A.
3. In the right pane, under **Actions** click the **LAN UpLinks Manager** tab.

4. Click **Create Port Channel** then Select Fabric A.
5. Specific ID <101> and Name <Tenant1CloudVPCNetworkA>.
6. On Add Ports, click **Ports** number <5> and <6>.
7. Click **Finish**.
8. By Default Port Channel < Tenant1CloudVPCNetworkB> is disabled, click **Enable**.
9. Click **Yes** as the warning message display asking to <Enable / Disable> on Fabric.
10. Select Fabric Interconnects B.
11. In the right pane, under **Actions** click the **LAN UpLinks Manager** tab.
12. Right-click **Create Port Channel** then select Fabric A.
13. Specific ID <102> and Name <Tenant1CloudVPCNetworkB>.
14. On Add Ports click **Ports** number <5> and <6>.
15. Click **Finish**.
16. By Default Port Channel < Tenant1CloudVPCNetworkB> is disabled, click **Enable**.
17. Click **Yes** as the warning message displays <Enable / Disable> on Fabric B.

**Note**

When configuring the Cisco Nexus 5548UP with vPCs, ensure that the status of all vPCs is “Up” for connected Ethernet ports.

Port Channel Status on Cisco Nexus 5548UP command output:

```
N5548-Switch-A(config-if)# sh vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 100
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
Type-2 consistency reason : QoSMgr Qos configuration incompatible
vPC role               : primary
Number of vPCs configured : 2
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
vPC Peer-link status
```

```
-----
id   Port   Status Active vlans
--   --
```

```
1    Po100  up      1,192-195,602-603,607,20
```

vPC status

```
-----
id   Port   Status Consistency Reason Active vlans
-----
101   Po101   up      success success 1,602-603,1000-1001,192-193,20
102   Po102   up      success success 1,602-603,1000-1001,192-193,20
103   Po103   up      success success 1,192-193,200-300
104   Po104   up      success success 1,192-193,200-300
```

```
N5548-Switch-B(config-if)# sh vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 100
Peer status            : peer adjacency formed ok
```

```

vPC keep-alive status          : peer is alive
Configuration consistency status: success
Per-vlan consistency status    : success
Type-2 consistency status      : success
Type-2 consistency reason      : QoSMgr Qos configuration incompatible
vPC role                       : primary
Number of vPCs configured      : 2
Peer Gateway                   : Disabled
Dual-active excluded VLANs     : -
Graceful Consistency Check     : Enabled
vPC Peer-link status
-----
id   Port   Status Active vlans
--   --
1    Po100  up     1,192-195,602-603,607,20
vPC status
-----
id   Port   Status Consistency Reason          Active vlans
-----
101  Po101   up     success    success    1,1000-10001,602-603,192-193,20
102  Po102   up     success    success    1,1000-10001,602-603,192-193,20
103  Po103   up     success    success    1,192-193,200,300
104  Po104   up     success    success    1,192-193,200,300

```

Creating the Uplink Ethernet Ports

LAN traffic enters and departs the Cisco UCS system on the Fabric Interconnects via the use of uplink ports. There are different types of uplink ports and here the uplink port channels are configured to connect to the NetApp storage.

To create and configure the necessary uplink port channels in the Cisco UCS environment login to the Cisco UCS Manager and follow these steps:

Fabric Interconnects A

1. Click the **LAN** tab in the left pane.

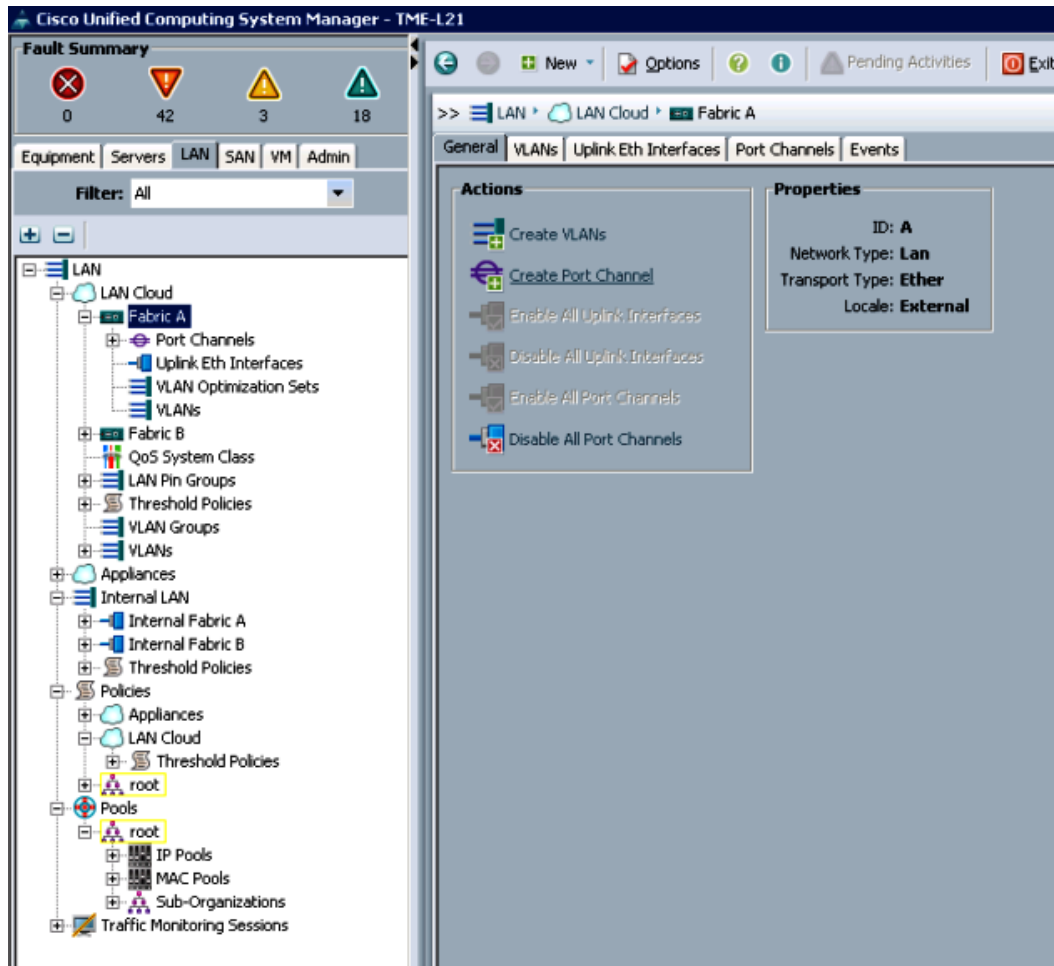


Note

Two Port Channels are created, one from Fabric A to both the Cisco Nexus 5548 switches, and one from Fabric B to both the Cisco Nexus 5548 switches.

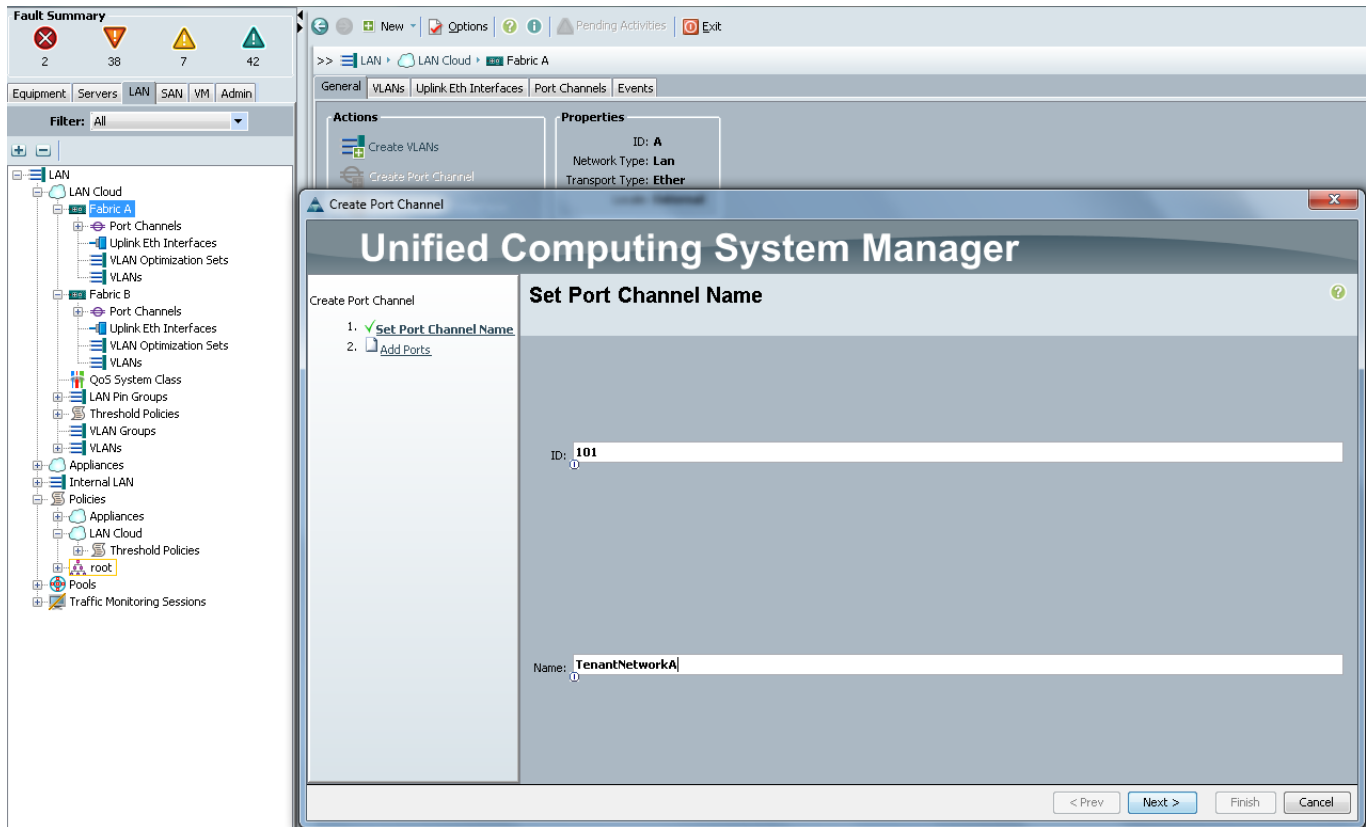
2. Under LAN Cloud, expand the Fabric A tree.
3. In the right pane under **General** tab click **Create Port Channel**.

Figure 10 **Creating Uplink Port Channels**



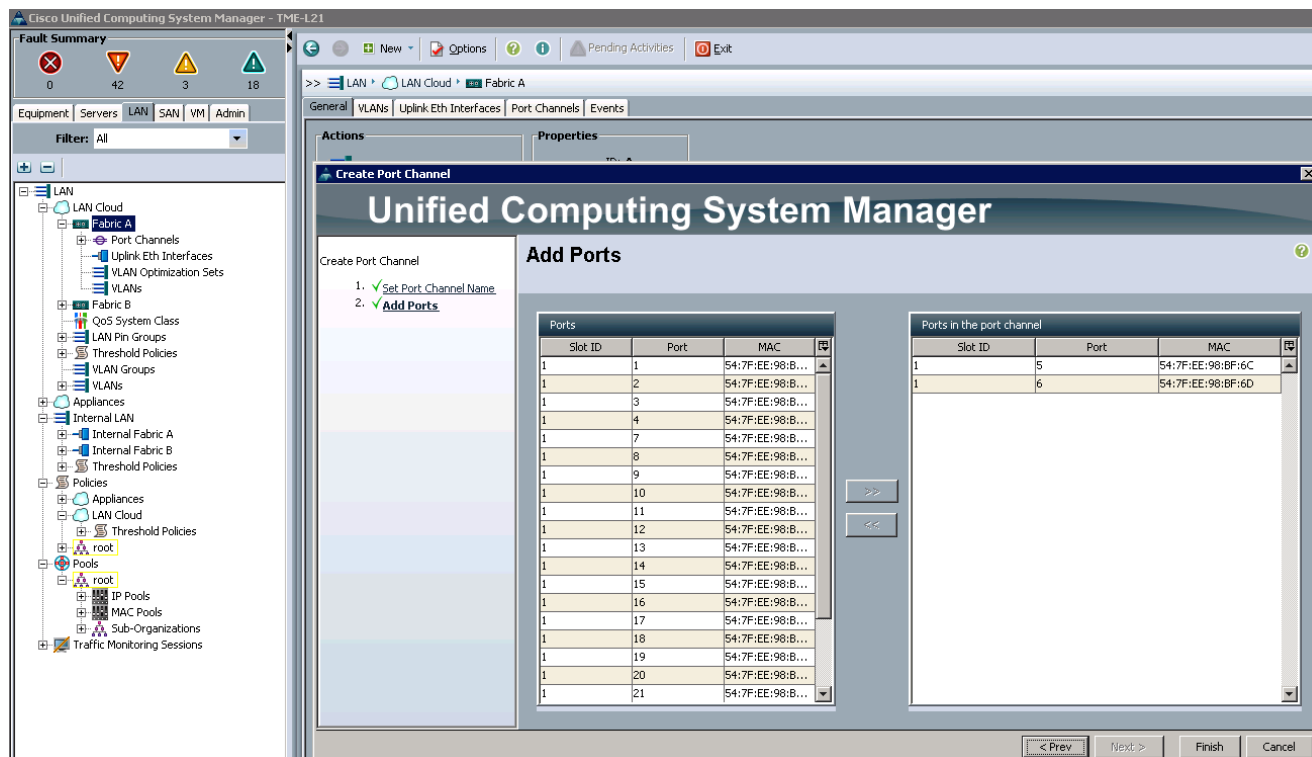
4. Enter 101 as the unique ID of the Port Channel.
5. Enter Tenant Network A as the name of the Port Channel.
6. Click **Next**.

Figure 11 **Naming the Port Channel**



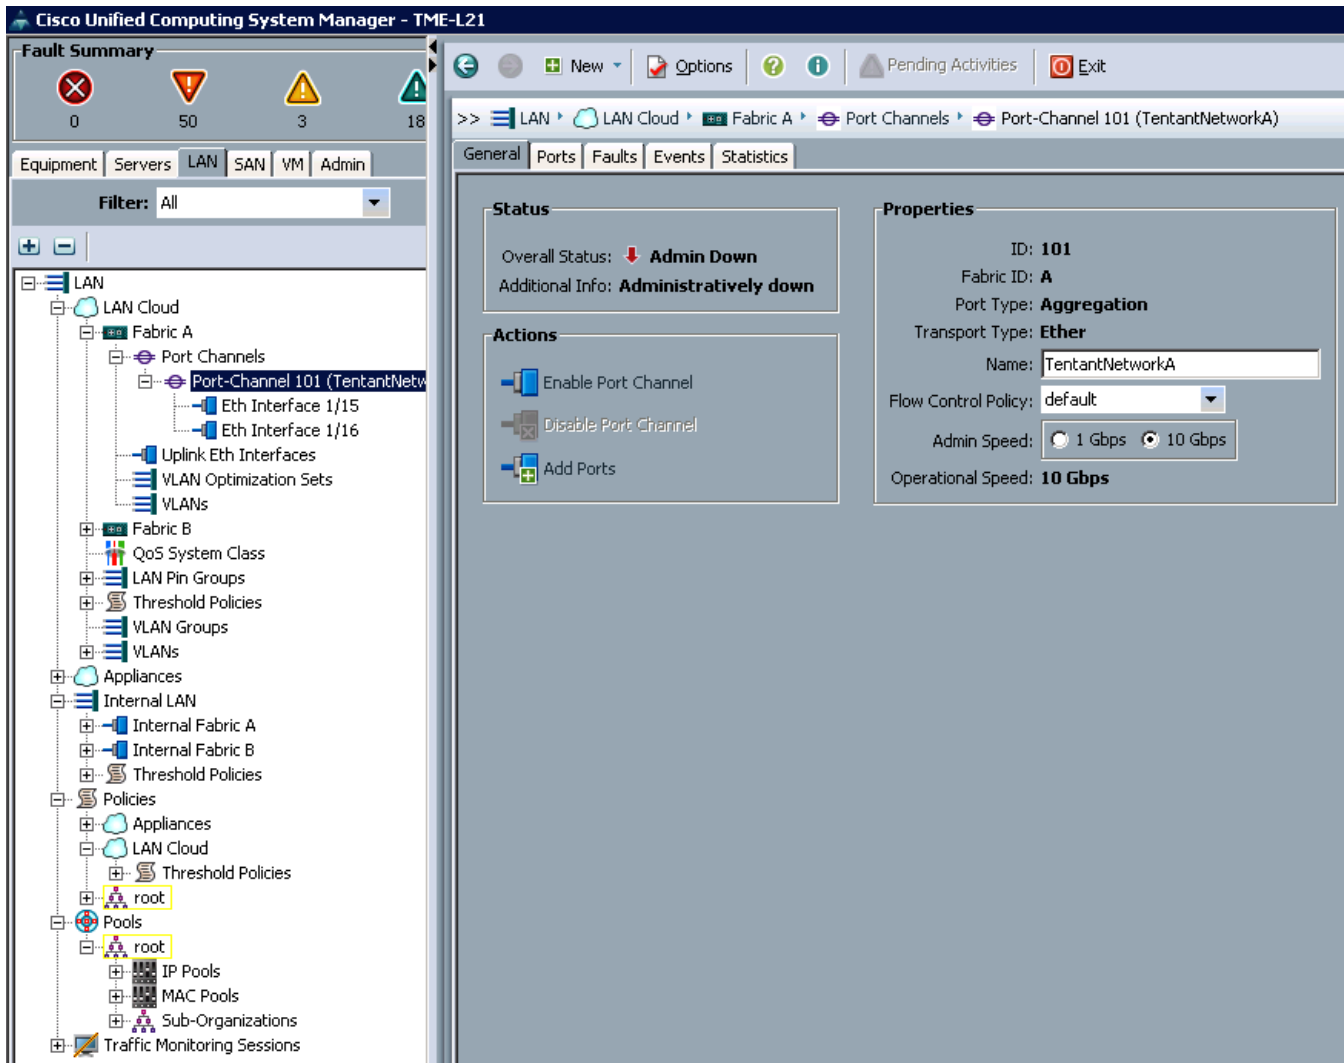
7. Select the port with slot ID: 1 and port: 5, and the port with slot ID:1 and port 6, from the Ports list.
8. Click >> to add the ports to the Port Channel.

Figure 12 **Selecting and Adding the Ports**



9. Click **Finish** to create the Port Channel.
10. Check the **Show navigator for Port-Channel 101 (Fabric A)** checkbox.
11. Click **OK**.
12. In **Actions** area, select **Enable Port Channel**.
13. In the pop-up box, click **Yes**, then click **OK** to enable port-channel.

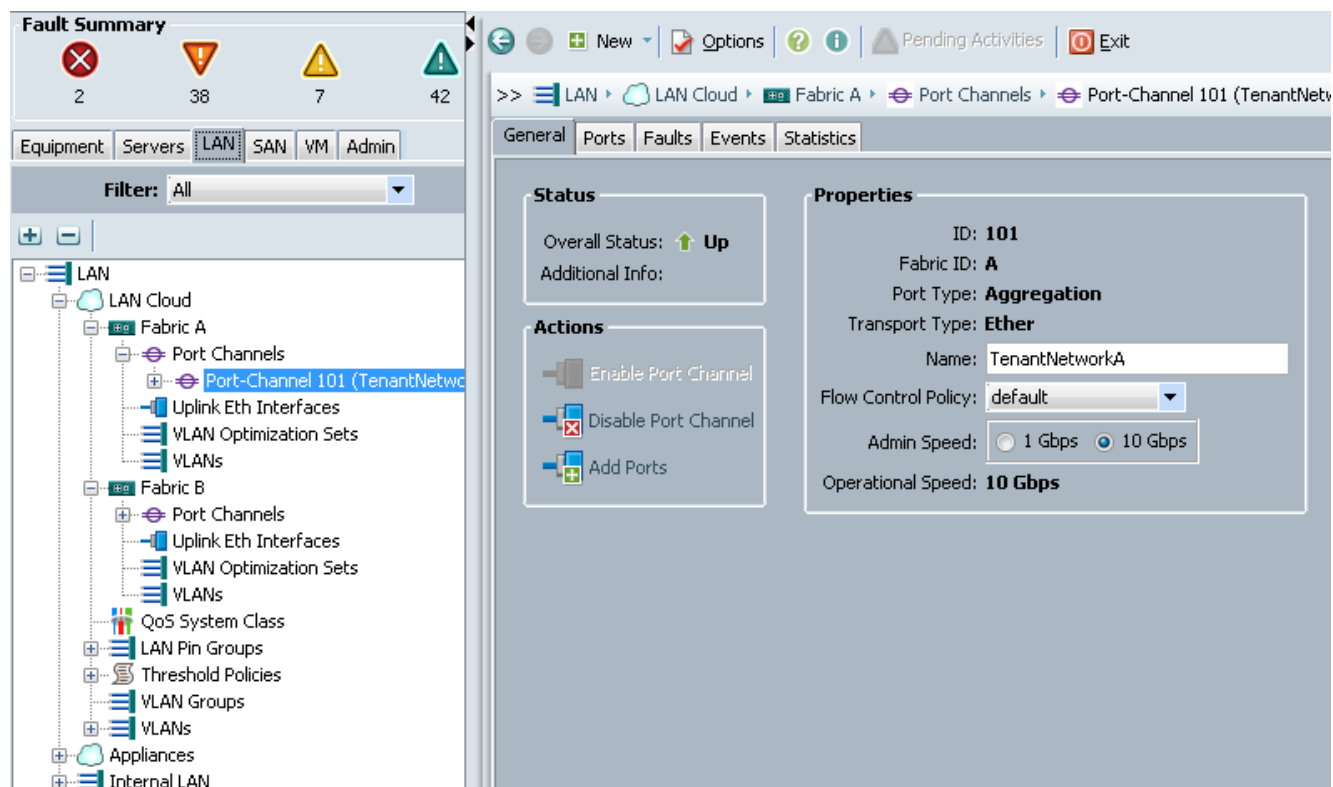
Figure 13 *Enabling the Port Channel*



14. Wait until the overall status of the Port Channel is Up.

15. Click **OK** to close the navigator.

Figure 14 Verifying the Port Channel Status



Fabric Interconnects B

1. Click the **LAN** tab in the left pane.

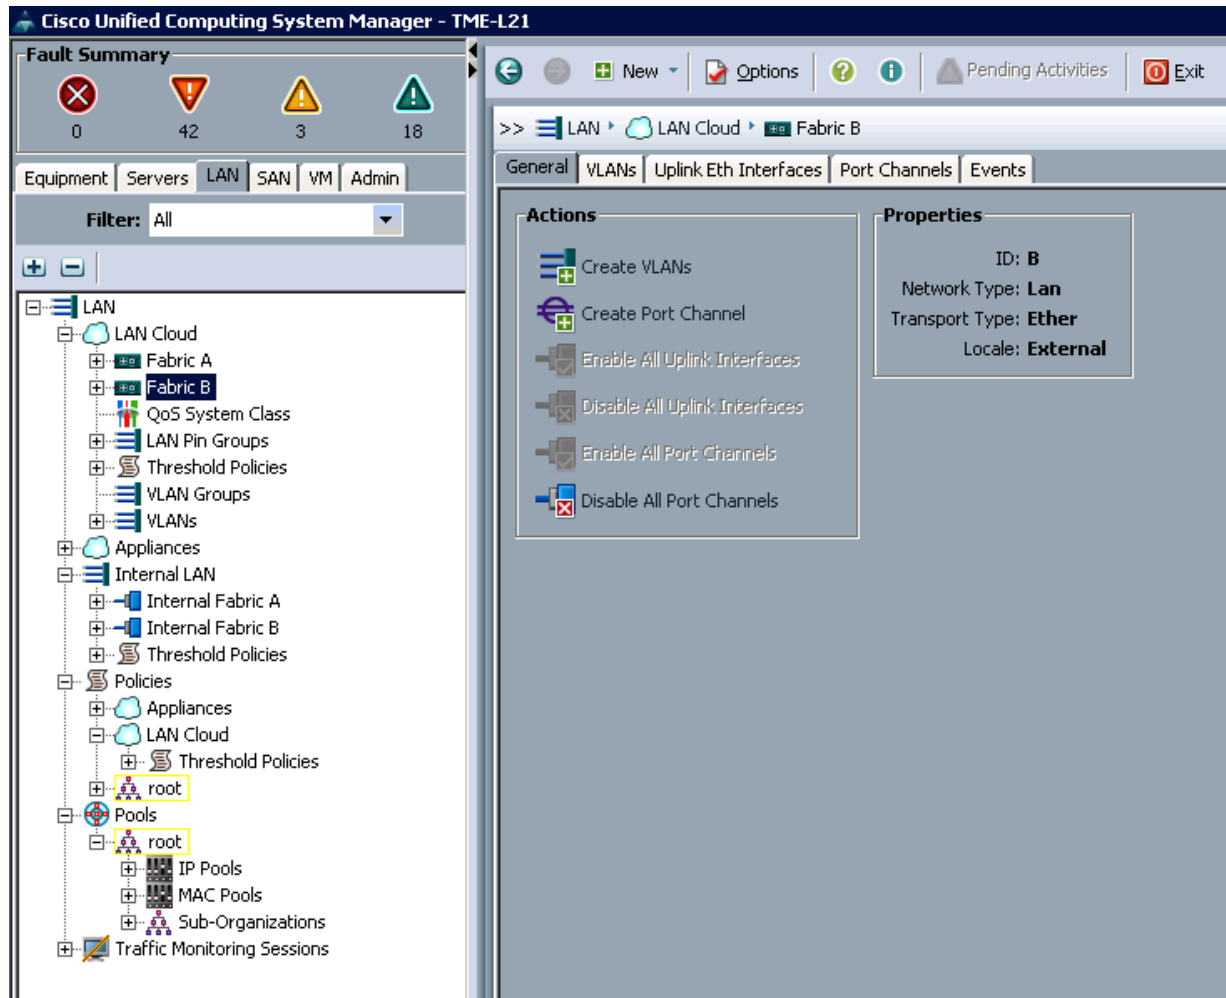


Note

Two Port Channels are created, one from Fabric A to both the Cisco Nexus 5548 switches and other from Fabric B to both the Cisco Nexus 5548 switches.

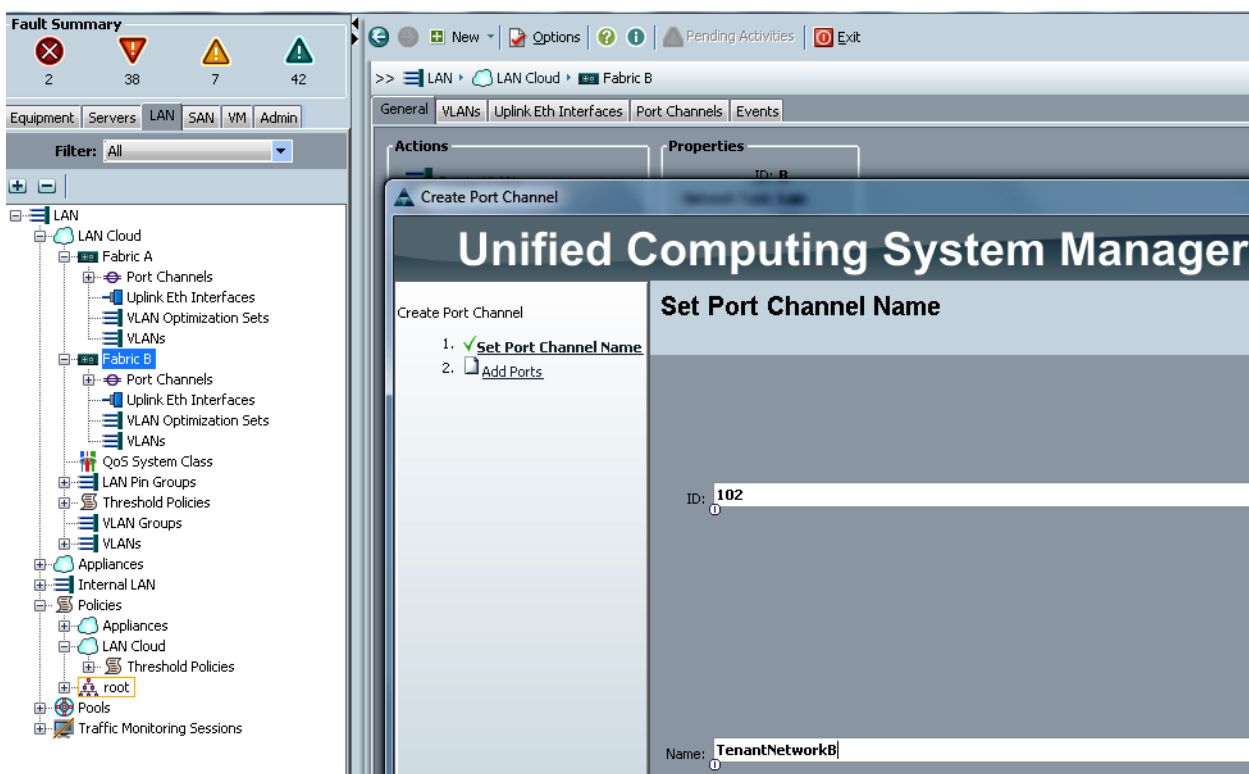
2. Under LAN Cloud, expand the Fabric B tree.
3. In the right pane under **General** tab click **Create Port Channel**.

Figure 15 *Creating the Port Channel*



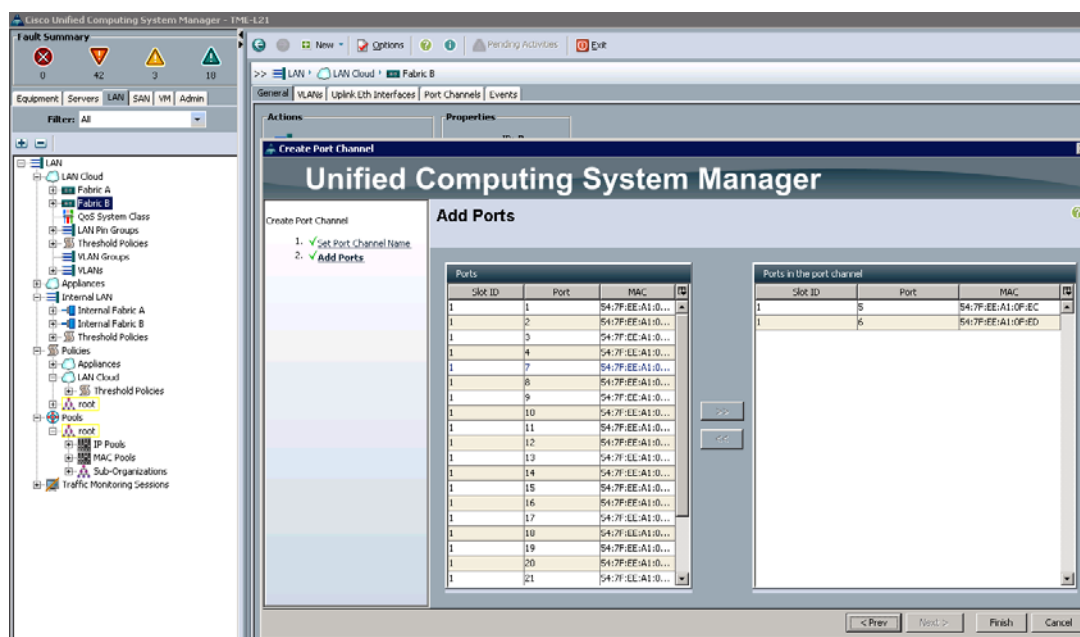
4. Enter 102 as the unique ID of the Port Channel.
5. Enter Tenant NetworkB as the name of the Port Channel.
6. Click **Next**.

Figure 16 **Setting the Port Channel ID and Name**



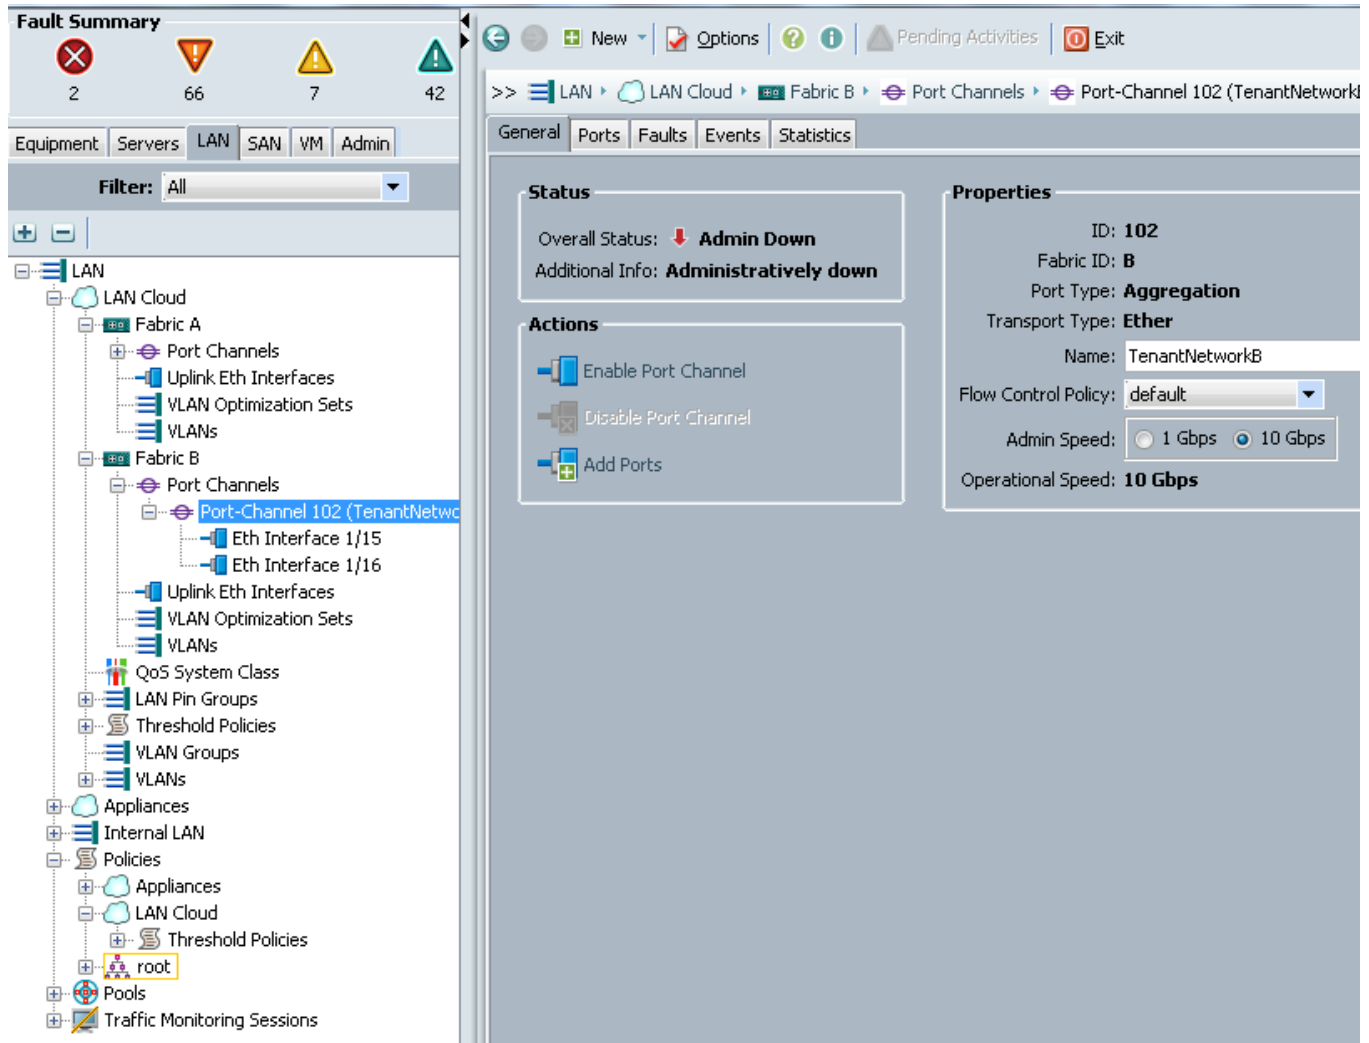
7. Select the port with slot ID: 1 and port: 5, and the port with slot ID: 1 and port 6, from the Ports list.
8. Click >> to add the ports to the Port Channel.

Figure 17 **Selecting and Adding Ports**



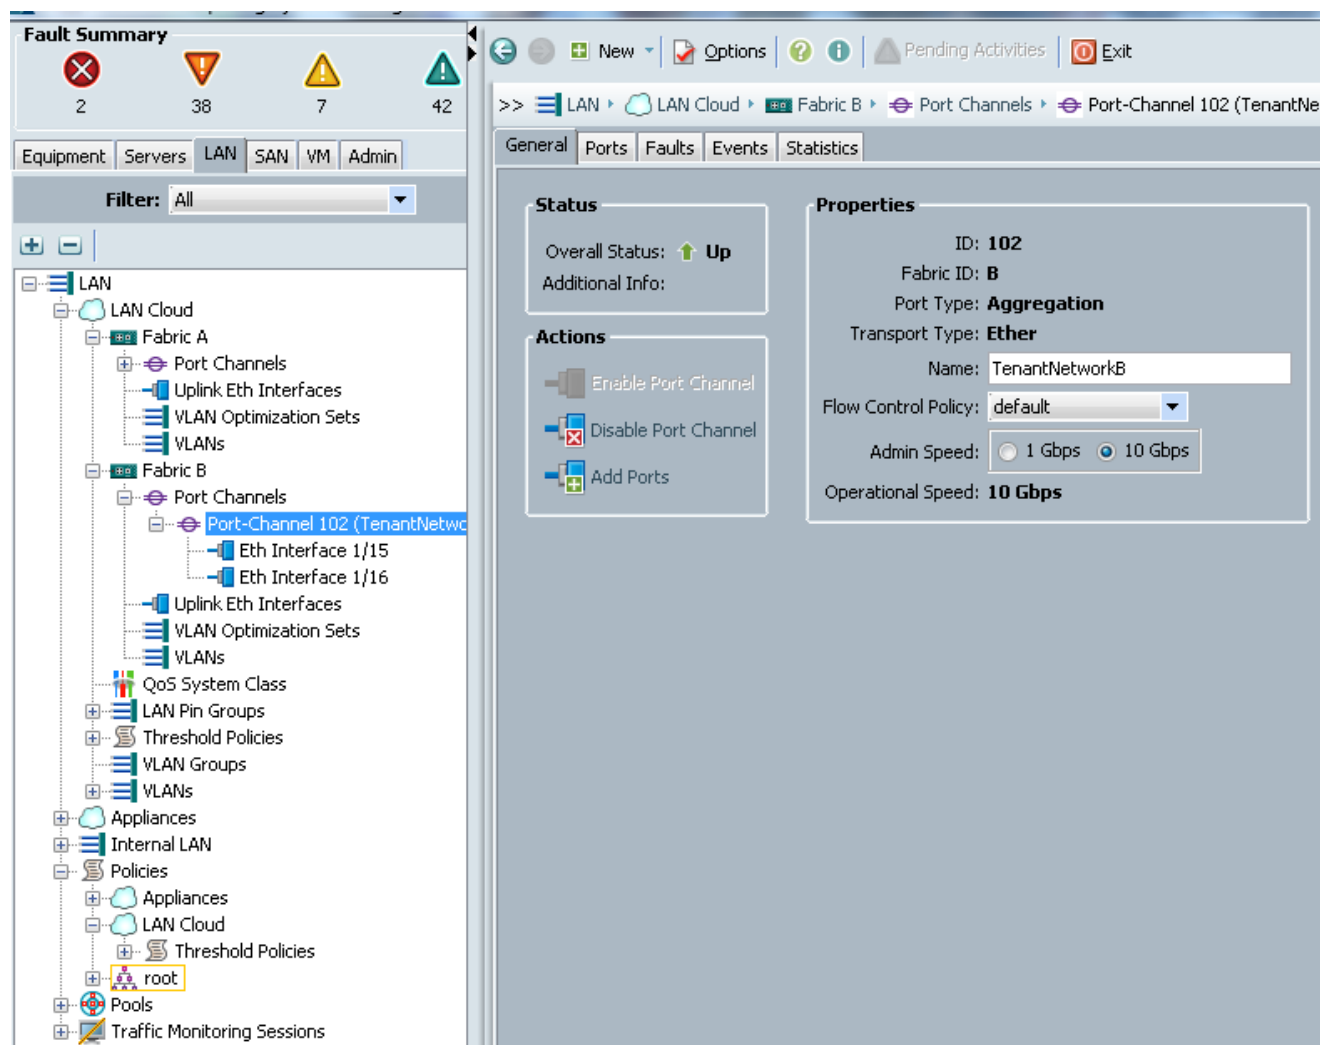
9. Click **Finish** to create the Port Channel.
10. Check the **Show navigator for Port-Channel 101 (Fabric A)** check box.
11. Click **OK** to continue.
12. In the **Actions** area, select **Enable Port Channel**.
13. In the pop-up box, click **Yes**, and then click **OK**.

Figure 18 *Enabling the Port Channel*



14. Wait until the overall status of the Port Channel is Up.
15. Click **OK** to close the navigator.

Figure 19 Verifying the Port Channel Status



Enabling Jumbo Frames

To set the Jumbo frames (9000-byte frame) for the defined Quality of Service classes in the Cisco UCS Fabric through the Cisco UCS Manager, follow these steps:

1. Click the **LAN** tab in the left pane.
2. Choose **LAN Cloud > QoS System Class**.
3. In the right pane, click the **General** tab.
4. Type 9000 in the MTU boxes for the Platinum, Gold, Silver, and Bronze priority value.
5. Click **Save Changes**.
6. Click **OK** to continue.

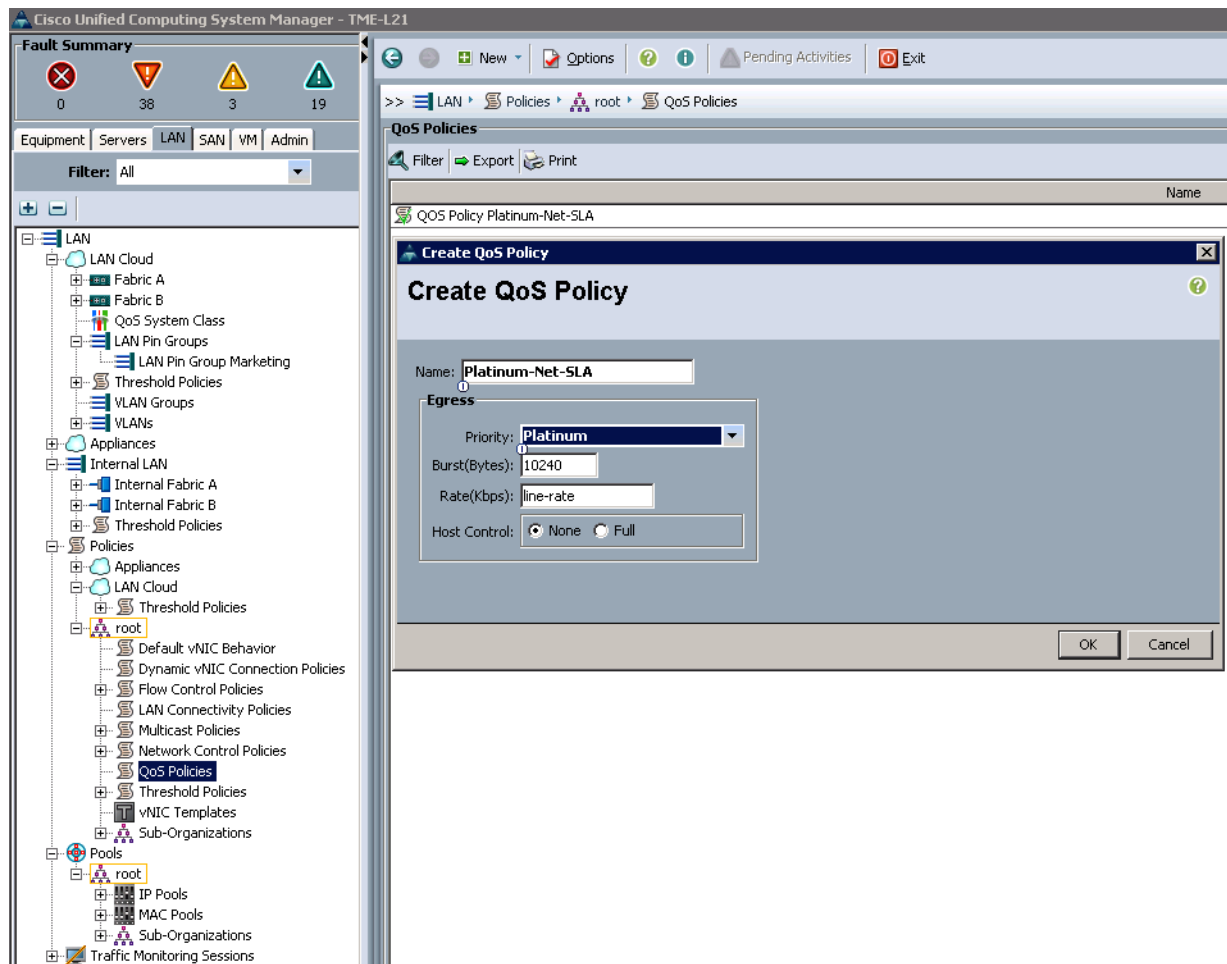
Enabling QoS on Cisco UCS 6248 Fabric Interconnects

In this section the Quality of Service (QoS) Policies-Platinum, Gold, Silver, Bronze, and Fibre Channel, will be created and enabled on the Cisco UCS Fabric in the solution design.

To enable the previously defined Quality of Service Policies on the Cisco UCS Fabric, login to the Cisco UCS Manager and follow these steps:

1. Choose **LAN > Policies > Root > QoS Policies**.
2. Right-click the **QoS Policies**.
3. Select **Create QoS Policy**.
4. Enter **Platinum-Net-SLA** as the QoS Policy name.
5. Change the Priority to **Platinum**. Retain the default values of Burst(Bytes)(10240), Rate (Kbps) (line-rate), and Host Control (None).
6. Click **OK**.

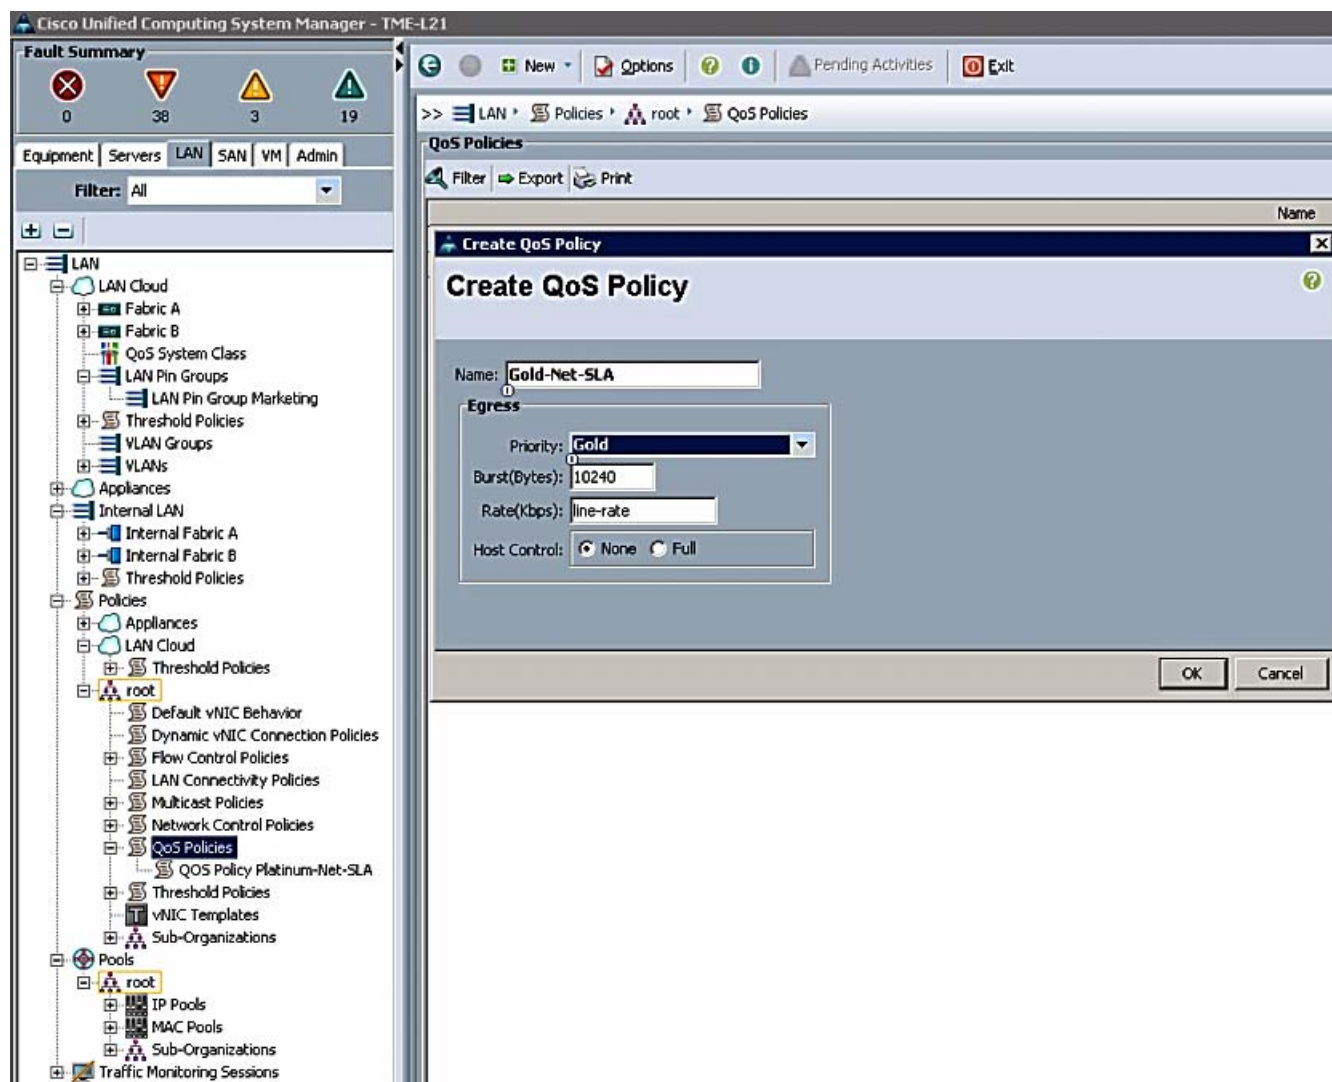
Figure 20 Naming and setting the PlatinumQoS Policy



7. Right-click the **QoS Policies**.
8. Select **Create QoS Policy**.

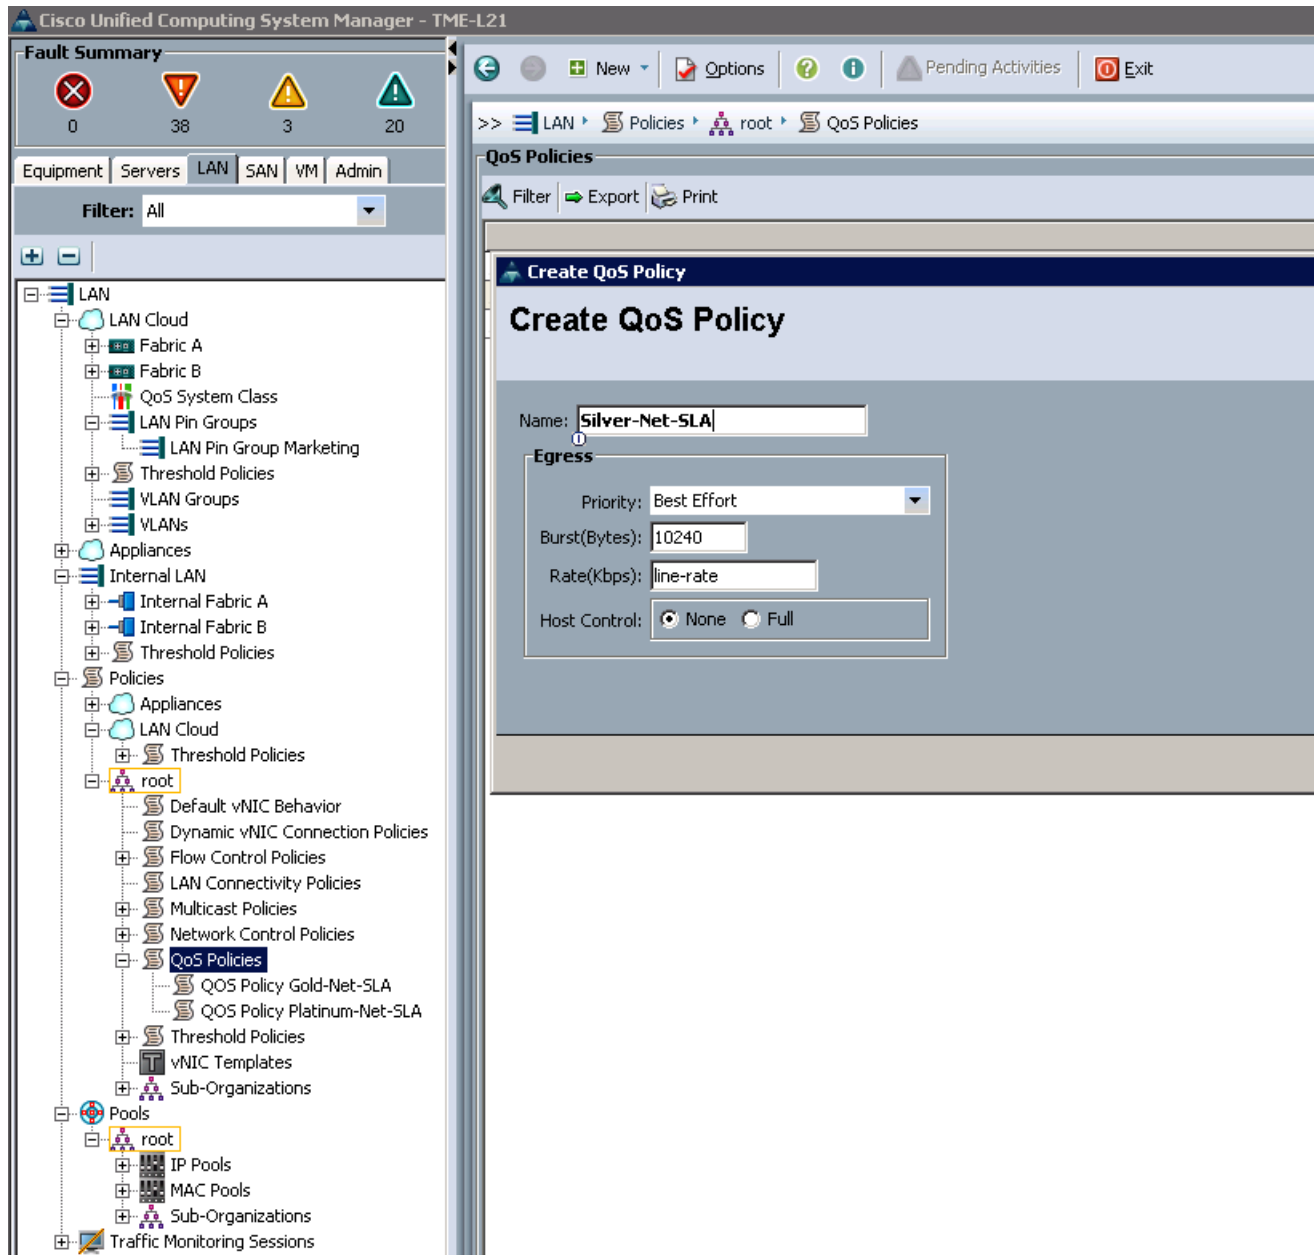
9. Enter Gold-Net-SLA as the QoS Policy name.
10. Change the Priority to Gold. Retain the default values of Burst(Bytes) (10240), Rate(Kbps) (line-rate), and Host Control (None).
11. Click **OK**.

Figure 21 Naming and Setting the Gold QoS Policy



12. Right-click the **QoS Policies**.
13. Select **Create QoS policy**.
14. Enter Silver-Net-SLA as the QoS policy name.
15. Change the Priority to Silver. Retain the default values of Burst(Bytes) (10240), Rate(Kbps) (line-rate), and Host Control (None).
16. Click **OK**.

Figure 22 Naming and Setting the Silver QoS Policy



17. Right-click the **QoS Policies**.
18. Select **Create QoS Policy**.
19. Enter FC Net-SLA as the QoS Policy name.
20. Change the Priority to FC. Retain the default values of Burst(Bytes) (10240), Rate(Kbps) (line-rate), and Host Control (None).
21. Click **OK**.

Figure 23 Naming and Setting the FC QoS Policy

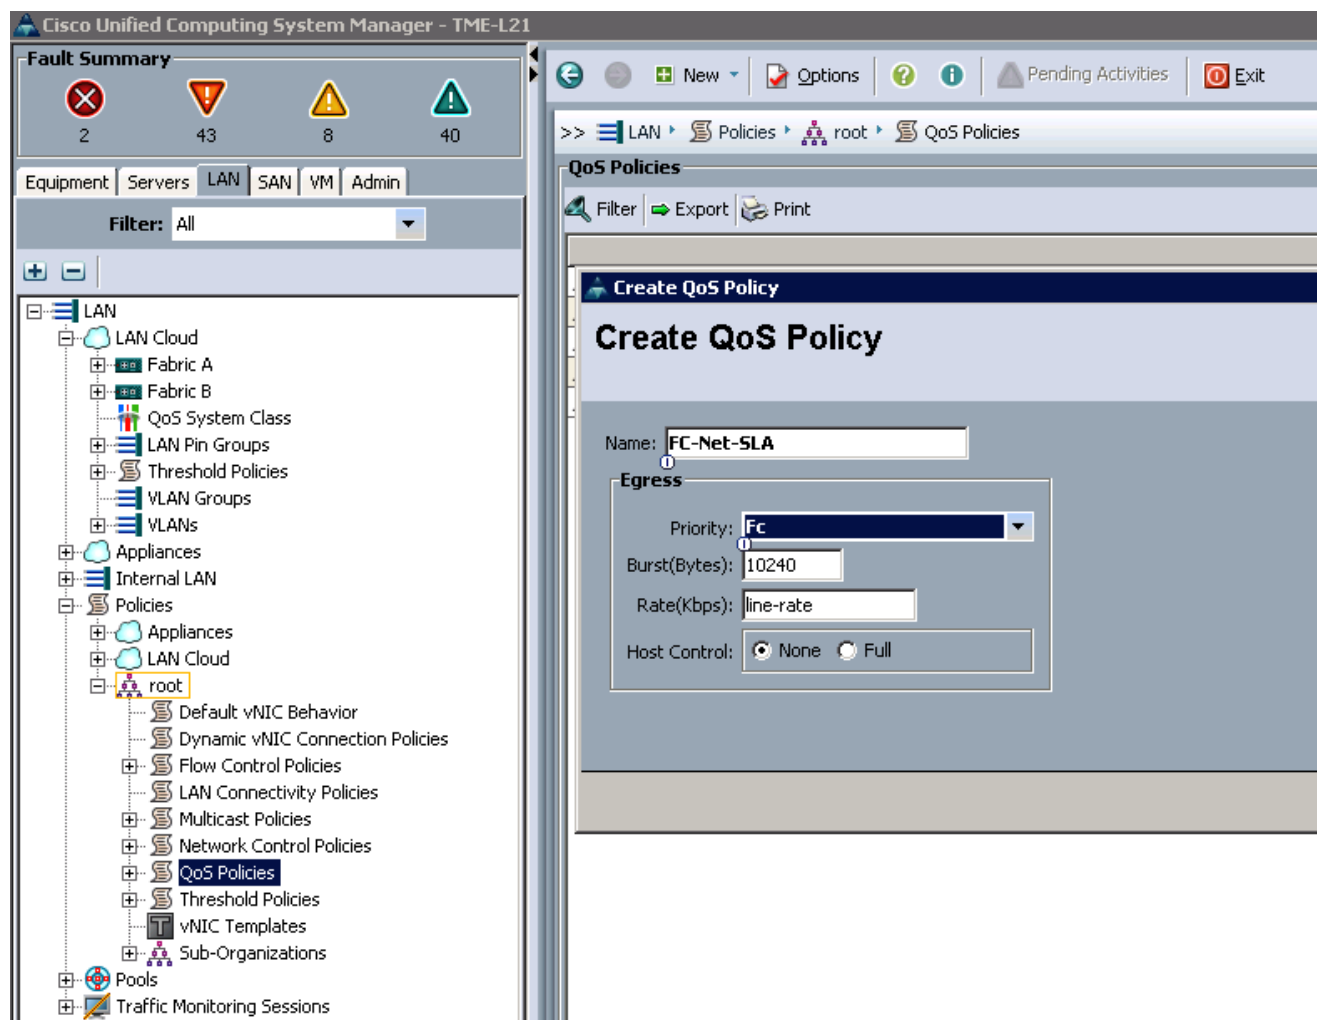
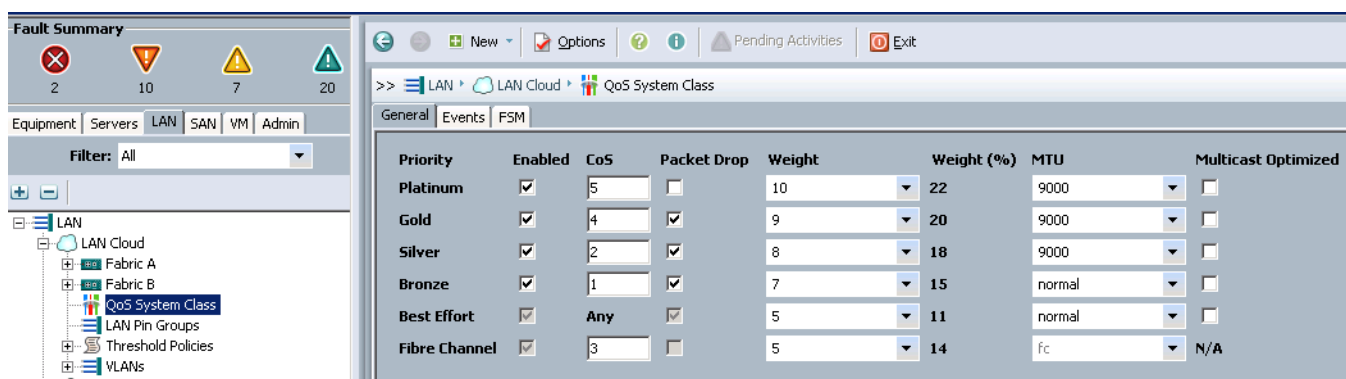


Figure 24 shows Cisco UCS QoS system class and QoS policy configurations defined for application on static and dynamic vNICs for accessing a Microsoft SQL Server iSCSI network.

Figure 24 Defining the Cisco UCS QoS Policy Configurations



QoS Configurations for Cisco Nexus 5548UP Switches

To apply QoS across the entire system, from Cisco UCS to the upstream switches (Cisco Nexus 5548UP Switches), it is essential to configure similar QoS class and policy types with the right class-of-service (CoS) values that correspond to the Cisco UCS QoS classes.

To set the QoS configurations for all the service classes, follow these steps on each Nexus 5548UP from the system console:

1. From the global configuration mode, type `spanning-tree port type network default`, to ensure that by default, the ports are considered as network ports with regard to spanning-tree.
2. Type `spanning-tree port type edge bpduguard default` to enable bpduguard on all edge ports by default.
3. Type `spanning-tree port type edge bpdufilter default` to enable bpdufilter on all edge ports by default.
4. Type `class-map type qos match-any class -- Platinum/Gold/Silver/Bronze/FCoE`.

Platinum class:

- a. Type `description Platinum Class` which is mapped to `Platinum-Network-SLA`.
- b. Type `match cos 5`.
- c. Type `exit`.

Gold class:

- a. Type `description Gold Class` which is mapped to `Gold-Network-SLA`.
- b. Type `match cos 4`.
- c. Type `exit`.

Silver class:

- a. Type `description Gold Class` which is mapped to `Silver-Network-SLA`.
- b. Type `match cos 2`.
- c. Type `exit`.

FCoE class:

- a. Type `match cos 3`.
- b. Type `exit`.

5. Type `policy-map type qos system_qos_policy`.

Platinum class:

- a. Type `class type qos class-Platinum`.
- b. Type `set qos-group 5`.
- c. Type `exit`.

Gold class:

- a. Type `class type qos class-Gold`.
- b. Type `set qos-group 4`.
- c. Type `exit`.

Silver class:

- a. Type `class type qos class-Silver`.
- b. Type `set qos-group 2`.

- c. Type exit.

FCoE class:

- a. Type class type qos class-fcoe.
- b. Type set qos-group 1
- c. Type exit

6. Type class-map type queuing match-all class-

Platinum class:

- a. Type class-map type queuing match-all class-Platinum.
- b. Type description Platinum Class which is mapped to Platinum-Network-SLA.
- c. Type match qos-group 5.
- d. Type exit.

Gold class:

- a. Type class-map type queuing match-all class-Gold.
- b. Type description Gold Class which is mapped to Gold-Network-SLA.
- c. Type match qos-group 4.
- d. Type exit

Silver class:

- a. Type class-map type queuing match-all class-Silver.
- b. Type description Gold Class which is mapped to Silver-Network-SLA.
- c. Type match qos-group 2.
- d. Type exit.

FCoE class:

- a. Type class-map type queuing match-all class-fcoe.
- b. Type match qos-group 1.
- c. Type exit.

7. Type policy-map type queuing system_q_in_policy.

Platinum class:

- a. Type class type qos class-Platinum.
- b. Type bandwidth percent 30.
- c. Type exit

Gold class:

- a. Type class type qos class-Gold.
- b. Type set bandwidth percent 20.
- c. Type exit

Silver class:

- a. Type class type qos class-Silver.
- b. Type set bandwidth percent 15.
- c. Type exit

Default class:

- a. Type class type qos class-default.
- b. Type set bandwidth percent 1.
- c. Type exit

FCoE class:

- a. Type class type qos class-FCoE.
- b. Type set bandwidth percent 34.
- c. Type exit
- d. Type exit

8. Type policy-map type queuing system_q_out_policy.

Platinum class:

- a. Type class type qos class-Platinum.
- b. Type bandwidth percent 30.
- c. Type exit

Gold class:

- a. Type class type qos class-Gold.
- b. Type set bandwidth percent 20.
- c. Type exit

Silver class:

- a. Type class type qos class-Silver.
- b. Type set bandwidth percent 15.
- c. Type exit

Default class:

- a. Type class type qos class-default.
- b. Type set bandwidth percent 1.
- c. Type exit

FCoE class:

- a. Type class type qos class-fcoe.
- b. Type set bandwidth percent 34.
- c. Type exit
- d. Type exit.

9. Type class-map type network-qos class-

Platinum class:

- a. Type class-map type network-qos class-Platinum.
- b. Type description Platinum Class which is mapped to Platinum-Network-SLA.
- c. Type match qos-group 5.
- d. Type exit.

Gold class:

- a. Type class-map type network-qos class-Gold.
- b. Type description Gold Class which is mapped to Gold-Network-SLA.
- c. Type match qos-group 4.
- d. Type exit.

Silver class:

- a. Type class-map type network-qos class-Silver.
- b. Type description Platinum Class which is mapped to Silver-Network-SLA.
- c. Type match qos-group 2.
- d. Type exit.

FCoE class:

- a. Type class-map type network-qos class-fcoe.
- b. Type match qos-group 1.
- c. Type exit.

10. Type policy-map type network-qos system_nq_policy.

Platinum class:

- a. Type class type qos class-Platinum.
- b. Type set.
- c. Type set cos 5.
- d. Type mtu 9000.
- e. Type exit.

Gold class:

- a. Type class type qos class-Gold.
- b. Type set.
- c. Type set cos 4.
- d. Type mtu 9000.
- e. Type exit.

Silver class:

- a. Type class type qos class-Silver.
- b. Type set.
- c. Type set cos 2.
- d. Type mtu 9000.
- e. Type exit.

FCoE class:

- a. Type class type qos class-fcoe.
- b. Type set.
- c. Type set cos 1.
- d. Type mtu 9000.
- e. Type pause no-drop.

- f. Type exit.
- Default class:**
 - a. Type class type qos class-default.
 - b. Type set.
 - c. Type mtu 9000.
 - d. Type exit.
11. Type system qos.
12. Type service-policy type qos input system_qos_policy.
13. Type service-policy type queuing input system_q_in_policy.
14. Type service-policy type queuing output system_q_out_policy.
15. Type service-policy type network-qos system_nq_policy.
16. Type exit.
17. Type interface port-channel 103.
18. Type service-policy type qos input system_qos_policy.
19. Type service-policy type queuing input system_q_in_policy.
20. Type service-policy type queuing output system_q_out_policy.
21. Type interface port-channel 104.
22. Type service-policy type qos input system_qos_policy.
23. Type service-policy type queuing input system_q_in_policy.
24. Type service-policy type queuing output system_q_out_policy.
25. Type copy run start.

**Note**

Configure the same QoS Policy on Nexus 5548UP Switch B which is part of VPC cluster, repeat steps 1-25 on Nexus 5548UP Switch B.

The commands enumerated below show the QoS configuration output at both Cisco Nexus 5548UP Switch A and B.

```
sh class-map type queuing
  Type queuing class-maps
  =====
  class-map type queuing class-Gold
    match qos-group 4
class-map type queuing class-fcoe
  match qos-group 1
  class-map type queuing class-Silver
    match qos-group 2
  class-map type queuing class-default
    match qos-group 0
  class-map type queuing class-Platinum
    match qos-group 5
  class-map type queuing class-all-flood
    match qos-group 2
  class-map type queuing class-ip-multicast
    match qos-group 2
N5548-L21-A(config)# sh class-map type network-qos
  Type network-qos class-maps
```



```

=====
class-map type network-qos class-Gold
  match qos-group 4
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-Silver
  match qos-group 2
class-map type network-qos class-default
  match qos-group 0
class-map type network-qos class-Platinum
  match qos-group 5
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
N5548-L21-A(config)# sh class-map type queuing
Type queuing class-maps
=====
class-map type queuing class-Gold
  match qos-group 4
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-default
  match qos-group 0
class-map type queuing class-Platinum
  match qos-group 5
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
sh policy-map type queuing
Type queuing policy-maps
=====
policy-map type queuing default-in-policy
  class type queuing class-default
    bandwidth percent 100
policy-map type queuing system_qos_policy
  class type queuing class-Silver
    priority
  class type queuing class-default
    bandwidth percent 100
policy-map type queuing default-out-policy
  class type queuing class-default
    bandwidth percent 100
policy-map type queuing system_q_in_policy
  class type queuing class-Platinum
    bandwidth percent 30
  class type queuing class-Gold
    bandwidth percent 20
  class type queuing class-Silver
    bandwidth percent 15
  class type queuing class-fcoe
    bandwidth percent 34
  class type queuing class-default
    bandwidth percent 1
policy-map type queuing system_q_out_policy
  class type queuing class-Platinum
    bandwidth percent 30
  class type queuing class-Gold
    bandwidth percent 20
  class type queuing class-Silver
    bandwidth percent 15
  class type queuing class-fcoe
    bandwidth percent 34

```

```

class type queuing class-default
    bandwidth percent 1
policy-map type queuing fcoe-default-in-policy
    class type queuing class-fcoe
        bandwidth percent 50
    class type queuing class-default
        bandwidth percent 50
policy-map type queuing fcoe-default-out-policy
    class type queuing class-fcoe
        bandwidth percent 50
    class type queuing class-default
        bandwidth percent 50
sh policy-map type qos
Type qos policy-maps
=====
policy-map type qos default-in-policy
    class type qos class-default
        set qos-group 0
policy-map type qos system_qos_policy
    class type qos class-Platinum
        set qos-group 5

    class type qos class_fcoe
        set qos-group 3
    class type qos class-default
        set qos-group 0
policy-map type qos fcoe-default-in-policy
    class type qos class-fcoe
        set qos-group 1
    class type qos class-default
        set qos-group 0
sh policy-map type network-qos
Type network-qos policy-maps
=====
policy-map type network-qos system_nq_policy
    class type network-qos class-Platinum
        set cos 5
        mtu 9000
    class type network-qos class-fcoe
        mtu 2158
        set cos 3
        pause no-drop
    class type network-qos class-Gold
        mtu 9000
    class type network-qos class-Silver
        mtu 9000
    class type network-qos class-default
        mtu 1500
    multicast-optimize
policy-map type network-qos default-nq-policy
    class type network-qos class-default
        mtu 9000
        multicast-optimize
policy-map type network-qos fcoe-default-nq-policy
    class type network-qos class-fcoe
        pause no-drop
        mtu 2158
    class type network-qos class-default
        mtu 1500
        multicast-optimize

```

NetApp Storage QoS Configuration



Note

On the Cisco Nexus 5548UP upstream switch, ensure that the correct QoS class and MTU value with policy types are applied to the Port Channel Ports (eth19 and eth 20). Port channels are connected to the NetApp FAS3270HA (Controllers Flexpod A and B), 10 Gigabit Ethernet interfaces (e1c and e1d), to allow network packets to be tagged from Nexus 5548 fabric. This is done because NetApp Storage will not tag any network packets with MTU and QoS values.

Following commands shows how to configure the CoS on Nexus 5548 for untagged packets originating from storage on the Port Channels.

CLI commands on Cisco Nexus 5548UP Application1

```
Switch# Configure Terminal
Switch(Config)# Interface port channel 105
Switch(Config-if)#untagged cos 2
Switch# sh policy-map type qos
Switch# Configure Terminal
Switch(Config)# Interface port channel 106
Switch(Config-if)#untagged cos 2
Switch# sh policy-map type qos
```

CLI commands on Cisco Nexus 5548UP Application2

```
Switch# Configure Terminal
Switch(Config)# Interface port channel 105
Switch(Config-if)#untagged cos 2
Switch# sh policy-map type qos
Switch# Configure Terminal
Switch(Config)# Interface port channel 106
Switch(Config-if)#untagged cos 2
Switch# sh policy-map type qos
```

Make sure that the MTU is set to 9000 and that jumbo frames are enabled on the Cisco UCS static and dynamic vNICs, and on the upstream Cisco Nexus 5548UP Switches.

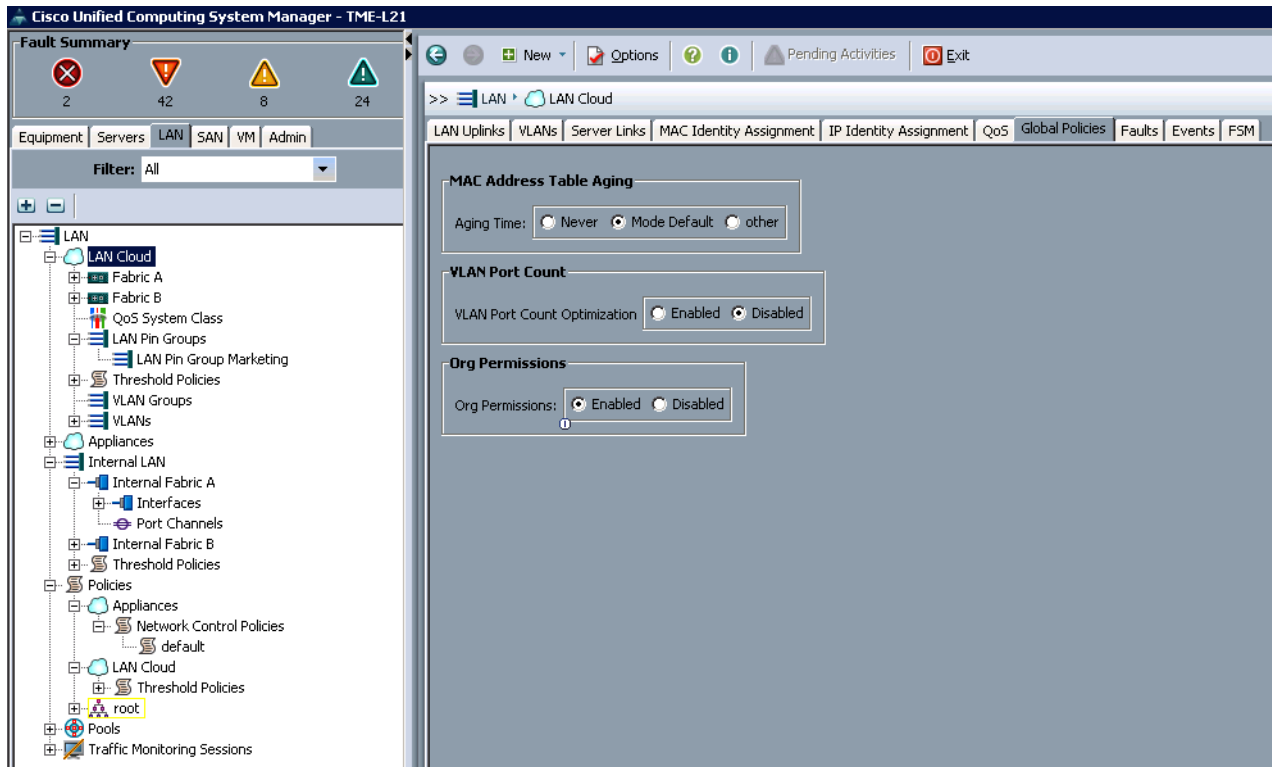
Configuring VLANs on Cisco UCS 6248UP Fabric Interconnects

In multitenant cloud environment, isolation of network data access is a critical requirement to adhere to security measure. Cisco UCS allows creation of logical VLANs that can be dedicated to, and accessed by a specific organization unit.

To configure the necessary dedicated VLANs to an organization unit in Cisco UCS, follow these steps:

1. Click the **LAN** tab in the right pane, click **Global Policies**.
2. In the right pane under **Org Permissions** area, click **Enabled** radio button.
3. Click **Save Changes**.

Figure 25 *Configuring the VLANs to the Organization Unit*



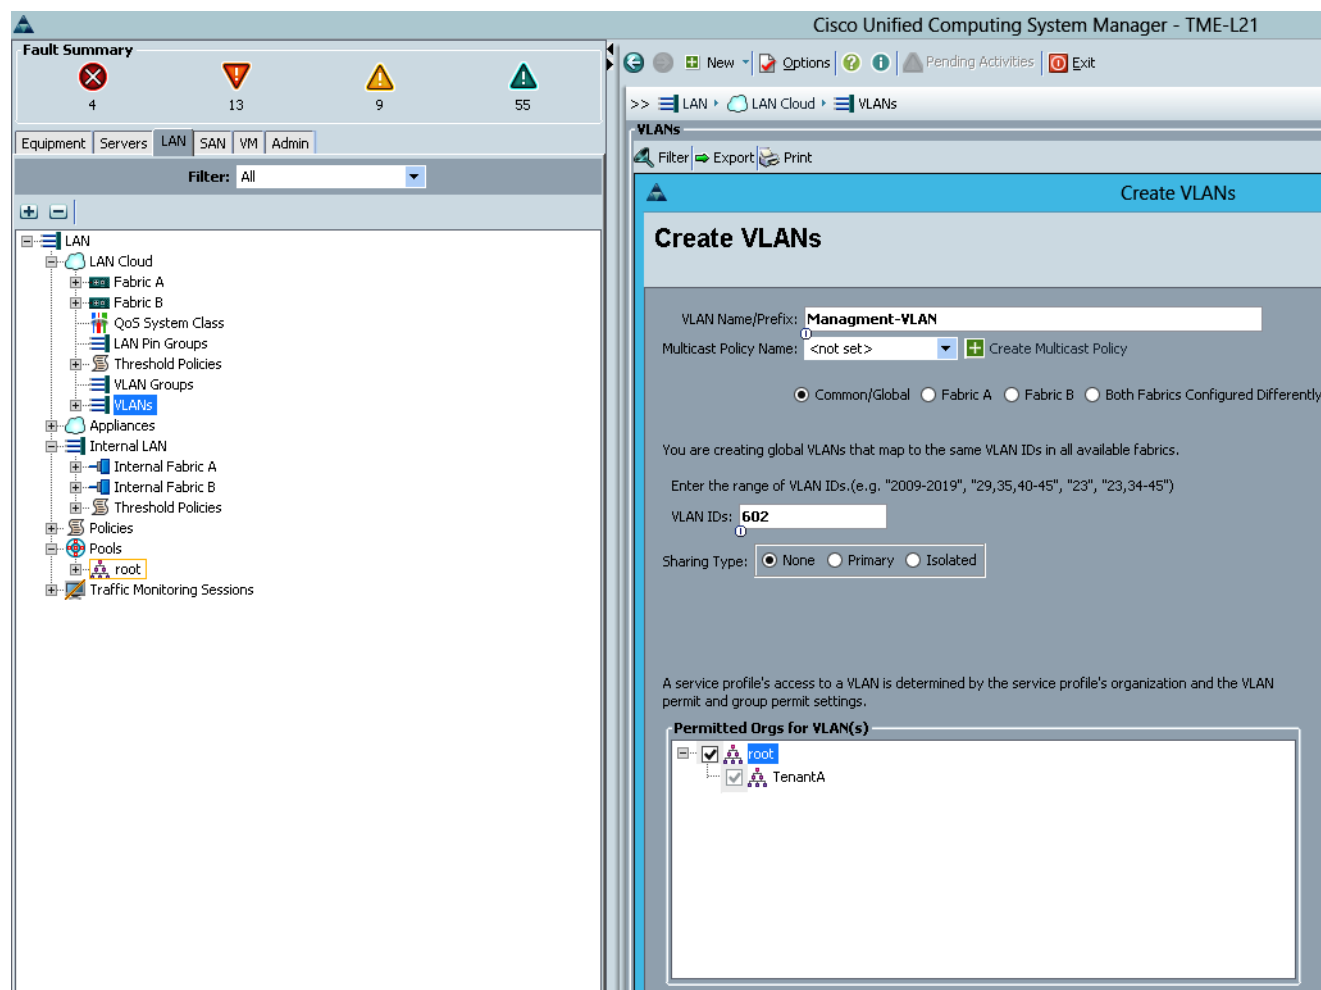
4. Click the **LAN** tab in the left pane.
5. Select **LAN Cloud**.
6. Right-click **VLANs**. (Creating Management VLAN)
7. Select **Create VLANs**.
8. Enter Management-VLAN in the Name field.



Note This VLAN will be used for management traffic.

9. Under **Multicast Policy Name** select default setting.
10. Keep the **Common/Global** option selected for the scope of the VLAN.
11. Enter the VLAN ID defined for the management VLAN. Retain the sharing type as none.
12. In **Permitted Orgs for VLAN(s)** select **Root & TenantA Organization Unit**.
13. Click **OK**.

Figure 26 Defining the Management VLAN Properties



14. Right-click **VLANs**. (Creating the BareMetal-PXE VLAN)
15. Select **Create VLANs**.
16. Enter BareMetal-VLAN in the Name field.



Note This VLAN will be used for the BareMetal host PXE Boot traffic.

17. Under **Multicast Policy Name** select default setting.
18. Keep the **Common/Global** option selected for the scope of the VLAN.
19. Enter the VLAN ID for the BareMetal-PXE VLAN.
20. In **Permitted Orgs for VLAN(s)** select **root** organization Unit.
21. Click **OK**.

Figure 27 *Defining BareMetal-PXE VLAN Properties*

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name: [Create Multicast Policy](#)

☒ Common/Global ☐ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: ☒ None ☐ Primary ☐ Isolated

A service profile's access to a VLAN is determined by the service profile's organization and the VLAN permit and group permit settings.

Permitted Orgs for VLAN(s)

- ☒ root
- ☒ TenantA

22. Right-click **VLANs**. (Creating the guest VLAN)

23. Select **Create VLANs**.

24. Enter guest-VLAN in the Name field.



Note This VLAN will be used for the guest VLAN traffic.

25. Under **Multicast Policy Name** select **Default** setting.

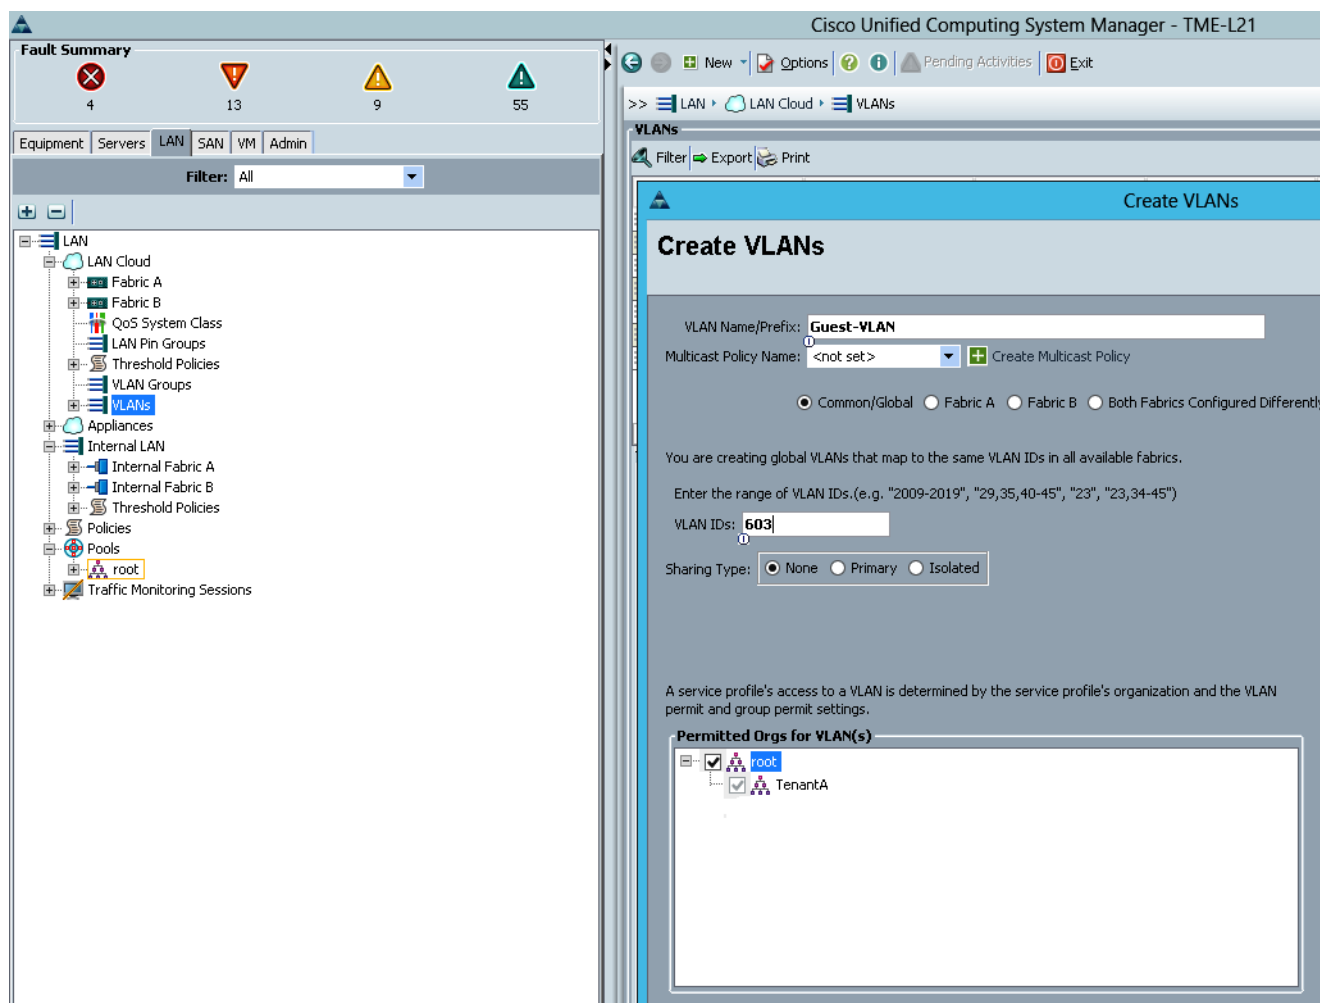
26. Keep the Common/Global option selected for the scope of the VLAN.

27. Enter the VLAN ID for the guest VLAN.

28. In Permitted Orgs for VLAN(s) select Root & TenantA Organization Unit.

29. Click **OK**.

Figure 28 Defining the guest VLAN Properties



30. Right-click **VLANs**. (Creating the guest VLAN)

31. Select **Create VLANs**.

32. Enter guest-VLAN-1000 in the Name field.



Note This VLAN will be used for the guest VLAN traffic.

33. Under Multicast Policy Name select Default setting.

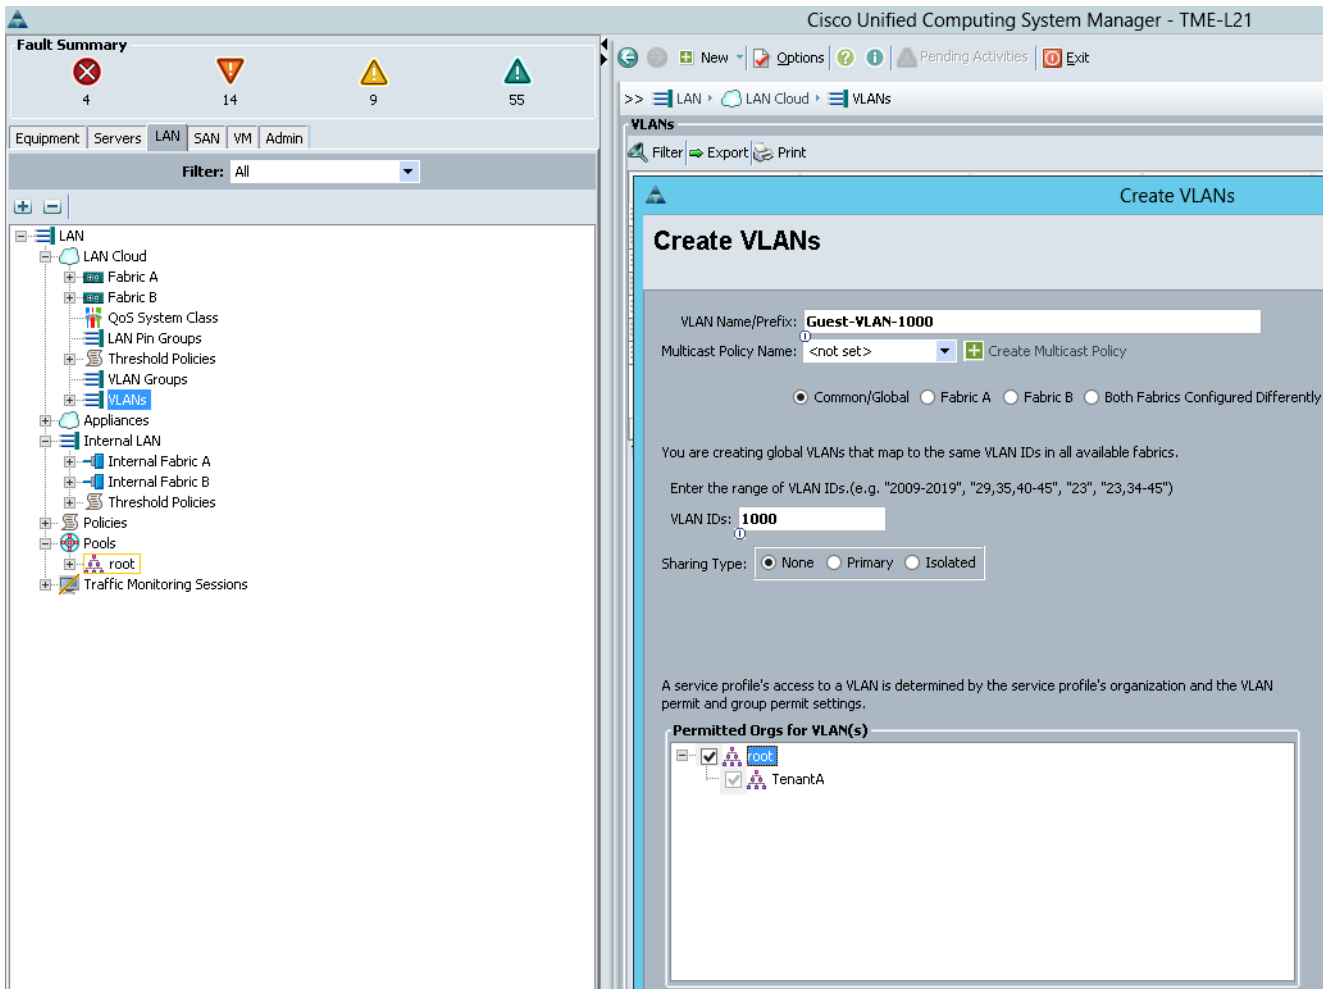
34. Keep the Common/Global option selected for the scope of the VLAN.

35. Enter the VLAN ID for the guest VLAN.

36. In Permitted Orgs for VLAN(s) select Root & TenantA Organization Unit.

37. Click **OK**.

Figure 29 *Defining the guest VLAN 1000 Properties*



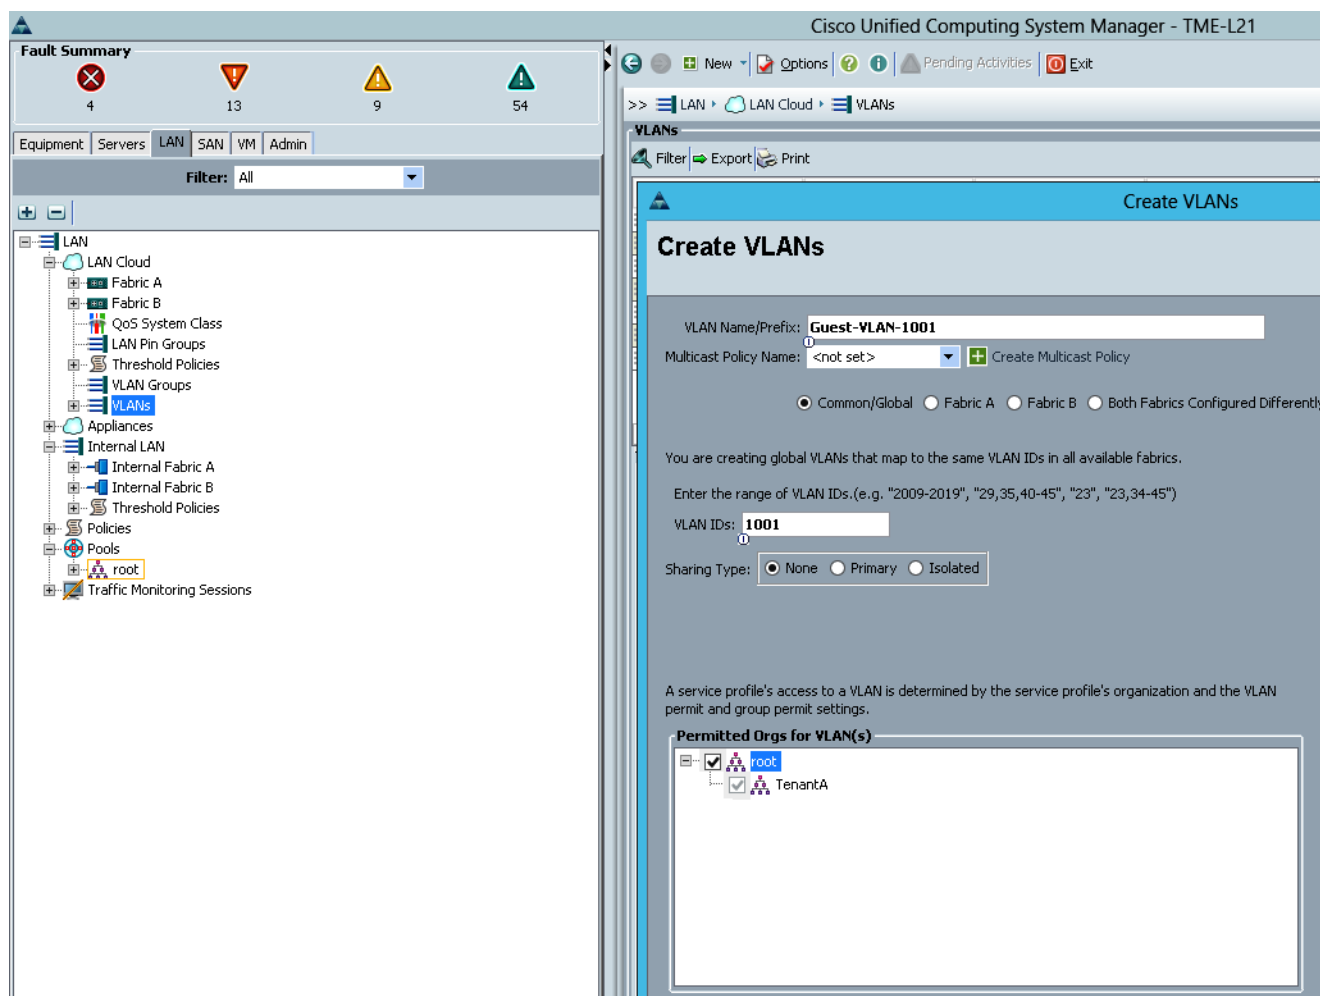
38. Right-click **VLANs**. (Creating the guest VLAN)
39. Select Create VLANs.
40. Enter guest-VLAN-1001 in the Name field.



Note This VLAN will be used for the guest VLAN traffic.

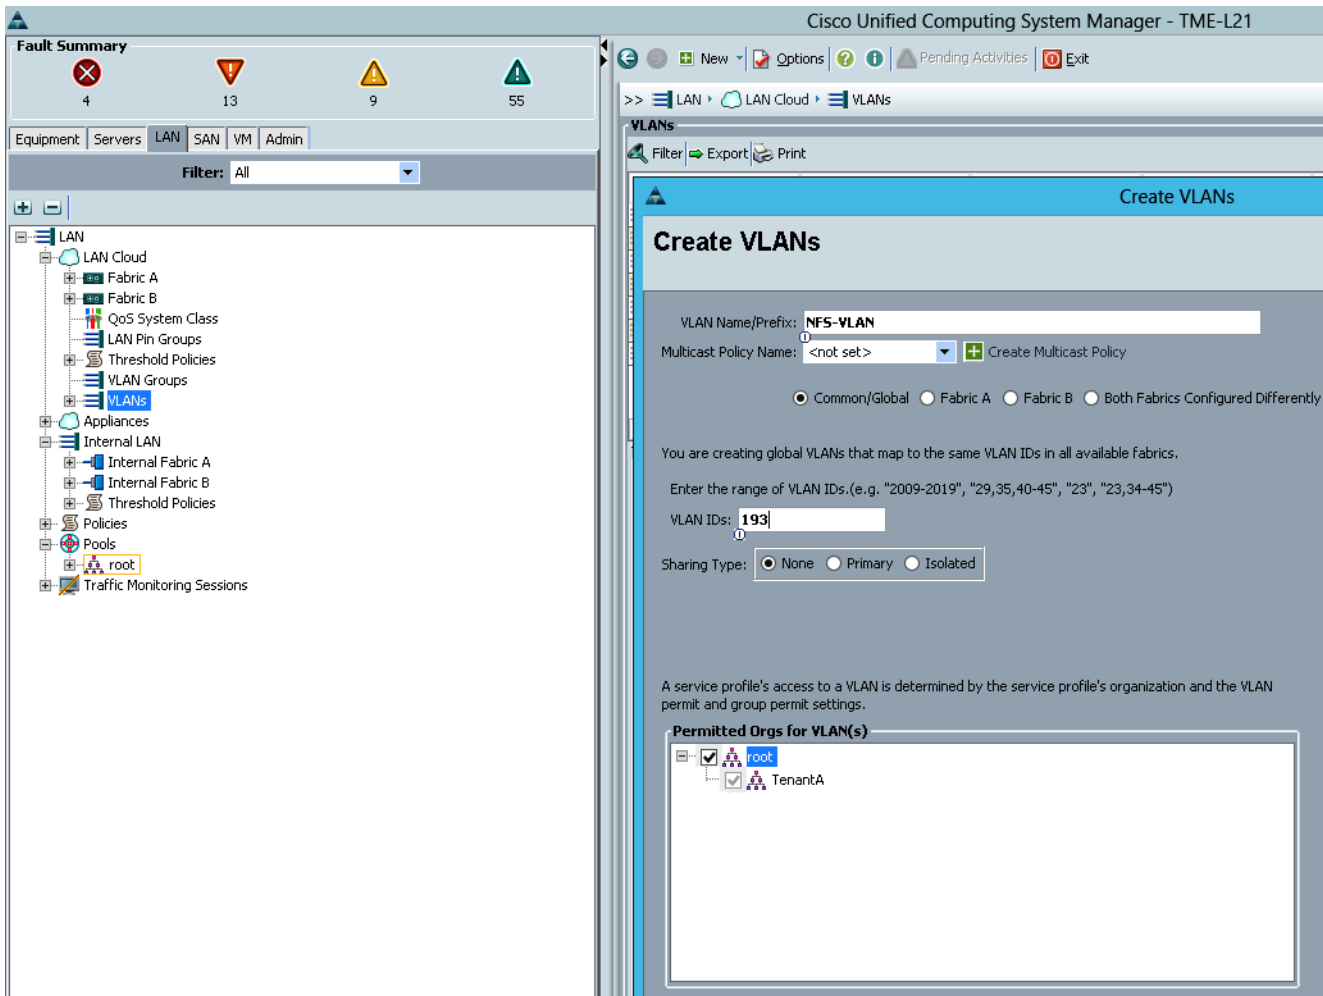
41. Under Multicast Policy Name select Default setting.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the VLAN ID for the guest VLAN.
44. In Permitted Orgs for VLAN(s) select Root & TenantA Organization Unit.
45. Click **OK**.

Figure 30 Defining the guest VLAN 1001 Properties



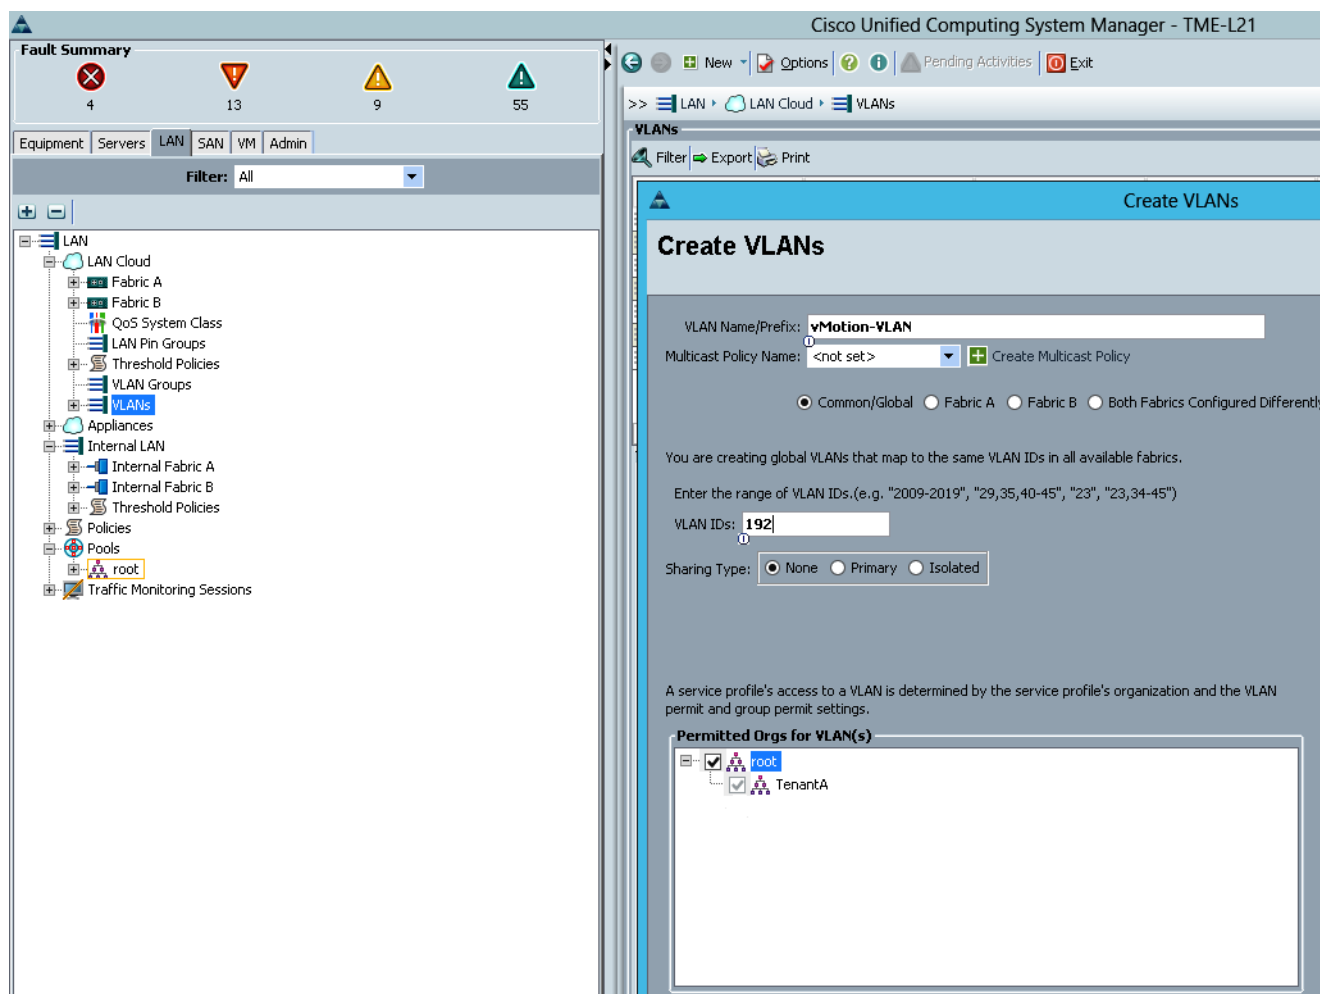
46. Right-click **VLANs**. (Creating the NFS VLAN)
47. Select **Create VLANs**.
48. Enter **NFS-VLAN** in the Name field. This VLAN will be used for the Secondary Storage NFS VLAN traffic.
49. Under Multicast Policy Name select Default setting.
50. Keep the Common/Global option selected for the scope of the VLAN.
51. Enter the VLAN ID for the NFS VLAN.
52. Keep the Common/Global option selected for the scope of the VLAN.
53. In Permitted Orgs for VLAN(s) select TenantA Organization Unit.
54. Click **OK**.

Figure 31 **Defining NFS VLAN Properties**



55. Right-click **VLANs**. (Creating the vMotion VLAN)
56. Select **Create VLANs**.
57. Enter vMotion-VLAN in the Name field. This VLAN will be used for the VM vMotion VLAN traffic.
58. Under Multicast Policy Name select Default setting.
59. Keep the Common/Global option selected for the scope of the VLAN.
60. Enter the VLAN ID for the vMotion VLAN.
61. In Permitted Orgs for VLAN(s) select TenantA Organization Unit.
62. Click **OK**.

Figure 32 **Defining vMotion VLAN Properties**



Storage Configurations

This section details all the configurations and setup steps required to meet the Cisco storage design considerations. In this study we will create Virtual Logical SAN Networks (VSANs) for isolating multi-tenants storage data traffic and create FCoE Port Channel for link aggregation.

The following tasks are described in this section;

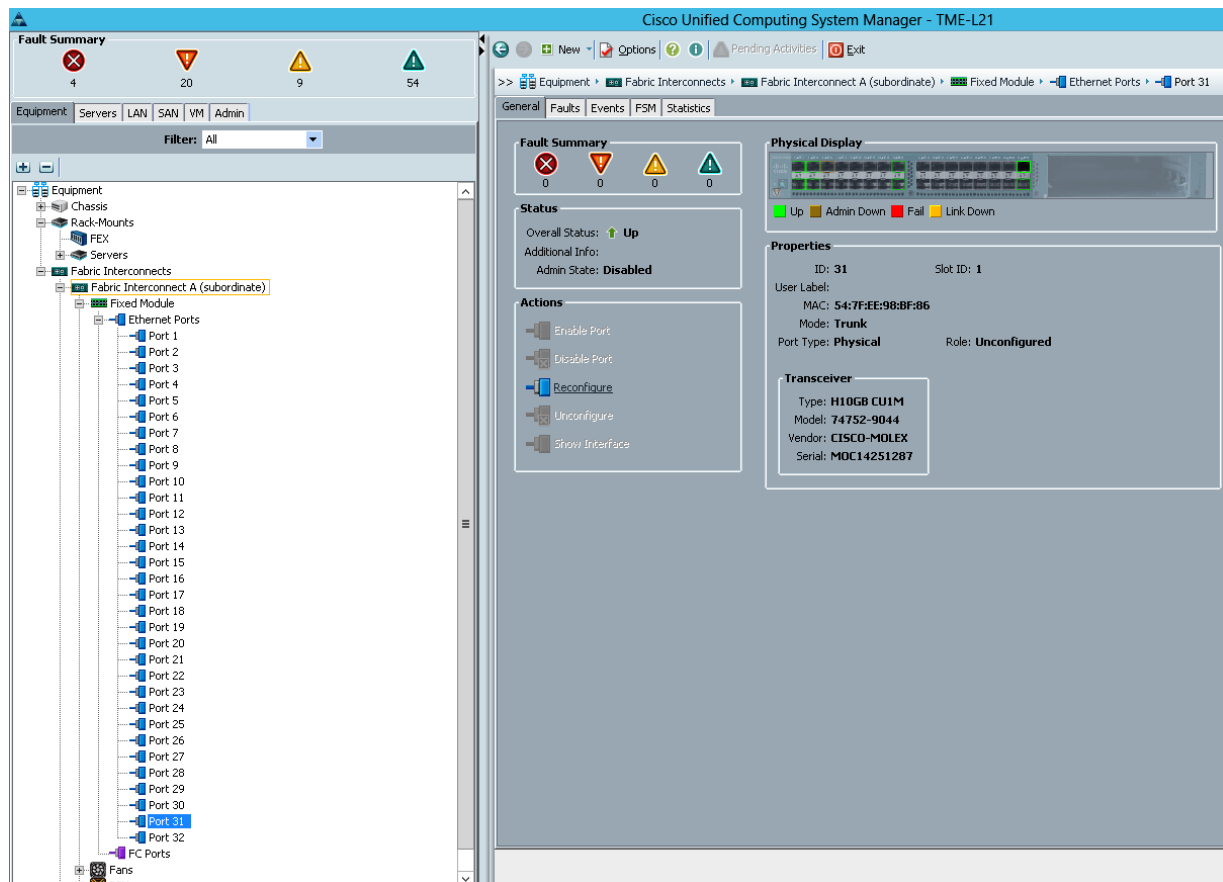
- Configure FCoE port
- Create VSAN
- Define VSAN and SAN port channel
- Configure FCoE Port channel on FI

Configuring FCoE Ports

FCoE ports are ports on the 6200 series Fabric Interconnects that can be configured to carry Fibre Channel traffic over Ethernet. These ports are not reserved. They cannot be used by a Cisco UCS domain until you configure them. To modify an unconfigured Ethernet port into a FCoE uplink port in the Cisco UCS environment, follow these steps:

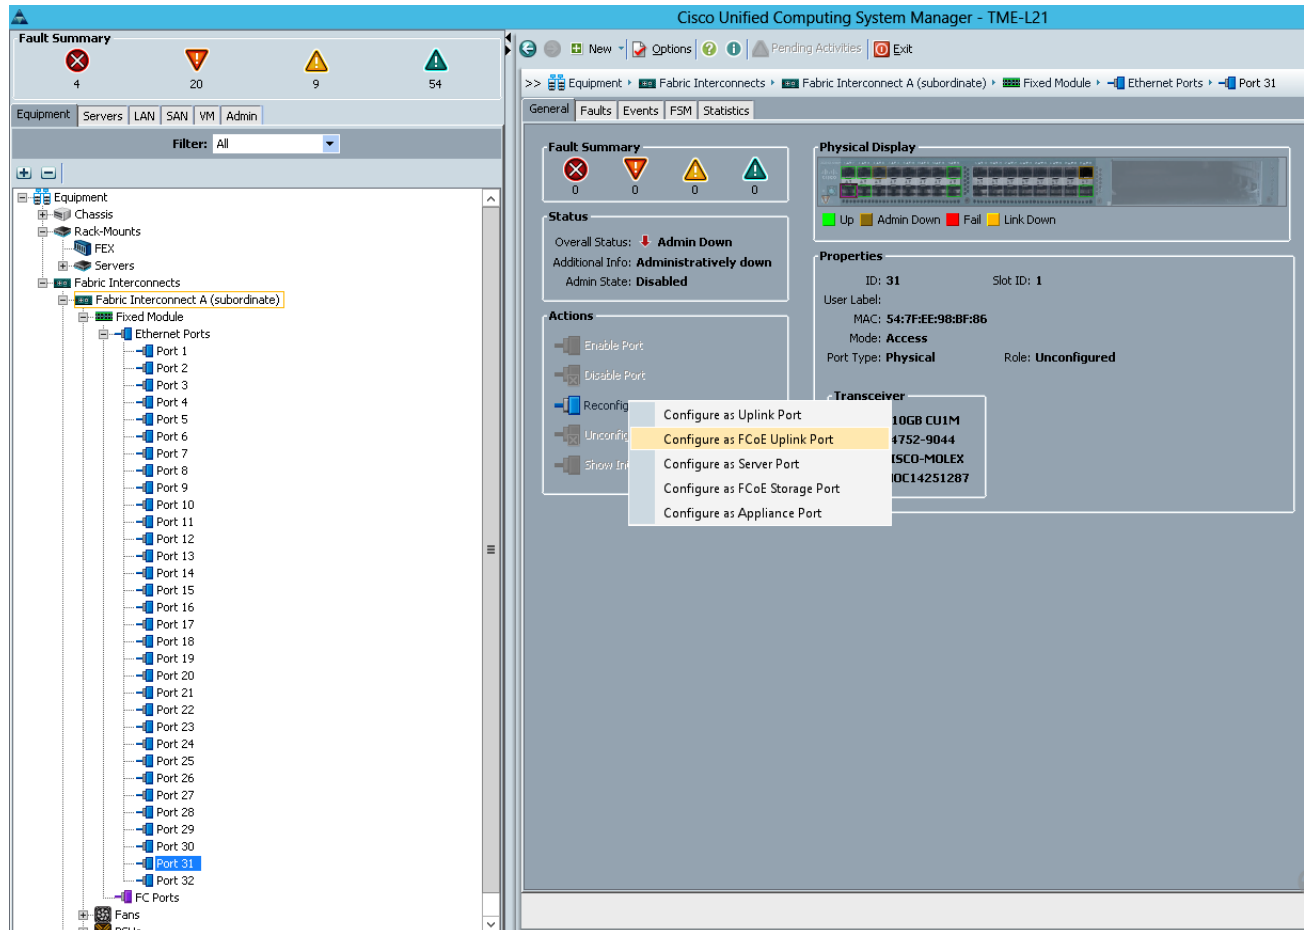
1. In the Cisco UCS Manager GUI, click the **Equipment** tab in the left pane. (Configure FCoE Ports on FI A)
2. Select Fabric Interconnects A.
3. Select Ethernet Ports.
4. Select Port 31.
5. In the right pane, click the **General** tab.
6. Select Reconfigure.

Figure 33 Reconfigure ethernet port to FCoE Uplink port



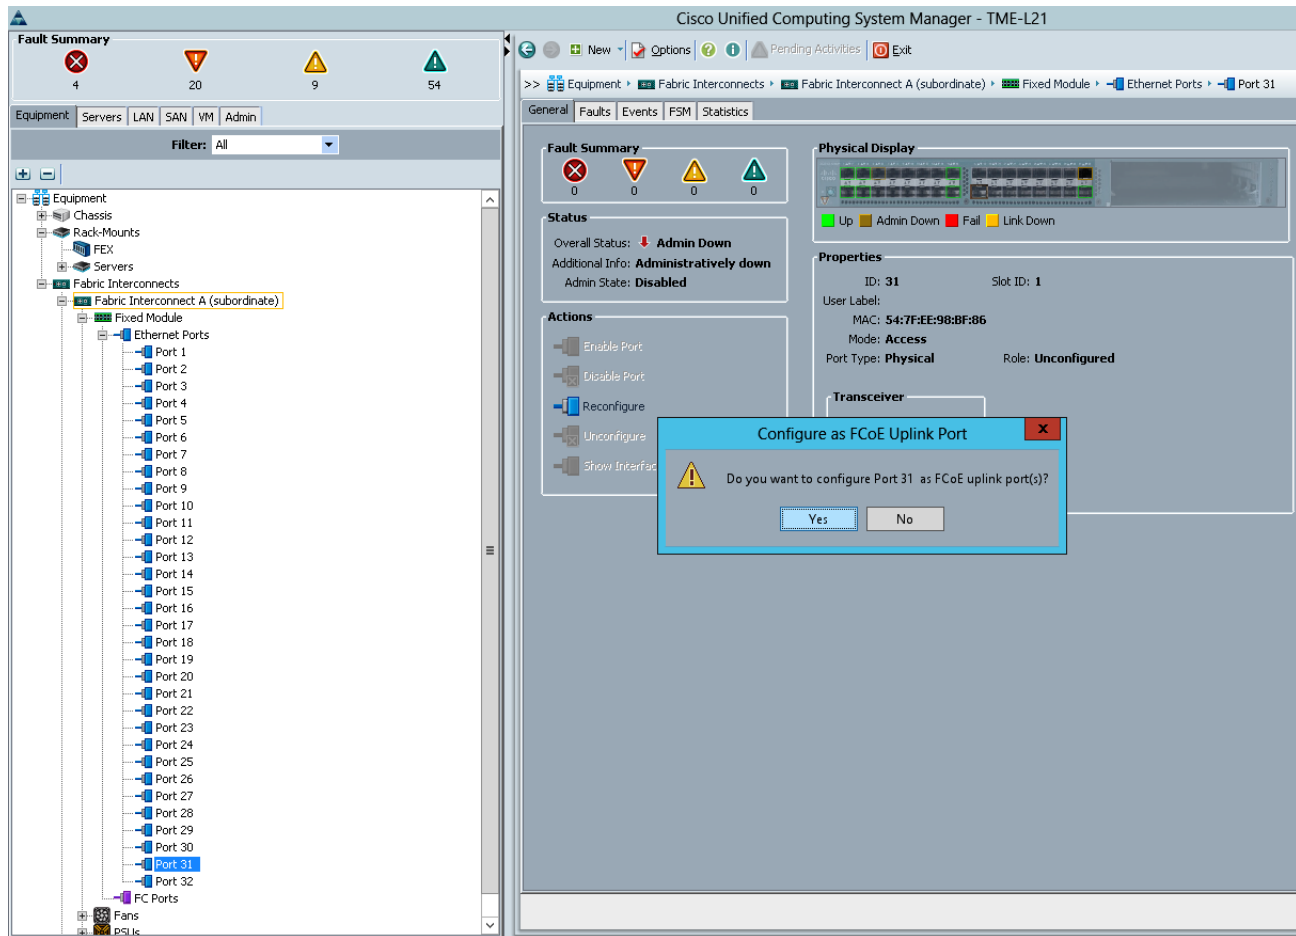
7. Select Configure as FCoE Uplink Port

Figure 34 *Configure as FCoE Uplink port*



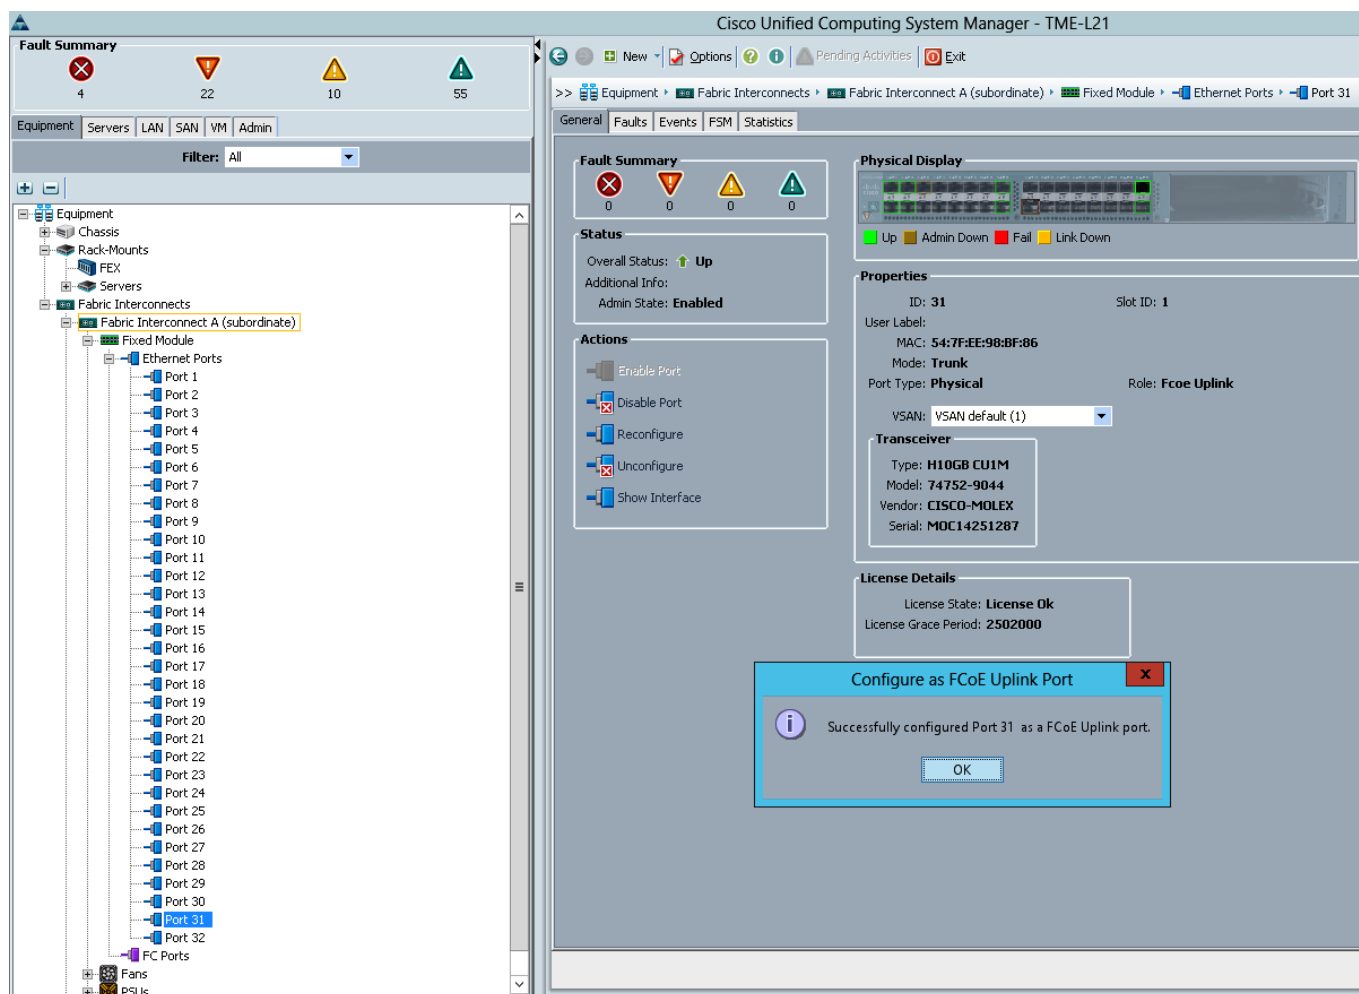
8. Click **Yes** to confirm.

Figure 35 *Configure FCoE Uplink port*



9. Click **OK**.
10. The Cisco UCS Manager GUI will close as the primary Fabric Interconnects reboot.
11. Click **OK**.

Figure 36 *Configure FCoE Uplink port*



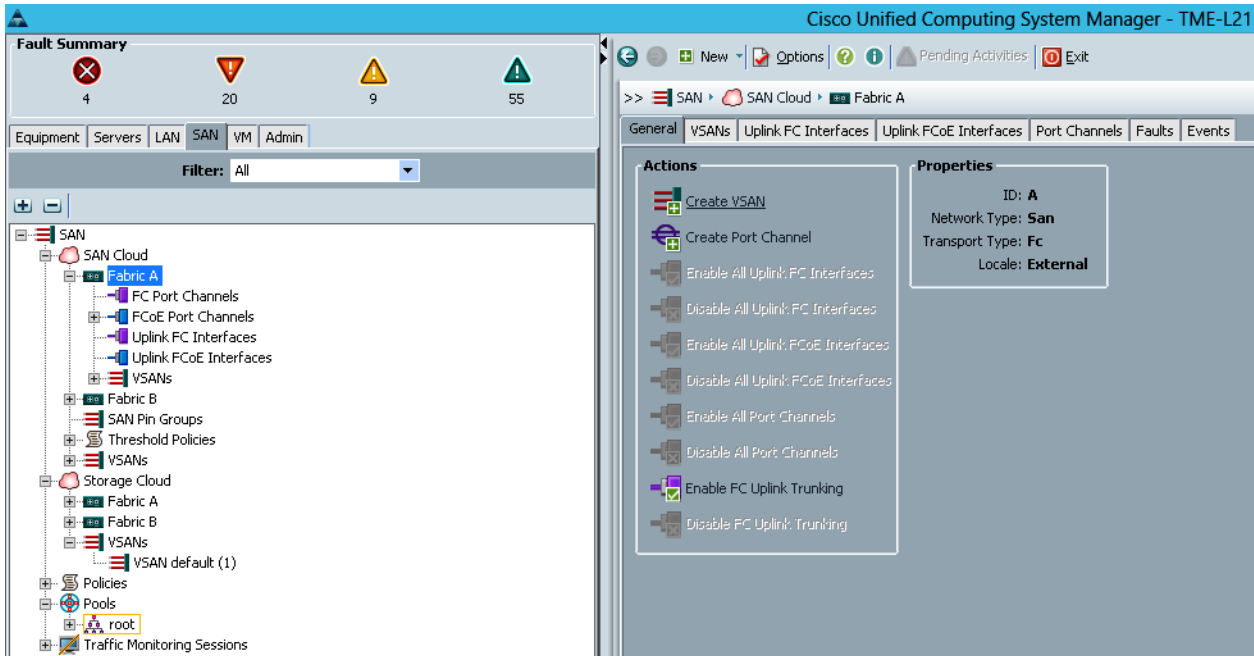
12. Repeat the same steps on Port 32 to configure FCoE Ports on FI A.
13. Repeat the same steps on Port 31 and 32 to configure FCoE Ports on FI Creating VSAN.

VSANs help you create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs. Any application-specific parameters can be configured for a VSAN before creating the VSAN. To create the VSAN, login to the Cisco UCS Manager, and follow these steps:

Fabric Interconnects A

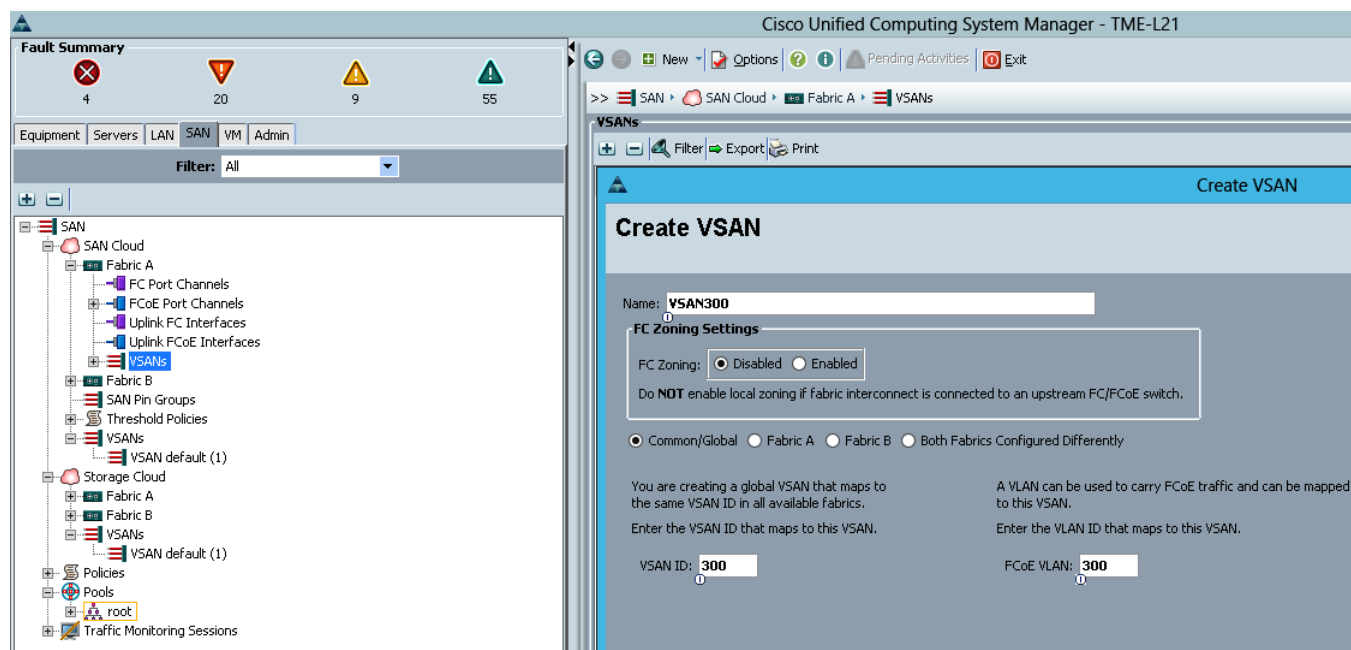
1. Click the **SAN** tab in the left pane.
2. Expand the SAN Cloud tree.
3. Click **Fabric A**.
4. On right pane, under **General** tab click on **Create VSAN**.

Figure 37 **Create VSAN**



5. Enter VSAN300 in the VSAN name field.
6. Keep the Disabled option selected for the Default Zoning.
7. Click **Common/Global** radio button.
8. Enter the VSAN ID 300.
9. Enter the FCoE VLAN ID 300.
10. Click **OK**.

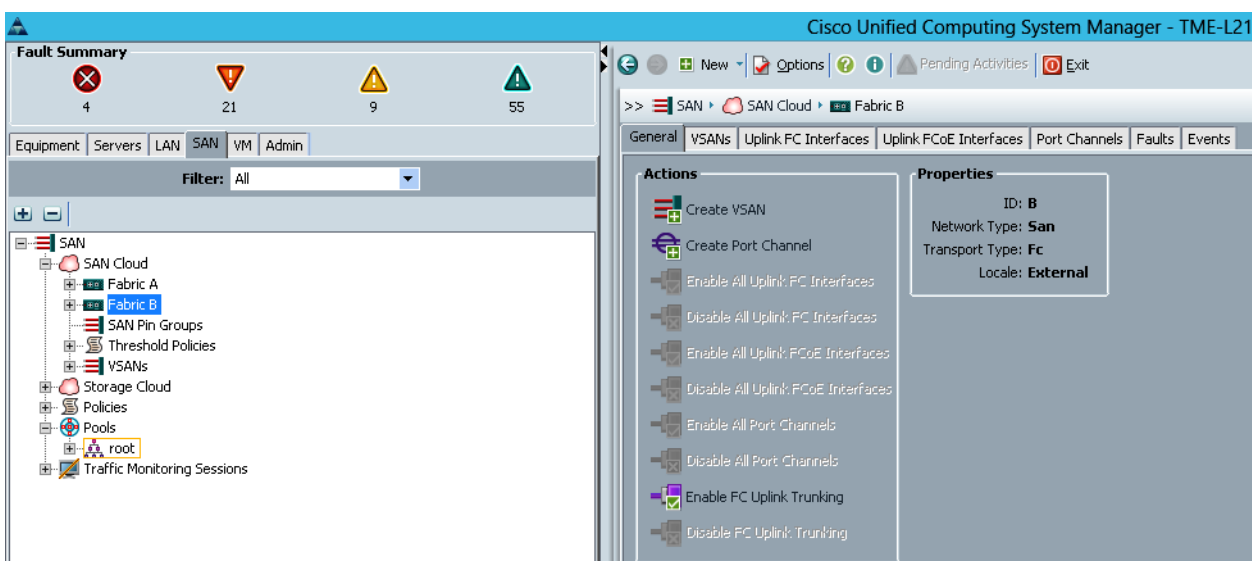
Figure 38 Defining the Common/Global VSAN



Fabric Interconnects B

1. Click the **SAN** tab in the left pane. Expand the SAN Cloud tree.
2. Click **Fabric A**.
3. On Right Pane Under General tab click on Create VSAN

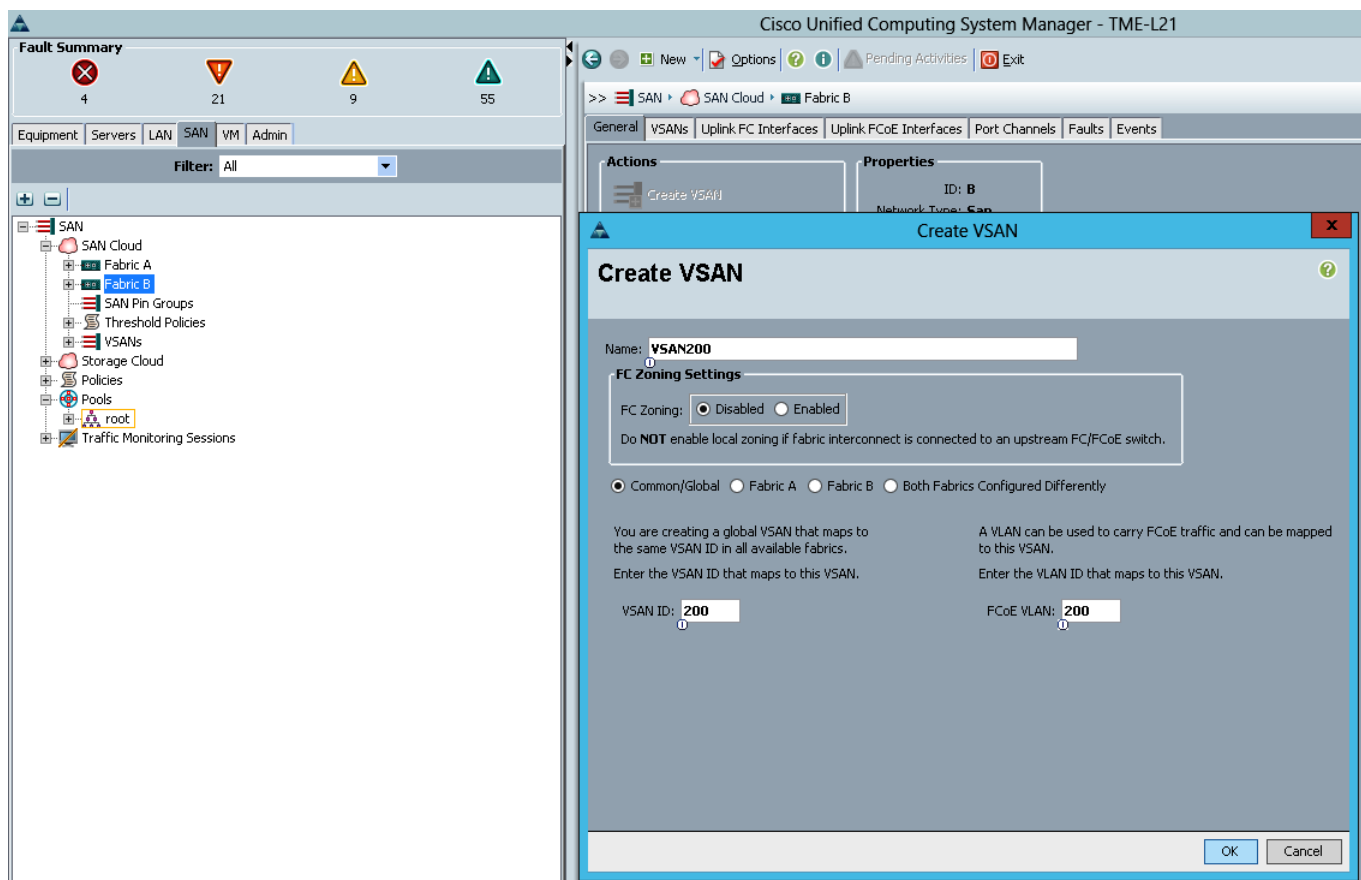
Figure 39 Create vSAN



4. Enter VSAN200 in the VSAN name field.
5. Keep the Disabled option selected for the Default Zoning.

6. Click **Common/Global** radio button.
7. Enter the VSAN ID 200.
8. Enter the FCoE VLAN ID 200.
9. Click **OK**.

Figure 40 *Defining the Common/Global VSAN*

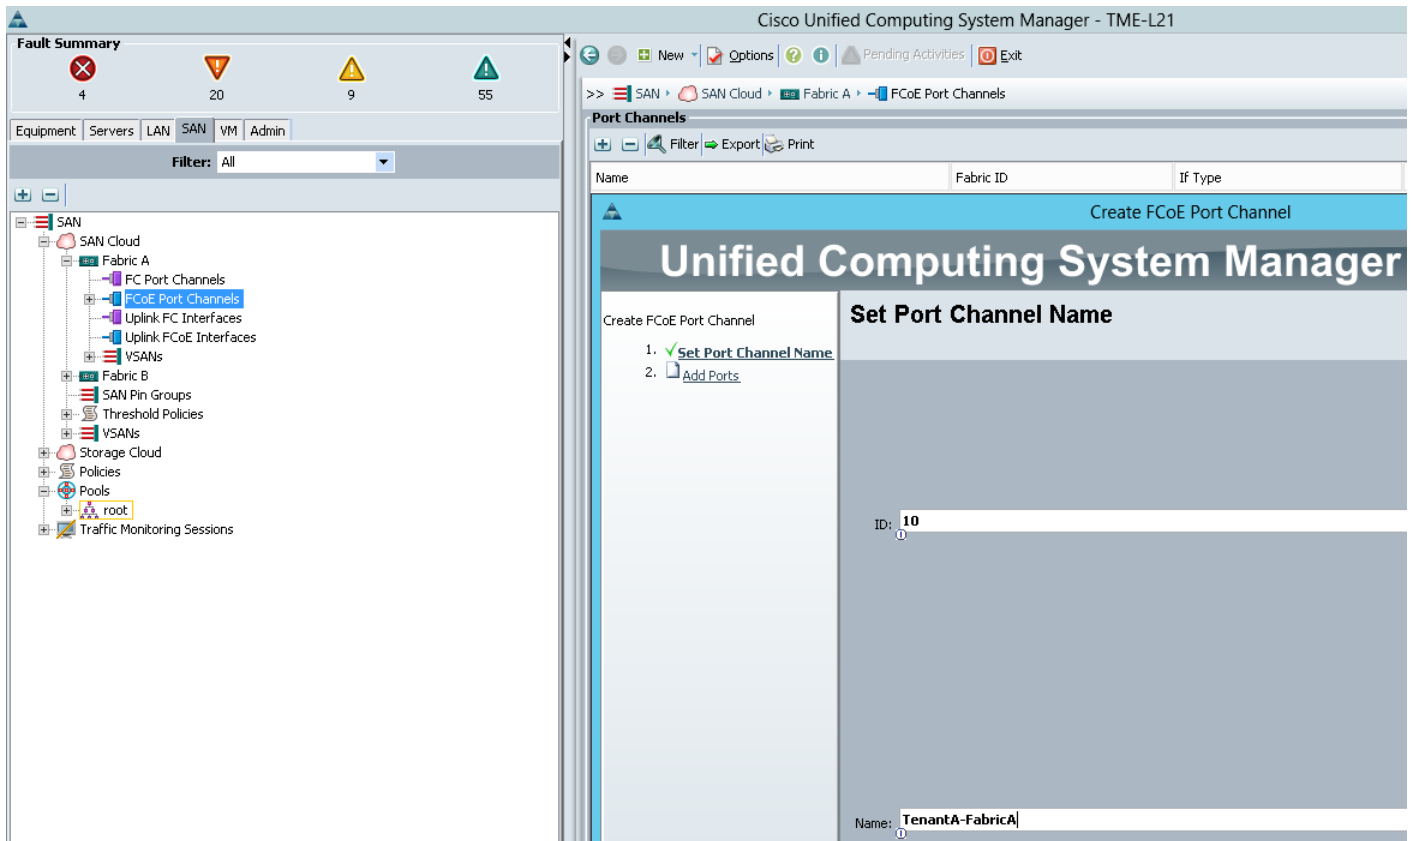


Creating Port Channel on Cisco UCS 6248 Fabric Interconnect

Fabric Interconnects A

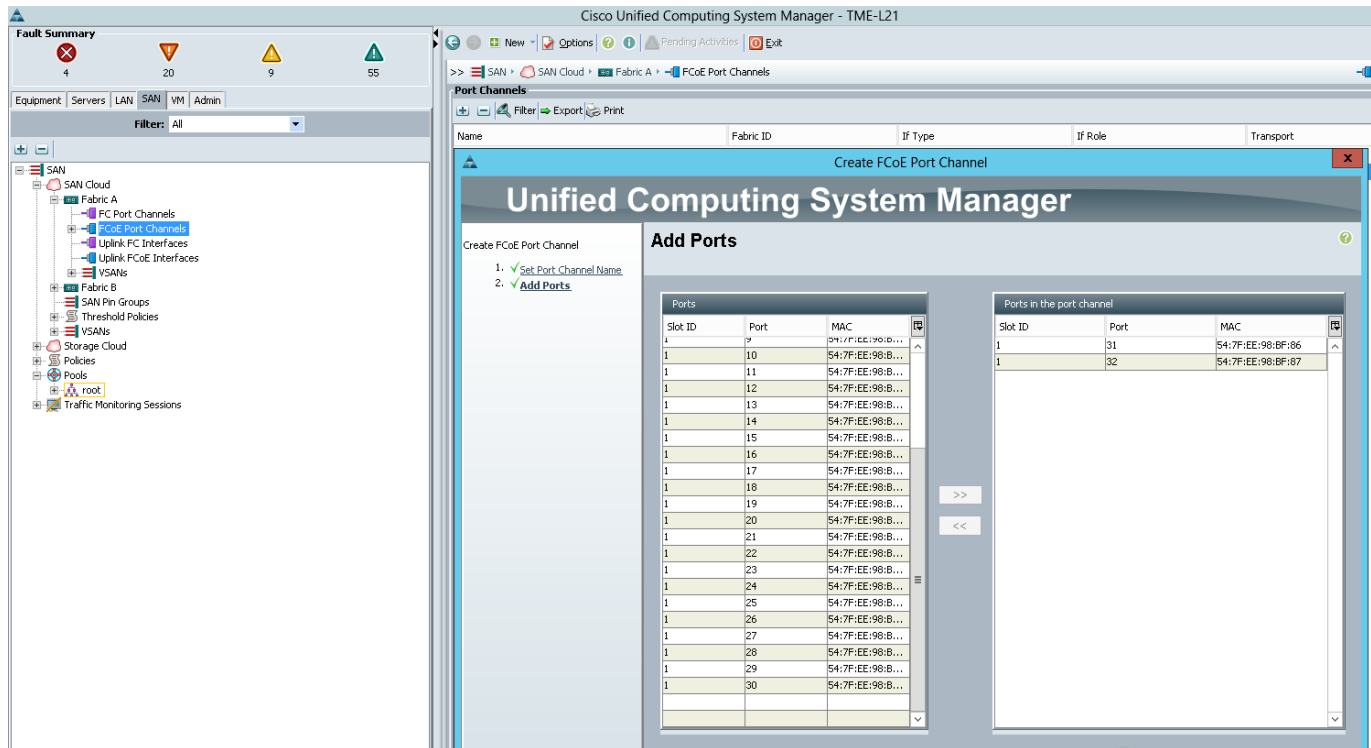
1. Under SAN Cloud, expand the Fabric A tree.
2. Right-click **FCoE Port Channels**.
3. Select **Create Port Channel**.
4. Click **Yes**.
5. Enter 10 in the Port Channel ID field and Tenant-Fabric-A in the Port Channel name field.
6. Click **Next**.

Figure 41 Defining the Name and ID for the FCoE Port Channel



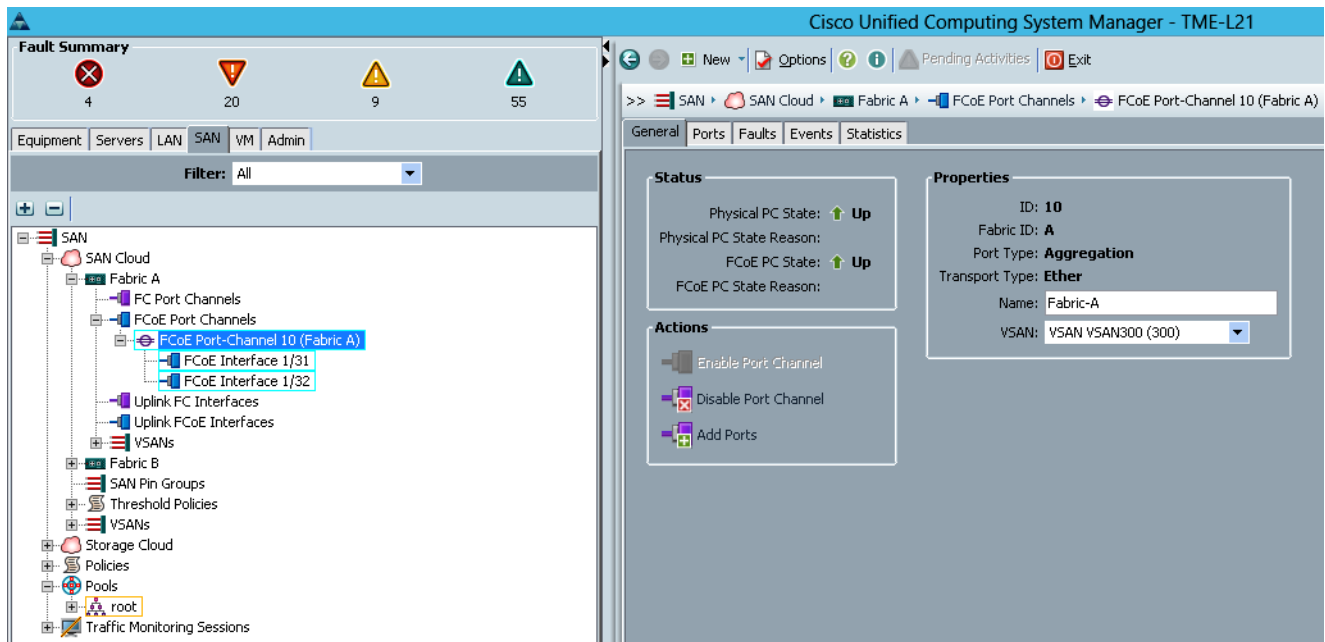
7. Select ports 31 and 32 and click >> to add the ports to the FCoE Port Channel.
8. Click **Finish**.

Figure 42 *Selecting and Adding Ports to the FCoE Port Channel*



9. Click **OK** to complete creating the FCoE Port Channel.
10. In the VSAN pull-down select **VSAN300**.
11. Click **Save Changes** and then click **OK**.

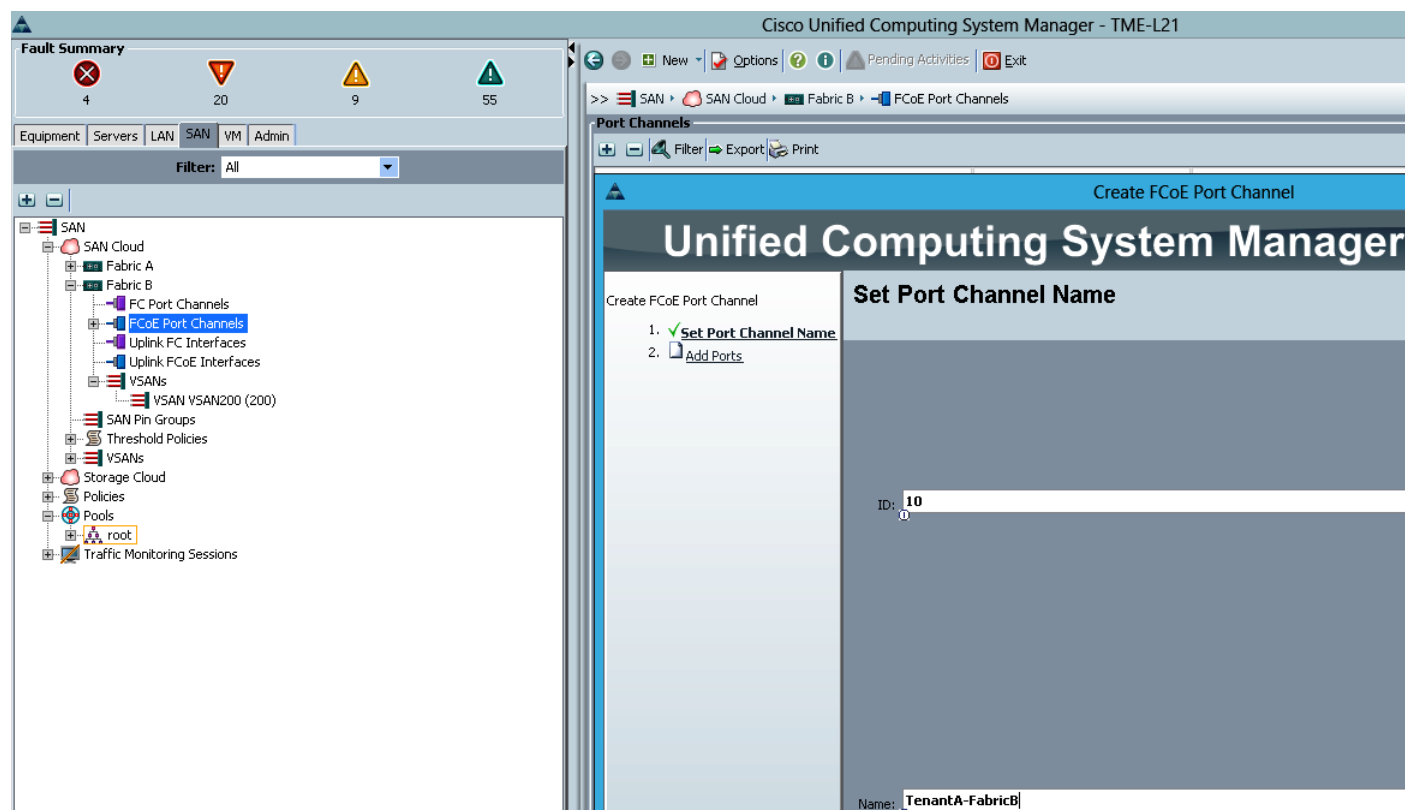
Figure 43 *Selecting VSAN ID for FCoE Port Channel*



Fabric Interconnects B

1. Under SAN Cloud, expand the Fabric B tree.
2. Right-click **FCoE Port Channels**.
3. Select **Create Port Channel**.
4. Click **Yes**, and then enter 10 in the Port Channel ID field and Tenant-Fabric-B in the FCoE Port Channel name field.
5. Click **Next**.

Figure 44 Entering the Name and ID for the FCoE Port Channel

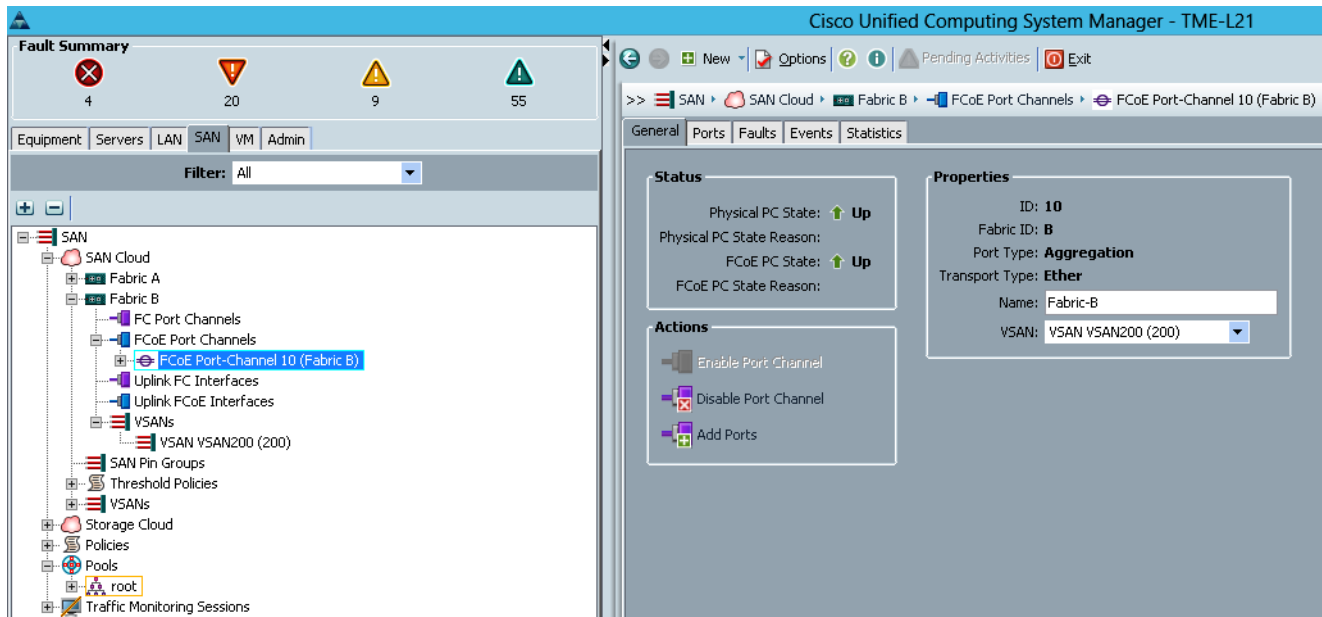


6. Select ports 31 and 32 and click >> to add the ports to the FCoE Port Channel.
7. Click **Finish**.

Figure 45 *Selecting and Adding FCoE Ports to the FCoE Port Channel*

8. Click **OK** to complete creating the FCoE Port Channel.
9. In the VSAN pull-down select VSAN200.
10. Click **Save Changes**.

Figure 46 *FCoE Port Channel Properties*



Creating SAN Port Channels

The same pair of Cisco Nexus 5548UP Switches can be used to for accessing Storage. The FCoE ports on Cisco UCS 6248UP Fabric Interconnects are configured with FCoE SAN Port Channel to carry SAN traffic over virtual SAN (vSAN) between Nexus 5548UP Switches. To provide SAN security Zoning was set up on the Cisco Nexus 5548UP switches to access right storage system's logical unit numbers (LUNs) visible to the infrastructure and test servers.

This section provides details for configuring the necessary VSANs and SAN Port Channels on Cisco UCS and Cisco Nexus network cloud environment.

In this deployment model we will create separate FCoE SAN Port Channel VSAN 10 on Tenant-A-Fabric-A, and Tenant-A-Fabric-B for handling cloud Tenant A storage traffic. In similar way various VSANs can be created for multi-tenant s in the Cloud to provide storage network multi-tendency.

Cisco Nexus 5548UP Switch A

SAN port channels are created with default values. The default configuration can be changed just as any other physical interface. To configure the SAN Port Channel in the global configuration mode, login to the Nexus Switch and, run the following commands.

1. Type feature fport-channel-trunk.
2. Type vsan database.
3. Type vsan 200.
4. Type exit.
5. Type interface port-channel 1.
6. Type switch mode trunk.
7. Type switch trunk allowed vlan 200.
8. Type exit
9. Type interface Ethernet 1/31.
10. Type switchport mode trunk.
11. Type switchport trunk allowed vlan 200.
12. Type channel-group 10 mode active.
13. Type exit.
14. Type interface Ethernet 1/32.
15. Type switchport mode trunk.
16. Type switchport trunk allowed vlan 200.
17. Type channel-group 10 mode active.
18. Type exit.

Cisco Nexus 5548UP Switch B

SAN port channels are created with default values. The default configuration can be changed just as any other physical interface. To configure the SAN Port Channel in the global configuration mode, run the following commands.

1. Type feature fport-channel-trunk.
2. Type vsan database.

3. Type vsan 300.
4. Type exit.
5. Type interface port-channel 10.
6. Type switch mode trunk.
7. Type switch trunk allowed vlan 300.
8. Type exit
9. Type interface Ethernet 1/31.
10. Type switchport mode trunk.
11. Type switchport trunk allowed vlan 300.
12. Type channel-group 10 mode active.
13. Type exit.
14. Type interface Ethernet 1/32.
15. Type switchport mode trunk.
16. Type switchport trunk allowed vlan 300.
17. Type channel-group 10 mode active.
18. Type exit.

NetApp Storage

This section presents design considerations of cloud storage layout for deploying storage infrastructure for building Citrix CloudPlatform 4.2.1. To support various service levels of shared storage infrastructure in a multi-tenant cloud environment, NetApp offers various features such as:

- Unified data storage architecture that supports multiple workloads.
- Seamless, multidimensional scaling that meets the dynamic demands of cloud computing.
- Storage efficiency that helps reduce capacity requirements and costs by 50% or more.
- Secure multi-tenancy segments, that isolates and delivers shared server, storage, and network resources.
- Service automation and analytics automate storage provisioning, comprehensive visibility, and monitoring.
- Nondisruptive, continuous operations that enable nonstop data availability for shared cloud storage resources.
- Integrated data protection that helps meet backup, disaster recovery, archiving, compliance, and security service-level agreements.
- Virtual storage tiering that automates data movement based on application affinity and workload.
- Embedded data security that protects data assets through role-based administration, encryption, and antivirus.

The NetApp aggregation layer provides a large virtualized pool of storage capacity and disk IOPS to be used on demand by Citrix CloudPlatform 4.2.1. The aggregation-layer sizing is based on the storage requirements for hosting Citrix CloudPlatform 4.2.1 to store tenant data, to meet the storage capacity, performance, and snapshot copy backup requirements of an assumed workload. When sizing for your environment, perform the necessary planning to determine the exact storage configuration for your

individual requirements. Aggregation layer 0 (Aggr0) is defined for hosting root NetApp flexible volumes (FlexVol® volumes) that use the NetApp clustered Data ONTAP operating system for handling NetApp storage configurations. For details on NetApp storage command options, see <http://now.netapp.com/NOW/public/knowledge/docs/ontap/rel732/pdfs/ontap/210-04499.pdf>.

**Note**

In this design, we will create aggregates, flexible volumes, FC and iSCSI LUNs, igroups, and NFS mount points for TenantA hosts to provision primary and secondary storage for hosting virtual machines in a Citrix cloud environment.

Storage Configuration

Controller FAS32xx Series

Table 7 **Controller FAS32XX series prerequisites**

Requirement	Reference	Comments
Physical site where storage system needs to be installed must be ready	Site Requirements Guide	Refer to the “Site Preparation” section.
Storage system connectivity requirements	Site Requirements Guide	Refer to the “System Connectivity Requirements” section.
Storage system general power requirements	Site Requirements Guide	Refer to the “Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements” section.
Storage system model-specific requirements	Site Requirements Guide	Refer to the “FAS32xx/V32xx Series Systems” section.

System Configuration Guides

System configuration guides provide supported hardware and software components for the specific clustered Data ONTAP version. These online guides provide configuration information for all NetApp storage appliances currently supported by the clustered Data ONTAP software. They also provide a table of component compatibilities. The documentation below applies to the clustered Data ONTAP 8.1.2 software that was deployed in this study.

1. Make sure that the hardware and software components are supported with the version of Data ONTAP that you plan to install by checking the [System Configuration Guides](#) at the [NetApp Support](#) site.
2. In the System Configuration Guide, click the appropriate NetApp storage appliance and then click the component you want to view. Alternatively, to compare components by storage appliance, click a component and then click the NetApp storage appliance you want to view.

Controllers

Follow the physical installation procedures for the controllers in the [FAS32xx documentation](#) at the [NetApp Support](#) site.

Disk Shelves DS2246 Series

DS2246 Disk Shelves

Follow the procedures in the [Disk Shelf Installation and Setup section of the DS2246 Disk Shelf Overview](#) to install a disk shelf for a new storage system.

Follow procedures for proper cabling with the controller model as described in [SAS Disk Shelves Universal SAS and ACP Cabling Guide](#).

The following information applies to DS2246 disk shelves:

- SAS disk drives use software-based disk ownership. Ownership of a disk drive is assigned to a specific storage system by writing software ownership information on the disk drive rather than by using the topography of the storage system's physical connections.
- Connectivity terms used: shelf-to-shelf (daisy-chain), controller-to-shelf (top connections), and shelf-to controller (bottom connections).
- Unique disk shelf IDs must be set per storage system (a number from 0 through 98).
- Disk shelf power must be turned on to change the digital display shelf ID. The digital display is on the front of the disk shelf.
- Disk shelves must be power cycled after the shelf ID is changed for it to take effect.
- Changing the shelf ID on a disk shelf that is part of an existing storage system running Data ONTAP requires that you wait at least 30 seconds before turning the power back on so that Data ONTAP can properly delete the old disk shelf address and update the copy of the new disk shelf address.
- Changing the shelf ID on a disk shelf that is part of a new storage system installation (the disk shelf is not yet running Data ONTAP) requires no wait; you can immediately power cycle the disk shelf.

Cisco Nexus 5596 Cluster Network Switch Configuration

Cisco Nexus 5596 cluster network switch configuration prerequisites.:

- Rack and connect power to the new Cisco Nexus 5596 switches.
- Provide a terminal session that connects to the switch's serial console port (9600, 8, n, 1).
- Connect the mgmt0 port to the management network and be prepared to provide IP address information.
- Obtain password for admin.
- Determine switch name.
- Identify SSH key type (dsa, rsa, or rsa1)
- Set up an e-mail server for Cisco Smart Call Home and IP connectivity between the switch and the e-mail server.
- Provide SNMP contact information for Cisco Smart Call Home (name, phone, street address).

- Identify a CCO ID associated with an appropriate Cisco SMARTnet® Service contract for Cisco Smart Call Home.

Initial Setup of Cisco Nexus 5596 Cluster Interconnect

The first time a Cisco Nexus 5596 cluster interconnect is accessed, it runs a setup program that prompts the user to enter an IP address and other configuration information needed for the switch to communicate over the management Ethernet interface. This information is required to configure and manage the switch. If the configuration must be changed later, the setup wizard can be accessed again by running the `setup` command in EXEC mode.

To set up the Cisco Nexus 5596 cluster interconnect, complete the following steps. These steps will need to be completed on both cluster interconnects.

1. Enter appropriate responses to the setup prompts displayed on the Cisco Nexus 5596 cluster interconnect.

```
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <switchname>
Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <ic_mgmt0_ip>
Mgmt0 IPv4 netmask: <ic_mgmt0_netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <ic_mgmt0_gw>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa): rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <ntp_server_ip>
Enter basic FC configurations (yes/no) [n]: Enter
```

2. At the end of the setup, the configuration choices are displayed. Verify the information and save the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: <n>
Use this configuration and save it? (yes/no) [y]: <y>
```

Download and Install NetApp Cluster Switch Software

When Cisco Nexus 5596 is being used as a cluster network switch with Data ONTAP 8.1.2, it should be running NX-OS version 5.2(1)N1(1). The `show version` command from the switch command line interface will show the switch version currently running on the switch. If the currently running version is not 5.2(1)N1(1), go to the [NetApp Support](#) site and download and install NX-OS 5.2(1)N1(1) for the Cisco Nexus 5596 switch. Make sure both cluster interconnects are running NX-OS version 5.2(1)N1(1).

Download and Merge of NetApp Cluster Switch Reference Configuration File

Cluster network and management network switches are shipped without the configuration files installed. These files must be downloaded to the switches during deployment. Configuration files must be downloaded when the cluster network and management network switches are first installed or after the Cisco switch software is updated or reinstalled.

After the initial setup is complete, the NetApp cluster network switch reference configuration must be transferred to the switch and merged with the existing configuration. Instructions for this task and the reference configuration files for the appropriate switches are available on the [NetApp Support](#) site.

To download configuration files to a host and install them on a Cisco Nexus 5596 switch, complete the following steps on both cluster interconnects:

1. Obtain a console connection to the switch. Verify the existing configuration on the switch by running the `show run` command.
2. Log in to the switch. Make sure that the host recognizes the switch on the network (for example, use the ping utility).
3. Enter the following command:

```
copy <transfer protocol>: bootflash: vrf management
```

4. Verify that the configuration file is downloaded.
5. Merge the configuration file into the existing running-config. Run the following command in which `<config file name>` is the file name for the switch type. A series of warnings regarding PortFast is displayed as each port is configured.

```
copy <config file name> running-config
```

6. Verify the success of the configuration merge by running the `show run` command and comparing its output to the contents of the configuration file (a .txt file) that was downloaded.
7. The output for both installed-base switches and new switches should be identical to the contents of the configuration file for the following items:
 - Banner (should match the expected version)
 - Switch port descriptions such as description Cluster Node x
 - The new ISL algorithm port-channel load-balance Ethernet source-dest-port
8. The output for new switches should be identical to the contents of the configuration file for the following items:
 - Port channel
 - Policy map
 - System QoS
 - Interface
 - Boot
9. The output for installed-base switches should have the flow control receive and send values on for the following items:
 - Interface port-channel 1 and 2
10. Ethernet interface 1/41 through Ethernet interface 1/48.
11. Copy the running-config to the startup-config.

```
copy running-config startup-config
```

Cisco Smart Call Home Setup

To configure Smart Call Home on a Cisco Nexus 5596 switch, follow these steps:

1. Enter the mandatory system contact using the `snmp-server contact` command in global configuration mode. Then run the `callhome` command to enter callhome configuration mode.

```
NX-5596#config t
```

- ```
NX-5596(config)#snmp-server contact <sys-contact>
NX-5596(config)#callhome
```
2. Configure the mandatory contact information (phone number, e-mail address, and street address).
 

```
NX-5596(config-callhome)#email-contact <email-address>
NX-5596(config-callhome)#phone-contact <+1-000-000-0000>
NX-5596(config-callhome)#streetaddress <a-street-address>
```
  3. Configure the mandatory e-mail server information. The server address is an IPv4 address, IPv6 address, or the domain-name of a SMTP server to which Call Home will send e-mail messages. Optional port number (default=25) and VRF may be configured.
 

```
NX-5596(config-callhome)#transport email smtp-server <ip-address> port 25 use-vrf <vrf-name>
```
  4. Set the destination profile CiscoTAC-1 e-mail address to [callhome@cisco.com](mailto:callhome@cisco.com).
 

```
NX-5596(config-callhome)#destination-profile CiscoTAC-1 email-addr callhome@cisco.com
vrf management
```
  5. Enable periodic inventory and set the interval.
 

```
NX-5596(config-callhome)#periodic-inventory notification
NX-5596(config-callhome)#periodic-inventory notification interval 30
```
  6. Enable callhome, exit, and save the configuration.
 

```
NX-5596(config-callhome)#enable
NX-5596(config-callhome)#end
NX-5596#copy running-config startup-config
```
  7. Send a callhome inventory message to start the registration process.
 

```
NX-5596#callhome test inventory
trying to send test callhome inventory message
successfully sent test callhome inventory message
```
  8. Watch for an e-mail from Cisco regarding the registration of the switch. Follow the instructions in the e-mail to complete the registration for Smart Call Home.

## SNMP Monitoring Setup

Configure SNMP by using the following example as a guideline. This example configures a host receiver for SNMPv1 traps and enables all link up/down traps.

```
NX-5596(config)# snmp-server host <ip-address> traps { version 1 } <community>
[udp_port <number>]
NX-5596(config)# snmp-server enable traps link
```

## Clustered Data ONTAP 8.1.2

### Node 1

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:
 

```
Starting AUTOBOOT press Ctrl-C to abort
```
2. From the Loader-A prompt:
 

```
printenv
```
3. If the last-OS-booted-ver parameter is not set to 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.
4. Allow the system to boot up.
 

```
boot_ontap
```
5. Press Ctrl-C when the Press Ctrl-C for Boot Menu message appears.


**Note**

If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 is the version being booted, then select option 8 and **yes** to reboot the node. Then proceed with step 5.

6. To install new software, first select option 7.  
7
7. Answer yes to perform a nondisruptive upgrade.  
y
8. Select e0M for the network port you want to use for the download.  
e0M
9. Select yes to reboot now.  
y
10. Enter the IP address, netmask, and default gateway for e0M in their respective places.  
<<var\_node01\_mgmt\_ip>> <<var\_node01\_mgmt\_mask>> <<var\_node01\_mgmt\_gateway>>
11. Enter the URL where the software can be found.


**Note**

This Web server must be pingable.

12. Press Enter for the user name, indicating no user name.  
Enter
13. Enter yes to set the newly installed software as the default to be used for subsequent reboots.  
y
14. Enter yes to reboot the node.  
y


**Note**

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the LOADER prompt. If these actions occur, the system might deviate from this procedure.

15. Press Ctrl-C to exit autoboot when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

16. From the LOADER-A prompt, enter:

```
printenv
```


**Note**

If bootarg.init.boot\_clustered true is not listed, the system is not set to boot in clustered Data ONTAP.

17. If the system is not set to boot in clustered Data ONTAP, at the LOADER prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. When you see Press Ctrl-C for Boot Menu:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

4  
21. Answer yes to Zero disks, reset config and install a new file system.

y  
22. Enter yes to erase all the data on the disks.

y

**Note**

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue to node 02 configuration while the disks for node 01 are zeroing.

## Node 2

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

2. From the Loader-A prompt, enter:

printenv

3. If the last-OS-booted-ver parameter is not set to 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.

4. Allow the system to boot up.

boot\_ontap

5. Press Ctrl-C when Press Ctrl-C for Boot Menu is displayed.

Ctrl-C

**Note**

If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 is the version being booted, then select option 8 and yes to reboot the node. Then proceed with step 15.

6. To install new software first select option 7.

7

7. Answer yes to perform a nondisruptive upgrade.

y

8. Select e0M for the network port you want to use for the download.

e0M

9. Select yes to reboot now.

y

10. Enter the IP address, netmask, and default gateway for e0M in their respective places.

<<var\_node02\_mgmt\_ip>> <<var\_node02\_mgmt\_mask>> <<var\_node02\_mgmt\_gateway>>

11. Enter the URL where the software can be found.

**Note**

This Web server must be reachable.

<<var\_url\_boot\_software>>

12. Press Enter for the user name, indicating no user name.

Enter

13. Select yes to set the newly installed software as the default to be used for subsequent reboots.

y

14. Select yes to reboot the node.

y



**Note**

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the LOADER prompt. If these actions occur, the system might deviate from this procedure.

15. Press Ctrl-C to exit autoboot when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

16. From the LOADER-A prompt, enter:

printenv



**Note**

If bootarg.init.boot\_clustered true is not listed, the system is not set to boot in clustered Data ONTAP.

17. If the system is not set to boot in clustered Data ONTAP, at the LOADER prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

autoboot

19. Press Ctrl-C for boot menu when prompted:

Ctrl - C

20. Select option 4 for clean configuration and initialize all disks.

4

21. Answer yes to Zero disks, reset config and install a new file system.

y

22. Enter yes to erase all the data on the disks.

y



**Note**

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

## Cluster Create in Clustered Data ONTAP

**Table 8** Cluster create in clustered Data ONTAP prerequisites

| Cluster Detail                    | Cluster Detail Value             |
|-----------------------------------|----------------------------------|
| Cluster name                      | <<var_clustername>>              |
| Clustered Data ONTAP base license | <<var_cluster_base_license_key>> |
| Cluster management IP address     | <<var_clustermgmt_ip>>           |
| Cluster management netmask        | <<var_clustermgmt_mask>>         |
| Cluster management port           | <<var_clustermgmt_port>>         |
| Cluster management gateway        | <<var_clustermgmt_gateway>>      |



**Table 8 Cluster create in clustered Data ONTAP prerequisites**

|                           |                             |
|---------------------------|-----------------------------|
| Cluster Node01 IP address | <<var_node01_mgmt_ip>>      |
| Cluster Node01 netmask    | <<var_node01_mgmt_mask>>    |
| Cluster Node01 gateway    | <<var_node01_mgmt_gateway>> |

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered Node01.

1. During the first node boot, the Cluster Setup wizard starts running on the console.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

**Note**

If a login prompt appears instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings and then enter the cluster setup command.

2. Enter the following command to create a new cluster:

```
create
```

3. The system defaults are displayed.

```
System Defaults:
Private cluster network ports [e1a,e2a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.

**Note**

The cluster is created; this can take a minute or two.

5. The steps to create a cluster are displayed.

```
Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter additional license key[]:
```

**Note**

For this validated architecture we recommend you install license keys for SnapRestore, NFS, FCP, FlexClone, and SnapManager Suite. After you finish entering the license keys, press Enter.

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster management interface port [e0a]: e0a
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

```
Enter the DNS domain names:<<var_dns_domain_name>>
Enter the name server IP addresses:<<var_nameserver_ip>>
```

**Note**

If you have more than one name server IP address, separate them with a comma.

### 7. Set up the node.

```
Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0b
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
enter the node management interface netmask:<<var_node01_mgmt_mask>>
Enter the node management interface default gateway:<<var_node01_mgmt_gateway>>
```

**Note**

The node management interface should be in a different subnet than the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

### 8. Press Enter to accept the AutoSupport message.

### 9. Reboot node 01.

```
system node reboot <<var_node01>>
y
```

### 10. Press Ctrl-C for boot menu when prompted:

```
Ctrl - C
```

### 11. Select 5 to boot into maintenance mode.

```
5
```

### 12. When prompted, Continue with boot?, enter y.

### 13. To verify the HA status of your environment, run the following command:

```
ha-config show
```

**Note**

If either component is not in HA mode, use the ha-config modify command to put the components in HA mode.

### 14. To see how many disks are unowned, enter:

```
disk show -a
```

**Note**

No disks should be owned in this list.

### 15. Assign disks.

**Note**

This reference architecture allocates half the disks to each controller. However, workload design could dictate different percentages.

```
disk assign -n <<var_#_of_disks>>
```

### 16. Reboot the controller.

```
halt
```

### 17. At the LOADER-A prompt, enter:

```
autoboot
```

## Cluster Join in Clustered Data ONTAP

**Table 9** Cluster join in clustered Data ONTAP prerequisites

| Cluster Detail                | Cluster Detail Value        |
|-------------------------------|-----------------------------|
| Cluster name                  | <<var_clustername>>         |
| Cluster management IP address | <<var_clustermgmt_ip>>      |
| Cluster Node02 IP address     | <<var_node02_mgmt_ip>>      |
| Cluster Node02 netmask        | <<var_node02_mgmt_mask>>    |
| Cluster Node02 gateway        | <<var_node02_mgmt_gateway>> |

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered Node01, and the node joining the cluster in this example is Node02.

1. During the node boot, the Cluster Setup wizard starts running on the console.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```



### Note

If a login prompt displays instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings, and then enter the cluster setup command.

2. Enter the following command to join a cluster:

```
join
```

3. The system defaults are displayed.

```
System Defaults:
Private cluster network ports [e1a,e2a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.



### Note

The cluster creation can take a minute or two.

5. The steps to create a cluster are displayed.

```
Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
```



### Note

The node should find the cluster name.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0b
Enter the node management interface IP address: <<var_node02_mgmt_ip>>
Enter the node management interface netmask: Enter
Enter the node management interface default gateway: Enter
```

7. The node management interface should be in a subnet different from the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

8. Press Enter to accept the AutoSupport message.

9. Log in to the Cluster Interface with the admin user id and <<var\_password>>.

10. Reboot node 02.

```
system node reboot <<var_node02>>
y
```

11. Press Ctrl-C for boot menu when prompted:

```
Ctrl - C
```

12. Select 5 to boot into maintenance mode.

```
5
```

13. When prompted, Continue with boot? enter:

```
y
```

14. To verify the HA status of your environment, enter:

```
ha-config show
```



#### Note

If either component is not in HA mode, use the `ha-config modify` command to put the components in HA mode.

15. To see how many disks are unowned, enter:

```
disk show -a
```

16. Assign disks.



#### Note

This reference architecture allocates half the disks to each controller. However, the workload design could dictate different percentages. Assign all remaining disks to node 02.

```
disk assign -n <<var_#_of_disks>>
```

17. Reboot the controller:

```
halt
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. Press Ctrl-C for boot menu when prompted:

```
Ctrl-C
```

## Log in to the Cluster

Open an SSH connection to cluster IP or host name and log in to the admin user with the password that was provided earlier.

## Zero All Spare Disks

Zero all spare disks in the cluster.

```
disk zerospares
```

## Set Auto-Revert on Cluster Management

To set the auto-revert parameter on the cluster management interface, enter:

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt -auto-revert true
```

## Failover Groups Management in Clustered Data ONTAP

Create a management port failover group.

```
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node01>> -port e0a
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node02>> -port e0a
```

## Assign Management Failover Group to Cluster Management LIF

Assign the management port failover group to the cluster management LIF.

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt -failover-group fg-cluster-mgmt
```

## Failover Groups Node Management in Clustered Data ONTAP

Create a management port failover group.

```
network interface failover-groups create -failover-group fg-node-mgmt-01 -node <<var_node01>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-01 -node <<var_node01>> -port e0M
network interface failover-groups create -failover-group fg-node-mgmt-02 -node <<var_node02>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-02 -node <<var_node02>> -port e0M
```

## Assign Node Management Failover Groups to Node Management LIFs

Assign the management port failover group to the cluster management LIF.

```
network interface modify -vserver <<var_node01>> -lif mgmt1 -auto-revert true -use-failover-group enabled -failover-group fg-node-mgmt-01
network interface modify -vserver <<var_node02>> -lif mgmt1 -auto-revert true -use-failover-group enabled -failover-group fg-node-mgmt-02
```

## Flash Cache in Clustered Data ONTAP

Complete the following steps to enable Flash Cache on each node.

Run the following commands from the cluster management interface:

```
system node run -node <<var_node01>> options flexscale.enable on
system node run -node <<var_node01>> options flexscale.lopri_blocks off
system node run -node <<var_node01>> options flexscale.normal_data_blocks on
system node run -node <<var_node02>> options flexscale.enable on
system node run -node <<var_node02>> options flexscale.lopri_blocks off
system node run -node <<var_node02>> options flexscale.normal_data_blocks on
```



### Note

Data ONTAP 8.1 and later does not require a separate license for Flash Cache.



### Note

For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, refer to [TR-3832: Flash Cache Best Practices Guide](#). Before customizing the settings, determine whether the custom settings are required or if the default settings are sufficient.

## 64-Bit Aggregates in Clustered Data ONTAP

A 64-bit aggregate containing the root volume is created during the Data ONTAP setup process. To create additional 64-bit aggregates, determine the aggregate name, the node on which to create it, and the number of disks it will contain.

1. Execute the following command to create new aggregates:

```
aggr create -aggregate aggr01 -nodes <<var_node01>> -B 64 -s <<var_raidsize>>
-diskcount <<var_num_disks>>
aggr create -aggregate aggr02 -nodes <<var_node02>> -B 64 -s <<var_raidsize>>
-diskcount <<var_num_disks>>
```



### Note

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.



### Note

Calculate the RAID group size to allow for roughly balanced (same size) RAID groups of from 12 through 20 disks (for SAS disks) within the aggregate. For example, if 52 disks were being assigned to the aggregate, select a RAID group size of 18. A RAID group size of 18 would yield two 18-disk RAID groups and one 16-disk RAID group. Keep in mind that the default RAID group size is 16 disks, and that the larger the RAID group size, the longer the disk rebuild time in case of a failure.



### Note

The aggregate cannot be created until disk zeroing completes. Use the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr01` and `aggr02` are online.

1. Disable Snapshot copies for the two data aggregates just created.

```
node run <<var_node01>> aggr options aggr01 nosnap on
node run <<var_node02>> aggr options aggr02 nosnap on
```

2. Delete any existing Snapshot copies for the two data aggregates.

```
node run <<var_node01>> snap delete -A -a -f aggr01
```

- ```
node run <<var_node02>> snap delete -A -a -f aggr02
```
3. Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02.
- ```
aggr show
aggr rename -aggregate aggr0 -newname <<var_node01_rootaggrname>>
```

## Service Processor

Gather information about the network and the AutoSupport settings before configuring the Service Processor (SP).

Configure the SP using DHCP or static addressing. If the SP uses a static IP address, verify that the following SP prerequisites have been met:

- An available static IP address
- The network netmask
- The network gateway IP
- AutoSupport information

A best practice is to configure the AutoSupport recipients and mail host before configuring the SP. Data ONTAP automatically sends AutoSupport configuration to the SP, allowing the SP to send alerts and notifications through an AutoSupport message to the system administrative recipients specified in AutoSupport. When configuring the SP, enter the name or the IP address of the AutoSupport mail host, when prompted.

A service processor needs to be set up on each node.

## Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.1.2, you must upgrade to the latest service processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Using a web browser, connect to <http://support.netapp.com/NOW/cgi-bin/fw>.
2. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
3. Proceed to the download page for the latest release of the SP firmware for your storage platform.
4. Using the instructions on this page, update the SPs on both nodes in your cluster. You will need to download the .zip file to a web server that is reachable from the cluster management interface. In step 1a of the instructions substitute the following command then execute steps 2–6 on each node.

```
system image get -node * -package http://web_server_name/path/SP_FW.zip. Also, instead of run local, use system node run <<var_nodename>>,
```

## Configure the Service Processor on Node 01

1. From the cluster shell, enter the following command:

```
system node run <<var_node01>> sp setup
```

2. Enter the following to set up the SP:

```
Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node01_sp_ip>>
Please enter the netmask of the SP[]: <<var_node01_sp_mask>>
```

Please enter the IP address for the SP gateway[]: <<var\_node01\_sp\_gateway>>

## Configure the Service Processor on Node 02

1. From the cluster shell, enter the following command:

```
system node run <<var_node02>> sp setup
```

2. Enter the following to set up the SP:

```
Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node02_sp_ip>>
Please enter the netmask of the SP[]: <<var_node02_sp_mask>>
Please enter the IP address for the SP gateway[]: <<var_node02_sp_gateway>>
```

## Storage Failover in Clustered Data ONTAP

Run the following commands in a failover pair to enable storage failover.

1. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_node01>> -enabled true
```



### Note

Enabling failover on one node enables it for both nodes.

2. Enable HA mode for two-node clusters only.

3. Do not run this command for clusters with more than two nodes because it will cause problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

4. Verify that hardware assist is correctly configured and if needed modify the partner IP address.

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <<var_node02_mgmt_ip>> -node
<<var_node01>>
storage failover modify -hwassist-partner-ip <<var_node01_mgmt_ip>> -node
<<var_node02>>
```

## IFGRP LACP in Clustered Data ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Therefore, make sure that the switch is configured properly.

Run the following commands on the command line to create interface groups (ifgrps).

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e4a
ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e4a
```



### Note

All interfaces must be in the down status before being added to an interface group.



**Note**

The interface group name must follow the standard naming convention of a0x.

## VLAN in Clustered Data ONTAP

### Create NFS VLANs

Run the following command:

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>
```

## Jumbo Frames in Clustered Data ONTAP

To configure a clustered Data ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), run the following command from the cluster shell:

```
network port modify -node <<var_node01>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.

Do you want to continue? {y|n}: y

```
network port modify -node <<var_node02>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.

Do you want to continue? {y|n}: y

## NTP in Clustered Data ONTAP

To configure time synchronization on the cluster, follow these steps:

1. Set the time zone for the cluster.

```
timezone <<var_timezone>>
```

**Note**

For example, in the Eastern United States, the time zone is `America/New_York`.

2. Set the date for the cluster.

```
date <<ccyyymmddhhmm>>
```

**Note**

The format for the date is <[Century][Year][Month][Day][Hour][Minute]>; for example, 201208081240.

3. Configure the Network Time Protocol (NTP) for each node in the cluster.

```
system services ntp server create -node <<var_node01>> -server
<<var_global_ntp_server_ip>> system services ntp server create -node <<var_node02>>
-server <<var_global_ntp_server_ip>>
```

4. Enable the NTP for the cluster.

```
system services ntp config modify -enabled true
```

## SNMP in Clustered Data ONTAP

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the sysLocation and sysContact variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <<var_oncommand_server_fqdn>>
```

## SNMPv1 in Clustered Data ONTAP

1. Set the shared secret plain-text password, which is called a community.

```
snmp community delete all
snmp community add ro <<var_snmp_community>>
```



### Note

Use the delete all command with caution. If community strings are used for other monitoring products, the delete all command will remove them.

## SNMPv3 in Clustered Data ONTAP

SNMPv3 requires that a user be defined and configured for authentication.

1. Create a user called snmpv3user.
 

```
security login create -username snmpv3user -authmethod usm -application snmp
```
2. Select all of the default authoritative entities and select md5 as the authentication protocol.
3. Enter an eight-character minimum-length password for the authentication protocol, when prompted.
4. Select des as the privacy protocol.
5. Enter an eight-character minimum-length password for the privacy protocol, when prompted.

## AutoSupport HTTPS in Clustered Data ONTAP

AutoSupport sends support summary information to NetApp through HTTPS.

Execute the following commands to configure AutoSupport:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>>
-transport https -support enable -noteto <<var_storage_admin_email>>
```

## Cisco Discovery Protocol in Clustered Data ONTAP

Enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers by using the following procedure.

**Note**

To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

To enable CDP on the NetApp storage controllers, complete the following step:

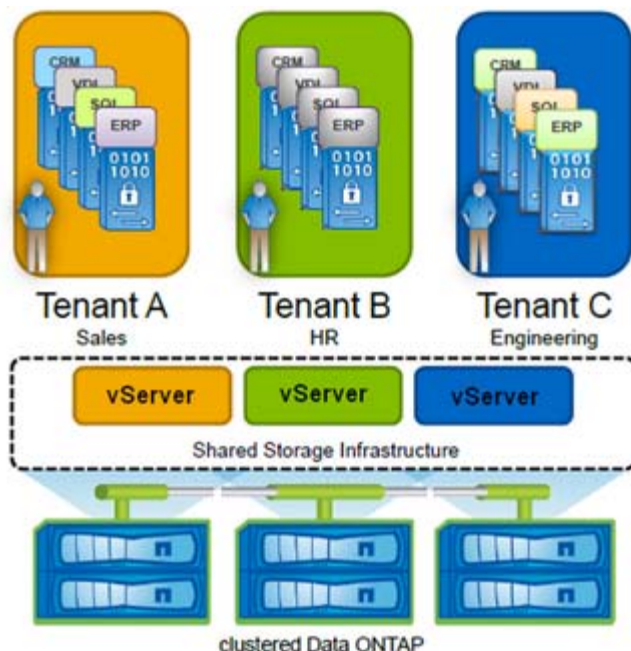
1. Enable CDP on Data ONTAP.

```
node run -node <<var_node01>> options cdpd.enable on
node run -node <<var_node02>> options cdpd.enable on
```

## Vserver - Storage Virtual Machine (SVM)

Secure multi-tenancy is the use of secure virtual partitions within a shared physical storage environment for the purpose of sharing the physical environment among multiple distinct tenants. For instance, a storage service provider might configure a storage array in such a way that each of three different customers is provisioned a certain portion of the array's disk capacity and network resources. In a secure multi-tenant environment, each customer would have access only to the resources explicitly provisioned to that customer. The customer would not have access to other customers' data or even be aware of the existence of the other customers or the fact that they share a common physical array.

The secure logical storage partition through which data is accessed in clustered Data ONTAP is known as a vServer, also called Storage Virtual Machine (SVM). A cluster serves data through at least one and possibly multiple vServer's. A vServer is a logical abstraction that represents a set of physical resources of the cluster. Data volumes and logical network interfaces (LIFs) are created and assigned to a vServer and may reside on any node in the cluster to which the vServer has been given access. A vServer may own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node to another. It is capable of supporting multiple data protocols concurrently. Volumes within the vServer can be junctioned together to form a single NAS namespace, which makes all of the vServers data available through a single share or mount point to NFS and CIFS clients.

**Figure 47** *vServer (Storage Virtual Machine)*

For more information on vServer (SVM) see [TR-4160: Secure Multi-tenancy in Clustered DATA Ontap](#).

The procedure to create a vServer with name `Infra_Vserver` is illustrated below. However, the same procedure can be followed to create additional vServer for isolation in Service Provider or shared environments as required.

To create an infrastructure vServer, follow these steps:

1. Run the vServer setup wizard.

```
vserver setup
```

Welcome to the Vserver Setup Wizard, which will lead you through the steps to create a virtual storage server that serves data to clients.

You can enter the following commands at any time:

"help" or "?" if you want to have a question clarified,  
 "back" if you want to change your answers to previous questions, and  
 "exit" if you want to quit the Vserver Setup Wizard. Any changes you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To accept a default or omit a question, do not enter a value.

Step 1. Create a Vserver.

You can type "back", "exit", or "help" at any question.

2. Enter the vServer name.

```
Enter the Vserver name:Infra_Vserver
```

3. Select the vServer data protocols to configure.

```
Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi}:nfs, fcp
```

4. Select the vServer client services to configure.

- Choose the Vserver client services to configure {ldap, nis, dns}:Enter
5. Enter the vServer's root volume aggregate:  
Enter the Vserver's root volume aggregate {aggr01, aggr02} [aggr01]:aggr01
  6. Enter the vServer language setting. English is the default [C].  
Enter the Vserver language setting, or "help" to see all languages [C]:Enter
  7. Enter the vServer's security style:  
Enter the Vservers root volume's security style {unix, ntfs, mixed}} [unix]: Enter
  8. Answer no to Do you want to create a data volume?  
Do you want to create a data volume? {yes, no} [Yes]: no
  9. Answer no to Do you want to create a logical interface?  
Do you want to create a logical interface? {yes, no} [Yes]: no
  10. Answer no to Do you want to Configure FCP? {yes, no} [yes]: no.  
Do you want to Configure FCP? {yes, no} [yes]: no
  11. Add the two data aggregates to the Infra\_Vserver aggregate list for NetApp Virtual Console.  
vserver modify -vserver Infra\_Vserver -aggr-list aggr01, aggr02

## Create Load Sharing Mirror of vServer Root Volume in Clustered Data ONTAP

1. Create a volume to be the load sharing mirror of the infrastructure vServer root volume on each node.  
volume create -vserver Infra\_Vserver -volume root\_vol\_m01 -aggregate aggr01 -size 20MB -type DP  
volume create -vserver Infra\_Vserver -volume root\_vol\_m02 -aggregate aggr02 -size 20MB -type DP
2. Create the mirroring relationships.  
snapmirror create -source-path //Infra\_Vserver/root\_vol -destination-path //Infra\_Vserver/root\_vol\_m01 -type LS  
snapmirror create -source-path //Infra\_Vserver/root\_vol -destination-path //Infra\_Vserver/root\_vol\_m02 -type LS
3. Initialize the mirroring relationship.  
snapmirror initialize-ls-set -source-path //Infra\_Vserver/root\_vol
4. Set an hourly (at 5 minutes past the hour) update schedule on each mirroring relationship.  
snapmirror modify -source-path //Infra\_Vserver/root\_vol -destination-path \* -schedule hourly

## FC Service in Clustered Data ONTAP

Create the FC service on each vServer. This command also starts the FC service and sets the FC alias to the name of the vServer.

```
fcv create -vserver Infra_Vserver
```

## HTTPS Access in Clustered Data ONTAP

Secure access to the storage controller must be configured.

1. Increase the privilege level to access the certificate commands.  
set -privilege advanced  
Do you want to continue? {y|n}: y
2. Generally, a self-signed certificate is already in place. Check it with the following command:

```
security certificate show
```

3. Run the following commands as one-time commands to generate and install self-signed certificates:

**Note**

You can also use the `security certificate delete` command to delete expired certificates

```
security certificate create -vserver Infra_Vserver -common-name
<<var_security_cert_vserver_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node01>> -common-name
<<var_security_cert_node01_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node02>> -common-name
<<var_security_cert_node02_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
```

4. Configure and enable SSL and HTTPS access and disable Telnet access.

```
system services web modify -external true -ssl3-enabled true
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -action allow
system services firewall policy create -policy mgmt -service http -action deny
-ip-list 0.0.0.0/0
system services firewall policy delete -policy mgmt -service telnet -action allow
system services firewall policy create -policy mgmt -service telnet -action deny
-ip-list 0.0.0.0/0
security ssl modify -vserver Infra_Vserver -certificate
<<var_security_cert_vserver_common_name>> -enabled true
y
security ssl modify -vserver <<var_clustername>> -certificate
<<var_security_cert_cluster_common_name>> -enabled true
y
security ssl modify -vserver <<var_node01>> -certificate
<<var_security_cert_node01_common_name>> -enabled true
y
security ssl modify -vserver <<var_node02>> -certificate
<<var_security_cert_node02_common_name>> -enabled true
y
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
vserver services web access create -name spi -role admin -vserver <<var_clustername>>
<<var_clustername>>
vserver services web access create -name ontapi -role admin -vserver
<<var_clustername>>
```

**Note**

It is normal for some of these commands to return an error message stating that the entry does not exist.

## NFSv3 in Clustered Data ONTAP

Run all commands to configure NFS on the vServer.

1. Secure the default rule for the default export policy and create the FlexPod export policy.

```
vserver export-policy rule modify -vserver Infra_Vserver -policyname default
-ruleindex 1 -rorule never -rwrule never -superuser never
vserver export-policy create -vserver Infra_Vserver FlexPod
```

2. Create a new rule for the FlexPod export policy.

**Note**

For each ESXi host being created, create a rule. Each host will have its own rule index. Your first ESXi host will have rule index 1, your second ESXi host will have rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra_Vserver -policyname FlexPod
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule
sys -superuser sys -allow-suid false
```

3. Assign the FlexPod export policy to the infrastructure vServer root volume.

```
volume modify -vserver Infra_Vserver -volume root_vol -policy FlexPod
```

## FlexVol in Clustered Data ONTAP

The following information is required to create a FlexVol: the volume's name and size, and the aggregate on which it will exist. Create two VMware datastore volumes, a server boot volume, and a volume to hold the OnCommand database LUN. Also, update the vServer root volume load sharing mirrors to make the NFS mounts accessible.

```
volume create -vserver Infra_Vserver -volume infra_datastore_1 -aggregate aggr02 -size
500g -state online -policy FlexPod -junction-path /infra_datastore_1 -space-guarantee
none -percent-snapshot-space 0
volume create -vserver Infra_Vserver -volume infra_swap -aggregate aggr01 -size 100g
-state online -policy FlexPod -junction-path /infra_swap -space-guarantee none
-percent-snapshot-space 0 -snapshot-policy none
```

```
volume create -vserver Infra_Vserver -volume TenantDataSecNFS_datastore_1 -aggregate
aggr02 -size 20g -state online -policy FlexPod -junction-path / TenantDataSecNFS 1
-space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra_Vserver -volume TenantA-PrimaryNFS-Storage datastore_1
-aggregate aggr02 -size 20g -state online -policy FlexPod -junction-path /
TenantDataSecNFS 1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra_Vserver -volume TenantA-SecondaryNFS-Storage datastore_1
-aggregate aggr02 -size 20g -state online -policy FlexPod -junction-path /
TenantDataSecNFS 1 -space-guarantee none -percent-snapshot-space 0
```

```
volume create -vserver Infra_Vserver -volume esxi_boot -aggregate aggr01 -size 100g
-state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra_Vserver -volume OnCommandDB -aggregate aggr02 -size 200g
-state online -policy default -space-guarantee none -percent-snapshot-space 0
snapmirror update-ls-set -source-path //Infra_Vserver/root_vol
```

## LUN in Clustered Data ONTAP

Create two boot LUNS: VM-Host-Infra-01 and VM-Host-Infra-02, and a LUN for Bare-metal: VM-BareMetalHost-Infra-02. Run the following commands:

```
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -size 20g
-ostype vmware -space-reserve disabled
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-02 -size 20g
-ostype vmware -space-reserve disabled
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-BareMetalHost-Infra-02
-size 20g -ostype vmware -space-reserve disabled
```

## Igroup in Clustered Data ONTAP

Create three igroups: TenantA-ESX-HostA-Platinum-FC-igroup, TenantA-ESX-HostB-Platinum-FC-igroup and TenantA-ESX-BareMetalHostA-Platinum-FC-igroup

```
igroup create -igroup TenantA-ESX-HostA-Platinum-FC-igroup -protocol fcp -ostype
vmware -initiator 20:00:00:25:b5:01:01:02,20:00:00:25:b5:01:01:03 -vserver
Infra_vserver
igroup create -igroup TenantA-ESX-HostB-Platinum-FC-igroup -protocol fcp -ostype
vmware -initiator 20:00:00:25:b5:01:01:04,20:00:00:25:b5:01:01:05 -vserver
Infra_vserver
igroup create -igroup TenantA-ESX-BareMetalHost-Platinum-FC-igroup -protocol fcp
-ostype vmware -initiator 20:00:00:25:b5:04:01:06,20:00:00:25:b5:04:01:05 -vserver
Infra_vserver
```

## Deduplication in Clustered Data ONTAP

Enable deduplication on appropriate volumes.

```
volume efficiency on -vserver Infra_Vserver -volume infra_datastore_1
volume efficiency on -vserver Infra_Vserver -volume TenantADataSecNFS
volume efficiency on -vserver Infra_Vserver -volume TenantA-PrimaryNFS-Storage
volume efficiency on -vserver Infra_Vserver -volume TenantA_SecondaryNFS-Storagevolume
efficiency on -vserver Infra_Vserver -volume esxi_boot
volume efficiency on -vserver Infra_Vserver -volume OnCommandDB
```

## Failover Groups NAS in Clustered Data ONTAP

Create an NFS port failover group.

```
network interface failover-groups create -failover-group fg-nfs-<<var_nfs_vlan_id>>
-node <<var_node01>> -port a0a-<<var_nfs_vlan_id>>
network interface failover-groups create -failover-group fg-nfs-<<var_nfs_vlan_id>>
-node <<var_node01>> -port a0a-<<var_nfs_vlan_id>>
```

## NFS LIF in Clustered Data ONTAP

Create an NFS logical interface (LIF).

```
network interface create -vserver Infra_Vserver -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>>
-address 193.191.1.20 -netmask 255.255.255.0 -status-admin up -failover-policy
nextavail -firewall-policy data -auto-revert true -use-failover-group enabled
-failover-group fg-nfs-<<var_nfs_vlan_id>>

network interface create -vserver Infra_Vserver -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>>
-address 193.191.1.30 -netmask 255.255.255.0 -status-admin up -failover-policy
nextavail -firewall-policy data -auto-revert true -use-failover-group enabled
-failover-group fg-nfs-<<var_nfs_vlan_id>>
```

## FCP LIF in Clustered Data ONTAP

Create four FCoE LIFs, two on each node.



```

network interface create -vserver Infra_Vserver -lif fcp_lif01a -role data
-data-protocol fcp -home-node <<var_node01>> -home-port 3a
network interface create -vserver Infra_Vserver -lif fcp_lif01b -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 4a
network interface create -vserver Infra_Vserver -lif fcp_lif02a -role data
-data-protocol fcp -home-node <<var_node01>> -home-port 3a
network interface create -vserver Infra_Vserver -lif fcp_lif02b -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 4a

```

## Add Infrastructure vSserver Administrator

Add the infrastructure vSserver administrator and vServer administration logical interface in the out-of-band management network with the following commands:

```

network interface create -vserver Infra_Vserver -lif vsmgmt -role data -data-protocol
none -home-node <<var_node02>> -home-port e0a -address <<var_vserver_mgmt_ip>>
-netmask <<var_vserver_mgmt_mask>> -status-admin up -failover-policy nextavail
-firewall-policy mgmt -auto-revert true -use-failover-group enabled -failover-group
fg-cluster-mgmt

network routing-groups route create -vserver Infra_Vserver -routing-group
d<<var_clustermgmt_ip>> -destination 0.0.0.0/0 -gateway <<var_clustermgmt_gateway>>
security login password -username vsadmin -vserver Infra_Vserver
Please enter a new password: <<var_vsadmin_password>>
Please enter it again: <<var_vsadmin_password>>

security login unlock -username vsadmin -vserver Infra_Vserver

```

## Cloud Infrastructure on Cisco UCS

This section explains in detail the design considerations for deploying Cisco UCS compute infrastructure for building Citrix CloudPlatform 4.2.1. In a multi-tenant cloud environment management, provisioning, availability, and service levels of shared infrastructure resources are some of the major tasks for the cloud provider.

To build and rebuild Cloud infrastructure within short duration of time is the key challenge in the cloud environment. Cisco Unified Computing System offers data center platform that unites compute, network, storage access, and virtualization into a cohesive system. The Cisco UCS Manager enables Stateless Computing where each compute node has no fixed set of configuration such as, MAC addresses, UUIDs, WWPN, WWNN, adapter policy, firmware and BIOS settings. All configured centrally and applied in the form of Service Profile to the servers. This reduces TCO and increase business agility, allowing consistent configuration and ease of repurposing compute in Citrix CloudPlatform 4.2.1.

This section details the steps to design and deploy the following Cloud components:

- Cloud Management Design and Deployment
- Cloud Compute Design and Deployment

## Cloud Management Design and Deployment

The Cisco UCS offers Organizational Units (OUs) which divides large physical infrastructure of the Cisco UCS domain into logical entities where each cloud tenant will be provided with a dedicated and an isolated management domain with unique resources such as policies, pools, service profiles,

templates and quality of service definitions. On compute security operations front, to provide access control Cisco UCS Manager offers OU locales feature where Users can be delegated with roles and privileges.



#### Note

In this paper we will create single Organization Unit Tenant A with access roles to delegate Cloud User with admin role for performing Cloud operations.

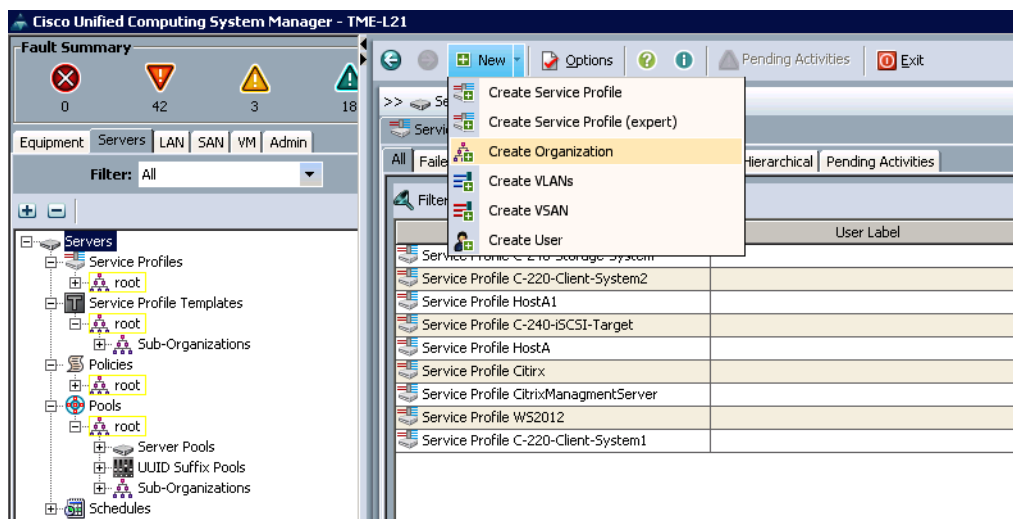
## Creating an Organization

These steps provide details for configuring an organization in the Cisco UCS environment. Organizations are used as a means to organize and control access to various resources like creating Pools (MAC, PWWN, Management IP, and Server), Policy (BIOS, Firmware, and Adapter), and define Service Profile Templates within the IT organization, thereby enabling multi-tenancy of the compute resources. The necessary steps to create these resources are included below.

To create an Organization, login to the Cisco UCS Manager and follow these steps:

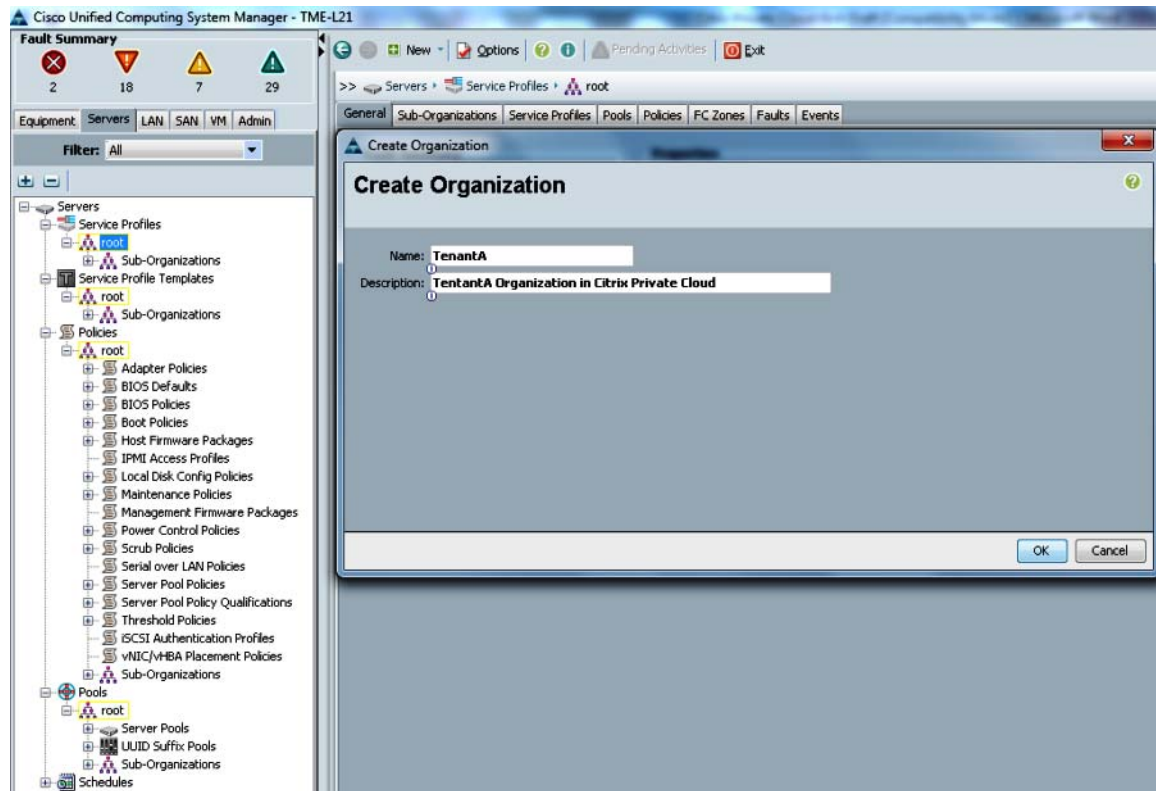
4. From the **Server** tab with the Servers node highlighted in the left pane, choose **New > Create Organization** from the top control bar which is shown in the right pane of the Cisco UCS Manager.

**Figure 48** *Creating and Configuring the Organization*



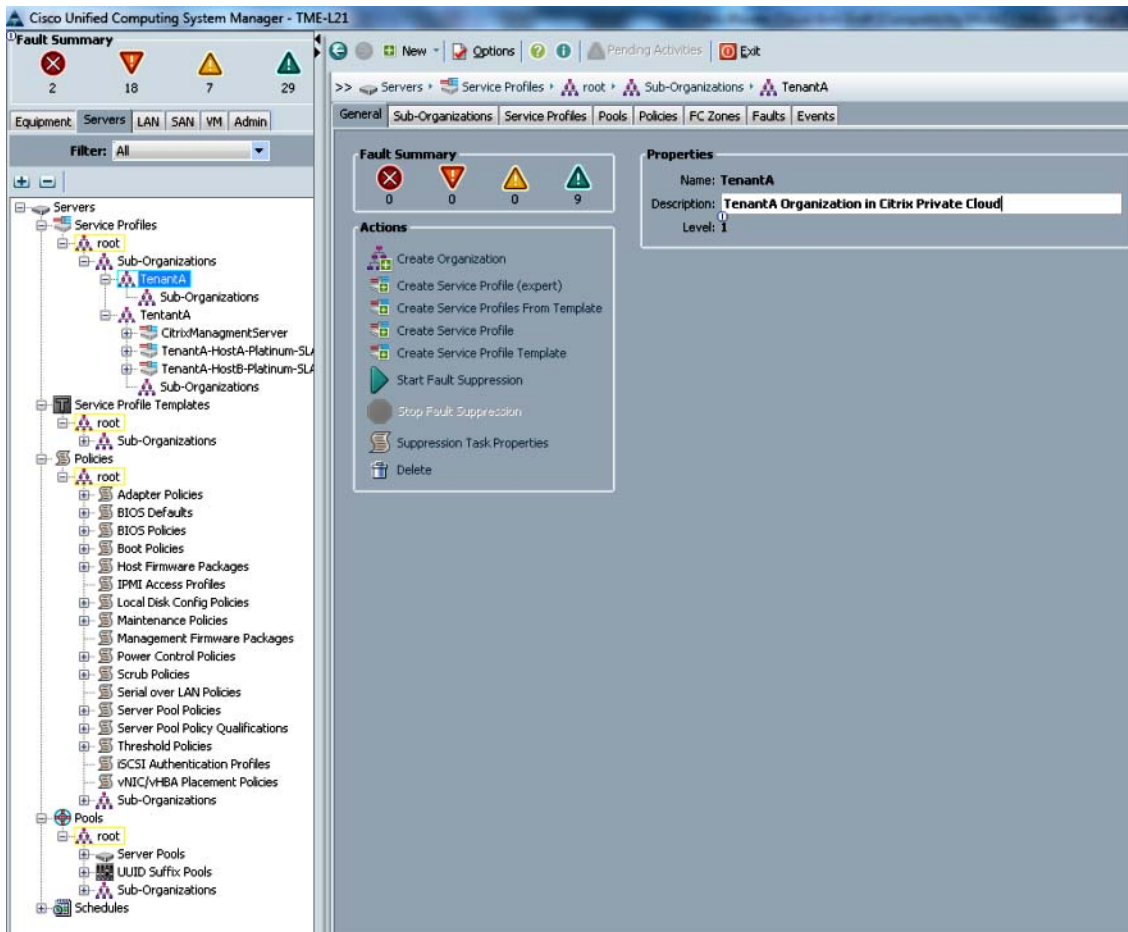
5. Enter a name for the organization in the Name field.
6. Enter a description for the organization (optional) in the Description field.

**Figure 49** Adding the Name and Description to the Organization



7. Click OK.
8. Click OK .

**Figure 50**      *Confirming Organization Properties*

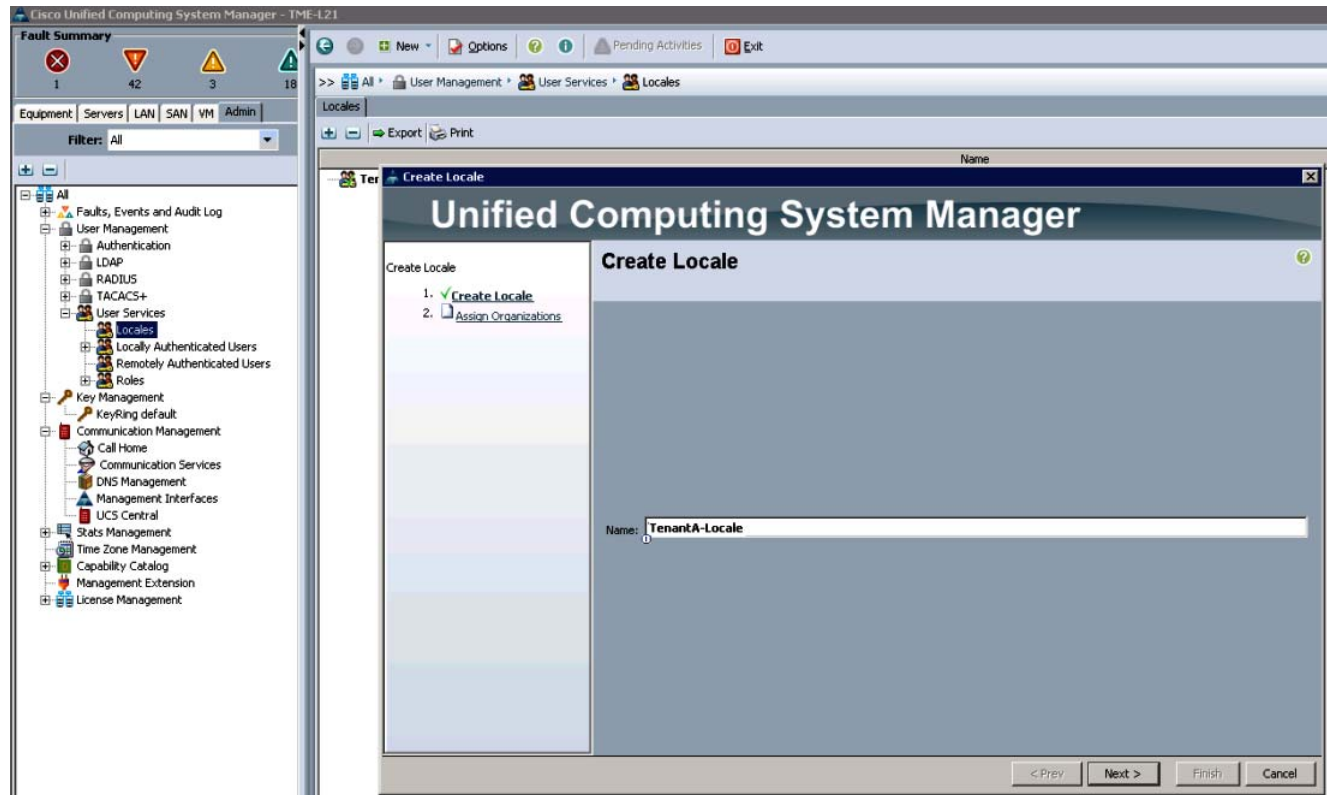


## Creating Locale for Organization

To create locales in the Cisco Cisco UCS Manager, follow these steps:

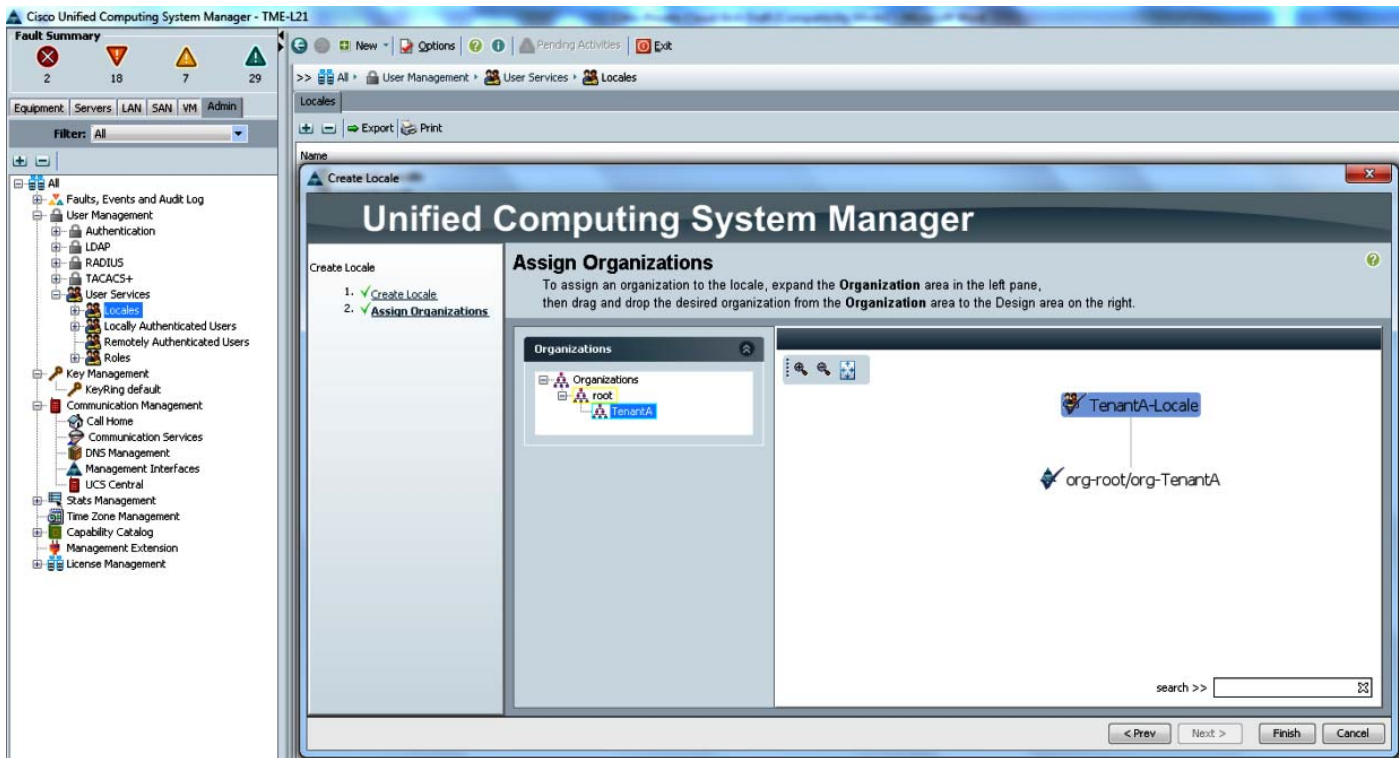
1. In the **Admin** tab in the left pane, click on **User Services** under **User Management**,
2. Right-click **Locales** and click **Create Locale**.
3. Enter a Name **TenantA-Locale**.

**Figure 51** Adding the Name for the Locale



4. Click **Next**.
5. Select **Organizations**, expand **root**, select **TenantA** and drag to **Org-root/org-Tenant A**.
6. Click **Finish**.

**Figure 52**      *Assigning Organization to the Locales*

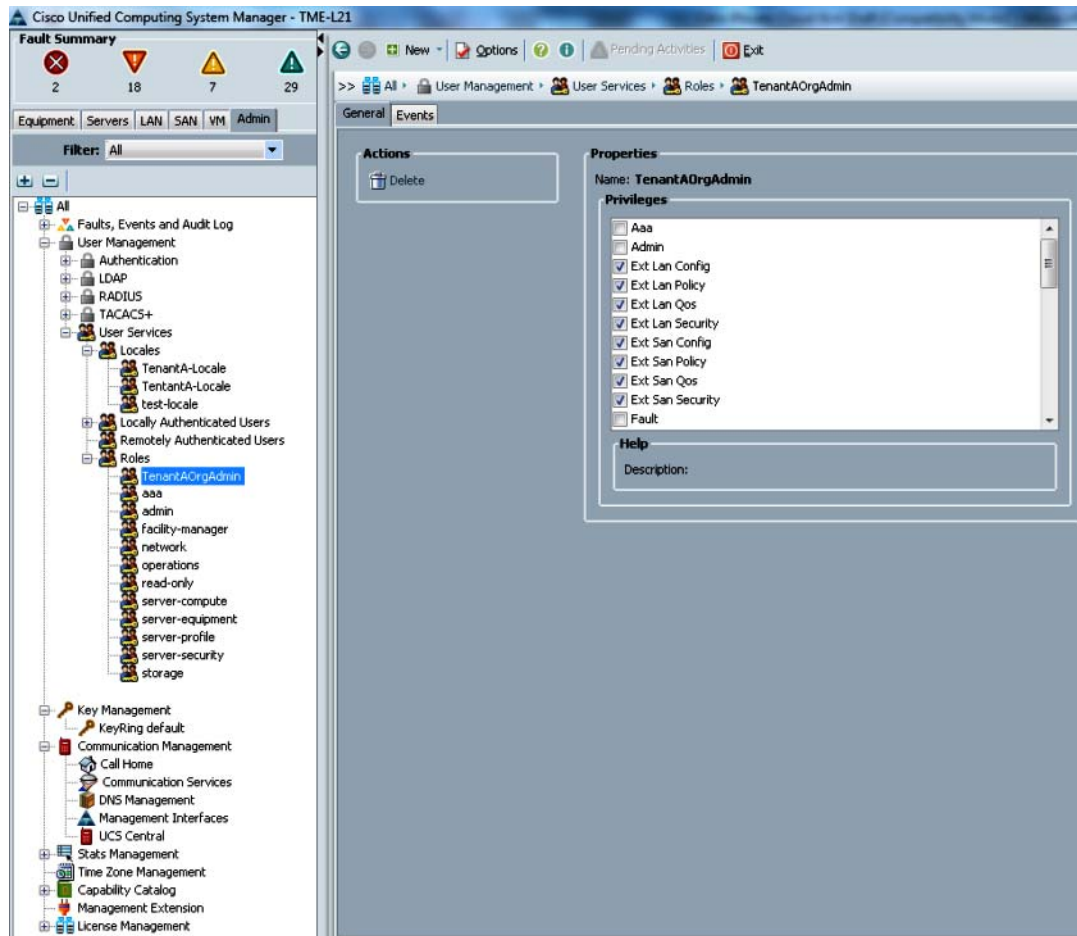


## Creating Role for Organization

To create user roles in the Cisco UCS Manager, follow these steps:

7. In the **Admin** tab in the left pane click **User Roles**.
8. Right-click **Create Role**.
9. Enter a Name **TenantAOrgAdmin** and select appropriate rights for managing **TenantA** Organization.
10. Do not check the **Aaa**, **Admin**, **Operations** and **Fault Privileges** check boxes.

**Figure 53** Defining the Properties and Privileges



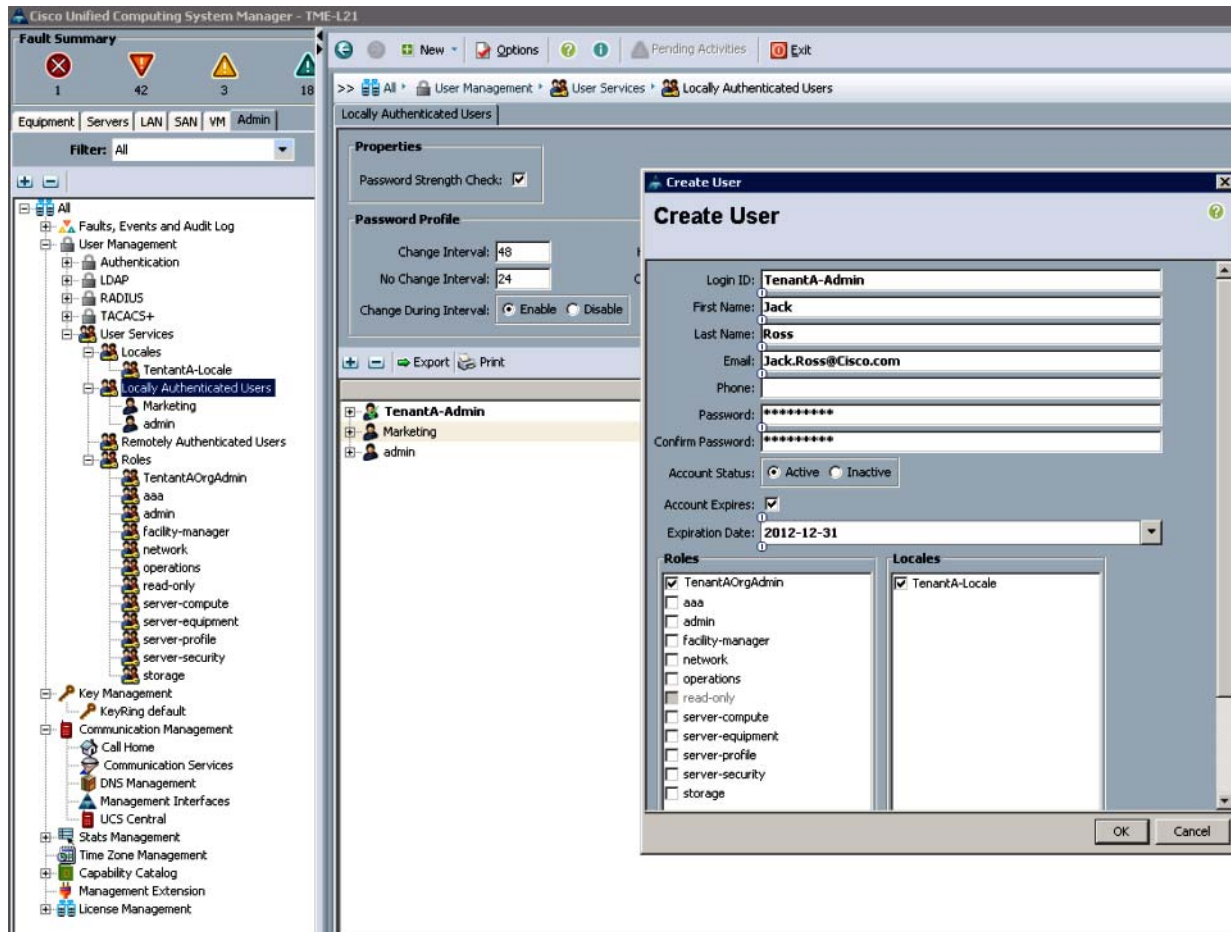
## Creating User and Assign Role for Organization

To create users and assign roles to the organizations, login to the Cisco UCS Manager, and follow these steps

1. In the **Admin** tab in the left pane, click on **LocallyAuthenticatedUsers**.
2. Click **Create User**.
3. Enter Tenant A-Admin in the Name field.
4. Choose the Roles (TenantAOrgAdmin), the Locales (TenantA-Locale).
5. Click **OK**.



**Figure 54** *Defining the Properties for the User Profile*



## Cloud Compute Design and Deployment

In multitenant cloud environment, the Cisco UCS Manager supports data center automation, which helps in increasing operational agility and scalability, while reducing risks. Through its unified, embedded, policy-based, and ecosystem friendly approach, Cisco UCS Manager helps reduce management and administration expenses, which are among the largest items in most IT budgets. The Cisco UCS Manager can manage up to 160 servers and thousands of Cisco UCS components in multiple chassis.

- A unified embedded management interface that integrates server, network, and storage access
- Policy and model-based management, with service profiles, that improves agility and reduces risk
- Auto-discovery to detect, inventory, manage, and provision system components that are added or changed
- Role-based administration that builds on existing skills and supports collaboration across disciplines

Cisco UCS Pools provide predetermined range for an attribute for an UCS blade, such as server UUID, MAC addresses, WWNNs, WWPNN and so on which can be defined at the global level with different attributes and applied to the physical blades in the form of service profile template to achieve faster deployment and state less cloud compute environment.



The service profile contains values for a server's property settings such as virtual network interface cards (vNICs), MAC addresses, virtual host bus adapters (vHBAs), WWPN and WWNN address, boot policies, adapter policies, placement policies, firmware revisions, server pools and other elements. By abstracting these settings from the blade server into a profile, a server can be deployed to any blade server in Cisco UCS. Furthermore, the profile can, at any time, migrate from one blade to another in a chassis.

In multi-tenant environment to logically separate compute identities based on the service levels and management configuration, Cisco UCS Manager allows defining UCS Pools for server UUID, MAC, WWPN, WWNN and Servers.

## Service Level Agreement- Definitions

To provide compute based Service Level Agreement and separate management configurations on Citrix CloudPlatform 4.2.1 environment for multi-tenants compute requirements, the following compute based offering SLAs have been defined. These SLAs are directly mapped to the cloud network offerings, which are provided to cloud tenants for choosing compute infrastructure based on cost.

- **Platinum-Compute**

To satisfy multi-tenant cloud compute service level requirements and to provide higher CPU, and memory, a UCS Server Pool is defined with corresponding management identifier for server UUID, MAC, WWPN, WWNN and IP Management.

- **Gold-Compute**

To satisfy multi-tenant cloud compute service level requirement and to provide slightly lower CPU and memory compared to Platinum-Compute, a UCS Server Pool is defined with corresponding management identifier for Server UUID, MAC, WWPN, WWNN and IP Management.

- **Silver-Compute**

To satisfy multi-tenant cloud compute service level requirement and to provide slightly lower CPU and Memory compared to Gold-Compute, a UCS Server Pool is defined with corresponding management identifier for Server UUID, MAC, WWPN, WWNN and IP Management.

- **Bronze-Compute**

To satisfy multi-tenant cloud compute service level requirement and to provide slightly lower CPU and Memory compared to Silver-Compute, a UCS Server Pool is defined with corresponding management identifier for Server UUID, MAC, WWPN, WWNN and IP Management.

In this section we will create different UCS pools, policies and LAN & SAN pin groups for various Compute-SLAs as defined above. These pools and policies will be applied to corresponding gold service profile template in Organization Unit TenantA, which we have already created for cloud TenantA customer.

This section describes the steps and procedures involved to accomplish the following tasks as per the service level definitions listed earlier:

- Creating UUID Suffix Pool
- Creating a MAC Address Pool
- Creating WWNN Pool
- Creating WWPN Pool
- Creating IP Pool
- Creating Server Pool

- Creating Server Pool Qualifications
- Creating Server BIOS Policy
- Creating Local Disk Configuration Policy
- Creating Boot Policies

## Creating UUID Suffix Pools

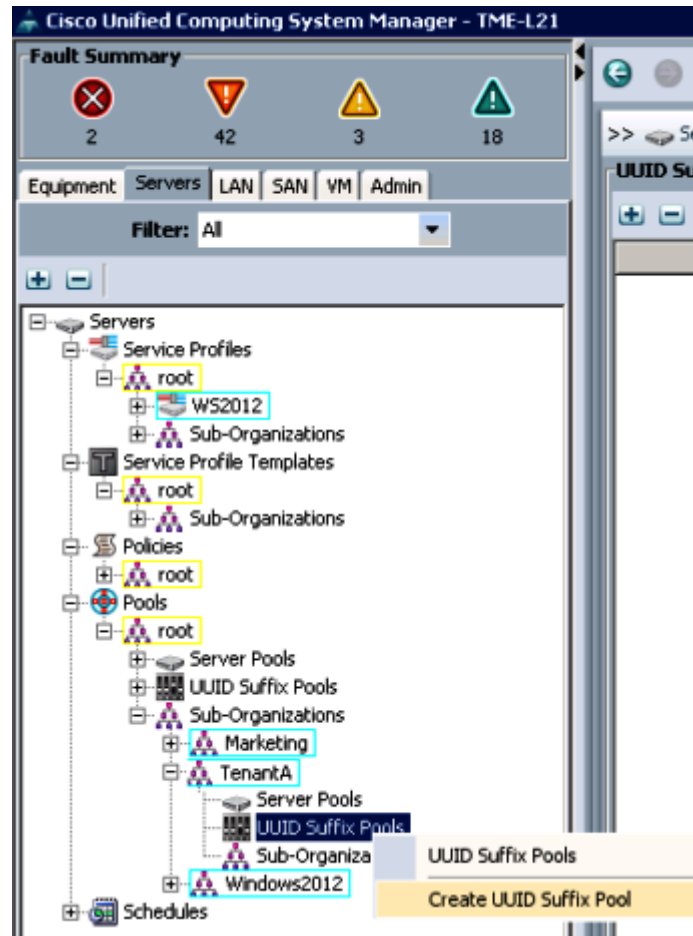
These steps provide details for creating and configuring the necessary UUID suffix pools for the Cisco UCS environment for the Platinum service class:

### Platinum-Compute-UUID

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization Tenant A

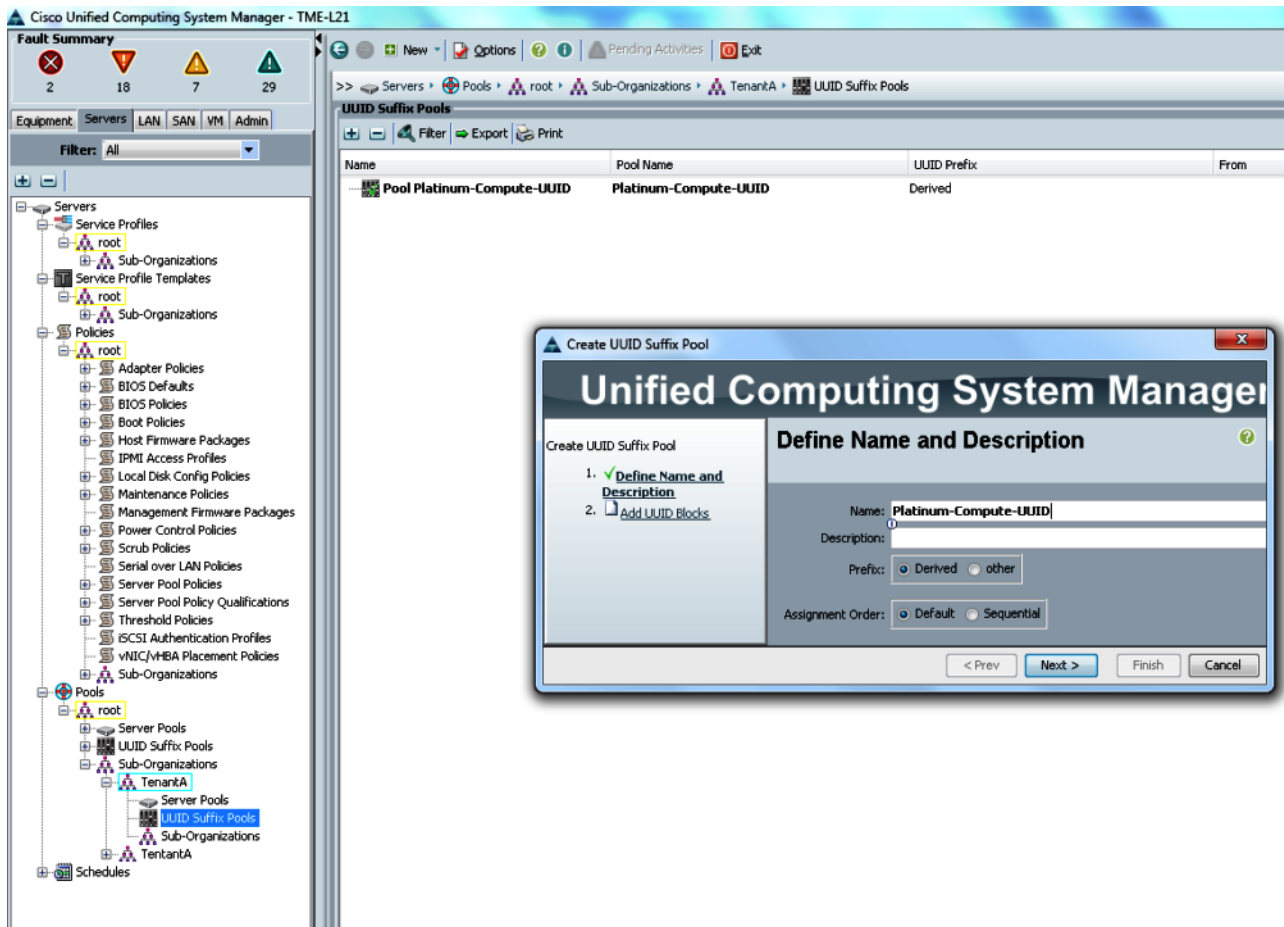
1. Click the **Servers** tab in the left pane.
2. Choose **Pools > Sub-Organizations**.
3. Expand the **Sub-Organizations**.
4. Right-click **UUID Suffix Pools**.
5. Select **Create UUID Suffix Pool**.

**Figure 55**      *Creating the UUID Suffix Pools*



6. Name the UUID suffix pool Platinum-Compute-UUID.
7. (Optional) Give the UUID suffix pool a description.
8. Leave the prefix at the derived option.
9. Click **Next** to continue.

**Figure 56** Adding Name and Description to UUID



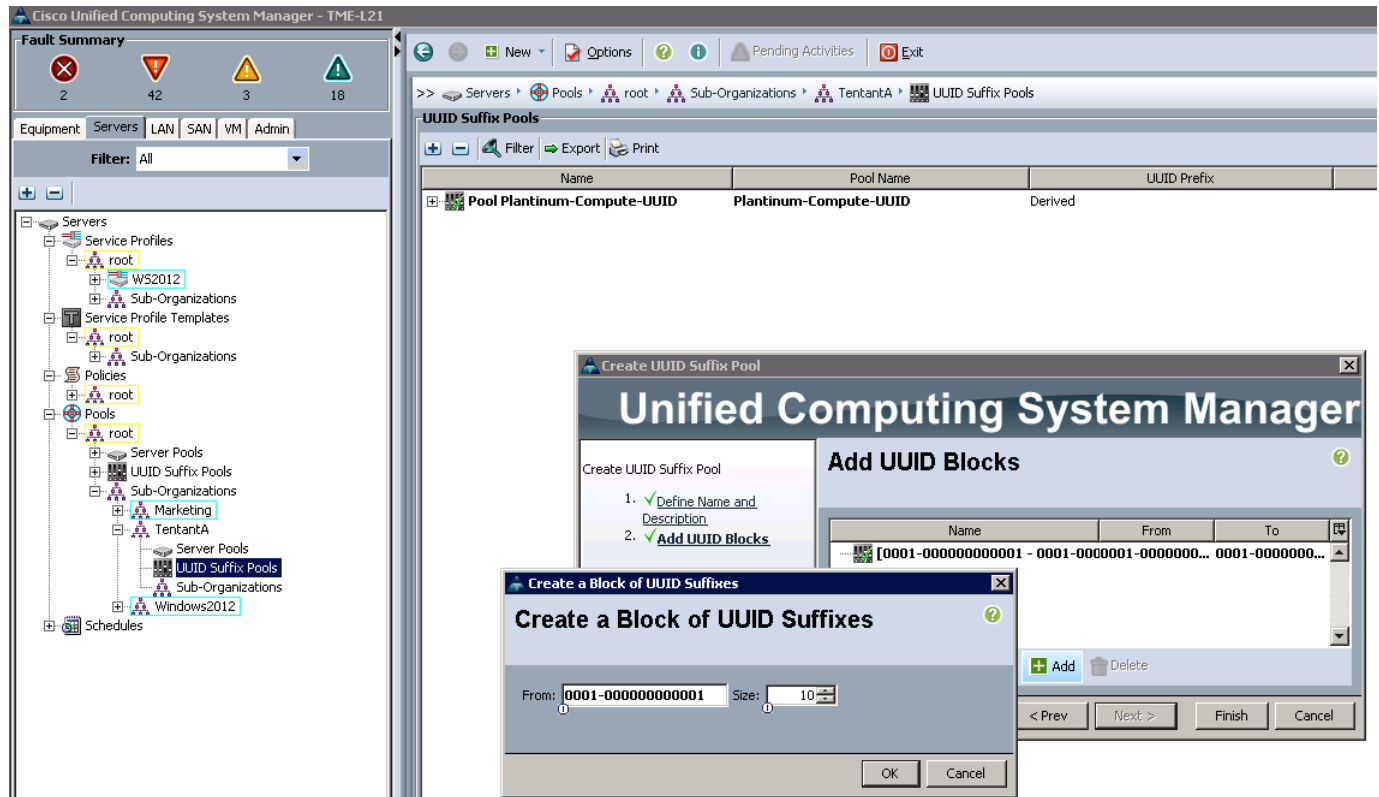
10. Click **Add** to add a block of UUIDs.
11. In the From field type, the UUID suffix value, 0001-000000000000.
12. Specify a size 10 of the UUID block



**Note** Ensure that sufficient number of blade resources are available.

13. Click **OK**.
14. Click **Finish**.

**Figure 57**      **Specifying the UUID Block Size**



Follow the steps described above to create the UUID Pool for Gold, Silver, and Bronze service classes as per the UUID and Size noted in the [Table 10](#).

**Table 10**      **The UUID ID and Size for Different Service Classes**

| UUID Pool Name        | UUID ID           | Size |
|-----------------------|-------------------|------|
| Platinum-Compute-UUID | 0001-000000000000 | 10   |
| Gold-Compute-UUID     | 0002-000000000000 | 10   |
| Silver-Compute-UUID   | 0003-000000000000 | 10   |
| Bronze-Compute-UUID   | 0004-000000000000 | 10   |

**Figure 58** Summary of all the UUID Pools created

| Name                       | Pool Name             | UUID Prefix        | From              | To                |
|----------------------------|-----------------------|--------------------|-------------------|-------------------|
| Pool Bronze-Compute-UUID   | Bronze-Compute-UUID   | E8328CBA-DFAA-11E1 | 0004-000000000001 | 0004-00000000000A |
| Pool Silver-Compute-UUID   | Silver-Compute-UUID   | E8328CBA-DFAA-11E1 | 0003-000000000001 | 0003-00000000000A |
| Pool Gold-Compute-UUID     | Gold-Compute-UUID     | E8328CBA-DFAA-11E1 | 0002-000000000001 | 0002-00000000000A |
| Pool Platinum-Compute-UUID | Platinum-Compute-UUID | E8328CBA-DFAA-11E1 | 0001-000000000001 | 0001-00000000000A |

## Creating a MAC Address Pool

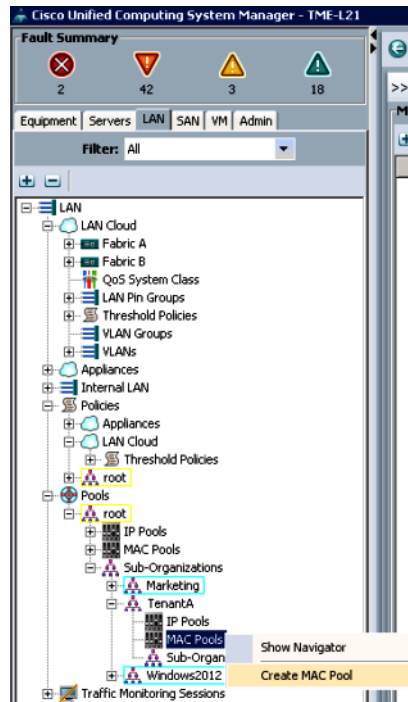
To configure the necessary MAC address pool for the Cisco UCS environment, follow these steps:

### Platinum-Compute-MAC

Login to the Cisco UCS Manager with User Tenant A-Admin created earlier for Organization Tenant A.

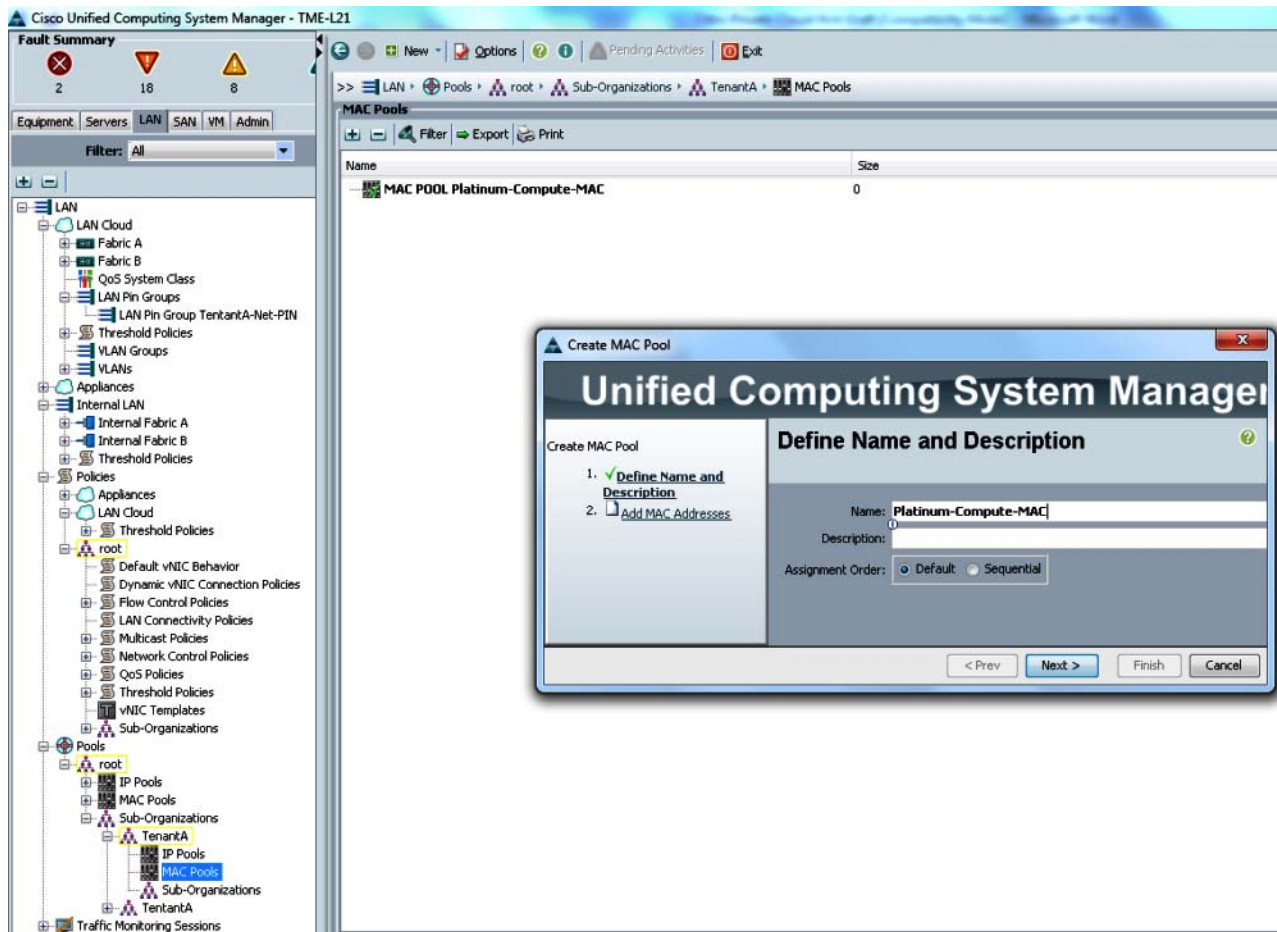
1. Click the **LAN** tab in the left pane.
2. Choose **Pools > Sub Organization**.
3. Expand **Sub-Organization** and click **TenantA**.
4. Right-click **MAC Pools** under the **root** organization.
5. Select **Create MAC Pool** to create the MAC address pool.

**Figure 59**      **Creating the MAC Pool for MAC Addresses**



6. Enter MAC\_Pool as the name of the MAC pool.
7. (Optional) Enter a description of the MAC pool.
8. Click **Next**.

**Figure 60**      *Adding Name and Description to MAC Pool*



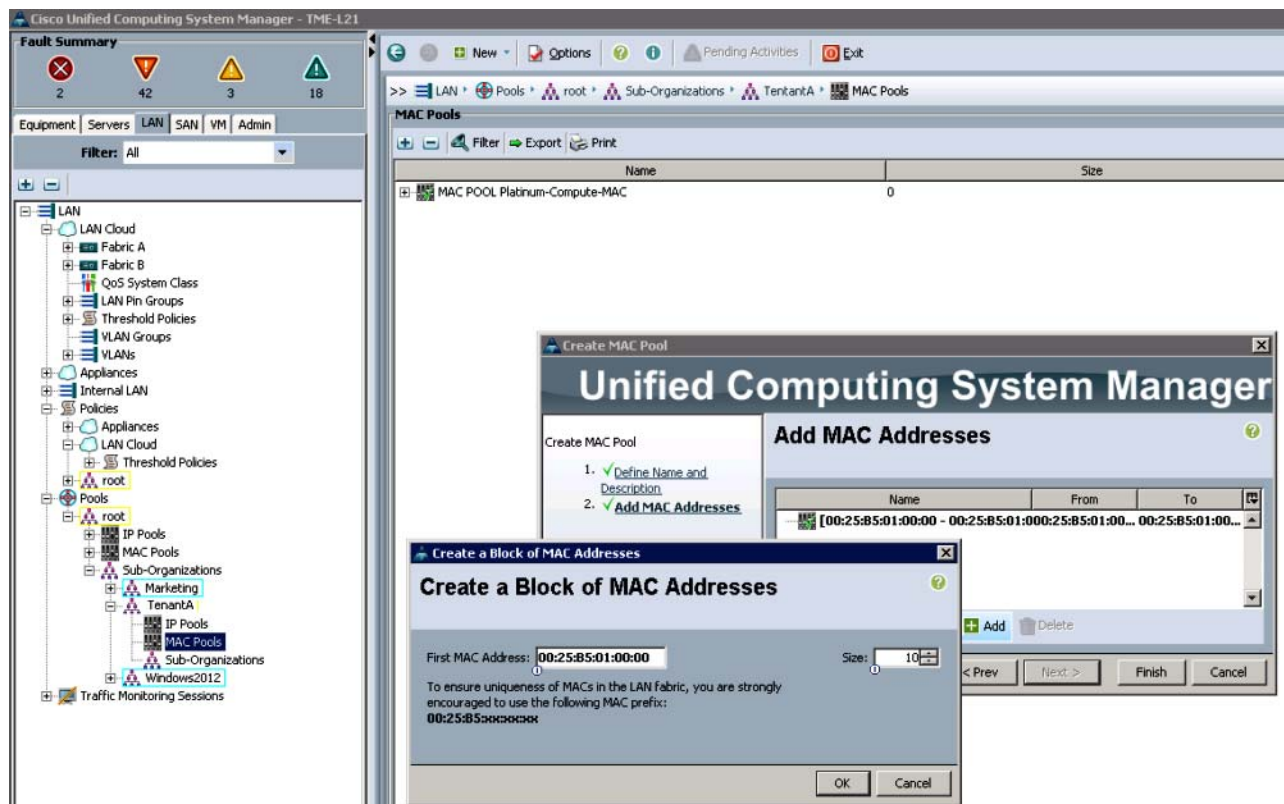
9. Click **Add**.
10. Specify a starting MAC address.
11. Specify 10 size of the MAC address pool.



**Note**      Ensure sufficient number of blade resources are available.



**Figure 61**      **Setting the Block Size for the MAC Address Pool**



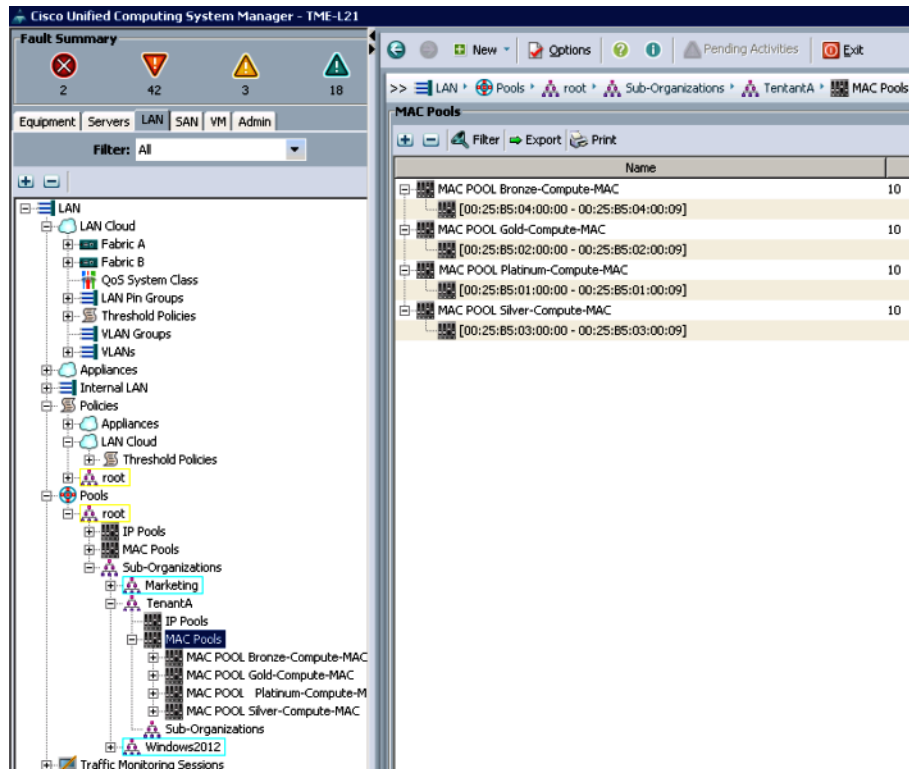
12. Click **OK**.
13. Click **Finish**.
14. In the message box that displays, click **OK**.

Follow the steps described above to create MAC Pool for Gold, Silver, and Bronze service classes as per the MAC and Size noted in [Table 11](#).

**Table 11**      **Block Size and ID for the MAC Pool**

| UUID Pool Name        | UUID ID           | Size |
|-----------------------|-------------------|------|
| Platinum-Compute-UUID | 00:25:B5:01:00:00 | 10   |
| Gold-Compute-UUID     | 00:25:B5:01:00:00 | 10   |
| Silver-Compute-UUID   | 00:25:B5:02:00:00 | 10   |
| Bronze-Compute-UUID   | 00:25:B5:03:00:00 | 10   |

**Figure 62** Summary of all the MAC Pools Created



## Creating WWNN Pools

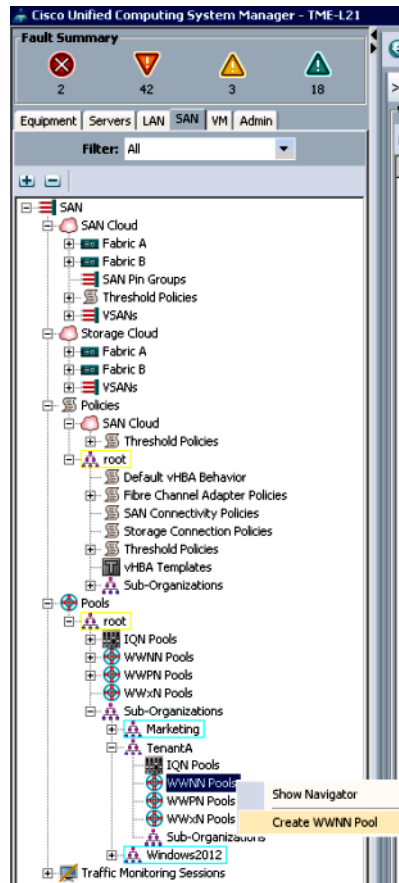
To configure the necessary WWNN pools for the Cisco UCS environment, follow these steps:

### Platinum-Compute-WWNN

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA.:

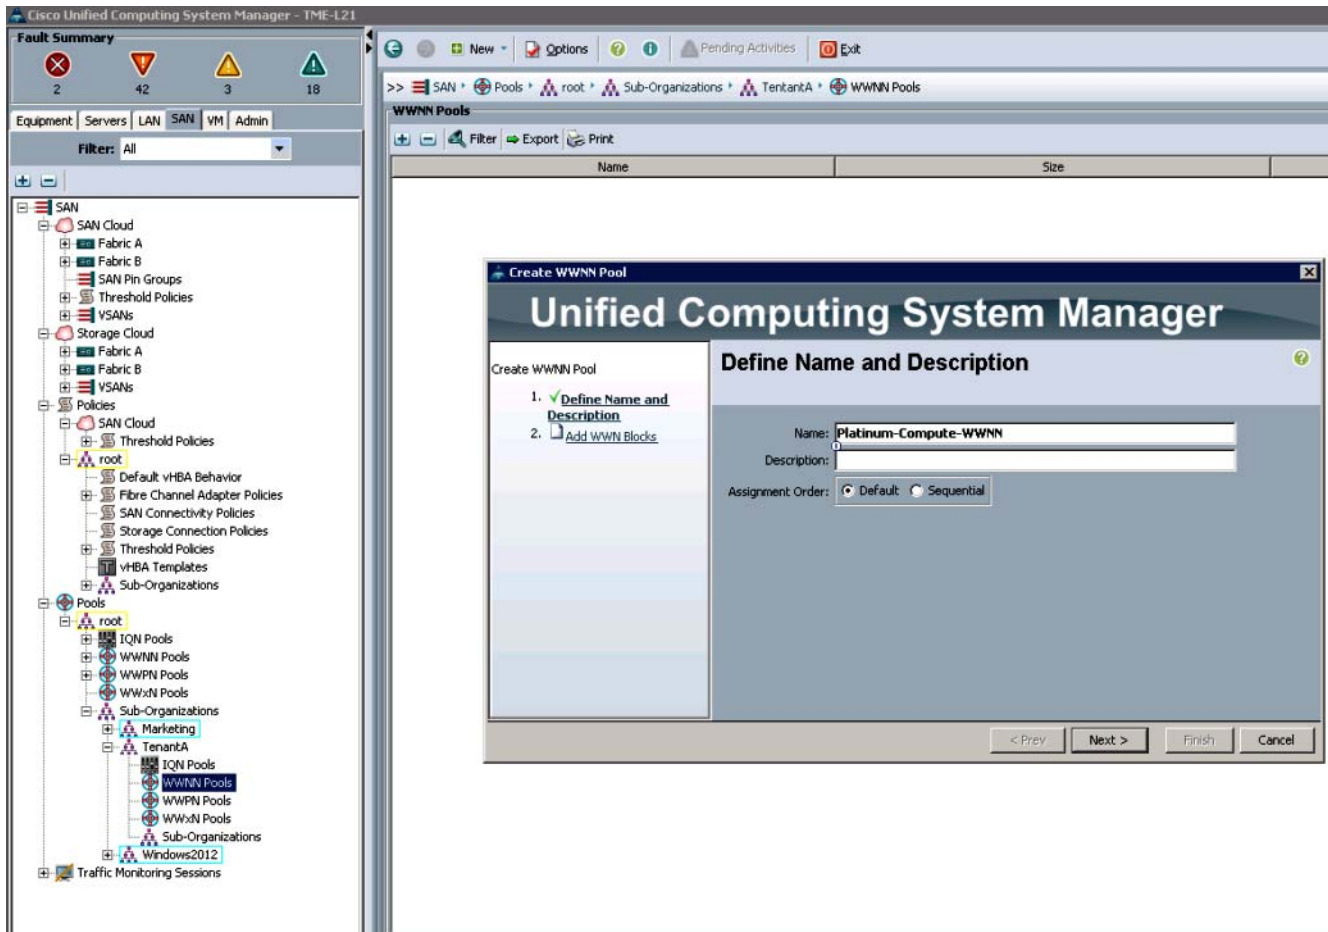
1. Click the **SAN** tab in the left pane.
2. Choose **Pools > Sub-Organization**. Expand TenantA.
3. Right-click **WWNN Pools**.
4. Select **Create WWNN Pool**.

**Figure 63**      **Creating the WWNN Pools**



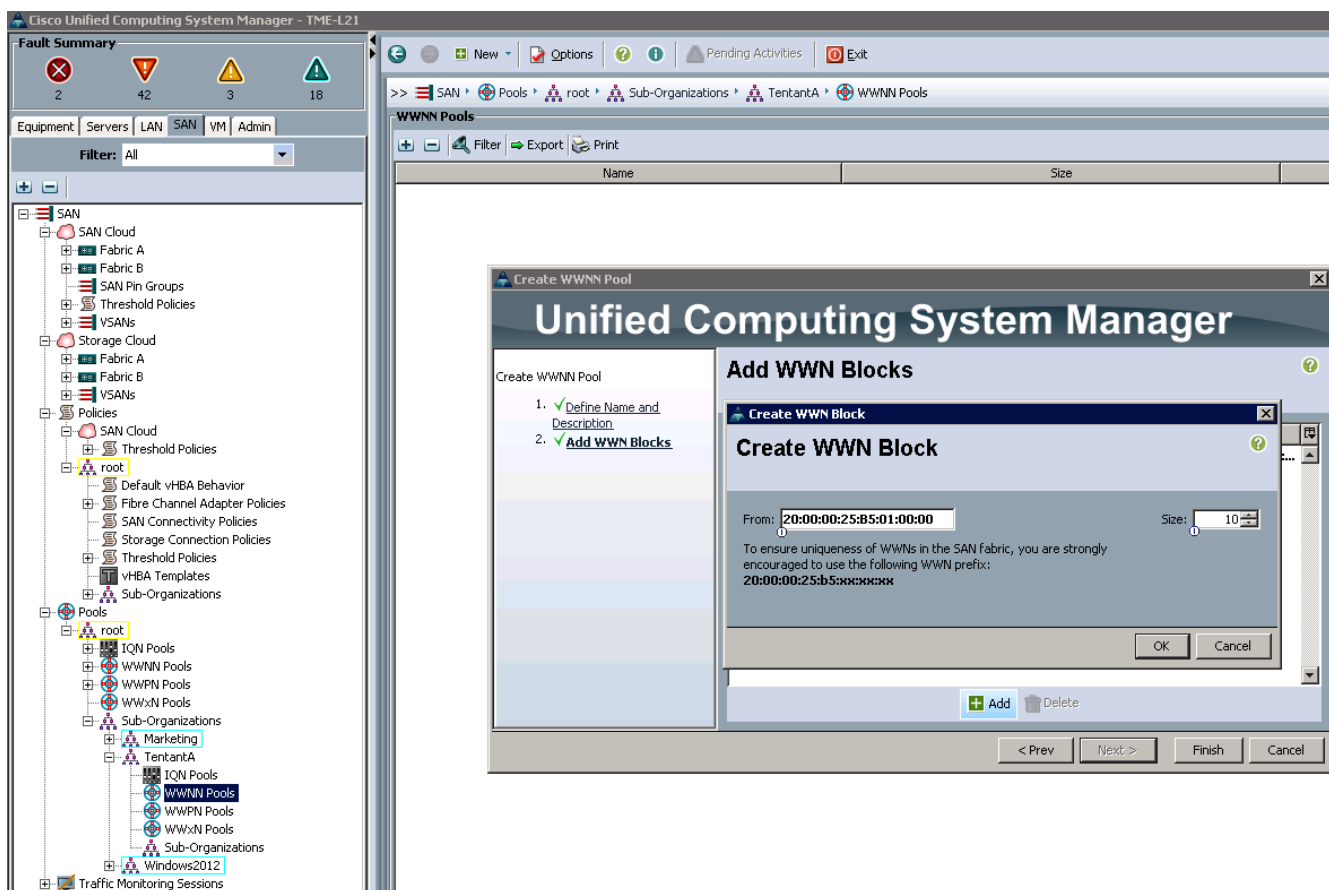
5. Enter **Platinum-Compute-WWNN** as the name of the WWNN pool.
6. (Optional) Add a description for the WWNN pool.
7. Click **Next** to continue.

**Figure 64** *Defining the Name and Description for the WWNN Pool*



8. Click **Add** to add a block of WWNNs 20:00:00:25:B5:01:00:00.
9. Specify size of 10 WWNN block. Ensure that there are sufficient number of blade resources is available.

**Figure 65** Specifying the Size to the WWN Block Created



10. Click **OK** to proceed.

11. Click **Finish** to proceed.

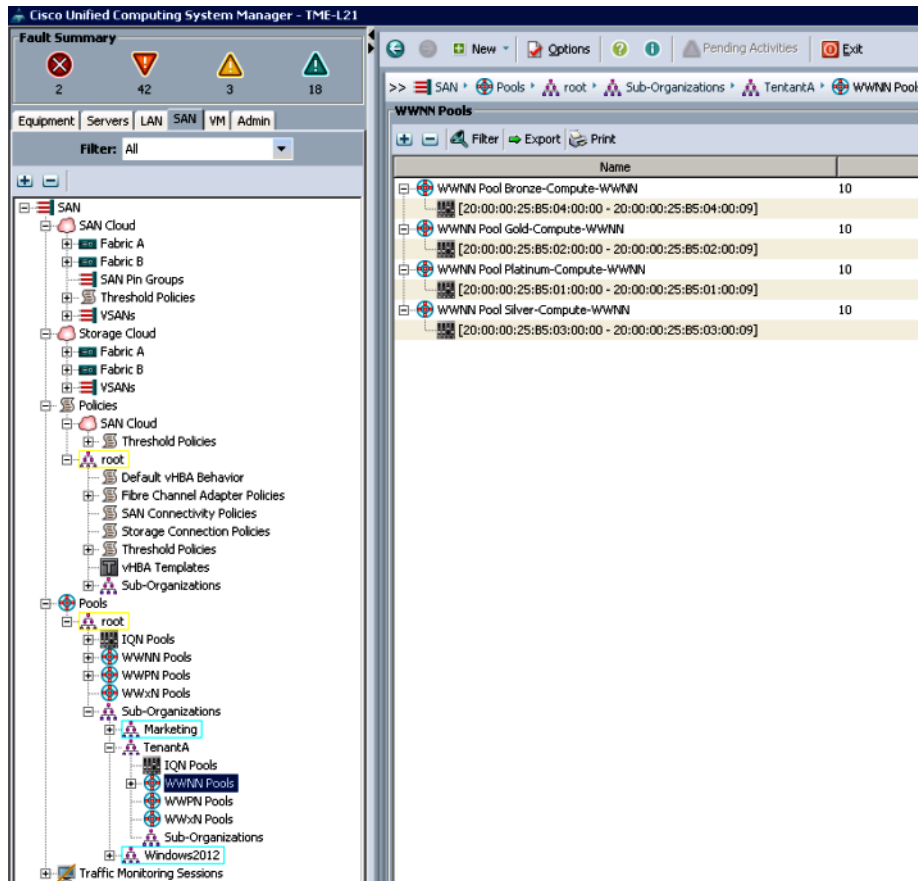
12. Click **OK** to finish.

Follow the same steps above to create WWNN Pool for Gold-Compute-WWNN, Silver-Compute-WWNN and Bronze-Compute-WWNN with below WWNN and Size noted in [Table 12](#).

**Table 12** UUID IDs and Size for all the Service Classes

| UUID Pool Name        | UUID ID                 | Size |
|-----------------------|-------------------------|------|
| Platinum-Compute-WWNN | 20:00:00:25:B5:01:00:00 | 10   |
| Gold-Compute-WWNN     | 20:00:00:25:B5:02:00:00 | 10   |
| Silver-Compute-WWNN   | 20:00:00:25:B5:03:00:00 | 10   |
| Bronze-Compute-WWNN   | 20:00:00:25:B5:04:00:00 | 10   |

**Figure 66** Summary of all the WWNN Pools Created



## Creating WWPN Pools

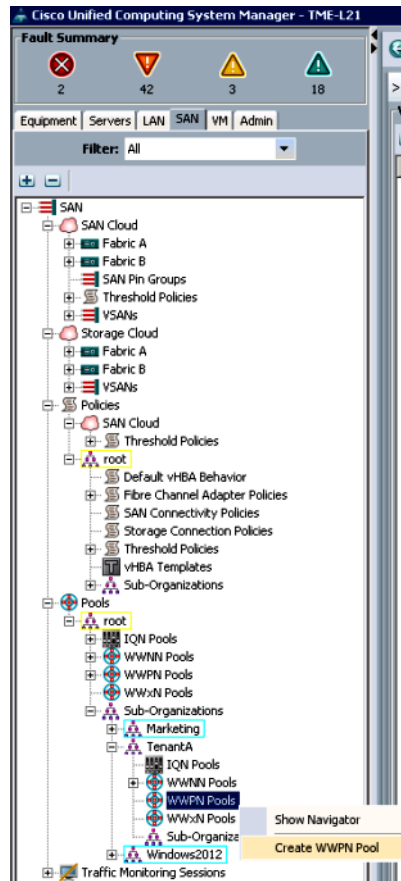
To configure necessary WWNN pools for the Cisco UCS environment, follow these steps:

### Platinum-Compute-WWPN

Login to Cisco UCS Manager with User TenantA-Admin which was created earlier for Organization TenantA

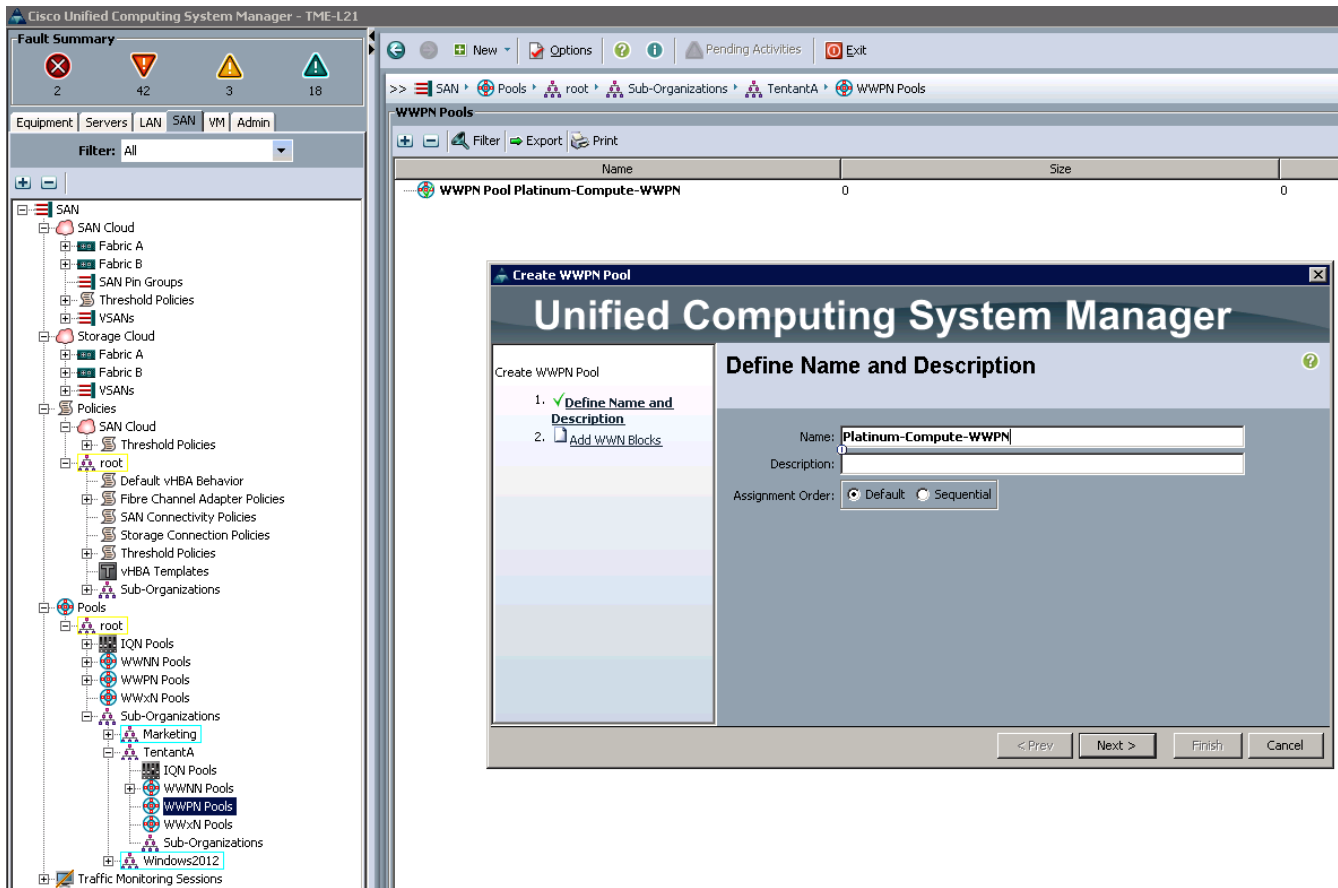
1. Click the **SAN** tab in the left pane.
2. Choose **Pools > Sub-Organization**. Expand TenantA.
3. Right-click **WWPN Pools**.
4. Select **Create WWPN Pool**.

**Figure 67**      **Selecting Create WWPN Pool option**



5. Enter Platinum-Compute-WWPN as the name of the WWPN pool.
6. (Optional) Add a description for the WWPN pool.
7. Click **Next** to continue.

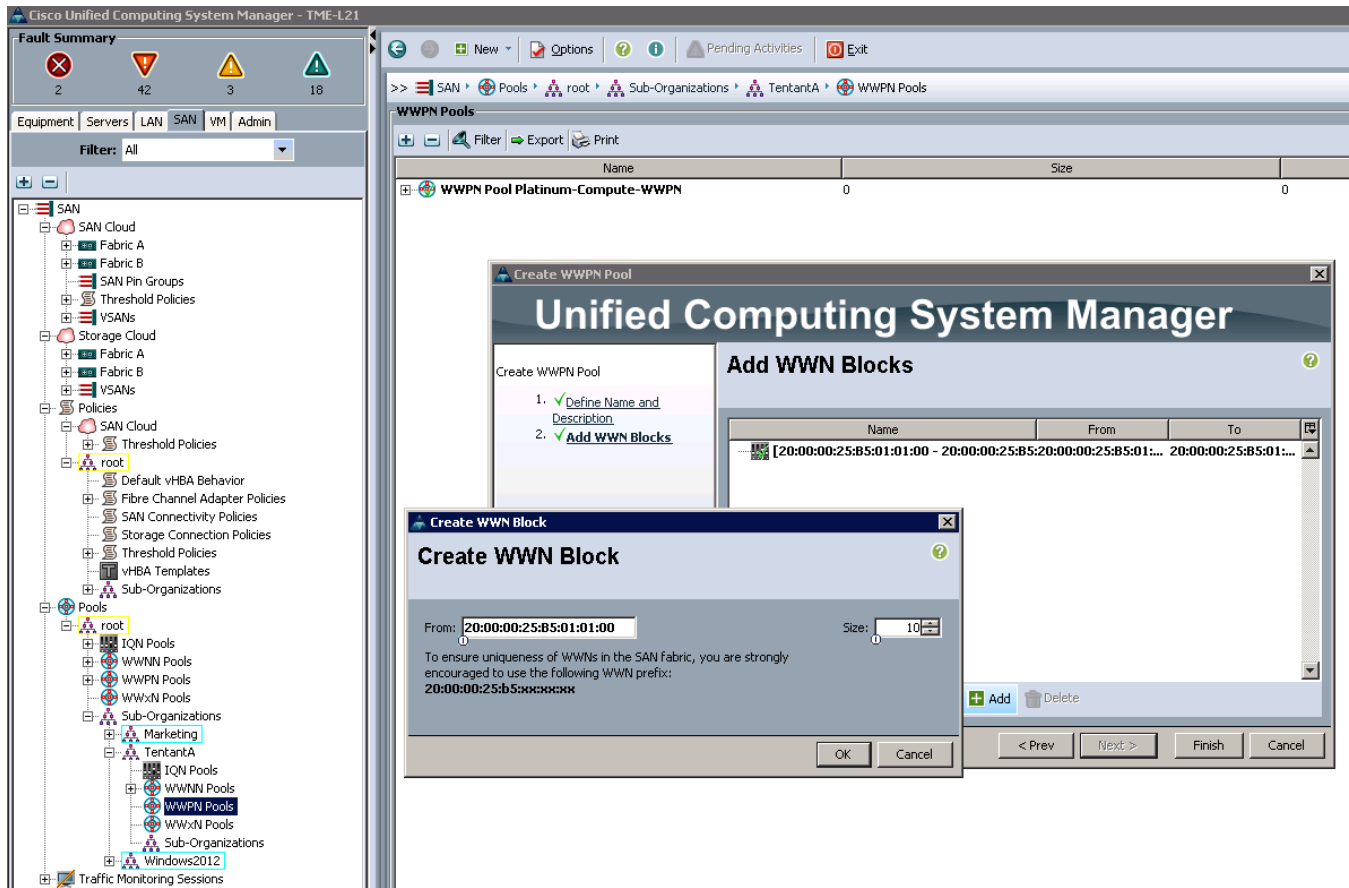
**Figure 68** Adding the Name and Description for the WWPNN Pool



8. Click **Add** to add a block of WWNNs 20:00:25:B5:01:01:00.
9. Specify a size of 10 WWNN block (Ensure that sufficient number of blade resources is available).



**Figure 69**      **Defining the Size of the WWN Block**



10. Click **OK** to proceed.

11. Click **Finish** to proceed.

12. Click **OK** to finish.

Follow the same steps above to create WWPN Pool for Gold-Compute-WWPN, Silver-Compute-WWPN and Bronze-Compute-WWPN with below WWNN and Size noted in [Table 13](#).

**Table 13**      **WWPN UUID Pool name and ID**

| UUID Pool Name      | UUID ID                 | Size |
|---------------------|-------------------------|------|
| Gold-Compute-WWPN   | 20:00:00:25:B5:02:01:00 | 10   |
| Silver-Compute-WWPN | 20:00:00:25:B5:03:01:00 | 10   |
| Bronze-Compute-WWPN | 20:00:00:25:B5:04:01:00 | 10   |

**Figure 70** Summary of all the four compute UUID Pools

The screenshot shows the Cisco Unified Computing System Manager interface. The left pane displays a tree view of the configuration hierarchy, with 'WWPN Pools' selected under 'Pools'. The right pane shows a table summarizing the four compute UUID pools.

| Name                                                                                   | Size | Assigned |
|----------------------------------------------------------------------------------------|------|----------|
| WWPN Pool Bronze-Compute-WWPN<br>[20:00:00:25:85:04:01:00 - 20:00:00:25:85:04:01:09]   | 10   | 0        |
| WWPN Pool Gold-Compute-WWPN<br>[20:00:00:25:85:02:01:00 - 20:00:00:25:85:02:01:09]     | 10   | 0        |
| WWPN Pool Platinum-Compute-WWPN<br>[20:00:00:25:85:01:01:00 - 20:00:00:25:85:01:01:09] | 10   | 1        |
| WWPN Pool Silver-Compute-WWPN<br>[20:00:00:25:85:03:01:00 - 20:00:00:25:85:03:01:09]   | 10   | 0        |

## Creating IP Pools

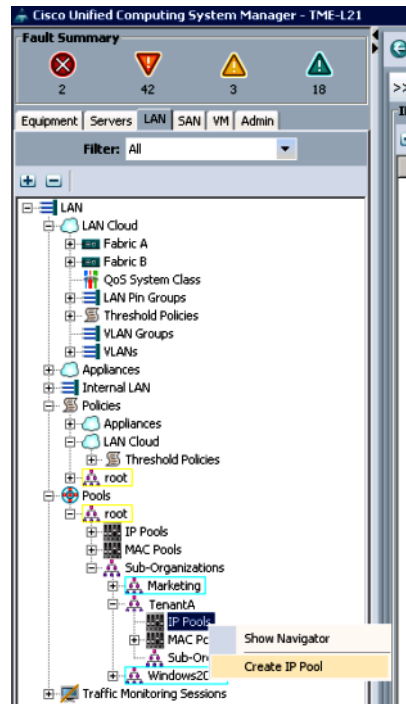
To configure the necessary IP Management pools for the Cisco UCS environment, follow these steps:

### Platinum-Compute-IP

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:.

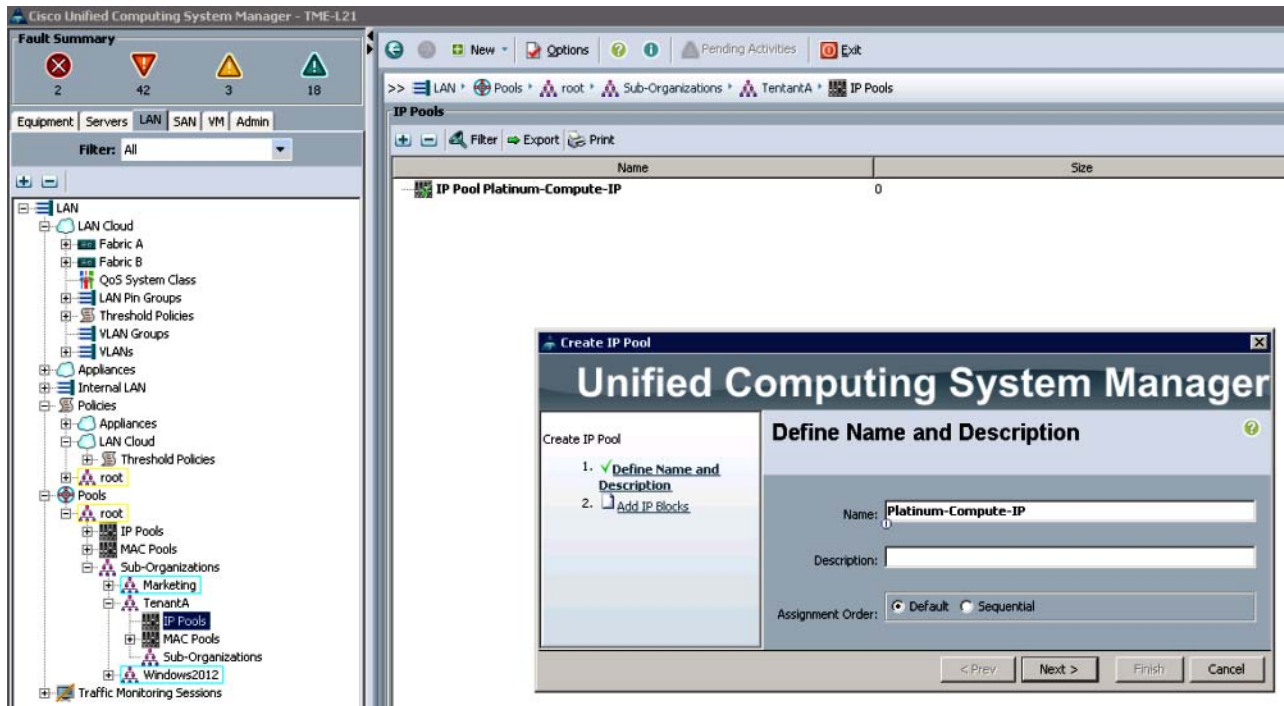
1. Click the **LAN** tab in the left pane.
2. Choose **Pools > Sub-Organization**.
3. Expand **TenantA**.
4. Right-click **IP Pools**.
5. Select **Create IP Pool**.

**Figure 71**      **Creating IP Pools for the Storage**



6. Enter Platinum-Compute-IP as the name of the IP pool.
7. (Optional) Add a description for the IP pool.
8. Click **Next** to continue.

**Figure 72**      *Defining the Name and Description to the IP Pool*



9. Click **Add** to add a block of IP 10.65.121.50.

10. Specify a size of 10 IP block.



**Note** Ensure that sufficient number of blade resources are available.

Specifying the IP Addresses and the Block Size and the DNS

11. Click **OK** to proceed.

12. Click **Finish** to proceed.

13. Click **OK** to finish.

Follow the steps described above to create IP Pool for Gold-Compute-IP, Silver-Compute-IP and Bronze-Compute-IP with WWPN and Size noted in [Table 14](#).

**Table 14**      *WWPN and Size for all UUID Pools*

| UUID Pool Name        | UUID ID      | Size |
|-----------------------|--------------|------|
| Platinum-Compute-WWPN | 10.65.121.50 | 10   |
| Gold-Compute-WWPN     | 10.65.121.61 | 10   |
| Silver-Compute-WWPN   | 10.65.121.71 | 10   |
| Bronze-Compute-WWPN   | 10.65.121.81 | 10   |

**Figure 73** Summary of all the IP Pools Created

| Name                                                         | Size | Assigned |
|--------------------------------------------------------------|------|----------|
| IP Pool Bronze-Compute-IP<br>[10.65.121.71 - 10.65.121.80]   | 10   | 0        |
| IP Pool Gold-Compute-IP<br>[10.65.121.61 - 10.65.121.70]     | 10   | 0        |
| IP Pool Platinum-Compute-IP<br>[10.65.121.50 - 10.65.121.59] | 10   | 3        |
| IP Pool Silver-Compute-IP<br>[10.65.121.71 - 10.65.121.80]   | 10   | 0        |

## Creating vNIC and vHBA Template

Cisco UCS Manager provides templates for the primary objects (vNICs and vHBAs) to facilitate reuse and rapid deployment. The vNIC and vHBA resources are always associated with specific FIs (A-side or B-side fabric interconnect) one bound to each side. vNIC (or vHBA) templates can be used to encapsulate both the MAC address pool (or WWPN pool) association.

In this study we will create two sets of vNIC and vHBA templates for Baremetal and Virtual Hosts Service Profile templates.



### Note

Create six vNIC templates for handling Management, Application, NFS and vMotion LAN traffic and two vHBA templates for handling SAN traffic to run VMware ESXi 5.1 hypervisor host for running Virtual Machines on TenantA Zone.

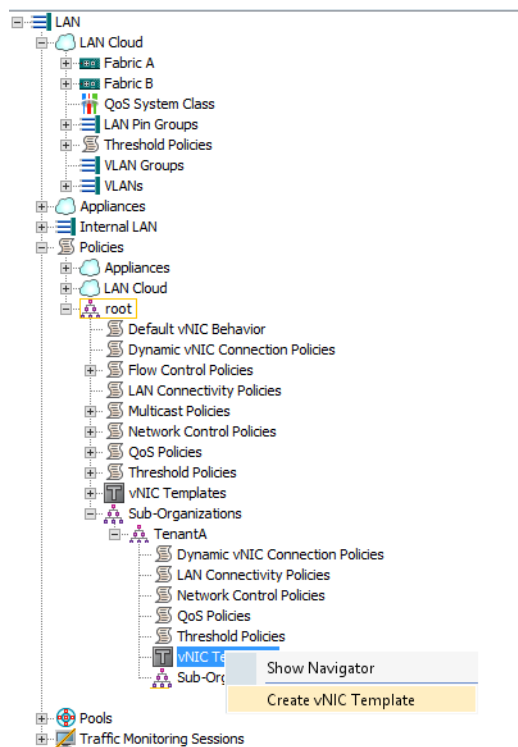
## Creating vNICs Templates

This section details the creation of vNIC templates created for BareMetal and VMware ESXi Hosts on Fabric A and Fabric B.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA::

1. Click the **LAN** tab in the left pane.
2. Choose **Pools > Sub-Organization**.
3. Expand **TenantA**.
4. Select **vNIC Templates** and right click **Create vNIC Template**.

**Figure 74**      **Creating vNIC Template**



5. Create **vNIC Template** window opens.
6. Enter Management-A-NIC in the Name field.
7. Enter Management Template in the **Description** field.
8. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
9. In **Target** check **Adapater** check box.
10. In **Template Type** select **Initial Template** radio button.
11. Under VLANs check **Management-VLAN**, **vMotion-VLAN** check box, and choose **Native VLAN** as Management-VLAN.
12. In the **MTU** field enter 9000.
13. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
14. Select **Silver-Net-SLA** in **QoS Policy** list box.
15. Select **default** in **Network Control Policy** list box.
16. Select **<not set>** in **Pin Group** list box.
17. Select **default** in **Stats Thershold Policy** list box.
18. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
19. Click **OK**.

**Figure 75** Defining the vNIC Properties for Management A-NIC

**Create vNIC Template**

Name: **Managment-A-NIC**

Description: **Managment Template**

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

**Target**

☒ Adapter ☐ VM

**Warning**  
If **VM** is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

**VLANs**

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input checked="" type="checkbox"/> | Management-VLAN      | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | NFS-VLAN             | <input type="radio"/>            |
| <input type="checkbox"/>            | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | vMotion-VLAN         | <input type="radio"/>            |

+ Create VLAN

MTU: **9000**

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: **Platinum-Compute-MAC(15/...**

QoS Policy: **Silver-Net-SLA**

Network Control Policy: **default**

Pin Group: **<not set>**

Stats Threshold Policy: **default**

OK Cancel

20. Create **vNIC Template** window opens.
21. Enter **Management-B-NIC** in the **Name** field.
22. Enter **Management Template** in the **Description** field.
23. In **Fabric ID** click **FabricB** radio button, do not check **Enable Failover** check box.
24. In **Target** check **Adapater** check box.
25. In **Template Type** select **Initial Template** radio button.
26. Under **VLANs** check **Management-VLAN**, **vMotion-VLAN** check box, and choose **Native VLAN** as **Management-VLAN**.
27. In the **MTU** field enter **9000**.
28. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
29. Select **Silver-Net-SLA** in **QoS Policy** list box.
30. Select **default** in **Network Control Policy** list box.
31. Select **<not set>** in **Pin Group** list box.
32. Select **default** in **Stats Thershold Policy** list box.
33. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.

34. Click **OK**.

**Figure 76** Defining the vNIC Properties for Management- B-NIC

**Create vNIC Template**

Name:

Description:

Fabric ID: ☐ Fabric A ☒ Fabric B ☐ Enable Failover

Target: ☒ Adapter ☐ VM

**Warning**  
If VM is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input checked="" type="checkbox"/> | Management-VLAN      | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | NFS-VLAN             | <input type="radio"/>            |
| <input type="checkbox"/>            | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | vMotion-VLAN         | <input type="radio"/>            |

+ Create VLAN

MTU:

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

35. Create **vNIC Template** window opens.
36. Enter guest-A-NIC in the Name field
37. Enter guest Template in the Description field.
38. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
39. In **Target** check **Adapater** check box.
40. In **Template Type** select **Initial Template** radio button.
41. Under **VLANs** check **guest-VLAN**, **guest-VLAN-1000**, **VLAN-guest-VLAN-1001**, and choose **Native VLAN** as guest-VLAN.
42. In the MTU field enter 9000.
43. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
44. Select **Platinum-Net-SLA** in **QoS Policy** list box.
45. Select **default** in **Network Control Policy** list box.
46. Select **<not set>** in **Pin Group** list box.



47. Select **default** in **Stats Thershold Policy** list box.
48. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
49. Click **OK**.

**Figure 77** Defining the vNIC Properties for guest- A-NIC

**Create vNIC Template**

Name:

Description:

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

**Target**

☒ Adapter ☐ VM

**Warning**  
If VM is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input checked="" type="checkbox"/> | Guest-VLAN           | <input checked="" type="radio"/> |
| <input checked="" type="checkbox"/> | Guest-Vlan-1000      | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input type="checkbox"/>            | NFS-VLAN             | <input type="radio"/>            |

[+ Create VLAN](#)

MTU:

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

50. Create vNIC Template window opens.
51. Enter guest-B-NIC in the Name field
52. Enter guest Template in the Description field.
53. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
54. In **Target** check **Adapater** check box.
55. In **Template Type** select **Initial Template** radio button.
56. Under **VLANs** check **guest-VLAN**, **guest-VLAN-1000**, **VLAN-guest-VLAN-1001**, and choose **Native VLAN** as guest-VLAN.
57. In the **MTU** field enter 9000.
58. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
59. Select **Platinum-Net-SLA** in **QoS Policy** list box.
60. Select **default** in **Network Control Policy** list box.

61. Select <not set> in **Pin Group** list box.
62. Select **default** in **Stats Thershold Policy** list box.
63. Select <not set> in **Dynamic vNIC Connection Policy** list box.
64. Click **OK**.

**Figure 78** Defining the vNIC Properties for guest- B-NIC

**Create vNIC Template**

Name:

Description:

Fabric ID: ☐ Fabric A ☒ Fabric B ☐ Enable Failover

**Target**

☒ Adapter ☐ VM

**Warning**  
If **VM** is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input checked="" type="checkbox"/> | Guest-VLAN           | <input checked="" type="radio"/> |
| <input checked="" type="checkbox"/> | Guest-Vlan-1000      | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input type="checkbox"/>            | NFS-VLAN             | <input type="radio"/>            |

[+ Create VLAN](#)

MTU:

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

**OK** **Cancel**

65. Create vNIC Template window opens.
66. Enter NFS-A-NIC in the Name field.
67. Enter NFSTemplate in the Description field.
68. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
69. In **Target** check **Adapater** check box.
70. In **Template Type** select **Initial Template** radio button.
71. Under **VLANs** check **NFS-VLAN** and choose **Native VLAN** as NFS-VLAN.
72. In the **MTU** field enter 9000.
73. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
74. Select **Gold-Net-SLA** in **QoS Policy** list box.
75. Select **default** in **Network Control Policy** list box.

76. Select **<not set>** in **Pin Group** list box.
77. Select **default** in **Stats Thershold Policy** list box.
78. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
79. Click **OK**.

**Figure 79** Defining the vNIC Properties for NFS- A-NIC

**Create vNIC Template**

Name: **NFS-A-NIC**

Description: **NFS Template**

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

**Target**

☒ Adapter  
☐ VM

**Warning**  
If **VM** is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

**VLANs**

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input type="checkbox"/>            | Managment-VLAN       | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | NFS-VLAN             | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input type="checkbox"/>            | vMotion-VLAN         | <input type="radio"/>            |

+ Create VLAN

MTU: **9000**

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: **Platinum-Compute-MAC(15/...**

QoS Policy: **Gold-Net-SLA**

Network Control Policy: **<not set>**

Pin Group: **<not set>**

OK Cancel

80. **Create vNIC Template** window opens.
81. Enter **NFS-B-NIC** in the **Name** field.
82. Enter **NFSTemplate** in the **Description** field.
83. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
84. In **Target** check **Adapater** check box.

85. In **Template Type** select **Initial Template** radio button.
86. Under **VLANs** check **NFS-VLAN** and choose **Native VLAN** as NFS-VLAN.
87. In the **MTU** field enter 9000.
88. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
89. Select **Gold-Net-SLA** in **QoS Policy** list box.
90. Select **default** in **Network Control Policy** list box.
91. Select **<not set>** in **Pin Group** list box.
92. Select **default** in **Stats Thershold Policy** list box.
93. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
94. Click **OK**.

**Figure 80** Defining the vNIC Properties for NFS- B-NIC

**Create vNIC Template**

Name:

Description:

Fabric ID: ☐ Fabric A ☒ Fabric B ☐ Enable Failover

**Target**

☒ Adapter  
☐ VM

**Warning**  
If **VM** is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

**VLANs**

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input type="checkbox"/>            | Management-VLAN      | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | NFS-VLAN             | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input type="checkbox"/>            | vMotion-VLAN         | <input type="radio"/>            |

[+ Create VLAN](#)

MTU:

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

### Baremetal Host vNIC Template

Create two vNIC templates for handling PXE and Management LAN traffic and two vHBA templates for handling SAN traffic to run Baremetal RHEL 6.3 host on TenantA-BareMetal-Zone.

1. Create **vNIC Template** window opens.
2. Enter PXE-A-NIC in the Name field.
3. Enter PXE Template in the **Description** field.
4. In **Fabric ID** click **FabricA** radio button, do not check **Enable Failover** check box.
5. In **Target** check **Adapter** check box.
6. In **Template Type** select **Initial Template** radio button.
7. Under **VLANs** check **Baremetal-VLAN20** and choose **Native VLAN** as BareMetal-VLAN20.

8. In the **MTU** field enter 9000.
9. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
10. Select **Platinum-Net-SLA** in **QoS Policy** list box.
11. Select **default** in **Network Control Policy** list box.
12. Select **<not set>** in **Pin Group** list box.
13. Select **default** in **Stats Thershold Policy** list box.
14. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
15. Click **OK**.

**Figure 81** Defining the vNIC Properties for PXE- A-NIC

**Create vNIC Template**

Name: **PXE-A-NIC**

Description: **PXE Template**

Fabric ID: ☒ Fabric A ☐ Fabric B ☒ Enable Failover

**Target**

☒ Adapter  
☐ VM

**Warning**  
If **VM** is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

| Select                              | Name             | Native VLAN                      |
|-------------------------------------|------------------|----------------------------------|
| <input type="checkbox"/>            | default          | <input type="radio"/>            |
| <input checked="" type="checkbox"/> | BareMetal-VLAN20 | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | Guest-VLAN       | <input type="radio"/>            |
| <input type="checkbox"/>            | Guest-Vlan-1000  | <input type="radio"/>            |

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: **Platinum-Compute-MAC(15/...**

QoS Policy: **Platinum-Net-SLA**

Network Control Policy: **<not set>**

Pin Group: **<not set>**

MTU: **9000**

**OK** **Cancel**

16. **Create vNIC Template** window opens.

17. Enter Bare-MGMT-A-NIC in the Name field.
18. Enter Baremetal Management Template in the Description field.
19. In **Fabric ID** click **Fabric A** radio button, do not check **Enable Failover** check box.
20. In **Target** check **Adapter** check box.
21. In **Template Type** select **Initial Template** radio button.
22. Under **VLANs** check **Baremetal-VLAN20** and choose **Native VLAN** as BareMetal-VLAN20.
23. In the **MTU** field enter 9000.
24. Select **Platinum-Compute-MAC** pool from the **MAC Address Assignment** list box.
25. Select **GoldNet-SLA** in **QoS Policy** list box.
26. Select **default** in **Network Control Policy** list box.
27. Select **<not set>** in **Pin Group** list box.
28. Select **default** in **Stats Thershold Policy** list box.
29. Select **<not set>** in **Dynamic vNIC Connection Policy** list box.
30. Click **OK**.

**Figure 82** Defining the vNIC Properties for Bare-MGMT-A-NIC

**Create vNIC Template**

Name: Bare-MGMT-A-NIC

Description: Baremetal Management Template

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

Target: ☒ Adapter ☐ VM

**Warning**  
If VM is selected, a port profile by the same name will be created.  
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

| Select                              | Name                 | Native VLAN                      |
|-------------------------------------|----------------------|----------------------------------|
| <input checked="" type="checkbox"/> | Management-VLAN      | <input checked="" type="radio"/> |
| <input type="checkbox"/>            | NFS-VLAN             | <input type="radio"/>            |
| <input type="checkbox"/>            | VLAN-Guest-VLAN-1001 | <input type="radio"/>            |
| <input type="checkbox"/>            | vMotion-VLAN         | <input type="radio"/>            |

[+ Create VLAN](#)

MTU: 9000

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: Platinum-Compute-MAC(15/...

QoS Policy: Gold-Net-SLA

Network Control Policy: <not set>

Pin Group: <not set>

OK Cancel

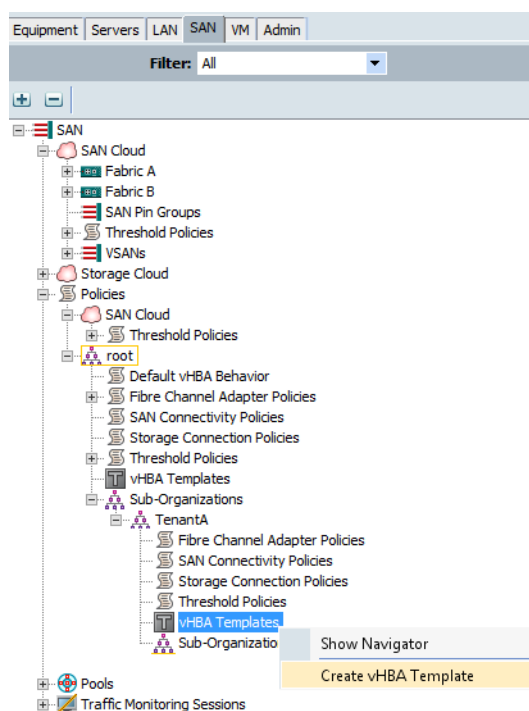
### Baremetal / Virtual Host vHBA Template

This section details the creation of two vHBA templates for handling SAN traffic to run Baremetal RHEL 6.3 or Hypervisor ESXi 5.1 hosts on TenantA-BareMetal-Zone and TenantA zones respectively.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA::

1. Click the **SAN** tab in the left pane.
2. Choose **Pools > Sub-Organization**.
3. Expand **TenantA**.
4. Select **vHBA Templates** and right click **Create vHBA Template**.

**Figure 83**      *Creating vHBA Template*



5. Create **vHBA Template** window opens.
6. Enter TenantAFabricA in the Name field.
7. Enter TenantA vHBA Fabric A Template in the Description field.
8. In **Fabric ID** click **FabricA** radio button.
9. Select **VSAN300** in **VSAN** list box.
10. In **Template Type** select **Initial Template** radio button.
11. In the **Max Data field Size** field enter 2048.
12. Select **Platinum-Compute-WWPN** pool from the **WWPN Pool** list box.
13. Select **FC-Net-SLA** in **QoS Policy** list box.
14. Select **<not set>** in **Pin Group** list box.
15. Select **default** in **Stats Thershold Policy** list box.
16. Click **OK**.



**Figure 84**      *Defining the vHBA Properties for TenantAFabricA*

**Create vHBA Template**

Name:

Description:

Fabric ID: ☒ A ☐ B

Select VSAN:  + Create VSAN

Template Type: ☒ Initial Template ☐ Updating Template

Max Data Field Size:

WWPN Pool:

QoS Policy:

Pin Group:

Stats Threshold Policy:

17. **Create vHBA Template** window opens.
18. Enter **TenantAFabricA** in the **Name** field.
19. Enter **TenantA vHBA Fabric A Template** in the **Description** field.
20. In **Fabric ID** click **FabricA** radio button.
21. Select **VSAN200** in **VSAN** list box.
22. In **Template Type** select **Initial Template** radio button.
23. In the **Max Data field Size** enter 2048.
24. Select **Platinum-Compute-WWPN** pool from the **WWPN Pool** list box.
25. Select **FC-Net-SLA** in **QoS Policy** list box.
26. Select **<not set>** in **Pin Group** list box.
27. Select **default** in **Stats Thershold Policy** list box.
28. Click **OK**.

**Figure 85**      *Defining the vHBA Properties for TenantAFabricB*

## Creating IPMI Policy

Intelligent Platform Management Interface (IPMI) is an open standard technology that defines how administrators monitor system hardware and sensors, control system components and retrieve logs of important system events to conduct remote management and recovery. IPMI runs on the BMC (Baseboard Management Controller) of the server blade and thus operates independently of the operating system. Since IPMI operates independent of the operating system, when sending commands to the BMC over IP, it provides administrators with the ability to monitor, manage, diagnose and recover systems, even if the operating system has hung or the server is powered down.



### Note

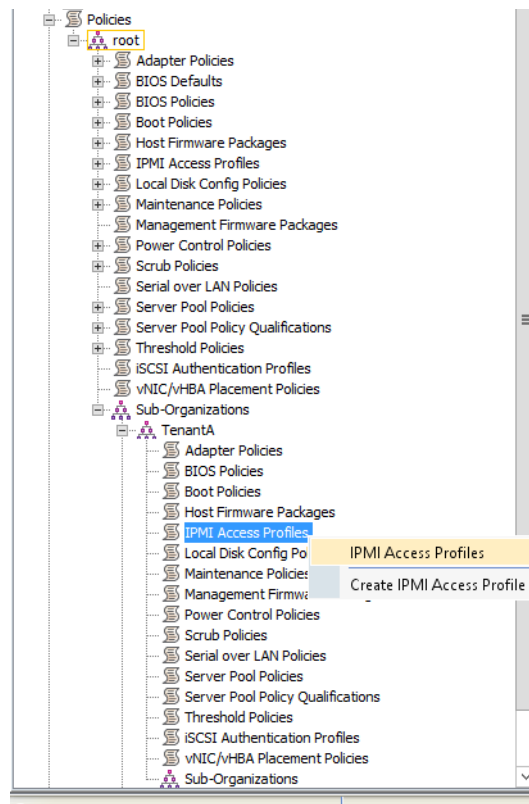
In the UCS system, the BMC supports IPMI version 2.0 only.

In this study we will create IPMI Access policy in Cisco UCS Manager for CloudPlatform 4.2.1 to perform Baremetal host operations like reboot, shutdown, resets and life cycle management in Cloud for Tenants.

Login to Cisco UCS Manager with root user, to configure TenantA Organization Unit:

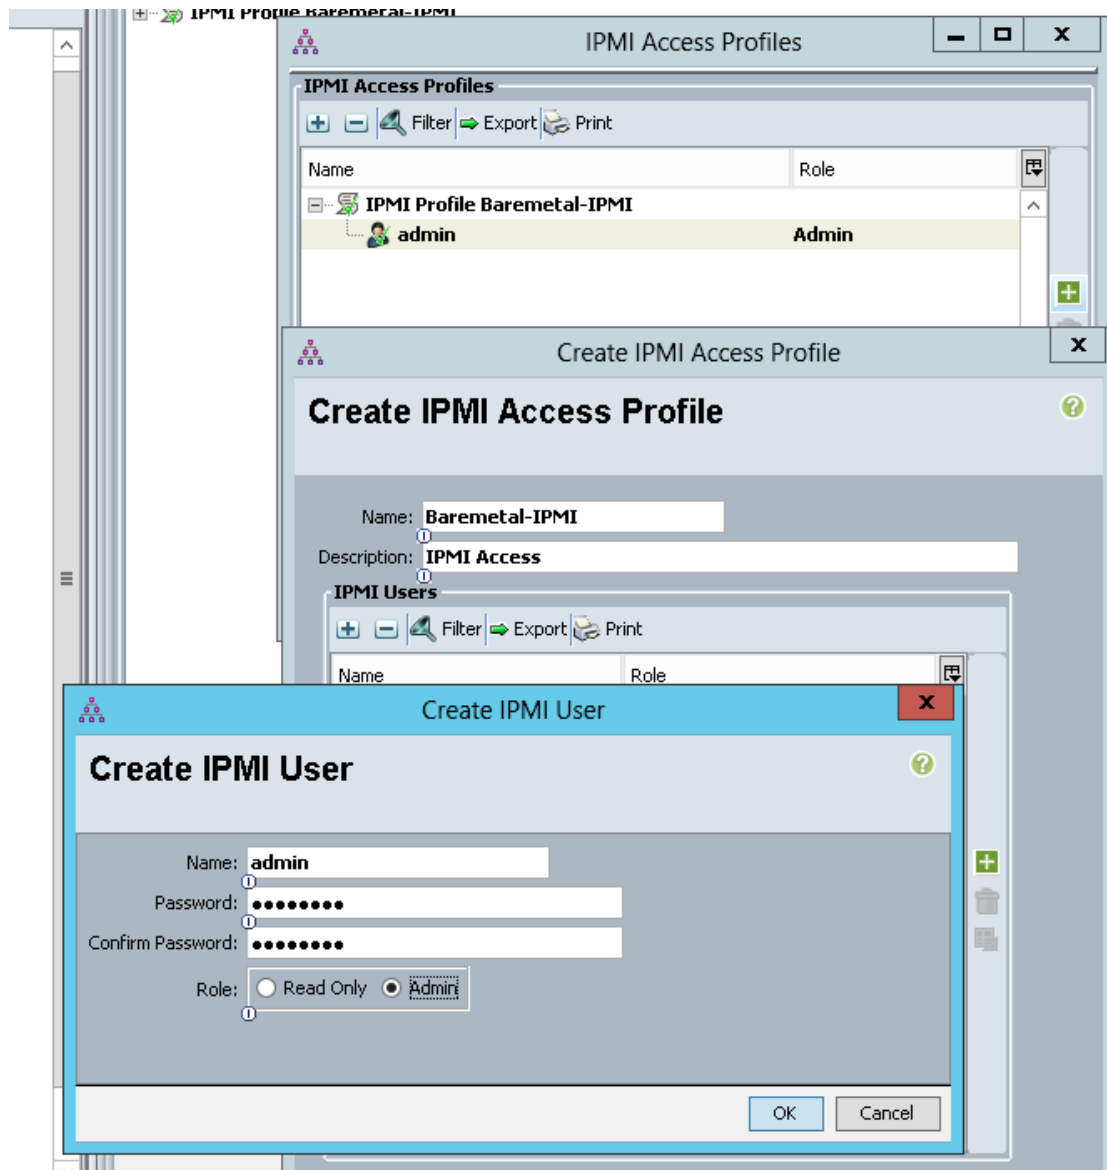
1. Click the **Policies servers** tab in the left pane.
2. Select **Policies**, expand **root** and then **Sub-Organization**.
3. Expand **TenantA**.
4. Right-click on **IPMI Access Profile** and select **Create IPMI Access Profiles**.

**Figure 86**      **Creating the IPMI Access Profile**



5. Click **Add** in the **IPMI Access Profiles** window.
6. Enter **Baremetal-IPMI** in the Name field.
7. Enter **IPMI Access** in the Description field.
8. Click **Add** in the **IPMI Access Profiles** window.
9. Enter **admin** in the Name field.
10. Enter **<XXXX>** in Password field.
11. Enter **<XXXX>** in confirm Password field.
12. Select **Admin** radio button in **Role** access.
13. Click **OK**.

**Figure 87** *Creating the IPMI Username and Password in Access Profile*



## Creating Server Pool Qualifications and Policy

A server pool contains a set of servers. These servers typically share the same characteristics, which can be their location in the chassis, or an attribute such as server type, amount of memory, local storage, type of CPU, or local drive configuration. A server can be assigned manually to a server pool, or the server pool policies and server pool policy qualifications can be used to automate the assignment.

To meet compute service levels based on the cloud services offering which includes cost, UCS offers a server pool that enables the Cloud administrator to designate one or more server pools to be used by a specific organization. UCS enables your business to meet the growth demands by facilitating automatic placement of these compute resources to specific server pools by defining Server Pool Qualifications parameters. These parameters include physical CPU and Memory types based on performance and capacity available on the compute blades.

In this study we will create four Server Pools with Policy qualifications based on cloud compute offering such as the Memory, CPU, and other attributes, which are automatically assigned to the concerned server pools as and when they are available on cloud compute system.

1. Platinum Server Pool
2. Gold Server Pool
3. Silver Server Pool
4. Bronze Server Pool

The steps described in this section provide details to configure the necessary Server Pool Qualifications for Compute-Server-Pool in the Cisco UCS environment.

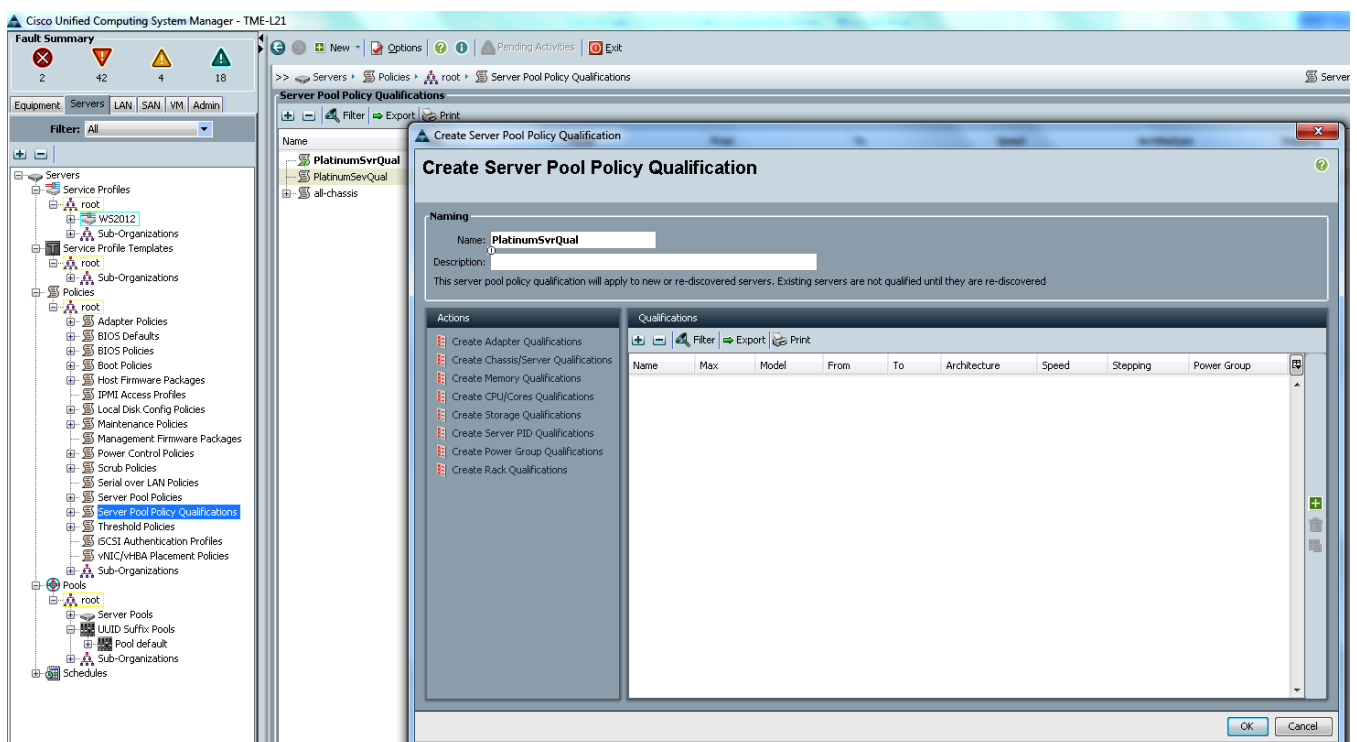
The Server Pool qualification data mentioned in this paper are based on general usage and can be altered to choose your specific compute needs in the cloud environment

### Platinum Compute Server Pool and Qualification

Login to Cisco UCS Manager with root user, to configure TenantA Organization Unit:

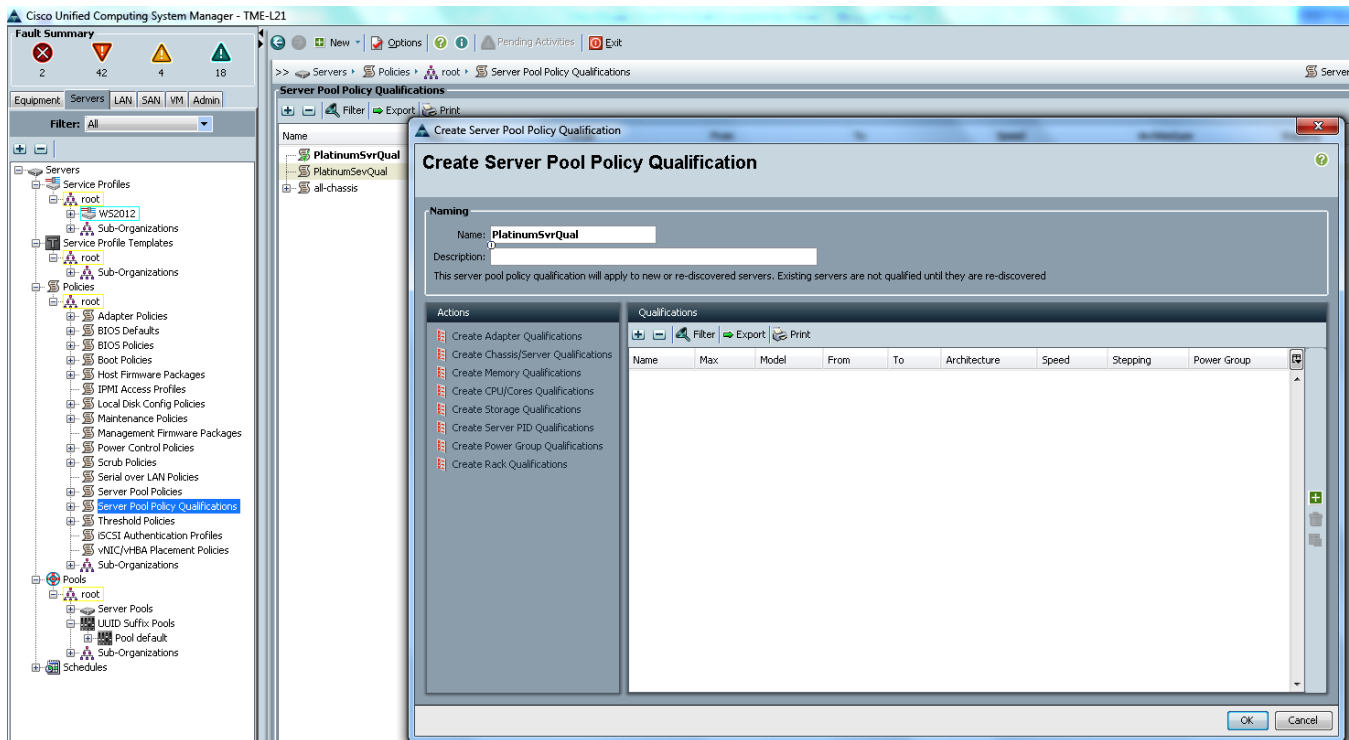
1. Click the **Servers** tab in the left pane.
2. Select **Policies**, and expand the **root**.
3. Right-click **Server Pool Policy Qualifications** and select **Create Server Pool Policy Qualification**.

**Figure 88** Creating the Server Pool Policy Qualification



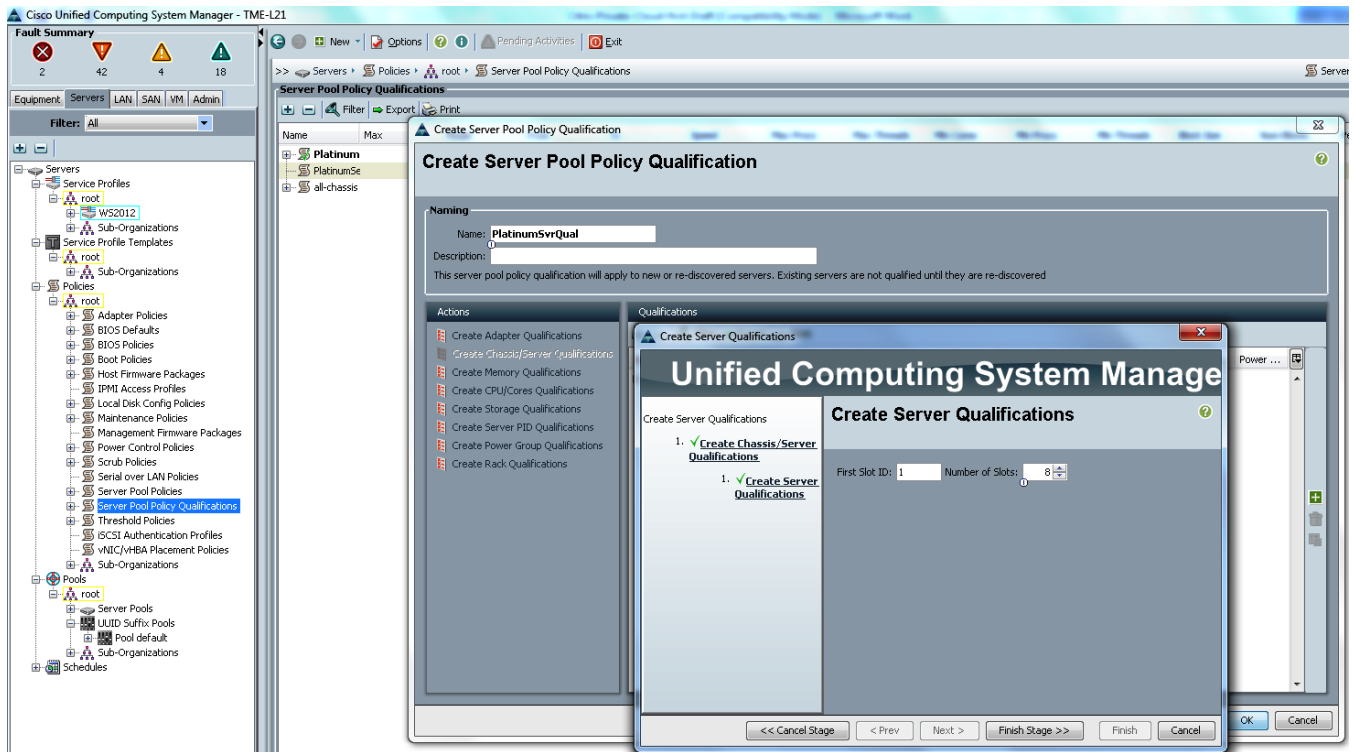
4. Name the server pool policy qualification PlatinumSvrQual.
5. (Optional) Give description.

**Figure 89** Adding the Name and Description to the Server Pool Policy Qualification



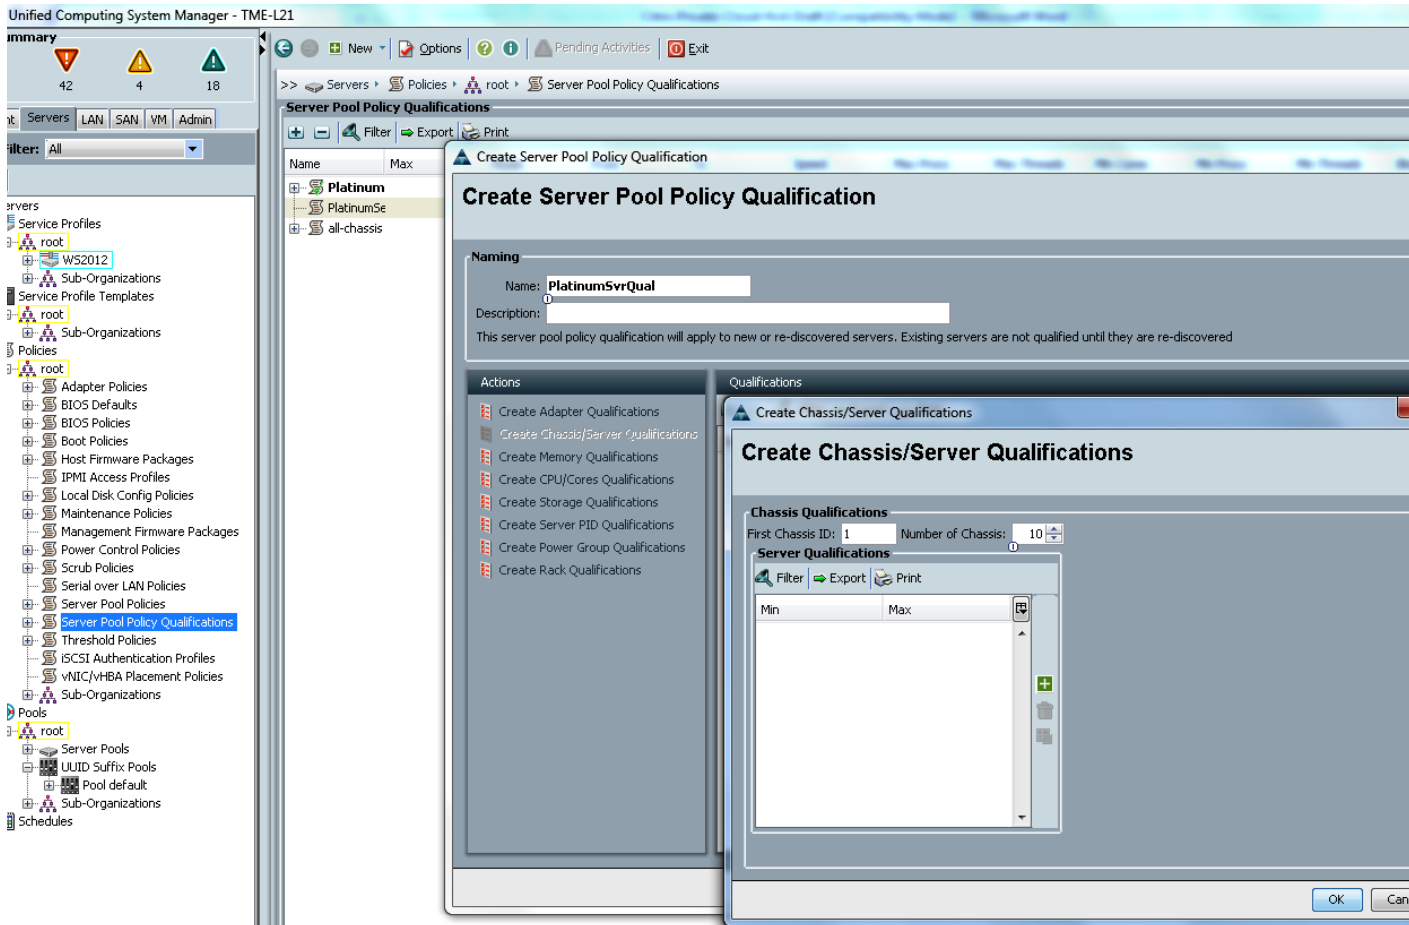
6. In the Actions pane, click **Create Chassis/Server Qualification**.
7. Type value 1 as the **First Chassis ID** and value 10 as the **Number of chassis**.

**Figure 90** Defining the ID and Slots for the Server Qualification



8. Under Server Qualifications click **ADD (+)**.

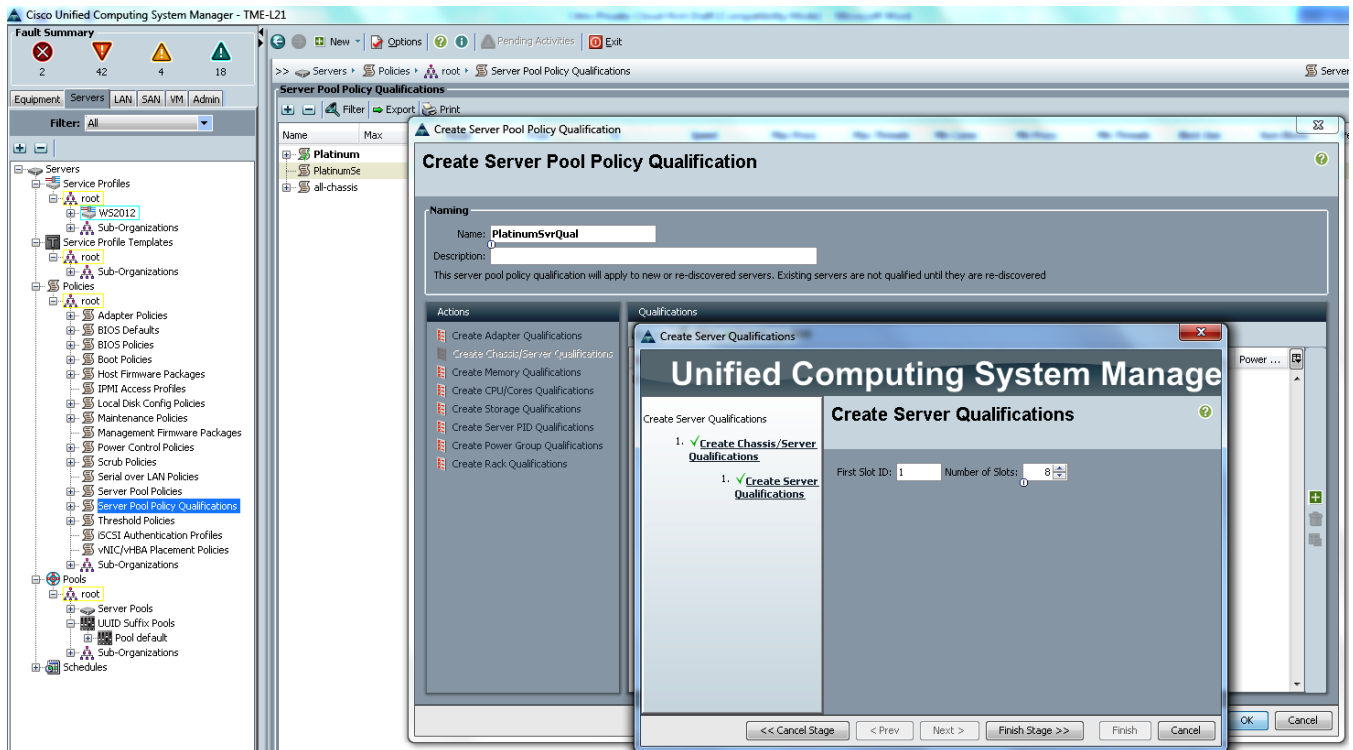
**Figure 91** Adding the Server Qualifications for Platinum Server



9. In First Slot ID enter 1 value, and Number of Slots enter 8.

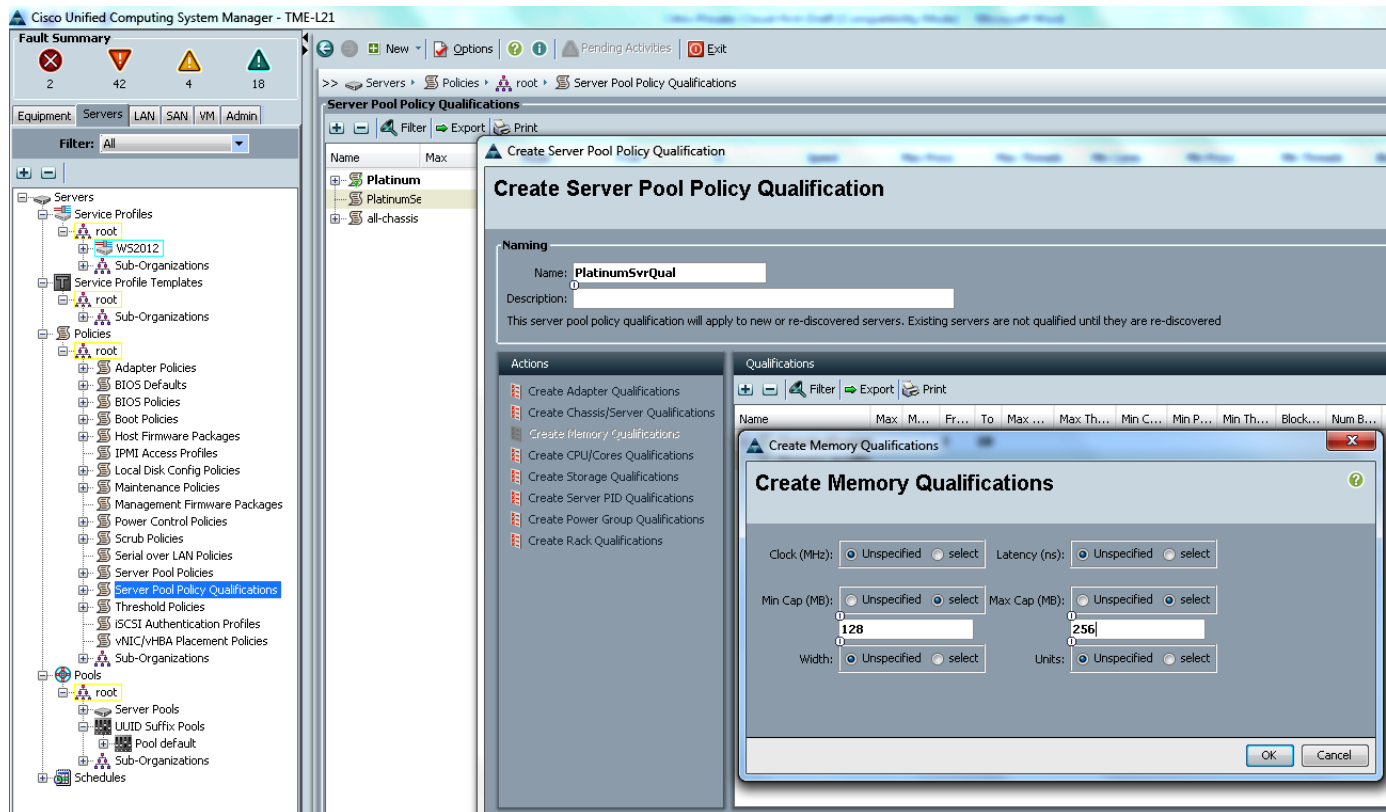


**Figure 92** Defining ID and Slots for Server Qualification

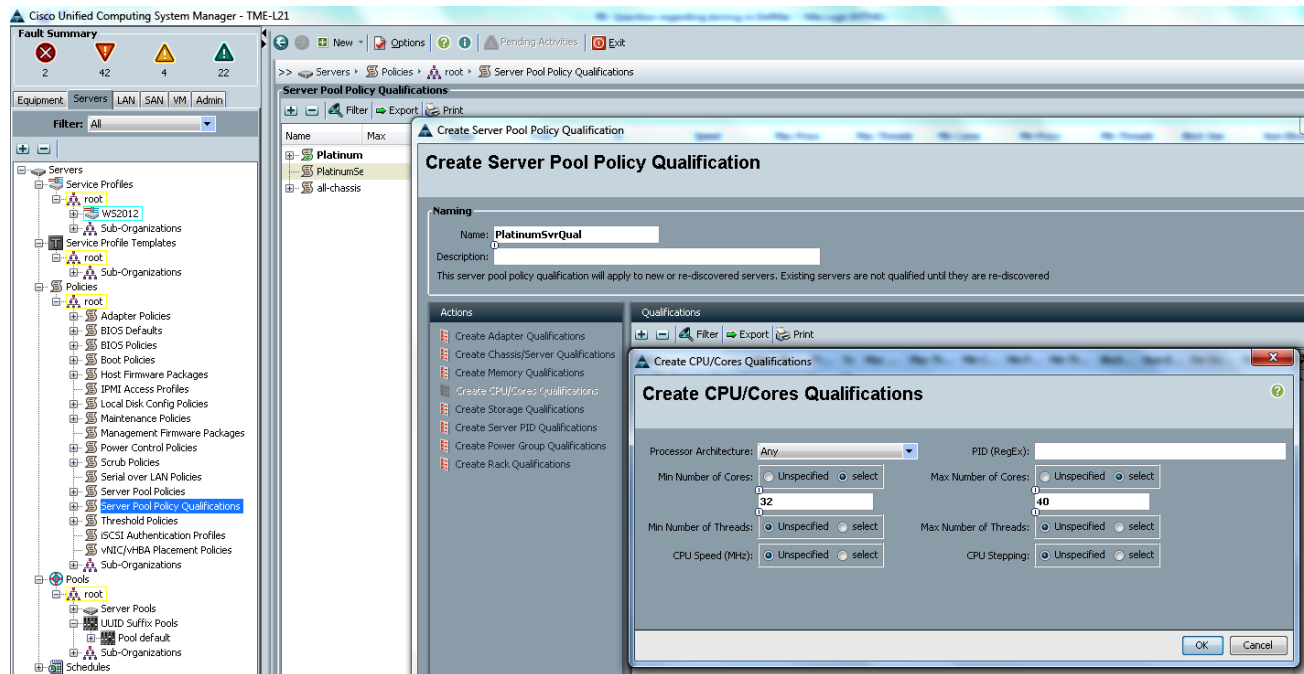


10. Click **Final Stage**, and then click **Finish**.
11. In the **Actions** pane, click **Create Chassis/Server Qualification**.
12. In the **First Chassis ID** field enter 1 and in **Number of Chassis** field enter 10.
13. In the **Actions** pane, click **Create Memory Qualification**.
14. For the **Clock (MHz)** field, click the **Unspecified** radio button.
15. For the **Latency (ns)** field, click the **Unspecified** radio button.
16. For the **Min Cap (MB)** field, click **select** radio button, and enter 128
17. For the **Max Cap (MB)** field click **select** radio button enter 256.
18. For the **Width** field, click the **Unspecified** radio button.
19. For the **Units** field, click the **Unspecified** radio button.
20. Click **OK**.

**Figure 93** Defining the parameters for the Memory Qualifications



21. In the **Actions** pane click **Create CPU Qualifications**.
22. For the **Min Number of Cores** field click **Select** radio button and enter 32
23. For the **Max Number of Cores** field click **Select** radio button and enter 40.
24. Select **Any for Processor Architecture**, and leave the **PID (RegEx)** field blank
25. For the **Min Number of Threads** field click **Unspecified** radio button.
26. For the **Max Number of Threads** field click **Unspecified** radio button.
27. For the **CPU Speed (MHz)** field click **Unspecified** radio button.
28. For the **CPU Stepping** field, click **Unspecified** radio button.
29. Click **Finish**.

**Figure 94**      **Defining the Parameters for the CPU Qualifications**

Follow the steps described above to create Server Pool Policy Qualifications Gold-Compute-Server, Silver-Compute-Server and Bronze-Compute-Server with Server Pool Policy Qualifications listed in [Table 15](#).

**Table 15**      **Server Pool Policy Qualification Values for all the Service Classes**

| Server Pool Policy Qualifications | Chassis                                       | Memory                    | CPU                                             |
|-----------------------------------|-----------------------------------------------|---------------------------|-------------------------------------------------|
| PlatinumSvrQual                   | ID Range 1 – 10<br>Server Slot ID range 1 – 8 | Min Cap 128 – Max Cap 256 | Min Number of cores 32 – Max Number of cores 40 |
| GoldSvrQual                       | ID Range 1 – 10<br>Server Slot ID range 1 – 8 | Min Cap 64 – Max Cap 128  | Min Number of cores 32 – Max Number of cores 40 |
| SilverSvrQual                     | ID Range 1 – 10<br>Server Slot ID range 1 – 8 | Min Cap 32 – Max Cap 64   | Min Number of cores 12 – Max Number of cores 16 |
| BronzeSvrQual                     | ID Range 1 – 10<br>Server Slot ID range 1 – 8 | Min Cap 16 – Max Cap 32   | Min Number of cores 12 – Max Number of cores 16 |

**Figure 95** Summary of the Server Pool Policy Qualifications Created

The screenshot shows the Cisco Unified Computing System Manager interface. The left pane displays a tree view of the configuration hierarchy, with 'Server Pool Policy Qualifications' selected under 'Policies'. The main pane shows a table of qualifications.

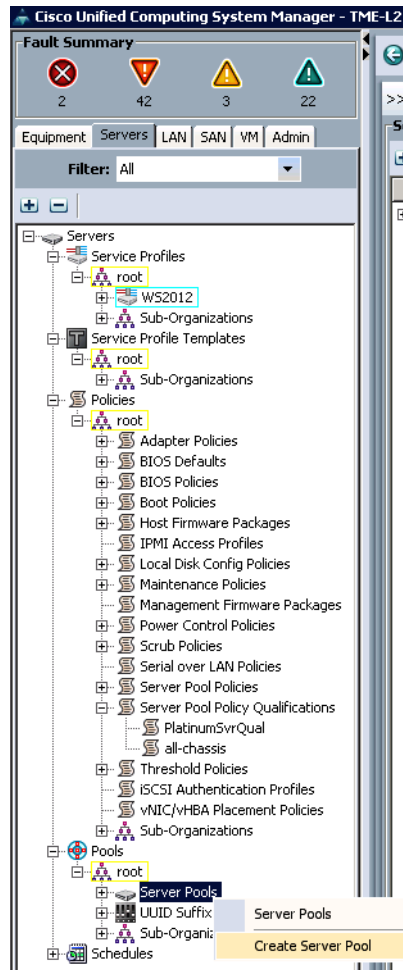
| Name                      | Max | Model | From | To | Speed       | Architecture | Stepping    |
|---------------------------|-----|-------|------|----|-------------|--------------|-------------|
| IronServerQual            |     |       |      |    |             |              |             |
| Chassis id range [1 - 10] |     |       | 1    | 10 |             |              |             |
| Memory qualification      |     |       |      |    | Unspecified |              |             |
| Processor qualification   |     |       |      |    | Unspecified | Any          | Unspecified |
| SilverSvrQual             |     |       |      |    |             |              |             |
| Chassis id range [1 - 10] |     |       | 1    | 10 |             |              |             |
| Memory qualification      |     |       |      |    | Unspecified |              |             |
| Processor qualification   |     |       |      |    | Unspecified | Any          | Unspecified |
| GoldSvrQual               |     |       |      |    |             |              |             |
| Chassis id range [1 - 10] |     |       | 1    | 10 |             |              |             |
| Memory qualification      |     |       |      |    | Unspecified |              |             |
| Processor qualification   |     |       |      |    | Unspecified | Any          | Unspecified |
| PlatinumSvrQual           |     |       |      |    |             |              |             |
| Chassis id range [1 - 10] |     |       | 1    | 10 |             |              |             |
| Memory qualification      |     |       |      |    | Unspecified |              |             |
| Processor qualification   |     |       |      |    | Unspecified | Any          | Unspecified |
| all-chassis               |     |       |      |    |             |              |             |

## Platinum Compute Server Pool

Login to Cisco UCS Manager with root user to configure Server Pools for TenantA Organization Unit:

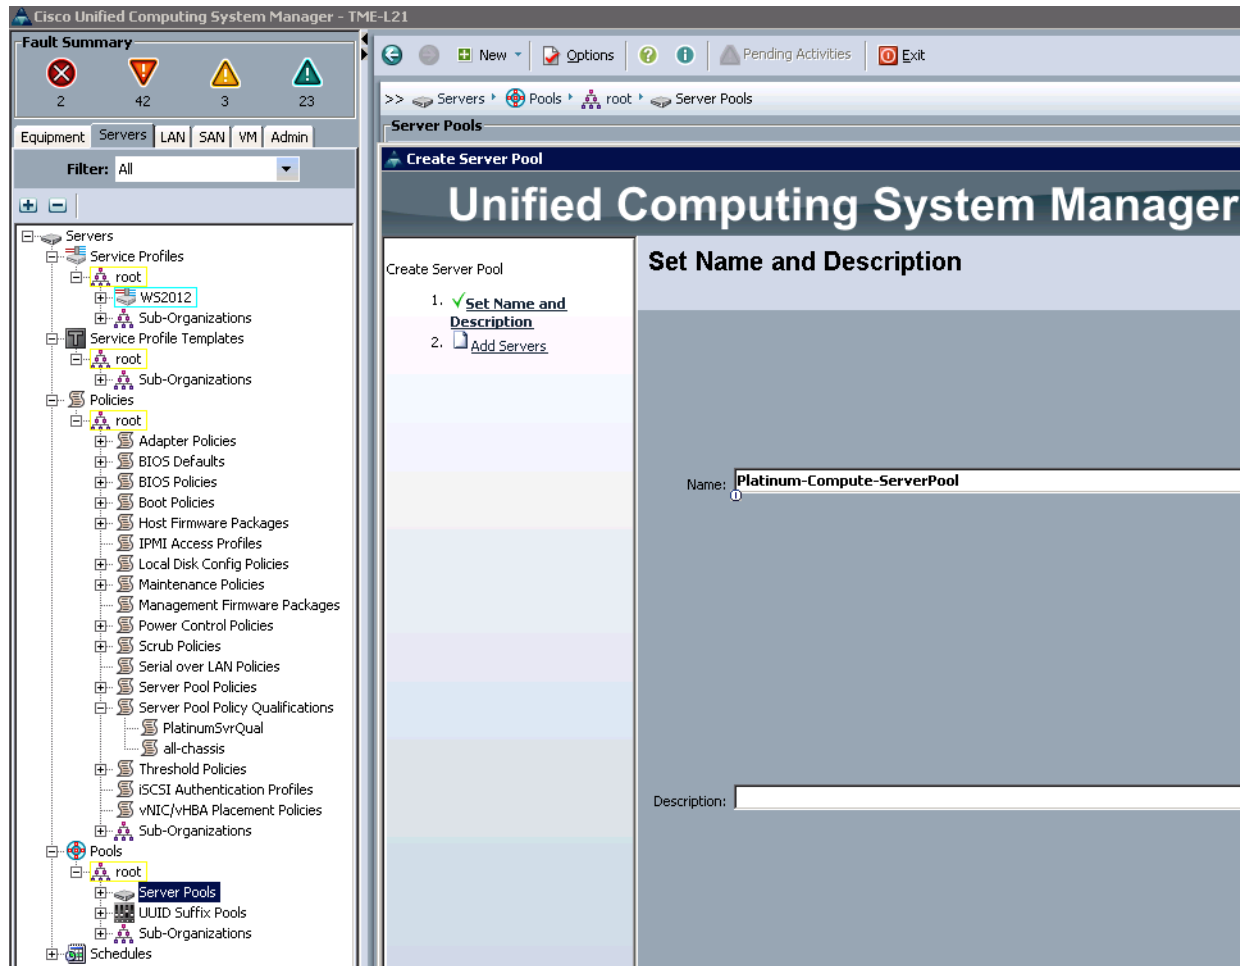
1. Click the **Servers** tab in the left pane.
2. Select **Pools**, expand **root**.
3. Right-click **Server Pool** and select **Create Server Pool**.

**Figure 96**      **Create Server Pool**



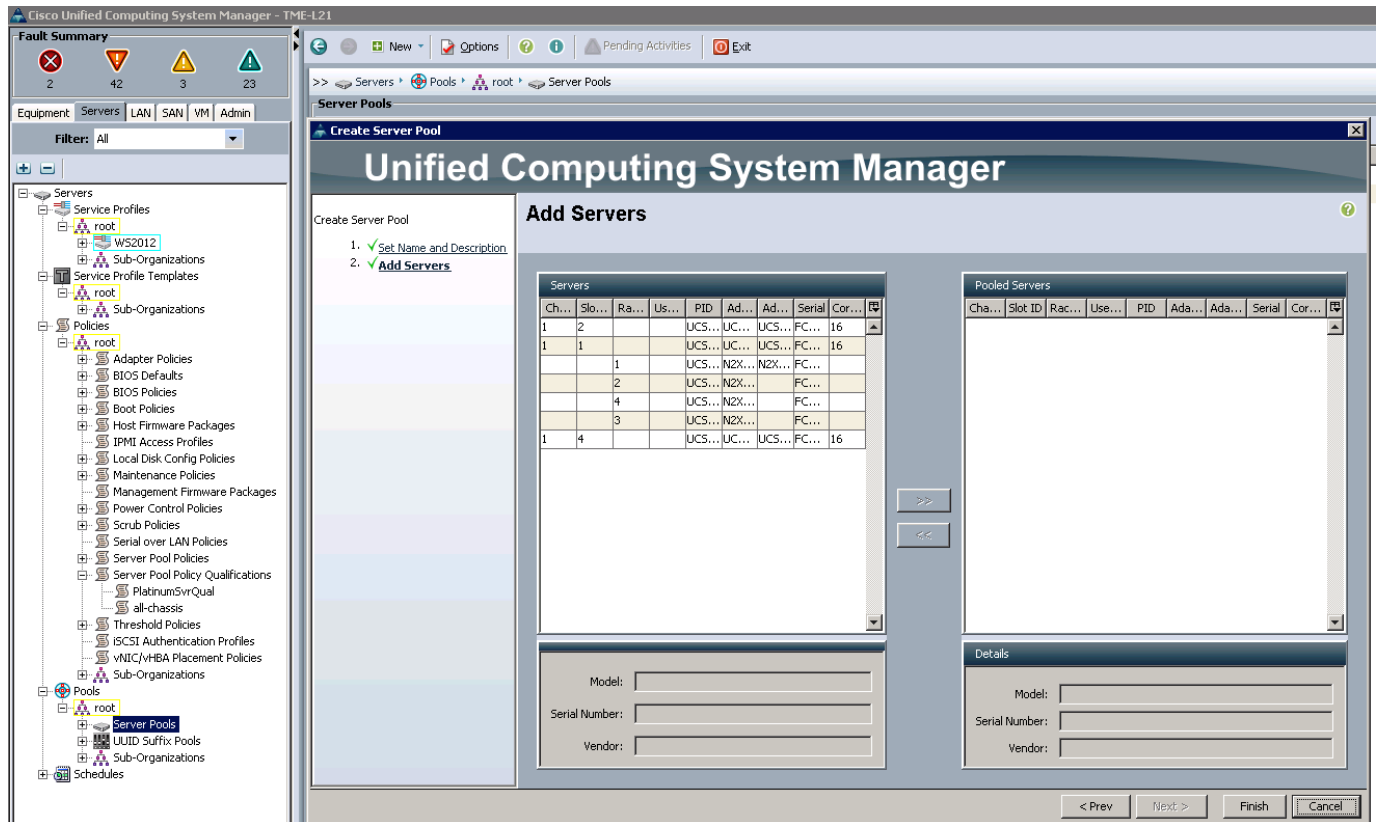
4. Enter Platinum-Compute-ServerPool in the Name field.
5. (Optional) Enter a description in the Description field.

**Figure 97** Adding the Name and Description for the Server Pool



6. Click Finish.

**Figure 98** Summary of the Servers in the Server Pool



Follow the steps described above to create Server Pools Gold-Compute-ServerPool, Silver-Compute-ServerPool and Bronze-Compute-ServerPool with the Server Pool Policy Qualifications Values listed in [Table 16](#)

**Table 16** Server Pool Policy Qualification values for the Server Pools

| Server Pool                 | Blades Values |
|-----------------------------|---------------|
| Platinum-Compute-ServerPool | 1 – 8         |
| Gold-Compute-ServerPool     | 1 – 8         |
| Silver-Compute-ServerPool   | 1 – 8         |
| Bronze-Compute-ServerPool   | 1 – 8         |

**Figure 99**      *Summary of all the Server Pools Created*

The screenshot shows the Cisco Unified Computing System Manager interface. The left pane displays a tree view of the system hierarchy, with 'Pools' selected. The main pane shows a table titled 'Server Pools' with columns for Name, Size, and Assigned. The table lists five server pools: Server Pool Bronze-Compute-ServerPool, Server Pool Silver-Compute-ServerPool, Server Pool Gold-Compute-ServerPool, Server Pool Platinum-Compute-ServerPool, and Server Pool default. The 'Server Pool default' has a size of 3 and is assigned 3 units.

| Name                                    | Size | Assigned |
|-----------------------------------------|------|----------|
| Server Pool Bronze-Compute-ServerPool   | 0    | 0        |
| Server Pool Silver-Compute-ServerPool   | 0    | 0        |
| Server Pool Gold-Compute-ServerPool     | 0    | 0        |
| Server Pool Platinum-Compute-ServerPool | 0    | 0        |
| Server Pool default                     | 3    | 3        |

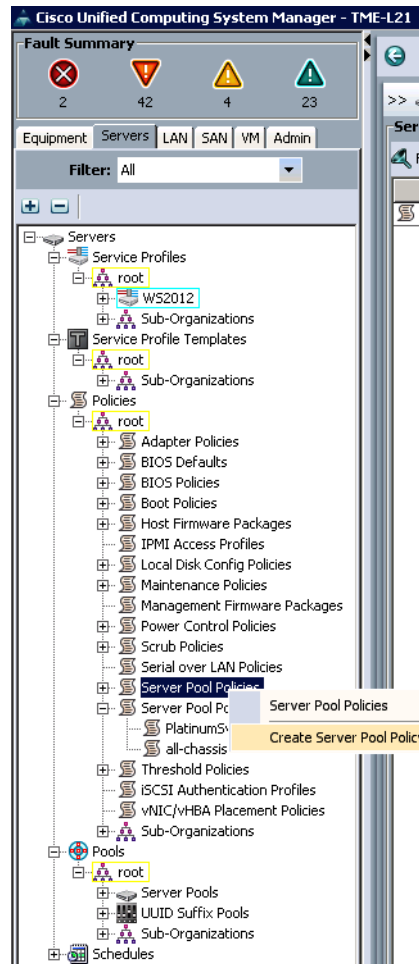
### Platinum Compute Server Pool Policies

Login to Cisco UCS Manager with root user to configure Server Pool Policies for TenantA Organization Unit:

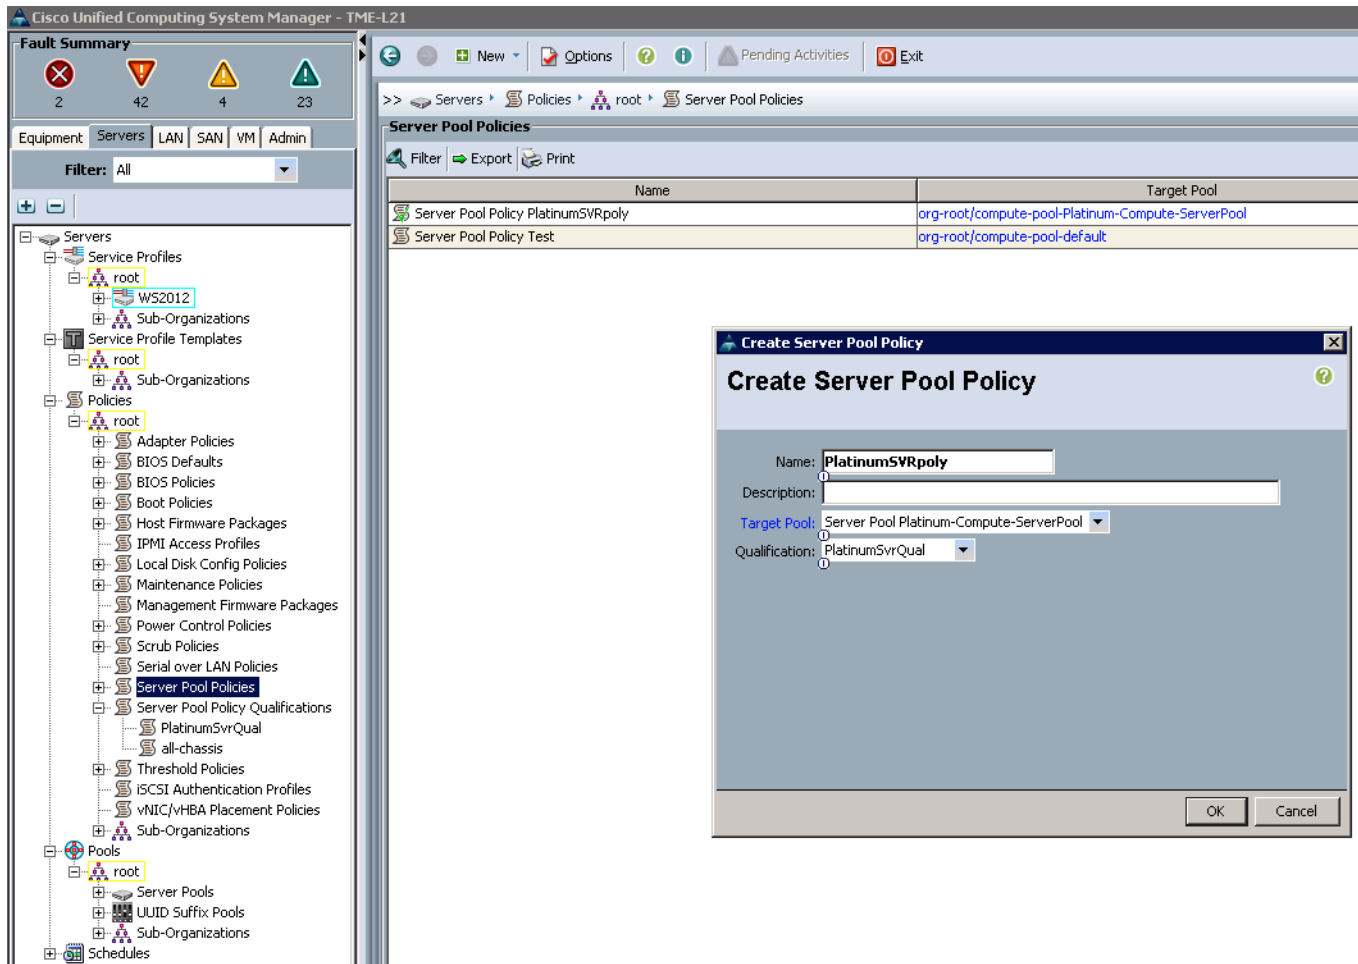
1. Click the **Servers** tab in the left pane.
2. Select **Policies**, expand **root**.
3. Right-click **Server Pool Policies** and select **Create Server Pool Policy**.



**Figure 100**      **Creating Server Pool Policies**



4. Enter **PlatinumSVRPply** in the Name field.
5. (Optional) Enter a description in the Description field.
6. In the **Target Pool** list box, select **Server Pool Platinum-Compute-ServerPool** (which was created earlier).
7. In the **Qualification** list box, select **PlatinumSvrQual** (which was created earlier).

**Figure 101** Defining the Server Pool Policy Attributes

Follow the steps described above to create Server Pools Policy Gold-Compute-ServerPool, Silver-Compute-ServerPool and Bronze-Compute-ServerPool with the Server Pool Policies listed in Table 17.

**Table 17** Server Pool Policy Values

| Server Pool Policies | Target Pool                             | Qualification   |
|----------------------|-----------------------------------------|-----------------|
| PlatinumSVRpoly      | Server Pool Platinum-Compute-ServerPool | PlatinumSvrQual |
| GoldSVRpoly          | Server Pool Gold-Compute-ServerPool     | GoldSvrQual     |
| SilverSVRpoly        | Server Pool Silver-Compute-ServerPool   | SilverSvrQual   |
| BronzeSVRpoly        | Server Pool Bronze-Compute-ServerPool   | BronzeSvrQual   |

**Figure 102** Summary of all the Server Pool Policies Created

The screenshot shows the Cisco Unified Computing System Manager (TME-421) interface. The left-hand navigation tree is expanded to 'Policies' > 'Server Pool Policies'. The main content area displays a table titled 'Server Pool Policies' with the following data:

| Name                              | Target Pool                                       | Qualification   |
|-----------------------------------|---------------------------------------------------|-----------------|
| Server Pool Policy BronzeSVRply   | org-root/compute-pool-Bronze-Compute-ServerPool   | BronzeSvrQual   |
| Server Pool Policy SilverSVRply   | org-root/compute-pool-Silver-Compute-ServerPool   | SilverSvrQual   |
| Server Pool Policy GoldSVRply     | org-root/compute-pool-Gold-Compute-ServerPool     | GoldSvrQual     |
| Server Pool Policy PlatinumSVRply | org-root/compute-pool-Platinum-Compute-ServerPool | PlatinumSvrQual |

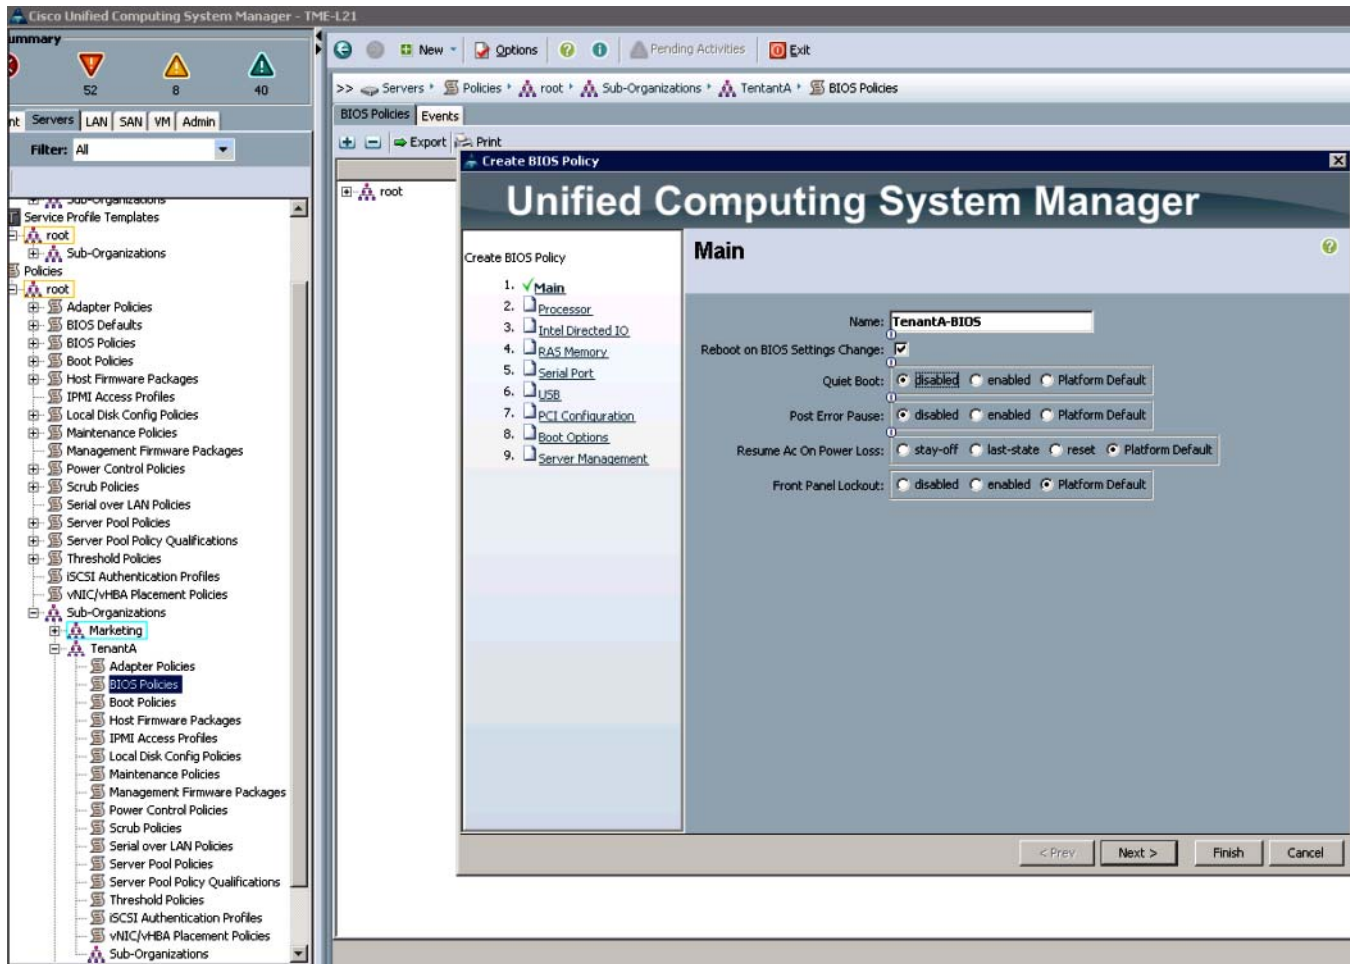
## Creating a Server BIOS Policy

These steps provide details for creating a server BIOS policy for the Cisco UCS environment.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA::

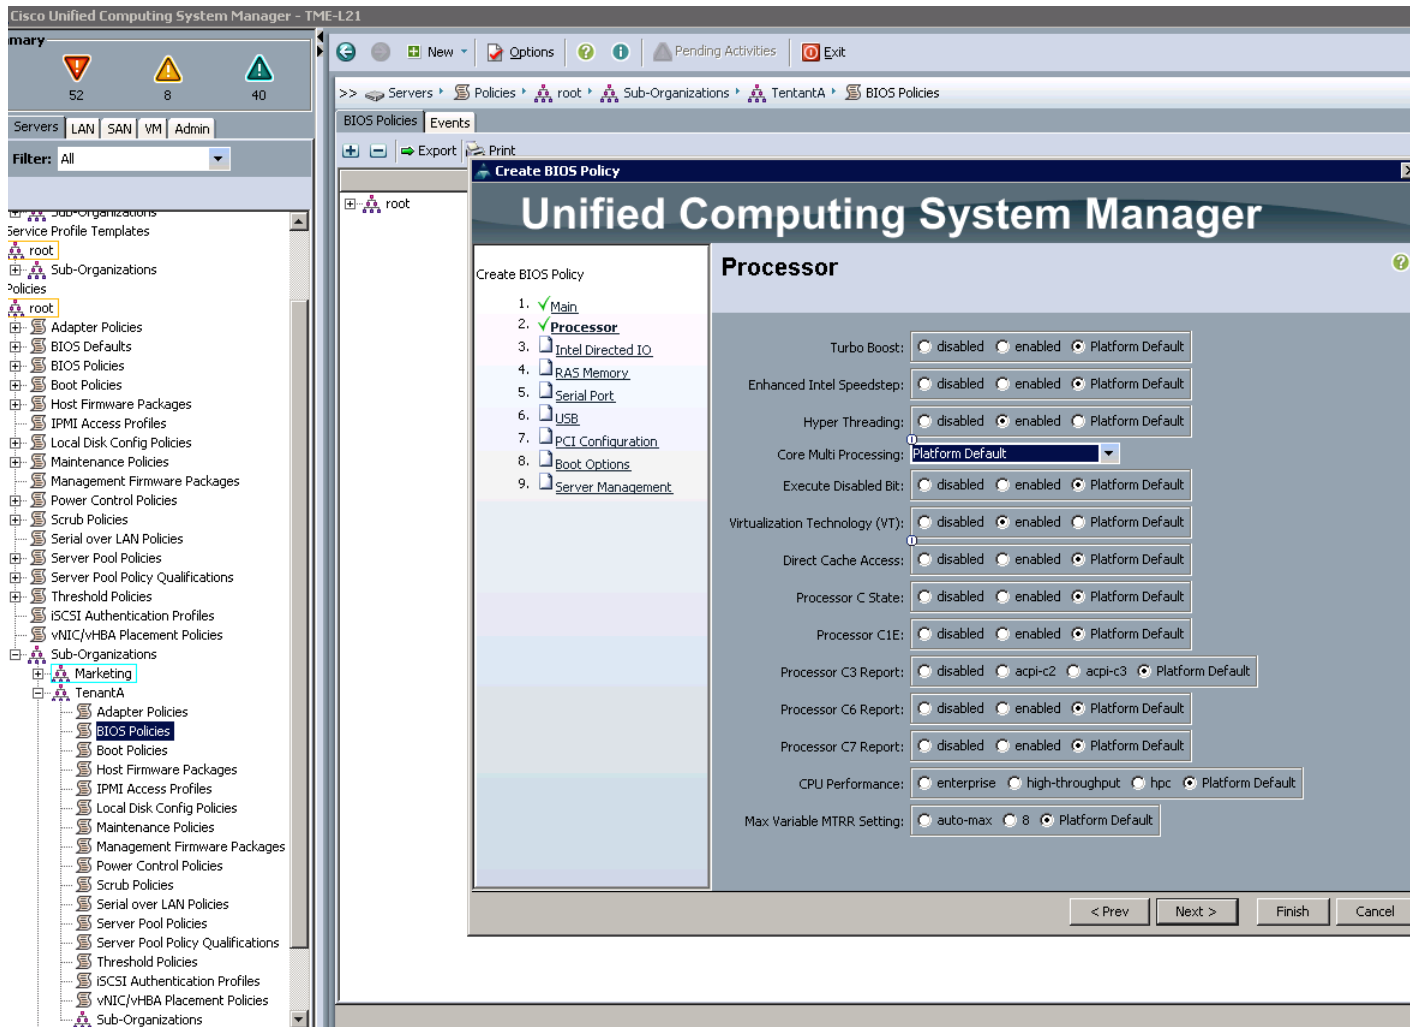
8. Click the **Server** tab.
9. Choose **Pools > Sub Organization**.
10. Expand TenantA and select **BIOS Policies**.
11. Right-click **BIOS Policies** and select **Create BIOS Policy**.
12. Enter TenantA-BIOS as the BIOS policy name.
13. Check **Reboot on BIOS Setting Change** check box.
14. Change the **Quiet Boot** property to **Disabled**.

Figure 103 Defining the BIOS Policy Attributes



15. Click **Next**.
16. On **Hyper Threading** click **Enabled** radio button.
17. On **Virtualization Technology (VT)** click **Enabled** radio button.
18. Accept default setting and click on **Finish**.
19. Click **OK**.

**Figure 104**      **Setting the Processor Attributes**



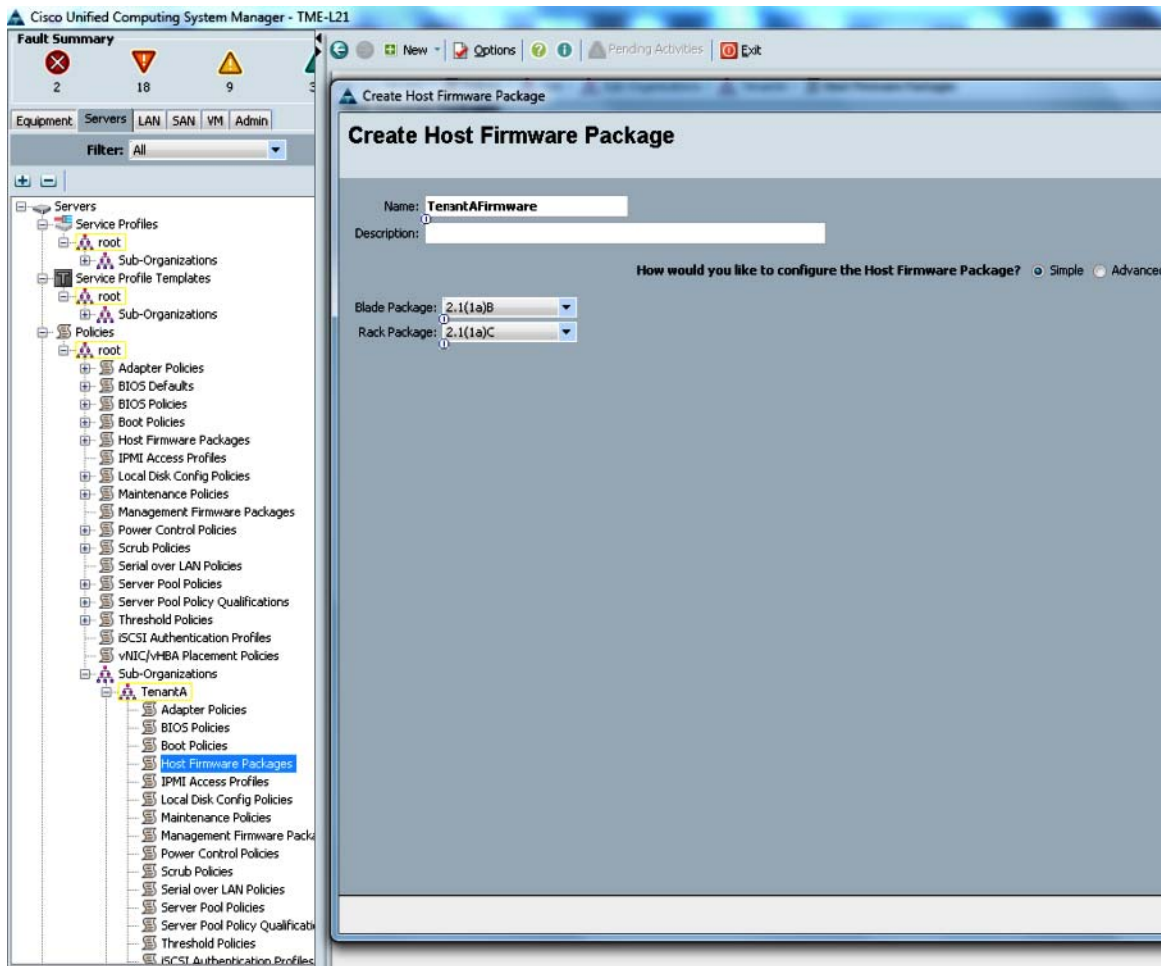
## Creating a Server Firmware Policy

These steps provide details for creating a server BIOS policy for the Cisco UCS environment.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA::

1. Click **Server** tab.
2. Choose **Pools > Sub Organization**.
3. Expand TenantA and select **Firmware Host Firmware Packages**.
4. Right-click **Host Firmware Packages** and select **Create**.
5. Enter TenantAFirmware as the BIOS policy name.
6. Click **Simple** radio button on How would you like to configure the Host Firmware Package.
7. On **Blade Package** and **Rack Package** select the latest Cisco UCS package.
8. Click **OK**.

**Figure 105**      *Configuring the Host Firmware Package*



## Creating Local Disk Configuration Policy

These steps provide details for creating a local disk configuration for the Cisco UCS environment, which is necessary if the servers in question do not have a local disk.



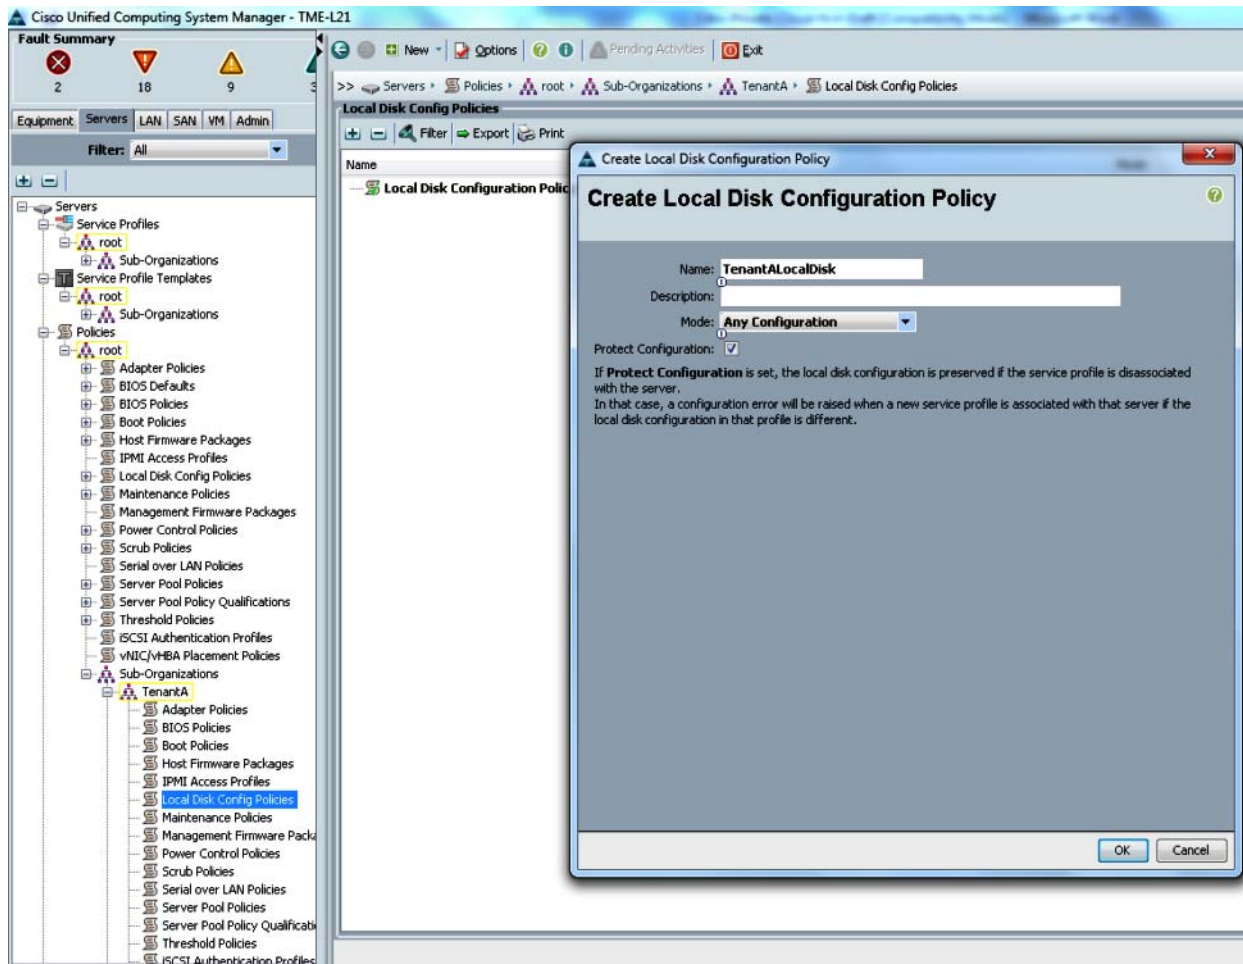
### Note

This policy should not be used on blades that contain local disks.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:

1. Click the **Server** tab.
2. Choose **Pools > Sub Organization**.
3. Expand TenantA and select **Local Disk Configuration Policy**.
4. Right-click **Local Disk Configuration Policy** and select **Create**.
5. Enter TenantALocalDisk as the local disk configuration policy name.
6. Change the **Mode** to **Any Configuration**, and check **Protect Configuration** check box.
7. Click **OK** to complete creating the host firmware package.

**Figure 106** *Creating the Local Disk Configuration Policy Properties*



## Creating Boot Policies

These steps provide details for creating boot policies for the Cisco UCS environment. These directions apply to an environment in which each storage controller 0c port is connected to fabric A and each storage controller 0d port is connected to fabric B.



### Note

In this paper we will create two Boot Policies one for ESXi 5.1 Host and second for Baremetal RHEL 6.3 Host.

The first boot policy will define SAN Boot options for booting ESXi 5.1 over FCoE Target

The second boot policy will define PXE Boot options for booting Baremetal RHEL 6.3 over FCoE Target  
SAN Boot Policy for ESXi 5.1 Host.

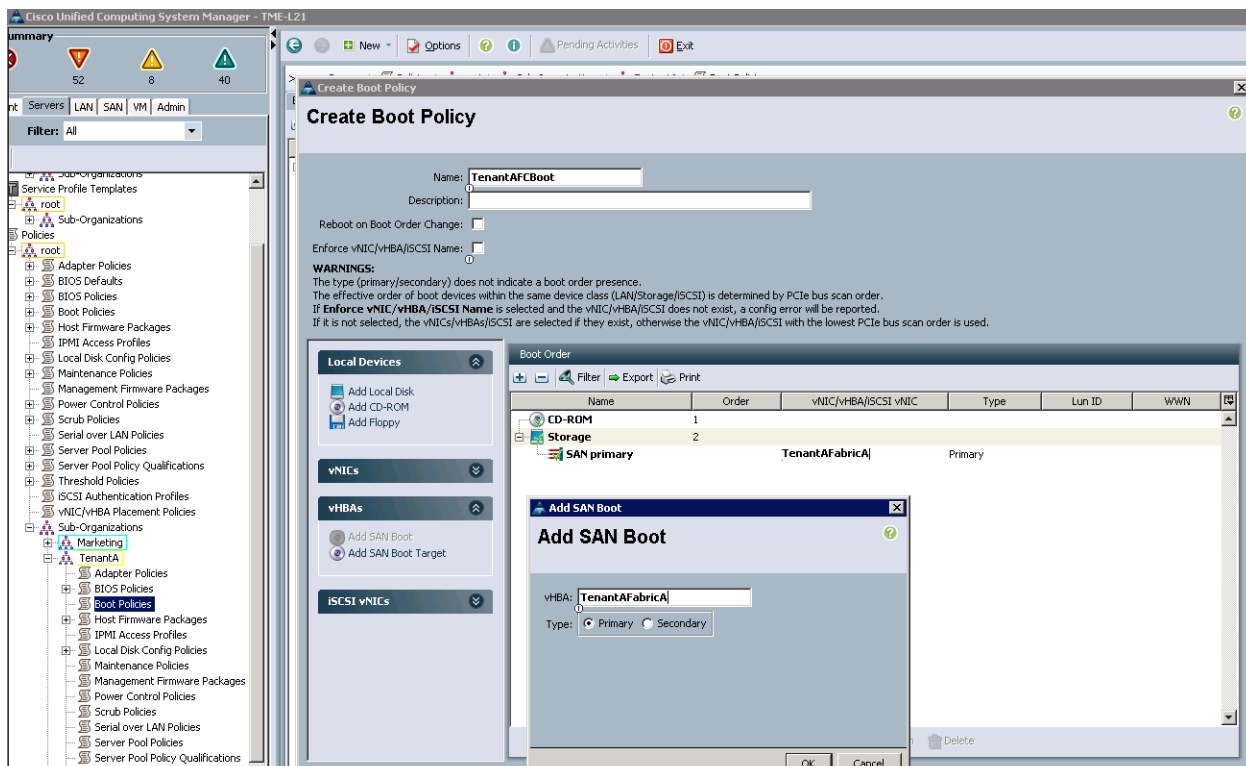
Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA::

1. Click the **Server** tab.
2. Select **Pools > Sub Organization**.



3. Expand TenantA and select the **Policies** tab in the right pane.
4. Click the **Boot Policies** sub-tab.
5. Click the + symbol at the right edge of the right pane to create a Boot Policy.
6. Enter TenantAFCBoot in the **Name** field.
7. (Optional) Give the boot policy a description.
8. Leave **Reboot on Boot Order Change** unchecked and uncheck **Enforce vNIC/vHBA Name**.
9. Expand the **Local Devices** drop-down list menu and select **Add CD-ROM**.
10. Expand **vHBAs** drop-down list menu and select **Add SAN Boot**.
11. Enter TenantAFabricA in the **vHBA** field of the **Add SAN Boot** window.
12. Ensure that **Primary** is selected as the type.
13. Click **OK** to add the SAN boot initiator.

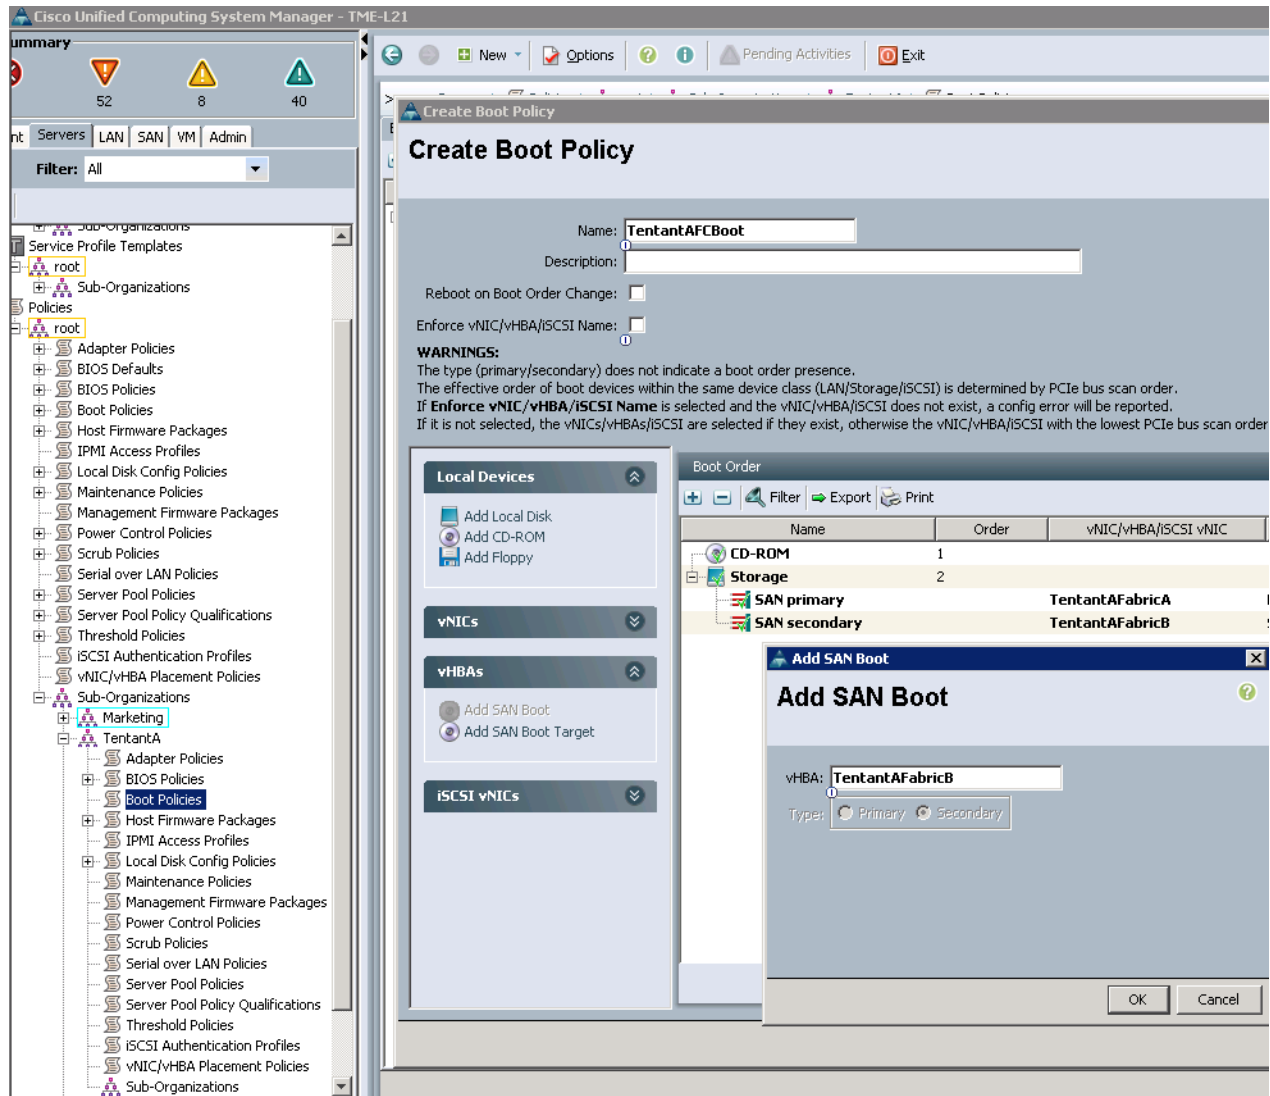
**Figure 107** Defining the vHBA Attributes in Primary SAN Boot



14. Click **Add SAN Boot** for the SAN boot initiator for secondary FC Boot Path.
15. Enter TenantAFabricB in the **vHBA** field of the **Add SAN Boot** window.
16. Click **OK** to add the SAN boot initiator.

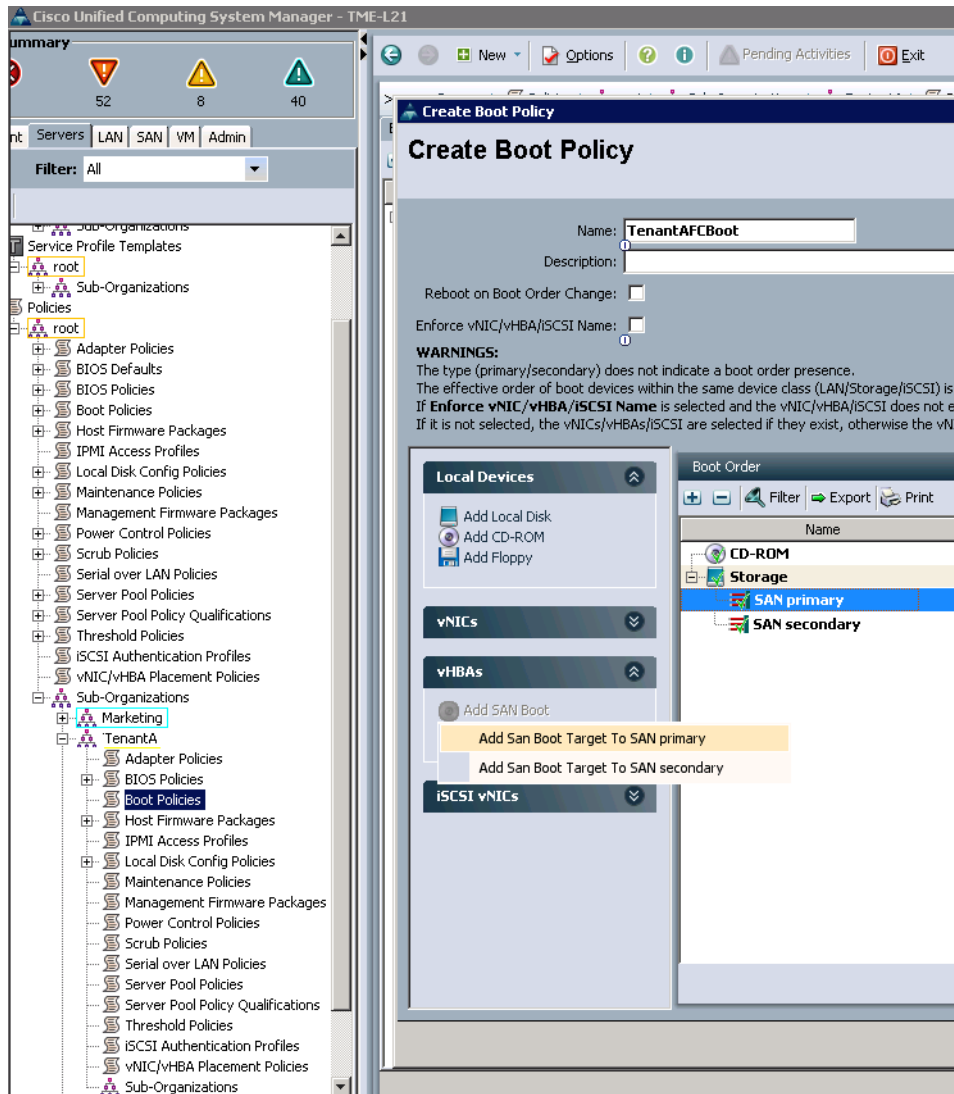


**Figure 108** Defining the vHBA attributes in Secondary SAN Boot



17. Under Boot Order in right pane select SAN Primary under Storage tree

18. From the vHBA dropdown menu, choose **Add SAN Boot > Add San Boot Target To SAN Primary**.

**Figure 109** Adding the SAN Boot Target

19. Type the value as 0 for the Boot Target LUN.
20. Enter the WWPN for the primary FC adapter interface 0c of controller A. To obtain the WWPN, login to controller A and in the command line type network interface show command.
21. Be sure to use the FC portname for 0c and not the FC node name.
22. Keep the type as Primary.
23. Click **OK** to add the SAN boot target.

**Figure 110** Adding the Boot Target LUN and WWPN Values

**Create Boot Policy**

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☐

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is selected.

**Local Devices**

- Add Local Disk
- Add CD-ROM
- Add Floppy

**vNICs**

- Add LAN Boot

**vHBAs**

- Add SAN Boot
- Add SAN Boot Target

**iSCSI vNICs**

**Boot Order**

| Name               | Order | vNIC/vHBA/iSCSI vNIC |
|--------------------|-------|----------------------|
| Storage            | 1     |                      |
| SAN primary        |       | TenantAFabricA       |
| SAN Target primary |       |                      |

**Add SAN Boot Target**

Boot Target LUN:

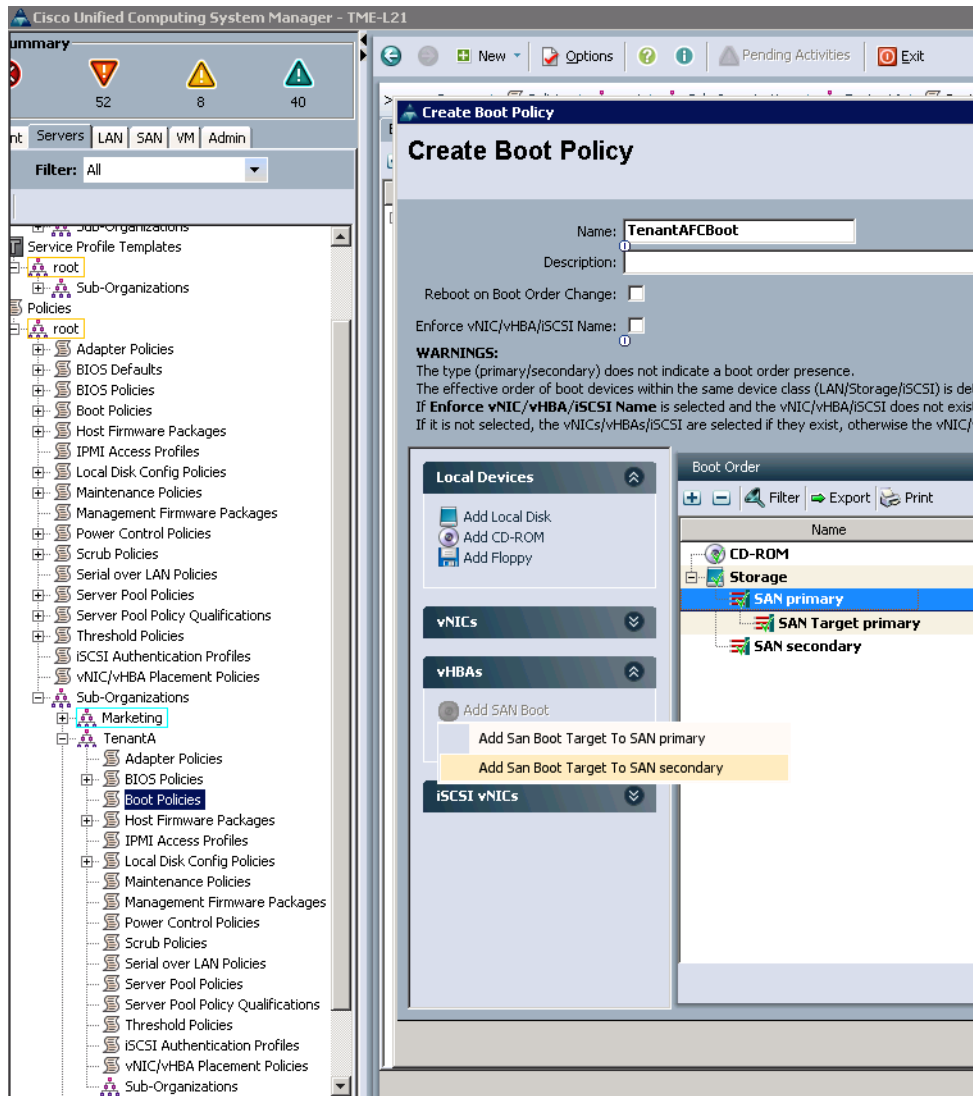
Boot Target WWPN:

Type: ☒ Primary ☐ Secondary

OK Cancel

24. Under Boot Order in right pane select SAN Primary under Storage tree.
25. From the vHBA drop down menu, choose **Add SAN Boot > Add San Boot Target to SAN Secondary**.

Figure 111 Adding the SAN Boot Target to Secondary SAN



26. Type the value 0 as the Boot Target LUN.
27. Enter the WWPN for the primary FC adapter interface 0c of controller B. To obtain WWPN log in to the controller B and in the command line type the network interface show command.

**Note**

Ensure to use the FC portname for port 0c and not the FC node name.

28. Click **OK** to add the SAN boot target.
29. Select **Add SAN Boot** under the vHBA drop-down menu.
30. Enter TenantAFabricB in the vHBA field in the Add SAN Boot window.
31. The type should automatically be set to Secondary and it should be grayed out. This is fine.
32. Click **OK** to add the SAN boot target.

**Figure 112** Defining the SAN Boot Secondary Target Properties

**Create Boot Policy**

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☐

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

- Add Local Disk
- Add CD-ROM
- Add Floppy

**vNICs**

- Add LAN Boot

**vHBAs**

- Add SAN Boot
- Add SAN Boot Target

**iSCSI vNICs**

**Boot Order**

| Name                 | Order | vNIC/vHBA/iSCSI vNIC | Type      |
|----------------------|-------|----------------------|-----------|
| Storage              | 1     |                      |           |
| SAN primary          |       | TenantAFabricA       | Primary   |
| SAN Target primary   |       |                      | Primary   |
| SAN Target secondary |       |                      | Secondary |

**Add SAN Boot Target**

Boot Target LUN:

Boot Target WWPN:

Type: ☒ Primary ☐ Secondary

OK Cancel

33. Under **Boot Order** in right pane select **SAN Secondary** under Storage tree.
34. From the vHBA drop down menu, choose **Add SAN Boot > Add San Boot Target To SAN primary**.
35. Type the value as 0 for Boot Target LUN.
36. Enter the WWPN for the primary FC adapter interface 0c of controller A. To obtain this information, log in to controller A and run the network interface show command.



**Note**

Ensure to use the FC portname for 0c and not the FC node name.

37. Keep the type as Primary.
38. Click **OK** to add the SAN boot target.

**Figure 113**      *Defining the SAN Boot Target Primary Properties*

**Create Boot Policy**

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☐

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

- Add Local Disk
- Add CD-ROM
- Add Floppy

**vNICs**

- Add LAN Boot

**vHBAs**

- Add SAN Boot
- Add SAN Boot Target

**iSCSI vNICs**

**Boot Order**

| Name                 | Order | vNIC/vHBA/iSCSI vNIC | Type      |
|----------------------|-------|----------------------|-----------|
| <b>Storage</b>       |       |                      |           |
| 1                    |       |                      |           |
| SAN primary          |       | TenantAFabricA       | Primary   |
| SAN Target primary   |       |                      | Primary   |
| SAN Target secondary |       |                      | Secondary |
| SAN secondary        |       | TenantAFabricB       | Secondary |
| SAN Target primary   |       |                      | Primary   |

**Add SAN Boot Target**

Boot Target LUN:

Boot Target WWPN:

Type: ☒ Primary ☐ Secondary

OK Cancel

39. Under **Boot Order** in right pane select **SAN Secondary** under **Storage** tab.
40. Choose **Add SAN Boot** > **Add San Boot Target To SAN secondary** from the vHBA menu.
41. Type the value as 0 for the Boot Target LUN.
42. Enter the WWPN for the primary FC adapter interface 0c of controller B. To obtain this information, log in to the controller B and run the network interface show command.



**Note**

Ensure to use the FC portname for port 0c and not the FC node name.

43. Click **OK** to add the SAN boot target.
44. Select **Add SAN Boot** under the vHBA drop-down menu.
45. Enter TenantAFabricB in the vHBA field in the **Add SAN Boot** window that displays.
46. The type should automatically be set to Secondary and it should be grayed out.
47. Click **OK** to add the SAN boot target.

**Figure 114** Defining SAN Boot Secondary Target Properties

**Create Boot Policy**

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☐

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

- Add Local Disk
- Add CD-ROM
- Add Floppy

**vNICs**

- Add LAN Boot

**vHBAs**

- Add SAN Boot
- Add SAN Boot Target

**iSCSI vNICs**

**Boot Order**

+ - Filter Export Print

| Name                 | Order    | vNIC/vHBA/iSCSI vNIC | Type      |
|----------------------|----------|----------------------|-----------|
| <b>Storage</b>       | <b>1</b> |                      |           |
| SAN primary          |          | TenantAFabricA       | Primary   |
| SAN Target primary   |          |                      | Primary   |
| SAN Target secondary |          |                      | Secondary |
| SAN secondary        |          | TenantAFabricB       | Secondary |
| SAN Target primary   |          |                      | Primary   |
| SAN Target secondary |          |                      | Secondary |

**Add SAN Boot Target**

Boot Target LUN:

Boot Target WWPN:

Type: ☒ Primary ☐ Secondary

OK Cancel

48. The final TenantAFCBoot policy.

Figure 115 TenantAFCBboot SAN Boot Policy Properties

**Create Boot Policy**

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☒

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

- Add Local Disk
- Add CD-ROM
- Add Floppy

**vNICs**

- Add LAN Boot

**vHBAs**

- Add SAN Boot
- Add SAN Boot Target

**iSCSI vNICs**

**Boot Order**

| Name                 | Order | vNIC/vHBA/iSCSI vNIC | Type      | Lun ID | WWN                     |
|----------------------|-------|----------------------|-----------|--------|-------------------------|
| CD-ROM               | 1     |                      |           |        |                         |
| Storage              | 2     |                      |           |        |                         |
| SAN primary          |       | TenantAFabricA       | Primary   |        |                         |
| SAN Target primary   |       |                      | Primary   | 0      | 20:05:00:A0:98:37:A0:70 |
| SAN Target secondary |       |                      | Secondary | 0      | 20:08:00:A0:98:37:A0:70 |
| SAN secondary        |       | TenantAFabricB       | Secondary |        |                         |
| SAN Target primary   |       |                      | Primary   | 0      | 20:06:00:A0:98:37:A0:70 |
| SAN Target secondary |       |                      | Secondary | 0      | 20:09:00:A0:98:37:A0:70 |

Move Up Move Down Delete

OK Cancel

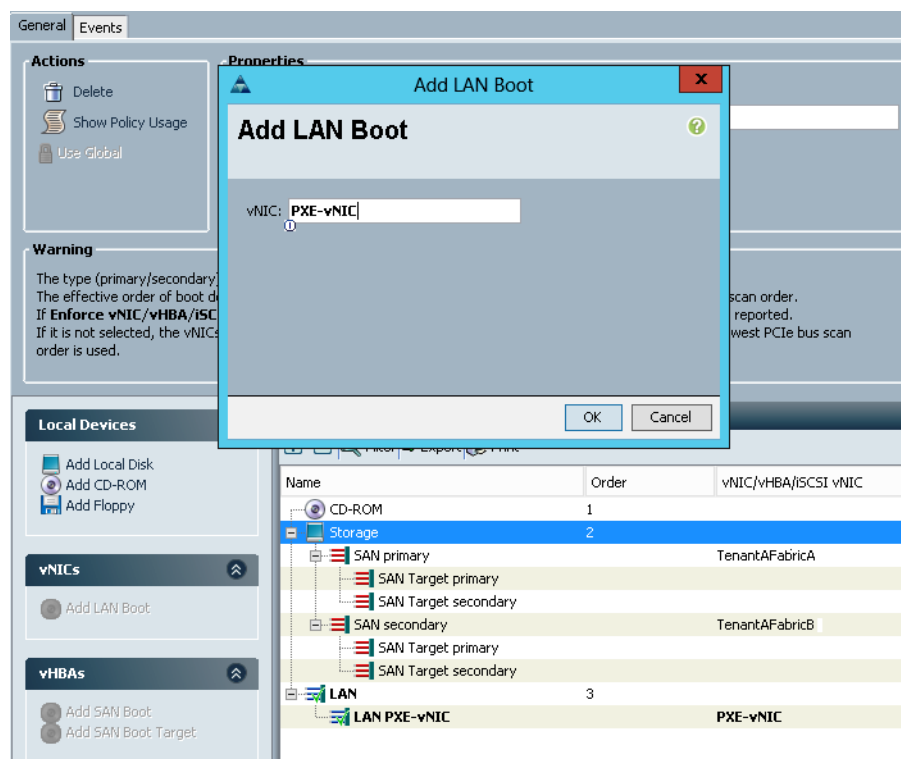
## PXE Boot Policy for Baremetal RHEL 6.3 Host

To configure PXE boot policy follow these setps :

1. Create Boot Policy TenantAPXEBoot on TenantA Org for SAN Boot Policy with primary and secondary type using two vHBAs as defined above in SAN Boot Policy for ESXi 5.1 host
2. Add LAN Boot option, type PXE-A-NIC in vNIC field ( Make sure PXE-A-NIC Static vNIC Name matches with LAN Boot Name (PXE-A-NIC)
3. Check **Enforce vNIC/vHBA/iSCSI Name** check box.



**Figure 116** Defining the FCoE Boot Policy Properties



4. Select SAN in boot order and Move Up to first
5. Click **Save**

**Figure 117**      **Defining Boot Order Properties**

**Actions**

- Delete
- Show Policy Usage
- Use Global

**Properties**

Name: **TenantAPXEBoot**

Description:

Owner: **Local**

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☒

**Warning**

The type (primary/secondary) does not indicate a boot order presence. The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order. If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported. If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

- vNICs
- vHBAs
- iSCSI vNICs

**Boot Order**

| Name                 | Order | vNIC/vHBA/iSCSI vNIC | Type      | Lun ID | WWN                     |
|----------------------|-------|----------------------|-----------|--------|-------------------------|
| Storage              | 1     |                      |           |        |                         |
| SAN primary          |       | TenantAFabricA       | Primary   |        |                         |
| SAN Target primary   |       |                      | Primary   | 0      | 20:05:00:A0:98:37:A0:70 |
| SAN Target secondary |       |                      | Secondary | 0      | 20:08:00:A0:98:37:A0:70 |
| SAN secondary        |       | TenantAFabricB       | Secondary |        |                         |
| SAN Target primary   |       |                      | Primary   | 0      | 20:06:00:A0:98:37:A0:70 |
| SAN Target secondary |       |                      | Secondary | 0      | 20:09:00:A0:98:37:A0:70 |
| LAN                  | 2     |                      |           |        |                         |
| LAN PXE-A-NIC        |       | PXE-A-NIC            | Primary   |        |                         |

## Configuring Storage Zoning

After the Cisco UCS service profiles have been created (in the previous steps), the infrastructure blades in the environment each have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information in the tables below.

**Table 18**      **NetApp Controller FC Port FC Portname**

| NetApp Controller | FC Port | FC Portname             |
|-------------------|---------|-------------------------|
| Controller A      | 0c      | 50:0a:09:83:8d:93:40:7f |
| Controller B      | 0c      | 50:0a:09:83:9d:93:40:7f |



### Note

On each NetApp controller CLI type the command `show fcp adapters` to gather the required information as listed in [Table 19](#).

**Table 19**      **WWPN Values for Tenants A and B**

| Cisco UCS Service Profile Name | Fabric-A WWPN           | Fabric-B WWPN           |
|--------------------------------|-------------------------|-------------------------|
| TenantA-HostA-Platinum-SLA1    | 20:00:00:25:b5:01:01:09 | 20:00:00:25:b5:01:01:08 |
| TenantA-HostB-Platinum-SLA1    | 20:00:00:25:b5:01:01:07 | 20:00:00:25:b5:01:01:06 |

To gather the WWPN values for both the fabrics, launch the Cisco UCS Manager GUI, and follow the steps below:

1. In the left pane select the **Servers** tab.
2. Choose **Servers > Service Profiles > root**.
3. Click **TenantA**.
4. Click each service profile and then click the **Storage** tab in the right.
5. Record the WWPN information in the right display window for both Fabric-A-HBA and Fabric-B-HBA for both the Service Profiles.

## Creating VSANs, Assign FC Ports, Turn on FC Ports

These steps provide details for configuring the necessary Zoneset and Zone members, to provide FC Initiator and target access to perform SAN Boot on UCS, NetApp Storage on Nexus network cloud environment



### Note

This procedure sets up Fibre Channel connections between the Cisco Nexus 5548s and the NetApp storage systems.

### Creating Zones and Zoneset

These steps provide details for configuring zones members and Zoneset for the primary and secondary boot path with all target ports.

#### Cisco Nexus 5548 A

To create the storage zone and the zoneset for each service profile, login to the Cisco Nexus switch and type the CLI commands listed in these steps:

1. Creating Zone.
2. Type zone name TenantA-HostA-Platinum-SLA1 vsan 200.
3. Type member pwwn 20:00:00:25:b5:01:01:01.
4. Type member pwwn 20:05:00:a9:98:37:a0:70.
5. Type member pwwn 20:06:00:a9:98:37:a0:70.
6. Type exit.
7. Type zone name TenantA-HostB-Platinum-SLA1 vsan 200.
8. Type member pwwn 20:00:00:25:b5:01:01:03.
9. Type member pwwn 20:05:00:a9:98:37:a0:70.
10. Type member pwwn 20:06:00:a9:98:37:a0:70.
11. Type exit.
12. Type zone name TenantA-BareMetalHostA-Platinum-SLA1 vsan 200.
13. Type member pwwn 20:00:00:25:b5:04:00:06.
14. Type member pwwn 20:05:00:a9:98:37:a0:70.
15. Type member pwwn 20:06:00:a9:98:37:a0:70.
16. Type exit.

**Note**

After the zone for the primary path of the first Cisco UCS service profiles has been created, create a zoneset to organize and manage them.

17. Create the zoneset and add the necessary members.
18. Type zoneset name TenantA vsan 200.
19. Type member TenantA-HostA-Plantinum-SLA1.
20. Type member TenantA-HostB-Plantinum-SLA2.
21. Type exit.
22. Type member TenantA-BareMetalHostA-Plantinum-SLA2.
23. Type exit.
24. Activate the zoneset.
25. Type zoneset activate name TenantA vsan 200.
26. Type exit.
27. Type copy run start. The command output is below.

```
N5548-L21-A# sh zoneset vsan 200
zoneset name TenantA vsan 200

zone name TenantA-HostA-Plantinum-SLA1 vsan 200
 pwn 20:00:00:25:b5:01:01:01
 pwn 20:05:00:a9:98:37:a0:70
 pwn 20:06:00:a9:98:37:a0:70
zone name TenantA-HostB-Plantinum-SLA1 vsan 200
 pwn 20:00:00:25:b5:01:01:03
 pwn 20:05:00:a9:98:37:a0:70
 pwn 20:06:00:a9:98:37:a0:70
zone name TenantA-BareMetalHostA-Plantinum-SLA1 vsan 200
 pwn 20:00:00:25:b5:04:00:06
 pwn 20:05:00:a9:98:37:a0:70
 pwn 20:06:00:a9:98:37:a0:70
```

**Cisco Nexus 5548 B**

1. Create the zone for each service profile.
2. Type zone name TenantA-HostA-Plantinum-SLA1 vsan 300.
3. Type member pwn 20:00:00:25:b5:01:01:02.
4. Type member pwn 20:05:00:a9:98:37:a0:70.
5. Type member pwn 20:06:00:a9:98:37:a0:70.
6. Type exit.
7. Type zone name TenantA-HostB-Plantinum-SLA1 vsan 300.
8. Type member pwn 20:00:00:25:b5:01:01:04.
9. Type member pwn 20:05:00:a9:98:37:a0:70.
10. Type member pwn 20:06:00:a9:98:37:a0:70.
11. Type exit.
12. Type zone name TenantA-BareMetalHostA-Plantinum-SLA1 vsan 200.
13. Type member pwn 20:00:00:25:b5:04:00:05.

14. Type member pwwn 20:05:00:a9:98:37:a0:70.
15. Type member pwwn 20:06:00:a9:98:37:a0:70.
16. Type exit.

**Note**

After the zone for the primary path of the first Cisco UCS service profiles has been created, create a zoneset to organize and manage them.

17. Create the zoneset and add the necessary members.
18. Type zoneset name TenantA vsan <300>.
19. Type member TenantA-HostA-Plantinum-SLA1.
20. Type member TenantA-HostB-Plantinum-SLA1.
21. Type exit
22. Type member TenantA-BareMetalHostA-Plantinum-SLA2.
23. Type exit.
24. Activate the zoneset.
25. Type zoneset activate name TenantA vsan <300>.
26. Type exit
27. Type copy run start. The command output is below.

```
N5548-L21-A# sh zoneset vsan 10
zoneset name TenantA vsan 10

zone name TenantA-HostA-Plantinum-SLA1 vsan 300
 pwwn 20:00:00:25:b5:01:01:02
 pwwn 20:05:00:a9:98:37:a0:70
 pwwn 20:06:00:a9:98:37:a0:70
zone name TenantA-HostB-Plantinum-SLA1 vsan 300
 pwwn 20:00:00:25:b5:01:01:04
 pwwn 20:05:00:a9:98:37:a0:70
 pwwn 20:06:00:a9:98:37:a0:70
zone name TenantA-BareMetalHostA-Plantinum-SLA1 vsan 200
 pwwn 20:00:00:25:b5:04:00:05
 pwwn 20:05:00:a9:98:37:a0:70
 pwwn 20:06:00:a9:98:37:a0:70
```

## Citrix CloudPlatform Host Definition

Hosts are like a single computer that provides the computing resources to run the guest virtual machines. Hypervisor software is installed on each host to manage the guest VMs. Cisco UCS Manager ideally defines the VMware ESXi Server Compute, Storage, and Network cloud infrastructure design required for multi-tenant s to host their services based on the Service Levels Agreements (SLA) defined by the cloud providers.

The host definition design components are discussed below.

## Compute

The compute definition includes CPU, Memory and IO components which are required to support multi-tenants services (Applications) function in terms of performance, dynamic scaling, high availability, manageability and provisioning based on these data points service levels have been defined with associated cost.

To meet these service levels UCS offers Server Pools with qualification policies like CPU speed, and memory capacity definitions for various service levels which are included with Cisco UCS Manager for faster provision time to handle dynamic scaling requirements. UCS offers extended memory technology and 80G IO Virtual Interface Adapters with ether channels to handle multi-tenants services data traffic at any time with no single point of failure it offers high availability in cloud environment.

## Storage

The storage definition involves network scalability, data isolation, high availability and performance to support multi-tenant storage requirements. Based on these data points service levels have been defined with associated cost.

To meet storage requirements, UCS offers virtual Logical SANs to provide storage data isolation in multi-tenant cloud environments. To provide high availability UCS offers SAN Port Channels which takes care of any physically FC ports failures and also can be dynamically expanded to support scalability. On the compute host side each virtual host bus adapter (vHBAs) storage data is pinned to specific SAN Port Channels using SAN PIN Groups depending on the bandwidth service level requirement.

## Network

The network definition involves quality of service, high availability, scalability and performance to support multi-tenant storage requirements. Based on these data points service levels have been defined with associated cost.

To meet network definitions UCS offers virtual port channels and Virtual LANs to provide network data isolations in multi-tenant s cloud environment. To provide high availability UCS offers Ethernet Port Channels which takes care of any physical port failures and also dynamic expansion to support scalability. On the compute host side each virtual Ethernet adapter (vNICs) network data is pinned to specific LAN Port Channels using LAN PIN Groups depending on the bandwidth service level requirement.

## Creating Service Profile Templates

The service profile represents a logical view of a single blade server, without the need to know exactly which blade is discussed. The profile object contains the server personality: for example, identity and network information. The profile can then be associated with a single blade at a time. The Compute, Storage, and Network infrastructure components are all provisioned, configured and applied on cloud host server using UCS Service Profile Templates.

The service profile template is a single management configuration window where all of infrastructure components can be defined, configured and applied to multiple cloud hosts in the form of service profile.

The instructions below show creation of different service profile templates based on the service levels which can be later used to instantiate service profiles applied to individual cloud tenant host.

In this section the following tasks are elaborated:

- Compute Definition
- Network Definition
- Storage Definition

## Compute Definition

### Creating Service Profile Templates

This section details the creation of two service profile templates: one for Hypervisor ESX Host and second for Baremetal RHEL Operating System.

Hypervisor ESX Host Service Profile Template

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:

1. Click the **Servers** tab in the left pane.
2. Choose **Service Profiles > Sub-Organization**. Highlight TenantA.
3. Select Create Service Profile Templates on the **General** tab in the right pane.
4. The **Create Service Profile Template** window is displayed.
5. Name the service profile template Platinum-Template.
6. Select Initial Template.
7. In the UUID section, select Platinum-Compute-UUID as the UUID pool.
8. Click **Next** to continue to the next section.

**Figure 118**      *Adding Unique Name to the Service Profile Template*

**Create Service Profile Template**

## Unified Computing System Manager

Create Service Profile Template

1. **Identify Service Profile Template**
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

### Identify Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root/org-TenantA**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ☒ Initial Template ☐ Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

**UUID**

UUID Assignment:

The UUID will be assigned from the selected pool.  
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev   Next >   Finish   Cancel

## Network Definition

### Creating vNICs

1. Leave the Dynamic vNIC Connection Policy field at the default.
2. Select Expert for the How would you like to configure LAN connectivity?
3. Click **Add** icon (Specify the desired number of vNICs that the server uses to connect to the LAN).
4. **Create vNIC** window opens.
5. Enter Management-A-NIC in the Name field.
6. Enable **Use vNIC Template** check box.
7. Select Management-A-NIC in vNIC Template list box.
8. Select **VMWare in Adapter Policy** list box.
9. Click **OK**.



**Figure 119**      *Defining the vNIC Properties for Management A-NIC*

The screenshot shows the 'Create vNIC' dialog box. The title bar is blue with a triangle icon and the text 'Create vNIC'. The main area has a light blue header with the title 'Create vNIC'. Below the header, there are several fields and controls:
 

- Name:** A text field containing 'Managment-A-NIC'.
- Use vNIC Template:** A checked checkbox.
- + Create vNIC Template:** A green button with a plus icon.
- vNIC Template:** A dropdown menu showing 'Managment-A-NIC'.
- Adapter Performance Profile:** A section header.
- Adapter Policy:** A dropdown menu showing 'VMWare'.
- + Create Ethernet Adapter Policy:** A green button with a plus icon.

10. Click **Add**.
11. Enter Management-B-NIC in the Name field.
12. Enable **Use vNIC Template** check box
13. Select **Management-B-NIC** in vNIC Template list box.
14. Select **VMWare** in Adapater Policy list box.
15. Click **OK**.

**Figure 120**      *Defining vNIC properties for the Management B-NIC*

The screenshot shows the 'Create vNIC' dialog box. The title bar is blue with a triangle icon and the text 'Create vNIC'. The main area has a light blue header with the title 'Create vNIC'. Below the header, there are several fields and controls:
 

- Name:** A text field containing 'Managment-B-NIC'.
- Use vNIC Template:** A checked checkbox.
- + Create vNIC Template:** A green button with a plus icon.
- vNIC Template:** A dropdown menu showing 'Managment-B-NIC'.
- Adapter Performance Profile:** A section header.
- Adapter Policy:** A dropdown menu showing 'VMWare'.
- + Create Ethernet Adapter Policy:** A green button with a plus icon.

16. Click **Add**
17. Enter guest-A-NIC in the Name field.
18. Enable **Use vNIC Template** check box.
19. Select **guest-A-NIC** in vNIC Template list box.
20. Select **VMWare** in Adapater Policy list box.
21. Click **OK**.

**Figure 121** *Defining vNIC Properties for the guest-A-NIC*

**Create vNIC**

Name:

Use vNIC Template: ☒

[+ Create vNIC Template](#)

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:  [+ Create Ethernet Adapter Policy](#)

22. Click **Add**.
23. Enter **Guest-A-NIC** in the Name field.
24. Enable **Use vNIC Template** check box.
25. Select **guest-A-NIC** in vNIC Template list box.
26. Select **VMWare** in Adapter Policy list box.
27. Click **OK**.

**Figure 122** *Defining vNIC Properties for the guest-B-NIC*

**Create vNIC**

Name:

Use vNIC Template: ☒

[+ Create vNIC Template](#)

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:  [+ Create Ethernet Adapter Policy](#)

28. Click **Add**.
29. Enter **NFS-A-NIC** in the Name field.
30. Enable **Use vNIC Template** check box
31. Select **NFS-A-NIC** in vNIC Template list box.
32. Select **VMWare** in Adapter Policy list box.
33. Click **OK**.

**Figure 123**      *Defining vNIC Properties for the NFS-A-NIC*

The screenshot shows the 'Create vNIC' dialog box. The title bar says 'Create vNIC'. The main title is 'Create vNIC'. The 'Name' field contains 'NFS-A-NIC'. Below it, 'Use vNIC Template' is checked. There is a '+ Create vNIC Template' button. The 'vNIC Template' dropdown menu is open, showing 'NFS-A-NIC' selected. Below that, the 'Adapter Performance Profile' section shows 'Adapter Policy' set to 'VMWare' with a '+ Create Ethernet Adapter Policy' button.

34. Click **Add**.
35. Enter **NFS-B-NIC** in the Name field.
36. Enable **Use vNIC Template** check box.
37. Select **NFS-B-NIC** in vNIC Template list box.
38. Select **VMWare** in Adapter Policy list box.
39. Click **OK**.

**Figure 124**      *Defining vNIC Properties for the NFS-B-NIC*

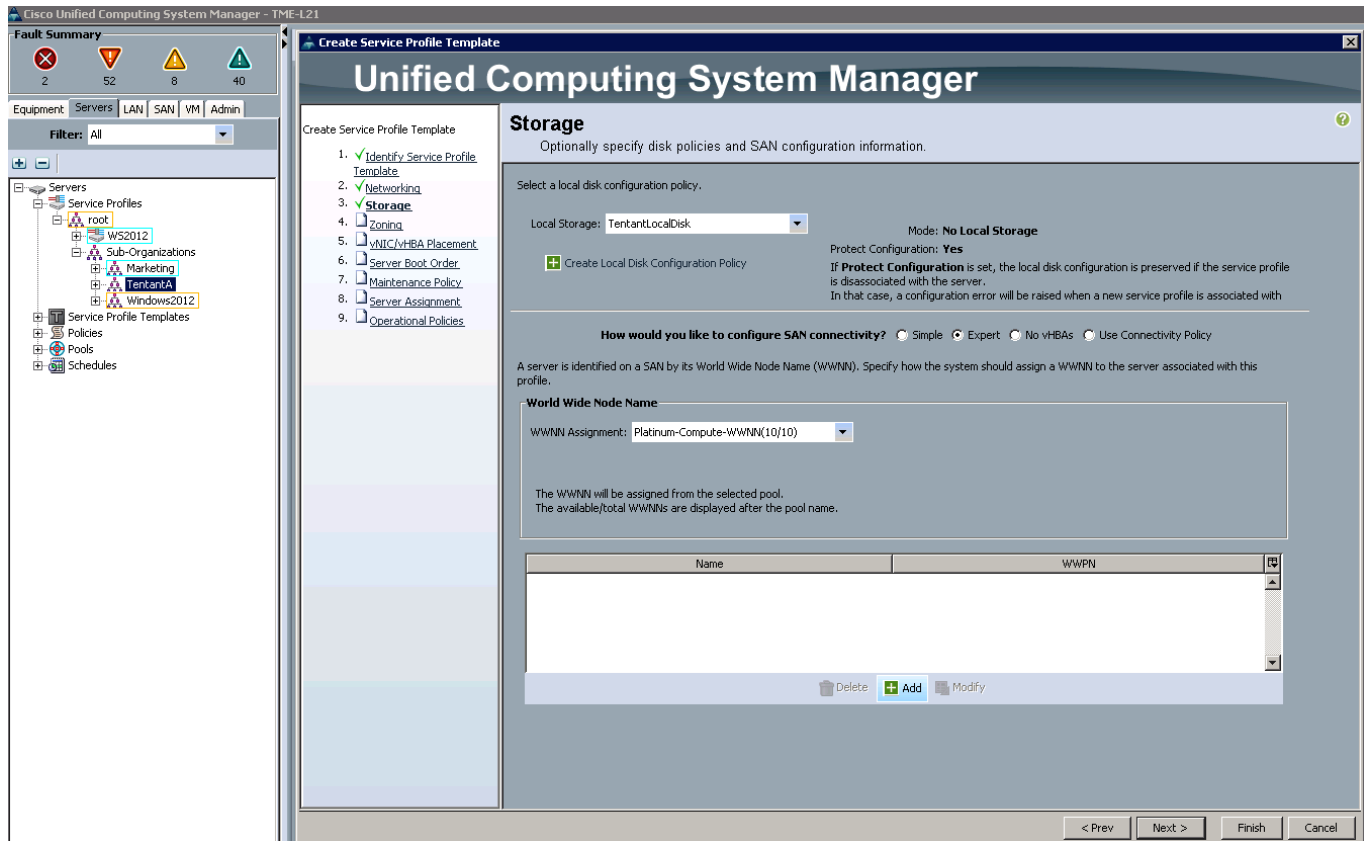
The screenshot shows the 'Create vNIC' dialog box. The title bar says 'Create vNIC'. The main title is 'Create vNIC'. The 'Name' field contains 'NFS-B-NIC'. Below it, 'Use vNIC Template' is checked. There is a '+ Create vNIC Template' button. The 'vNIC Template' dropdown menu is open, showing 'NFS-B-NIC' selected. Below that, the 'Adapter Performance Profile' section shows 'Adapter Policy' set to 'VMWare' with a '+ Create Ethernet Adapter Policy' button.

## Storage Definition

### Creating vHBAs


1. Select TenantLocalDisk for Local Storage field.
2. Click **Expert** to define how would you like to configure SAN connectivity field.
3. In the WWNN Assignment field, select Platinum-Compute-WWNN.

**Figure 125** *Defining the Disk Policies and SAN Storage Properties*



- Click **Add** button to add vHBAs to the template.
- Enter **TenantAFabricA** in the Name field.
- Select **Use vHBA Template** check box.
- Select **TenantAFabricA** in vHBA Template.
- Select VMWare Adapter Policy list box.
- Click **OK**.

**Figure 126** *Defining the vHBA Values for the Fabric-A-HBA*



- g. Click **Add** button to add vHBAs to the template.
- h. Enter **TenantAFabricB** in the Name field.
- i. Select **Use vHBA Template** check box.
- j. Select **TenantAFabricB** in vHBA Template.
- k. Select VMWare Adapter Policy list box.
- l. Click **OK**.

**Figure 127** Defining the vHBA Values for the Fabric-B-HBA

**Create vHBA**

Name:

Use vHBA Template: ☒

[+ Create vHBA Template](#)

vHBA Template:

**Adapter Performance Profile**

Adapter Policy:  [+ Create Fibre Channel Adapter Policy](#)

4. Click **Next** to continue to the next section.
5. Zoning section
  - a. Accept all values as default.
6. Click **Next**.

**Figure 128** Define the Zoning Information

**Unified Computing System Manager**

**Create Service Profile Template**

**Zoning**

Specify zoning information

**WARNING: Switch in end-host mode. In end-host mode, zoning configuration will NOT be applied.**

Zoning configuration involves the following steps:

1. Select vHBA Initiator(s) (vHBAs are created on storage page)
2. Select vHBA Initiator Group(s)
3. Add selected Initiator(s) to selected Initiator Group(s)

**Select vHBA Initiators**

| Name         |
|--------------|
| Fabric-A-HBA |
| Fabric-B-HBA |

[>> Add To >>](#)

**Select vHBA Initiator Groups**

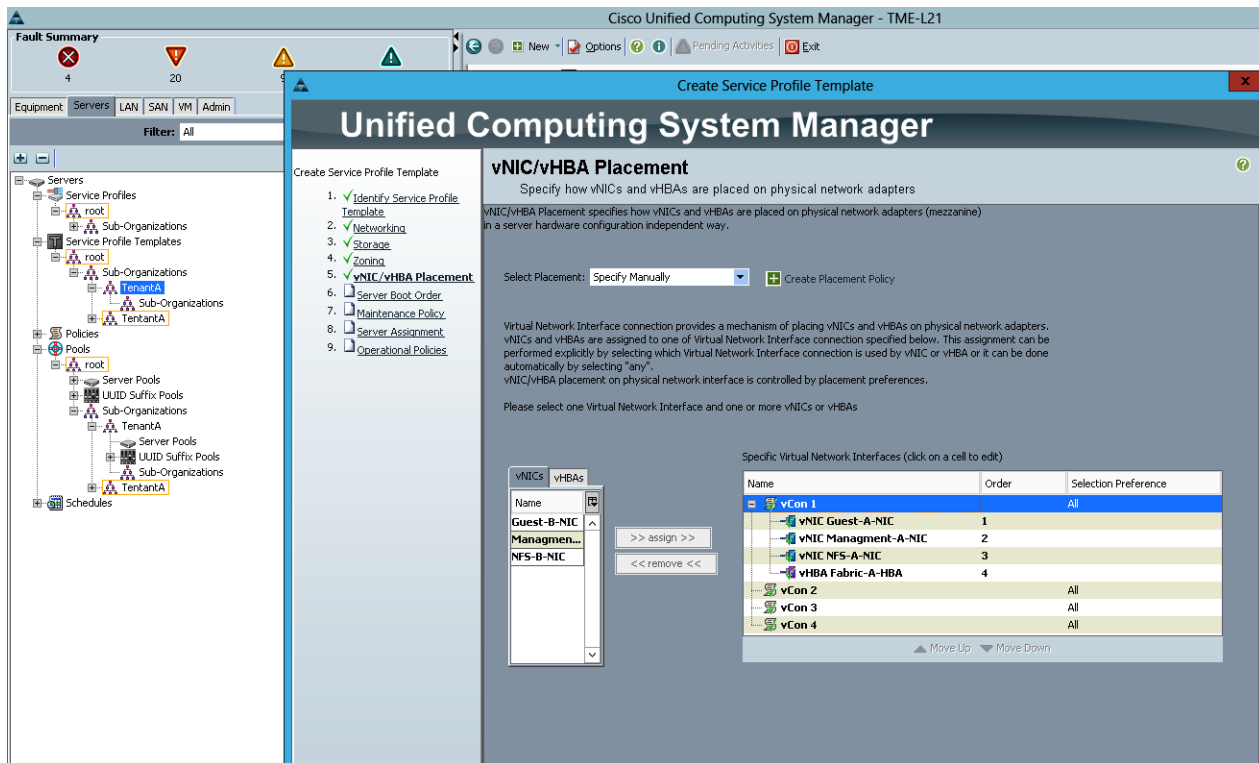
| Name | Storage Connection Policy Name |
|------|--------------------------------|
|------|--------------------------------|

[Delete](#) [Add](#) [Modify](#)

7. Click **Next** to continue to the next section.
8. vNIC/vHBA Placement Section

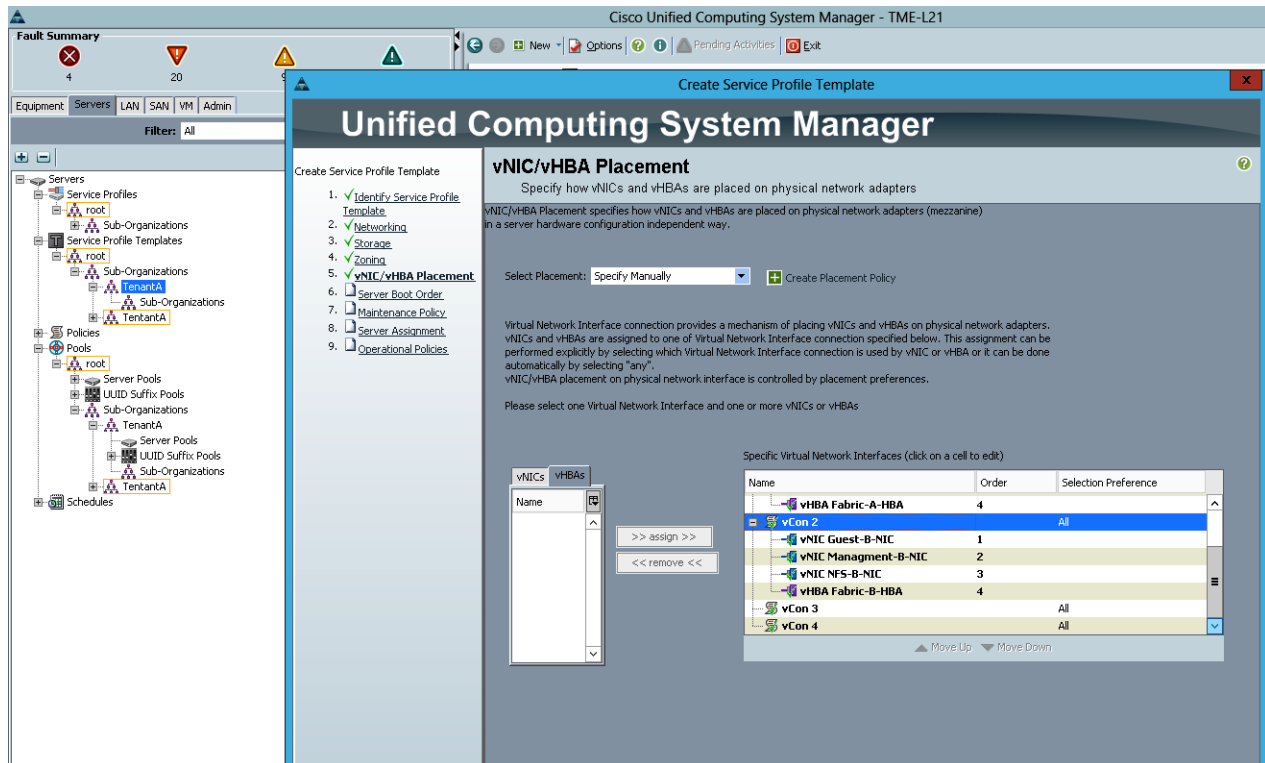
- b. Select the Manual Placement Policy in the Select Placement field.
- c. Select vCon1 assign the vNICs in the following order:
  - Management-A-NIC
  - guest-A-NIC
  - NFS-A-NIC
  - Fabric-A-HBA

**Figure 129**      *Defining the vNIC and vHBA Placement*



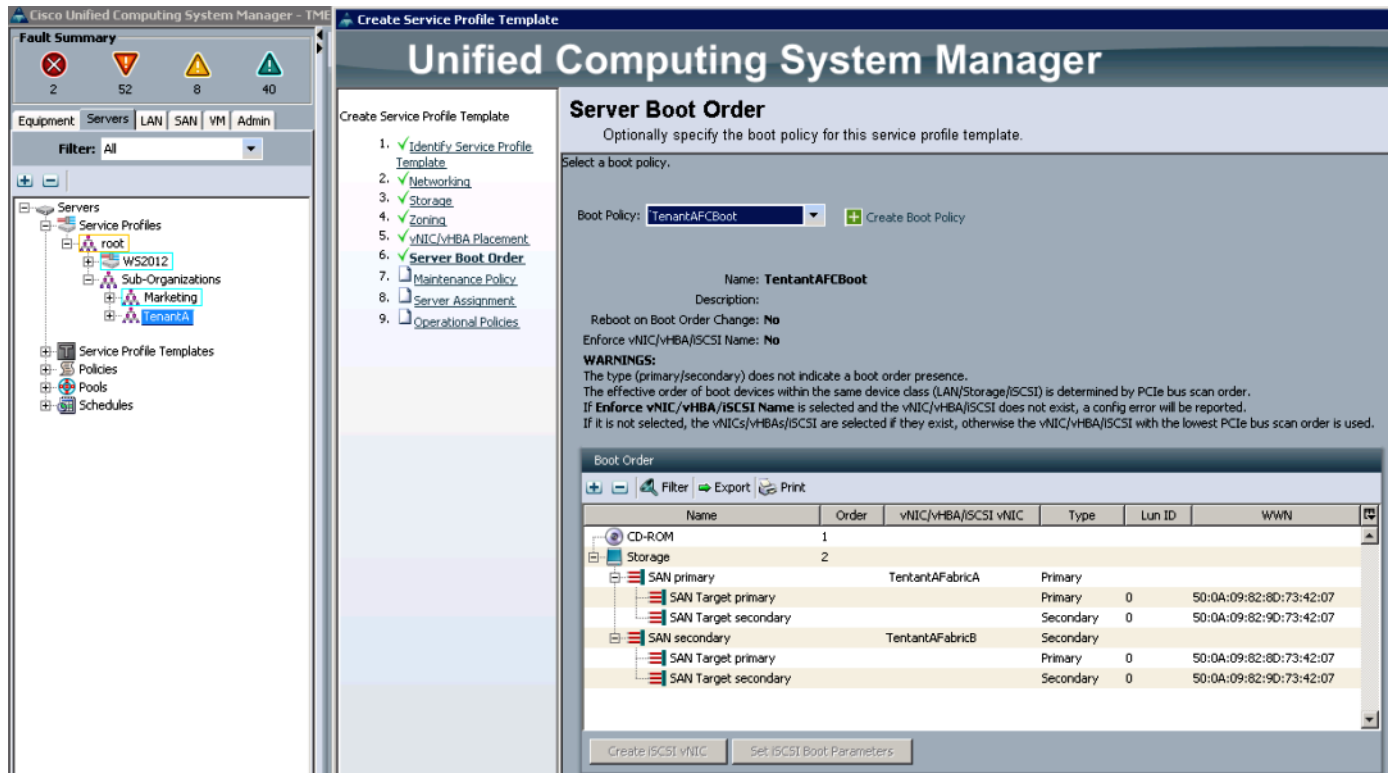
- d. Select vCon2, and assign the vNICs in the following order:
  - Management-B-NIC
  - guest-B-NIC
  - NFS-B-NIC
  - Fabric-B-HBA
9. Click **Next**.

**Figure 130**      **Assigning the vNICs in the vCon2**



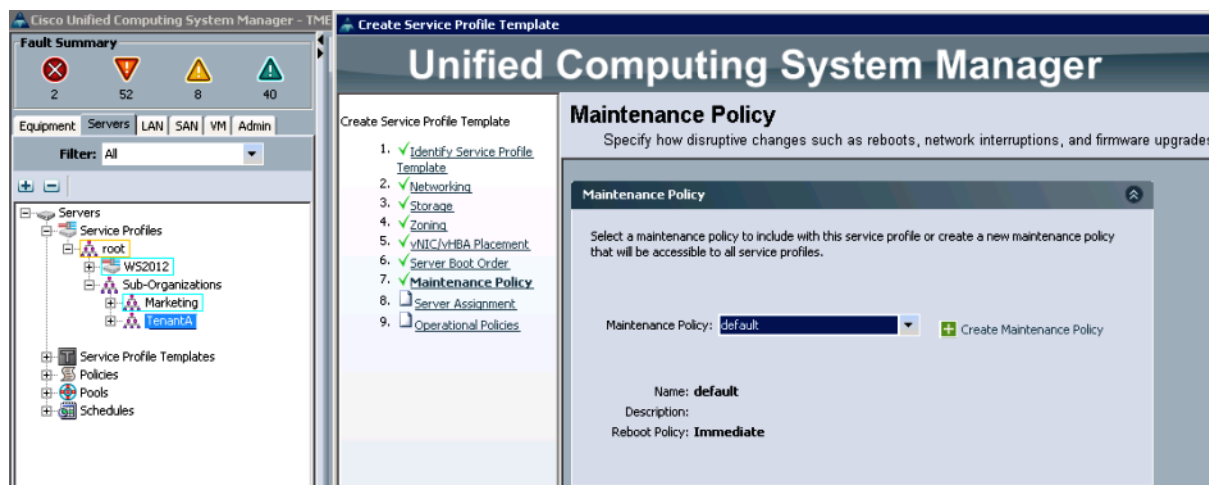
10. Server Boot Order Section
  - a. Select TenantAFCBoot Boot Policy.
11. Click Next.

**Figure 131** Applying TenantFCBoot Boot Policy for the Service Profile Template



12. Maintenance Policy Section
  - a. Select default Maintenance Policy.
13. Click Next.

**Figure 132** Applying default Policies applied for Maintenance Policy During Service Disruptions



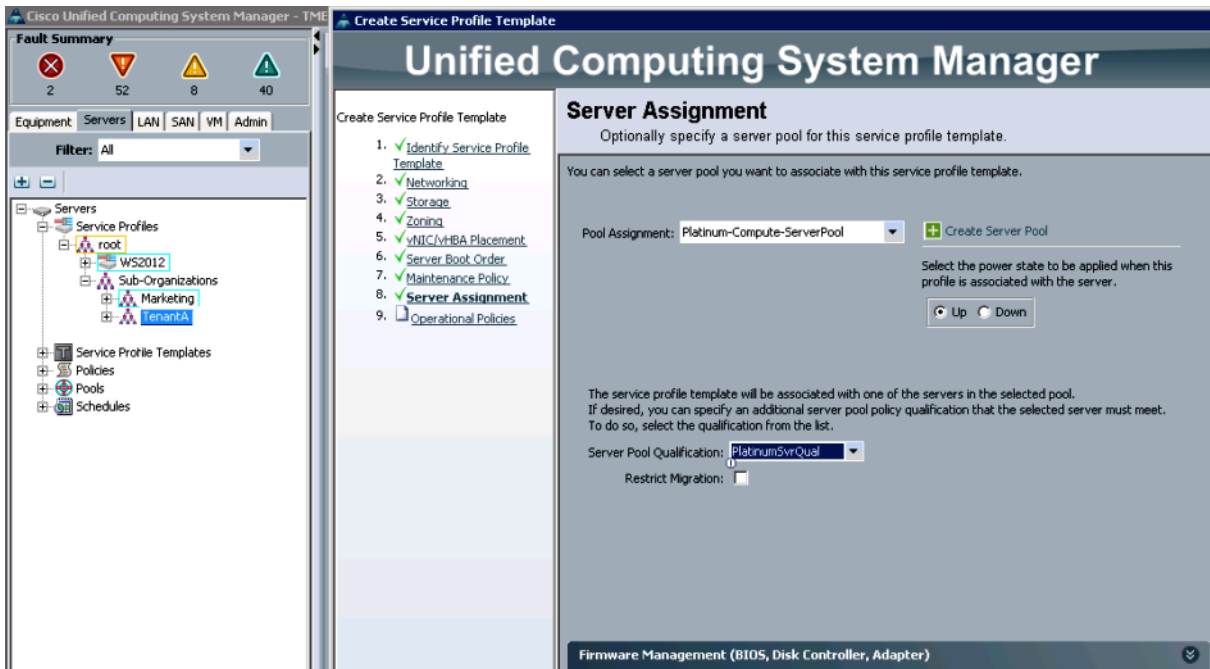
14. Server Assignment Section
  - a. Select Platinum-Compute-ServerPool for Pool Assignment.



- b. Click **Up** radio button to select the power state to be applied when this profile is associated with the server.
- c. Select **PlatinumSvrQual** for Server Pool Qualification.
- d. Do not check **Restrict Migration**.

15. Click **Next**.

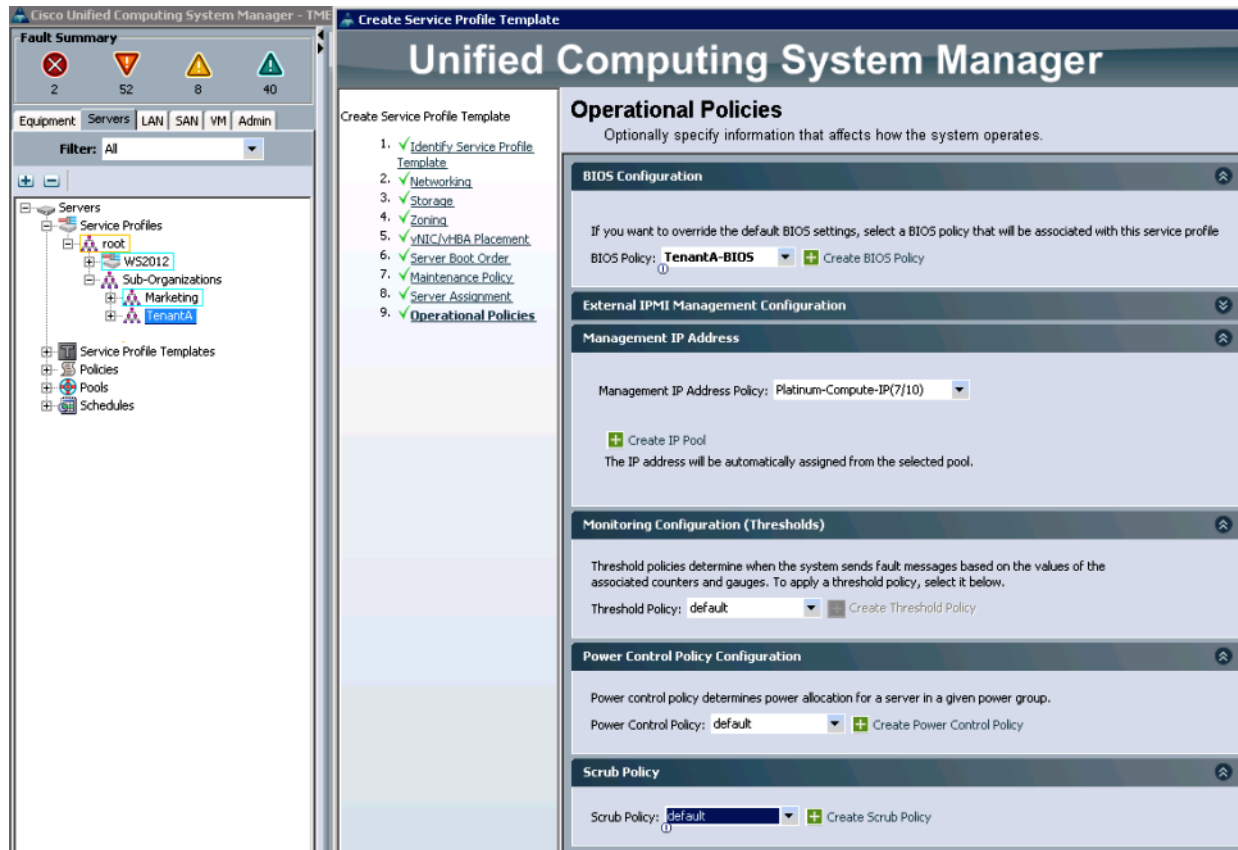
**Figure 133** Associating the Server Pool with the Service Profile



16. Operational Policies Section

- a. Select **TenantA-BIOS** for BIOS Policy.
- b. Select default for External IPMI Management Configuration
- c. Select **Platinum-Compute-IP** for Management IP Address Policy.
- d. Select default for Monitoring Configuration.
- e. Select default for Power Control Policy Configuration.
- f. Select default for Scrub Policy.

17. Click **Next**.

**Figure 134**      *Defining the attributes for the Operational Policy*

Follow the steps described above to create Service Profile Templates for Gold-Template, Silver-Template and Bronze-Template with data values listed in [Table 20](#) and [Table 21](#).

**Table 20**      *Service Profile Templates*

| Name            | Type             | UUID                |
|-----------------|------------------|---------------------|
| Gold-Template   | Initial Template | Gold-Compute-UUID   |
| Silver-Template | Initial Template | Silver-Compute-UUID |
| Bronze-Template | Initial Template | Bronze-Compute-UUID |

**Table 21**      *Networking Configuration*

| Dynamic vNIC Connection Policy    | Configure LAN Connectivity |
|-----------------------------------|----------------------------|
| No Dynamic vNIC Policy by default | Expert                     |

Create vNIC For Gold-Template as per the values defined in [Table 22](#).

**Table 22** *vNICs and their Properties for the Gold Service Class*

| Name             | MAC Address      | Fabric ID | VLAN Native VLAN | MTU  | Pin Group        | Adapter policy | Dynamic vNIC Connection Policy | QoS Policy       | Network Control Policy |
|------------------|------------------|-----------|------------------|------|------------------|----------------|--------------------------------|------------------|------------------------|
| Management-A-NIC | Gold-Compute-MAC | A         | Management-VLAN  | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| Management-B-NIC | Gold-Compute-MAC | B         | Management-VLAN  | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| guest-A-NIC      | Gold-Compute-MAC | A         | guest-VLAN       | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |
| guest-B-NIC      | Gold-Compute-MAC | B         | guest-VLAN       | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |
| NFS-A-NIC        | Gold-Compute-MAC | A         | guest-VLAN       | 9000 | Tenant A-Net-PIN | default        | Not set                        | Silver-Net-SLA   | default                |
| NFS-B-NIC        | Gold-Compute-MAC | B         | guest-VLAN       | 9000 | Tenant A-Net-PIN | default        | Not set                        | Silver-Net-SLA   | default                |

Create vNIC For Silver-Template using the values defined in [Table 23](#).

**Table 23** *vNICs and their Properties for the Silver Service Class*

| Name             | MAC Address        | Fabric ID | VLAN Native VLANs | MTU  | Pin Group        | Adapter policy | Dynamic vNIC Connection Policy | QoS Policy       | Network Control Policy |
|------------------|--------------------|-----------|-------------------|------|------------------|----------------|--------------------------------|------------------|------------------------|
| Management-A-NIC | Silver-Compute-MAC | A         | Management-VLAN   | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| Management-B-NIC | Silver-Compute-MAC | B         | Management-VLAN   | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| guest-A-NIC      | Silver-Compute-MAC | A         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |
| guest-B-NIC      | Silver-Compute-MAC | B         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |

**Table 23** *vNICs and their Properties for the Silver Service Class*

|           |                    |   |            |      |                  |         |         |                |         |
|-----------|--------------------|---|------------|------|------------------|---------|---------|----------------|---------|
| NFS-A-NIC | Silver-Compute-MAC | A | guest-VLAN | 9000 | Tenant A-Net-PIN | default | Not set | Silver-Net-SLA | default |
| NFS-B-NIC | Silver-Compute-MAC | B | guest-VLAN | 9000 | Tenant A-Net-PIN | default | Not set | Silver-Net-SLA | default |

Create vNIC For Bronze-Template using the values defined in [Table 24](#).

**Table 24** *vNICs and Their Properties for the Bronze Class*

| Name             | MAC Address        | Fabric ID | VLAN Native VLANs | MTU  | Pin Group        | Adapter policy | Dynamic vNIC Connection Policy | QoS Policy       | Network Control Policy |
|------------------|--------------------|-----------|-------------------|------|------------------|----------------|--------------------------------|------------------|------------------------|
| Management-A-NIC | Bronze-Compute-MAC | A         | Management-VLAN   | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| Management-B-NIC | Bronze-Compute-MAC | B         | Management-VLAN   | 9000 | Tenant A-Net-PIN | default        | Not set                        | Platinum-Net-SLA | default                |
| guest-A-NIC      | Bronze-Compute-MAC | A         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |
| guest-B-NIC      | Bronze-Compute-MAC | B         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Gold-Net-SLA     | default                |
| NFS-A-NIC        | Bronze-Compute-MAC | A         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Silver-Net-SLA   | default                |
| NFS-B-NIC        | Bronze-Compute-MAC | B         | guest-VLAN        | 9000 | Tenant A-Net-PIN | default        | Not set                        | Silver-Net-SLA   | default                |

**Table 25** *Storage Configurations*

| Local Storage   | Configure SAN Connectivity | WWNN Assignment   |
|-----------------|----------------------------|-------------------|
| TenantLocalDisk | Expert                     | Gold-Compute-WWNN |

**Table 26** *vSAN values for the Fabric A and B*

| Name         | WWPN              | Fabric ID | VSAN | Pin Group       | Persistent Binding | Max Data field Size | Adapter Policy | QoS Policy |
|--------------|-------------------|-----------|------|-----------------|--------------------|---------------------|----------------|------------|
| Fabric-A-HBA | Gold-Compute-WWPN | A         | 10   | TenantA-SAN-PIN | Enabled            | 2048                | default        | FC-Net-SLA |
| Fabric-B-HBA | Gold-Compute-WWPN | B         | 10   | TenantA-SAN-PIN | Enabled            | 2048                | default        | FC-Net-SLA |

**Table 27** *Zoning and vNIC / vHBA Placement*

| Zoning  | vHBA Placement |
|---------|----------------|
| Default | Manual         |

**Table 28** *vNIC and vHBA Placement on vCon1 and vCon2*

| vCon1            | vCon2            |
|------------------|------------------|
| Management-A-NIC | Management-B-NIC |
| guest-A-NIC      | guest-B-NIC      |
| NFS-A-NIC        | NFS-B-NIC        |
| Fabric-A-HBA     | Fabric-B-HBA     |

**Table 29** *Values for Server Boot Order and the Maintenance Policy*

| Boot Policy    | Maintenance Policy |
|----------------|--------------------|
| TenantAFC-Boot | Default            |

**Table 30** *Server Assignment*

| Pool Assignment           | Power State | Server Pool Qualification | Restrict Migration |
|---------------------------|-------------|---------------------------|--------------------|
| Gold-Compute-ServerPool   | Up          | GoldSvrQual               | Un Check           |
| Silver-Compute-ServerPool | Up          | SilverSvrQual             | Un Check           |
| Bronze-Compute-ServerPool | Up          | BronzeSvrQual             | Un Check           |

**Table 31**      **Operational Policies**

| <b>BIOS Policy</b> | <b>External IPMI Management</b> | <b>Management IP Address</b> | <b>Monitoring Configuration</b> | <b>Power Control Policy</b> | <b>Scrub Policy</b> |
|--------------------|---------------------------------|------------------------------|---------------------------------|-----------------------------|---------------------|
| TenantA-BIOS       | default                         | Gold-Compute-IP              | default                         | default                     | default             |
| TenantA-BIOS       | default                         | Silver-Compute-IP            | default                         | default                     | default             |
| TenantA-BIOS       | default                         | Bronze-Compute-IP            | default                         | default                     | default             |

### Baremetal RHELHost Service Profile Template

This section details the creation of Baremetal Service Profile Template defined to provision blade to host RHEL 6.X Operating System using CloudPlatform 4.2.1.

For Baremetal as a Service the Service profile Templates, Pools (UUID, MAC, IP, PWWN, WWNN, and Server), vNIC & vHBA Templates, Policies (Boot, IPMI, Adapter, BIOS, Firmware) must be created on root organization.



#### Note

In this study we do not show steps to create below Pools, Policies, vHBA Templates we expect all are created under root Organization, for more details in creating them see Cloud Compute Design and Deployment section

1. Platinum-Compute-UUID
2. Baremetal-Disk
3. Platinum-Compute-WWNN
4. Platinum-Compute-WWPN

Login to Cisco UCS Manager with User Admin:

1. Click the **Servers** tab in the left pane.
2. Choose **Service Profiles > root**.
3. Select Create Service Profile Templates on the General tab in the right pane.
4. The Create Service Profile Template window is displayed.
5. Name the service profile template Platinum-Baremetal-Template.
6. Select Initial Template.
7. In the UUID section, select BarePlatinum-Compute-UUID as the UUID pool.
8. Click **Next** to continue to the next section.

**Figure 135**      **Adding Unique Name to the Service Profile Template**

**Create Service Profile Template**

**Unified Computing System Manager**

Create Service Profile Template

1. **Identify Service Profile Template**
2. [Networking](#)
3. [Storage](#)
4. [Zoning](#)
5. [vNIC/vHBA Placement](#)
6. [Server Boot Order](#)
7. [Maintenance Policy](#)
8. [Server Assignment](#)
9. [Operational Policies](#)

### Identify Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root/org-TenantA**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ☒ Initial Template ☐ Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

**UUID**

UUID Assignment:

The UUID will be assigned from the selected pool.  
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev   Next >   Finish   Cancel

## Network Definition

### Creating vNICs

1. Leave the Dynamic vNIC Connection Policy field at the default.
2. Select Expert for the How would you like to configure LAN connectivity?
3. Click **Add** icon (Specify the desired number of vNICs that the server uses to connect to the LAN).
4. **Create vNIC** window opens.
5. Enter Bare-MGMT-A-NIC in the Name field.
6. Enable **Use vNIC Template** check box.
7. Select **Bare-MGMT-A-NIC** in vNIC Template list box.
8. Select **Linux** in Adapter Policy list box.
9. Click **OK**.

**Figure 136**      *Defining the vNIC Properties for Bare-MGMT-A-NIC*

**Create vNIC**

Name:

Use vNIC Template: ☒

+ Create vNIC Template

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:  + Create Ethernet Adapter Policy

10. Click **Add**.
11. Enter **PXE-A-NIC** in the Name field.
12. Enable **Use vNIC Template** check box.
13. Select **PXE-A-NIC** in vNIC Template list box.
14. Select **Linux** in Adapter Policy list box.
15. Click **OK**.

**Figure 137**      *Defining vNIC properties for the PXE-A-NIC*

**Create vNIC**

Name:

Use vNIC Template: ☒

+ Create vNIC Template

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:  + Create Ethernet Adapter Policy

## Storage Definition

### Creating vHBAs

1. Select **TenantLocalDisk** for Local Storage field.
2. Click **Expert** to define How would you like to configure SAN connectivity field.
3. In the WWNN Assignment field, select **Platinum-Compute-WWNN**
4. Click **Add**.



5. Enter TenantAFabricA in the Name field
6. Enable Use vHBA Template check box
7. Select TenantAFabricA in vHBA Template
8. Select Linux in Adapter Policy

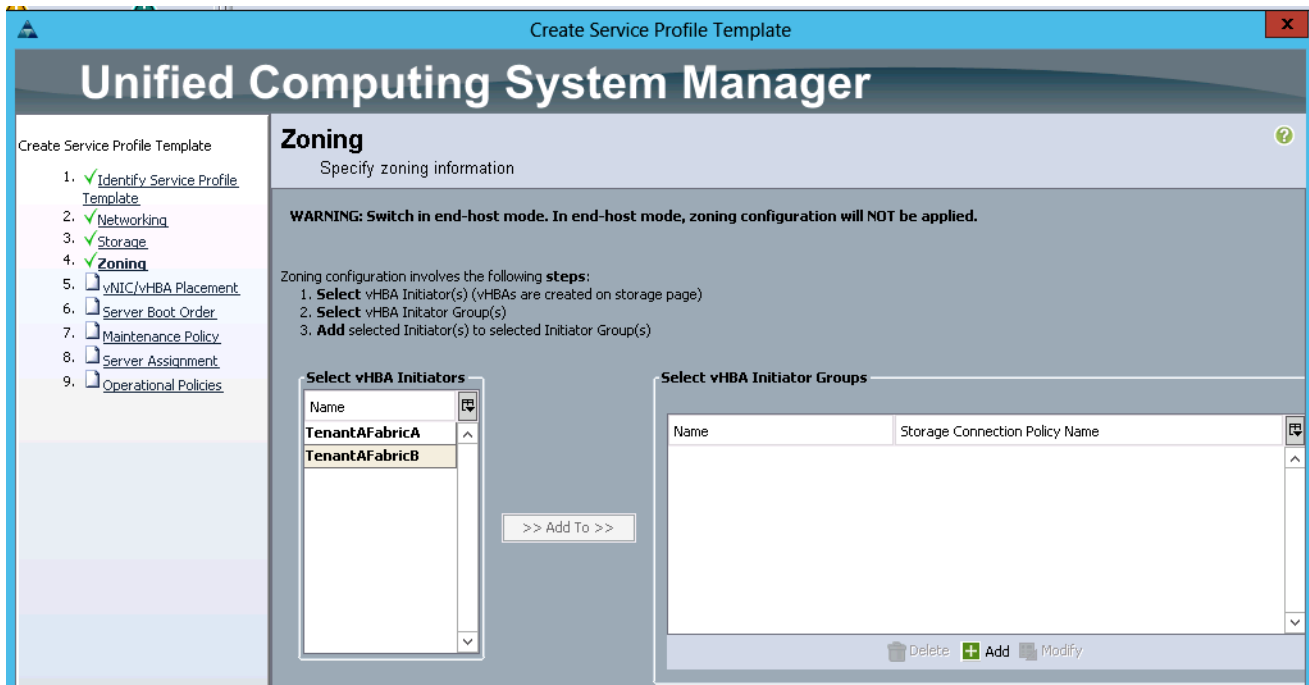
**Figure 138** *Defining the TenantAFabricA vHBA for the Service Profile Template*

9. Click **Add**.
10. Enter TenantAFabricB in the Name field.
11. Enable Use vHBA Template check box.
12. Select TenantAFabricA in vHBA Template.
13. Select Linux in Adapter Policy.

**Figure 139** *Defining the TenantAFabricB vHBA for the Service Profile Template*

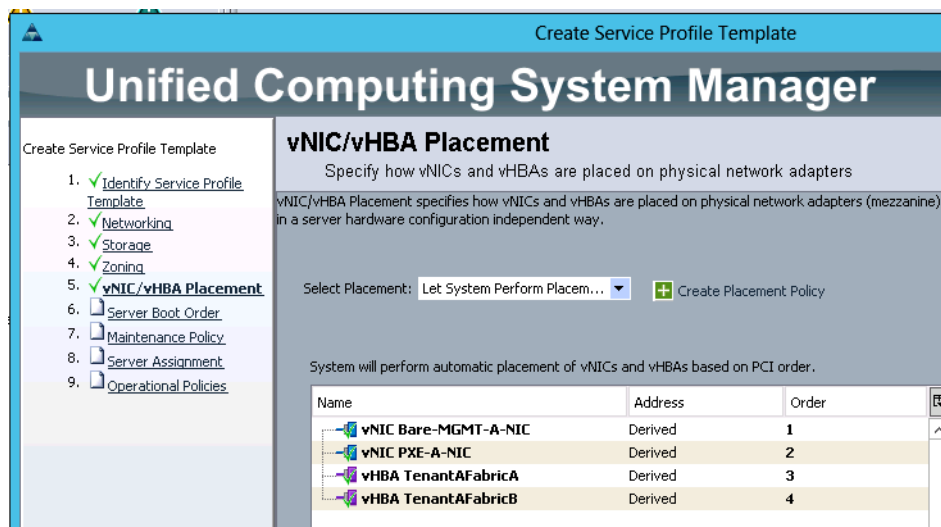
14. Click **Next** to continue to the next section.
15. Zoning section
  - a. Accept all values as default.
16. Click **Next**.

Figure 140 Define the Zoning Information



17. Click **Next** to continue to the next section.
18. vNIC/vHBA Placement Section
  - a. Select Let System Perform Placement Policy in the Select Placement field.
19. Click **Next**.

Figure 141 Defining the vNIC and vHBA Placement



20. Server Boot Order Section
  - a. Select TenantAPXEBoot Boot Policy.

21. Click **Next**.

**Figure 142** Defining the Boot Policy for the Service Profile Template

**Modify Boot Policy**

Boot Policy: **TenantAPXEBoot** + Create Boot Policy

Name: **TenantAPXEBoot**

Description:

Reboot on Boot Order Change: **No**

Enforce vNIC/vHBA/iSCSI Name: **Yes**

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Boot Order**

| Name                 | Order    | vNIC/vHBA/iSCSI vNIC | Type      | Lun ID | WWN             |
|----------------------|----------|----------------------|-----------|--------|-----------------|
| <b>Storage</b>       | <b>1</b> |                      |           |        |                 |
| SAN primary          |          | TenantAFabricA       | Primary   |        |                 |
| SAN Target primary   |          |                      | Primary   | 0      | 20:05:00:A0:... |
| SAN Target secondary |          |                      | Secondary | 0      | 20:08:00:A0:... |
| SAN secondary        |          | TenantAFabricB       | Secondary |        |                 |
| SAN Target primary   |          |                      | Primary   | 0      | 20:06:00:A0:... |
| SAN Target secondary |          |                      | Secondary | 0      | 20:09:00:A0:... |
| <b>LAN</b>           | <b>2</b> |                      |           |        |                 |
| LAN PXE-A-NIC        |          | PXE-A-NIC            | Primary   |        |                 |

Create iSCSI vNIC Set iSCSI Boot Parameters

22. Maintenance Policy Section

a. Select default Maintenance Policy

23. Click **Next**.

**Figure 143** *Defining the Policies applied to Server During Service Disruptions*

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The left sidebar lists the steps: 1. Identify Service Profile Template, 2. Networking, 3. Storage, 4. Zoning, 5. vNIC/vHBA Placement, 6. Server Boot Order, 7. Maintenance Policy (highlighted), 8. Server Assignment, and 9. Operational Policies. The main panel is titled 'Maintenance Policy' and contains the following text: 'Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades associated with this service profile.' Below this is a sub-section titled 'Maintenance Policy' with the instruction: 'Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.' There is a dropdown menu for 'Maintenance Policy' set to 'default' and a '+ Create Maintenance Policy' button. Below the dropdown, the details for the 'default' policy are shown: Name: default, Description: (empty), and Reboot Policy: Immediate.

**24. Server Assignment Section**

- a. Select Platinum-Compute-ServerPool for Pool Assignment.
- b. Click **Up** radio button to select the power state to be applied when this profile is associated with the server.
- c. Select PlatinumSvrQual for Server Pool Qualification.
- d. Do not check **Restrict Migration**.

**25. Click Next.**

**Figure 144** *Associating the Server Pool with the Service Profile*

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The left sidebar lists the steps: 1. Identify Service Profile Template, 2. Networking, 3. Storage, 4. Zoning, 5. vNIC/vHBA Placement, 6. Server Boot Order, 7. Maintenance Policy, 8. Server Assignment (highlighted), and 9. Operational Policies. The main panel is titled 'Server Assignment' and contains the following text: 'Optionally specify a server pool for this service profile template.' Below this is a sub-section titled 'Server Assignment' with the instruction: 'You can select a server pool you want to associate with this service profile template.' There is a dropdown menu for 'Pool Assignment' set to 'Platinum-Compute-ServerPool' and a '+ Create Server Pool' button. Below the dropdown, there is a section for power state selection: 'Select the power state to be applied when this profile is associated with the server.' with radio buttons for 'Up' (selected) and 'Down'. Below this is a section for server pool qualification: 'The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.' There is a dropdown menu for 'Server Pool Qualification' set to 'PlatinumSvrQual'. At the bottom, there is a checkbox for 'Restrict Migration' which is unchecked.

**26. Operational Policies Section**

- a. Select TenantA-BIOS for BIOS Policy.
  - b. Select BareMetal for External IPMI Management Configuration.
  - c. Select Platinum-Compute-IP for Management IP Address Policy.
  - d. Select default for Monitoring Configuration.
  - e. Select default for Power Control Policy Configuration.
  - f. Select default for Scrub Policy.
27. Click **Next**.

**Figure 145**      *Defining the attributes for the Operational Policy*

Create Service Profile Template

## Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. ✓ Server Assignment
9. ✓ **Operational Policies**

### Operational Policies

Optionally specify information that affects how the system operates.

#### BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy: TenantA-BIOS + Create BIOS Policy

#### External IPMI Management Configuration

If you want to access the CIMC on the server externally, select an IPMI access profile. The users and passwords in that profile will be populated into the CIMC when the profile is associated with the server.

IPMI Access Profile: BareMetal + Create IPMI Access Profile

To enable Serial over LAN access to the server, select an SoL configuration profile.

SoL Configuration Profile: Select a Policy to use (no SoL Access by defa...) + Create Serial over LAN Policy

This service profile will not have Serial over LAN access.

#### Management IP Address

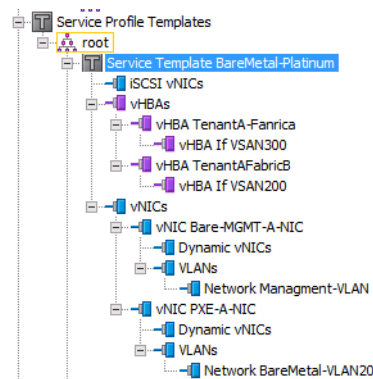
Management IP Address Policy: Platinum-Compute-IP(0/0)

+ Create IP Pool

The IP address will be automatically assigned from the selected pool.

**WARNING:** The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

**Figure 146** Summary of the BareMetal-Platinum Service Profile Template



## CloudPlatform Host Preparations

This section outlines VMWare ESXi 5.1 Host preparations required for multi-tenant to host cloud services with cloud based service level definition.

The cloud host preparations can be divided into three major sections as following:

- Creating Host
- Installing Host
- Configuring Host

### Creating Host

This section outline steps required to create cloud host compute, network, and storage infrastructure based on the service levels definition required to host cloud services in multi-tenantcloud environment. Cisco UCS offers service profile templates with pre-defined cloud host compute, network and storage resources based on service levels. Four Service Profiles Templates were created to meet the Cloud service level agreement requirements.

The steps below show cloud host creation based on Platinum service levels defined by service profile template Platinum-Template for TenantA Organization unit.

In this section the following tasks are described in detail:

- Create Service Profile based on a Service Profile Template
- Associate the Service Profile with Server

### Creating a Service Profile based on a Service Profile Template

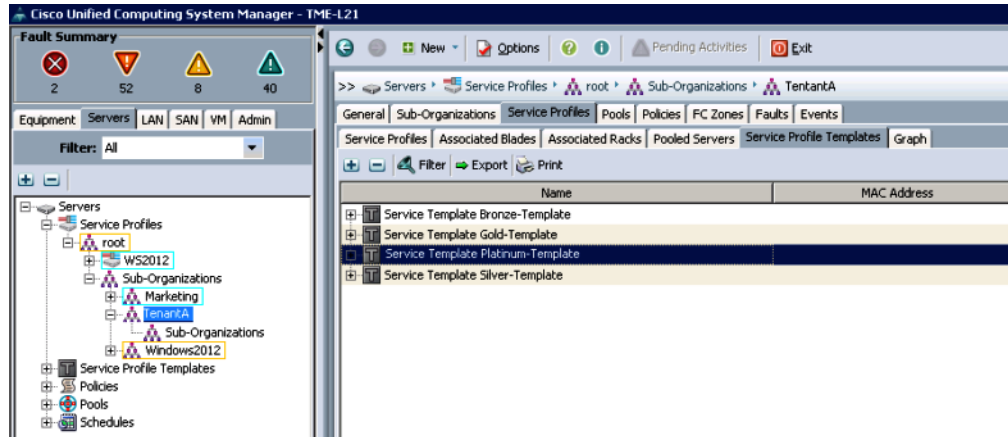
This section details the creation of service profile based on the Service Profile Template we created above.

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:

1. Click the **Server** tab in the left pane.
2. Choose **Service Profiles > Sub-Organization**. Expand TenantA.

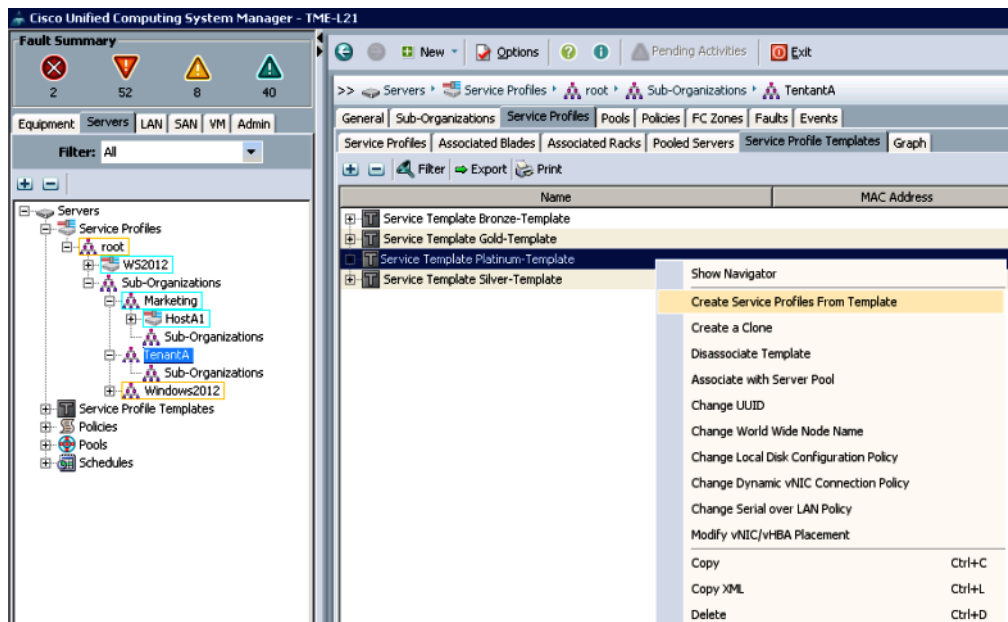
3. In the right pane choose **Service Profiles > Service Profile Template > Service Template Platinum-Template**.

**Figure 147** *Selecting the Service Profile Template*



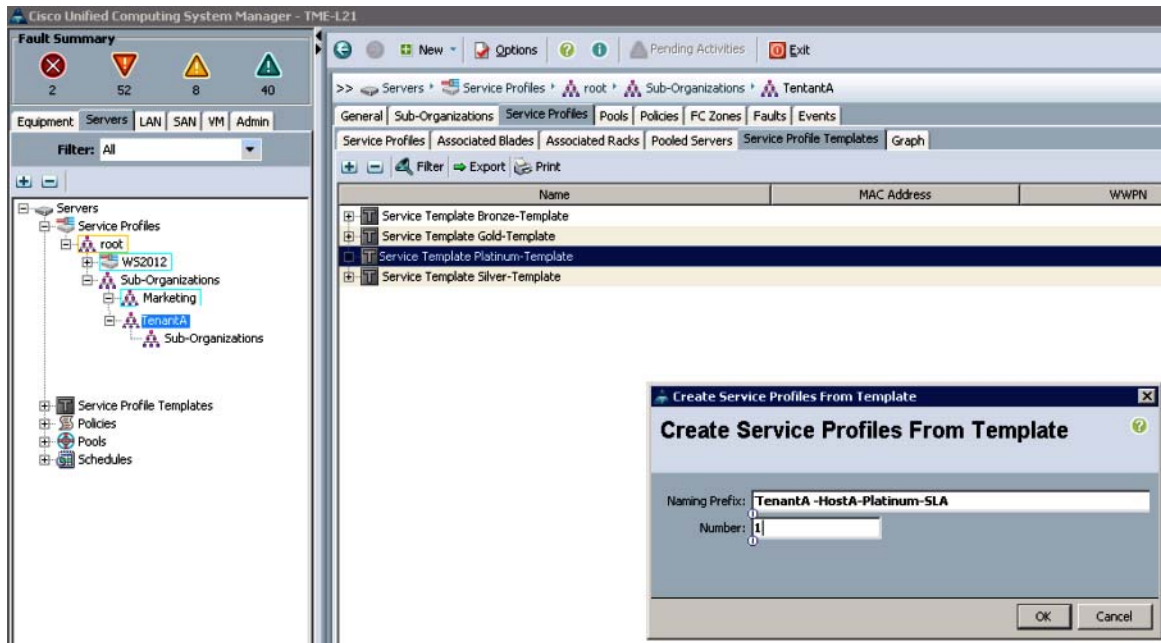
4. Right-click **Create Service Profile from Template**.

**Figure 148** *Selecting the Service Profile Form Template*



5. Type **TenantA-HostA-Platinum-SLA** in Naming Prefix, and type a number corresponding to the number of Service Profiles to be created in the Number text box. In our case, we entered 1.
6. Click **OK**.

**Figure 149**      *Adding the Name and Number to the Profile*



## Associating the Service Profile with Server

1. Under the TenantA the newly created Service Profile TenantA-HostA-Platinum-SLA1 is displayed and will be in associating phase.



**Figure 150**      *Displaying the Service Profile Status Details*

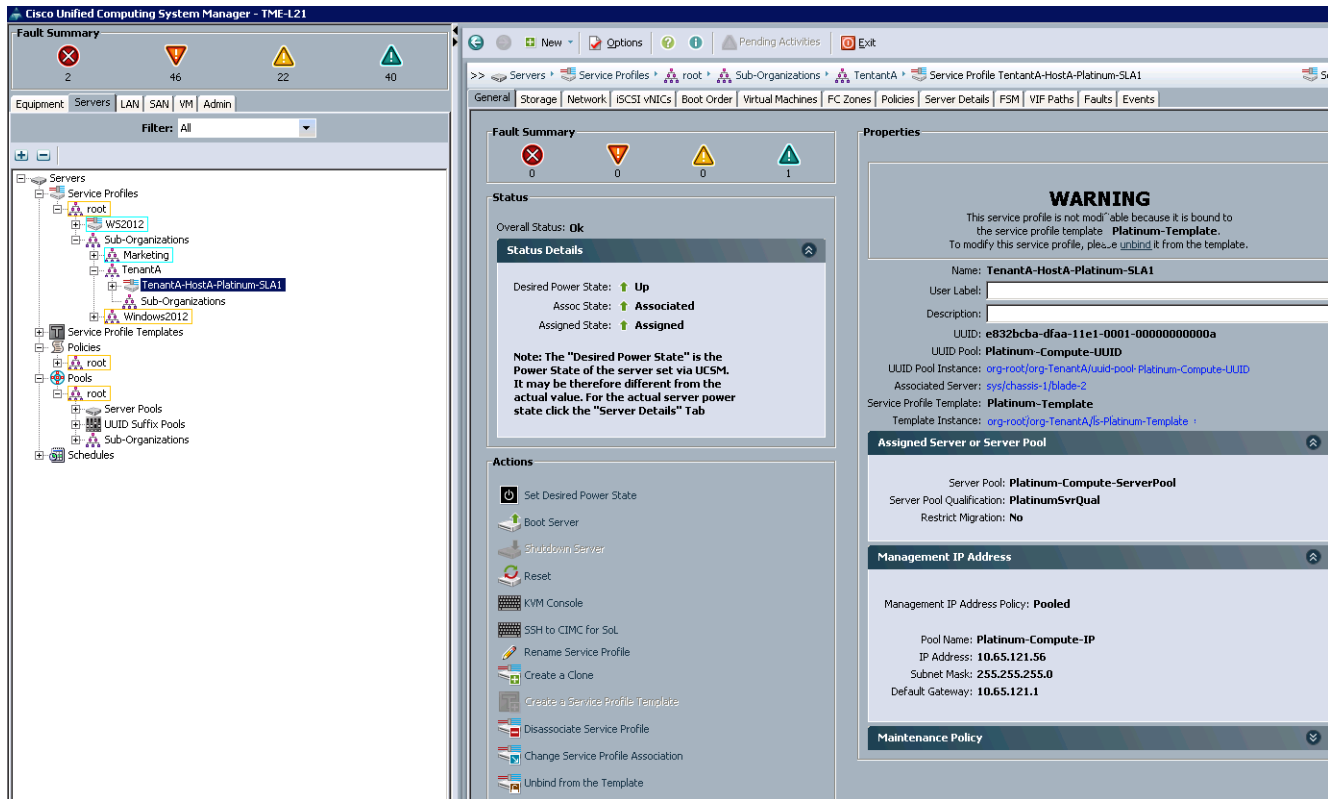
The screenshot displays the Cisco Unified Computing System Manager (UCSM) interface. On the left, a navigation tree shows the hierarchy: Servers > Service Profiles > root > Sub-Organizations > Marketing > TenantA > TenantA-HostA-Platinum-SLA1. The main panel is divided into three sections:

- Fault Summary:** Shows overall status as 'Ok' with 0 faults. Status details include:
  - Desired Power State: Up
  - Assoc State: Associated
  - Assigned State: Assigned
 A note states: "The 'Desired Power State' is the Power State of the server set via UCSM. It may be therefore different from the actual value. For the actual server power state click the 'Server Details' Tab".
- Properties:** Contains a 'WARNING' message: "This service profile is not modifiable because it is bound to the service profile template Platinum-Template. To modify this service profile, please unbind it from the template." Below this, it lists:
  - Name: TenantA-HostA-Platinum-SLA1
  - User Label: (empty)
  - Description: (empty)
  - UUID: e832bcba-dfaa-11e1-0001-00000000000a
  - UUID Pool: Platinum-Compute-UUID
  - UUID Pool Instance: org-root/org-TenantA/uuid-pool-Platinum-Compute-UUID
  - Associated Server: sys/chassis-1/blade-2
  - Service Profile Template: Platinum-Template
  - Template Instance: org-root/org-TenantA/ls-Platinum-Template
- Assigned Server or Server Pool:** Shows:
  - Server Pool: Platinum-Compute-ServerPool
  - Server Pool Qualification: PlatinumSvrQual
  - Restrict Migration: No
- Management IP Address:** Shows:
  - Management IP Address Policy: Pooled
  - Pool Name: Platinum-Compute-IP
  - IP Address: 10.65.121.56
  - Subnet Mask: 255.255.255.0
  - Default Gateway: 10.65.121.1
- Maintenance Policy:** (Section header, details not visible)

The bottom section, **Actions**, lists various operations: Set Desired Power State, Boot Server, Shutdown Server, Reset, KVM Console, SSH to CIMC for Sol, Rename Service Profile, Create a Clone, Create a Service Profile Template, Disassociate Service Profile, Change Service Profile Association, and Unbind from the Template.

2. After few minutes Service Profile TenantA-HostA-Platinum-SLA will be associated with the available blade defined in the server pool.

**Figure 151** Verifying the Service Profile Associated with the Blade



## Installing Host

This section outlines steps required to install VMware vSphere ESXi 5.1 operating system on a cloud host. Cisco UCS offers a KVM console which is a video over IP representation of the video output on the blade. The KVM console access to server blades in Cisco UCS is conceptually similar to any industry standard KVM console access to the blade. Once the server profile association is complete; KVM console can be accessed. The KVM console is an interface accessible from the Cisco UCS Manager GUI or the KVM Launch Manager that emulates a direct KVM connection.

The steps below show installation of the ESXi 5.1 hypervisor on newly associated blade serve using service profile TenantA-HostA-Platinum-SLA on TenantA Organization unit.

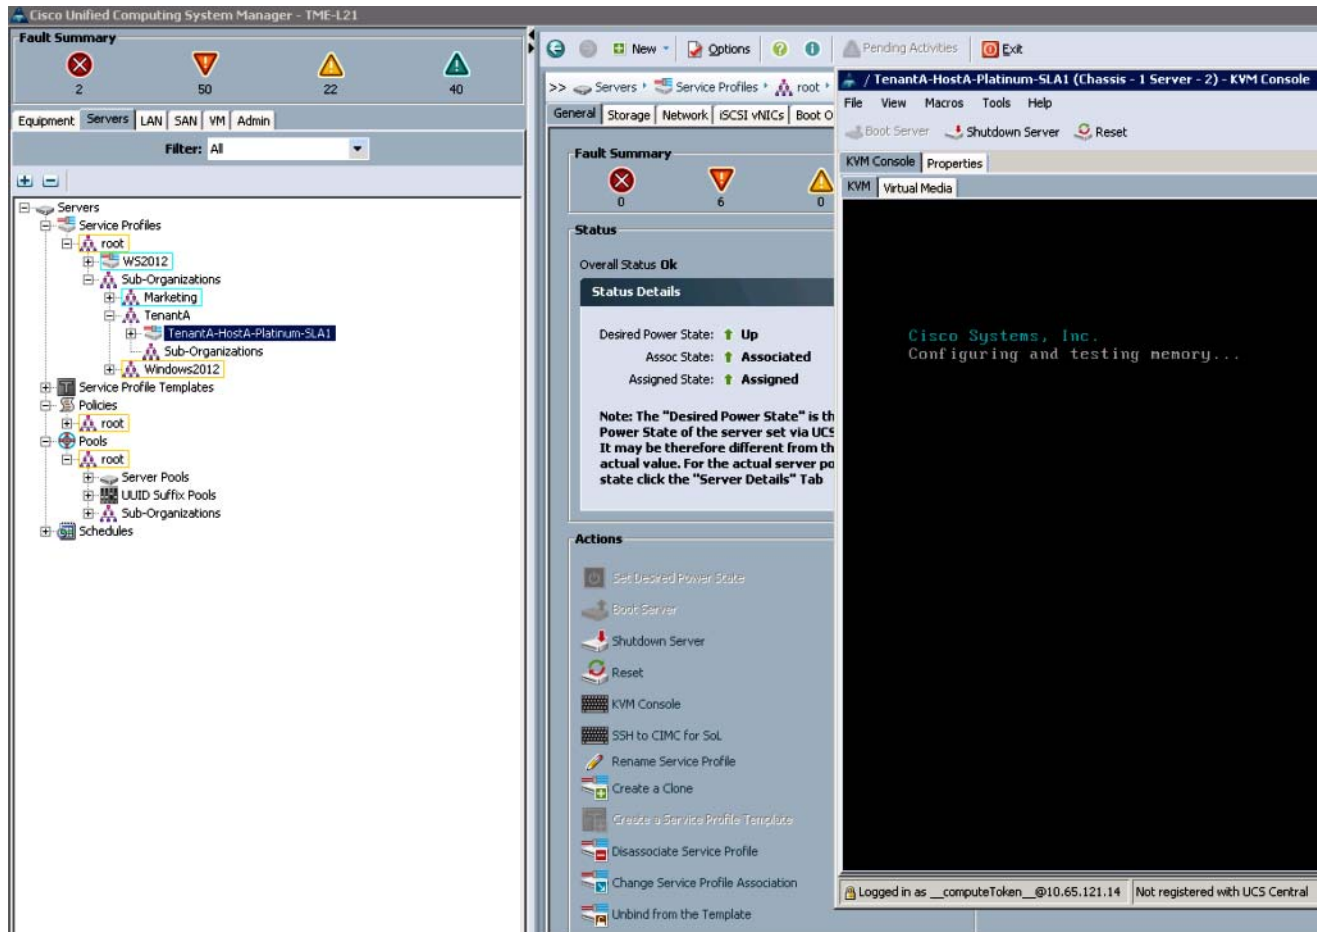
In this section the following tasks are described:

- Connecting to the blade server KVM console
- Installing the VMware ESXi 5.1 hypervisor

Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:

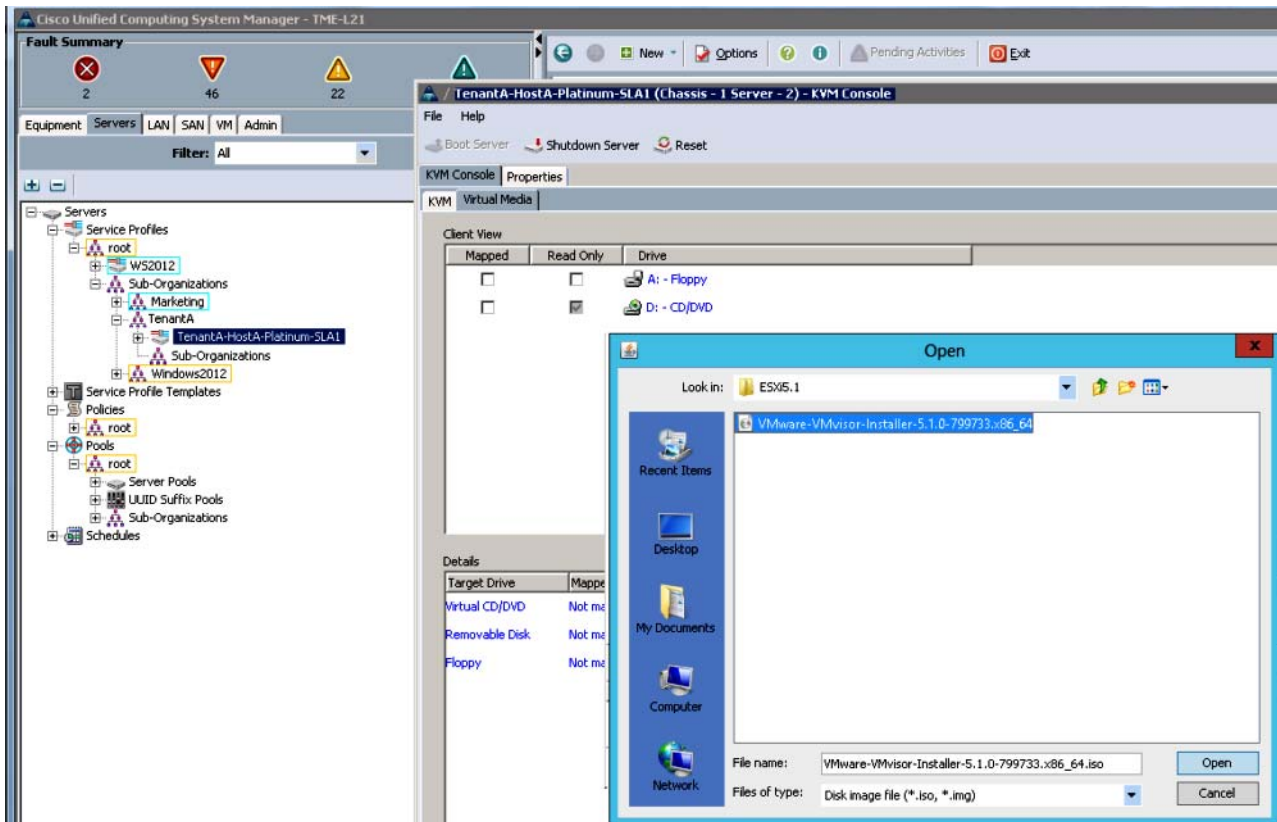
1. Click the **Server** tab in the left pane.
2. Choose **Service Profiles expand > Sub-Organization**. Expand TenantA.
3. Expand TenantA. Select Tenant A-HostA-Platinum-SLA1. In the right pane under Actions: click **KVM Console**.
4. A pop up window /TenantA-HostA-Platinum-SAL1 (Chassis -1 server -2) KVM Console launches.

**Figure 152**      **Launching the KVM Console**



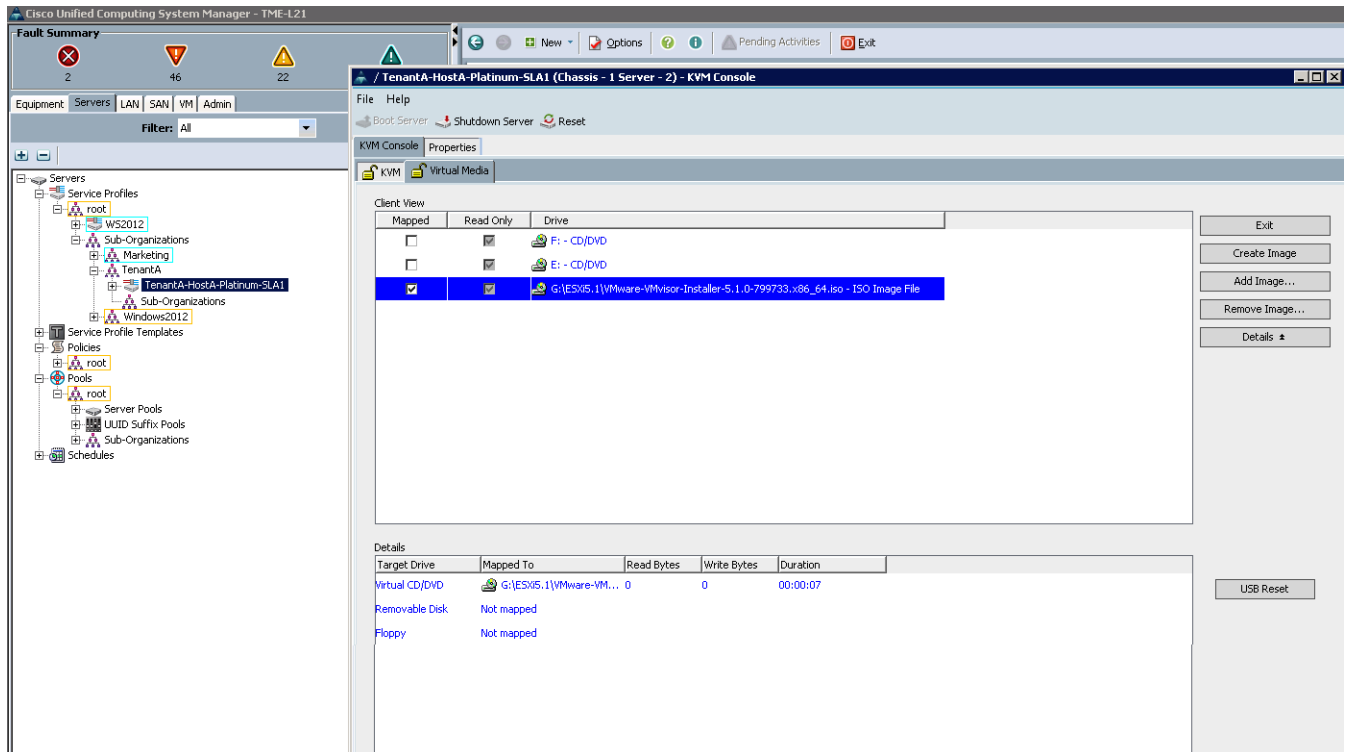
5. On the KVM Console window choose **Virtual Media > Add Image > Browse ESX Server 5.1 ISO** image. Click **Open**.

**Figure 153** Adding the VMWare ESXi Server Image



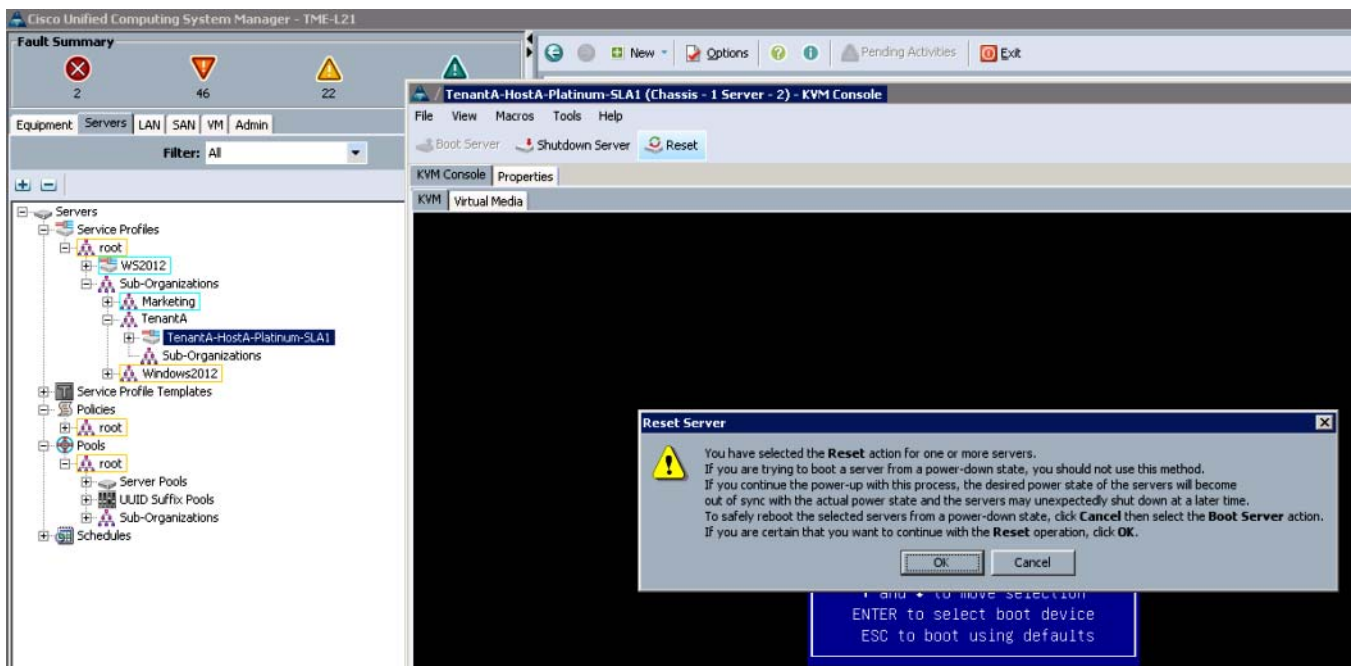
6. Check Mapped check box, and click the **KVM** tab in the KVM Console.

**Figure 154**      **Selecting the VMWare ESXi Image**



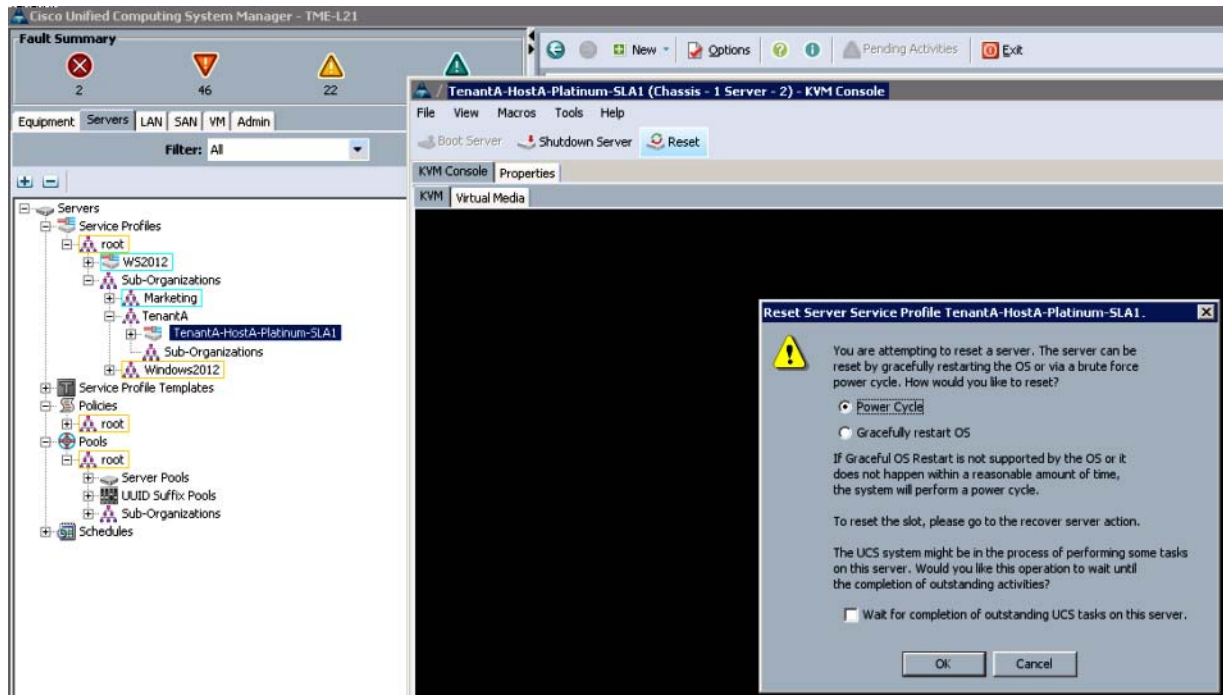
7. Click **Reset** button and click **OK**.

**Figure 155**      **Resetting the Server**




8. Click **Power Cycle** radio button and then click **OK**.

**Figure 156**      *Resetting the Power Cycle for the Server Service Profile*




9. At the boot menu select Cisco Virtual CD/DVD 1.22 option and press Enter.

**Figure 157**      *Selecting the Boot Device*



10. Press enter at the boot prompt. The VMWare ESXi Server is launched.

**Figure 158**      *Launching the VMWare ESXi Server*

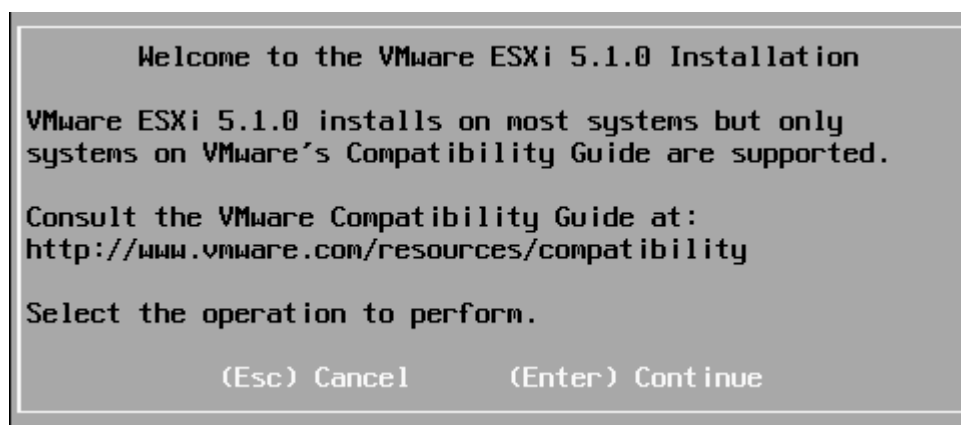


## Installing the VMWare ESXi Server

To install the VMWare ESXi Server launched at the KVM console, follow these steps:

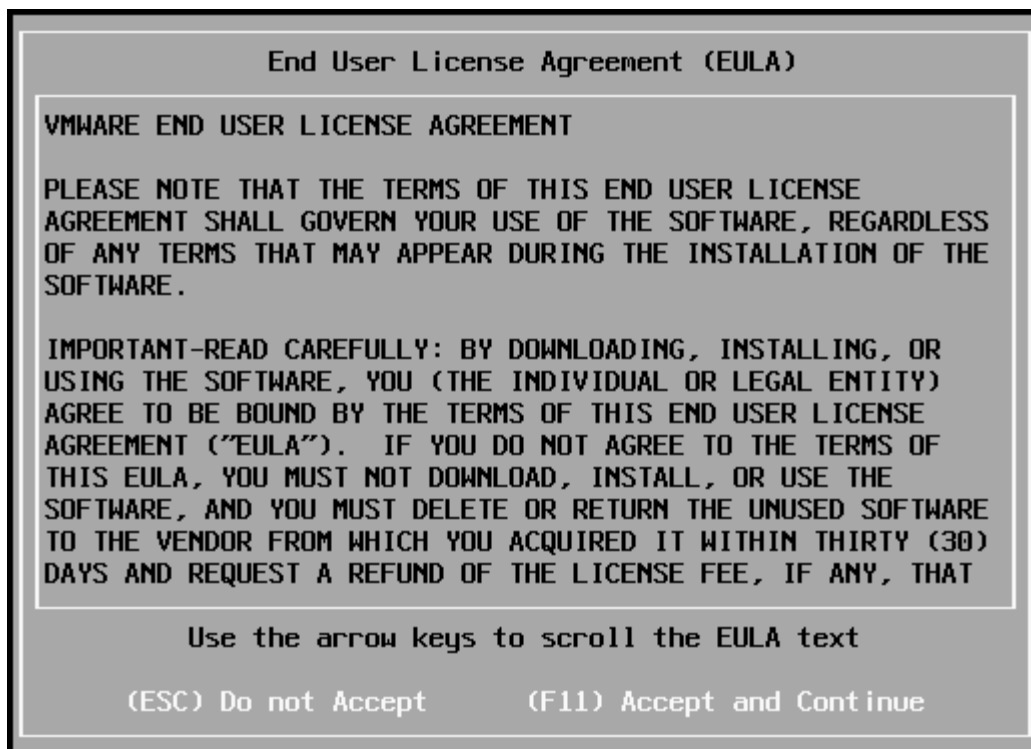
1. Installing the VMWare ESXi Server.
2. Press Enter (Continue) key.

**Figure 159**      *Selecting the Enter Option*



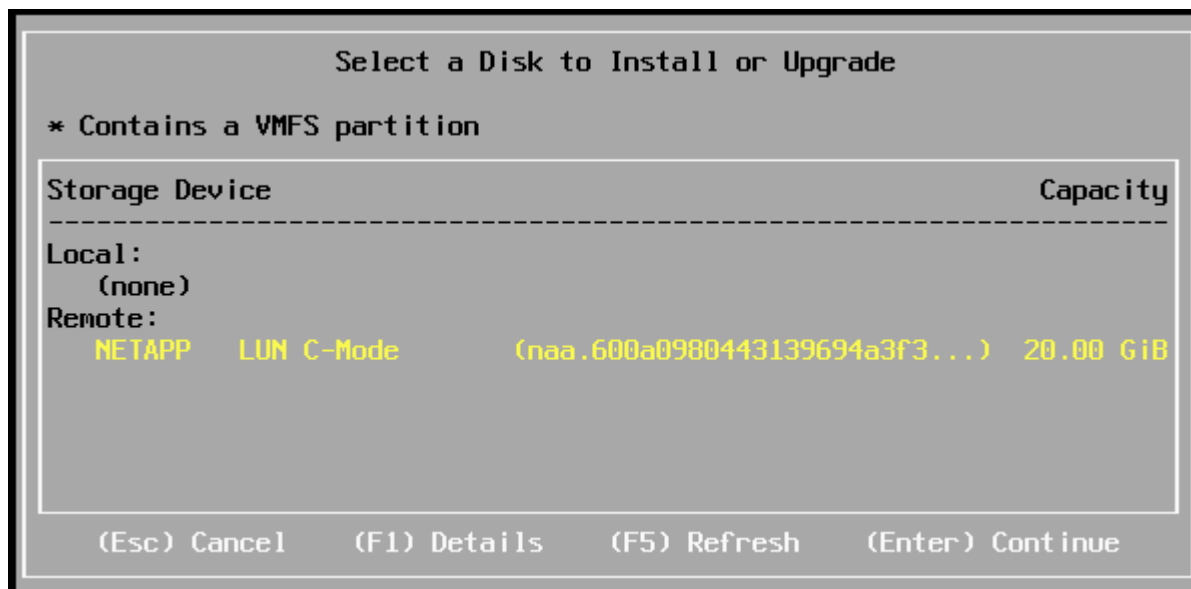
3. Select F11 to Accept and Continue

Figure 160 End User License



4. Select NetApp LUN C-Mode 20GB
5. Press Enter (Continue).

Figure 161 Selecting Disk



6. Select US Default.



7. Press (Enter) to Continue.

**Figure 162**      *Selecting Keyboard layout*




8. Enter Password.
9. Press (Enter) to Continue.

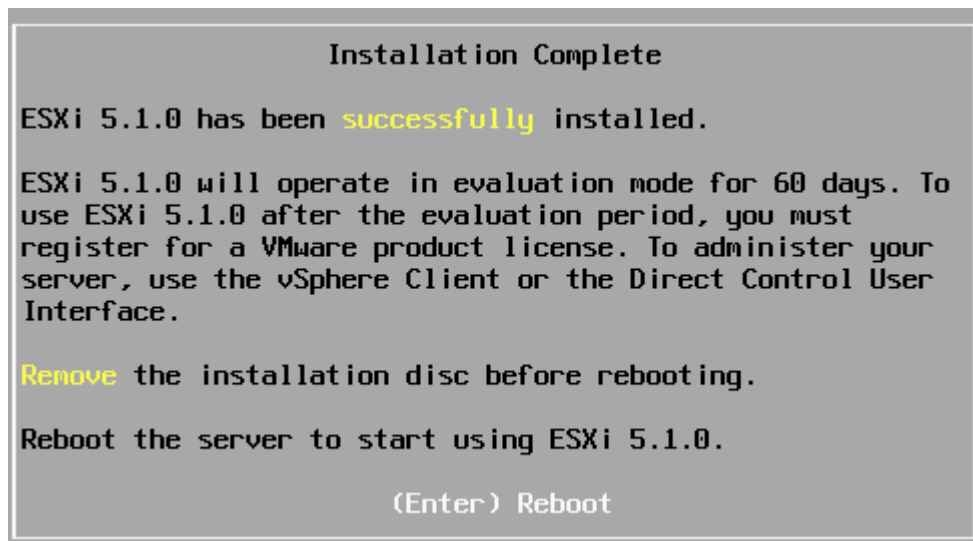
**Figure 163**      *Entering Password*



10. Enter F11 (Install) to confirm installation

**Figure 164** *Defining the Disk for ESXi Install*

11. After installation a reboot is required and VMWare ESXi 5.1 Host Installation Complete information is displayed.

**Figure 165** *Displaying the VMWare ESXi Installation Status*

## Configuring Host

This section outlines steps required to configure VMWare ESXi 5.1 Server on cloud host after installation is completed with service levels definition required to host cloud services in multi-tenant cloud environment.

The host configuration is divided into two tasks.

- Network task

This will define design and configuration of ESXi Host Networking infrastructure to carry Tenants Management, guest and Primary or Secondary Storage (NFS) network data traffic. Also, it offers Physical and Logical isolation with high availability and load balancing using NIC Teaming and Load balancing features.

- Storage task

This will define design and configuration of ESXi Host Storage infrastructure to create tenants Primary Storage that are later referred by Citrix CloudPlatform to define Zone wide Primary Storage infrastructure for multi-tenants in cloud.

In this section the following tasks have been described in detail:

- Configuring VMWare ESXi Host Network
- Creating VMFS based DataStore on Fibre Channel Storage LUN

## Configuring VMWare ESXi Host Network

This section details the configuration of Virtual Network Switches with VMKernel Interfaces for Management, guest and NFS Storage network traffic.

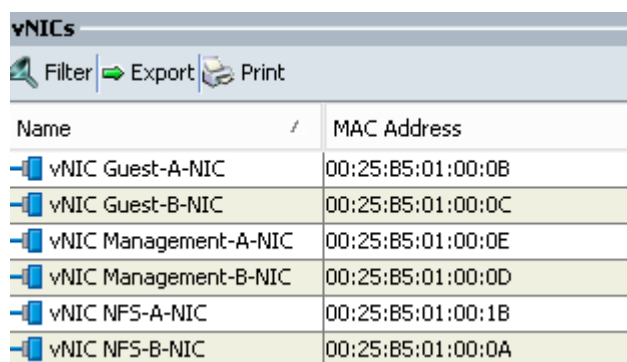
Each Network Adapters Uplinked to Virtual Network Switch should be matched with the correct Cisco UCS virtual NIC (vNICs), VLANs defined in Service Profile TenantA-HostA-Platinum-SLA1.

Table 32 provides Citrix ESXi Host Network Bond Ethernet Interfaces mapping with UCS vNICs.

**Table 32** ESXi vSwitch NIC Interfaces Mapping with UCS vNICs

| Virtual Network Switch (vSwitch) | VNIC Ethernet Interfaces               | UCS Service Profile vNICs                                                |
|----------------------------------|----------------------------------------|--------------------------------------------------------------------------|
| TenantA-Mgmt-Uplink              | 00:25:B5:01:00:0E<br>00:25:B5:01:00:0D | Management-A-NIC 00:25:B5:01:00:0E<br>Management-B-NIC 00:25:B5:01:00:0D |
| TenantA-guest-Uplink             | 00:25:B5:01:00:0B<br>00:25:B5:01:00:0C | guest-A-NIC 00:25:B5:01:00:0B<br>guest-B-NIC 00:25:B5:01:00:0C           |
| TenantA-NFS-Uplink               | 00:25:B5:01:00:1B<br>00:25:B5:01:00:0A | NFS-A-NIC 00:25:B5:01:00:1B<br>NFS-B-NIC 00:25:B5:01:00:0A               |

**Figure 166** MAC Addresses Corresponding to the Cisco UCS Service Profile vNICs



| vNICs                 |                   |
|-----------------------|-------------------|
| Name                  | MAC Address       |
| vNIC Guest-A-NIC      | 00:25:B5:01:00:0B |
| vNIC Guest-B-NIC      | 00:25:B5:01:00:0C |
| vNIC Management-A-NIC | 00:25:B5:01:00:0E |
| vNIC Management-B-NIC | 00:25:B5:01:00:0D |
| vNIC NFS-A-NIC        | 00:25:B5:01:00:1B |
| vNIC NFS-B-NIC        | 00:25:B5:01:00:0A |

## Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, follow these steps on each ESXi host:

**ESXi Host TenantA-ESX-HostA-Platinum-SLA1**

To configure the VM-Host-Infra-01 ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Choose the Configure the Management Network option and press Enter.
4. Choose the VLAN (Optional) option and press Enter.
5. Enter the <<var\_mgmt\_vlan ID>> 603 and press Enter.




---

**Note** In this guide, we input 603 for var\_mgmt\_vlan ID label

---

6. From the Configure Management Network menu, choose IP Configuration and press Enter.
7. Choose the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the first ESXi host: <<10.65.121.163>> <<var\_vm\_host\_infra\_01\_ip>>.




---

**Note** In this guide, we input 10.65.121.163 for var\_vm\_host\_infra\_01\_ip label

---

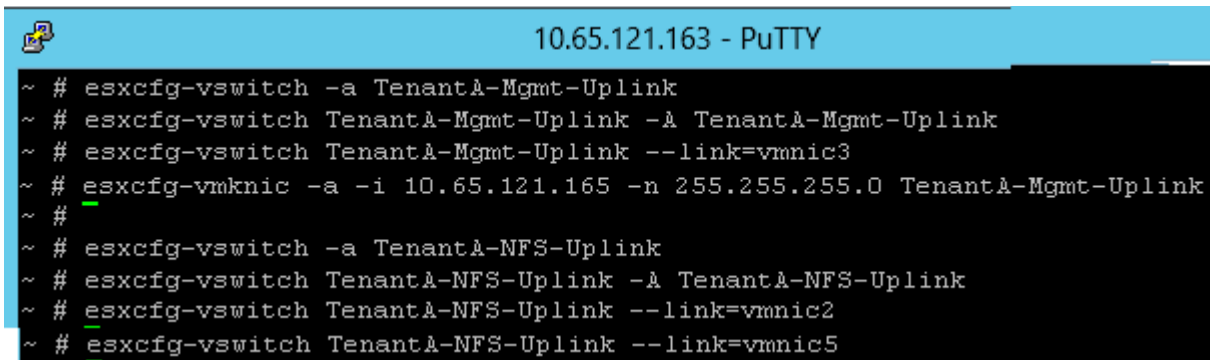
9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Choose the IPv6 Configuration option and press Enter.
13. Using the spacebar, deselect Enable IPv6 (restart required) and press Enter.
14. Choose the DNS Configuration option and press Enter.
15. Enter the DNS IP Address for the first ESXi host.

To Configure Management and NFS Virtual Switches, follow these steps:

1. Perform SSH to ESXi Host <<10.65.121.163 var\_vm\_host\_infra\_01\_ip>>.
2. Provide Root and <Password> credentials
3. Create Virtual Switches <TenantA-Mgmt-Uplink> & <TenantA-NFS-Uplink> with appropriate Network Adapters defined in Service profile.

```
~ # esxcfg-vswitch -aTenantA-Mgmt-Uplink
~ #esxcfg-vswitch -a TenantA-NFS-Uplink
~ # esxcfg-vswitch TenantA-Mgmt-Uplink -A TenantA-Mgmt-Uplink
~ # esxcfg-vswitch TenantA-Mgmt-Uplink --link=vmnic3
~ # esxcfg-vmknics -a -i 10.65.121.165 -n 255.255.255.0 TenantA-Mgmt-Uplink
~ # esxcfg-vswitch TenantA-NFS-Uplink -A TenantA-NFS-Uplink
~ # esxcfg-vswitch TenantA-NFS-Uplink --link=vmnic2 vmnic5
```

**Figure 167** ESXi vSwitch TenantA-Mgmt-Uplink and TenantA-NFS-Uplink Creation Commands

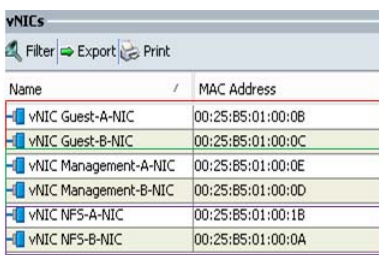


```

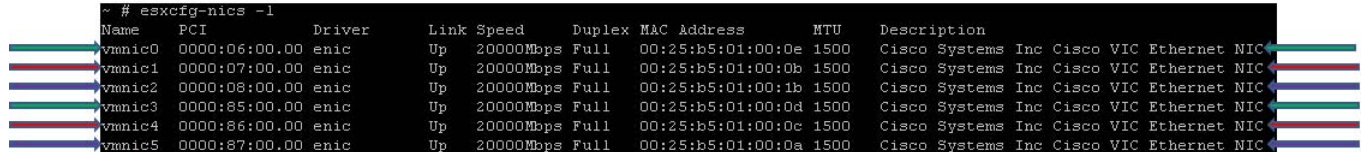
10.65.121.163 - PuTTY
~ # esxconfig-vswitch -a TenantA-Mgmt-Uplink
~ # esxconfig-vswitch TenantA-Mgmt-Uplink -A TenantA-Mgmt-Uplink
~ # esxconfig-vswitch TenantA-Mgmt-Uplink --link=vmnic3
~ # esxconfig-vmknic -a -i 10.65.121.165 -n 255.255.255.0 TenantA-Mgmt-Uplink
~ #
~ # esxconfig-vswitch -a TenantA-NFS-Uplink
~ # esxconfig-vswitch TenantA-NFS-Uplink -A TenantA-NFS-Uplink
~ # esxconfig-vswitch TenantA-NFS-Uplink --link=vmnic2
~ # esxconfig-vswitch TenantA-NFS-Uplink --link=vmnic5

```

**Figure 168** ESXi vSwitch TenantA-Mgmt-Uplink and TenantA-NFS-Uplink Pnic Mapping with UCS Service Profile static vNIC



| Name                  | MAC Address       |
|-----------------------|-------------------|
| vNIC Guest-A-NIC      | 00:25:b5:01:00:0B |
| vNIC Guest-B-NIC      | 00:25:b5:01:00:0C |
| vNIC Management-A-NIC | 00:25:b5:01:00:0E |
| vNIC Management-B-NIC | 00:25:b5:01:00:0D |
| vNIC NFS-A-NIC        | 00:25:b5:01:00:1B |
| vNIC NFS-B-NIC        | 00:25:b5:01:00:0A |

```

~ # esxconfig-nics -l
Name PCI Driver Link Speed Duplex MAC Address MTU Description

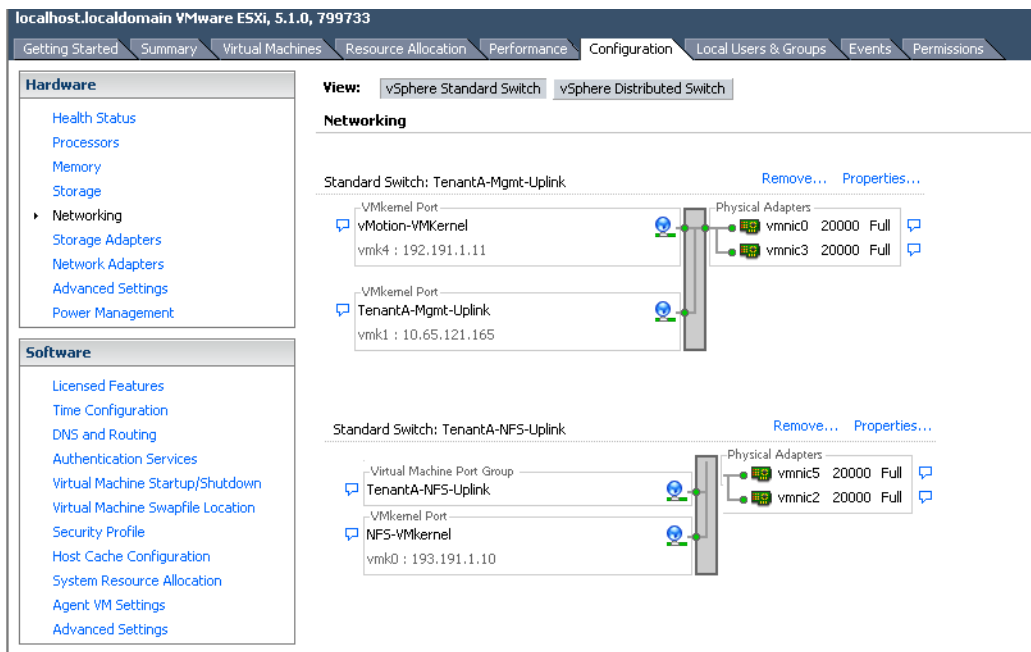
vmnic0 0000:06:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:0e 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic1 0000:07:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:0b 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic2 0000:08:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:1b 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic3 0000:85:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:0d 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic4 0000:86:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:0e 1500 Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic5 0000:87:00.00 enic Up 20000Mbps Full 00:25:b5:01:00:0a 1500 Cisco Systems Inc Cisco VIC Ethernet NIC

```

4. Download VMware vSphere Client and vSphere Remote CLI
  - a. Open a Web browser on the management workstation and navigate to the VM-Host-Infra-01 newly created vmkernel management IP address <10.65.121.165> on TenantA-Mgmt-Uplink Port Group.
  - b. Download and install both the vSphere Client and the Windows version of vSphere Remote Command Line.
5. Log in to VMware ESXi Hosts by Using VMware vSphere Client
  - a. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Infra-01 as the host you are trying to connect to: <<10.65.121.165 var\_vm\_host\_infra\_01\_ip>>.
  - b. Enter root for the user name.
  - c. Enter the root password.
  - d. Click **Login** to connect.
6. Delete default vSwitch0 and add network adapter vmnic0 to TenantA-Mgmt-Uplink
  - a. From each vSphere Client, choose the host in the inventory.
  - b. Click the **Configuration** tab.
  - c. Click **Networking** pane.

- d. On the right pane, click **Remove** on vSwitch0.
  - e. Click **Properties**.
  - f. Click **Network Adapters**.
  - g. Click **Add** and select **vmnic0** under Unclaimed Adapters.
  - h. Click Next and again Next and click Finish to add network adapter
7. Add VMKernel interface on TenantA-Mgmt-Uplink virtual switch to access vMotion data traffic
    - a. From each vSphere Client, choose the host in the inventory.
    - b. Click the **Configuration** tab.
    - c. Click **Networking**.
    - d. On right pane click **Properties** on TenantA-Mgmt-Uplink
    - e. Click **Add** and select VMKernel Connection type radio button.
    - f. Click **Next**.
    - g. Type vMotion-VMKernel in Network Label text box.
    - h. Type 192 (vMotion ID) in VLANID text box.
    - i. Select **Use this port group for vMotion** check box.
    - j. Click **Use the following IP Setting** radio button.
    - k. Type 192.191.1.10 in **IP Address** text box, and 255.255.255.0 in **Subnet Mask**.
    - l. Keep the Default Gateway IP Address.
    - m. Click **Next**.
    - n. Click **Finish**.

**Figure 169** Final ESX Virtual Switch Configuration



8. Add VMKernel interface on TenantA-NFS-Uplink virtual switch to access NFS data traffic

- a. From each vSphere Client, choose the host in the inventory.
  - b. Click the **Configuration** tab.
  - c. Click **Networking**.
  - d. On right pane click Properties on TenantA-NFS-Uplink.
  - e. Click **Add** and Select VMKernel Connection type Radio Button and click Next
  - f. Type NFS-VMKernel in Network Label: text box and click Next
  - g. Select **Use the following IP Setting** radio button
  - h. Type 193.191.1.1.10 in IP Address text box, 255.255.255.0 in Subnet Mask.
  - i. Keep the Default Gateway IP Address and click **Next**.
  - j. Click **Finish**.
9. To load the updated versions of the enic and fnic drivers for the Cisco VIC, follow these steps for the hosts on each vSphere Client:
- a. From each vSphere Client, choose the host in the inventory.
  - b. Choose the **Summary** tab to view the environment summary.
  - c. From **Resources Configuration** tab choose **Storage**.
  - d. Right-click **datastore1** and choose **Browse Datastore**.
  - e. Click the fourth button and choose **Upload File**.
  - f. Navigate to the saved location for the downloaded enic driver version and choose net-enic-2.1.2.38-1OEM.500.0.0.472560.x86\_64.zip.
  - g. Click **Open** to open the file.
  - h. Click **Yes** to upload the .zip file to datastore1.
  - i. Click the fourth button and choose **Upload File**.
  - j. Navigate to the saved location for the downloaded fnic driver version and choose scsi-fnic-1.5.0.20-1OEM.500.0.0.472560.x86\_64.zip.
  - k. Click **Open** to open the file.
  - l. Click **Yes** to upload the .zip file to datastore1.
  - m. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
  - n. At the command prompt, run the following commands to account for each host (enic):
 

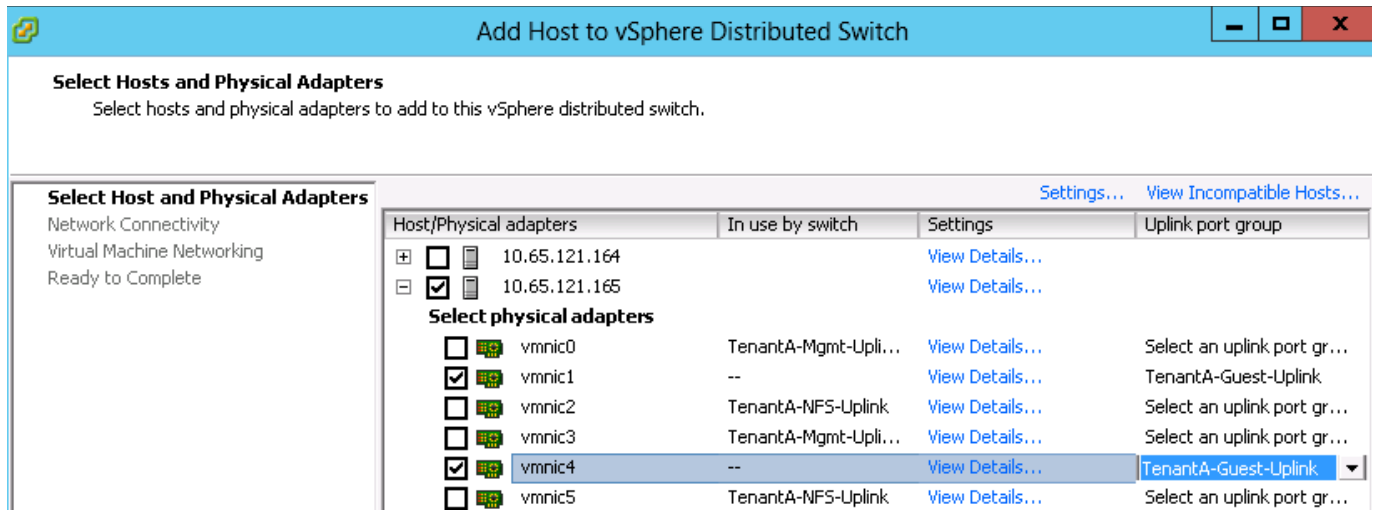
```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/net-enic-2.1.2.38-1OEM.500.0.0.472560.x86_64.zip
```
  - o. At the command prompt, run the following commands to account for each host (fnic):
 

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/scsi-fnic-1.5.0.20-1OEM.500.0.0.472560.x86_64.zip
```
  - p. From the vSphere Client, right-click each host in the inventory and choose Reboot.
  - q. Click Yes to continue.
  - r. Enter a reason for the reboot and click OK.
  - s. After the reboot is complete, log back in to both hosts using the vSphere Client.
10. To install the Virtual Ethernet Module (VEM) on the ESXi hosts, follow these steps:

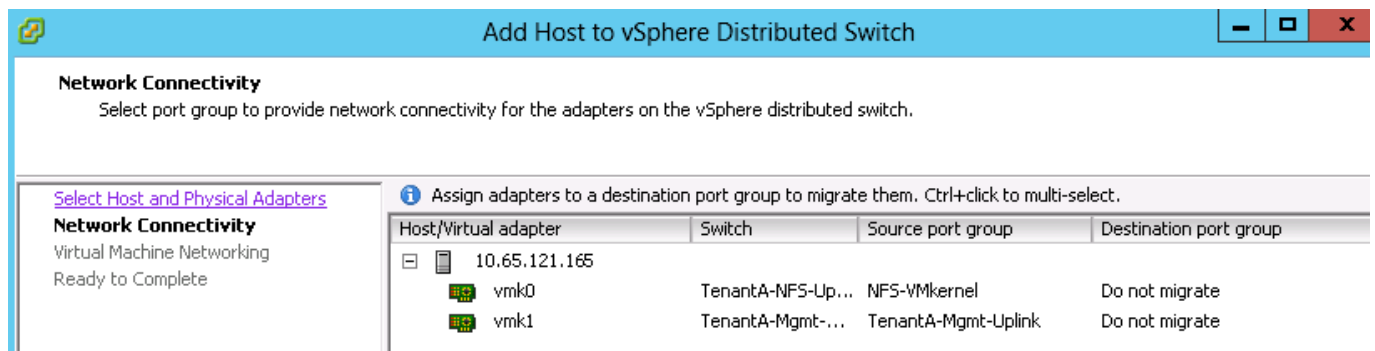
- a. Launch a Web browser to `http://<var_vsm_mgmt_ip>>`.
- b. Right-click the **cross\_cisco-vem-v152-4.2.1.1.2.1.1a.0-3.1.1.vib** hyperlink and choose **Save target as**.
- c. Save the file as `cross_cisco-vem-v152-4.2.1.1.2.1.1a.0-3.1.1.vib`, type All Files, on the desktop of the management workstation.
- d. From the main window in the vSphere Client connected to vCenter, click the first server in the list under the FlexPod Management cluster.
- e. Click the **Summary** tab.
- f. Under Storage on the right, right-click **infra\_datastore\_1** and choose **Browse Datastore**.
- g. Choose the root folder (/) and click the third button at the top to add a folder.
- h. Name the folder VEM and click **OK**.
- i. On the left, select the **VEM** folder.
- j. Click the fourth button at the top and choose **Upload File**.
- k. Navigate to the `cross_cisco-vem-v152-4.2.1.1.2.1.1a.0-3.1.1.vib` file and click **Open**.
- l. Click **Yes**. The VEM file should now appear in the VEM folder in the datastore.
- m. Open the VMware vSphere CLI command prompt.
- n. For each ESXi host in the VMware vSphere CLI, run the following command:
 

```
esxcli -s <Host Server IP> -u root -p <Root Password> software vib install -v
/vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.1.2.1.1a.0-3.1.1.vib
```
11. Add TenantA-guest-Uplink VMKernel PortGroup on Cisco Nexus 1000V DVS switch on ESX host
  - a. In the VMware vSphere Client connected to vCenter, choose **Home > Networking**.
  - b. Expand the vCenter, DataCenter, and Cisco Nexus 1000V folders.
  - c. Choose the Cisco Nexus 1000V switch.
  - d. Under **Basic Tasks** for the vSphere distributed switch, choose **Add a Host**.
  - e. Select host `<var_vm_host_infra_01_ip label> 10.65.121.165`.
  - f. Select `vmnic1` & `vmnic4` network adapters.
  - g. Select TenantA-guest-Uplink under Uplink port group list box.



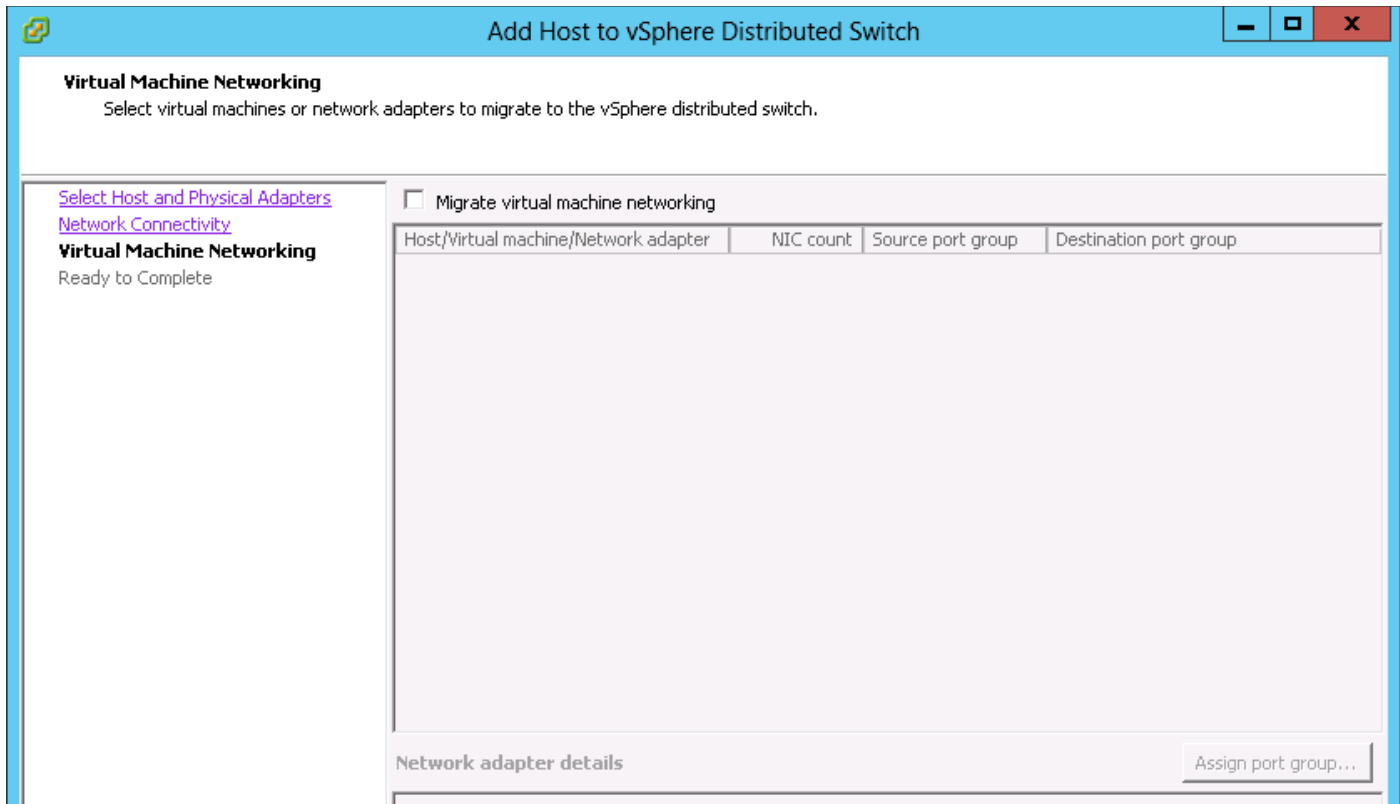
**Figure 170 Adding Host Network Adapters to Nexus 1000V DVS Switch**

- h. Click Next.
- i. Do not migrate VMKernel adapters vmk0 and vmk1 interfaces to Nexus 1000V DVS switch.
- j. Click Next.

**Figure 171 Do not Migrate VMKernel interfaces**

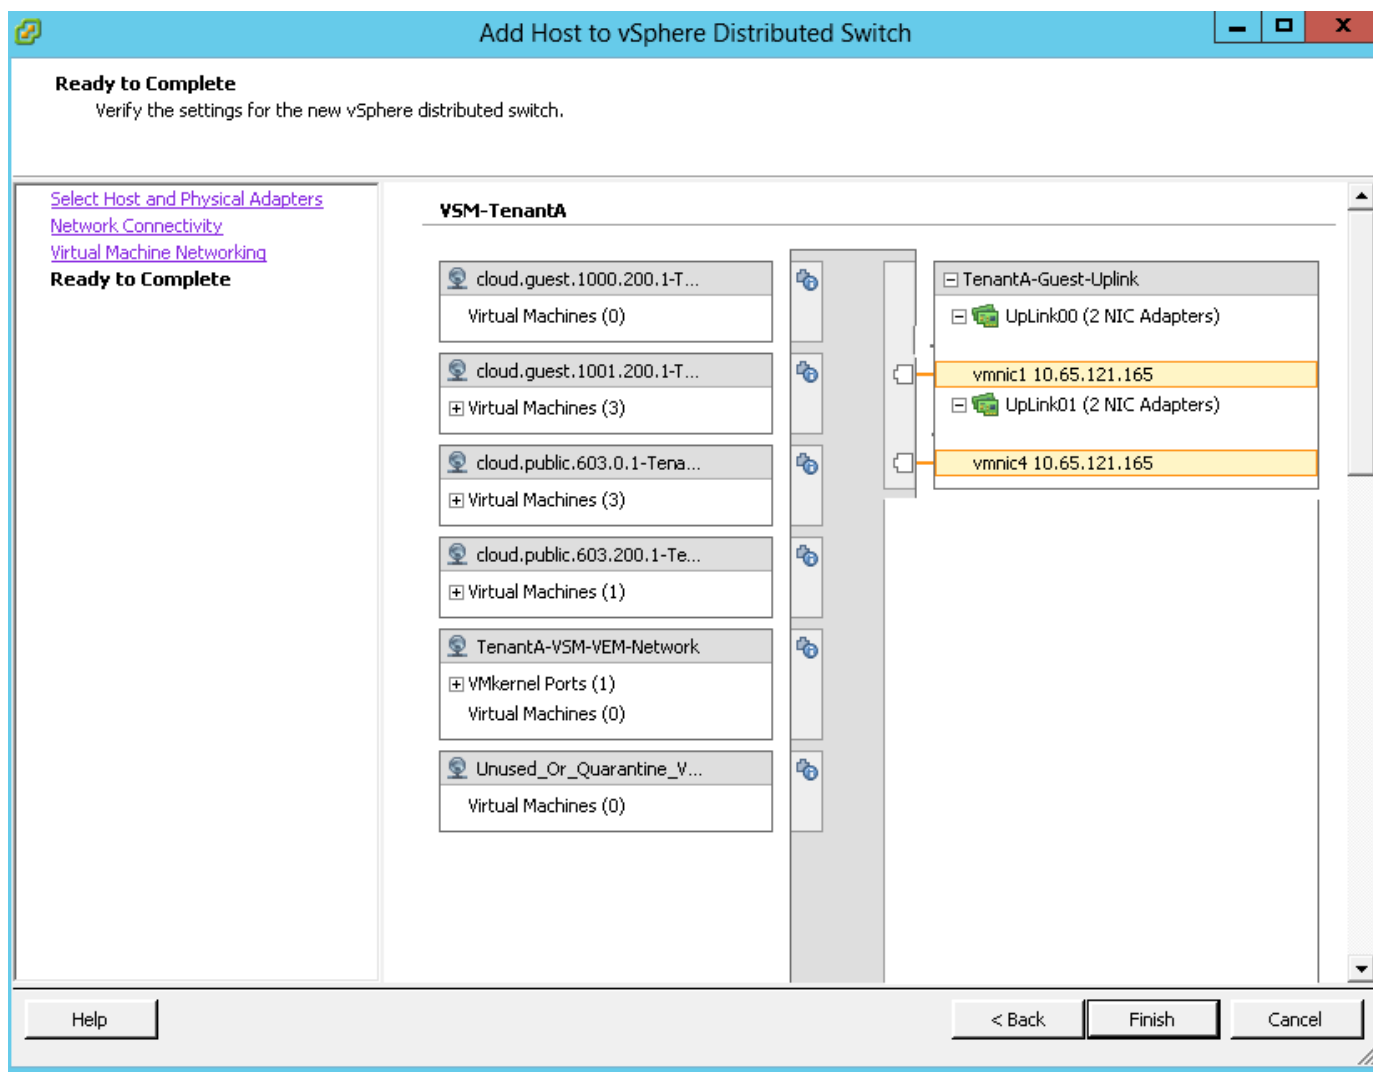
- k. Click Next.

**Figure 172**      **No Virtual Machine network to Migrate**

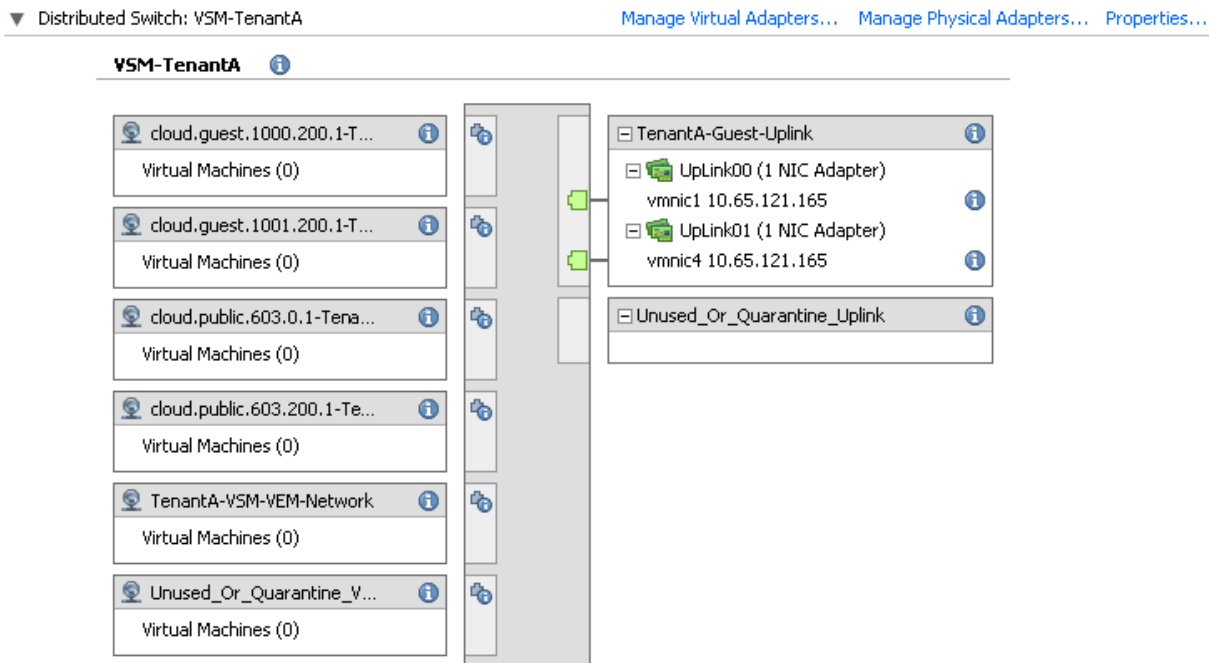


- l. Make sure you have two network adapters listed under TenantA-guest-Uplink Port Group.
- m. Click **Finish**.

**Figure 173**      **Add Network Adapters**



- n. Verify virtual adapters link status is Up and ESX Host <<10.65.121.165>> is been successfully added to Nexus 1000V DVS switch.

**Figure 174**      **Nexus 1000V DVS switch Configuration**

1. Add VMKernel interface Nexus 1000V DVS << VSM-TenantA>> on TenantA-guest-Uplink Port Group to access Tenant guest data traffic
  - a. From each vSphere Client, choose the host in the inventory.
  - b. Click the **Configuration** tab.
  - c. Click **Networking**.
  - d. On right pane view click **vSphere Distributed Switch Properties on TenantA-NFS-Uplink**.
  - e. Click **Manage Virtual Adapters**.
  - f. Click **Add**.
  - g. Click **New Virtual adapter** radio button.
  - h. Click **Next**.
  - i. Select VMKernel virtual Adapter Types Click Next
  - j. Select port group TenantA-VSM-VEM-Network in the list box.
  - k. Type IP Address <10.65.121.169>, Subnet Mask <255.255.255.0> and leave default gateway field.
  - l. Click **Next**.
  - m. Click **Finish**.
  - n. Click **Close**.
2. Verfiy VEM module on Nexus 1000V DVS switch <<10.65.121.170>>
  - a. Telnet to VSM IP Address <<10.65.121.70>>.
  - b. At command line type show module vem.

**Figure 175** VEM module on Nexus 1000V DVS

```

VSM-TenantA(config)#
VSM-TenantA(config)#
VSM-TenantA(config)# sh module vem
Mod Ports Module-Type Model Status
--- ---
3 248 Virtual Ethernet Module NA ok

Mod Sw Hw
--- ---
3 4.2(1)SV2(1.1a) VMware ESXi 5.1.0 Releasebuild-799733 (3.1)

Mod MAC-Address(es) Serial-Num
--- ---
3 02-00-0c-00-04-00 to 02-00-0c-00-04-80 NA

Mod Server-IP Server-UUID Server-Name
--- ---
3 10.65.121.165 babc32e8-aadf-e111-0001-000000000008 10.65.121.165

```

## NetApp Virtual Storage Console for VMware vSphere

### VSC 4.1 Preinstallation Considerations


The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.1.2:

- Protocol licenses (NFS and FCP)
- FlexClone (for provisioning and cloning only)
- Install VSC 4.1

To install the VSC 4.1 software, follow these steps:

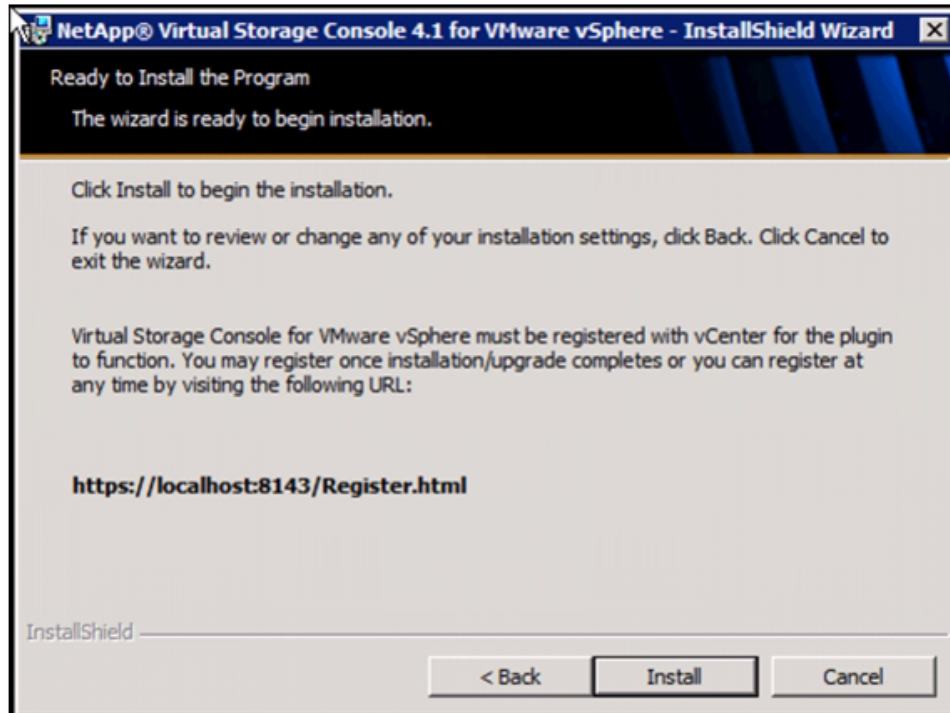
1. Using the instructions in section "Build Microsoft SQL Server VM," build a VSC and an OnCommand virtual machine with 4GB RAM, two CPUs, and one virtual network interface in the <<var\_ib-mgmt\_vlan\_id>> VLAN. The virtual network interface should be a VMXNET 3 adapter. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain. Install the current version of Adobe Flash Player on the VM. Install all Windows updates on the VM.
2. Log in to the VSC and OnCommand VM as the FlexPod admin user.
3. Download the x64 version of the Virtual Storage Console 4.1 at:  
<http://support.netapp.com/NOW/cgi-bin/software/?product=Virtual+Storage+Console&platform=VMware+vSphere> from the NetApp Support site.
4. Right-click the file downloaded and choose Run As Administrator.
5. Click **Yes** at the User Access Control warning.
6. On the Installation wizard welcome page, click **Next**.
7. Choose the backup and recovery capability.
8. Click **Next**.
9. Click **Next** to accept the default installation location.

**Figure 176** VSC Install Location



10. Click **Install**.

**Figure 177**      *Ready to install VSC 4.1*



11. Click **Finish**.

#### Register VSC with vCenter Server

To register the VSC with the vCenter Server, follow these steps:

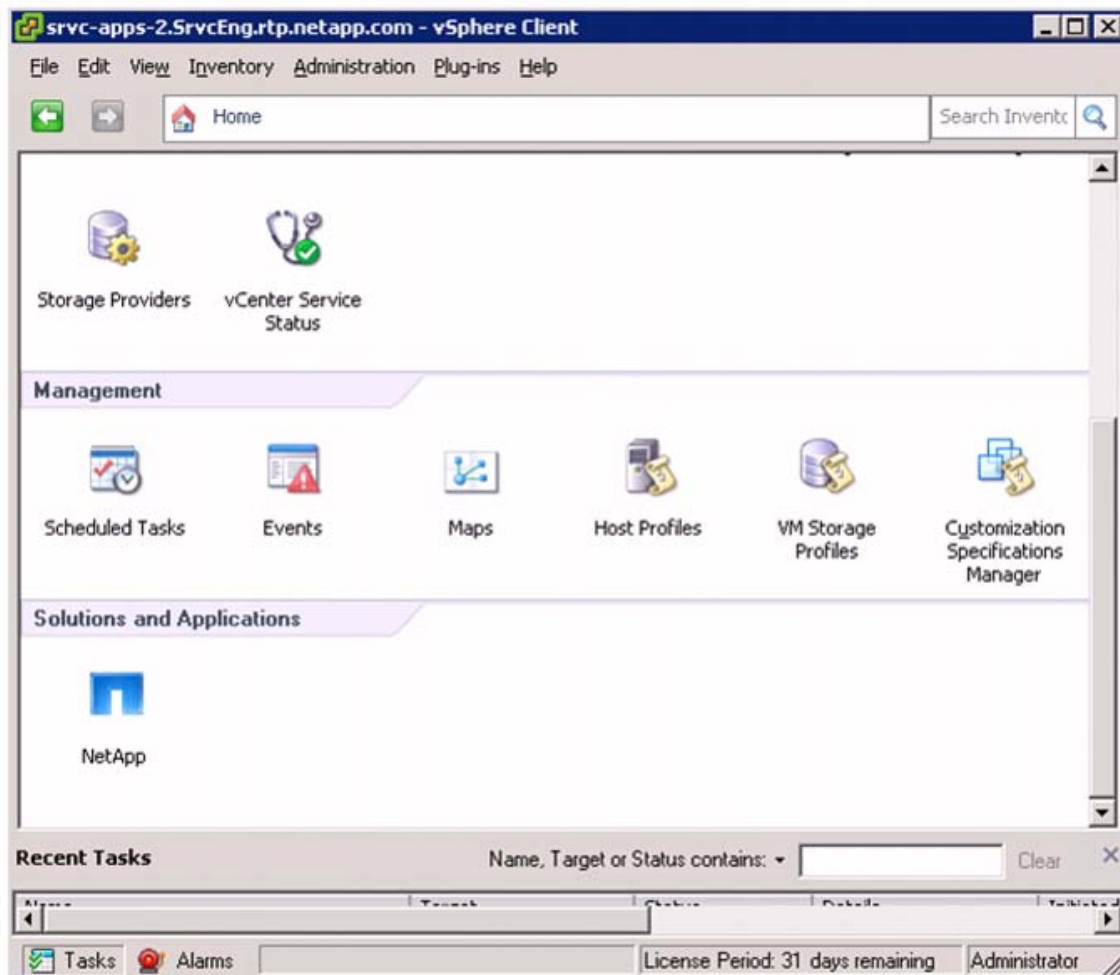
1. A browser window with the registration URL opens automatically when the installation phase is complete.
2. Click **Continue** to this website (not recommended).
3. In the Plug-in Service Information section, choose the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.
4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user), and user password for the vCenter Server.
5. Click **Register** to complete the registration.

#### Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, follow these steps:

1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user. If the vSphere Client was previously opened, close it and then reopen it.
2. Click the **Home** tab in the left side of the vSphere Client window.
3. Under Solutions and Applications, click **NetApp**.

Figure 178 Adding Storage Resource



4. Click Yes when the security certificate warning appears. To view the certificate, click View Certificate.
5. In the navigation pane, choose **Monitoring and Host Configuration** if it is not selected by default.
6. In the list of storage controllers, right-click the first controller listed and choose **Modify Credentials**.
7. Enter the storage cluster management IP address in the Management IP address field. Enter admin for the user name, and the admin password for the password. Ensure that Use SSL is selected.
8. Click **OK**.
9. Click **OK** to accept the controller privileges.
10. For additional information on storage provisioning, storage efficiency tools, reporting and backup/recovery refer to the [VSC Administration Guide](#)

## Creating Primary Storage on Fibre Channel LUN

1. In the VMware vSphere Client connected to vCenter, choose **Home > Hosts and Clusters**.
2. Expand the vCenter, DataCenter, and ESX Host <var\_vm\_host\_infra\_01\_ip label > 10.65.121.165.



3. Click **Storage** under Hardware
4. Click **Add Storage**.
5. Select **Disk/LUN** radio button.
6. Click **Next**.
7. Select NetApp LUN =2.
8. Click **Next**.
9. Select **VMFS-5** radio button.
10. Click **Next**.
11. Click **Next** to create partition.
12. Click **Next**.
13. Type TenantA-FC-PrimaryStorage.
14. Click **Next**.
15. Select **Maximum available space** radio button.
16. Click **Next**.
17. Click **Finish**.
18. Verify Datastore <<TenantA-FC-PrimaryStorage>> is created under Darastores.

**Figure 179** Datastore TenantA-FC-PrimaryStorage is mapped to NetApp Fiber Channel LUN

View: **Datstores** **Devices**

| Identification            | Status | Device             | Drive Type | Capacity  | Free      | Type  | Last Update          | Alarm Actions | Storage I/O Control |
|---------------------------|--------|--------------------|------------|-----------|-----------|-------|----------------------|---------------|---------------------|
| datastore1 (2)            | Normal | NETAPP Fibre Ch... | Non-SSD    | 10.00 GB  | 9.14 GB   | VMFS5 | 8/26/2013 4:43:46 AM | Enabled       | Disabled            |
| TenantA-FC-PrimaryStorage | Normal | NETAPP Fibre Ch... | Non-SSD    | 179.75 GB | 176.68 GB | VMFS5 | 8/26/2013 4:43:46 AM | Enabled       | Disabled            |

**Datastore Details**

**TenantA-FC-PrimaryStorage** 179.75 GB Capacity

Location: /vmfs/volumes/52038279-9397ac35-f9ee-0025b5000001

Hardware Acceleration: Unknown

3.07 GB Used

176.68 GB Free

[Refresh Storage Capabilities](#)

System Storage Capability: N/A

User-defined Storage Capability: N/A

| Path Selection      | Properties                                                     | Extents                                                                       | Storage I/O Control |
|---------------------|----------------------------------------------------------------|-------------------------------------------------------------------------------|---------------------|
| Round Robin (VM...) | Volume Label: TenantA-FC-...<br>Datastore Name: TenantA-FC-... | NETAPP Fibre Channel Disk ... 180.00 GB<br>Total Formatted Capacity 179.75 GB | Disabled            |

**Paths**

Total: 4  
Broken: 0  
Disabled: 0

**Formatting**


File System: VMFS 5.58  
Block Size: 1 MB

# Cloud Deployment Design

The Management Server manages one or more zones (typically, datacenters) containing host computers where guest virtual machines will run. The cloud infrastructure is organized as follows:

- Zone  
Typically, a zone is equivalent to a single datacenter. A zone consists of one or more pods and secondary storage.
- Pod  
A pod is usually one rack of hardware that includes a layer-2 switch and one or more clusters.
- Cluster  
A cluster consists of one or more hosts and primary storage.
- Host  
A single compute node within a cluster. The hosts are where the actual cloud services run in the form of guest virtual machines.
- Primary storage is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster.
- Secondary storage is associated with a zone, and it stores templates, ISO images, and disk volume snapshots.

**Figure 180**      *Logical View of CloudPlatform Cloud Provisioning*



This section outlines Citrix CloudPlatform Compute, Network and Storage infrastructure design required to create hosted private cloud for multi-tenants.

CloudPlatform is a multi-hypervisor, multi-tenant, high-availability Infrastructure as a Service cloud management platform. It provides a cloud infrastructure orchestration layer; giving automation of the creation, provisioning and configuration of IaaS components (such as virtual servers).

It provides role based access features for enterprises to set up utilization charge-backs, user self-administration, approved virtual machine template libraries, and security controls to govern how users utilize the cloud.

To manage and configure cloud infrastructure, Citrix provides CloudPlatform, a software platform that pools computing resources to build public, private, and hybrid cloud network. CloudPlatform manages the network, storage, and compute nodes that make up a cloud infrastructure. Use CloudPlatform to deploy, manage, and configure cloud computing environments.

In this guide, Citrix CloudPlatform 4.2.1 software version is installed on RHEL 6.0.3 Operating System on Cisco B-200 M3 blade server. This paper shows the installation procedure of CloudPlatform 4.2.1 in single user mode.


**Note**

It is recommended installing CloudPlatform in multi user mode to support high availability under any disaster

The below table shows Network Interfaces and their corresponding IP Address (VLAN ID) configured to access Management, NFS and PXE networks.

Make sure the CloudPlatform host is able to access internet, pingable and route has been configured correctly to access networks defined below in [Table 33](#).

**Table 33** *Ethenert Interfaces with IP Address (VLAN ID) details*

| Ethernet Interface Name | IP Address (VLAN ID)                                  | Network Access    |
|-------------------------|-------------------------------------------------------|-------------------|
| Eth0                    | <var_vm_host_infra_01_ip_mngnt><br>10.65.121.70 (603) | Management Access |
| Eth1                    | <var_host_infra_01_NFS_ip><br>193.191.1.10 (193)      | NFS Access        |
| Eth2                    | <var_host_infra-01_PXE_ip><br>20.1.1.5 (20)           | PXE Access        |

## Install CloudPlatform in single user mode

This section describes installing a single Management Server and installing MySQL on the same node. To prepare The RHEL 6.0.3 OS to host the Management Server follow these steps:

1. Log in to your OS as root.
2. Check for a fully qualified hostname: `# hostname --fqdn`.
3. This should return a fully qualified hostname such as "kvm1.lab.example.org". If it does not, edit `/etc/hosts` so that it does.
4. Set SELinux to be permissive by default.
5. In RHEL or CentOS, SELinux are installed and enabled by default. You can verify this with: `# rpm -qa | grep selinux`.

6. Set the SELINUX variable in /etc/selinux/config to “permissive”. This ensures that the permissive setting will be maintained after a system reboot.

```
vi /etc/selinux/config
```

7. Then set SELinux to permissive starting immediately, without requiring a system reboot.

```
setenforce 0
```

8. Make sure that the machine can reach the internet

```
ping www.google.com
```

9. Download the CloudPlatform Management Server onto the host from <https://www.citrix.com/English/ss/downloads/>.

10. Choose CloudPlatform 4.2.1 under the CloudStack 4.2 listing.




---

**Note** You will need a MyCitrix account.

---

11. Install the CloudPlatform packages. You should have a file in the form of “CloudStack-VERSION-N-OSVERSION.tar.gz”. Untar the file and then run the install.sh script inside it. Replace the file and directory names below with those you are using:

```
tar xzf CloudStack-VERSION-N-OSVERSION.tar.gz
cd CloudStack-VERSION-N-OSVERSION
./install.sh
```

12. You should see a few messages as the installer prepares, followed by a list of choices.

```
Choose M to install the Management Server software.
Install the Agent
Install BareMetal Agent
Install the Usage Monitor
Install the CloudPlatform packages installed on this computer
Stop any running CloudPlatform services and remove the CloudPlatform packages from
this computer
Remove the MySQL server (will not remove the MySQL databases)
Quit
```

13. Wait for a message such as “Complete! Done.”

14. When the installation is finished, run the following commands to start essential services (the commands might be different depending on your OS).

```
service rpcbind start
service nfs start
chkconfig nfs on
chkconfig rpcbind on
```

15. Continue with Install and Configure the Database, If you already have a version of MySQL installed on the Management Server node, make one of the following choices, depending on what version of MySQL it is.

- a. If you already have installed MySQL version 5.1.58 or later, skip to step 18
- b. If you have installed a version of MySQL earlier than 5.1.58, you can either skip to step 18 or uninstall MySQL or proceed to step 10 to install a more recent version.

16. On the same machine where you have installed the CloudPlatform Management Server, re-run install.sh

```
./install.sh
```

17. Choose D to install MySQL server from the distribution’s repository.

```
> D
```

18. Edit the MySQL configuration (/etc/my.cnf or /etc/mysql/my.cnf, depending on your OS) and insert the following lines in the [mysqld] section. You can put these lines below the datadir line. The max\_connections parameter should be set to 350 multiplied by the number of Management Servers you are deploying. This example assumes one Management Server.

```
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
max_connections=350
log-bin=mysql-bin
binlog-format = 'ROW'
```

19. Restart the MySQL service, then invoke MySQL as the root user.

```
service mysqld restart
mysql -u root
```

20. On RHEL and CentOS, MySQL does not set a root password by default. It is very strongly recommended that you set a root password as a security precaution. Run the following commands, and substitute your own desired root password.

```
mysql> SET PASSWORD = PASSWORD('password');
From now on, start MySQL with mysql -p so it will prompt you for the password.
```

21. To grant access privileges to remote users, perform the following steps.

- a. Run the following commands from the mysql prompt:

```
mysql> GRANT ALL PRIVILEGES ON *.* TO 'root'@'%' WITH GRANT OPTION;
mysql> exit
```

- b. Restart the MySQL service.

- c. # service mysqld restart

- d. Open the MySQL server port (3306) in the firewall to allow remote clients to connect.

```
iptables -I INPUT -p tcp --dport 3306 -j ACCEPT
```

22. Edit the /etc/sysconfig/iptables file and add the following line at the beginning of the INPUT chain.

```
-A INPUT -p tcp --dport 3306 -j ACCEPT
```

23. Set up the database. The following command creates the cloud user on the database for performing DBA activity.

- a. In dbpassword, specify the password to be assigned to the cloud user. You can choose to provide no password.

- b. In deploy-as, specify the username and password of the user deploying the database. In the following command, it is assumed the root user is deploying the database and creating the cloud user.

- c. (Optional) For encryption\_type, use file or web to indicate the technique used to pass in the database encryption password. Default: file..

- d. (Optional) For management\_server\_key, substitute the default key that is used to encrypt confidential parameters in the CloudPlatform properties file. Default: password. It is highly recommended that you replace this with a more secure value.

- e. (Optional) For database\_key, substitute the default key that is used to encrypt confidential parameters in the CloudPlatform database. Default: password. It is highly recommended that you replace this with a more secure value.

```
cloud-setup-databases cloud:<dbpassword>@localhost --deploy-as=root:<password>
-e <encryption_type> -m <management_server_key> -k <database_key>
```

24. If you are running the KVM hypervisor on the same machine with the Management Server, edit /etc/sudoers and add the following line:

```
Defaults:cloud !requiretty
```

25. Now that the database is set up, you can finish configuring the OS for the Management Server. This command will set up iptables, sudoers, and start the Management Server.

```
cloud-setup-management
```

## 26. Prepare the System VM Template.

- Secondary storage must be seeded with a template that is used for CloudPlatform system VMs.
- On the Management Server, run one or more of the following cloud-install-sys-tmplt commands to retrieve and decompress the system VM template. Run the command for each hypervisor type that you expect end users to run in this Zone.
- Mount NetApp NFS share on Management Server named /TenantADDataSecNFS mount point name.
- If you set the CloudPlatform database encryption type to "web" when you set up the database, you must use the parameter -s <management-server-secret-key>. This process will require approximately 5 GB of free space on the local file system and up to 30 minutes each time it runs.

For vSphere:

```
/usr/share/cloudstack-common/scripts/storage/secondary/cloud-install-sys-tmplt -m
/export/TenantADDataSecNFSVol -u
http://download.cloud.com/templates/4.2.1/systemvmtemplate-4.2.1-vh7.ova -h vmware -o
<<var_vm_host_infra_01_ip_mngnt> 10.65.121.70 -r admin -d xxxxxx -F
```

Prepare BareMetal Agent and configure PXE on RHEL Linux Operating System for CloudPlatform 4.2.1 to perform Baremetal as a Service in cloud



### Note

In this paper we have used same system which we have installed CloudPlatform 4.2.1.

- Login in as root to RHEL 6.X download and Install PXE, DHCP, TFTP and NFS packages
- The CloudPlatform-VERSION-N-OSVERSION.tar.gz." is already downloaded and untar as part of CloudPlatform 4.2.1 installation which you have prepared as shown in install CloudPlatform section
- To install BareMetal Agent run the install.sh script found in extracted CloudPlatform 4.2.1 directory

```
[root@PXE CloudPlatform-4.2.1.1-rhel6.3]# ./install.sh
Setting up the temporary repository...
Cleaning Yum cache...
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
Updating certificate-based repositories.
Unable to read consumer identity
Cleaning repos: cloud-temp
0 metadata files removed
Welcome to the CloudPlatform Installer. What would you like to do?
NOTE: For installing KVM agent, please setup EPEL<http://fedoraproject.org/wiki/EPEL>
yum repo first;
For installing CloudPlatform on RHEL6.x, please setup distribution yum repo either
from ISO or from your registration account.
3.We detect you already have MySQL server installed, you can bypass mysql install
chapter in CloudPlatform installation guide.
Or you can use E) to remove current mysql then re-run install.sh selecting D)
to reinstall if you think existing MySQL server has some trouble.
For MySQL downloaded from community, the script may not be able to detect it.
A) Install the Agent
B) Install BareMetal Agent
S) Install the Usage Monitor
U) Upgrade the CloudPlatform packages installed on this computer
R) Stop any running CloudPlatform services and remove the CloudPlatform packages
from this computer
E) Remove the MySQL server (will not remove the MySQL databases)
Q) Quit
Type B option to install BareMetal Agent
> B
```

```

Installing the BareMetal Agent...
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
Updating certificate-based repositories.
Unable to read consumer identity
cloud-temp
| 1.3 kB 00:00 ...
cloud-temp/primary
| 2.5 kB 00:00 ...
cloud-temp
7/7
Setting up Install Process
Package cloudstack-baremetal-agent-4.2.1.0-2.el6.x86_64 already installed and latest
version
Nothing to do
Done

```

- e. Run the BareMetal setup script after installation of BareMetal Agent. Make sure to note down TFTP root directory**

```

[root@PXE CloudPlatform-4.2.1.1-rhel6.4]# cloudstack-setup-baremetal
Checking is root
[OK]
Checking tftp-server
[OK]
Checking syslinux
[OK]
Checking xinetd
[OK]
Checking chkconfig
[OK]
Checking dhcp
[OK]
Executing 'chkconfig --level 345 tftp on'
[OK]
Executing 'chkconfig --level 345 xinetd on'
[OK]
Executing 'chkconfig --level 345 dhcpd on'
[OK]
Executing '/etc/init.d/xinetd restart'
[OK]
Detected iptables is running, need to open tftp port 69
Executing 'iptables -I INPUT 1 -p udp --dport 69 -j ACCEPT'
[OK]
Executing '/etc/init.d/iptables save'
[OK]
Executing 'cp -f /usr/share/syslinux/pxelinux.0 /export'
[OK]
Setup BareMetal PXE server successfully
TFTP root directory is: /export

```

- f. Configure NFS configuration file /etc/exports. Make sure /export (TFTP root) and /var/www/html/RHEL63 (http root) directories are NFS exported for BareMetal Host to access PXE Image**

```

[root@PXE ~]# cat /etc/exports
/export *(rw,async,no_root_squash)
/var/www/html/RHEL63 *(rw,async,no_root_squash)

```

- g. Verify that NFS is running on the NFS server**

```

[root@PXE ~]# vi /export/RHEL63.ks
[root@PXE ~]# rpcinfo -p
 program vers proto port service
 100000 4 tcp 111 portmapper
 100024 1 udp 662 status
 100011 1 udp 875 rquotad
 100005 1 udp 892 mountd

```

```
100003 2 tcp 2049 nfs
```

- h. Download PXE files pxelinux.0, pxelinux.cfg and RHEL63.ks (kickstart image) from RHEL 6.0.3 Operating System ISO image to shared director /export (TFTP root directory)

```
[root@PXE export]# ll
total 112
drwxr-xr-x. 2 root root 4096 Oct 11 04:52 pxeimages
-rw-r--r--. 1 root root 26828 Oct 22 04:16 pxelinux.0
-rw-r-xr-x. 2 root root 4096 Oct 22 04:55 pxelinux.cfg
-rwxr-xr-x. 1 root root 1478 Oct 22 04:56 RHEL63.ks
```

- i. Verify that DHCPD service is started and running on the server. The sample DHCPD.conf file can be created to get DHCPD service to run.

```
[root@PXE export]# service dhcpd status
dhcpd (pid 7306) is running...
```

#### Sample file

```
DHCP Server Configuration file.
see /usr/share/doc/dhcp*/dhcpd.conf.sample
see 'man 5 dhcpd.conf'
ignore client-updates;
authoritative;
allow booting;
allow bootp;
subnet 10.65.121.0 netmask 255.255.255.0
{
 option routers 10.65.121.1;
 option subnet-mask 255.255.0.0;
 option broadcast-address 10.65.121.255;
 option time-offset -18000;
 default-lease-time 21600;
 max-lease-time 43200;
}
```

- j. Copy RHEL 6.3 OS Images by extracting ISO to shared directory /export/RHEL6-3.

```
[root@PXE RHEL6-3]# ls
EFI RELEASE-NOTES-es-ES.html RELEASE-NOTES-si-LK.html
EULA RELEASE-NOTES-fr-FR.html RELEASE-NOTES-ta-IN.html
GPL RELEASE-NOTES-gu-IN.html RELEASE-NOTES-te-IN.html
HighAvailability RELEASE-NOTES-hi-IN.html RELEASE-NOTES-zh-CN.html
images RELEASE-NOTES-it-IT.html RELEASE-NOTES-zh-TW.html
isolinux RELEASE-NOTES-ja-JP.html repodata
LoadBalancer RELEASE-NOTES-kn-IN.html ResilientStorage
media.repo RELEASE-NOTES-ko-KR.html RPM-GPG-KEY-redhat-beta
Packages RELEASE-NOTES-ml-IN.html RPM-GPG-KEY-redhat-release
README RELEASE-NOTES-mr-IN.html ScalableFileSystem
RELEASE-NOTES-as-IN.html RELEASE-NOTES-or-IN.html Server
RELEASE-NOTES-bn-IN.html RELEASE-NOTES-pa-IN.html TRANS.TBL
RELEASE-NOTES-de-DE.html RELEASE-NOTES-pt-BR.html
RELEASE-NOTES-en-US.html RELEASE-NOTES-ru-RU.html
```

- k. Start and configure http and tftp server to provide access for CloudPlatform BareMetal agent to access RHEL 6.3 kickstart, kernel and initrd files for performing PXE boot.

```
[root@PXE RHEL63]# service httpd start
Starting httpd:
[root@PXE export]# service xinetd status
xinetd (pid 3748) is running...
```

4. Create softlinks paths under /var/www/html/RHEL63 which have links created for RHEL 6.3 kickstart image, initrd.img and vmlinuz kernel files for BareMetal Host PXE support.

```
[root@PXE RHEL63]# ln -s /export/RHEL6-3/images/pxeboot/initrd.img initrd.img
[root@PXE RHEL63]# ln -s /export/RHEL63.ks RHEL63.ks
[root@PXE RHEL63]# ln -s /export/RHEL63.ks RHEL6-3/images/pxeboot/vmlinuz vmlinuz
[root@PXE RHEL63]# ll
total 1780
```



```
lrwxrwxrwx. 1 root root 41 Oct 21 05:45 initrd.img ->
/export/RHEL6-3/images/pxeboot/initrd.img
lrwxrwxrwx. 1 root root 17 Oct 21 07:11 RHEL63.ks -> /export/RHEL63.ks
lrwxrwxrwx. 1 root root 38 Oct 17 02:52 vmlinuz ->
/export/RHEL6-3/images/pxeboot/vmlinuz
```

- I. Configure RHEL 6.3 kickstart image file NGS installation media option with NFS IP Address and RHEL 6.3 OS Image.



#### Note

In this paper we have created **Eth2 (20.1.1.5) VLAN 20** PXE Network for access PXE boot and have NFS export path **/export** which stores RHEL 6.3 Operating System Installtions files, kickstart, initrd and kernel files.

#### Sample RHEL kickstart Image file

```
[root@PXE export]# cat RHEL63.ks
#platform=x86, AMD64, or Intel EM64T
#version=DEVEL
Firewall configuration
firewall --disabled
Install OS instead of upgrade
install
Use NFS installation media
nfs --server=20.1.1.5 --dir=/export/RHEL6-3
Root password
rootpw --iscrypted 1rWDxD1lu$iqJek3Co0ivKUEhl68CmA/
System authorization information
auth --useshadow --passalgo=nbv12345
Use graphical install
graphical
Run the Setup Agent on first boot
firstboot --enable
System keyboard
keyboard us
System language
lang en_US
Driver Disk
#driverdisk
--source=nfs:20.1.1.50:/var/www/html/RHEL63/dd-fnic-rhel6.3-1.6.0.5.iso
#driverdisk
--source=nfs:192.85.0.2:/home/Delmar-MR1-Driver/RHEL6.3/enic-2.1.1.41-rhel6u3-dd.i
so
SELinux configuration
selinux --disabled
Installation logging level
logging --level=info
Reboot after installation
#reboot
System timezone
timezone Asia/Kolkata
Network information
network --bootproto=dhcp --device=eth0 --onboot=yes --activate
#network --bootproto=dhcp --device=eth2 --onboot=yes --activate
System bootloader configuration
#bootloader --location=mbr
Clear the Master Boot Record
#zerombr
Partition clearing information
#clearpart --none
Disk partitioning information
#part /boot --fstype="ext4" --size=100
#part swap --fstype="swap" --size=2048
```

```
#part / --fstype="ext4" --grow --size=1
%packages
@base
@network-file-system-client
@network-tools
@server-platform
%end
install
#xconfig --defaultdesktop=GNOME --depth=8 --resolution=800x600
```

## CloudPlatform Advanced Design

The Citrix Cloud Design has been split into the following tasks:

1. Creating Zones
2. Defining Network
3. Adding Pods
4. Adding Cluster
5. Adding Hosts
6. Defining Storage

## Creating Zones

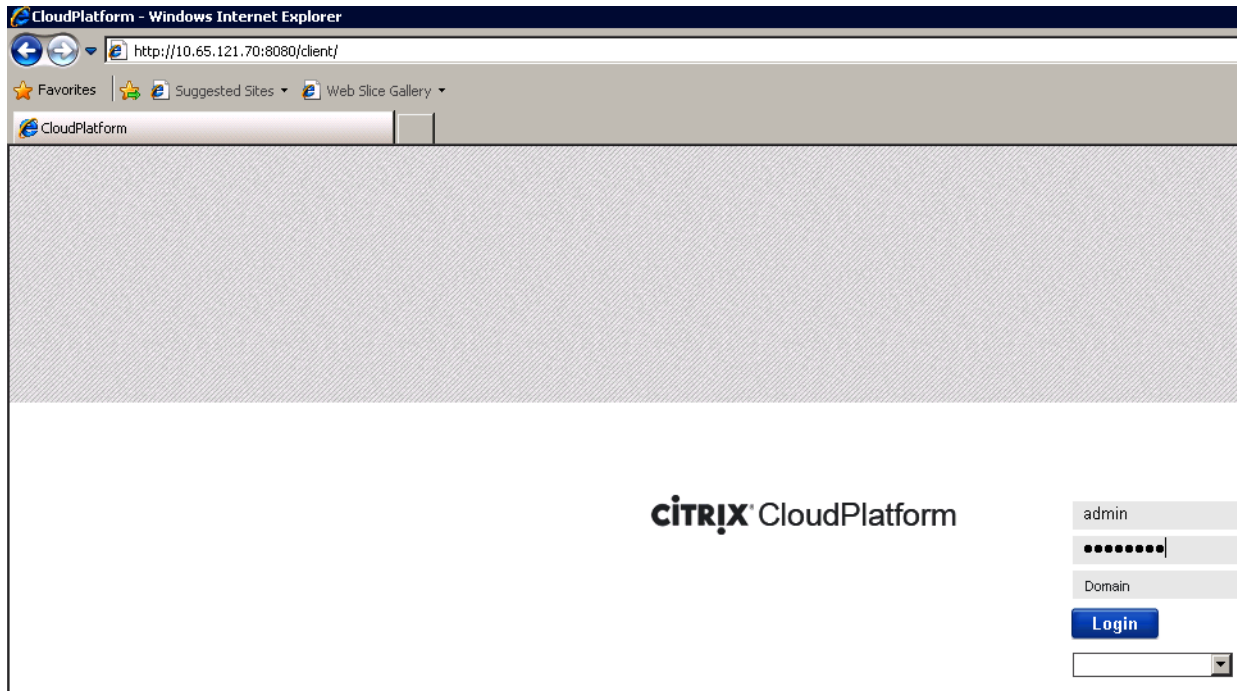
A zone is the largest organizational unit within a CloudPlatform deployment. A zone typically corresponds to a single datacenter, although it is permissible to have multiple zones in a datacenter. The benefit of organizing cloud compute, network and storage infrastructure into zones is to provide physical isolation and redundancy.

In order to provide true multi tenant cloud services Tenants cloud infrastructure Compute, Storage and Network for hosting virtual machines resources like CPU, Memory, Network and Storage can be granularly grouped into different isolated logical partitions by creating Zones, Pod, Clusters and Hosts.

To create zones using the Citrix CloudPlatform application to host cloud services, login to CloudPlatform with user credentials and follow these steps:

1. Access CloudPlatform 4.2.1 Management server IP Address on web browser with URL <http://10.65.121.70:8080/client/>.
2. Type the User Name <root>, Password <XXXXXX>, and Domain.
3. Click **Login**.

**Figure 181**      *Logging to the Citrix CloudPlatform*

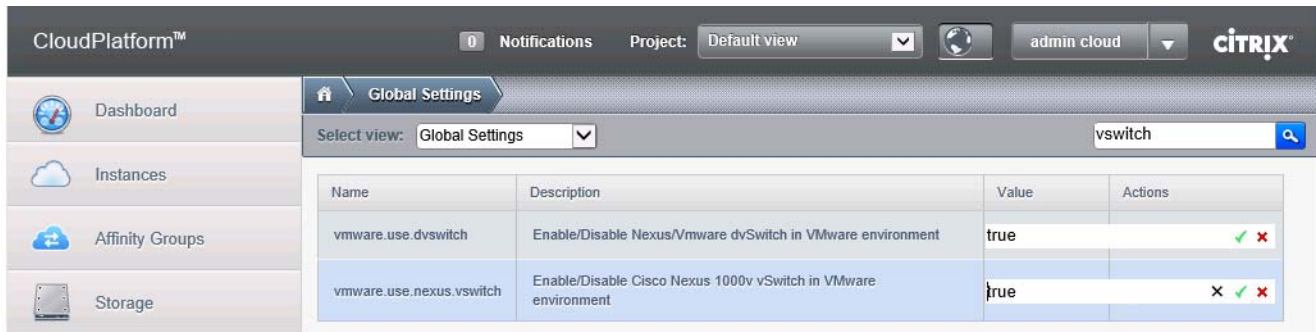


4. Click **Agree** to confirm the End User License Agreement.
5. Click **I have used CloudPlatform before, skip this guide** button.
6. Click **Infrastructure** in the left pane. This shows the currently unconfigured CloudPlatform environment.

**Figure 182**      *Displaying the Infrastructure Page*



7. To handle TenantA Virtual Machines Network access enable vmware use nexus vswitch DVS on CloudPlatform 4.2.1. Click **Global Settings** in the left pane. In the right pane, with Global Settings in the Select View box selected, type vSwitch in the Search box.
8. Press Enter.
9. Change Value to True on VMWare.use.dvswitch.
10. Change Value to True on VMWare.use.nexus.vswitch.
11. Press Enter.

**Figure 183**      **Changing Value to True on Global Setting**

12. In search box Type vmware.
13. Press Enter.
14. Change Value to TenantA-Mgmt-Uplink on vmware.managment.portgroup.
15. Press Enter.
16. Restart cloudstack Management service on console by ssh to 10.65.121.70.

```
[root@cloudstack CloudPlatform-UCS-for-cisco-4.2.1.1-5-rhel6.3]# service
cloudstack-management status
Stopping cloudstack-management: [OK]
Starting cloudstack-management: [OK]
```

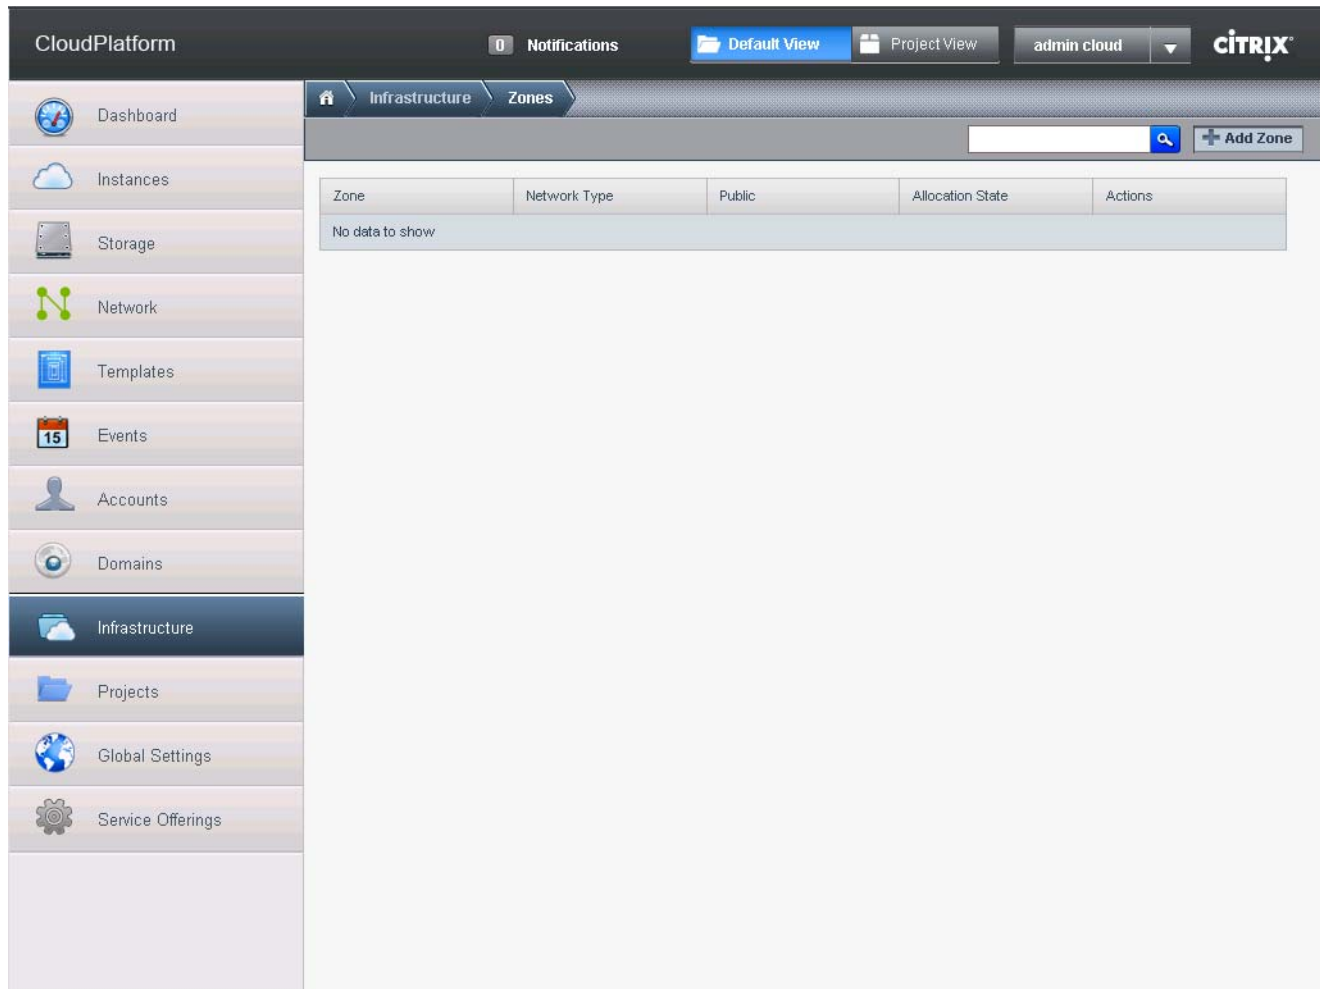
**Figure 184** Changing Value to TenantA-Mgmt-Uplink on Global Setting

The screenshot shows the Citrix CloudPlatform™ interface. The left sidebar contains navigation links: Dashboard, Instances, Affinity Groups, Storage, Network, Templates, Events, and Projects. The main area is titled 'Global Settings' with a search bar set to 'vmware'. A table lists various settings:

| Name                                  | Description                                                                                      | Value                       | Actions            |
|---------------------------------------|--------------------------------------------------------------------------------------------------|-----------------------------|--------------------|
| router.template.vmware                | Name of the default router template on Vmware.                                                   | SystemVM Template (vSphere) | [Edit]             |
| vmware.additional.vnc.portrange.size  | Start port number of additional VNC port range                                                   | 1000                        | [Edit]             |
| vmware.additional.vnc.portrange.start | Start port number of additional VNC port range                                                   | 50000                       | [Edit]             |
| vmware.create.full.clone              | If set to true, creates VMs as full clones on ESX hypervisor                                     | false                       | [Edit]             |
| vmware.management.portgroup           | Specify the management network name(for ESXi hosts)                                              | TenantA-Mgmt-Uplink         | [Edit] [X] [✓] [✗] |
| vmware.nested.virtualization          | When set to true this will enable nested virtualization when this is supported by the hypervisor | false                       | [Edit]             |

Below the table, two network diagrams are shown. The first diagram, titled 'Standard Switch: TenantA-Mgmt-Uplink', shows a 'Virtual Machine Port Group' with 'cloud.private.untagged.0.1-TenantA-Mgmt-Uplink' and 'TenantA-Mgmt-Uplink' (vmk1: 10.65.121.165) connected to 'Physical Adapters' vmnic0 and vmnic3. The second diagram, titled 'Standard Switch: TenantA-NFS-Uplink', shows a 'Virtual Machine Port Group' with 'TenantA-NFS-Uplink' and 'NFS-VMkernel' (vmk0: 193.191.1.10) connected to 'Physical Adapters' vmnic5 and vmnic2.

17. Login to CloudPlatform 4.2.1 Management server by providing User Name and Password
18. Click **Infrastructure** in the left pane, and then **View all** under the **Zones** tile in the right pane.
19. Create Zone for Tenants to provide multitenancy by clicking **+ Add Zone** at the top right corner of the right pane.

**Figure 185**     *Displaying all the Zones in the Infrastructure*

20. Click **Advanced zone type** radio button.
21. Click **Next**.

**Figure 186**      **Defining the Configuration for the Zone**

Add zone

1 Zone Type   2 Setup Zone   3 Setup Network   4 Add Resources   5 Launch

**Set up zone type**  
Please select a configuration for your zone.

☒ **Basic**  
Provide a single network where each VM instance is assigned an IP directly from the network. Guest isolation can be provided through layer-3 means such as security groups (IP address source filtering).

☐ **Advanced**  
For more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

Cancel   Next

22. Enter TenantA-Zone in the Name field.
23. Enter 171.70.168.183 to access public internet in DNS1 field.
24. Enter 10.65.121.70 running on CloudPlatform Management server in Internal DNS 1 field.
25. In Hypervisor select VMware in list box.
26. Enter 10.1.1.0/24 guest CIDR field which will be assigned to Virtual Machines communications inside Pod by DHCP server running in Router system VM in CloudPlatform 4.2.1
27. Uncheck **Public** check box.
28. Uncheck **Local Storage Enabled** check box.
29. Click **Next**.

**Figure 187**     *Defining the Zone Attributes*

**+ Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

A zone is the largest organizational unit in CloudPlatform™, and it typically corresponds to a single datacenter. Zones provide physical isolation and redundancy. A zone consists of one or more pods (each of which contains hosts and primary storage servers) and a secondary storage server which is shared by all pods in the zone.

\* Name: TenantA-Zone

\* IPv4 DNS1: 171.70.168.183

IPv4 DNS2:

\* Internal DNS 1: 10.65.121.70

Internal DNS 2:

\* Hypervisor: VMware

Network Offering: DefaultSharedNetworkOffering

Network Domain:

Dedicate: ☒

Domain: ROOT

Account:

Local storage enabled: ☐

Previous Cancel Next



## Defining Network

After defining zone the next step is to design and define cloud network. There are two types of network definitions-- Basic and Advanced. In the basic networking only one physical network can exist in a zone, however, the advanced networking allows multiple physical networks in a zone. Each physical network can carry one or more traffic types, and CloudPlatform lets you define which type of network traffic will be carried by each network. This solution design implements the advanced networking.

There are three types of traffic which are defined and each is carried using separate Virtual Switches with corresponding Ethernet Interfaces and Port Group names defined on ESXi Host:

- **Management Traffic**

It is generated within the CloudPlatform as the internal resources communicate with each other. This includes communication between hosts, system VMs (VMs used by CloudPlatform to perform various tasks in the cloud), and any other component that communicates directly with the CloudPlatform Management Server.

To access Management Traffic a separate ESX Virtual Switch <<TenantA-Mgmt-Uplink>> with VMKernel Port Group <<TenantA-Mgmt-Uplink>> is created with corresponding Network Adapters of appropriate network access and QoS definition are created on ESXi 5.1 Host which is managed by vCenter Host.

- **guest Private Traffic**

It is used and generated by end users running VMs. The guest network results when VMs communicate with each other over a network. guest Public traffic is generated when VMs in the cloud access the Internet. Publicly accessible IPs must be allocated for this purpose. End users can use the CloudPlatform UI to acquire the IPs to implement NAT between their guest network and the public network.

To access guest Public and Private Traffic Cisco Nexus 1000V DVS Switch <<VSM-TenantA>> with VMKernel Port Groups <<TenantA-guest-Uplink>> Uplink is created with corresponding Network Adapters of appropriate network access and QoS definition

- **Storage Traffic**

It is used for VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudPlatform uses a separate Network Interface Controller (NIC) named storage NIC for storage network traffic. Use of storage NIC that always operates on a high bandwidth network allows rapid template and snapshot copying.

To access Storage NFS Traffic a separate ESX Virtual Switch <<TenantA-NFS-Uplink>> with VMKernel Port Group <<TenantA-NFS-Uplink>> is created with corresponding Network Adapters of appropriate network access and QoS definition

**Table 34**      **Network Types and Information**

| Network Type      | ESX vSwitch / Cisco Nexus 1000V DVS | Tag                  |
|-------------------|-------------------------------------|----------------------|
| Management        | TenantA-Mgmt-Uplink                 | TenantA-Mgmt-Uplink  |
| guest             | TenantA-guest-Uplink                | TenantA-guest-Uplink |
| Secondary Storage | TenantA-NFS-Uplink                  | TenantA-NFS-Uplink   |

To configure the zone advanced network using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Drag Public and Guest traffic Types icon to Physical Network 2.

**Figure 188** Defining the Advanced Networking Configuration for the Zone



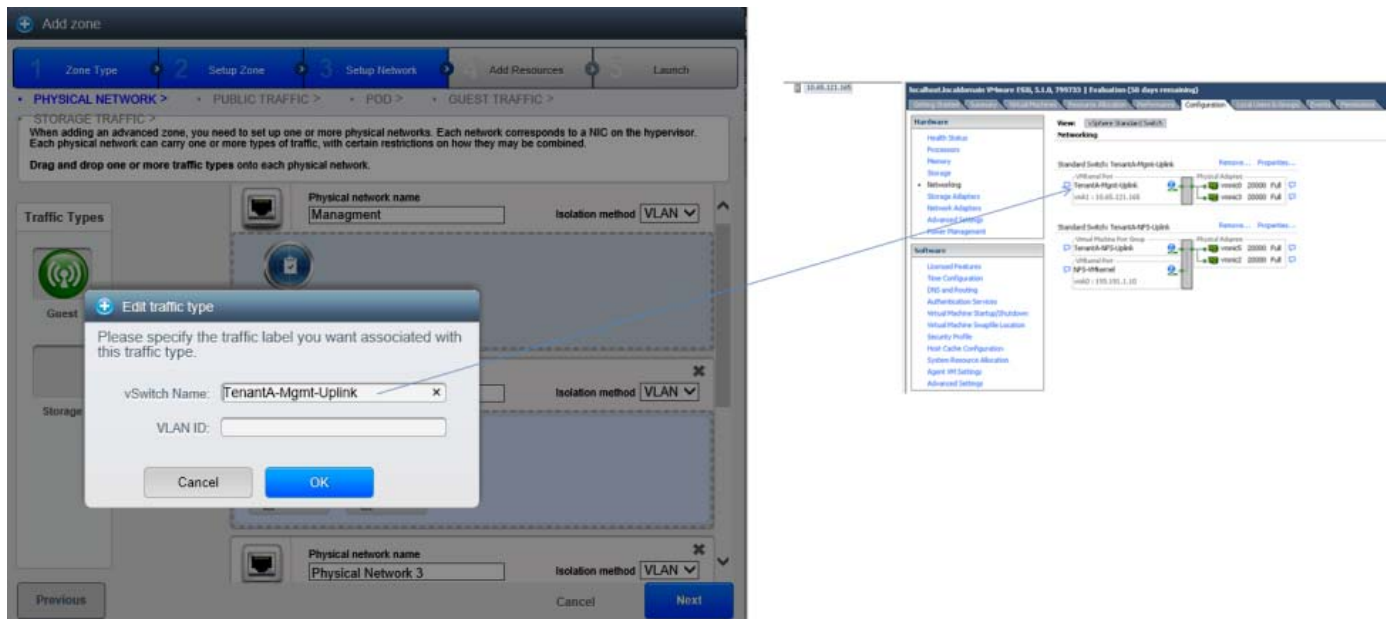
2. Drag Storage traffic Type icon to Physical Network 3.

**Figure 189** Adding the Traffic Type to Physical Network



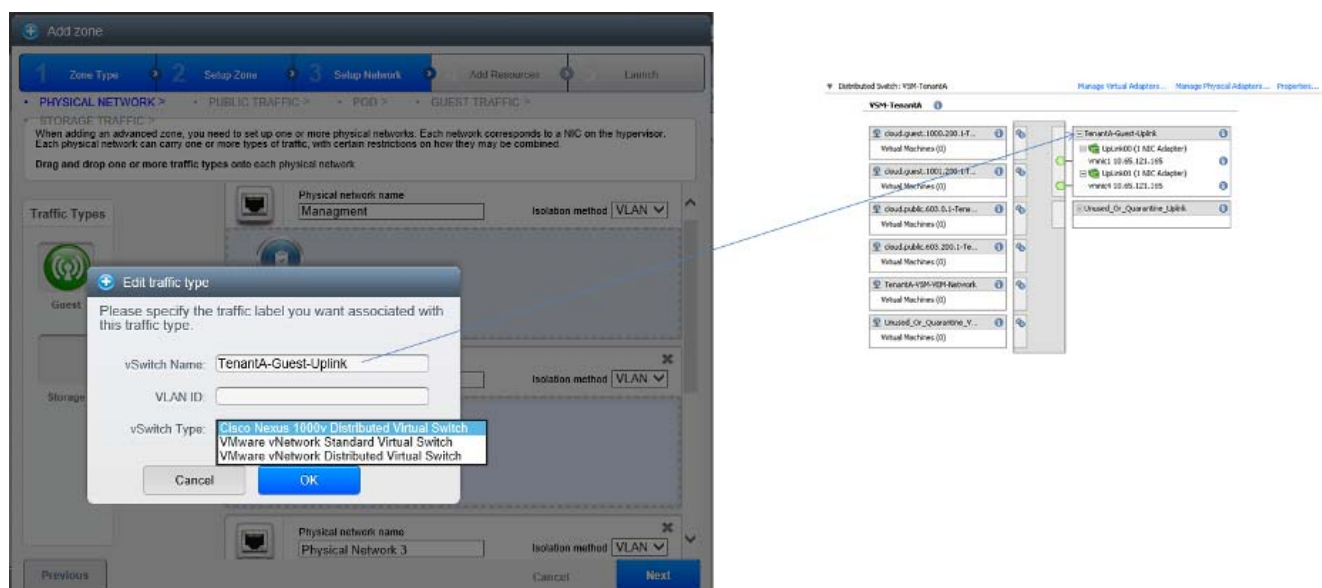
3. Rename Physical network Network 1 to <Management>.
4. Click **Edit** enter <TenantA-Mgmt-Uplink> in CloudPlatform Traffic label field. In VLAN ID leave blank for Untagged VLANs.
5. Select **VLAN** under Isolation method list box.
6. Click **OK**.

Figure 190 Labeling the Traffic Type on the Management Network



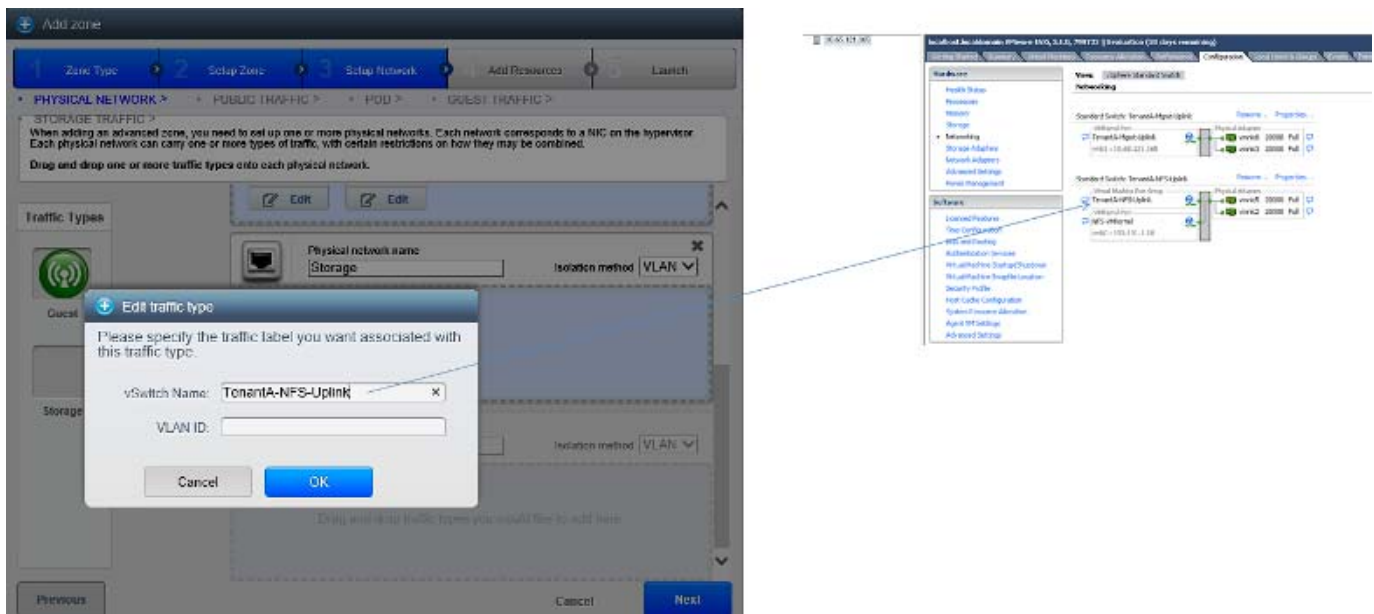
7. Rename Physical network as Network 2 to <guest>.
8. Click **Public Edit**, enter <TenantA-guest-Uplink> in CloudPlatform Traffic label field. Leave the VLAN ID field blank for Untagged VLANs.
9. Click **Guest Edit**, enter <TenantA-guest-Uplink> in CloudPlatform Traffic label field. Leave the VLAN ID field blank for Untagged VLANs.
10. Select VLAN under Isolation method list box.
11. Click OK.

Figure 191 Labeling the Traffic Type for guest Network



12. Rename Physical network Network 3 to <Storage>.
13. Click **Public Edit** enter <TenantA-NFS-Uplink> in CloudPlatform Traffic label field. Leave the VLAN ID field blank for Untagged VLANs.
14. Select VLAN under Isolation method list box.
15. Click **OK**.

**Figure 192** Labeling the Traffic Type for the Storage



16. Under VM Public TRAFFIC enter the below IP Network details.
17. Enter <10.65.122.1 in Gateway field.
18. Enter 255.255.255.0 in Network field.
19. Leave blank in VLAN (UCS tags VLAN 603).
20. Enter 10.65.122.161 in Start IP.
21. Enter 10.65.122.170 in End IP.
22. Click **ADD** button.
23. Click **Next**.

**Figure 193**      **Defining VM Public Traffic Network Details**

**Add zone**

1 Zone Type   2 Setup Zone   **3 Setup Network**   4 Add Resources   5 Launch

• **PUBLIC TRAFFIC >**   • POD >   • GUEST TRAFFIC >   • STORAGE TRAFFIC >

Public traffic is generated when VMs in the cloud access the internet. Publicly-accessible IPs must be allocated for this purpose. End users can use the CloudPlatform™ UI to acquire these IPs to implement NAT between their guest network and their public network.

Provide at least one range of IP addresses for internet traffic.

| Gateway              | Netmask              | VLAN                 | Start IP             | End IP               | Add                                | Actions                          |
|----------------------|----------------------|----------------------|----------------------|----------------------|------------------------------------|----------------------------------|
| <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="button" value="Add"/> |                                  |
| 10.65.122.1          | 255.255.255.0        |                      | 10.65.122.181        | 10.65.122.170        |                                    | <input type="button" value="X"/> |

## Adding Pods

Pod is the second-largest organizational unit within a CloudPlatform deployment. Pods are contained within zones. Each zone can contain one or more pods. A pod consists of one or more clusters of hosts and one or more primary storage servers

To add pod network in a zone advanced network configuration using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Enter TenantA-Pod in Pod Name field.
2. Enter <Pod\_VLAN\_603\_Gateway>10.65.121.1 in Reserved System Gateway field.
3. Enter 255.255.255.0 in Reserved System netmask field.
4. Enter 10.65.121.205 in Start Reserved system IP field.
5. Enter 10.65.121.215 in End Reserved system IP field.

6. Click **Next**.

**Figure 194**      **Adding the Pod to the Zone**

**Add zone**

1 Zone Type > 2 **Setup Zone** > 3 Setup Network > 4 Add Resources > 5 Launch

• PUBLIC TRAFFIC > • **POD** > • GUEST TRAFFIC > • STORAGE TRAFFIC >

Each zone must contain in one or more pods, and we will add the first pod now. A pod contains hosts and primary storage servers, which you will add in a later step. First, configure a range of reserved IP addresses for CloudStack's internal management traffic. The reserved IP range must be unique for each zone in the cloud.

\* Pod name:

\* Reserved system gateway:

\* Reserved system netmask:

\* Start Reserved system IP:

End Reserved system IP:

**Previous** **Cancel** **Next**

7. Enter 1000–1001 in guest VLAN Range field.
8. Click **Next**.



**Figure 195**      **Setting the Guest VLAN Range**

**Add zone**

1 Zone Type > 2 Setup Zone > **3 Setup Network** > 4 Add Resources > 5 Launch

• PUBLIC TRAFFIC > • POD > • **GUEST TRAFFIC >** • STORAGE TRAFFIC >

Guest network traffic is communication between end-user virtual machines. Specify a range of VLAN IDs to carry guest traffic for each physical network.

**Guest**

VLAN Range:   x

Previous Cancel Next

### Adding Storage Traffic

To add storage network in a zone advanced network configuration using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Enter <Storage\_NFS\_VLAN\_Gateway>193.191.1.1 in Gateway field.
2. Enter 255.255.255.0 in netmask field.
3. Enter 193 in the VLAN field.
4. Enter 193.191.1.50 in Start IP field.
5. Enter 193.191.1.80 in End IP field.
6. Click **Add**.
7. Click **Next**.



**Figure 196** Defining the Storage Traffic Configurations

The screenshot shows the 'Add zone' wizard with five steps: 1. Zone Type, 2. Setup Zone, 3. Setup Network, 4. Add Resources, and 5. Launch. Step 3, 'Setup Network', is currently active. Below the steps, there are navigation links: PUBLIC TRAFFIC >, POD <>, GUEST TRAFFIC >, and STORAGE TRAFFIC > (which is highlighted). A text box explains: 'Traffic between CloudPlatform™'s internal resources, including any components that communicate with the Management Server, such as hosts and CloudPlatform™ system VMs. Please configure storage traffic here.'

| Gateway     | Netmask       | VLAN | Start IP     | End IP       | Add                  | Actions |
|-------------|---------------|------|--------------|--------------|----------------------|---------|
| 193.191.1.1 | 255.255.255.0 | 193  | 193.191.1.50 | 193.191.1.80 | <button>Add</button> |         |
| 193.191.1.1 | 255.255.255.0 | 193  | 193.191.1.50 | 193.191.1.80 |                      |         |

At the bottom, there are three buttons: 'Previous' (disabled), 'Cancel', and 'Next' (active).

## Adding Cluster

Cluster provides a way to group hosts. Here to be precise, we have a cluster with a set of VMware clusters, preconfigured in vCenter. By definition all hosts part of cluster should have identical hardware, run the same hypervisor, are on the same subnet, and access the same shared primary storage. Virtual machine instances (VMs) can be live-migrated from one host to another within the same cluster, without interrupting service to the user.

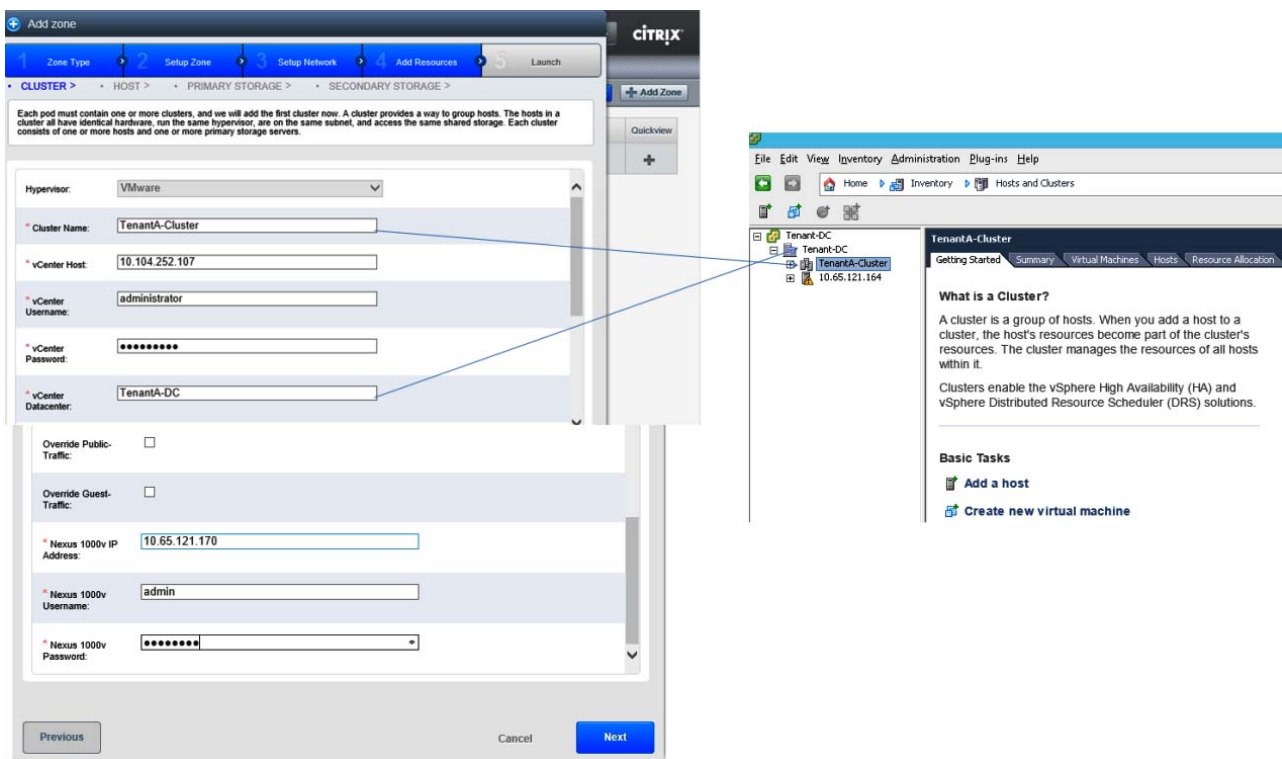
A cluster is the third-largest organizational unit within a CloudPlatform deployment. Clusters are contained within pods, and pods are contained within zones. Size of the cluster is limited by the underlying hypervisor, although the CloudPlatform recommends less in most cases. A cluster consists of one or more hosts and one or more primary storage servers.

To add a cluster in a zone using Citrix CloudPlatform for TenantA, follow these steps:

1. Enter <TenantA-Cluster> in Cluster Name field.

2. Enter <vCenter\_IP\_Address>10.104.252.107 in vCenter Host.
3. Enter <administrator> in vCenter Username.
4. Enter <xxxxxx> in vCenter Password.
5. Enter <Tenant-DC> in vCenter Datacenter Name. Make sure naming convention is followed in vCenter Datacenter name.
6. Uncheck **Override Public** and **Guest Traffic** check boxes.
7. Enter <VSM\_IP\_Address>10.65.121.170> Nexus 1000V IP Address.
8. Enter <admin> in Nexus 1000V Username.
9. Enter <xxxxx> in Nexus 1000V Password.
10. Click **Next**.

**Figure 197** Adding the Cluster to the Zone



## Adding Hosts

A host provides the computing resources that run the guest virtual machines. Each host has hypervisor software installed on it to manage the guest VMs. The Host here is an ESXi server. The host is the smallest organizational unit within a CloudPlatform deployment. Hosts are contained within clusters, clusters are contained within pods, and pods are contained within zones.

During initial zone creation CloudPlatform will automatically sync with VMware vCenter to automatically add hosts which are part of vCenter Cluster.

Later to add or modify hosts in zones CloudPlatform allows you to add / remove manually hosts which is explained in Cloud Infrastructure Design section.

## Defining Storage

### Primary Storage

CloudPlatform allows to define Primary Storage on Zone wide bases. It is associated with a cluster, and it stores the Operating System and Data disk volumes for all the VMs running on hosts in that cluster. The Primary Storage can File or Block based Protocol to storage VM data.

To add Primary Storage for zone using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Enter <TenantA-PrimaryNFS-Storage> in the Name field.
2. Enter Zone-Wide in Scope field.
3. Select nfs under Protocol list box.
4. Enter <NetApp\_NFS\_LIF\_IP\_Address>193.191.1.20 in Server field.
5. Enter </TenantA\_PrimaryStorage> in Path field.
6. Click **Next**.

**Figure 198**     *Defining the Primary Storage Properties*

The screenshot shows a multi-step wizard titled "Add zone". The steps are: 1. Zone Type, 2. Setup Zone, 3. Setup Network, 4. Add Resources, and 5. Launch. Step 4 is currently active. Below the steps, there are breadcrumb links: CLUSTER > HOST > PRIMARY STORAGE > SECONDARY STORAGE >. A message box states: "Each cluster must contain one or more primary storage servers, and we will add the first one now. Primary storage contains the disk volumes for all the VMs running on hosts in the cluster. Use any standards-compliant protocol that is supported by the underlying hypervisor." The form fields are as follows:

- \* Name: TenantA-PrimaryNFS-Storage
- Scope: Zone-Wide (dropdown)
- \* Protocol: nfs (dropdown)
- \* Server: 193.191.1.20
- \* Path: /TenantA\_PrimaryStorage
- Storage Tags: (empty text box)

At the bottom, there are buttons for "Previous", "Cancel", and "Next".

## Secondary Storage

It is associated with a zone, and it stores VM templates used for OS images that can be used to boot VMs and can include additional configuration information, such as installed applications. It provides ISO repository for storing OS images or disc images containing data or bootable media for operating systems. The storage VM disk volume snapshots stores the saved copies of VM data which can be used for data recovery or to create new templates

The items in secondary storage are available to all hosts in the zone.

To add Secondary Storage for zone using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Select <NFS> in Provider field
2. Enter </TenantA\_SecondaryNFS-Storage> in Name feild.
3. Enter <NetApp\_NFS\_LIF\_IP\_Address>193.191.1.20 in NFS Server feild.

4. Enter </TenantA\_SecondaryStorage> in Path field.
5. Click **Next**.

**Figure 199** Defining the Name and Path for the Secondary Storage

**Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• CLUSTER > • HOST > • PRIMARY STORAGE > • SECONDARY STORAGE >

Each zone must have at least one NFS or secondary storage server, and we will add the first one now. Secondary storage stores VM templates, ISO images, and VM disk volume snapshots. This server must be available to all hosts in the zone.  
Provide the IP address and exported path.

Provider: NFS

Name: TenantA-SecondaryNFS-Storage

\* NFS Server: 193.191.1.20

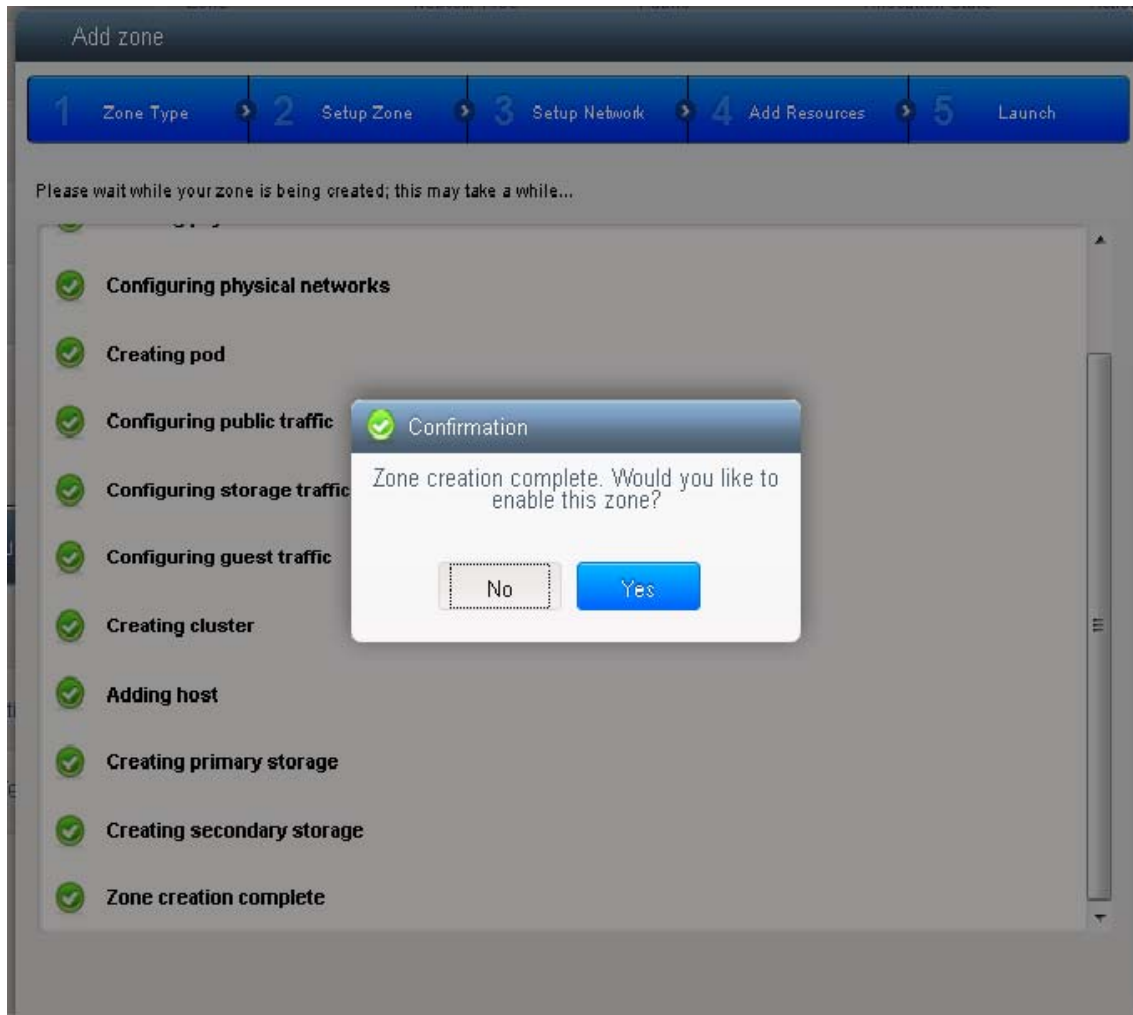
\* Path: /TenantA\_SecondaryStorage

Previous Cancel Next



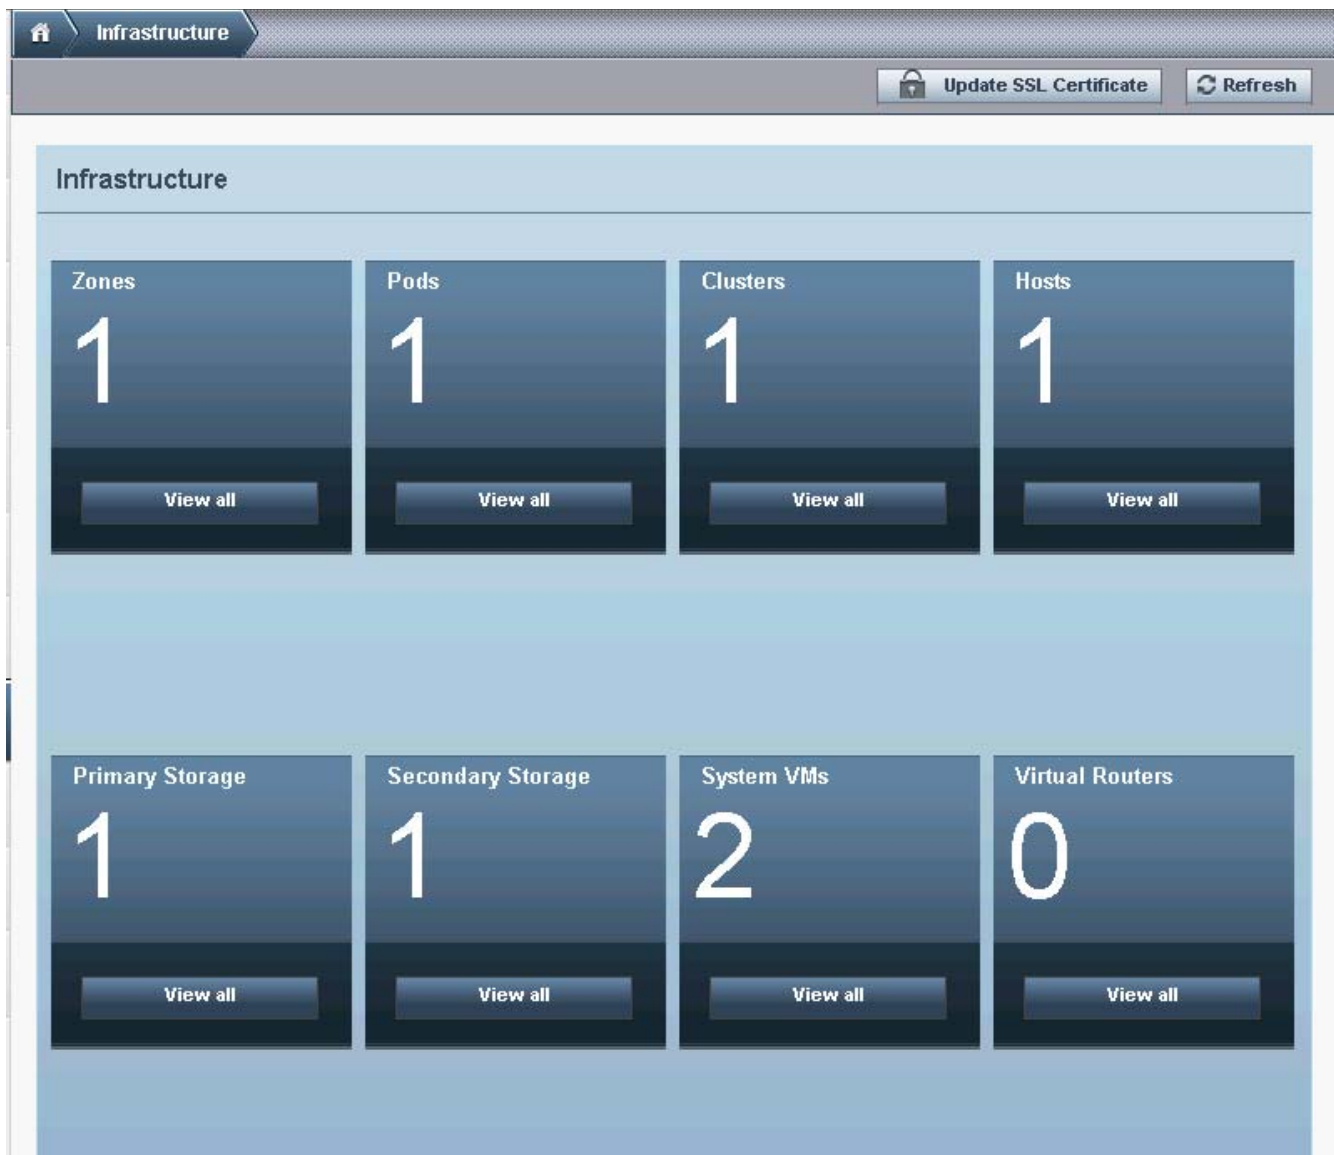
**Note**

After the Zones are created a window appears with message Zone is ready to launch; proceed to the next step. Click **Launch Zone** button, which then launches a window to confirm the Enabling of the Zone. Click **Yes** to enable the Zone.

**Figure 200**      *Enabling the Zone***Note**

The Cloud Infrastructure Window appears displaying all the components defined in the Zone.

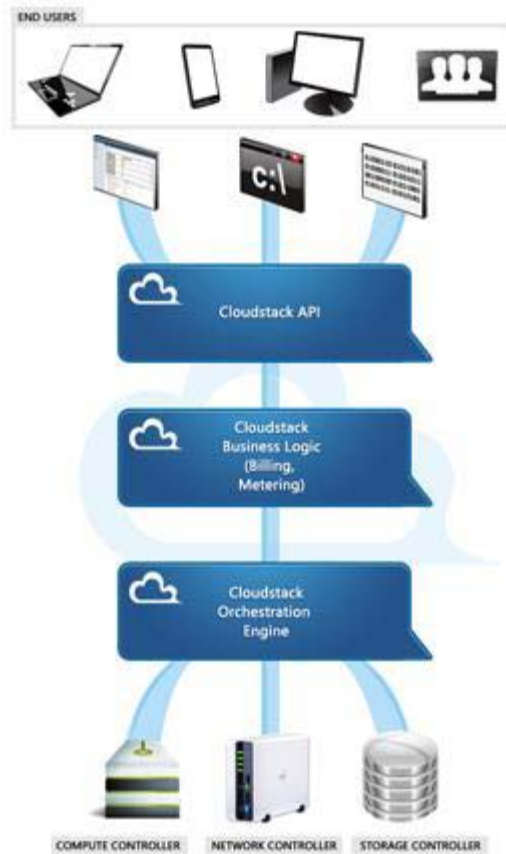
**Figure 201**      *Displaying the Zones and Other Components*



## Cloud Infrastructure Design

This section outlines Citrix CloudPlatform infrastructure offering with service levels provided to multiple tenants to host cloud services. The infrastructure offering includes Compute, Network and Storage resources which are provided in a shared model to tenants; therefore providing Service Levels to these resources with cost based on quality of services.

**Figure 202**      **CloudPlatform Infrastructure Design**



Citrix CloudPlatform, Cisco UCS and NetApp Storage offer cloud infrastructure resources which are provisioned, configured and deployed based on the service levels definitions with complete integrated, multi-tenant solution in a cloud. In this section the following infrastructure offerings are enabled, configured and deployed in the Citrix Cloud Platform:

- Cloud Compute
- Cloud Storage
- Cloud Network

## Cloud Compute

Cisco UCS System offers Service Profile Templates and Service Profiles with compute server pools and server pool qualification policy definitions based on CPU and Memory requirements which can be related to service level costs definition in service catalog. These SLA definitions can be mapped on Compute by defining server pool qualification definitions, Cisco UCS System will automatically assign compute resources to the correct Service Profile and place the server into its respective server pool. Cisco UCS offers Service Profile Templates and Profiles which provide a single window to configure compute, network and storage definition on host, derived using server pools to provision multi-tenant cloud requirement.



CloudPlatform provides single pane of management window to manage, configure and deploy Cloud Compute on Cisco Unified Computing System using a plugin which calls Cisco UCS Manager API to perform compute operations. Once the Cisco UCS Manager Plugin is registered with CloudPlatform, the Cloud Admin can quickly provision and associate Service Profile to relevant B-Series Blade for installing ESXi hypervisor OS server.

Host management for vSphere is done through a combination of vCenter and the CloudPlatform admin UI. CloudPlatform requires all hosts added in a CloudPlatform cluster, are part of VMware vCenter Cluster defined under DataCenter. Clusters of multiple hosts allow for features like vMotion migration. CloudPlatform allows to perform VM Migration operation which will internally call vCenter API to perform vMotion operations on VMs part of Clusters. To perform VM vMotion you require shared FC or NFS storage.

Based on the Service Level definitions Citrix CloudPlatform can define multiple clusters in a Zone to support high availability, resource on demand, performance and capacity. Table 35 shows four sample cluster definitions that can be created on TenantA Zone.

**Table 35** *Cloud Platform Hosts in Cluster Definitions for all the Service Levels*

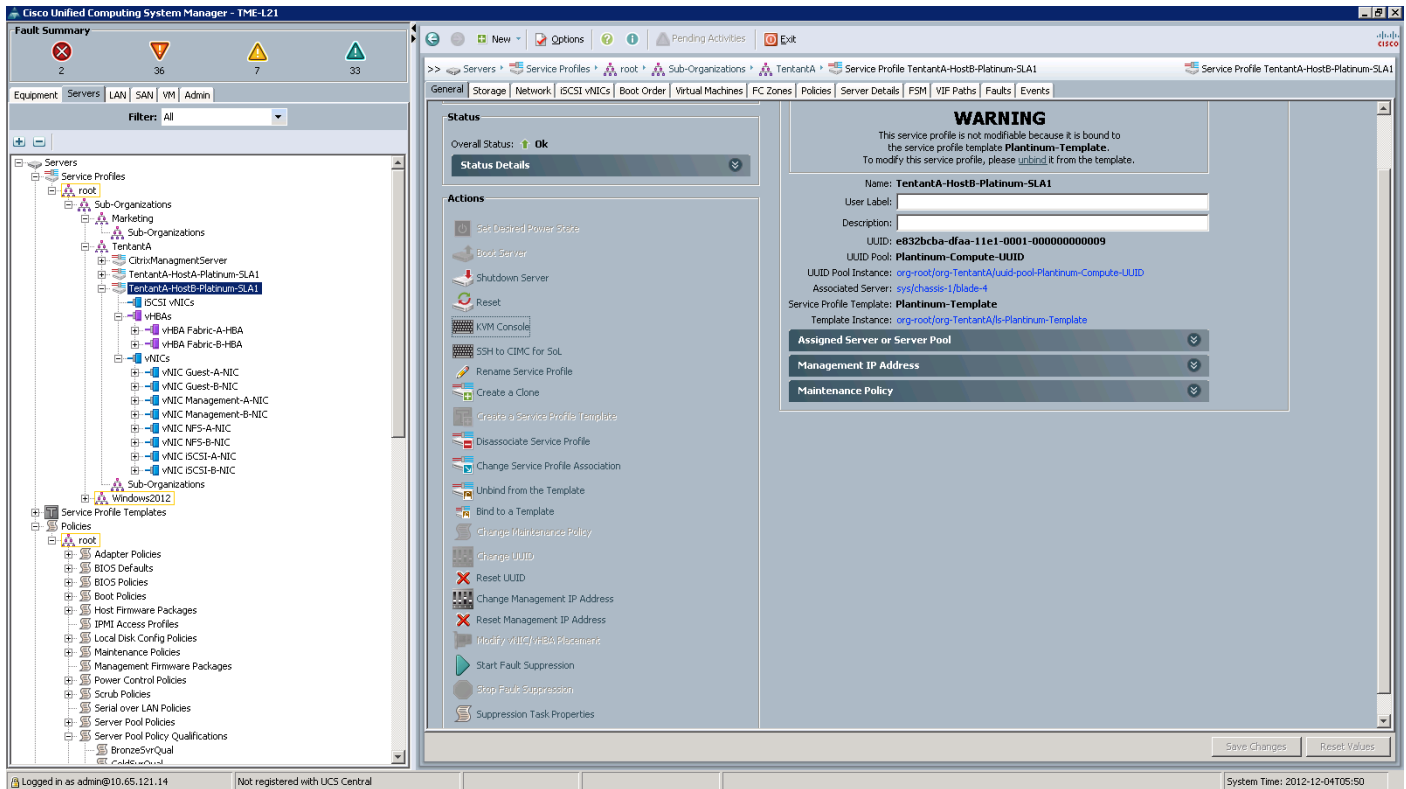
| Cluster Name             | Host Resources                                                 | UCS Service Profile Template | Compute Qualification Min - Max         |
|--------------------------|----------------------------------------------------------------|------------------------------|-----------------------------------------|
| TenantA-Platinum-Cluster | First Highest CPU & Memory Performance<br>High Availability    | Platinum-Template            | 32 – 40 CPU Core<br>128 – 256 GB Memory |
| TenantA-Gold-Cluster     | Second Highest CPU and Memory Performance<br>High Availability | Gold-Template                | 32 – 40 CPU Core<br>64 – 128 GB Memory  |
| TenantA-Silver-Cluster   | Third Highest CPU and Memory Performance<br>High Availability  | Silver-Template              | 12 – 16 CPU Core<br>32 – 64 GB Memory   |
| TenantA-Bronze-Cluster   | Fourth Highest CPU and Memory Performance<br>High Availability | Bronze-Template              | 12 – 16 CPU Core<br>16 – 32 GB Memory   |

In this paper VMware Cluster TenantA-Cluster part of vCenter Data Center TenantA-DC with two ESXi host TenantA-HostA and TenantA-HostB ESXi are added to provide high availability, and define Compute Offering for allocating Virtual Machines using Citrix CloudPlatform. The steps below show creation of Compute Offering with CPU, Memory, Network, Host and High Availability definitions for allocating Virtual Machine using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services.

To deploy ESXi Server on second host TenantA-HostB refer Cloud Host Preparations for creating service profile based on Platinum-Template with compute, network and storage infrastructure definitions with service levels.

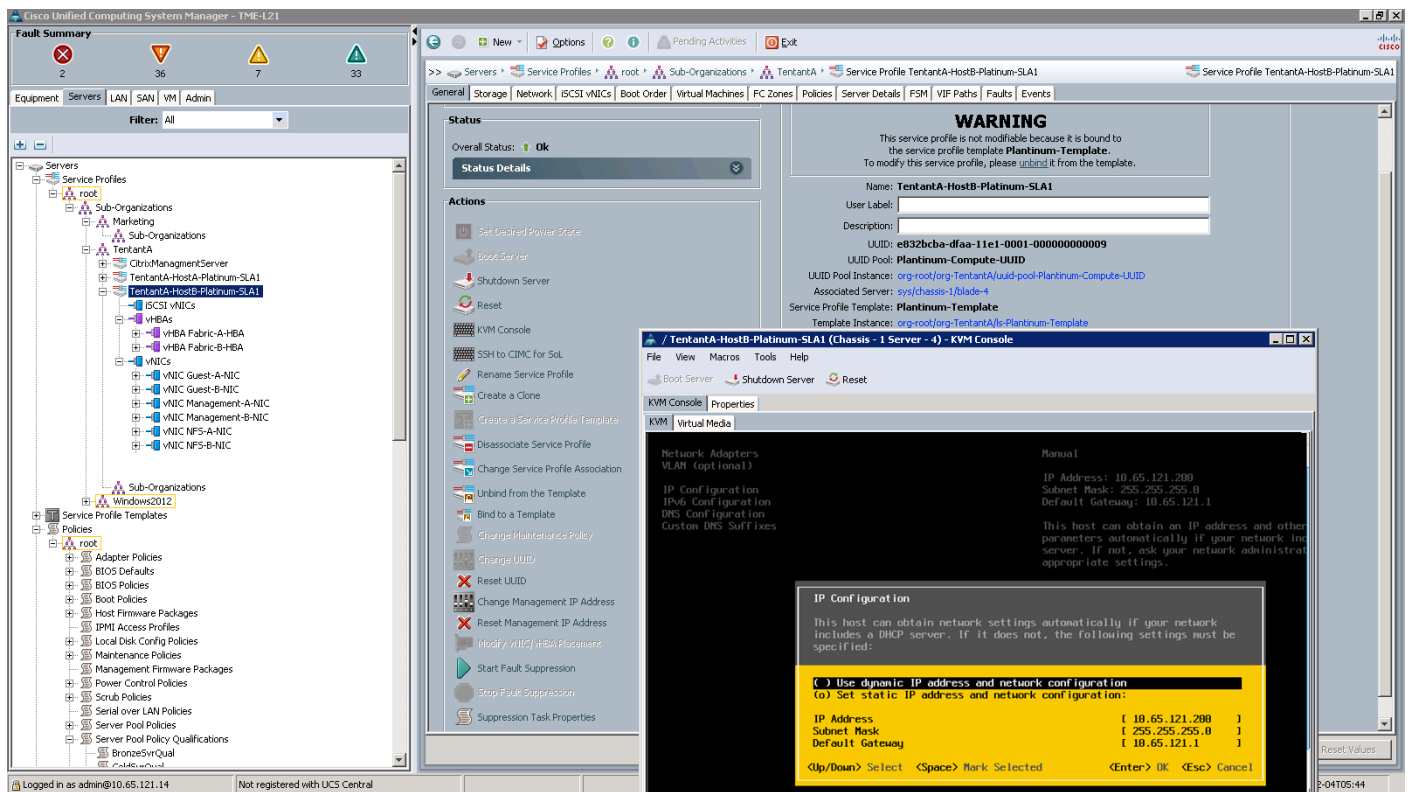
Login to Cisco UCS Manager with User TenantA-Admin created earlier for Organization TenantA:

**Figure 203**      *Displaying the Service Profile of the Tenant A*



1. Click the **Server** tab in the left pane.
2. Choose **Service Profiles > Sub-Organization**. Expand **TenantA**.
3. Follow the steps mentioned in CloudPlatform Host Preparation section to create Service Profile <TenantA-HostB-Platinum-SLA2>.
4. Click **Service Profile <TenantA-HostB-Platinum-SLA2>**.
5. Click **KVM Console** in the right pane.

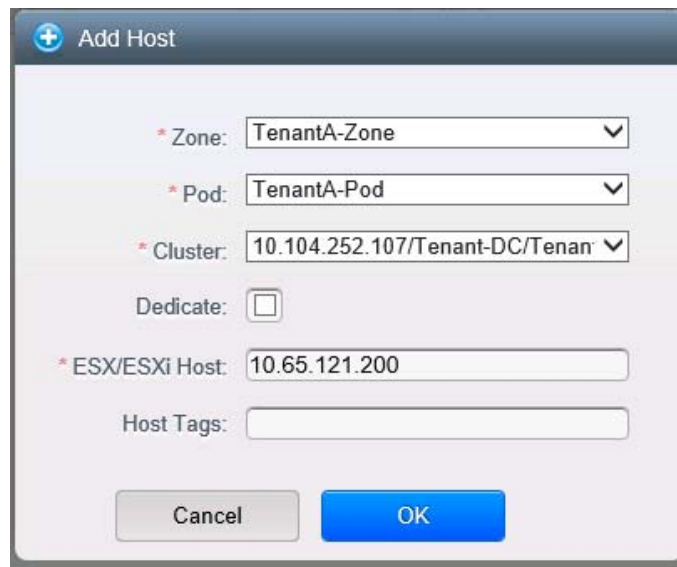
**Figure 204**      *Displaying the Status of the ESXi Server*



6. Add ESXi Host <10.65.121.200> to vCenter Cluster <TenantA-Cluster> which is part of vCenter DataCenter <Tenant-DC>. CloudPlatform will automatically detect ESXi Host <10.65.121.200> under Hosts tab.
  - a. In the VMware vSphere Client connected to vCenter, choose Home > Host and Clusters.
  - b. Expand the vCenter, DataCenter <Tenant-DC>, and Cluster <TenantA-Cluster> .
  - c. Click **Add a host** on Basic Tasks.
  - d. Enter <10.65.121.200> in Host text box, Enter <root> in Username text box <xxxx> in Password and click **Next**.
  - e. Click **Yes** to confirm Security Alert.
  - f. Click **Next** in Host summary.
  - g. Click **Assign a new license key to host** radio button.
  - h. Click **Next**.
  - i. Uncheck **Enable Lockdown Mode**.
  - j. Click **Next**.
  - k. Click **Finish**.
7. Add ESXi Host <10.65.121.200> under Hosts tab in CloudPlatform once Cluster <TenantA-Cluster> is added to DataCenter <Tenant-DC> in vCenter.
  - a. Login to CloudPlatform with UserName and Password.
  - b. Click **Infrastructure** tab.

- c. On left pane click on **Hosts**.
- d. Click **Add Host**.
- e. Select TenantA-Zone in Zone list box.
- f. Select Tenant-Pod in Pod list box.
- g. Select <10.104.252.107/Tenant-DC/TenantA-Cluster>in Cluster list box.
- h. Enter <10.65.121.200> in ESX/ESXi Host Name field.
- i. Click **OK**.

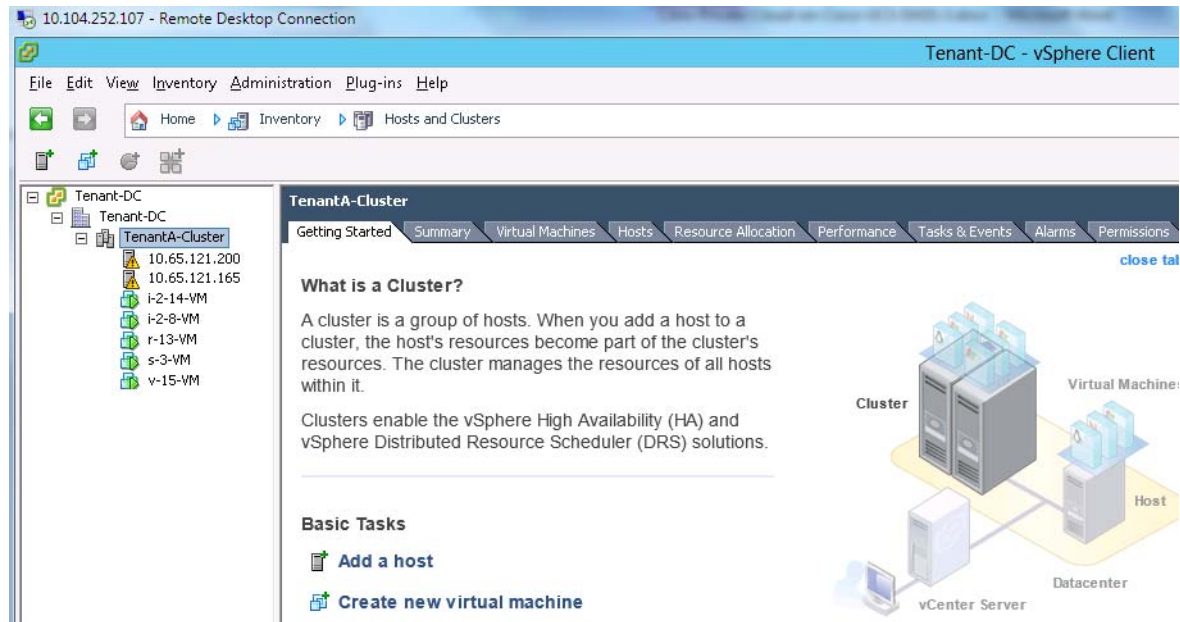
**Figure 205**      **Adding second cluster host in CloudPlatform**



The screenshot shows the 'Add Host' dialog box in the CloudPlatform interface. The dialog is titled 'Add Host' and contains the following fields and controls:

- Zone:** A dropdown menu with 'TenantA-Zone' selected.
- Pod:** A dropdown menu with 'TenantA-Pod' selected.
- Cluster:** A dropdown menu with '10.104.252.107/Tenant-DC/Tenan' selected.
- Dedicate:** A checkbox that is currently unchecked.
- ESX/ESXi Host:** A text field containing the IP address '10.65.121.200'.
- Host Tags:** An empty text field.
- Buttons:** 'Cancel' and 'OK' buttons at the bottom.

8. Verify host <10.65.121.200> is added on vCenter and CloudPlatform.

**Figure 206** VMWare vCenter DataCenter and Cluster Configuration**Figure 207** CloudPlatform Hosts displaying ESXi Cluster Hosts

## Defining Tags

A tag is a text string attribute associated with primary storage, a disk offering, or a compute Offering. Tags allow administrators to provide additional information about the cloud service offering and they are matched against tags placed on compute, service and disk offerings. CloudPlatform allows defining required tags on compute, service and disk offerings to allocate Virtual Machines Host placement and data disks on the primary storage.

### Host Tags

1. Click **TenantA-HostA**.
2. Click **Edit** icon.
3. Enter TenantA-HostA name in Host Tags field.
4. Click **Apply**.

**Figure 208**      **Adding the Tags to the Host**

The screenshot displays the 'TentantA-Host' configuration page in the Cloud Infrastructure Design interface. The left pane shows a list of hosts with 'TentantA-Host' selected. The right pane shows the 'Details' tab for this host, with fields for Name, ID, Resource state, State, Type, Host Tags, OS Preference, Zone, and Pod. The 'Host Tags' field is currently empty, and the 'OS Preference' dropdown is set to 'None'. The 'Apply' button is highlighted in blue.

| Name          | Zone          |
|---------------|---------------|
| TentantA-Host | TentantA-Zone |

Details
Statistics

View Instances

Name
TentantA-Host

ID
e313faed-4e0d-40d5-85fb-0ecee8b8798c

Resource state
Enabled

State
Up

Type
Routing

Host Tags
TentnatA-HostA

OS Preference
None

Zone
TentnatA-Zone

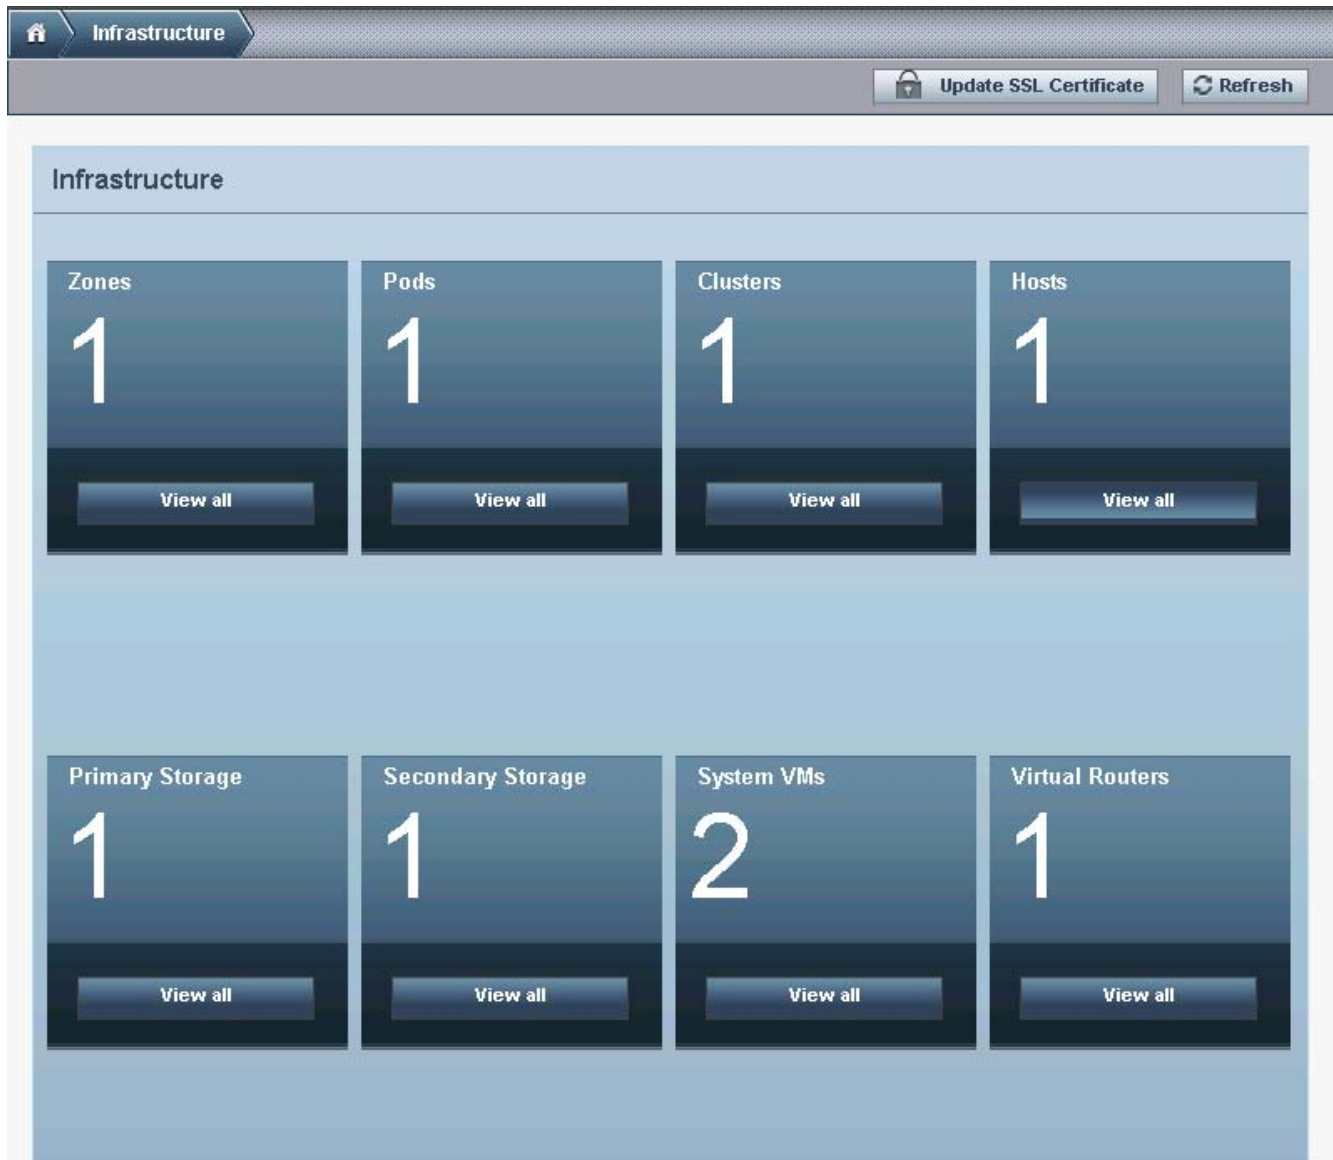
Pod
TentantA-Pod

Apply
Cancel

## Storage Tags

1. Click the **Infrastructure** tab.
2. Click the **Primary Storage** in the right pane.

**Figure 209**      **Selecting the Primary Storage**



3. Click **TenantAPlatinumNFSStorage**.
4. Click **Edit** icon
5. Enter **Platinum-NFS-Storage** in Storage Tags field.
6. Click **Apply**.

**Figure 210**      *Applying the Tags to the Storage*

The screenshot shows the CloudPlatform web interface for configuring storage. The breadcrumb navigation at the top indicates the path: Infrastructure > Primary Storage > Primary Storage > TenantAPrimaryNFSSStorage. A 'Refresh' button is located in the top right corner. Below the navigation, there are two tabs: 'Details' (selected) and 'Settings'. The 'Details' tab contains a table of storage properties. The 'Storage Tags' property is currently being edited, with a text input field containing 'Platinum-NFS-Storage'. At the bottom of the interface, there are 'Apply' and 'Cancel' buttons.

|              |                                          |
|--------------|------------------------------------------|
| Name         | TenantAPrimaryNFSSStorage                |
| ID           | 548f325f-eb3b-3f5c-8ffb-0c2a82c997c9     |
| State        | Up                                       |
| Storage Tags | Platinum-NFS-Storage                     |
| Zone         | TenantA-Zone                             |
| Pod          | TenantA-Pod                              |
| Cluster      | 10.104.252.107/Tenant-DC/TenantA-Cluster |
| Type         | NetworkFilesystem                        |

## Service Offering

A service offering is a set of virtual hardware features such as CPU core count and speed, memory, and disk size. The CloudPlatform administrator can set up various offerings, and then end users choose from the available offerings when they create new Virtual Machines with compute, network and storage resources.

CloudPlatform separates service offerings into compute offerings and disk offerings.

The computing service offering specifies:

- Guest CPU
- Guest RAM



- Host Tags
- Storage Tags
- Network rate

The disk offering specifies:

- Disk Size
- Storage Tags

## Defining the Compute Offering

1. Click **Service Offering** tab.

**Figure 211**      *Displaying the Compute Service Offerings*

| Name             | Description                 | Order     |
|------------------|-----------------------------|-----------|
| Platinum-Compute | TenantA-Platinum-Compute VM | ▲ ▼ ▲ ▼ ≡ |
| Platinum         | Platinum-NFS                | ▲ ▼ ▲ ▼ ≡ |
| Medium Instance  | Medium Instance             | ▲ ▼ ▲ ▼ ≡ |
| Small Instance   | Small Instance              | ▲ ▼ ▲ ▼ ≡ |

2. Select Compute Offering in Select Offering list box.
3. Click the **Add compute offering** button.
4. Enter Platinum-Compute in the Name field.
5. Enter TenantA-Platinum-Compute VM in the Description field.
6. Select shared in Storage Type list box.
7. Enter 16 in No of CPU Cores field.
8. Enter 2GHz in CPU field.
9. Enter 32GB in Memory field.
10. Enter 8G in Network rate field.
11. Check **Offer HA** check box.
12. Enter Platinum-NFS-Disk in Storage Tags.
13. Enter TenantA-HostA in Host Tags.
14. Check **CPU Cap** check box.
15. Click **OK**.

**Figure 212** Define the Compute Offerings Attributes

**Add compute offering**

\* Name:  \* Description:

Storage Type:  \* # of CPU Cores:

\* CPU (in MHz):  \* Memory (in MB):

Network Rate (Mb/s):  Disk Read Rate (BPS):

Disk Write Rate (BPS):  Disk Read Rate (IOPS):

Disk Write Rate (IOPS):

Storage Tags:  Host Tags:

CPU Cap: ☒ Offer HA: ☐

isVolatile: ☒ Public: ☒

Deployment Planner:

Planner Mode:

Based on the Service Level definitions Citrix CloudPlatform can define multiple compute offering in a Zone for provisioning tenant's virtual machines to support high availability, resource on demand, performance and capacity. Table 36 shows four compute offering definition that can be created on TenantA Zone.

**Table 36** Compute Offering Definitions for Four Service Classes

| Compute Offering | VM CPU Cores and MHz | Memory | Network Rate | Disk Write / Read IOPS / Rate | Storage / Host Tags                       |
|------------------|----------------------|--------|--------------|-------------------------------|-------------------------------------------|
| Platinum-Compute | 8 Core<br>2 GHz      | 32 GB  | 8GB          | 110000000                     | Platinum-Storage<br>TenantA-Platinum-Host |
| Gold-Compute     | 6 Core<br>2 GHz      | 16 GB  | 6 GB         | 810000000                     | Gold-Storage<br>TenantA-Gold-Host         |
| Silver-Compute   | 4 Core<br>1 GHz      | 8 GB   | 4Gb          | 540000000                     | Silver-Storage<br>TenantA-Silver-Host     |
| Bronze-Compute   | 2 Core<br>1 GHz      | 4 GB   | 2 Gb         | 270000000                     | Bronze-Storage<br>TenantA-Bronze-Host     |

## Cloud Storage

This section provides storage design based on service levels definition for hosting multi-tenants data on the cloud. The storage spaces are provisioned by NetApp Storage System based on the service levels cost and exposed to CloudPlatform application as primary storage which allows defining and configuring storage to guest Virtual Machines to store operating system and data files to host multi-tenants.

The Primary storage is associated with a cluster, and it stores the OS and disk volumes for all the virtual machines running on hosts in that cluster which are part of zone. These primary stores can be accessed over Fibre Channel, NFS or iSCSI protocol.

Based on the Service Level definitions Citrix CloudPlatform can define multiple primary storage in a Zone to support high availability, resource on demand, performance and capacity, with these definition as listed in [Table 37](#) shows three primary storage definitions that can be created on clusters on TenantA Zone.

**Table 37**      **Primary Storage Definitions for the Service Classes**

| Cluster Name             | Primary Storage (Protocol)                |
|--------------------------|-------------------------------------------|
| TenantA-Platinum-Cluster | TenantA-FC-PrimaryStorage (Fibre Channel) |
|                          | TenantA-PrimaryNFS-Storage (NFS)          |
| TenantA-Gold-Cluster     | TenantA-Gold-FC-Storage (Fibre Channel)   |
|                          | TenantA-Gold-NFS-Storage (NFS)            |
| TenantA-Silver-Cluster   | TenantA-Silver-FC-Storage (Fibre Channel) |
|                          | TenantA-Silver-NFS-Storage (NFS)          |
| TenantA-Bronze-Cluster   | TenantA-Bronze-FC-Storage (Fibre Channel) |
|                          | TenantA-Bronze-NFS-Storage (NFS)          |

In this paper we will create TenantA-Platinum-Storage, TenantA-Platinum-NFS-Storage and TenantA-Platinum-FC--Storage primary storage in TenantA-Platinum-Cluster for allocating data disks to Virtual Machines using CloudPlatform.

The steps below shows creation of primary storage on Fibre Channel and NFS for allocating data disks to Virtual Machine using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services.

Login to CloudPlatform with User credentials to configure Primary Storage on TenantA Zone.

### Adding the FC Primary Storage

1. Click **Add Primary Storage** button.
2. Select Cluster in Scope list box.
3. Select TenantA-Zone in Zone list box.
4. Select TenantA-Pod in Pod list box.
5. Select <vCenter\_IP\_Address>10.104.252.107/Tenant-DC-/TenantA-Cluster.
6. Enter TenantA-FC-PrimaryStorage in the Name field.
7. Choose vmfs under Protocol list box.

8. Enter <vCenter\_IP\_Address>10.104.252.107 in Server field.
9. Enter Tenant-DC in vCenter Datacenter.
10. Enter TenantA-FC-PrimaryStorage in vCenter Datastore.
11. Click OK.

**Figure 213** *Defining the FC Primary Storage Information*

**View:** Datastores Devices

**Datastores**

| Identification                   | Status | Device              | Drive Type | Capacity  |
|----------------------------------|--------|---------------------|------------|-----------|
| 548f325feb3b3f5c8ffb0c2a82c997c9 | Normal | 193.191.1.20:/Te... | Unknown    | 285.00 GB |
| datastore1 (2)                   | Normal | NETAPP Fibre Ch...  | Non-SSD    | 10.00 GB  |
| TenantA-FC-PrimaryStorage        | Normal | NETAPP Fibre Ch...  | Non-SSD    | 179.75 GB |

CloudPlatform™

Notifications Project: Default view admin cloud CITRIX

Infrastructure Primary Storage

+ Add Primary Storage

| Name                      | Server                                               | Path                                 | Cluster | Scope   | Quickview |
|---------------------------|------------------------------------------------------|--------------------------------------|---------|---------|-----------|
| TenantA-FC-PrimaryStorage | VMFS datastore: /Tenant-DC/TenantA-FC-PrimaryStorage | /Tenant-DC/TenantA-FC-PrimaryStorage |         | CLUSTER | +         |
| TenantA-PrimaryNFS        |                                                      |                                      |         | ZONE    | +         |

**+ Add Primary Storage**

Scope: Cluster

\* Zone: TenantA-Zone

\* Pod: TenantA-Pod

\* Cluster: 10.104.252.107/Tenant-DC/Tenan

\* Name: TenantA-FC-PrimaryStorage

\* Protocol: vmfs

\* Server: 10.104.252.107

\* vCenter Datacenter: Tenant-DC

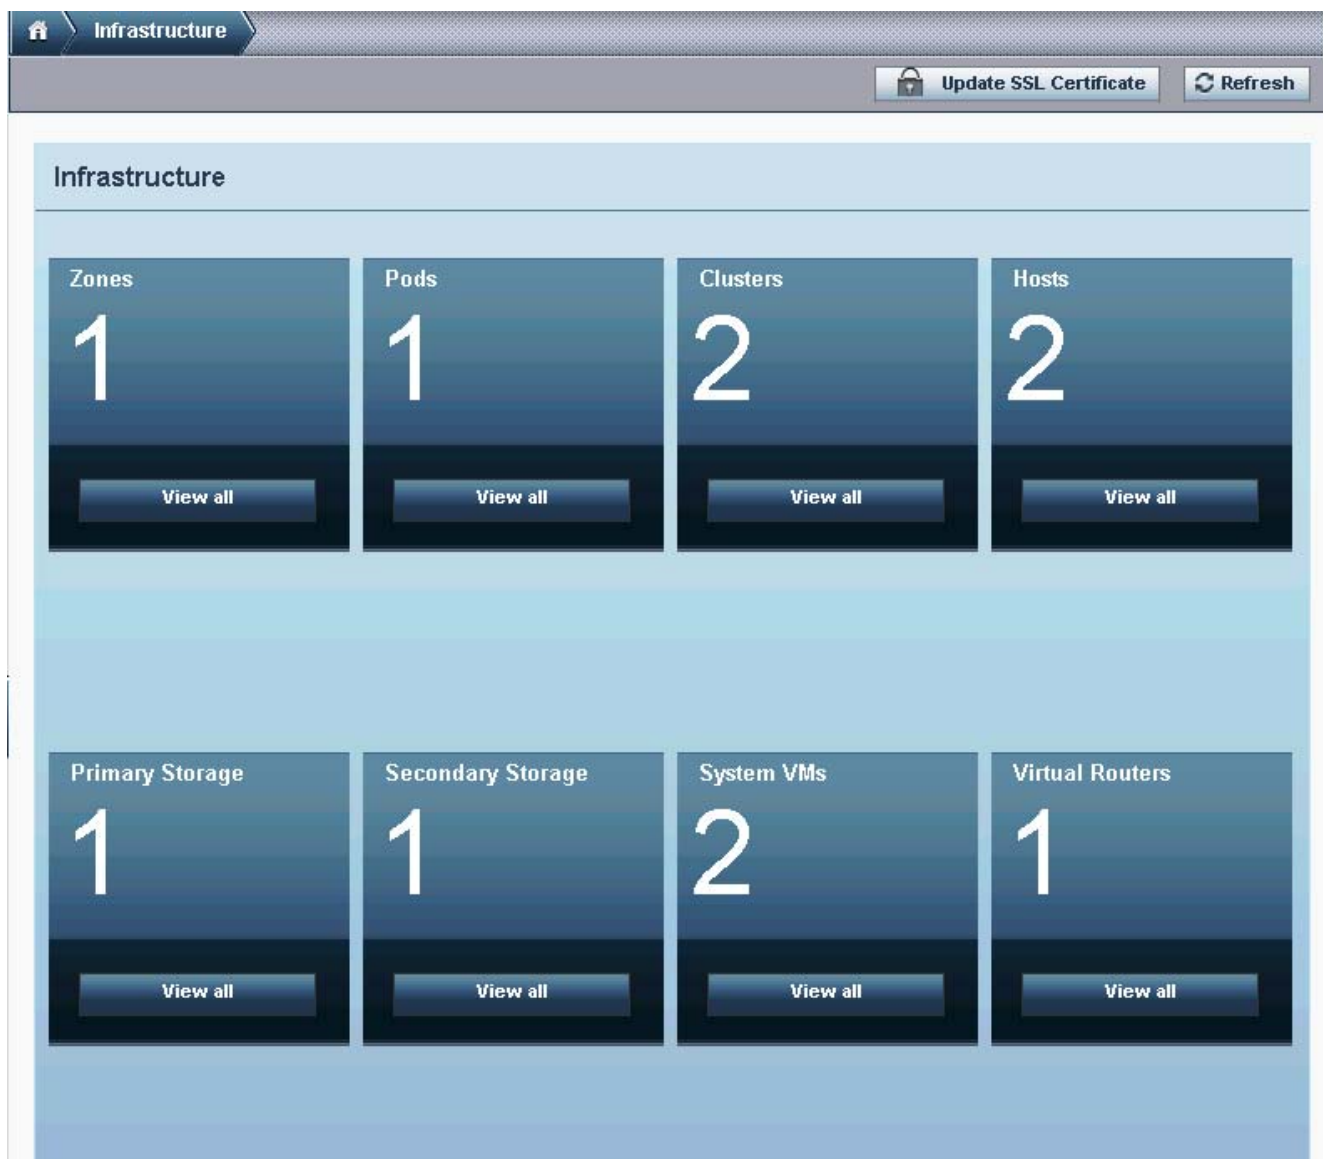
\* vCenter Datastore: TenantA-FC-PrimaryStorage

Storage Tags:

## Adding the NFS Primary Storage

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Infrastructure** tab on the left Side Pane.
4. Click the **View All** under Primary Storage in the right pane.

**Figure 214**      **View Primary Storage**



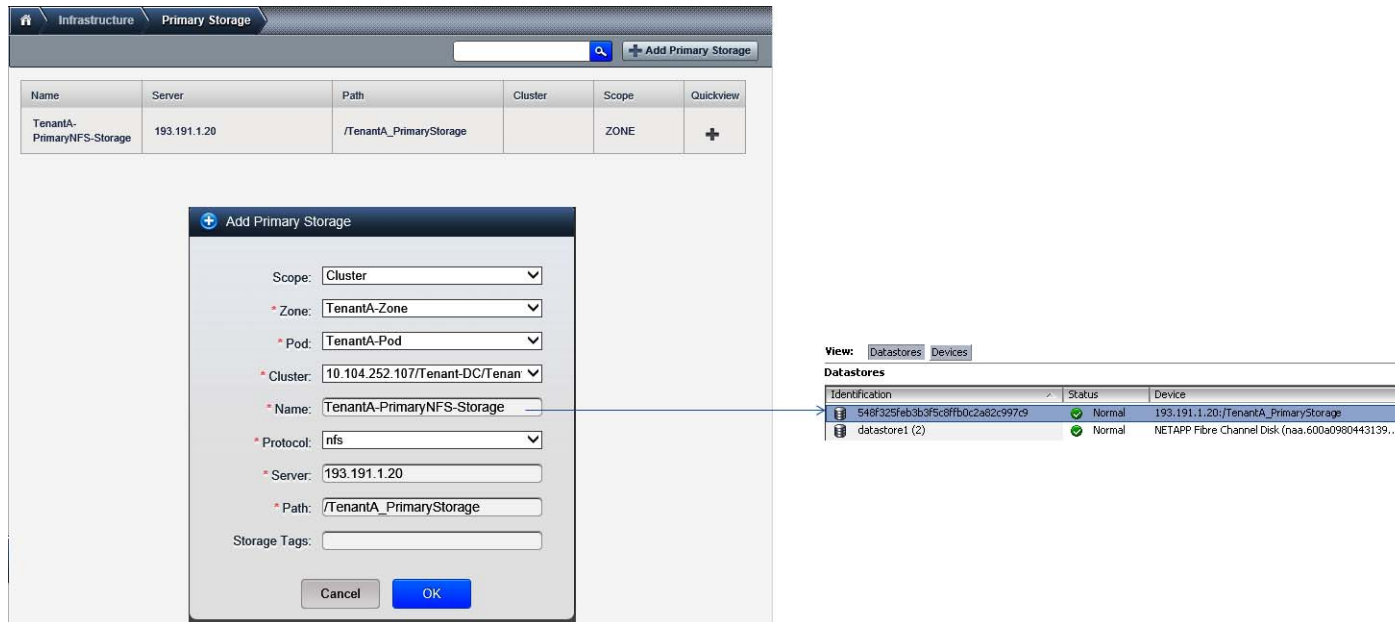
## Add NFS Primary Storage

1. Click **Add Primary Storage**.
2. Select Cluster in Scope list box..
3. Select TenantA-Zone in Zone list box.
4. Select TenantA-Pod in Pod list box.
5. Select <vCenter\_IP\_Address> 10.104.252.107/Tenant-DC-/TenantA-Cluster.
6. Enter TenantA-PrimaryNFS-Storage in the Name field.
7. Choose nfs under Protocol list box.
8. Enter <NetApp\_NFS\_LIF\_IP\_Address>193.191.1.20 in Server field.

9. Enter /TenantA\_PrimaryStorage in Path.
10. Enter **OK**.

On successful addition of NFS Primary Storage on CloudPlatform, corresponding NFS Datastore will be automatically configured on ESX Hosts part of VMWare Cluster under DataCenter.

**Figure 215** Defining NFS Primary Storage Information



## Cloud Storage Volume

To provide data storage access to Tenant Virtual Machines based on service levels definition created on primary storage, CloudPlatform offers volume, a unit of storage which can provide additional data disk for a specific VM. CloudPlatform defines a volume as a unit of storage available to a guest VM.

Data disks provide for additional storage (e.g. As “/opt” or “D:”). Every guest VM has a root disk, and VMs can also optionally have a data disk. End users can mount multiple data disks to guest VMs. Users choose data disks from the disk offerings created by cloud administrators.

Login to CloudPlatform with user credentials to configure Primary Storage on TenantA Zone

### Defining the Disk Offering

1. Click the **Service Offering** tab.
2. Select Offering in Select Offering list box.
3. Click **Add Disk Offering** button.

**Figure 216**      **Selecting the Disk Offering**

| Name   | Description        | Custom Disk Size | Disk Size (in GB) | Order     |
|--------|--------------------|------------------|-------------------|-----------|
| Small  | Small Disk, 5 GB   | No               | 5                 | ▲ ▼ ▲ ▼ ≡ |
| Medium | Medium Disk, 20 GB | No               | 20                | ▲ ▼ ▲ ▼ ≡ |
| Large  | Large Disk, 100 GB | No               | 100               | ▲ ▼ ▲ ▼ ≡ |
| Custom | Custom Disk        | Yes              | N/A               | ▲ ▼ ▲ ▼ ≡ |

4. Enter Platinum-NFS-Storage in the Name field.
5. Enter Platinum-NFS-Disk in the Description field.
6. Select shared in Storage Type list box.
7. Check Custom Tags check box.
8. Leave QoS Type blank
9. Enter Platinum-NFS-Storage in Storage Tags.
10. Check the Public check box.
11. Click **OK**.

**Figure 217**      **Defining the Disk Offering Information**

**+ Add Disk Offering**

\* Name:

\* Description:

Storage Type:

Custom Disk Size: ☒

QoS Type:

Storage Tags:

Public: ☒

Based on the Service Level definitions Citrix CloudPlatform can define multiple disk offerings in a Zone for provisioning tenant's virtual machines to support high availability, resource on demand, performance and capacity. [Table 38](#) shows four disk offering definitions that can be created on TenantA Zone.

**Table 38** Disk Offering Definitions for all the Service Classes

| Disk Offering         | Storage Tags         |
|-----------------------|----------------------|
| Platinum-Storage-Disk | Platinum-NFS-Storage |
| Gold-Storage-Disk     | Gold-NFS-Storage     |
| Silver-Storage-Disk   | Silver-NFS-Storage   |
| Bronze-Storage-Disk   | Bronze-NFS-Storage   |

## Creating Storage Volume

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Storage** tab on the left side pane.
4. Click **Add Volume** in the right pane.
5. Enter Platinum-NFS-Data-Vol in the Name field.
6. Select TenantA-Zone in Availability Zone list box.
7. Select Platinum-NFS-Disk in Disk Offering list box.
8. Enter 20 value in Disk Size field.

**Figure 218** Defining the Volume Data

**Add Volume**

Please fill in the following data to add a new volume.

\* Name:

Availability Zone:

Disk Offering:

\* Disk Size (in GB):

Based on the Service Level definitions Citrix CloudPlatform can define multiple volumes based on the disk offering requirement in a Zone to support high availability, resource on demand, performance and capacity, with these definitions. [Table 39](#) shows four volume definitions that can be created for allocating data disk to Virtual Machines on TenantA Zone.



**Table 39**      **Volume Definition Based on Disk Offering**

| Volume Name           | Disk Offering     |
|-----------------------|-------------------|
| Platinum-NFS-Data-Vol | Platinum-NFS-Disk |
| Gold-NFS-Data-Vol     | Gold-NFS-Disk     |
| Silver-NFS-Data-Vol   | Silver-NFS-Disk   |
| Bronze-NFS-Data-Vol   | Bronze-NFS-Disk   |

## Cloud Network

To provide network isolation and to offer custom network services based on the service levels definitions for specific Tenants network, CloudPlatform offers advanced zone network traffic types and network offering which can be applied on zone bases.

In this paper the Zones have been configured to use advanced zone networking where multiple physical networks can carry one or more traffic types, and to the admin must define in the CloudPlatform as to what type of network traffic is carried by the network. The traffic types in an advanced zone are mentioned below.

- Guest

When end users run VMs, they generate guest traffic. The guest VMs communicate with each other over a network that can be referred to as the guest network. This network can be isolated or shared. In an isolated guest network, the administrator needs to reserve VLAN ranges to provide isolation for each CloudPlatform account's network (potentially a large number of VLANs). In a shared guest network, all guest VMs share a single network.

- Management

When CloudPlatform's internal resources communicate with each other, they generate management traffic. This includes communication between hosts, system VMs (VMs used by CloudPlatform to perform various tasks in the cloud), and any other component that communicates directly with the CloudPlatform Management Server. You must configure the IP range for the system VMs to use.

- Public

This traffic is generated when VMs in the cloud access the Internet. Publicly accessible IPs must be allocated for this purpose. End users can use the CloudPlatform UI to acquire these IPs to implement NAT between their guest network and the public network

- Storage

This traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudPlatform uses a separate Network Interface Controller (NIC) named storage NIC for storage network traffic.



### Note

Refer Cloud Host Preparation and Cloud Host Deployment sections for more details on designing and configuring these network traffic types on different physical network devices on Cisco UCS Compute System.

## Network Offering

The CloudPlatform administrator can create any number of custom network offerings, in addition to the default network offerings provided by CloudPlatform. By creating multiple custom network offerings, the cloud is equipped to offer different classes of service on a single multi-tenant physical network and this also ensures various levels of services based on the service levels cost definition.

Example: Considering there are two tenants-- tenant A may require only the simple firewall protection for their website, while tenant B may be running a web server farm and require a scalable firewall solution, load balancing solution, and alternate networks for accessing the database backend.

In this paper default isolated network offering is used that will be defined during zone advanced network configuration. Based on the Service Level definitions and network offering services Citrix CloudPlatform can define multiple networks offering in a Zone.

Login to CloudPlatform with User credentials to configure Primary Storage on TenantA Zone

### Creating Network Offering

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Service Offering** tab.
4. Select Network Offering in Select Offering list box.
5. Click **Add network offering**.
6. Enter TenantA-Zone-Network in the Name field.
7. Enter TenantA Network in the Description field.
8. Select Isolated in Guest Traffic list box.
9. Check **DHCP**, **DNS**, **UserData** check boxes in Required Services list box.
10. Click **OK**.

**Figure 219**      **Defining the Network Offering Properties**

**Add network offering**

\* Name:

\* Description:

Network Rate (MB/s):

Guest Type:

Specify VLAN: ☐

VPC: ☐

Supported Services:

- VPN: ☐
- DHCP: ☒
- DHCP Provider:
- DNS: ☒
- DNS Provider:
- Firewall: ☐
- Load Balancer: ☐
- User Data: ☒












System Offering:

Conserve mode: ☐

Tags:

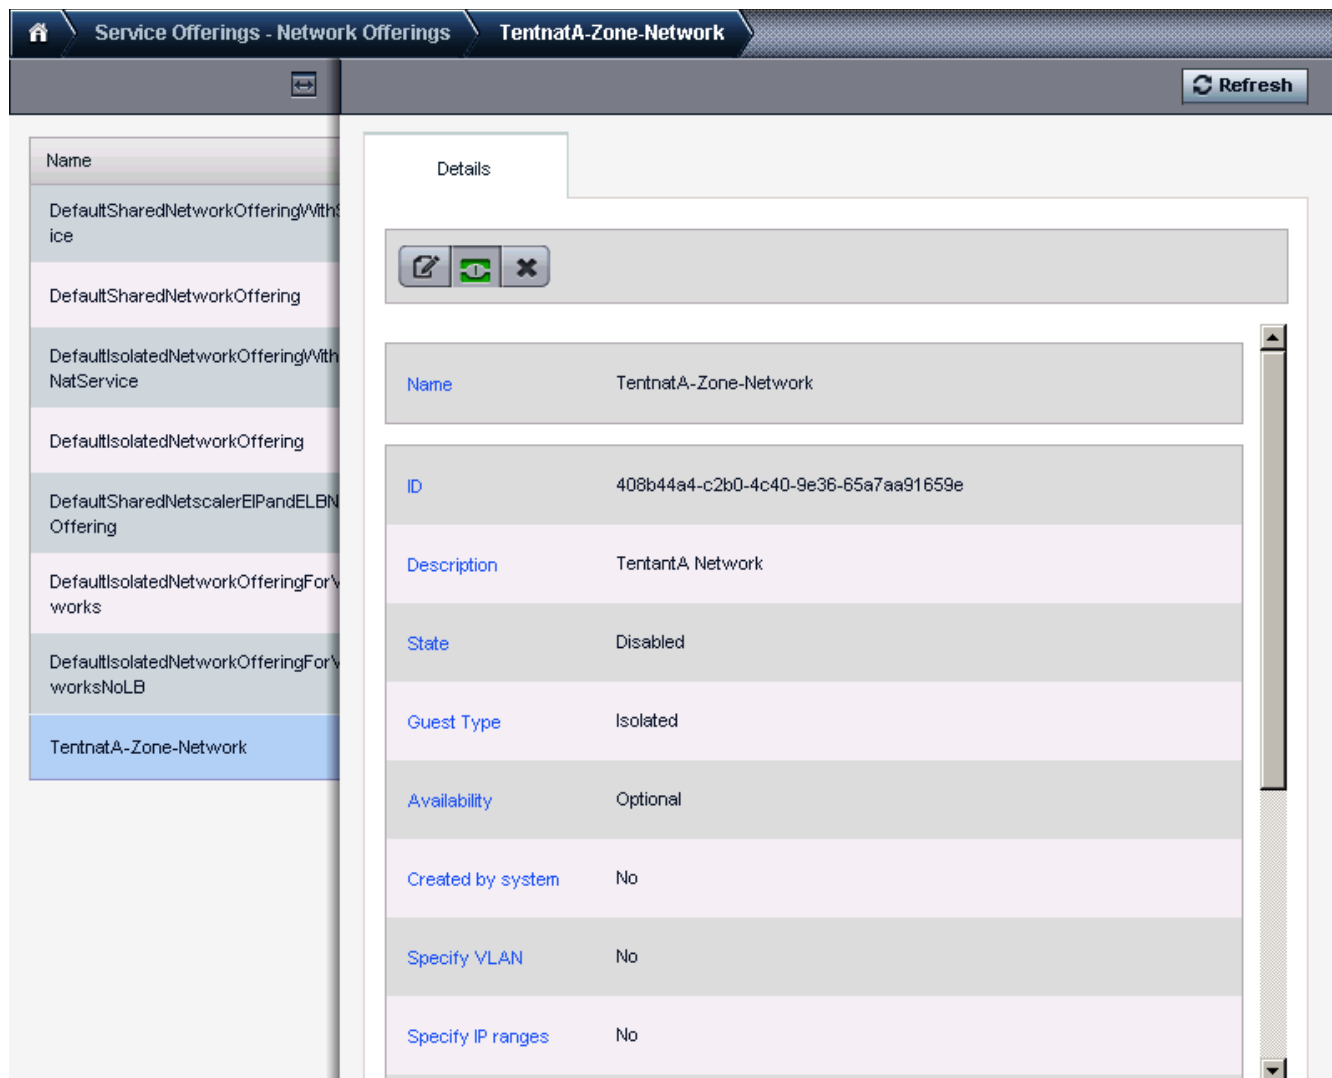
11. Click **TenantA-Zone-Network** radio button, which is disabled.

Figure 220 Displaying Network Offerings Status

| Service Offerings - Network Offerings              |                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Select offering: <span>Network Offerings</span>    | <input type="text"/>                                                                       | <input type="button" value="+ Add network offering"/>                                                                                                                                                                                                                                                                                                                                                                               |
| Name                                               | State                                                                                      | Order                                                                                                                                                                                                                                                                                                                                                                                                                               |
| TentnatA-Zone-Network                              |  Disabled |      |
| DefaultSharedNetworkOfferingWithSGService          |  Enabled  |      |
| DefaultSharedNetworkOffering                       |  Enabled  |      |
| DefaultIsolatedNetworkOfferingWithSourceNatService |  Enabled  |      |
| DefaultIsolatedNetworkOffering                     |  Enabled  |      |
| DefaultSharedNetscalerEIPandELBNetworkOffering     |  Enabled  |      |
| DefaultIsolatedNetworkOfferingForVpcNetworks       |  Enabled  |      |
| DefaultIsolatedNetworkOfferingForVpcNetworksNoLB   |  Enabled  |      |


12. Click **Enable** icon.

**Figure 221**      **Enabling the TenantA Zone Network**



13. Click OK.

**Figure 222**      **Confirming the Network Offering Enabling**



# Virtual Machine Configuration and Management

CloudPlatform provides administrators with complete control over the life cycle of all guest VMs executing in the cloud. CloudPlatform provides several guest management operations for end users and administrators. VMs may be stopped, started, rebooted, and destroyed.

This section explains Virtual Machine (VM) life cycle management design process followed to provision multi-tenants virtual machines on shared cloud infrastructure. The VM management features offered by Citrix CloudPlatform includes, live snapshots, checkpoint, migration, backup, and deletion.

The tasks involved in managing the virtual machine life cycle are listed below:

- VM Infrastructure
- Installing the Operating System on VM
- Configuring VM
- VM Backup and Restore
- VM Migration
- VM Deletion

## VM Infrastructure

This section explains how to create compute, network and storage infrastructure for virtual machines which are dedicated for multi-tenant s based on the service levels cost defined. The initial cloud infrastructure designs for compute, network and storage cloud requirements are defined in the previous sections.

In this study we will add Windows 2008 R2 ISO image and create virtual machine with Platinum level compute, network and storage cloud infrastructure using CloudPlatform.

Login to CloudPlatform with User credentials to create ISO Image TenantA Zone

## Adding HTTP where Windows 2008 R2 ISO is been placed for CloudPlatform to download

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click **Global Setting** tab under search type secstorage.allowed.internal.sites.
4. Under secstorage.allowed.internal.sites enter <<Var\_HTTP\_Server>> 10.65.121.70.
5. Restart CloudPlatform service.

## Adding ISO Image

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Templates** tab.
4. Select ISO in Select view list box.

**Figure 223**      *Displaying the ISO Templates*

| Name              | Zone         | Order     | Quickview |
|-------------------|--------------|-----------|-----------|
| Linux-VMwareTools | TenantA-Zone | ▲ ▼ ▲ ▼ ≡ | +         |
| vmware-tools.iso  |              | ▲ ▼ ▲ ▼ ≡ | +         |
| xs-tools.iso      |              | ▲ ▼ ▲ ▼ ≡ | +         |

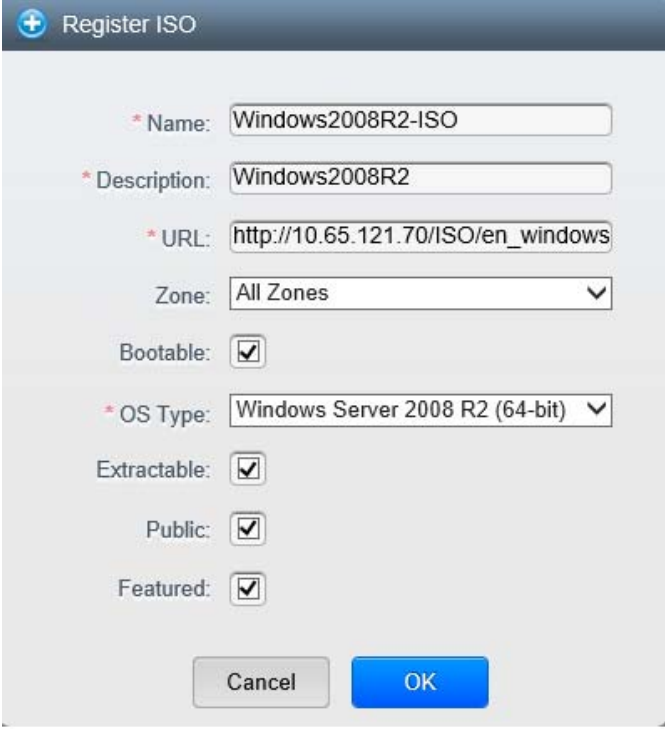
5. Click **Register ISO**.
6. Enter Windows2008R2-ISO in the Name field.
7. Enter Windows2008R2 in the Description field.
8. Enter  
[http://10.65.121.70/ISO/en\\_windows\\_server\\_2008\\_r2\\_standard\\_enterprise\\_datacenter\\_and\\_web\\_with\\_sp1\\_x64\\_dvd\\_617601.iso](http://10.65.121.70/ISO/en_windows_server_2008_r2_standard_enterprise_datacenter_and_web_with_sp1_x64_dvd_617601.iso) in URL field.



**Note** Ensure to setup http server and stored Windows 2008 R2 ISO image under /etc/www/html path.

9. Select TenantA-Zone under Zone list box.
10. Check **Bootable** check box.
11. Select Windows 2008 R2 (64-bit) under OS Type list box.
12. Select **Extractable** check box.
13. Select **Public** check box.
14. Select **Featured** check box.
15. Click **OK**.

**Figure 224**      *Registering the ISO Properties*

A screenshot of a 'Register ISO' dialog box. The dialog has a title bar with a blue plus icon and the text 'Register ISO'. It contains several fields and checkboxes. The fields are: 'Name' with the value 'Windows2008R2-ISO', 'Description' with 'Windows2008R2', 'URL' with 'http://10.65.121.70/ISO/en\_windows', and 'Zone' with a dropdown menu showing 'All Zones'. There are four checkboxes: 'Bootable' (checked), 'OS Type' (dropdown showing 'Windows Server 2008 R2 (64-bit)'), 'Extractable' (checked), and 'Public' (checked). At the bottom, there are two buttons: 'Cancel' and 'OK'.

**Register ISO**

\* Name: Windows2008R2-ISO

\* Description: Windows2008R2

\* URL: http://10.65.121.70/ISO/en\_windows

Zone: All Zones

Bootable: ☒

\* OS Type: Windows Server 2008 R2 (64-bit)

Extractable: ☒

Public: ☒

Featured: ☒

Cancel OK

## Creating VM

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instance** tab.
4. Select **ISO** in Select View list box.
5. Click **Add Instance**.
6. Select TenantA-Zone in Select a Zone list box.
7. Select **ISO** radio button.
8. Click **Next**.



**Figure 225**      **Defining the Name and Image of the Instance**

Add Instance

1 Setup 2 Select a template 3 Compute offering 4 Data Disk Offering 5 Network 6 Review

**Select a zone**  
A zone typically corresponds to a single datacenter. Multiple zones help make the cloud more reliable by providing physical isolation and redundancy.

TenantA-Zone

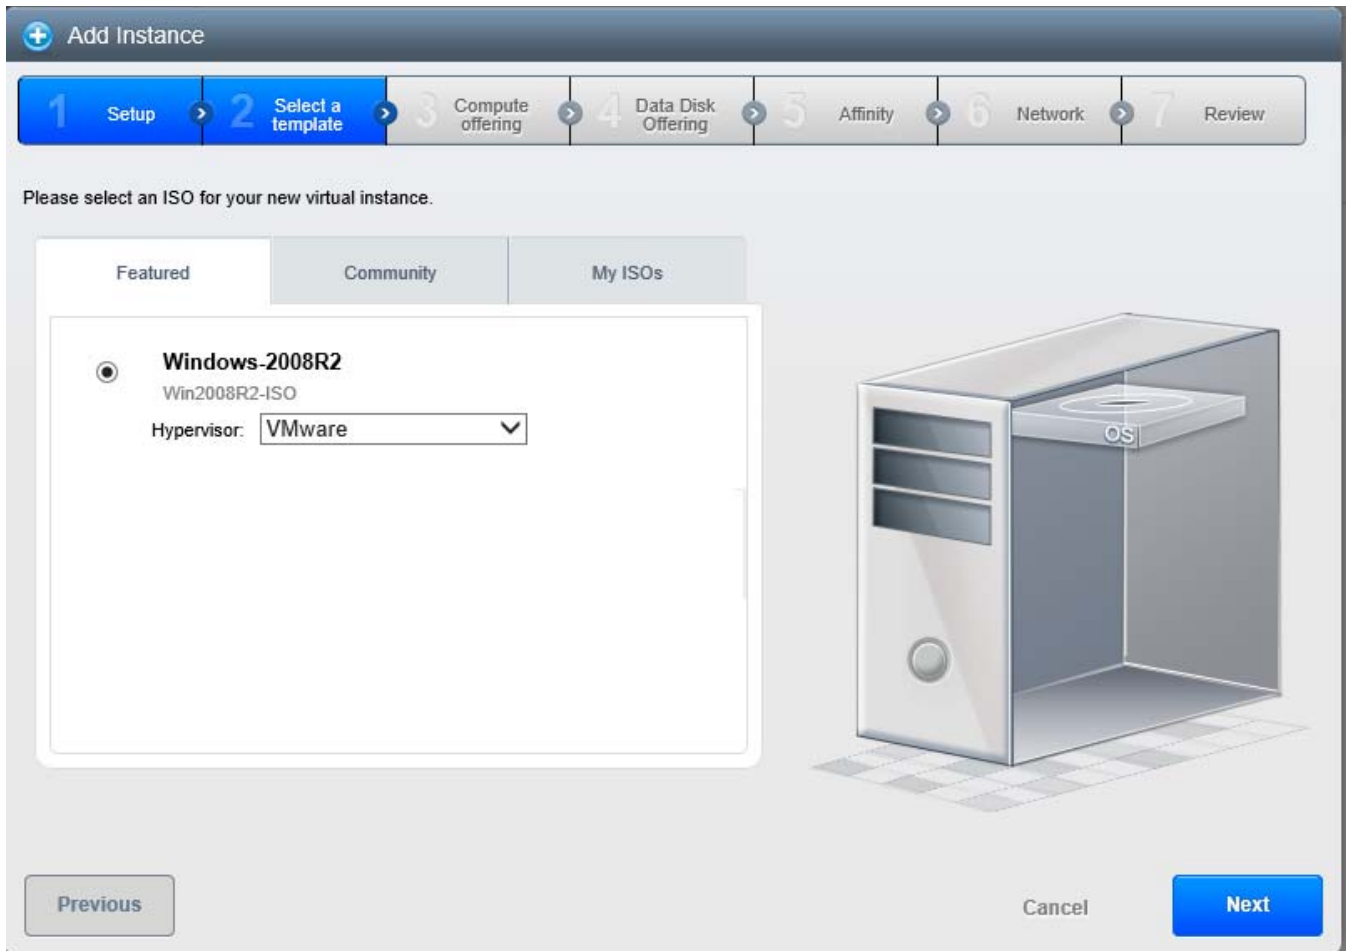
**Select ISO or template**

☒ **Template** OS image that can be used to boot VMs

☐ **ISO** Disc image containing data or bootable media for OS

Cancel Next

9. Click **Windows-2008R2 Image** radio button.
10. Select ESXiServer in Hypervisor list box.
11. Click **Next**.

**Figure 226** *Defining the Hypervisor for the Instance*

12. Select **Platinum-Compute**.
13. Click **Next**.

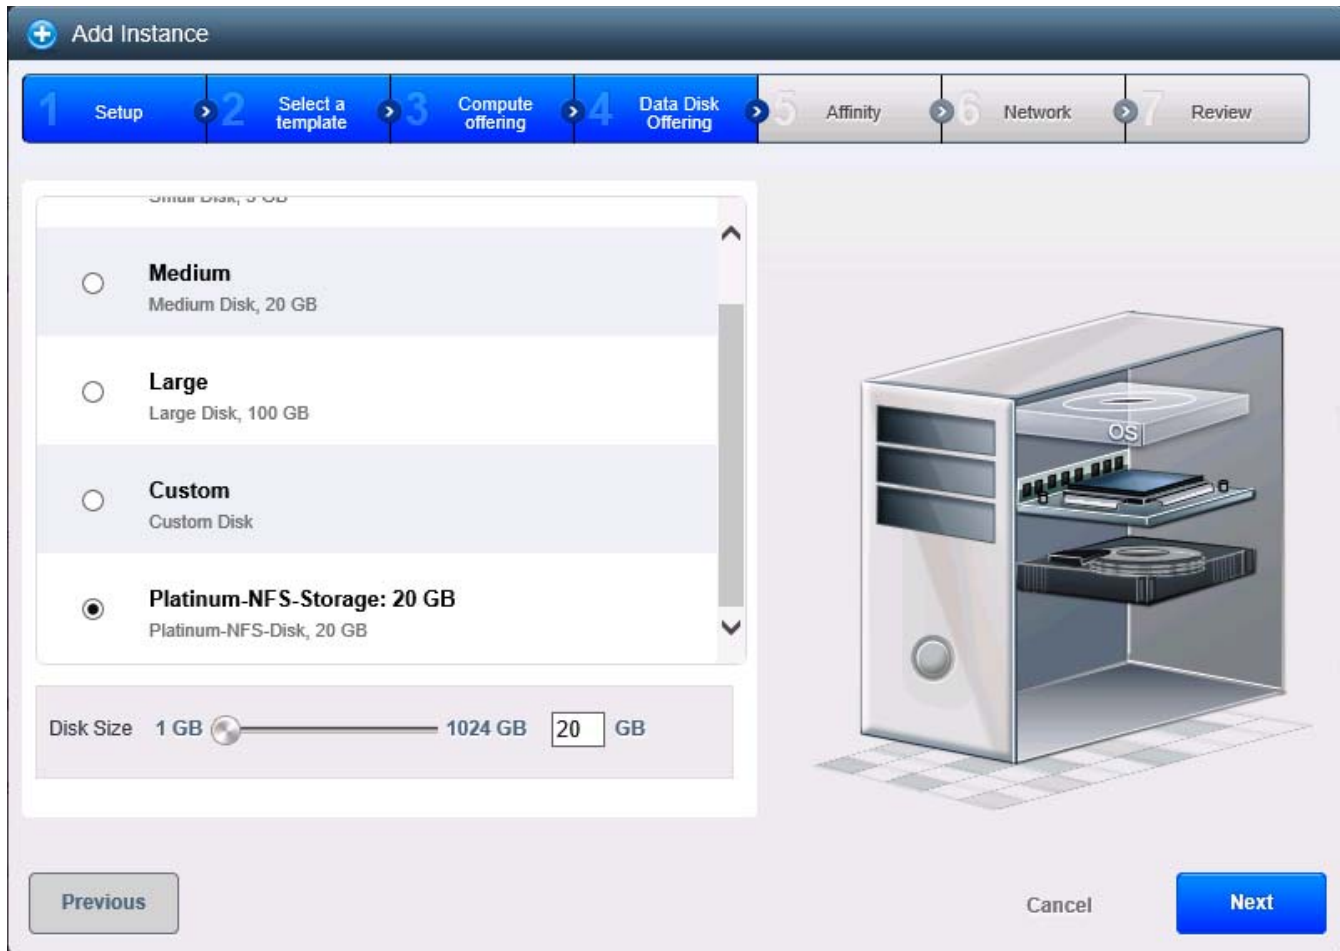
**Figure 227** Defining the Compute Offering for the Instance

The screenshot shows the 'Add Instance' wizard with the following components:

- Header:** 'Add Instance'
- Progress Bar:** 1 Setup, 2 Select a template, 3 Compute offering (active), 4 Data Disk Offering, 5 Network, 6 Review.
- Instance Selection List:**
  - ☐ **Medium Instance**  
Medium Instance
  - ☐ **Small Instance**  
Small Instance
  - ☒ **Platinum-Compute**  
TentatA-Platinum-Compute
- 3D Illustration:** A server rack with a transparent front panel showing internal components like a hard drive and a network card. The label 'OS' is visible on the top drive.
- Navigation Buttons:** 'Previous' (disabled), 'Cancel', and 'Next' (active).

14. Select **Platinum-NFS-Storage-Disk**.
15. Enter 20-GB Value in Disk Size field.
16. Click **Next**.

**Figure 228**      *Selecting the Disk Offering and Size*



17. Check **TenantA-Network** and **Default** check boxes under Network.
18. Click **Next**.

**Figure 229**      **Selecting Networks for the Virtual Machine**

Add Instance

1 Setup   2 Select a template   3 Compute offering   4 Data Disk Offering   5 Network   6 Review

Please select networks for your virtual machine. vpc: None

| Networks                            |                  |                               |
|-------------------------------------|------------------|-------------------------------|
| <input checked="" type="checkbox"/> | TentnatA-Network | Isolated <span>Default</span> |

Add Network

☐ New

Previous   Cancel   Next

19. Enter TenantA-Platinum-VM1 in VM Name field.
20. Click **Launch VM**.

**Figure 230**      *Verifying the Virtual Instance Information*

Add Instance

1 Setup 2 Select a template 3 Compute offering 4 Data Disk Offering 5 Network 6 Review

Please review the following information and confirm that your virtual instance is correct before launch.

|                         |                       |
|-------------------------|-----------------------|
| Name (Optional)         | TentnatA-Platinum-VM1 |
| Add to group (Optional) |                       |
| Zone                    | <a href="#">Edit</a>  |
| Hypervisor              | <a href="#">Edit</a>  |
| Template                | <a href="#">Edit</a>  |
| Compute offering        | <a href="#">Edit</a>  |
| Data Disk Offering      | <a href="#">Edit</a>  |
| Network                 | <a href="#">Edit</a>  |

Previous

Cancel [Launch VM](#)

## Installing the Operating System

This section explains how to install guest operating system on virtual machines after provisioning compute, network and storage infrastructure which are dedicated for multi-tenant s to host data.

In this study we will install RHEL 6.2 guest operating system on newly created TenantA-Platinum-VM1 virtual machine with Platinum level compute, network and storage cloud infrastructure using CloudPlatform.

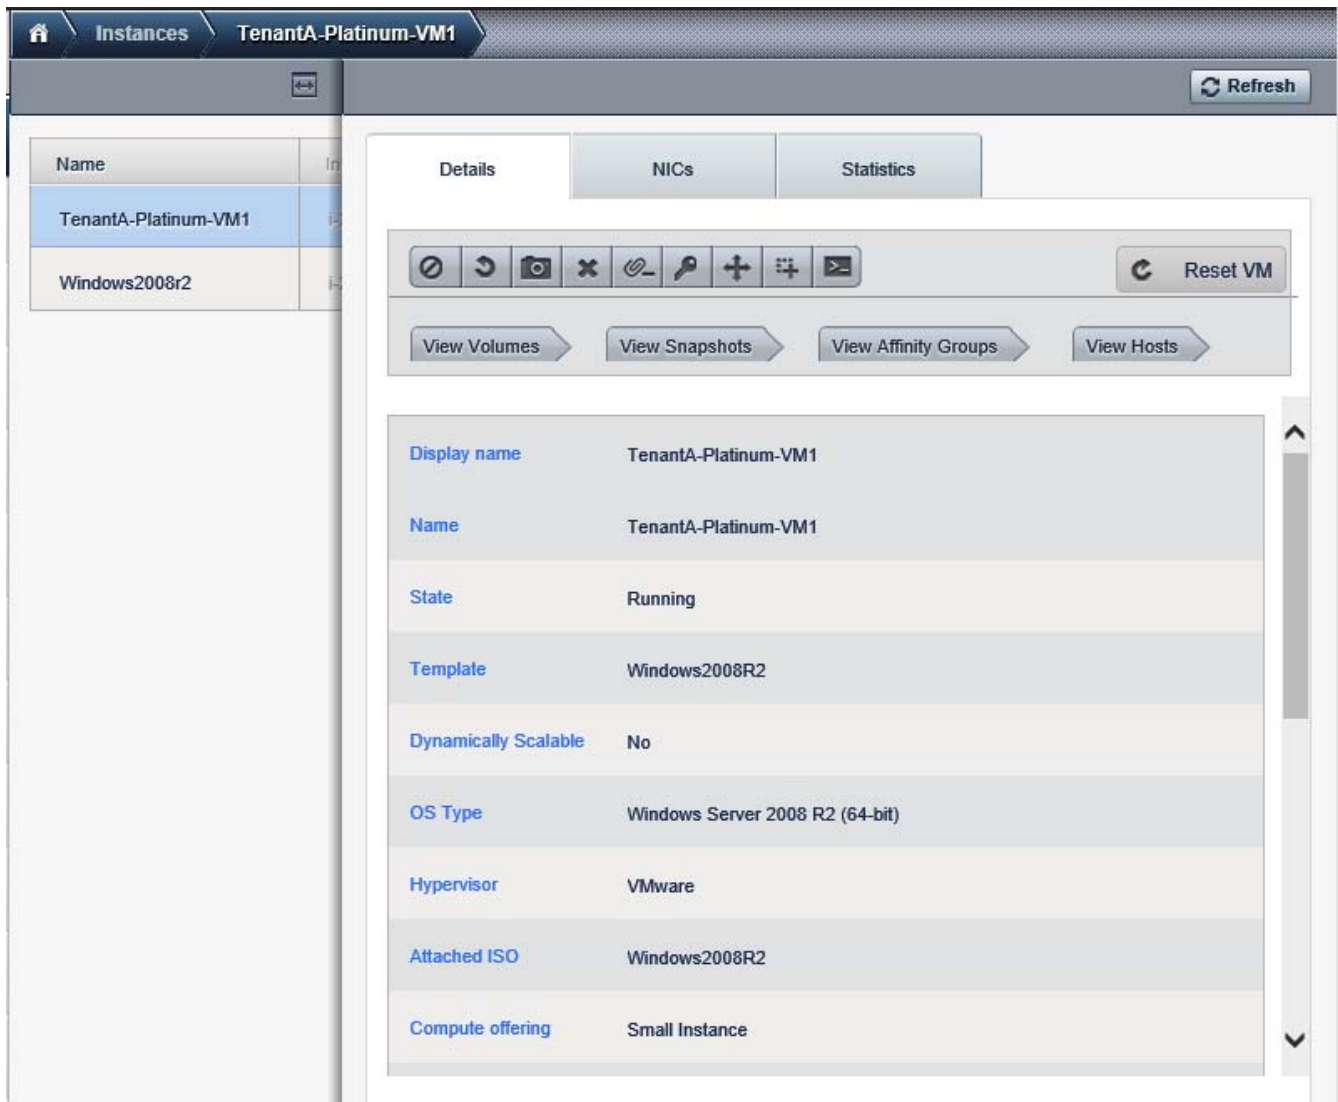
Login to CloudPlatform with User credentials to create ISO Image TenantA Zone:

## Installing Windows OS

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instances** tab.
4. Click **TenantA-Platinum-VM1** under Display name.

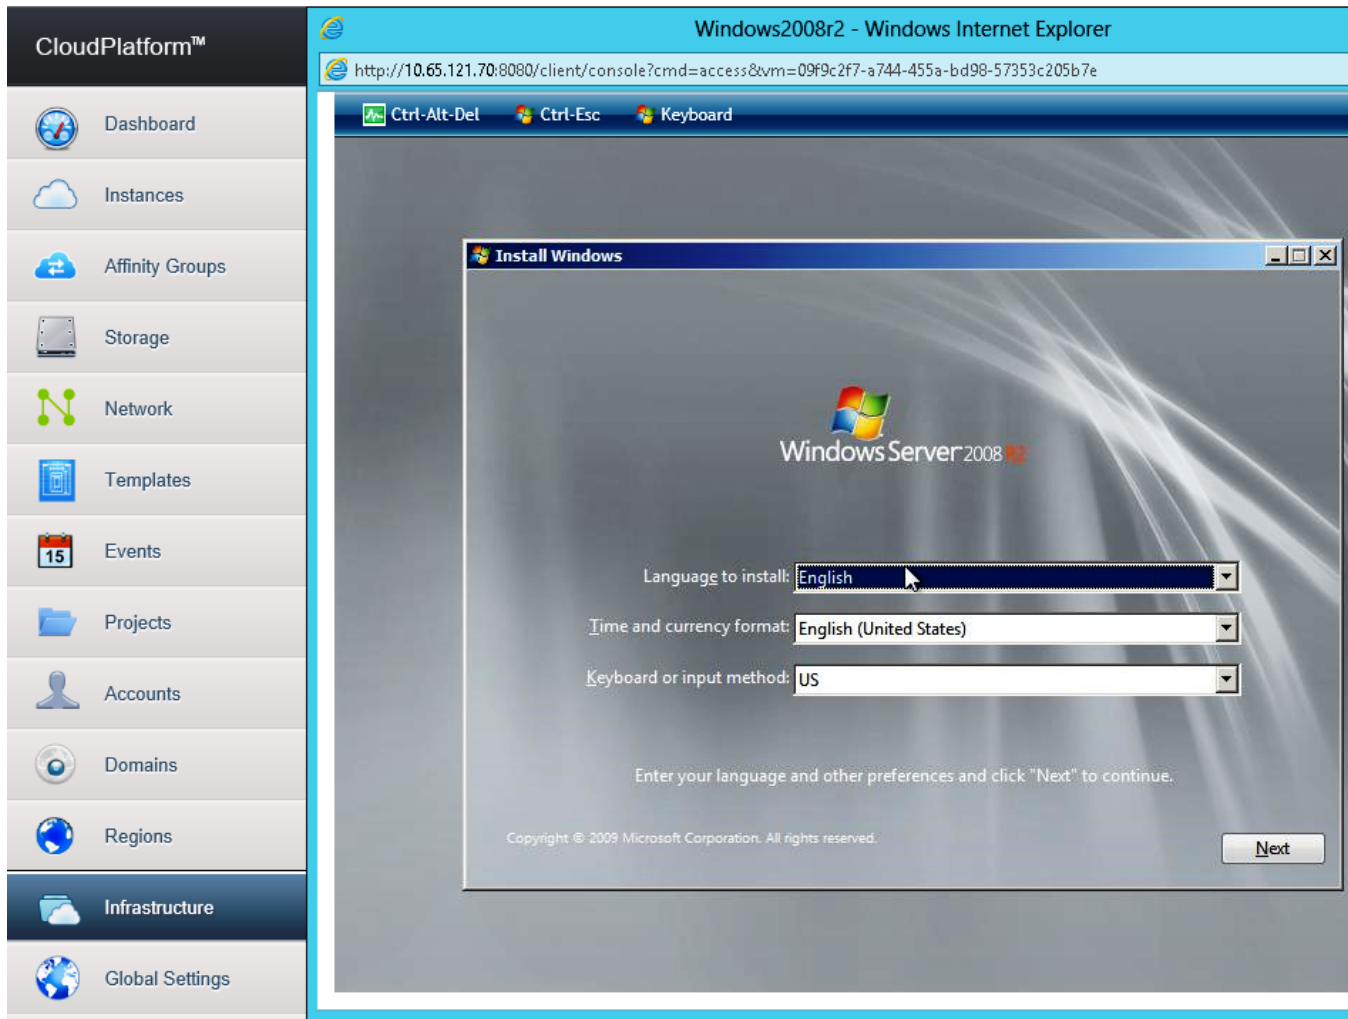
5. Click **View Console** icon.

**Figure 231**      *Displaying the VM Details*



6. Click Console button to launch for installation of Windows OS.

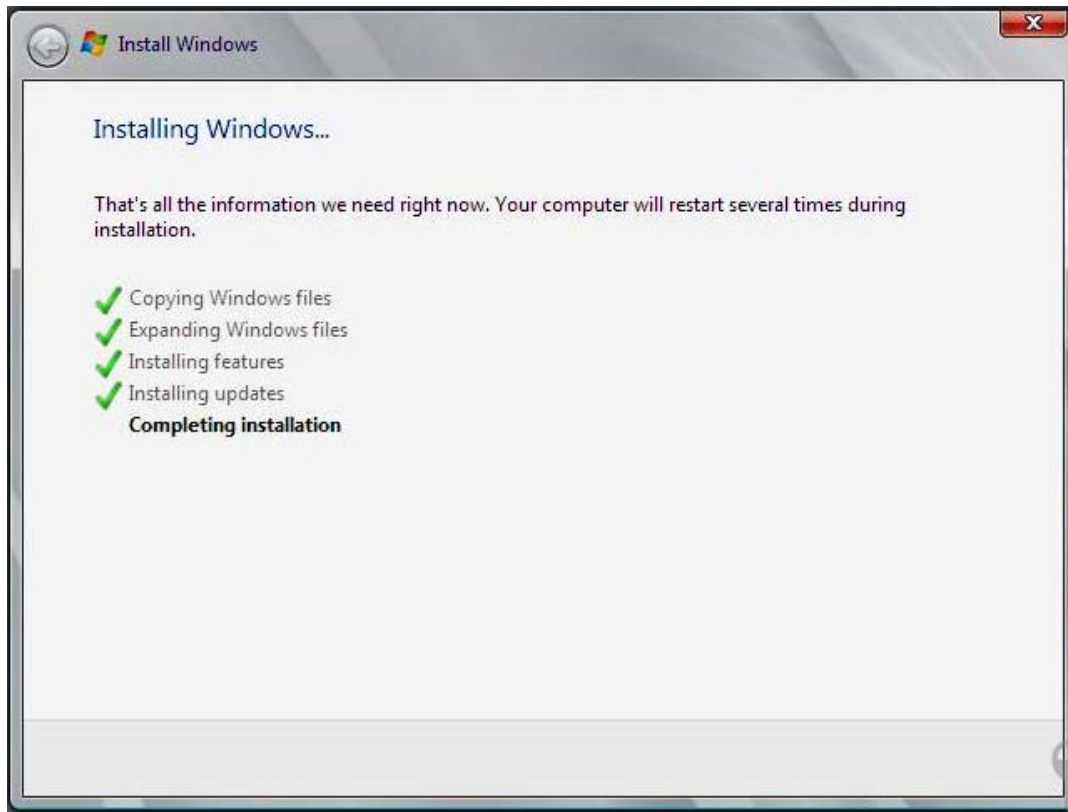
Figure 232 Launching the VM



7. Click **English** in Language to install, Choose **English (United States)** and **US** in keyboard or input method.
8. Click **Next** button.
9. Click **Install Now** button.
10. Select **Windows Server 2008 R2 Enterprise (Full Installation)** OS type.
11. Click **Next** button.
12. **Accept license** check box
13. Click **Next** button
14. Click **Custom Installation** method.
15. Click **Next** button.
16. Select **Disk0**.
17. Click **Next** for OS installation.




**Figure 233**      *Windows 2008 R2 OS installation completion into the VM*



**18.** Enter **Password** <xxxxxx> and **Confirm Password** <xxxxxx>

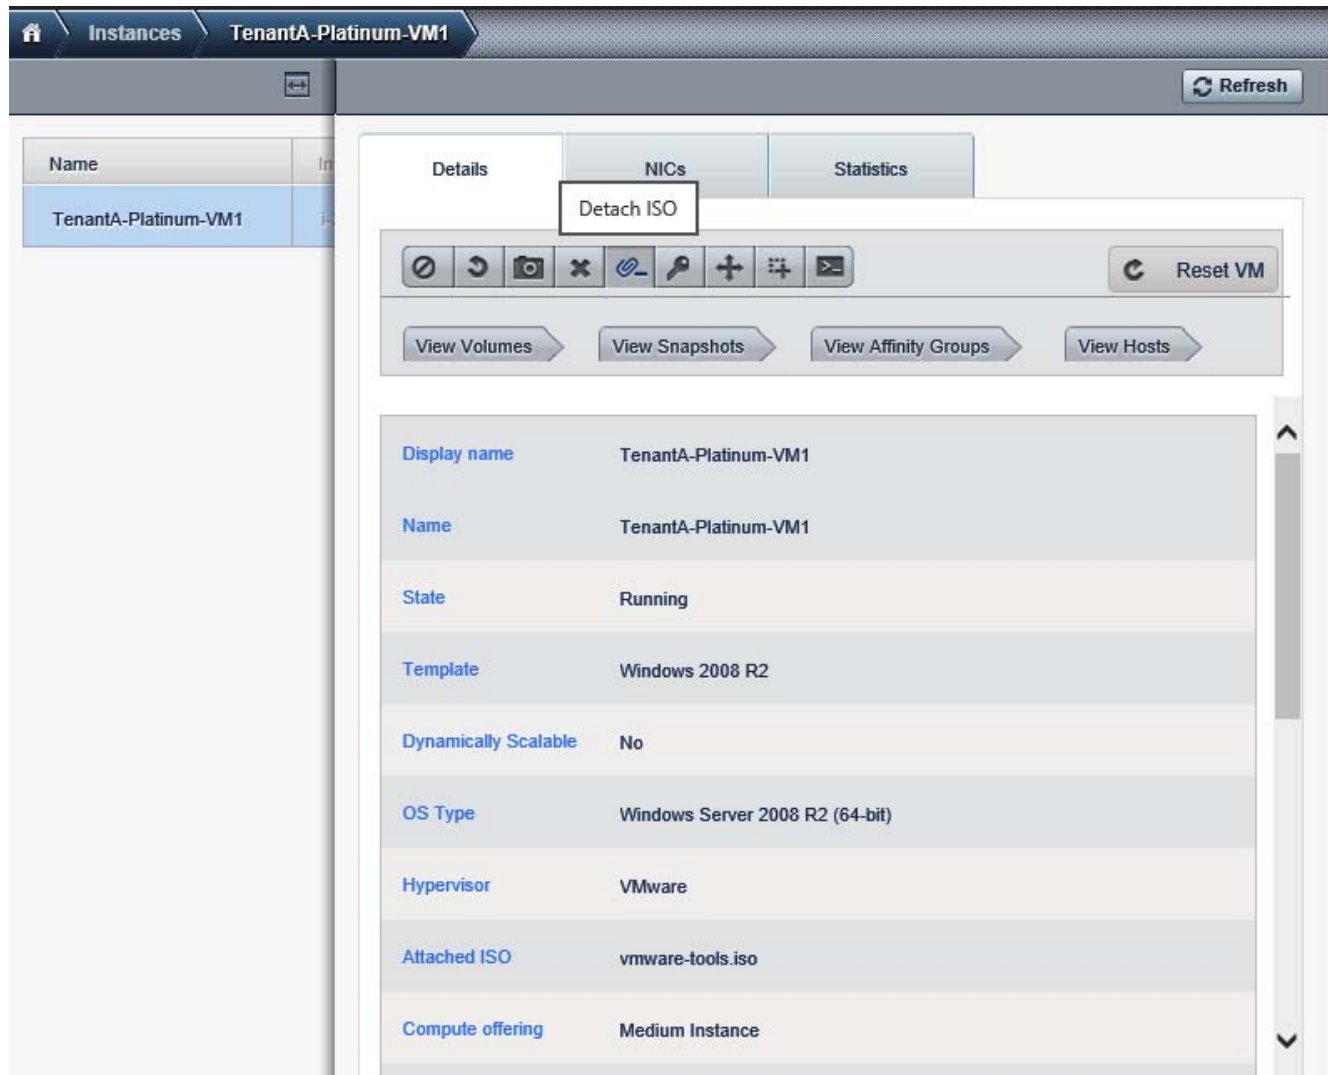
**Figure 234**      *Logging into the VM*



## Detaching ISO on TenantA-Platinum-VM1

1. Click the **Instance** tab.
2. Click **Detach ISO** icon.

**Figure 235**      *Detaching the ISO Image*



3. Click **Yes** in the confirmation window.

## VM Configuration

This section explains Virtual Machine cloud network address translation (NAT) for accessing the public network and provisioning external data disk with ESXi PV Drivers for storing tenants hosted data information.

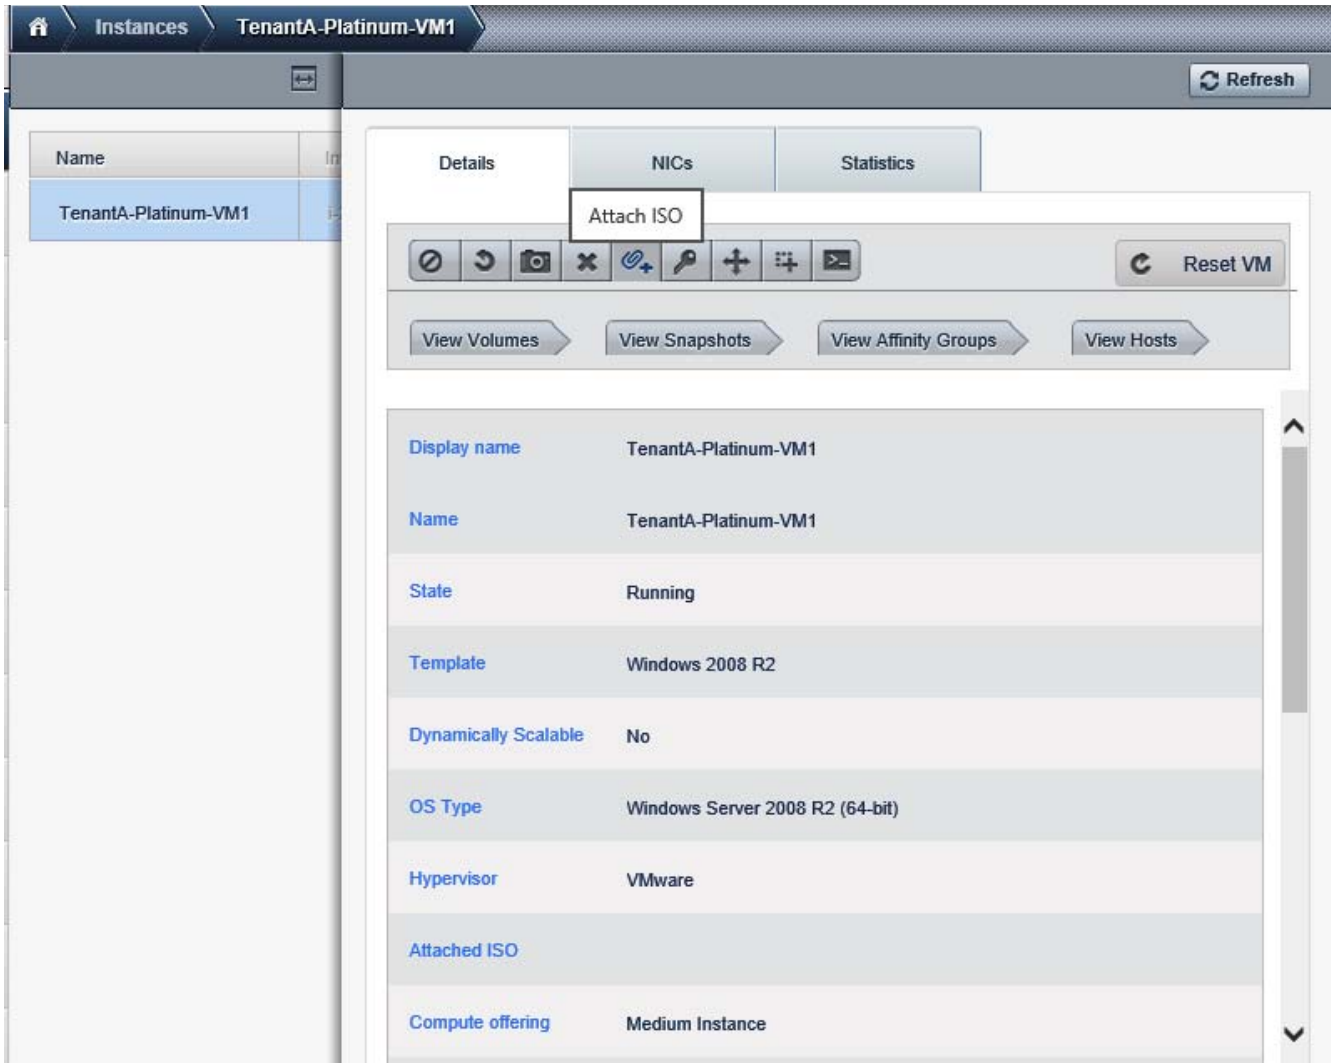
In this study we will install ESXi PV Driver, provision data disk and attach the newly created virtual machine TenantA-Platinum-VM1, and configure source NAT for accessing Virtual Machine from public access using CloudPlatform.

Login to CloudPlatform with User credentials to install ESXi PV Driver on Virtual Machine on TenantA Zone

## ESXi PV Driver Installation

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instances** tab.
4. Click **Attach ISO** icon.

**Figure 236**      *Attaching the ISO*



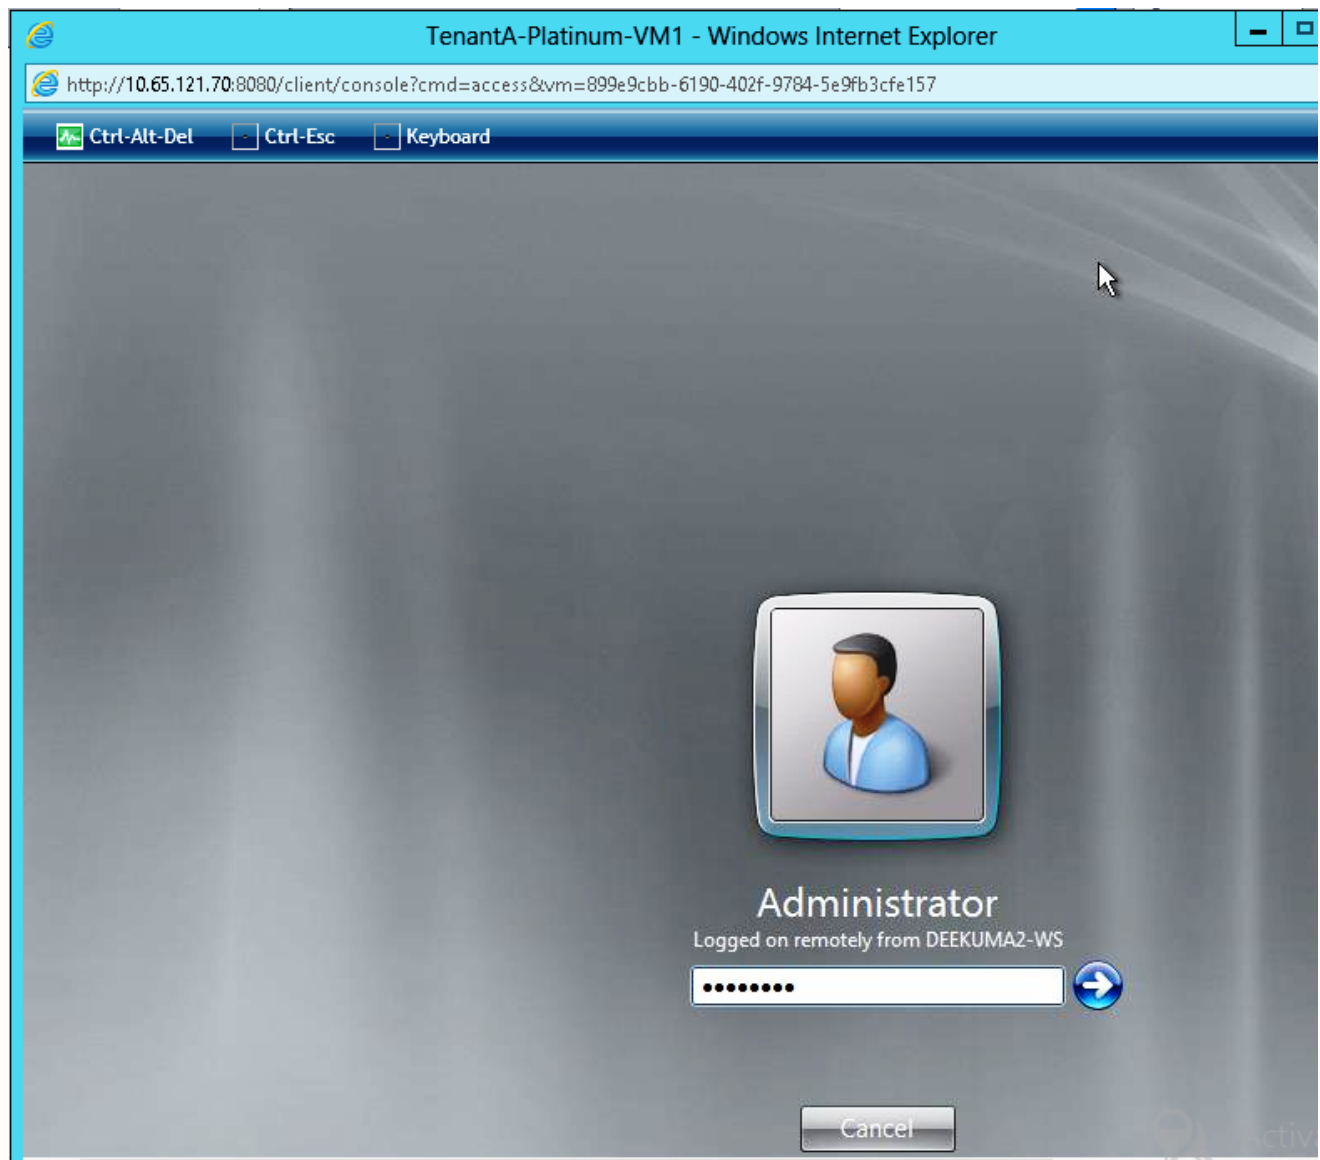
5. Select **VMware Tools Installer ISO** from ISO list box.
6. Click **OK**.

**Figure 237**      **Selecting the ISO**



7. Launch VM Console and login to <TenantA-Platinum-VM1>.
8. Enter User Name <Root> and password <XXXXXX>.

**Figure 238**      **Logging into the VM**



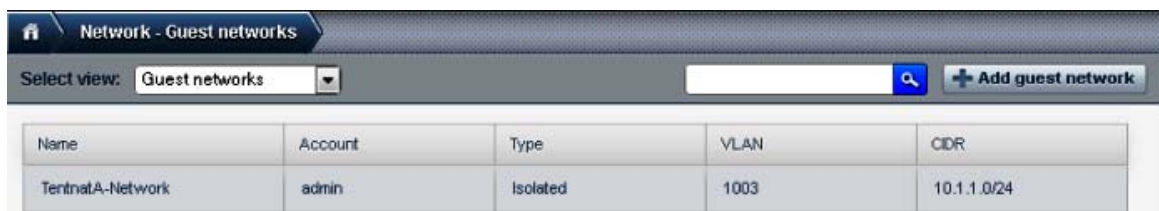
9. Click **Run setup64.exe installer**.
10. Click **Next**.
11. Click **Complete** radio button.
12. Click **Install**.
13. Click **Finish**.
14. Click **Yes** to reboot.

Login to CloudPlatform with User credentials to configure VM NAT on TenantA Zone:

## Sourcing the NAT to access VM from Public Access

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Network** tab.

**Figure 239**      *Displaying the Guest Network*



| Name              | Account | Type     | VLAN | CIDR        |
|-------------------|---------|----------|------|-------------|
| TenatnatA-Network | admin   | Isolated | 1003 | 10.1.1.0/24 |

4. Click **TenantA-Network** tab.
5. Click **Egress rules**.
6. Enter 0.0.0.0/0 in Source CIDR and select **All** in Protocol list box.
7. Click **ADD** button.

**Figure 240**      *Displaying the Guest Network Egress Rules*



| Source CIDR | Protocol | Add |
|-------------|----------|-----|
| 0.0.0.0/0   | All      | Add |

8. Click **TenantA-Network** tab.
9. Click **View IP Address**

**Figure 241**      *Displaying the Network Details*

The screenshot shows the 'Details' tab for 'TentnatA-Network'. At the top, there are navigation tabs: 'Network - Guest networks' and 'TentnatA-Network'. A 'Refresh' button is in the top right. Below the tabs, there are three icons (edit, refresh, delete) and a 'View IP Addresses' button. The main content is a table of network details:

|                  |                                                                |
|------------------|----------------------------------------------------------------|
| Name             | TentnatA-Network                                               |
| ID               | b7ea6861-1a64-4422-b8ab-0d3c93073454                           |
| Zone             | TenantA-Zone                                                   |
| Description      | TentnatA-Network                                               |
| Type             | Isolated                                                       |
| State            | Implemented                                                    |
| Restart required | No                                                             |
| VLAN ID          | 1003                                                           |
| Network Offering | Offering for isolated networks with Source Nat service enabled |

10. Click **+Acquire New IP** button.

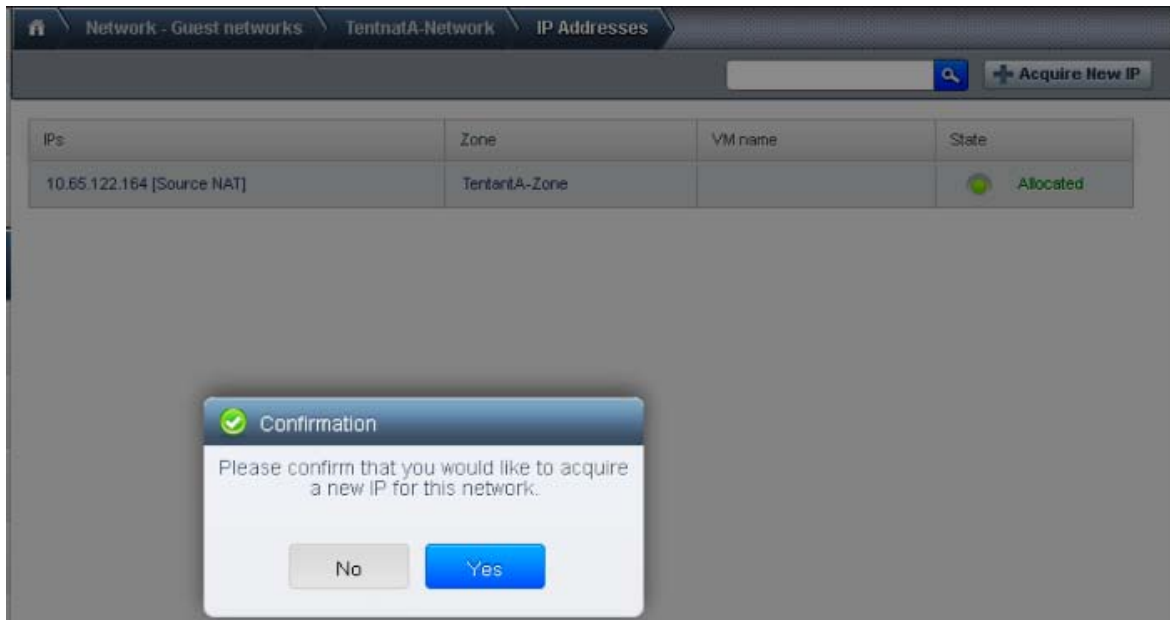
**Figure 242**      *Displaying the IP Address for the Guest Network*

The screenshot shows the 'IP Addresses' tab for 'TentnatA-Network'. At the top, there are navigation tabs: 'Network - Guest networks', 'TentnatA-Network', and 'IP Addresses'. A search bar and a '+ Acquire New IP' button are in the top right. Below the tabs is a table of IP addresses:

| IPs                        | Zone         | VM name | State     |
|----------------------------|--------------|---------|-----------|
| 10.65.122.164 [Source NAT] | TenantA-Zone |         | Allocated |

11. Click **Yes** button.

**Figure 243**      *Confirming the IP Address*



12. IP Address <10.65.122.166> Public IP Address is allocated.

**Figure 244**      *Displaying the IP Address*

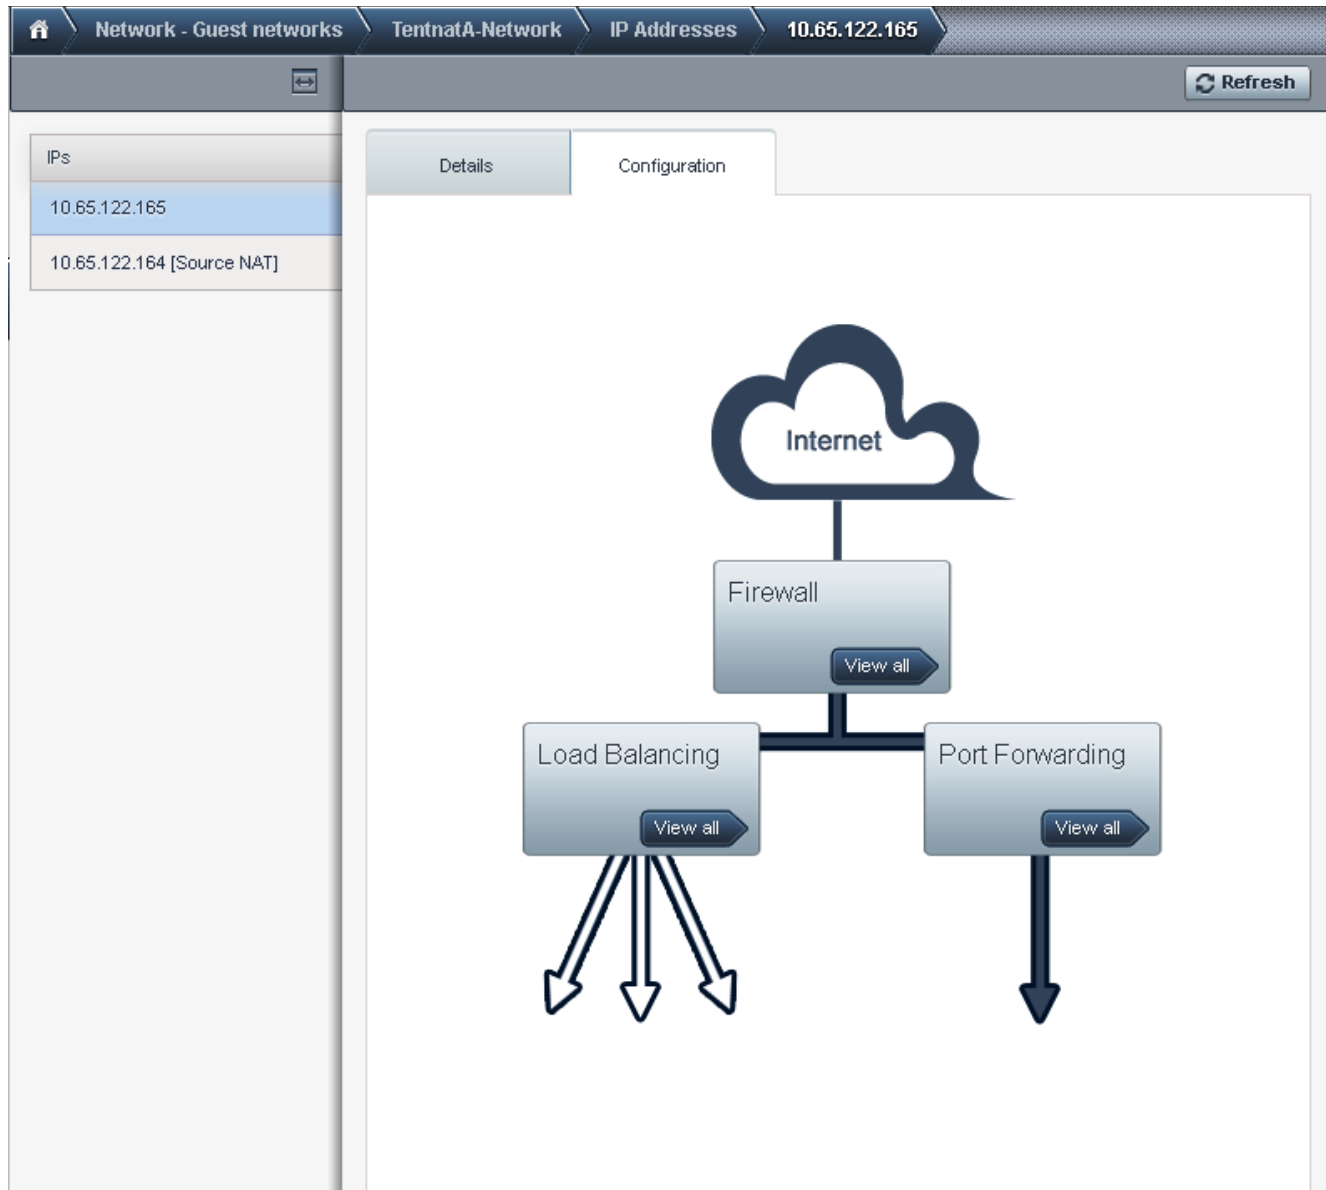
The screenshot shows the 'IP Addresses' tab in the 'TenantA-Network' section. A table displays the allocated IP addresses. The table has five columns: 'IPs', 'Zone', 'VM name', 'State', and 'Quickview'.

| IPs                        | Zone         | VM name | State     | Quickview |
|----------------------------|--------------|---------|-----------|-----------|
| 10.65.122.166              | TenantA-Zone |         | Allocated | +         |
| 10.65.122.165 [Source NAT] | TenantA-Zone |         | Allocated | +         |

13. Click **IP Address** <10.65.122.166>.
14. Click the **Configuration** tab.



**Figure 245**      *Displaying the Configuration Flow diagram*



15. Click **Firewall** tab.
16. Enter <0.0.0.0/0> in Source CIDR field.
17. Select TCP Protocol List Value.
18. Enter 3389 in (RDP) in Start Port field.
19. Enter 3389 in (RDP) End Port field.
20. Click **Add** button.

**Figure 246** *Defining the Firewall Configurations*

| Source CIDR | Protocol | Start Port | End Port | ICMP Type | ICMP Code | Add rule | Actions |
|-------------|----------|------------|----------|-----------|-----------|----------|---------|
| 0.0.0.0/0   | TCP      | 3389       | 3389     |           |           | Add      |         |

21. Click **Port Forwarding** in the flow diagram.

**Figure 247** *Selecting the Port Forwarding Option*

| Private Port | Public Port | Protocol | Add VM | Actions |
|--------------|-------------|----------|--------|---------|
| 3389 3389    | 3389 3389   | TCP      | Add    |         |

22. Enter 3389 in Private Port field.
23. Enter 3389 in Public Port field.
24. Select TCP under Protocol list box.

**Figure 248** *Defining the Port Forwarding Configurations*

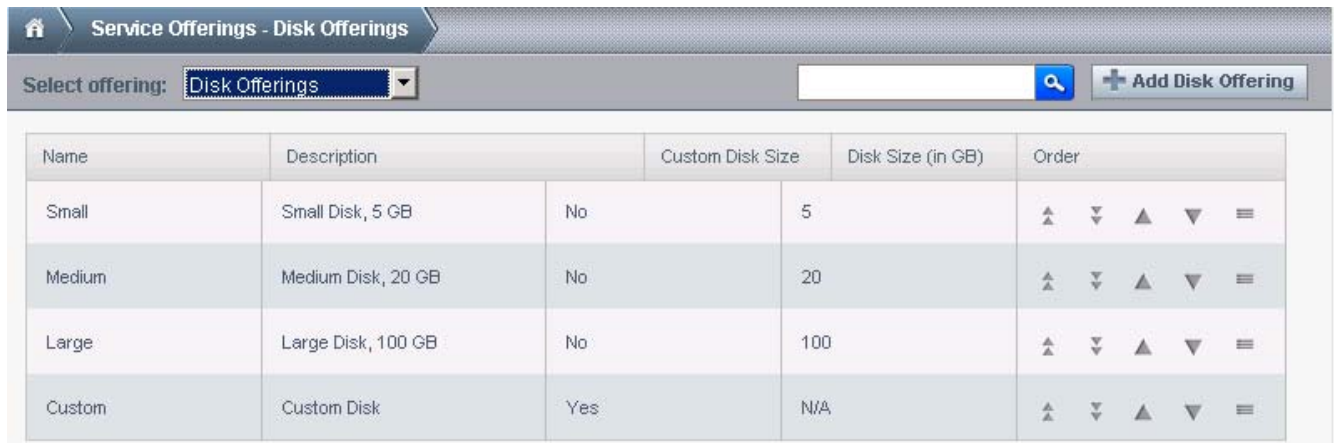
25. Click **Add**.
26. Click **Select** radio button on TenantA-Platinum-VM1.
27. Click **Apply**.

**Figure 249 Adding VMs**

To configure VM Data Disk on TenantA Zone, login to the CloudPlatform:

#### Attaching Data Disk to TenantA-Platinum-VM1

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Service Offerings** tab.
4. Select Disk Offering in Select offering list box.

**Figure 250 Displaying the Disk Offerings**

5. Enter Platinum-NFS-Storage in the Name field.
6. Enter Platinum-NFS-Disk in the Description field.
7. Select shared in Storage Type list box.
8. Check **Custom Tags** check box.
9. On QoS Type List do not select.
10. Enter Platinum-NFS-Storage in Storage Tags.
11. Select **Public** check box.
12. Click **OK**.

**Figure 251**      *Defining the Disk Offerings Information*


The 'Add Disk Offering' dialog box contains the following fields and controls:

- Name:** Platinum-NFS-Storage
- Description:** Platinum-NFS-Disk
- Storage Type:** shared (dropdown menu)
- Custom Disk Size:** ☒
- QoS Type:** (empty dropdown menu)
- Storage Tags:** Platinum-NFS-Storage
- Public:** ☒
- Buttons:** Cancel, OK

13. Click the **Storage** tab.
14. Select Volumes in Select view list box.

**Figure 252**      *Displaying the Storage Volume*

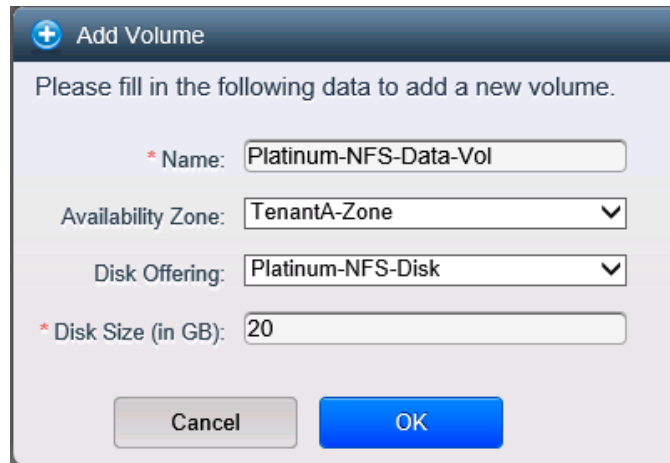

The 'Storage - Volumes' section shows a table with the following data:

| Name    | Type | Hypervisor | VM display name        |
|---------|------|------------|------------------------|
| ROOT-11 | ROOT | XenServer  | TenitantA-Platinum-VM1 |

At the top of the section, there is a 'Select view:' dropdown menu set to 'Volumes', a search bar, and two buttons: 'Upload volume' and '+ Add Volume'.

15. Click **Add Volume** button.
16. Select Volumes in Select view list box.
17. Enter Platinum-NFS-Data-Vol.
18. Select <TenantA-Zone> in Availability Zone list box.
19. Select Platinum-NFS-Disk in Disk Offering list box.
20. Enter 20 in Disk Size (in GB) field.

**Figure 253**      **Defining the Volume Information**



The image shows a dialog box titled "Add Volume" with a plus icon in the top-left corner. The text inside says "Please fill in the following data to add a new volume." There are four input fields: "Name" with a red asterisk, "Availability Zone", "Disk Offering", and "Disk Size (in GB)" with a red asterisk. The values entered are "Platinum-NFS-Data-Vol", "TenantA-Zone", "Platinum-NFS-Disk", and "20" respectively. At the bottom are "Cancel" and "OK" buttons.

| Field               | Value                 |
|---------------------|-----------------------|
| * Name              | Platinum-NFS-Data-Vol |
| Availability Zone   | TenantA-Zone          |
| Disk Offering       | Platinum-NFS-Disk     |
| * Disk Size (in GB) | 20                    |

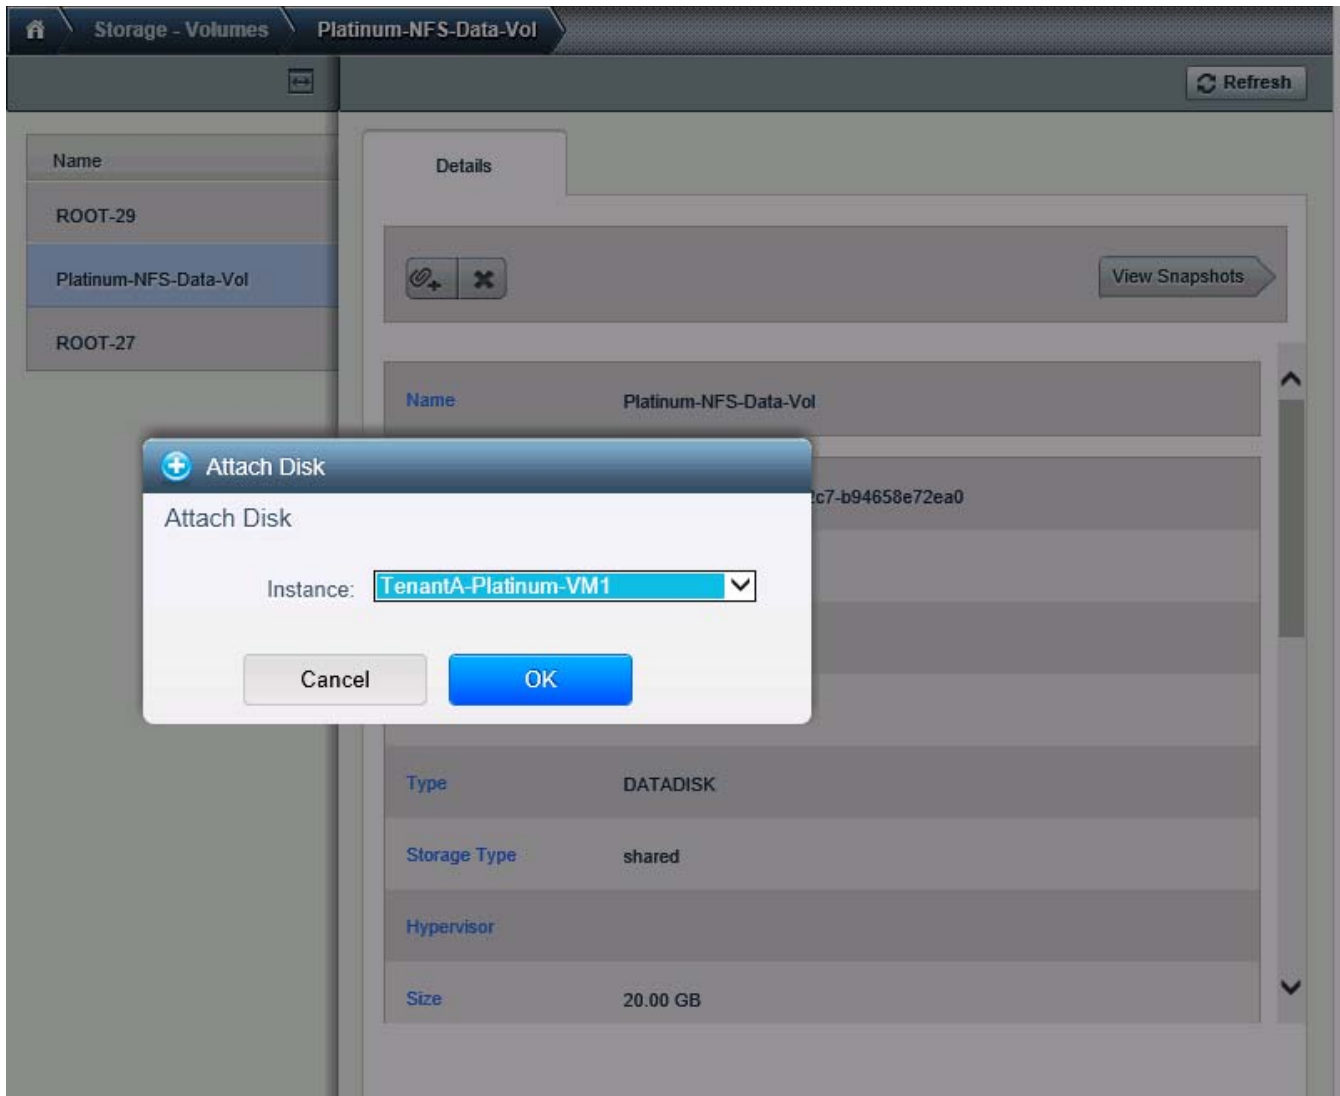
21. Click <Platinum-NFS-Data-VOL>.
22. Click **Attach Disk** icon.

Figure 254 Displaying Data Disk Details

The screenshot displays a web-based management interface for storage volumes. The left sidebar shows a list of volumes: 'Name', 'ROOT-29', 'Platinum-NFS-Data-Vol' (selected), and 'ROOT-27'. The main panel is titled 'Platinum-NFS-Data-Vol' and includes a 'Refresh' button. Below the title, there is a 'Details' tab and an 'Attach Disk' button. A 'View Snapshots' button is also present. The details section contains a table with the following information:

|              |                                      |
|--------------|--------------------------------------|
| Name         | Platinum-NFS-Data-Vol                |
| ID           | 7d558754-041f-4f98-a2c7-b94658e72ea0 |
| Zone         | TenantA-Zone                         |
| State        | Allocated                            |
| Status       |                                      |
| Type         | DATADISK                             |
| Storage Type | shared                               |
| Hypervisor   |                                      |
| Size         | 20.00 GB                             |

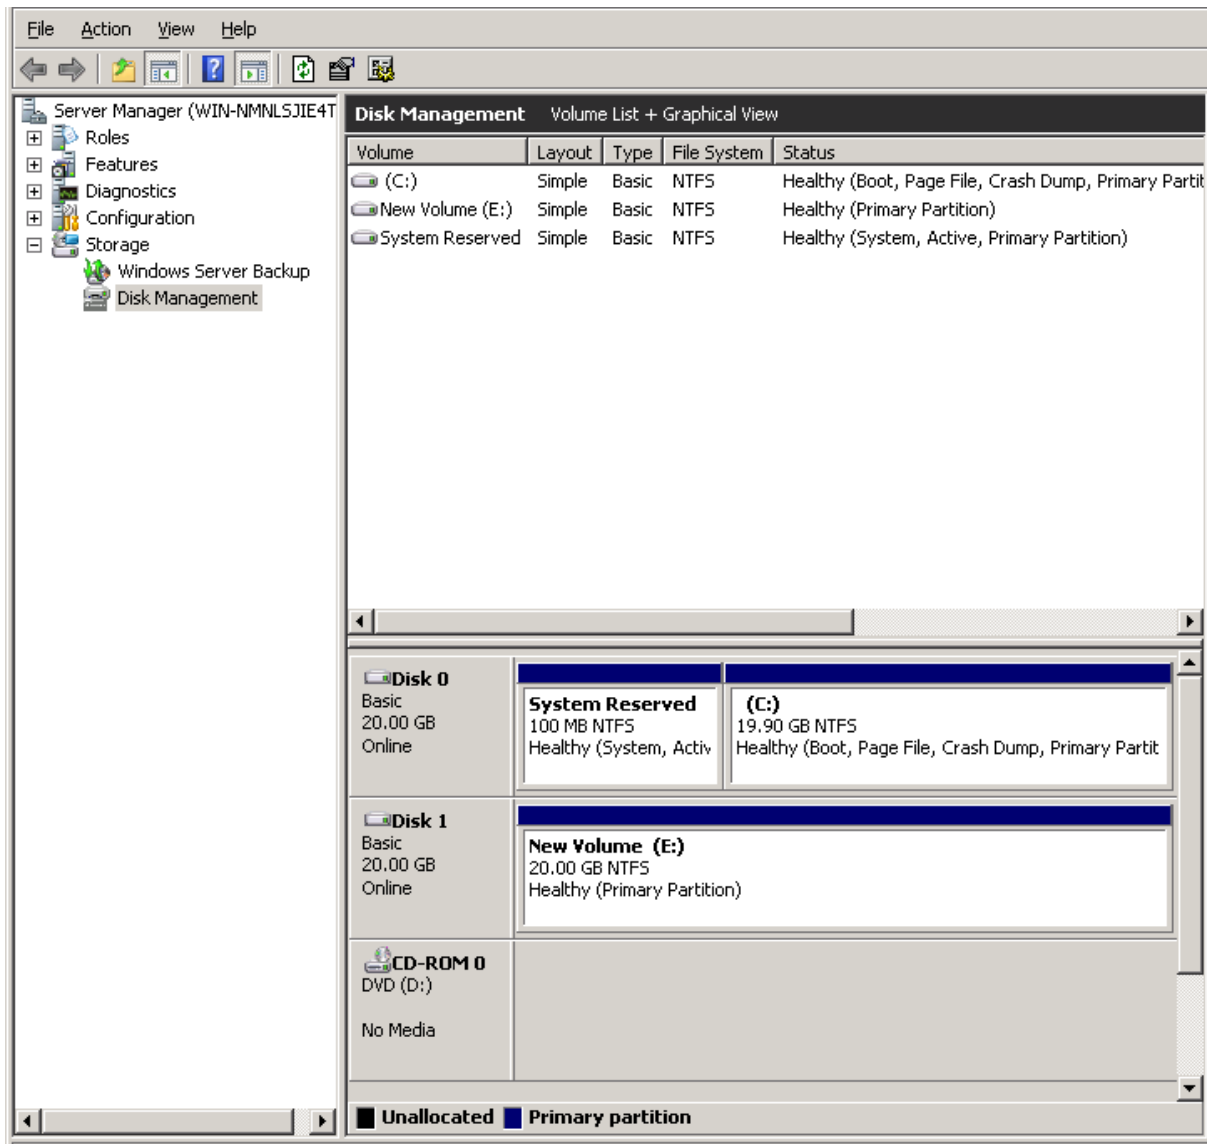
23. Select <TenantA-Platinum-VM1> in Instance list box.
24. Click **OK**.

**Figure 255**      **Attaching Instance to Disk**

25. Remotely connect to TenantA-Platinum-VM1 <10.65.122.166> (IP Acquired).
26. Login with <administrator> and Password <xxxxxxxxxx>.
27. Open **Server Manager**.
28. Click **Storage Under Disk Management**.
29. Right click **Disk 1 Initilize and Online Disk**.
30. Select **Create simple volume**.
31. Click **Next**.
32. Select Assign drive Letter (E).
33. Click **Next**.
34. Select default options.
35. Click **Next**.

36. Click **Finish**.

**Figure 256** Windows VM displaying Attached Disk info



## VM Backup and Restore

This section explains how to configure virtual machine backup and restore policy based on the service levels offering in cloud. The CloudPlatform performs backup of VM OS and Data disks by taking snapshots of disk volumes. Snapshots are a point-in-time capture of virtual machine disks where OS and Data information is stored.

- Backup Methods



There are two methods to take a VM backup, the first method is manual where cloud administrator has to manually take a backup of the disk volumes, and the second method is automatic wherein the recurring snapshots policies can be configured allowing schedule backup time, based on Hourly, Daily, Weekly and Monthly basis with retention numbers. More than one backup policy can be defined according to the requirement.

- **Restore Methods**

There are two ways to perform snapshot restore, the first method is by creating volume of the snapshot, attaching the snapshot in the form of disk to instance VM. The second method is by creating a template and access the backup VM.


**Note**

It is recommended to create snapshot restore of Data disk using first method and root disk (Operating System) using second method.

In this study we will create manual backup of TenantA-Platinum-VM1 root and automatic backup of Platinum-NFS-Data-Vol data disk and restore root disk with template and Data disk with volume method using CloudPlatform.

Login to CloudPlatform with User credentials to create TenantA-Platinum-VM1 manual and automatic backup on TenantA Zone

## Manual Backup

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Storage** tab.

Figure 257 Displaying the Storage Volume Data

The screenshot shows a web interface for managing storage volumes. The left sidebar has a 'Name' section with a list containing 'Platinum-NFS-Data-Vol' and 'ROOT-27'. The 'ROOT-27' item is selected. The main area is titled 'Details' and contains a table of properties for the selected volume.

| Details      |                                      |
|--------------|--------------------------------------|
| Name         | ROOT-27                              |
| ID           | 2e86c19a-4112-43d6-8562-7d685edac131 |
| Zone         | TenantA-Zone                         |
| State        | Ready                                |
| Status       |                                      |
| Type         | ROOT                                 |
| Storage Type | shared                               |
| Hypervisor   | VMware                               |
| Size         | 20.00 GB                             |

4. Click **Root-27** under Name in right pane.
5. (Manual backup of TenantA-Platinum-VM1 root disk (Operating System))
6. Click **Take Snapshot** icon with single camera.

**Note**

Ensure you Virtual Machines are powered down to take a Consistent Snapshot copy on ESXi Hypervisor Host.

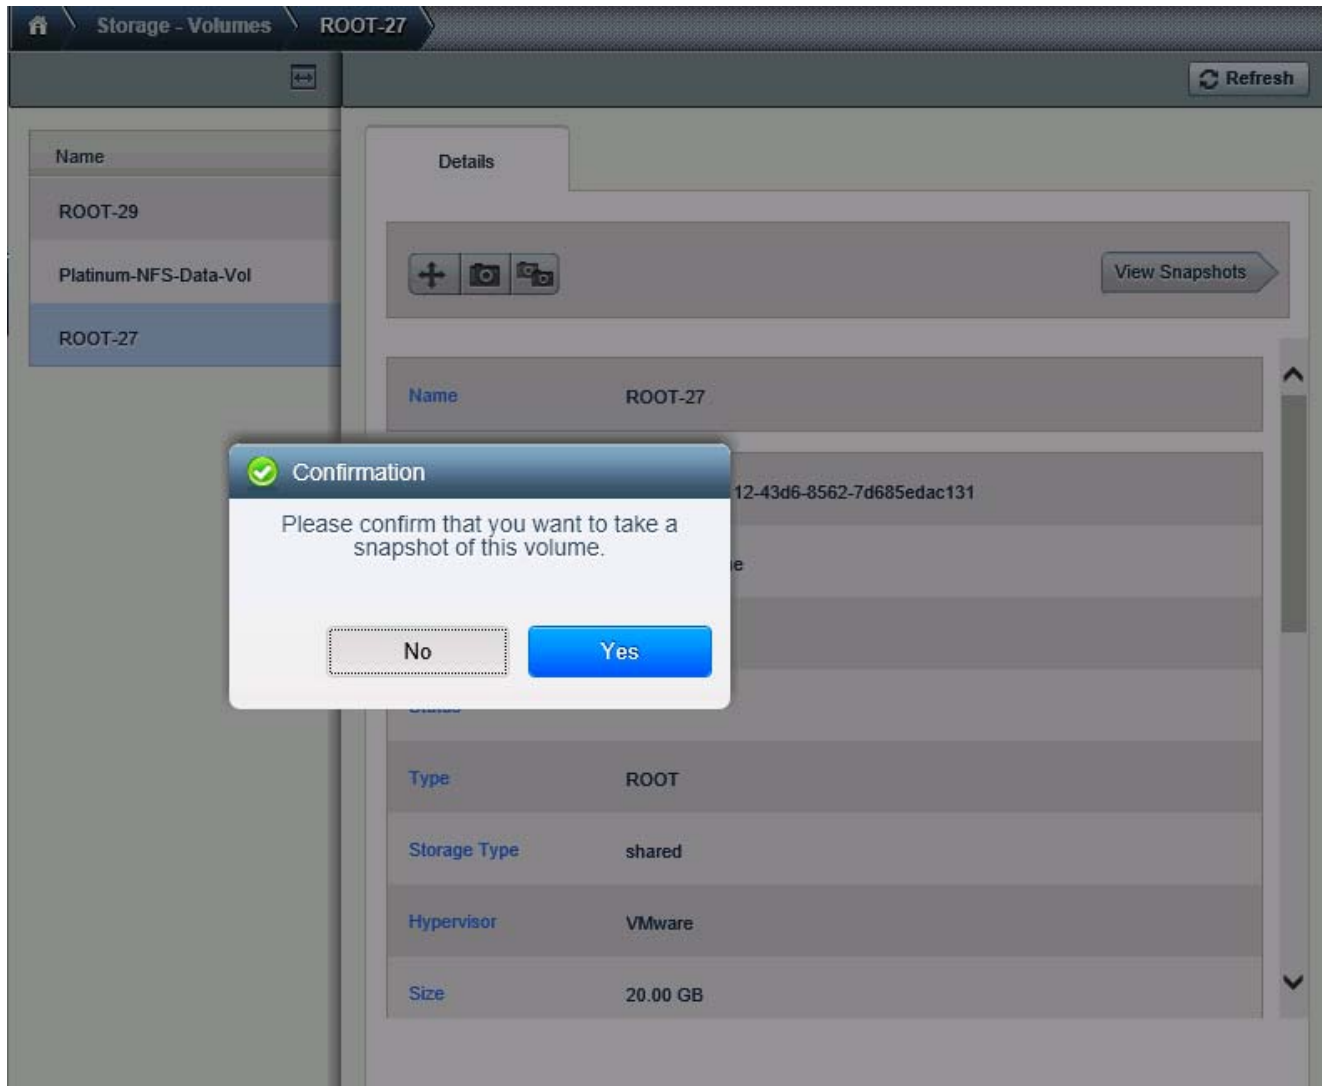
**Figure 258**      **Displaying Snapshots of Storage Volume**

The screenshot displays a web-based management interface for storage volumes. The top navigation bar includes a home icon, 'Storage - Volumes', and 'ROOT-27'. A 'Refresh' button is located in the top right corner. On the left, a sidebar lists 'Name' with entries 'Platinum-NFS-Data-Vol' and 'ROOT-27'. The main area shows the 'Details' for 'ROOT-27'. A 'Take Snapshot' button is visible, along with icons for adding, deleting, and cloning snapshots, and a 'View Snapshots' button. Below this, a table lists the volume's properties.

|              |                                      |
|--------------|--------------------------------------|
| Name         | ROOT-27                              |
| ID           | 2e86c19a-4112-43d6-8562-7d685edac131 |
| Zone         | TenantA-Zone                         |
| State        | Ready                                |
| Status       |                                      |
| Type         | ROOT                                 |
| Storage Type | shared                               |
| Hypervisor   | VMware                               |
| Size         | 20.00 GB                             |

7. Click **Yes** to Confirm.

Figure 259 Taking the Snapshot



## Automatic Backup

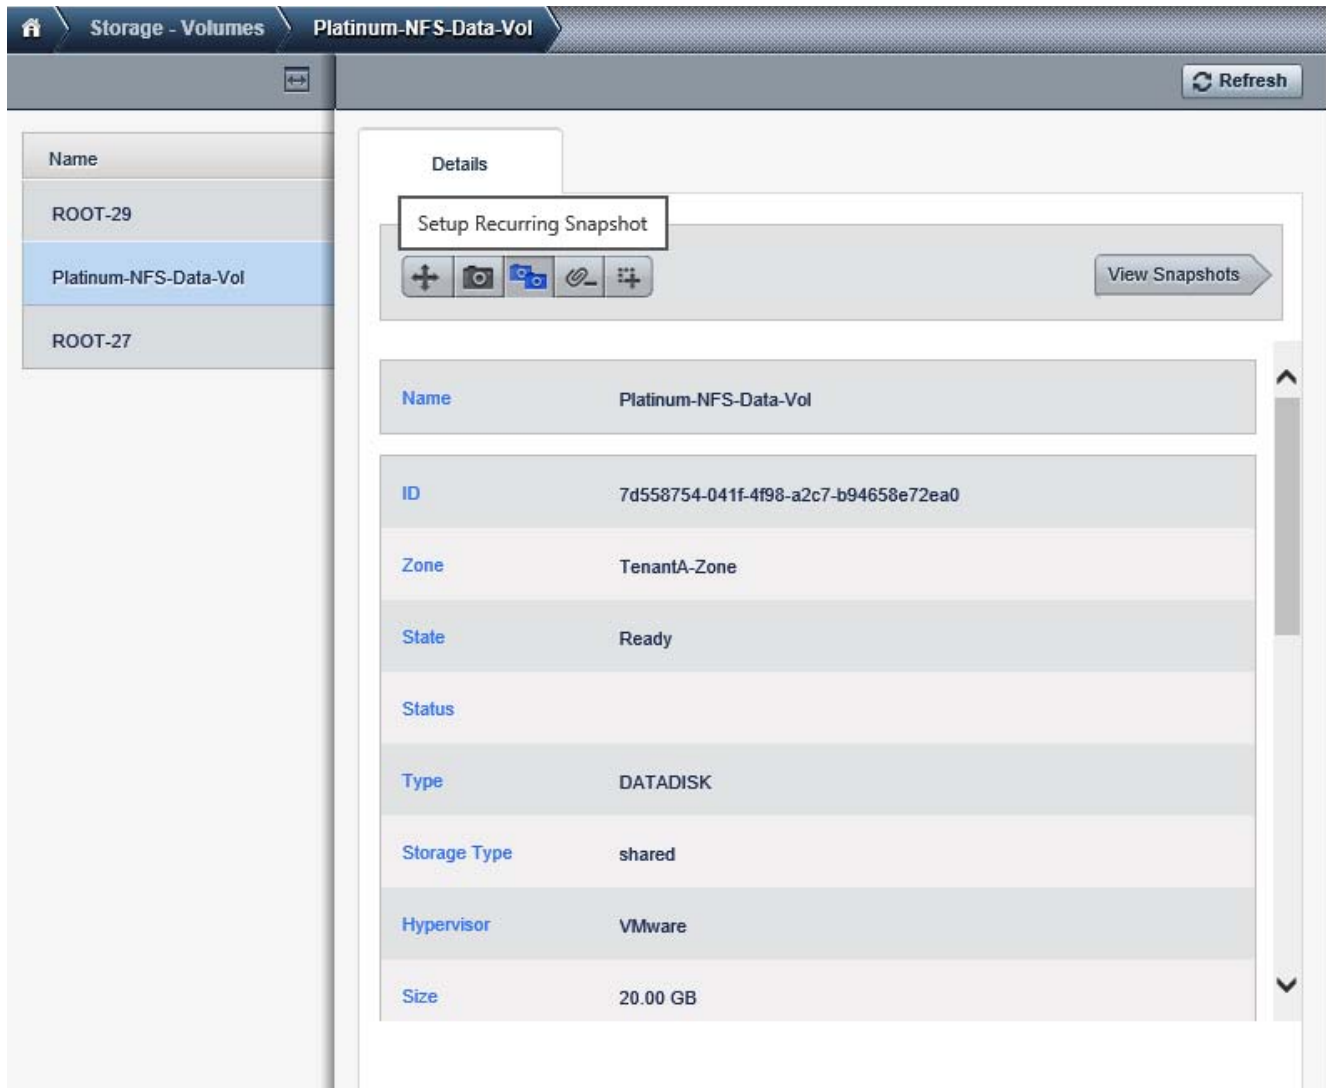
1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Storage** tab.

**Figure 260**      *Displaying the Storage Volume*

| Storage - Volumes     |          |            |                      |            |
|-----------------------|----------|------------|----------------------|------------|
| Select view:          | Volumes  |            | Upload volume        | Add Volume |
| Name                  | Type     | Hypervisor | VM display name      | Quickview  |
| Platinum-NFS-Data-Vol | DATADISK | VMware     | TenantA-Platinum-VM1 | +          |
| ROOT-27               | ROOT     | VMware     | TenantA-Platinum-VM1 | +          |

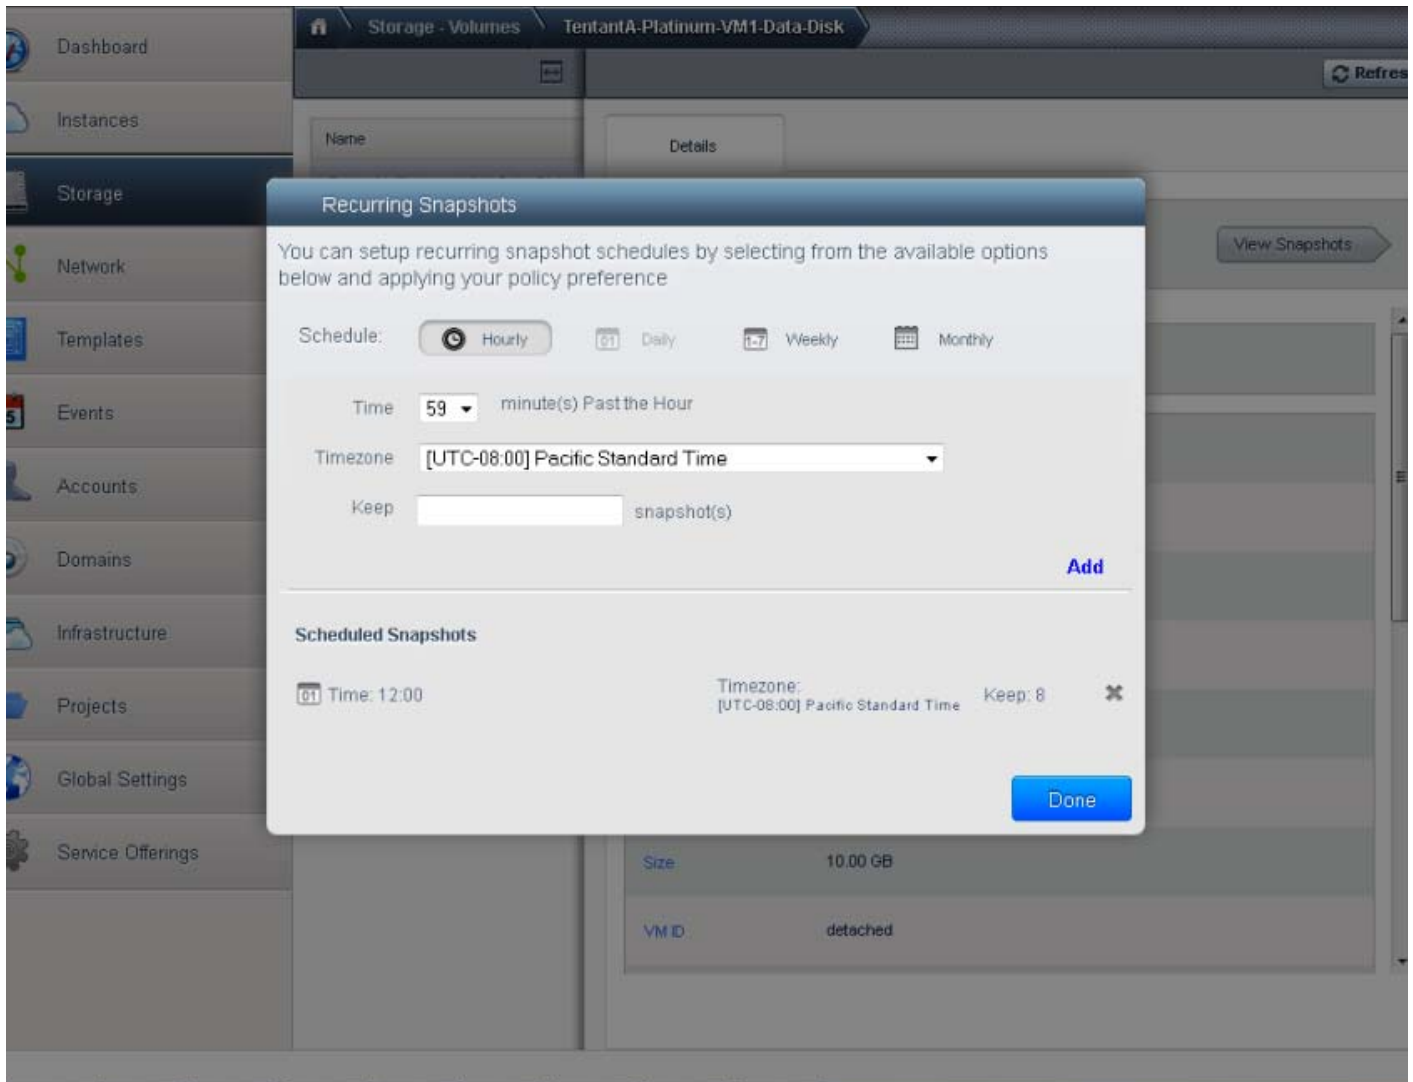
4. Click **Platinum-NFS-Data-Vol** Under Name in the right pane
5. (Automatic backup of TenantA-Platinum-VM1 data disk based on every day at midnight and keep 2 weeks backup data).
6. Click **Recurring Snapshot** icon with double camera.

**Figure 261**      **Setting up the Recurring Snapshot Feature**



7. Click **Daily** button in Schedule field.
8. Select 12 Hour 00 minutes and PM in Time list box.
9. Select Pacific Standard Time in Timezone list box.
10. Enter 8 value in Keep field.
11. Click **Add**.
12. Click **Done**.

**Figure 262**      **Defining the Schedule for Recurring Snapshot**



## Template based Snapshot Restore

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Storage** tab.

**Figure 263**      *Displaying the Storage Volume*

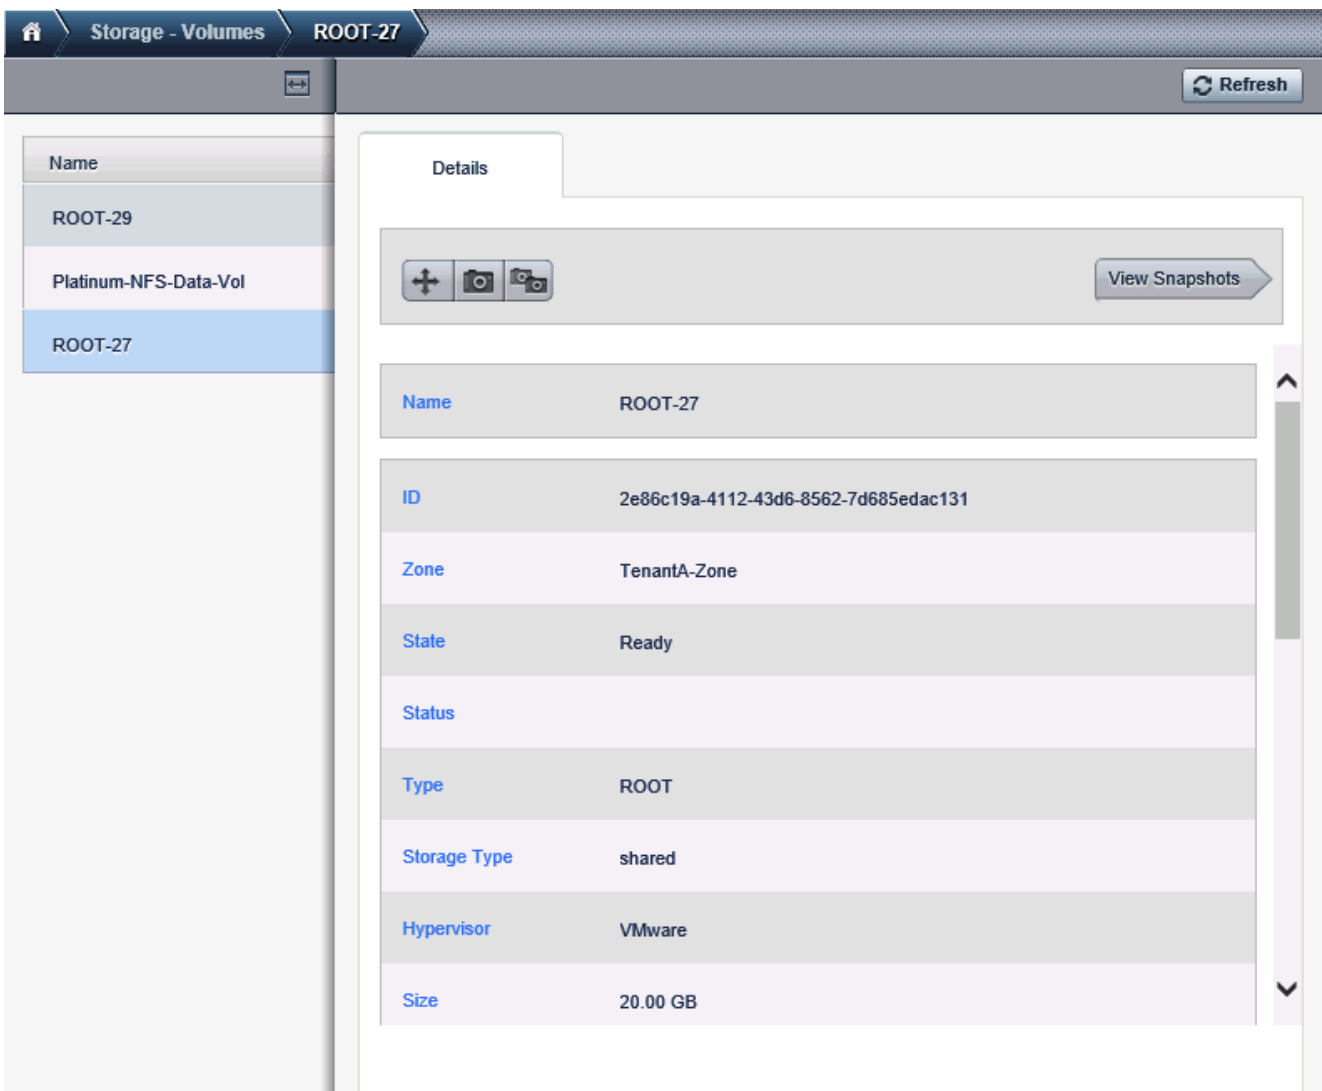


The screenshot shows the 'Storage - Volumes' interface. At the top, there is a 'Select view:' dropdown set to 'Volumes', a search bar, and buttons for 'Upload volume' and '+ Add Volume'. Below this is a table with the following data:

| Name                  | Type     | Hypervisor | VM display name      | Quickview |
|-----------------------|----------|------------|----------------------|-----------|
| Platinum-NFS-Data-Vol | DATADISK | VMware     | TenantA-Platinum-VM1 | +         |
| ROOT-27               | ROOT     | VMware     | TenantA-Platinum-VM1 | +         |

4. Click **Root-27** under Storage-Volumes.
5. Click **View Snapshots**.

**Figure 264**      *Displaying the Snapshot Details*



The screenshot shows the 'Storage - Volumes' interface with the 'ROOT-27' volume selected. The left sidebar shows a list of volumes: 'Name', 'ROOT-29', 'Platinum-NFS-Data-Vol', and 'ROOT-27' (highlighted). The main area displays the details for 'ROOT-27'.

**Details**

Buttons: +, Camera, Refresh, View Snapshots


|                     |                                      |
|---------------------|--------------------------------------|
| <b>Name</b>         | ROOT-27                              |
| <b>ID</b>           | 2e86c19a-4112-43d6-8562-7d685edac131 |
| <b>Zone</b>         | TenantA-Zone                         |
| <b>State</b>        | Ready                                |
| <b>Status</b>       |                                      |
| <b>Type</b>         | ROOT                                 |
| <b>Storage Type</b> | shared                               |
| <b>Hypervisor</b>   | VMware                               |
| <b>Size</b>         | 20.00 GB                             |

6. Click **Root-27 Snapshot**.



7. Click **View Snapshots**.

**Figure 265**      *Viewing the Snapshot Status*

| Storage - Volumes   ROOT-27   Snapshots |               |                               |                                                                                              |           |
|-----------------------------------------|---------------|-------------------------------|----------------------------------------------------------------------------------------------|-----------|
| Volume                                  | Interval Type | Created                       | State                                                                                        | Quickview |
| ROOT-27                                 | MANUAL        | Fri, 20 Sep 2013 08:50:47 UTC |  BackedUp | +         |

8. Click **Create Template** icon with + symbol.

**Figure 266**      *Creating the New Template*

| Storage - Volumes   ROOT-27   Snapshots |               |         |       |           |
|-----------------------------------------|---------------|---------|-------|-----------|
| Volume                                  | Interval Type | Created | State | Quickview |
| ROOT-27                                 | MANUAL        |         |       | +         |

Quickview: ROOT-27

---




**Name**      TenantA-Platinum-VM1\_ROOT-27\_20130920085047

**ID**      c4bc5358-45c1-411b-b05c-b00eda11184b

**Volume Name**      ROOT-27

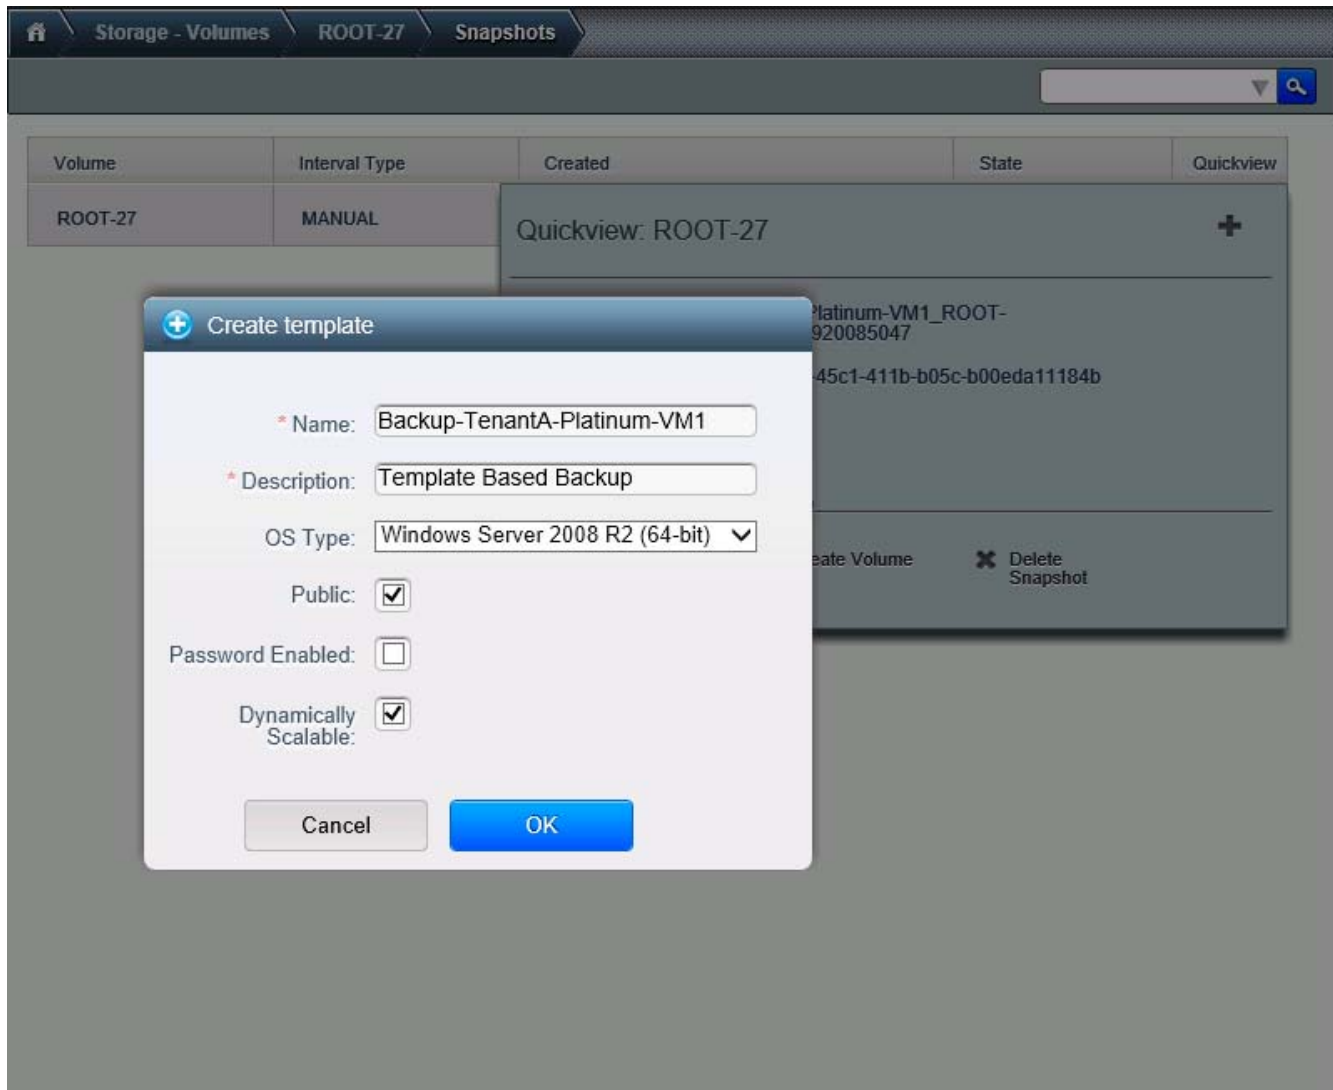
**State**      Create template      BackedUp

---

 Create template  
  Create Volume  
  Delete Snapshot

9. Click **Create Template** icon with + symbol.
10. Enter Backup-TenantA-Platinum-VM1.
11. Enter Template based Restore.
12. Select Windows 2008 R2(64-bit) in OS Type list box.
13. Select **Public** check box.
14. Do not select **Password Enabled** check box.

Figure 267 Defining the Template Properties



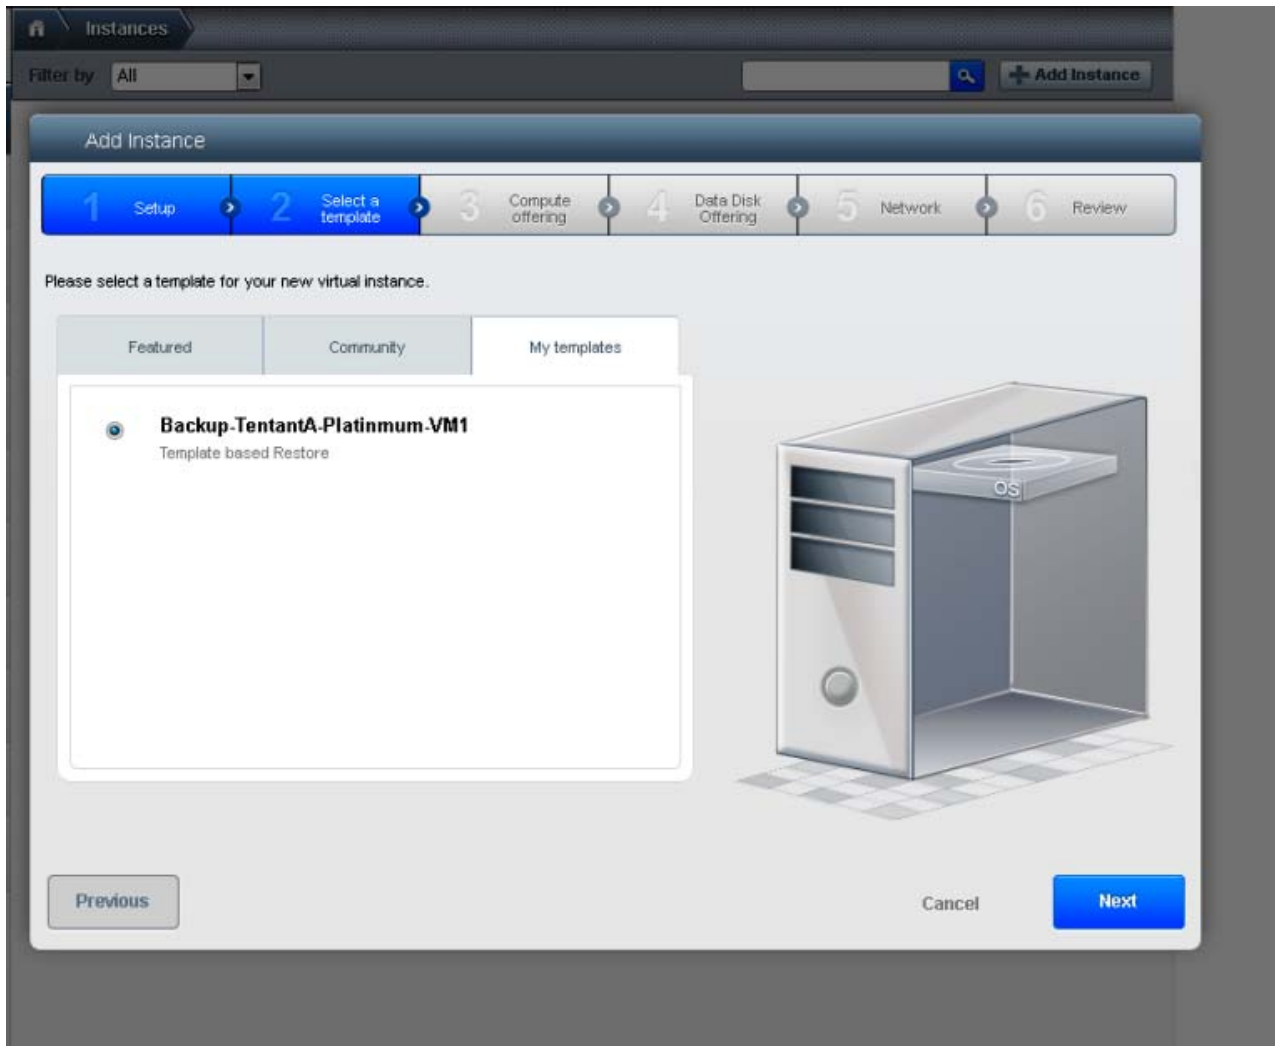
15. Click the **Instances** tab.
16. Click **Add Instance** button in the right pane.
17. Select TenantA-Zone in Zone list box.
18. Click **Template** radio button.
19. Click **Next**.

**Figure 268**      **Selecting the Zone for the Template**



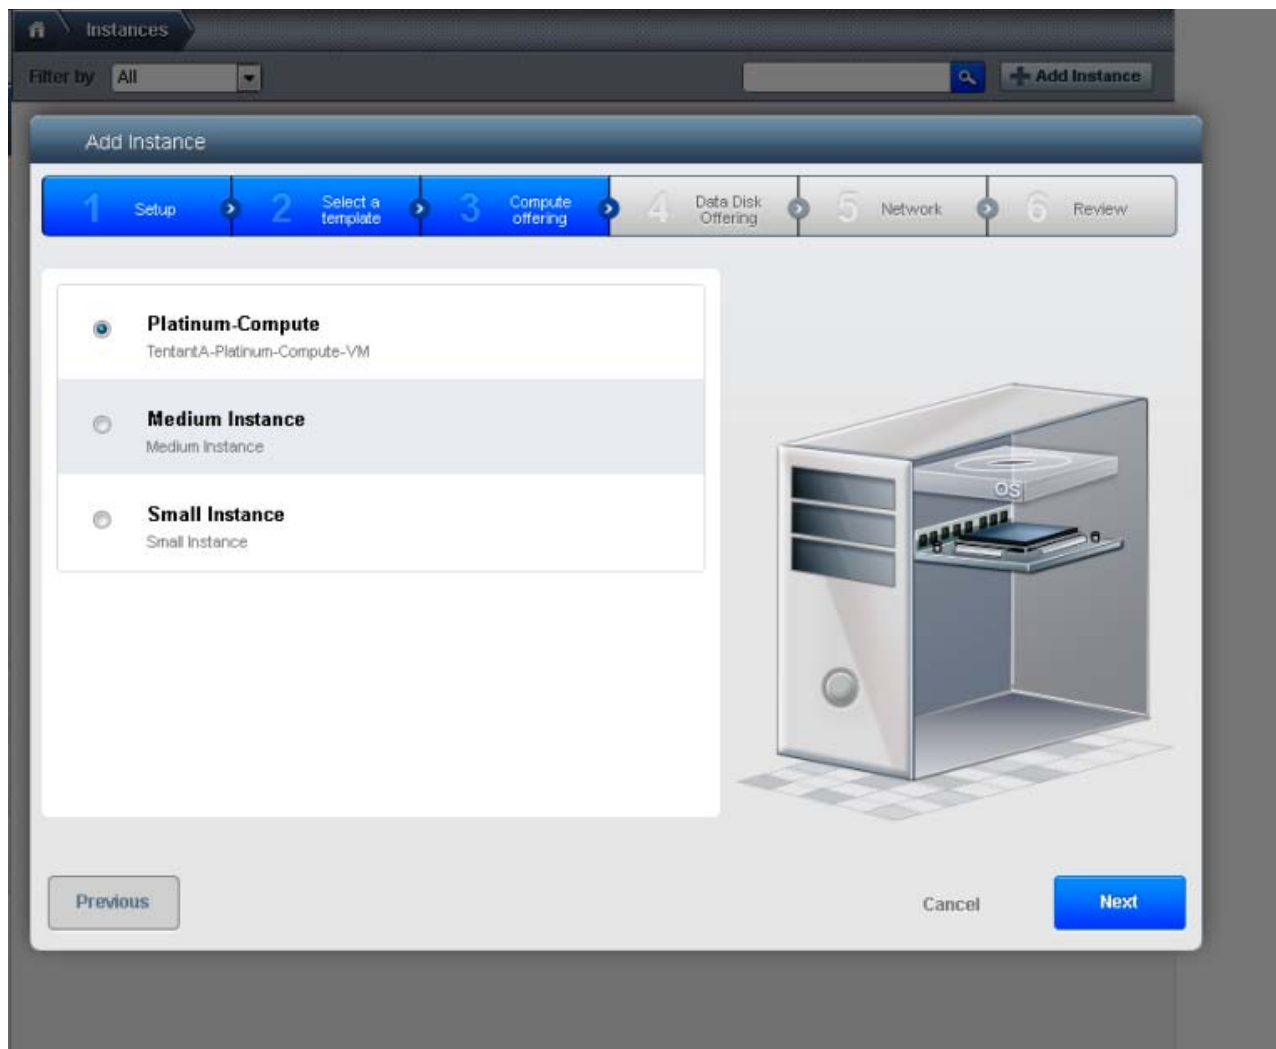
20. Click **My Templates**.
21. Click **Backup-TenantA-Platinum-VM1** radio button.
22. Click **Next**.

**Figure 269**      *Selecting the Template for Virtual Interface*



23. Click **Platinum-Compute** radio button.
24. Click **Next**.

**Figure 270**      **Selecting the Compute Offering**



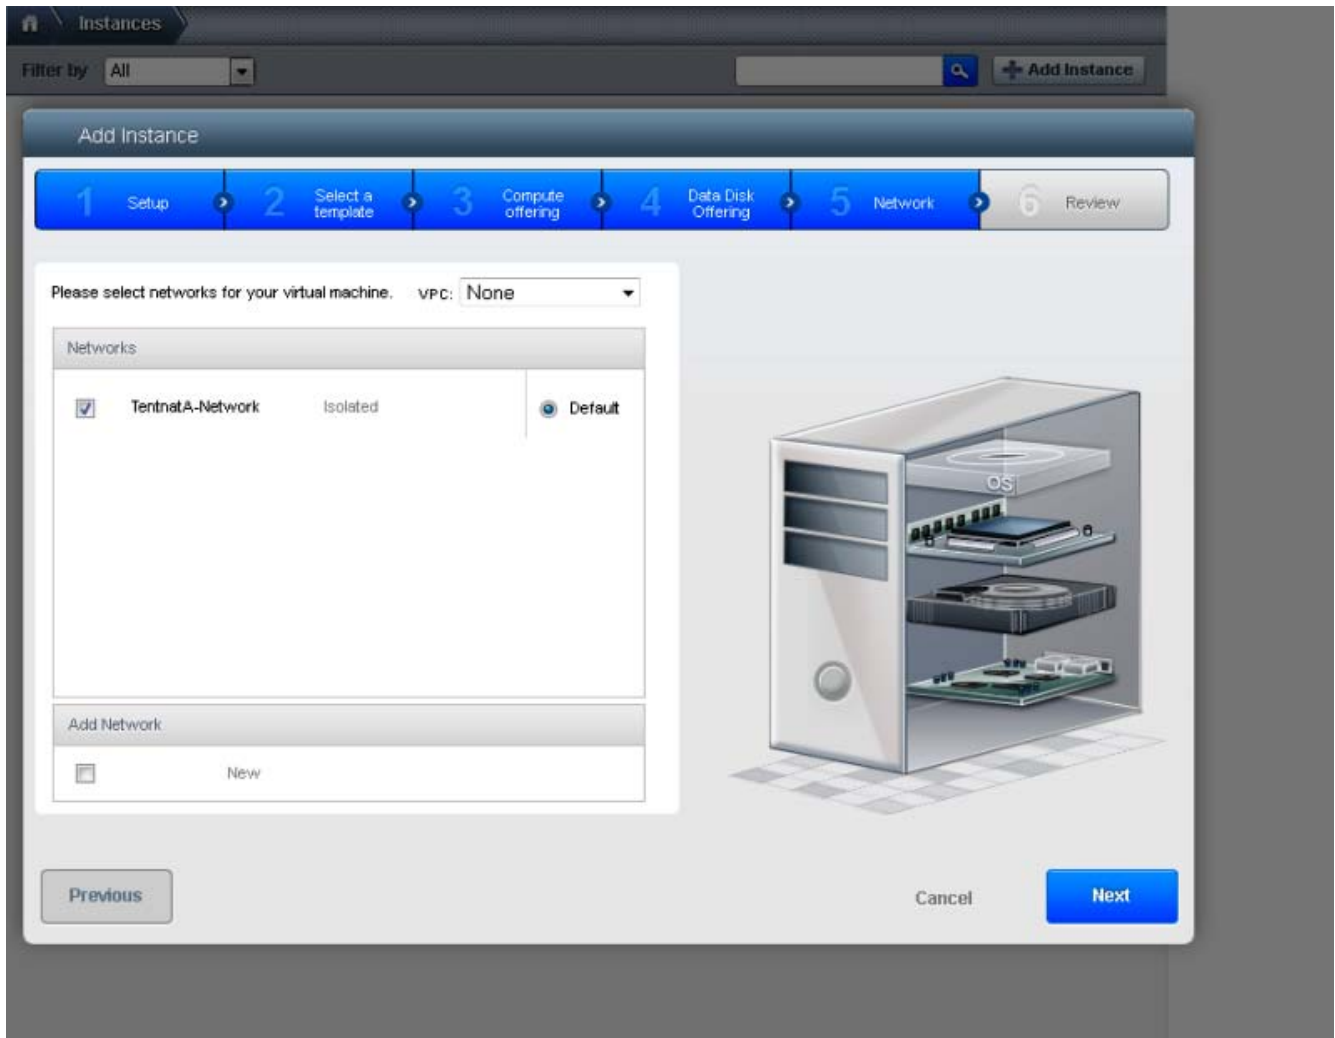
25. Click **No Thanks** radio button.
26. Click **Next**.

**Figure 271** Not selecting any Disk Offering



27. Check TenantA-Network check box, and click **Default** radio button.
28. Click **Next**.

**Figure 272**      **Selecting the Disk Networks for VM**



29. Enter Backup-TenantA-Platinum-VM1 in Name (Optional) field.
30. Click **Launch VM**.

**Figure 273**      *Verifying the Virtual Instance Information*

**Instances**

Filter by All

**+ Add Instance**

1 Setup > 2 Select a template > 3 Compute offering > 4 Data Disk Offering > 5 Affinity > 6 Network > 7 Review

Please review the following information and confirm that your virtual instance is correct before launch.

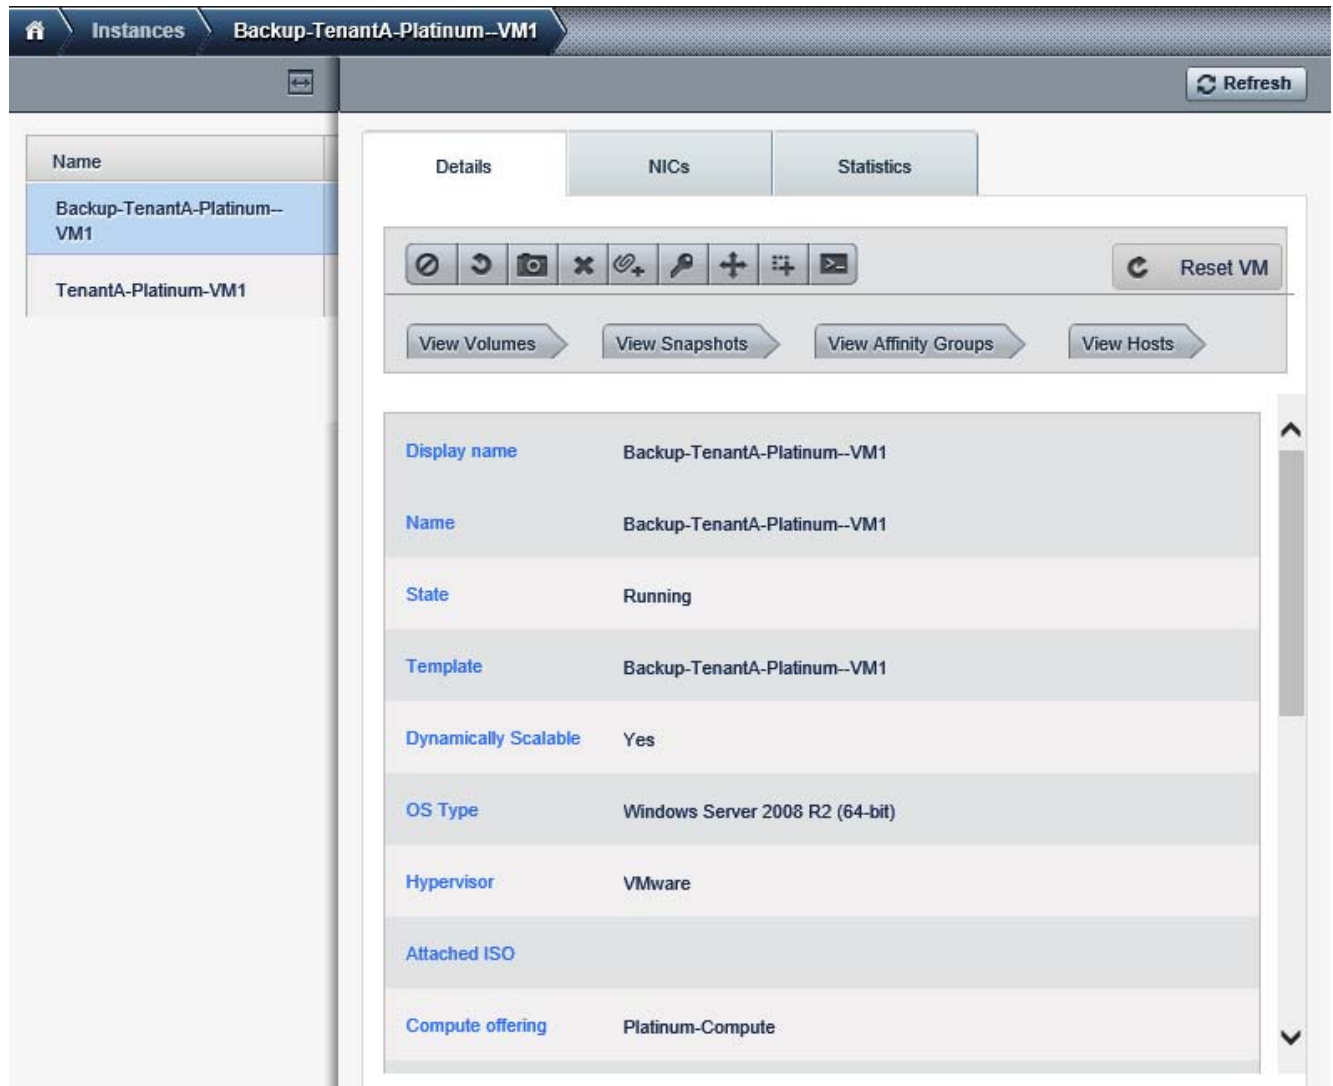
|                         |                             |                      |
|-------------------------|-----------------------------|----------------------|
| Name (Optional)         | Backup-TenantA-Platinun     |                      |
| Add to group (Optional) |                             |                      |
| Zone                    | TenantA-Zone                | <a href="#">Edit</a> |
| Hypervisor              | VMware                      | <a href="#">Edit</a> |
| Template                | Backup-TenantA-Platinum-VM1 | <a href="#">Edit</a> |
| Compute offering        | Platinum-Compute            | <a href="#">Edit</a> |
| Data Disk Offering      | (None)                      | <a href="#">Edit</a> |
| Affinity Groups         | (None)                      | <a href="#">Edit</a> |
| Network                 | TenantA-Network             | <a href="#">Edit</a> |

[Previous](#) [Cancel](#) [Launch VM](#)

31. Click the **Instances** tab.
32. Click **Backup-TenantA-Platinum-VM1**.

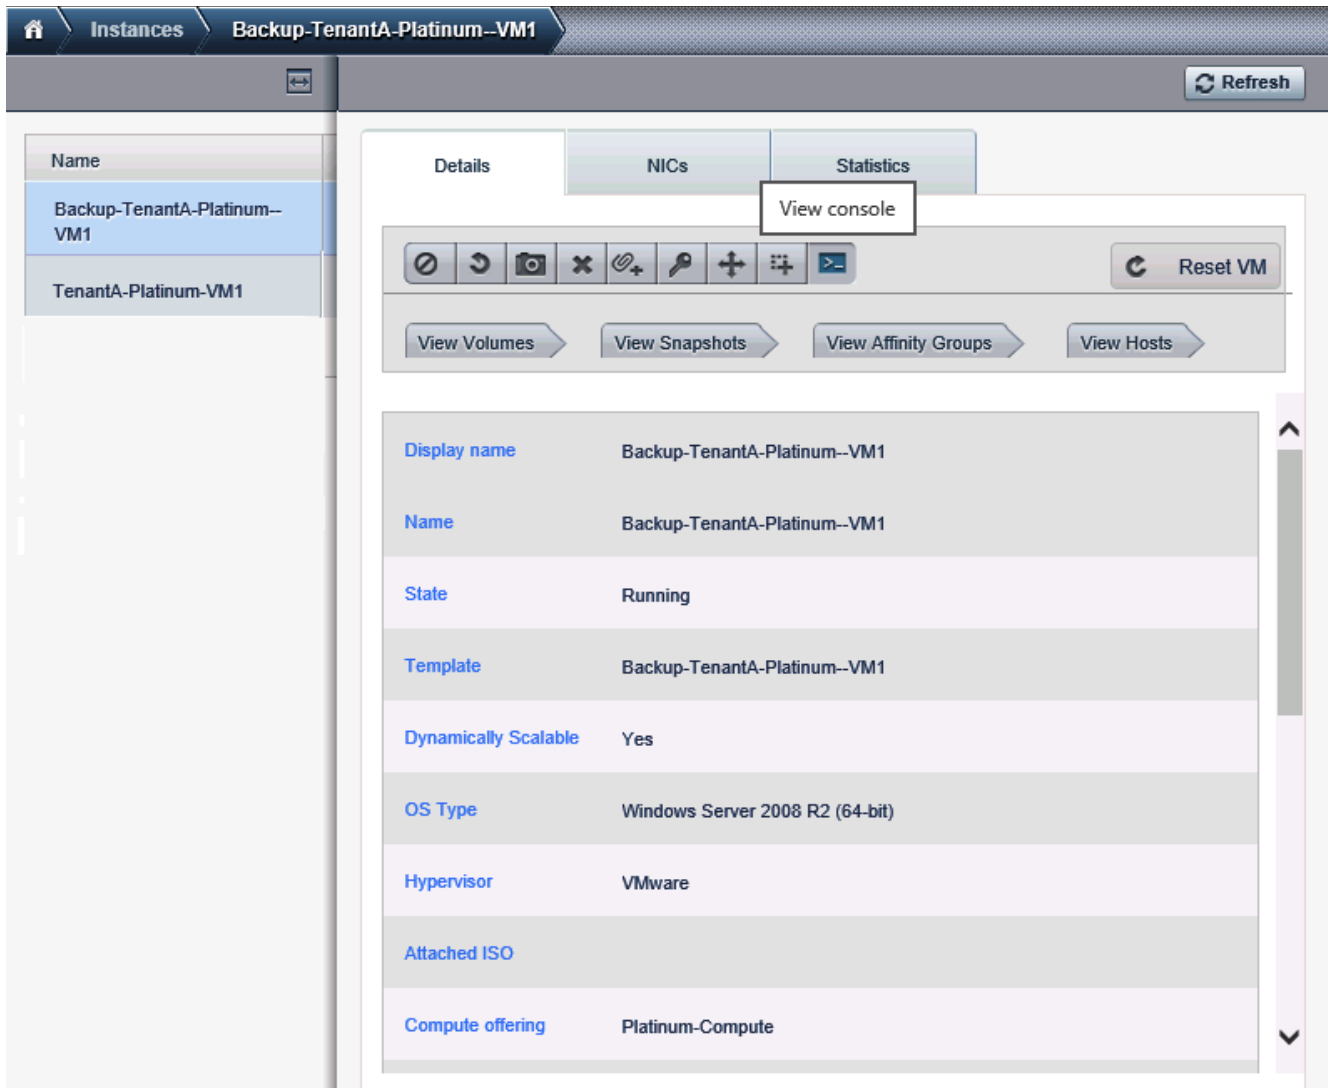


**Figure 274**      **Selecting the Backup VM**



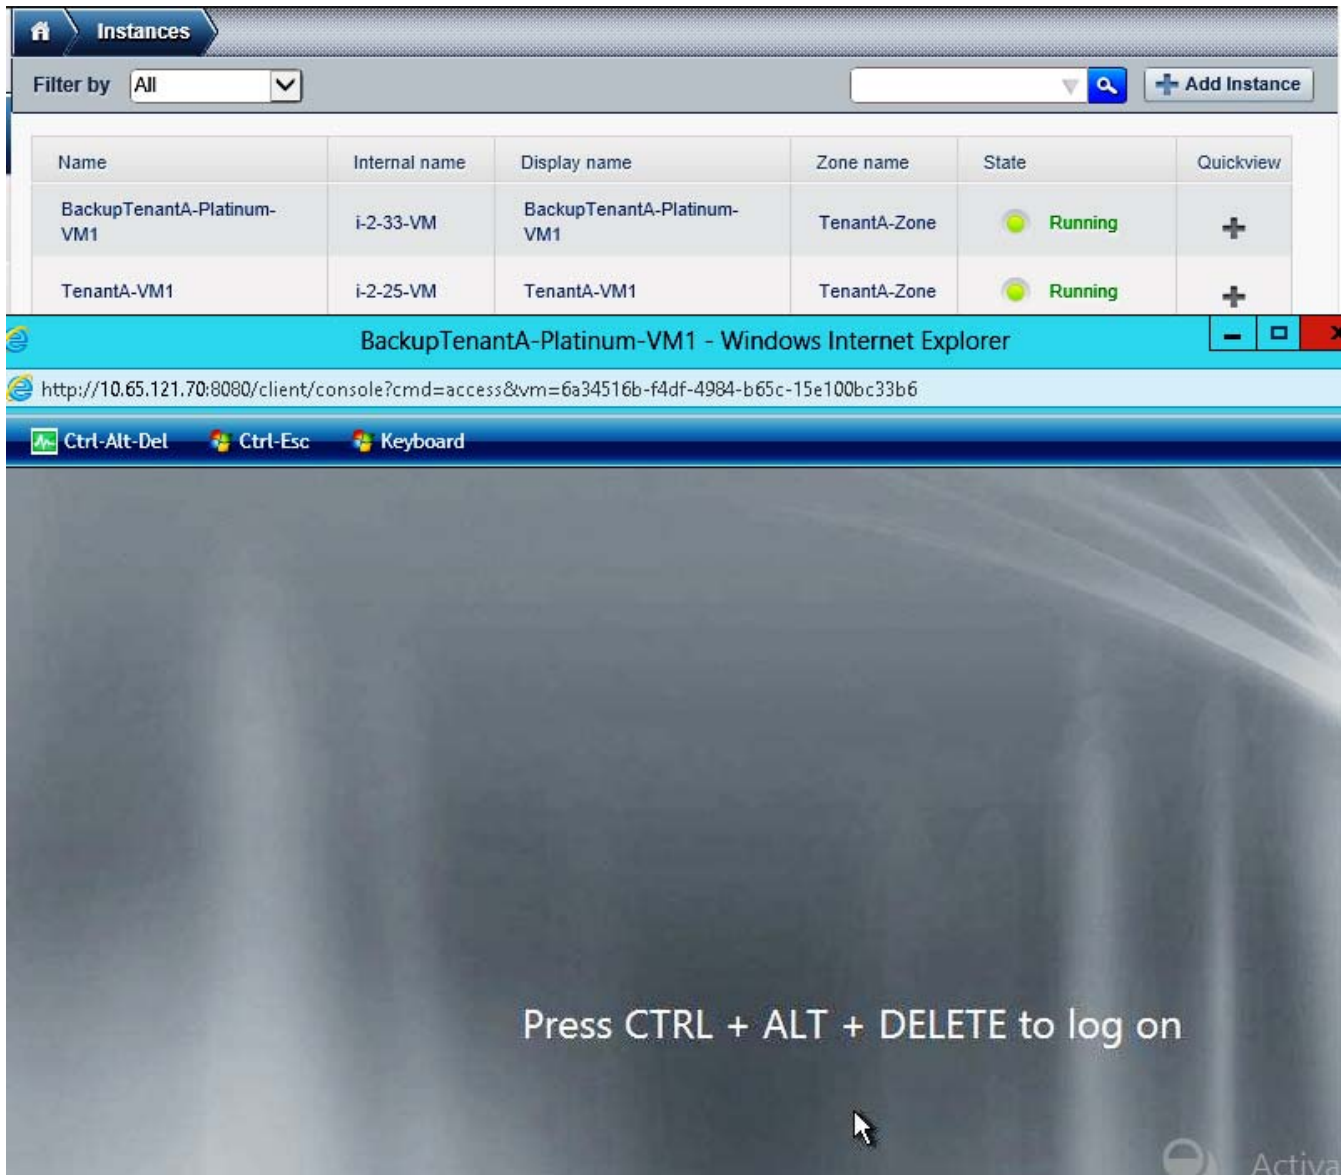
33. Click View Console.

**Figure 275**      *Viewing the Backup VM Details*



34. Enter Username <administrator> and password <XXXXXX> and login
35. Note: The Password is same has Parent Instance VM is set.

**Figure 276**      **Logging into the VM**



## Volume based Snapshot Restore

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Storage** tab.
4. Select **Volumes** under Select View list box.

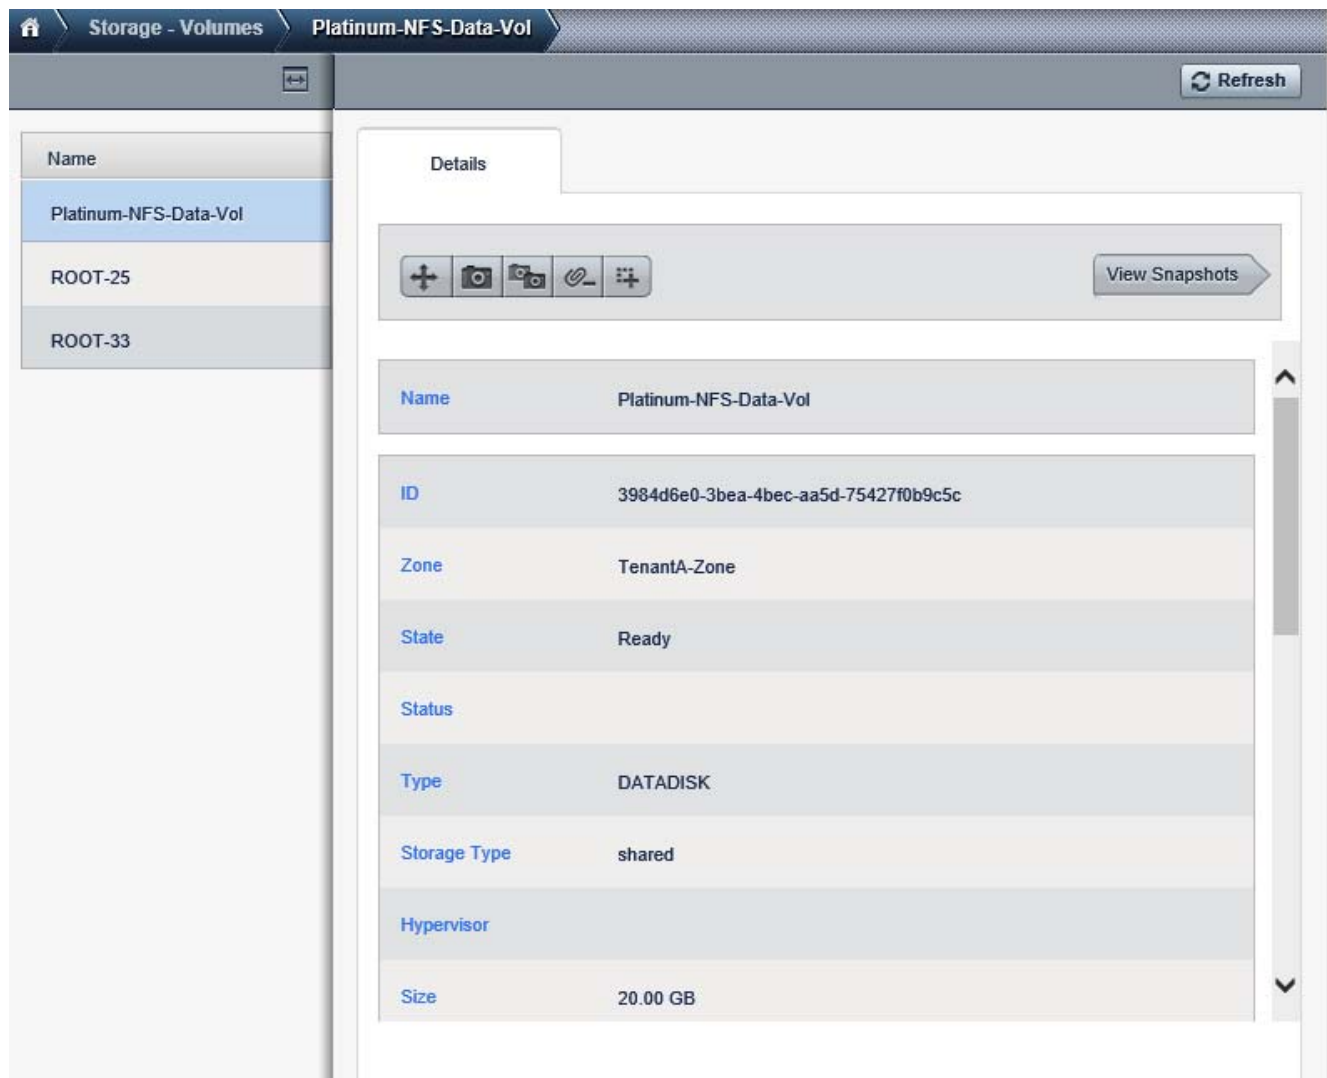
**Figure 277**      *Displaying Storage Volume*

The screenshot shows a web interface titled "Storage - Volumes". At the top, there is a "Select view:" dropdown menu set to "Volumes", a search bar, and two buttons: "Upload volume" and "+ Add Volume". Below this is a table with five columns: "Name", "Type", "Hypervisor", "VM display name", and "Quickview". The table contains three rows of data.

| Name                  | Type     | Hypervisor | VM display name            | Quickview |
|-----------------------|----------|------------|----------------------------|-----------|
| Platinum-NFS-Data-Vol | DATADISK |            | TenantA-VM1                | +         |
| ROOT-25               | ROOT     |            | TenantA-VM1                | +         |
| ROOT-33               | ROOT     |            | BackupTenantA-Platinum-VM1 | +         |

5. Click **TenantA-Platinum-Data-Disk** Volume.
6. Click **View Snapshots**.

**Figure 278** Viewing Snapshot of Virtual Data Disk



7. Click **TenantA-Platinum-Data-Disk Volume**.
8. Click **View Snapshots**.
9. Click **Create Volume** icon.

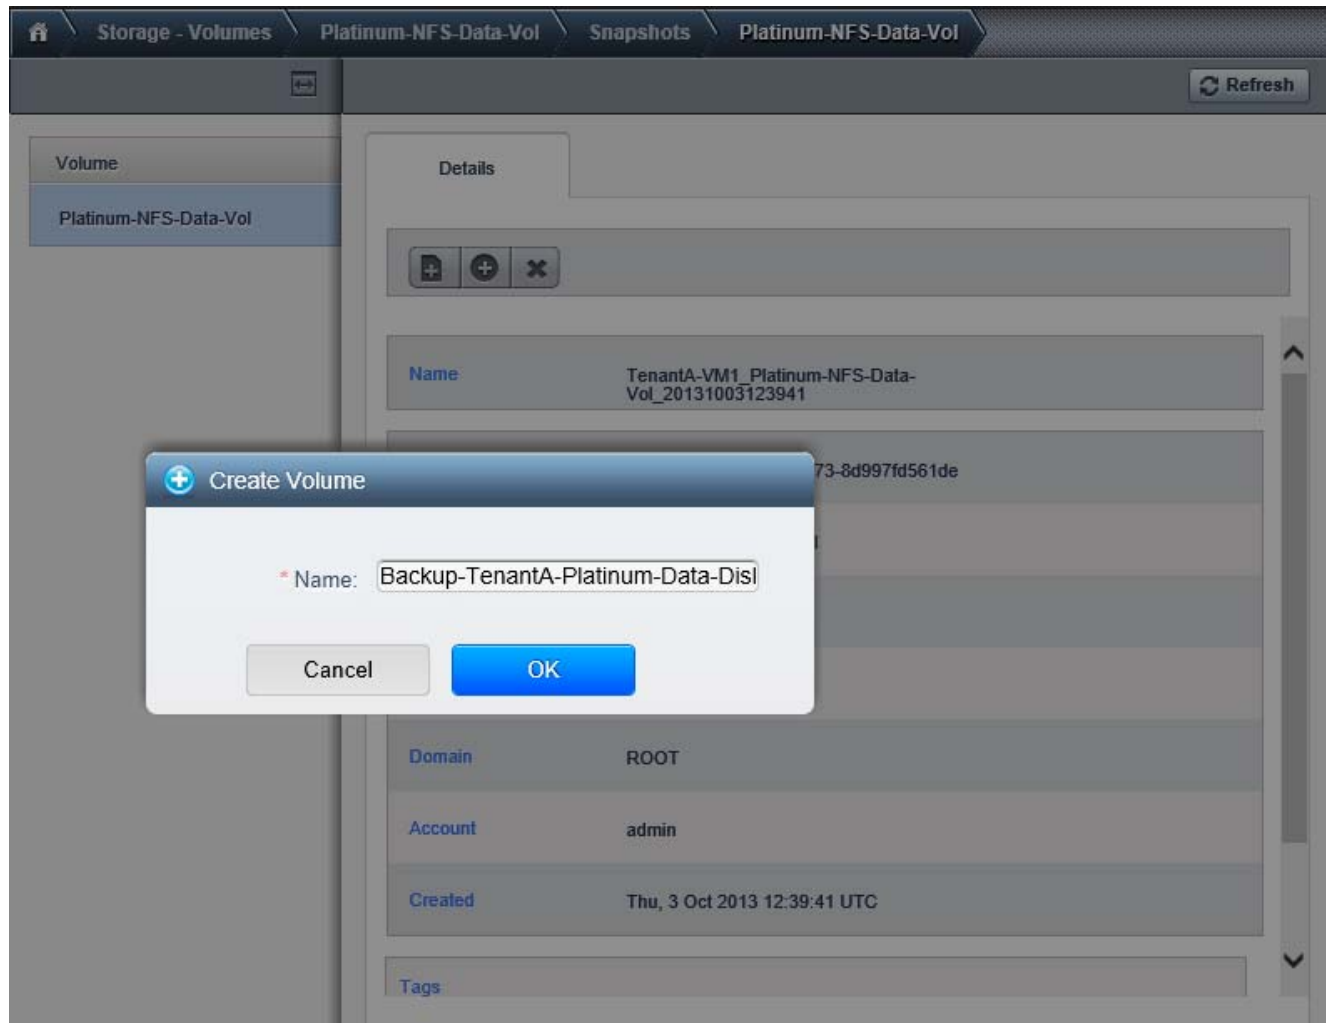
**Figure 279**      **Creating Volume**

The screenshot shows a web-based storage management interface. At the top, there is a navigation bar with tabs: 'Storage - Volumes', 'Platinum-NFS-Data-Vol', 'Snapshots', and 'Platinum-NFS-Data-Vol'. Below the navigation bar, on the left, is a sidebar with a 'Volume' section containing a link to 'Platinum-NFS-Data-Vol'. The main area is titled 'Details' and contains a 'Create Volume' dialog box. The dialog box has a 'Name' field with the text 'TenantA-VM1\_Platinum-NFS-Data-Vol\_20131003123941', an 'ID' field with 'dcc5fdc0-102f-4c66-9673-8d997fd561de', a 'Volume Name' field with 'Platinum-NFS-Data-Vol', a 'State' field with 'BackedUp', an 'Interval Type' field with 'MANUAL', a 'Domain' field with 'ROOT', an 'Account' field with 'admin', and a 'Created' field with 'Thu, 3 Oct 2013 12:39:41 UTC'. There is also a 'Tags' field at the bottom. The dialog box has a '+', a '+', and an 'x' button. A 'Refresh' button is located in the top right corner of the main area.

| Details          |                                                  |
|------------------|--------------------------------------------------|
| Create Volume    |                                                  |
| <div>+ + x</div> |                                                  |
| Name             | TenantA-VM1_Platinum-NFS-Data-Vol_20131003123941 |
| ID               | dcc5fdc0-102f-4c66-9673-8d997fd561de             |
| Volume Name      | Platinum-NFS-Data-Vol                            |
| State            | BackedUp                                         |
| Interval Type    | MANUAL                                           |
| Domain           | ROOT                                             |
| Account          | admin                                            |
| Created          | Thu, 3 Oct 2013 12:39:41 UTC                     |
| Tags             |                                                  |

10. Enter <Backup-TenantA-Platinum-Data-Disk> in the Name field.
11. Click **OK**.

**Figure 280** Adding Name for the Volume



12. Click **Storage** tab.
13. Click **Backup-TenantA-Platinum-Data-Disk**.

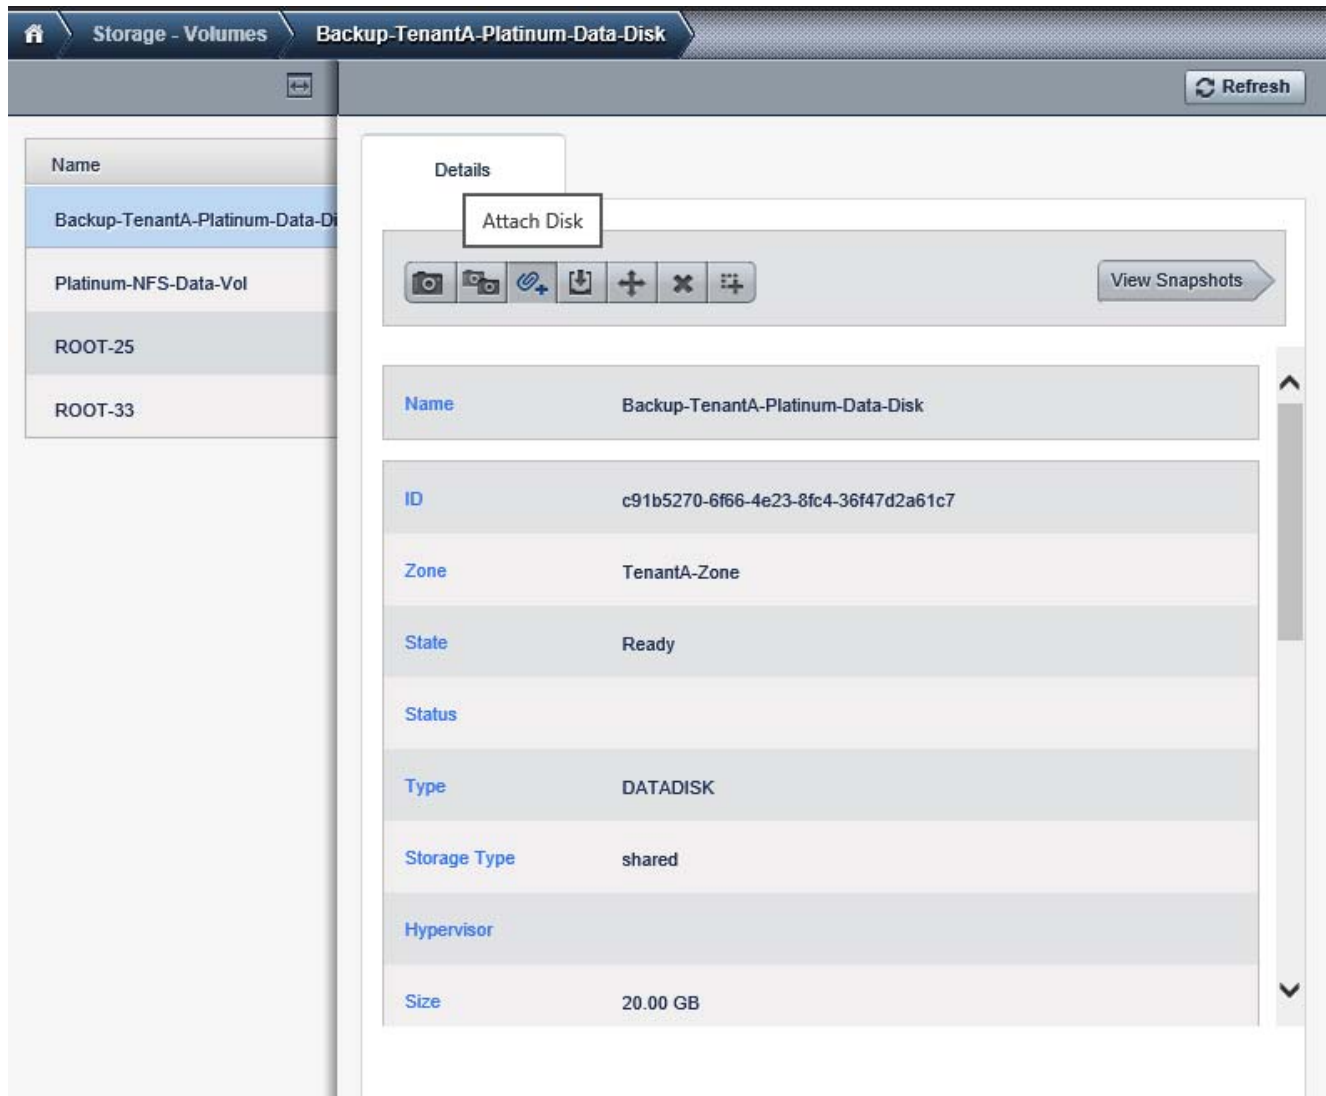
**Figure 281**      *Displaying the Storage Volume*

| Storage - Volumes                 |          |            |                            |              |
|-----------------------------------|----------|------------|----------------------------|--------------|
| Select view:                      | Volumes  |            | Upload volume              | + Add Volume |
| Name                              | Type     | Hypervisor | VM display name            | Quickview    |
| Backup-TenantA-Platinum-Data-Disk | DATADISK |            |                            | +            |
| Platinum-NFS-Data-Vol             | DATADISK |            | TenantA-VM1                | +            |
| ROOT-25                           | ROOT     |            | TenantA-VM1                | +            |
| ROOT-33                           | ROOT     |            | BackupTenantA-Platinum-VM1 | +            |

14. Click **Attach Disk** icon.



**Figure 282**      **Attaching the Disk**



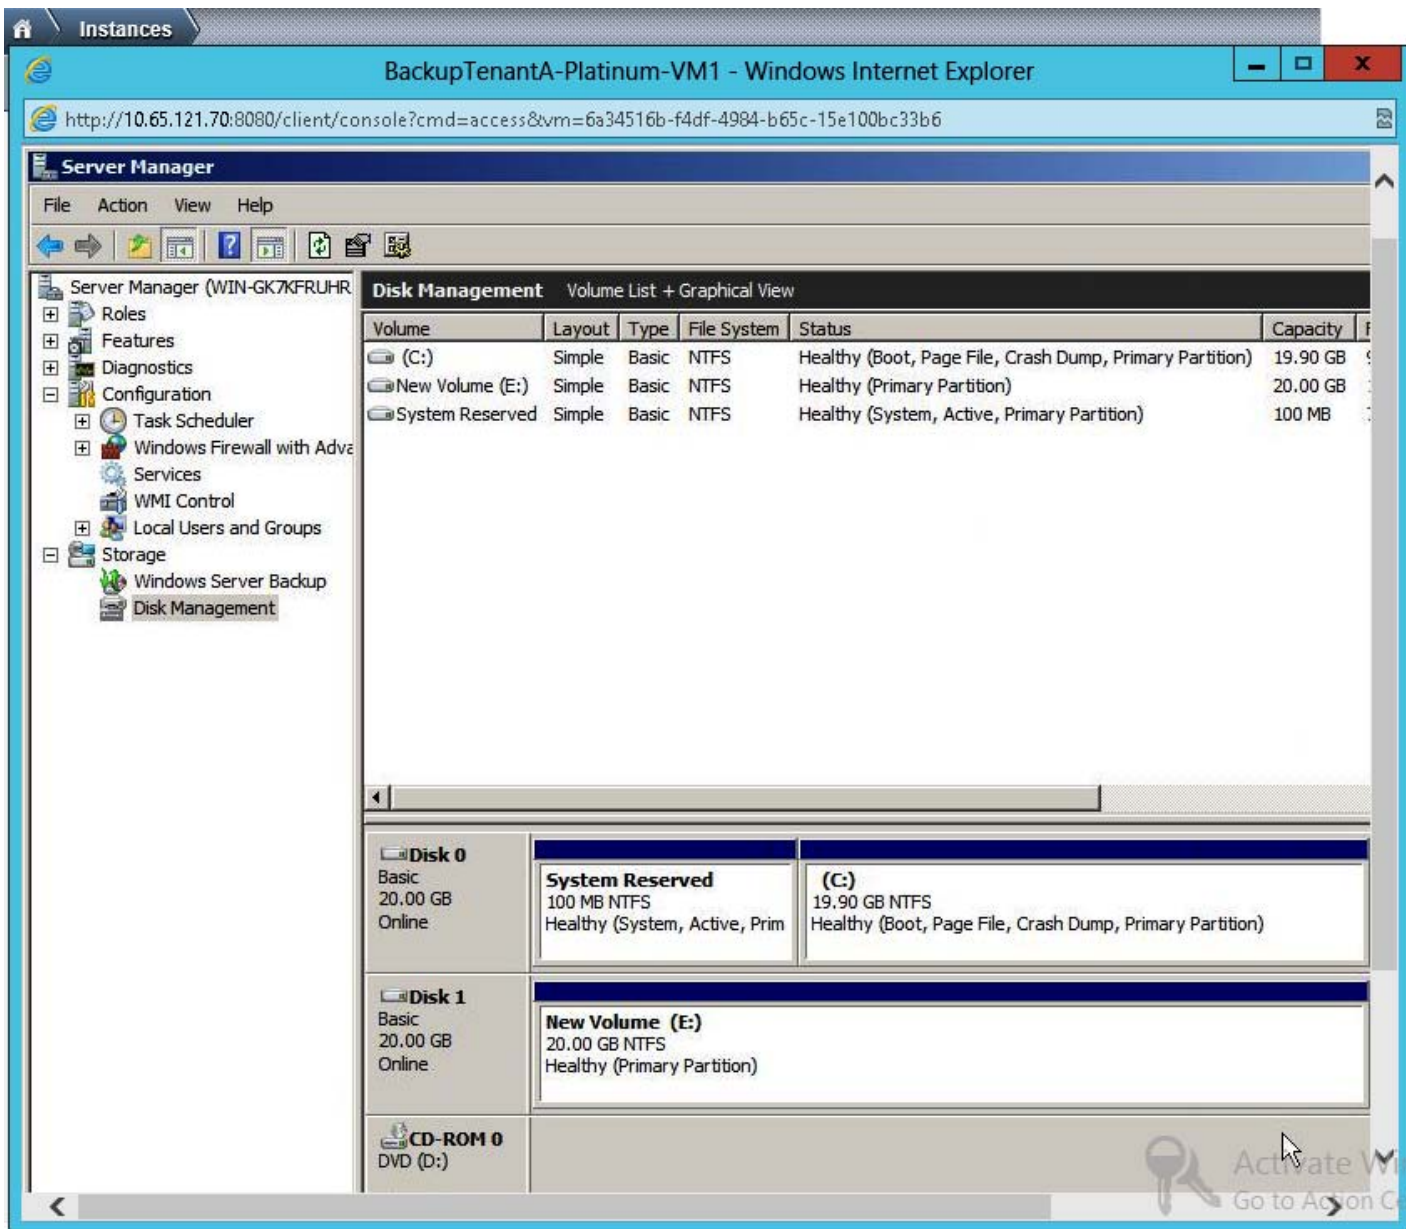
15. Select **Backup-TenantA-Platinum-VM1** under Instance list box.
16. Click **OK**.

**Figure 283**      **Selecting the Disk to the Instance**



17. Select Backup-TenantA-Platinum-VM1 under Instance list box.
18. Click **View console**.
19. Provide User Name <Administrator> Password <XXXXXX>
20. Access Disk Manumit Online Disk1.

**Figure 284**      *Online snap restored Disk*



## VM Migration

This section explains the configuration steps required to perform the manual migration of the VM from one host to another without interrupting service to users or going into maintenance mode. This is called manual live migration, and can be done under the following conditions:

- The root administrator is logged in. Domain admins and users cannot perform manual live migration of VMs.
- The VM is running. Stopped VMs cannot be live migrated.
- The destination host must be in the same cluster as the original host.
- The VM must not be using local disk storage.
- The destination host must have enough available capacity. If not, the VM will remain in the “migrating” state until memory becomes available.



### Note

In this study VM instance TenantA-Platinum-VM1 on Cloud Host TenantA-HostA will be migrated to TenantA-HostB.

Login to CloudPlatform with User credentials to manually live migrate TenantA-Platinum-VM1 to TenantA-Host B:

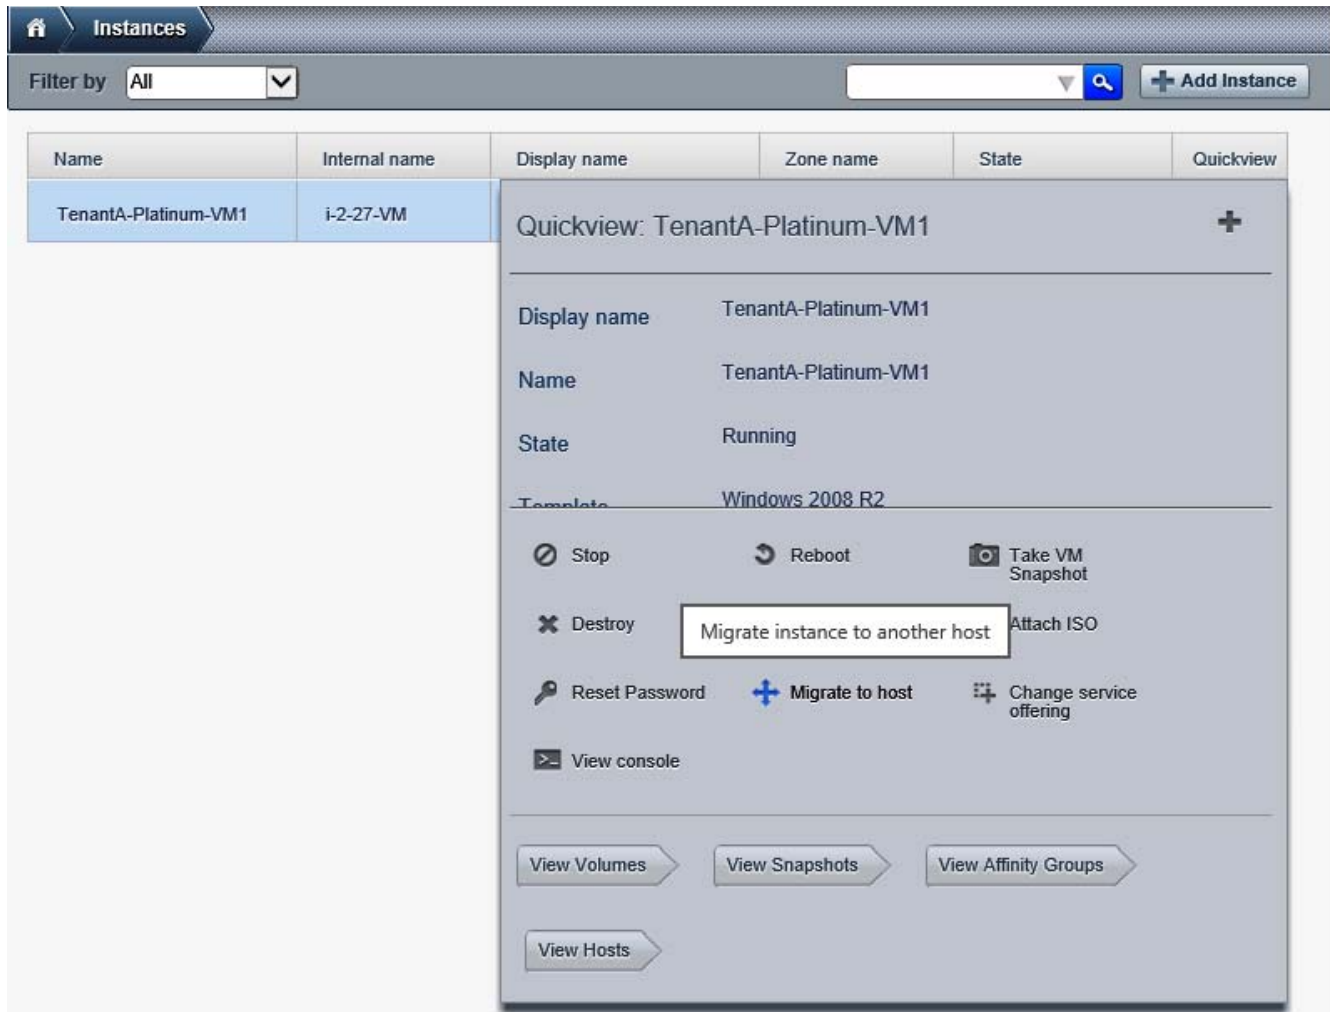
1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instance** tab.
4. Click **TenantA-Platinum-VM1**.

**Figure 285**      *Selecting the Tenant A Platinum Host*

| Name                 | Internal name | Display name         | Zone name    | State                                        | Quickview |
|----------------------|---------------|----------------------|--------------|----------------------------------------------|-----------|
| TenantA-Platinum-VM1 | i-2-27-VM     | TenantA-Platinum-VM1 | TenantA-Zone | <span style="color: green;">●</span> Running | +         |

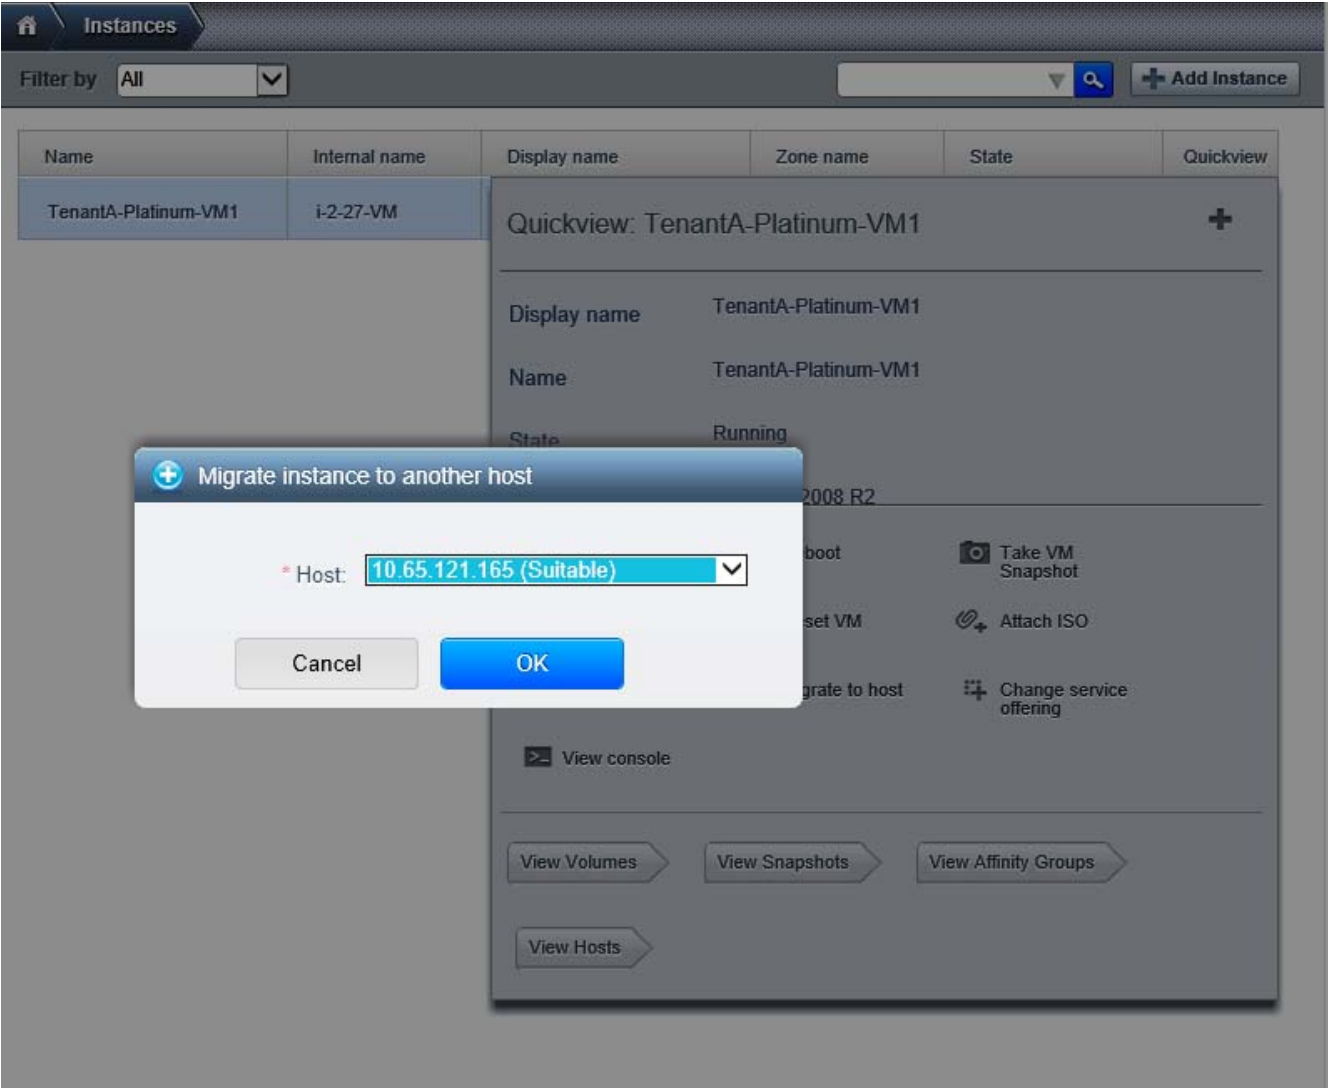
5. Click the **migrate instance to another host** button.

Figure 286 Migrate Instance



6. Select 10.65.121.165 (TenantA-HostB) Host in list box.
7. Click OK.

Figure 287      Selecting Destination Host for Migration



8. VM Migrate task is completed.

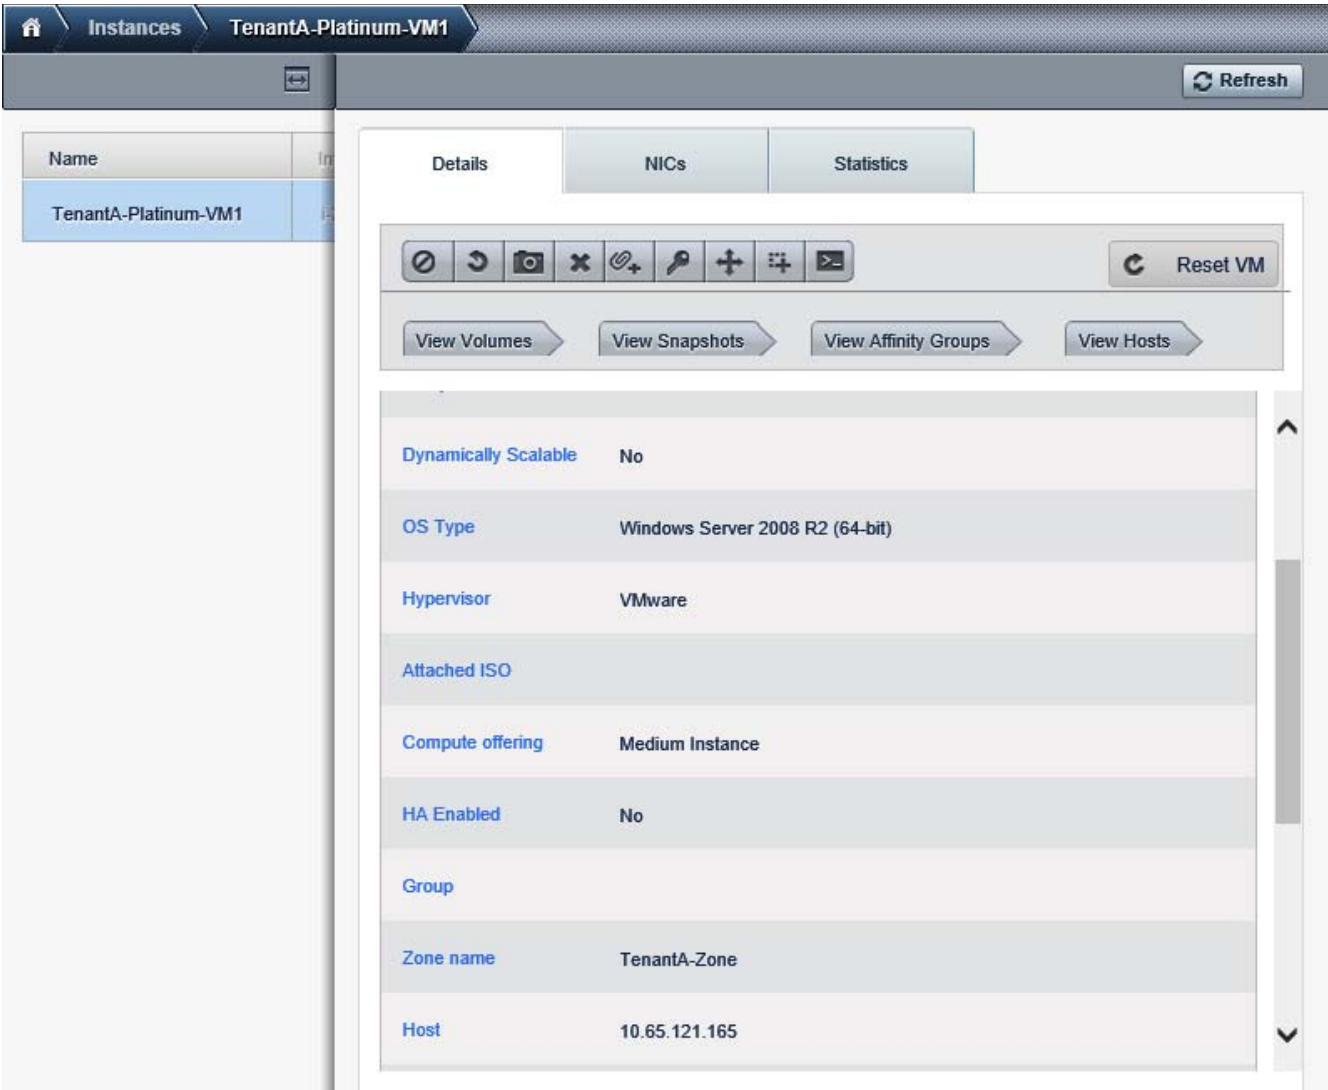
**Figure 288**      **Migration Completion**

The screenshot displays the management console for a virtual machine named "TenantA-Platinum-VM1". The interface includes a top navigation bar with "Instances" and "TenantA-Platinum-VM1" tabs, and a "Refresh" button. Below this is a tabbed interface with "Details", "NICs", "Statistics", and "Volumes". The "Details" tab is active, showing a list of VM properties. Above the list is a toolbar with icons for various actions and a "View Volumes" button. A notification banner at the bottom right indicates that the migration task is completed.

| Property         | Value                                 |
|------------------|---------------------------------------|
| Display name     | TenantA-Platinum-VM1                  |
| Internal name    | i-2-3-VM                              |
| State            | Running                               |
| Hypervisor       | XenServer                             |
| Template         | RHEL6-0                               |
| OS Type          | Red Hat Enterprise Linux 6.0 (64-bit) |
| Attached ISO     | RHEL6-0                               |
| Compute offering | Platinum-Compute                      |
| HA Enabled       | Yes                                   |

**Task completed**  
Migrate instance to another host

Figure 289 VM Migration Summary



VM Storage Migration

This section explains the configuration steps required to perform the manual migration of the Virtual Machine Storage from the primary storage to the other primary storage for maintenance, or load balancing mode. The storage migration can be performed on stopped instance only and no disks should be attached to the VM instance.



Note

In current CloudPlatform 4.2.1.1 version Virtual Machine Storage Migration with same Cluster in Zone is not supported using GUI, however you can achieve this by making API Call.

To accomplish this use API “migrateVirtualMachineWithVolume” by specifying a target host for VM & target storage pool for each volume of the VM.

Here is example:

```
command=migrateVirtualMachineWithVolume&hostid=<TARGET_HOST>&virtualmachineid=<V
M_IN_QUESTION>&migrateto[<I>].volume=<VOLUME_ID>&migrateto[<I>].pool=<TARGET_POOL
_ID>
<TARGET_HOST> - UUID of target host for the VM to be relocated to. This host can
be within or out of cluster.
<VM_IN_QUESTION> - UUID of VM which need to be migrated out of the current
storage pool.
<VOLUME_ID> - UUID of volume of the VM
<TARGET_POOL_ID> - UUID of target primary storage pool. This pool can belong to
current cluster or other cluster)
```

## VM Deletion

This section explains the steps taken for cleaning up virtual machines as part of maintenance efforts by the cloud administrator which releases compute, network and storage resources back to cloud pool for reusability.

Once a virtual machine is destroyed, all the resources used by the virtual machine will be reclaimed by the system. This includes the virtual machine's IP address.



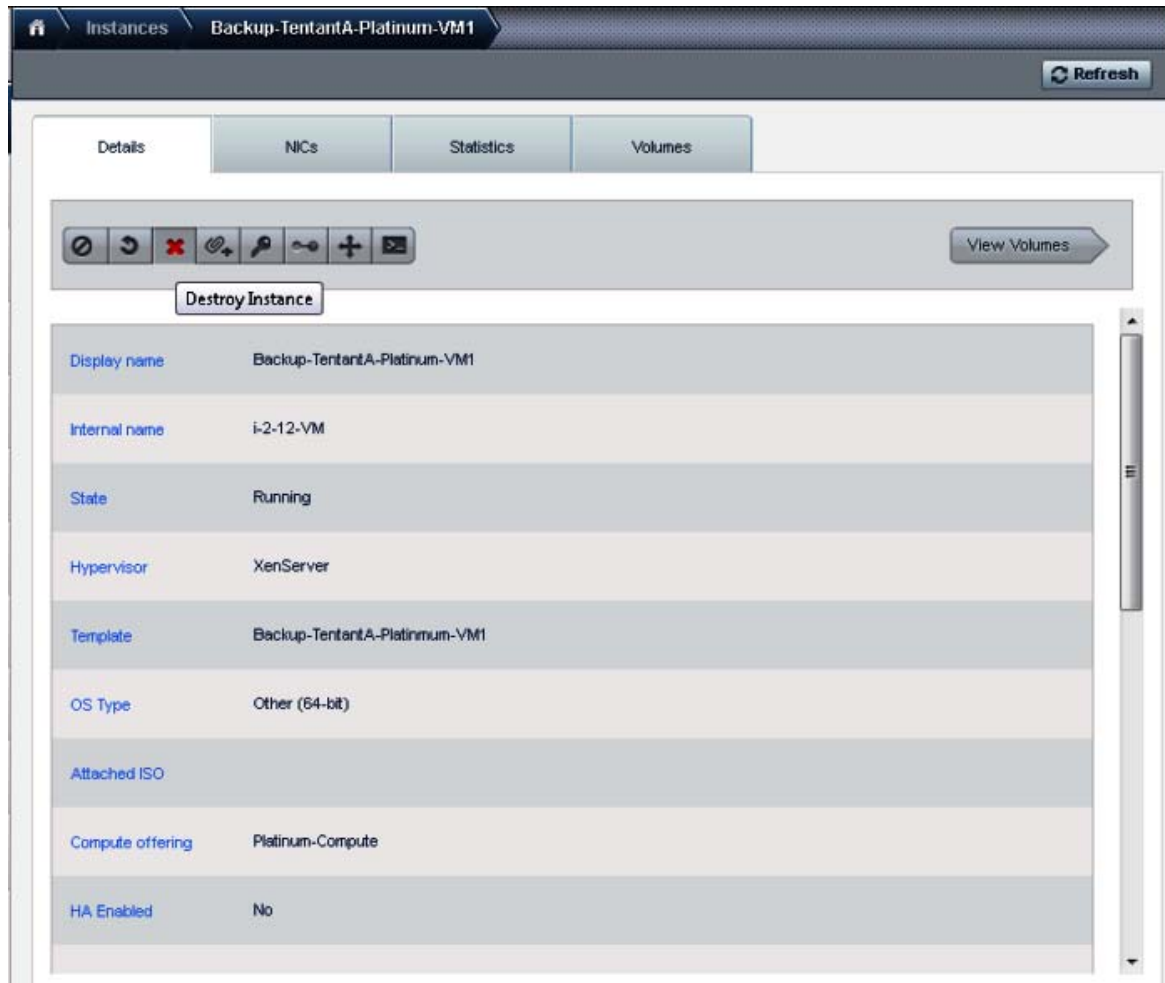
### Note

In this study we will delete Backup-TenantA-Platinum-VM1 using CloudPlatform.

Login to CloudPlatform with User credentials to create TenantA-Platinum-VM1 manual and automatic backup on TenantA Zone:

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Global Setting** tab under search type secstorage.allowed.internal.sites.
4. Under secstorage.allowed.internal.sites enter <<Var\_HTTP\_Server>> 10.65.121.705.
5. Restart CloudPlatform service.



**Figure 290**      *Destroying the Instance to check the Backup*

## Baremetal Host Cloud Configuration Design

This section outlines Citrix CloudPlatform 4.2.1 BareMetal host as a service offering which can be provided in private cloud environment. Multiple tenants can install LINUX based operating system on physical compute to run work loads which have performance and security requirement, CloudPlatform 4.2.1 provides multi-tenancy and life cycle management of these BareMetal hosts.

CloudPlatform 4.2.1 supports the kick start installation method for RPM-based Linux operating system on baremetal hosts in basic zones. Users can provision a baremetal host managed by CloudPlatform as long as they have the kick start file and corresponding OS installation ISO ready.

Kickstart installation eliminates manual intervention during OS installation. It uses a text file as a script to automate installation. The kickstart file contains responses to all the user input prompts that are displayed when you install an operating system. With kickstart installation, you can automate the installation of operating system software on large numbers of hosts.

## Limitation of Kick Start Baremetal Installation

1. Use in advanced zones is not supported. Use in basic zones only.
2. CloudPlatform storage concepts: primary storage, secondary storage, volume, snapshot.
3. System VMs: SSVM, CPVM, VR.
4. Template copy or template download.
5. VM migration.
6. Multiple NICs.
7. Using host tag for allocating host, capacity (cpu, memory) specifying in service offering.
8. A stopped VM (the OS running on host) can only start on the host it was most recently on.

The Citrix BaremetalCloud Design has been split into the following tasks:

1. Defining Network and Compute Offerings.
2. Defining Zones.
3. Defining Network.
4. Adding Pods.
5. Adding Cluster.
6. Adding Hosts.
7. Baremetal Host Deployment.

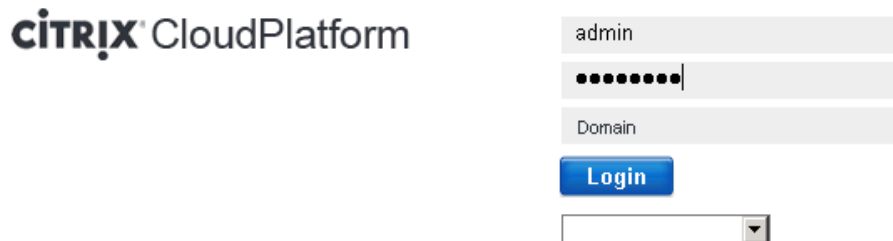
## Defining Network and Compute Offerings

A zone is the largest organizational unit within a CloudPlatform deployment. A zone typically corresponds to a single datacenter, although it is permissible to have multiple zones in a datacenter. The benefit of organizing cloud compute, network and storage infrastructure into zones is to provide physical isolation and redundancy.

To create zones using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, login to CloudPlatform with user credentials and follow these steps:

1. Type the User Name <root>, Password <XXXXXX>, and Domain.
2. Click **Login**.

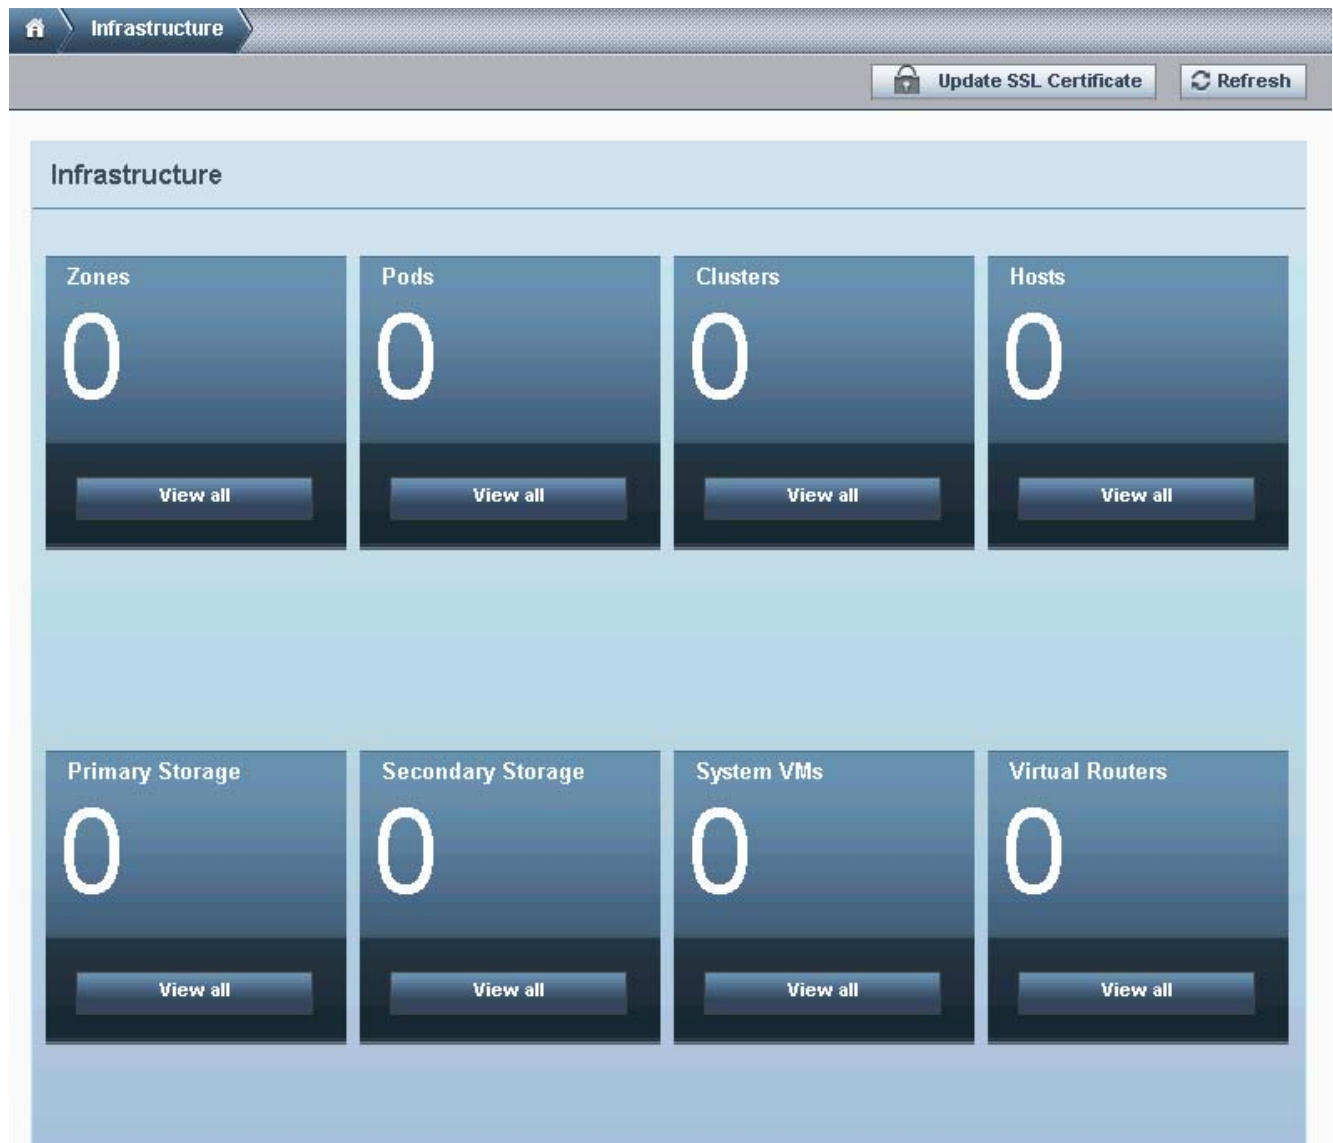
**Figure 291**      *Logging to the Citrix CloudPlatform*



3. Click **Agree** to confirm the End User License Agreement.

4. Click **I have used CloudPlatform before, skip this guide** button.
5. Click **Infrastructure** tab in the left pane.

**Figure 292**      *Displaying the Infrastructure Page*



6. Click **Service Offerings** in the left pane.
7. On the right pane, under Select Offering list box, select **Network Offering**.
8. Click **Add Network offering**.
9. Enter TenantA-BareMetal-Network in the Name field.
10. Enter Network for BareMetal Hosts in the Description field.
11. Select **Shared** in Guest Type list box.
12. Check **Specify VLAN** check box.
13. In Supported Services, select **DHCP** check box.

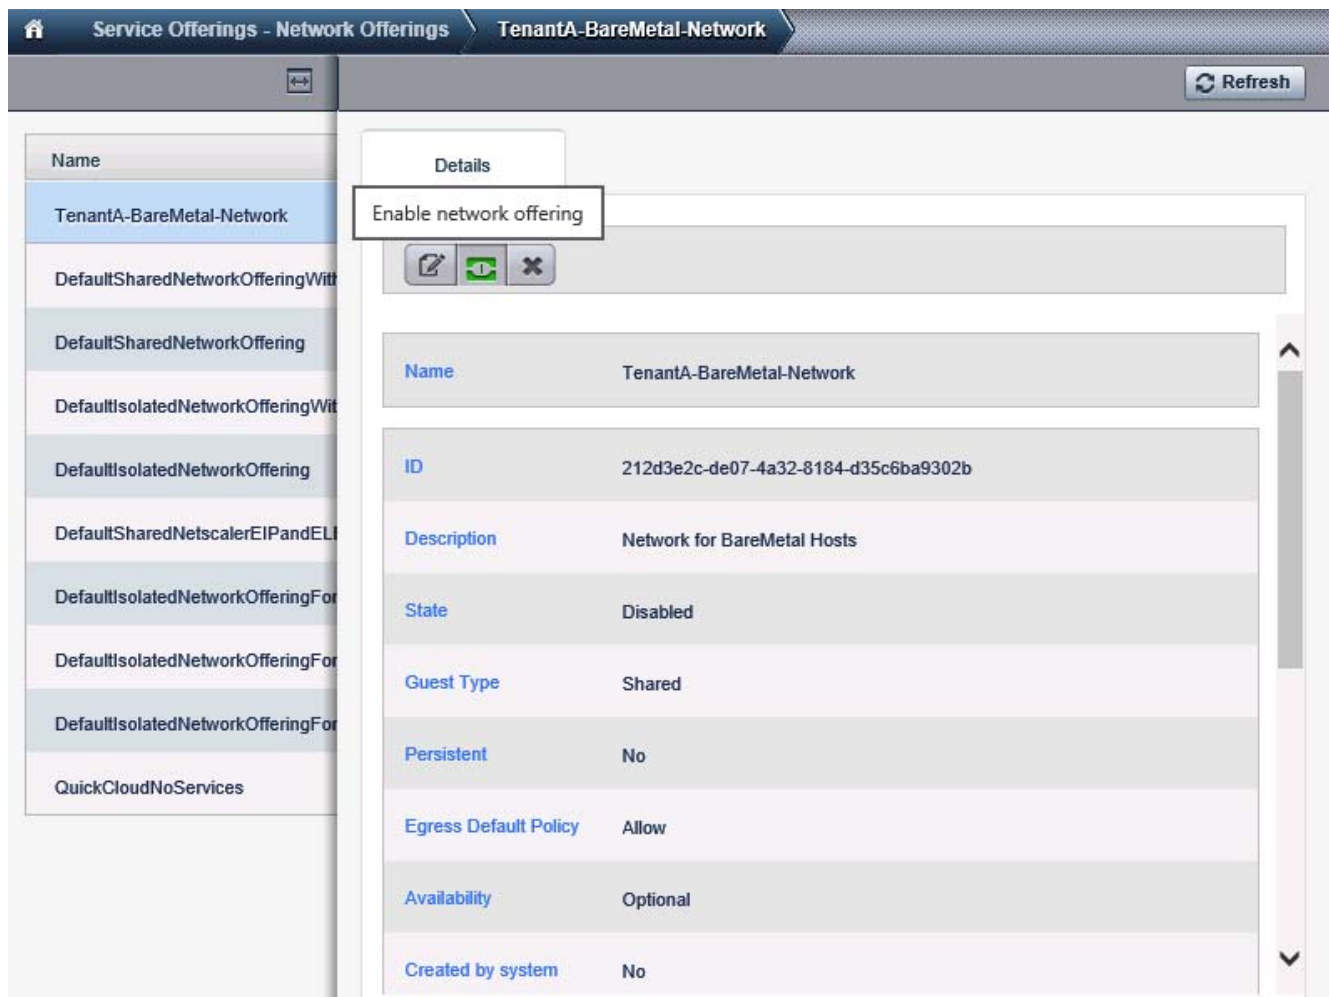
14. Select BaremetalDhcpProvider under DHCP Provider list box.
15. Select UserData check box.
16. Select BaremetalUserdataProvider under User Data Provider list box.
17. Check **BaremetalPxeService** check box.
18. Click **OK**.

**Figure 293**      *Configuring Baremetal Network Offering*

The screenshot shows the 'Add network offering' dialog box in the Citrix Cloud Platform interface. The dialog is titled 'Add network offering' and is open over the 'Service Offerings - Network Offerings' page. The 'Name' field is 'TenantA-BareMetal-Network' and the 'Description' is 'Network for BareMetal Hosts'. The 'Network Rate (Mb/s)' field is empty. The 'Guest Type' is set to 'Shared'. The 'Persistent' checkbox is unchecked. The 'Specify VLAN' checkbox is checked. In the 'Supported Services' section, 'VPN' is unchecked, 'DHCP' is checked, 'DHCP Provider' is set to 'Baremeta', 'DNS' is unchecked, 'Firewall' is unchecked, 'Load Balancer' is unchecked, 'User Data' is checked, and 'User Data' is set to 'Baremeta'. The 'Conserve mode' checkbox is unchecked. The 'Tags' field is empty. The 'Cancel' and 'OK' buttons are at the bottom of the dialog.

19. Select TenantA-BareMetal-Network and Enable Network.

**Figure 294**      **Enable TenantA-BareMetal-Network offering**



20. Click **Service Offerings** in the left pane.
21. On right pane, under Select Offering list box, select **Compute Offering**.
22. Click **Add Compute offering**.
23. Enter TenantA-BareMetal-Compute in the Name field.
24. Enter Compute Offering in the Description field.
25. Enter Shared in Storage Type field.
26. Enter <Physical\_Blade\_CPU\_Core> 16 in CPU Cores field.
27. Enter <Physical\_Blade\_CPU\_Mhz> 1333 in CPU Mhz field.
28. Enter <Physical\_Blade\_Memory> 212992 in Memory (MB) field.
29. Click **OK**.

Figure 295 Configuring Baremetal Compute Offering

**+ Add compute offering**

\* Name:  \* Description:

Storage Type:  # of CPU Cores:

\* CPU (in MHz):  \* Memory (in MB):

Network Rate (Mb/s):  Disk Read Rate (BPS):

Disk Write Rate (BPS):  Disk Read Rate (IOPS):

Disk Write Rate (IOPS):  Offer HA: ☐

Storage Tags:  Host Tags:

CPU Cap: ☐ Public: ☐

isVolatile: ☐ Deployment Planner:

Planner Mode:

### Defining Zones

1. With **Infrastructure** tab selected in the left pane, click **View all** on **Zones** in the right pane.
2. Click **+ Add Zone**.

Figure 296 Displaying the Infrastructure Page

**Infrastructure** **Zones**

| Zone            | Network Type | Public | Allocation State | Actions |
|-----------------|--------------|--------|------------------|---------|
| No data to show |              |        |                  |         |

3. Click **Basic zone type** radio button.
4. Click **Next**.

**Figure 297**      **Defining the Basic Zone Configuration**

**+ Add zone**

1 **Zone Type** > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

**Set up zone type**  
Please select a configuration for your zone.

☒ **Basic**  
Provide a single network where each VM instance is assigned an IP directly from the network. Guest isolation can be provided through layer-3 means such as security groups (IP address source filtering).

☐ **Advanced**  
For more sophisticated network topologies. This network model provides the most flexibility in defining guest networks and providing custom network offerings such as firewall, VPN, or load balancer support.

**Isolation Mode**

☐ **Security Groups**  
Choose this if you wish to use security groups to provide guest VM isolation.

Cancel **Next**

5. Enter TenantA-BareMetal-Zone in the Name field.
6. Enter <Public\_DNS\_IP\_Address>72.163.128.140 in DNS1 field.
7. Enter 171.70.168.183 in Internal DNS 1 field.
8. In Hypervisor select BareMetal in list box.
9. In Network Offering select TenantA-BareMetal-Network.
10. Click **Next**.



**Figure 298**      *Defining the Zone Attributes*

**Add zone**

1 Zone Type    2 **Setup Zone**    3 Setup Network    4 Add Resources    5 Launch

A zone is the largest organizational unit in CloudPlatform™, and it typically corresponds to a single datacenter. Zones provide physical isolation and redundancy. A zone consists of one or more pods (each of which contains hosts and primary storage servers) and a secondary storage server which is shared by all pods in the zone.

\* Name:

\* IPv4 DNS1:

IPv4 DNS2:

\* Internal DNS 1:

Internal DNS 2:

\* Hypervisor:

Network Offering:

Previous      Cancel      Next

## Defining Network

After defining zone the next step is to design and define cloud network. This solution design implements the basic networking, since Baremetal hosts can only be added to zones with basic network support. We will have one physical network defined in a zone which will carry Guest and Management traffic types.

- Management Traffic - is generated within the CloudPlatform as the internal resources communicate with each other. This includes communication between hosts
- Guest Private Traffic - is used and generated by Baremetal hosts. The guest network results when Baremetal hosts communicate with each other over a network.

To configure the zone basic network using Citrix CloudPlatform application dedicated for TenantA-BareMetal multi-tenant to host cloud services, follow these steps:

1. Drag guest Public and Private traffic and Storage Types icon to Physical Network 2.
2. Rename Physical network Network 1 TenantA-BareMetal-Network in Physical network name field.



3. Click **Next**.

**Figure 299** Defining the Basic Networking Configuration for the Zone



## Adding Pods

Pod is the second-largest organizational unit within a CloudPlatform deployment. Pods are contained within zones. Each zone can contain one or more pods. A pod consists of one or more clusters of hosts and one or more primary storage servers.

To add pod network in a zone advanced network configuration using Citrix CloudPlatform application dedicated for TenantA-BareMetal multi-tenant to host cloud services, follow these steps:

1. Enter TenantA-BareMetal-Pod in Pod Name field.
2. Enter <Management\_VLAN\_602\_GW\_IP\_Address>10.65.121.1 in Reserved System Gateway field.
3. Enter 255.255.255.0 in Reserved System netmask field.

4. Enter 10.65.121.190 in Start Reserved system IP field.
5. Enter 10.65.121.200 in End Reserved system IP field.
6. Click **Next**.

**Figure 300**      *Adding the Pod to the Zone*

**Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• POD > • GUEST TRAFFIC > • STORAGE TRAFFIC >

Each zone must contain in one or more pods, and we will add the first pod now. A pod contains hosts and primary storage servers, which you will add in a later step. First, configure a range of reserved IP addresses for CloudPlatform™'s internal management traffic. The reserved IP range must be unique for each zone in the cloud.

\* Pod name:

\* Reserved system gateway:

\* Reserved system netmask:

\* Start Reserved system IP:

End Reserved system IP:

## Adding Guest Traffic

To add guest network in a zone basic network configuration using Citrix CloudPlatform application dedicated for TenantA-BareMetal multi-tenant to host cloud services, follow these steps:

1. Enter <guest\_VLAN20\_GW\_IP\_Address>20.1.1.1 in Gateway field.
2. Enter 255.255.255.0 in netmask field.
3. Enter 20.1.1.20 in Start IP field.
4. Enter 20.1.1.30 in End IP field.
5. Click **Next**.

**Figure 301**      **Setting the Guest VLAN Range**

**+ Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• POD > • **GUEST TRAFFIC >** • STORAGE TRAFFIC >

Guest network traffic is communication between end-user virtual machines. Specify a range of IP addresses that CloudPlatform™ can assign to guest VMs. Make sure this range does not overlap the reserved system IP range.

Guest Gateway: 20.1.1.1

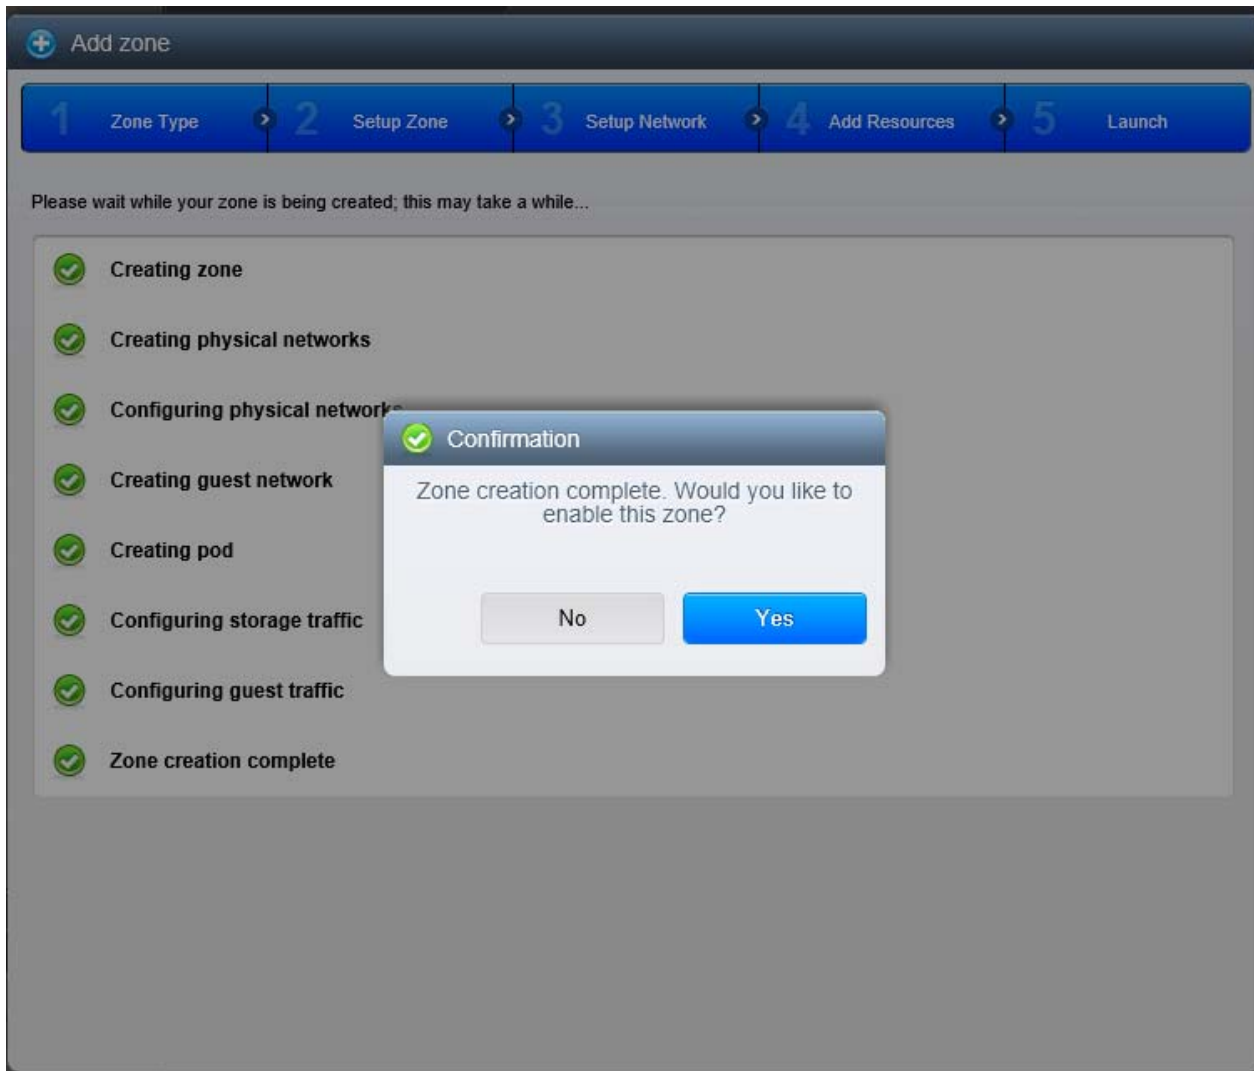
Guest Netmask: 255.255.255.0

Guest start IP: 20.1.1.20

Guest end IP: 20.1.1.30

Previous Cancel Next

6. Click **Yes** to Zone creation.

**Figure 302**      **Zone creation Confirmation**

## Adding Cisco UCS Manager Plugin

You can provision Cisco UCS server blades into CloudPlatform for use as bare metal hosts. The goal is to enable easy expansion of the cloud by leveraging the programmability of the UCS converged infrastructure and CloudPlatform's knowledge of the cloud architecture and ability to orchestrate. CloudPlatform can automatically understand the UCS environment, server profiles and so on such that CloudPlatform administrators can deploy a bare metal OS on a Cisco UCS.

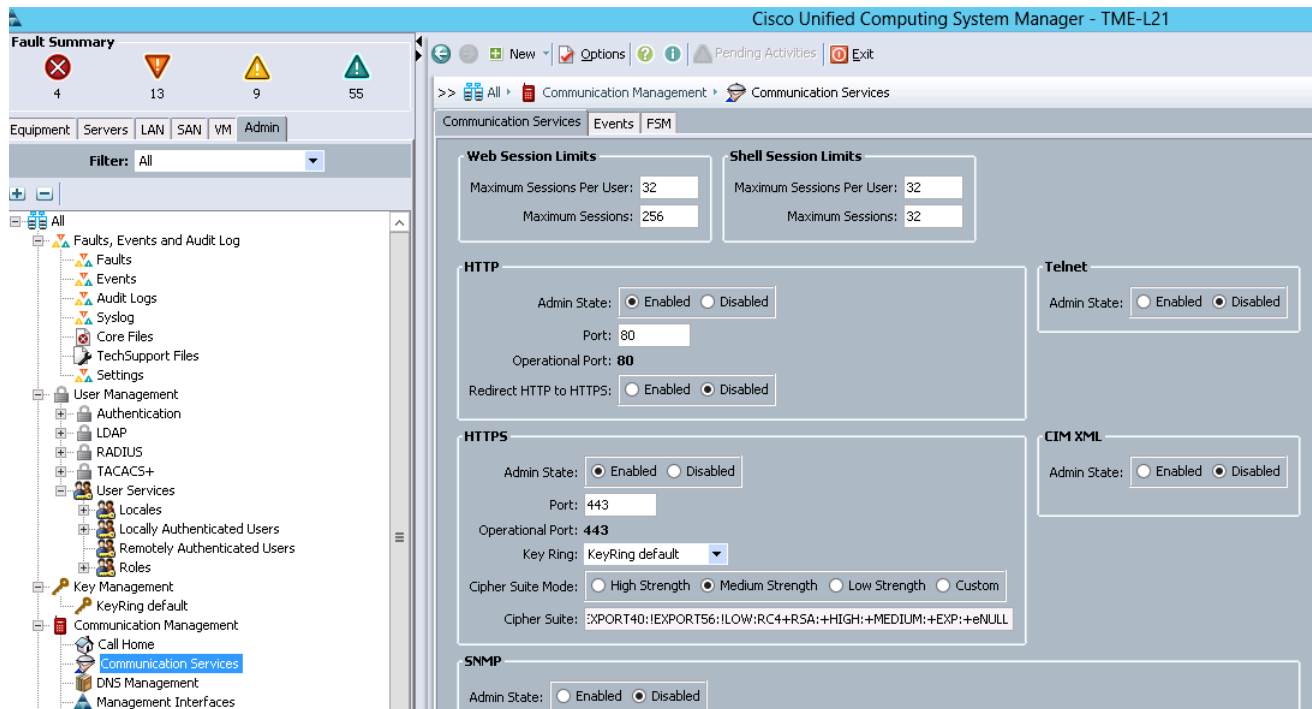
In this study we will register Cisco UCS Manager to CloudPlatform using Cisco UCS Manager Plugin. From hardware inventory to display available blades and their status (Service Profile associated / disassociated).

The CloudPlatform admin now can choose blade which is free (unassociated), create Service Profile based on Service Profile Template already defined by Cisco UCS Manager which trigger Cisco UCS Manager API calls to create Service Profile based on Service Profile Template associate blade.

On successfully association of blade CloudPlatform performs Baremetal as a Service provisioning by installing RHEL 6.3 Operating System using PXE Boot.

1. The CloudPlatform needs HTTP and HTTPS protocol to be enabled for accessing Cisco UCS Manager, so make sure Cisco UCS Manager Communication Services for HTTP Admin State is **<enabled>** & Redirect HTTP to HTTPS is **<disabled>**.

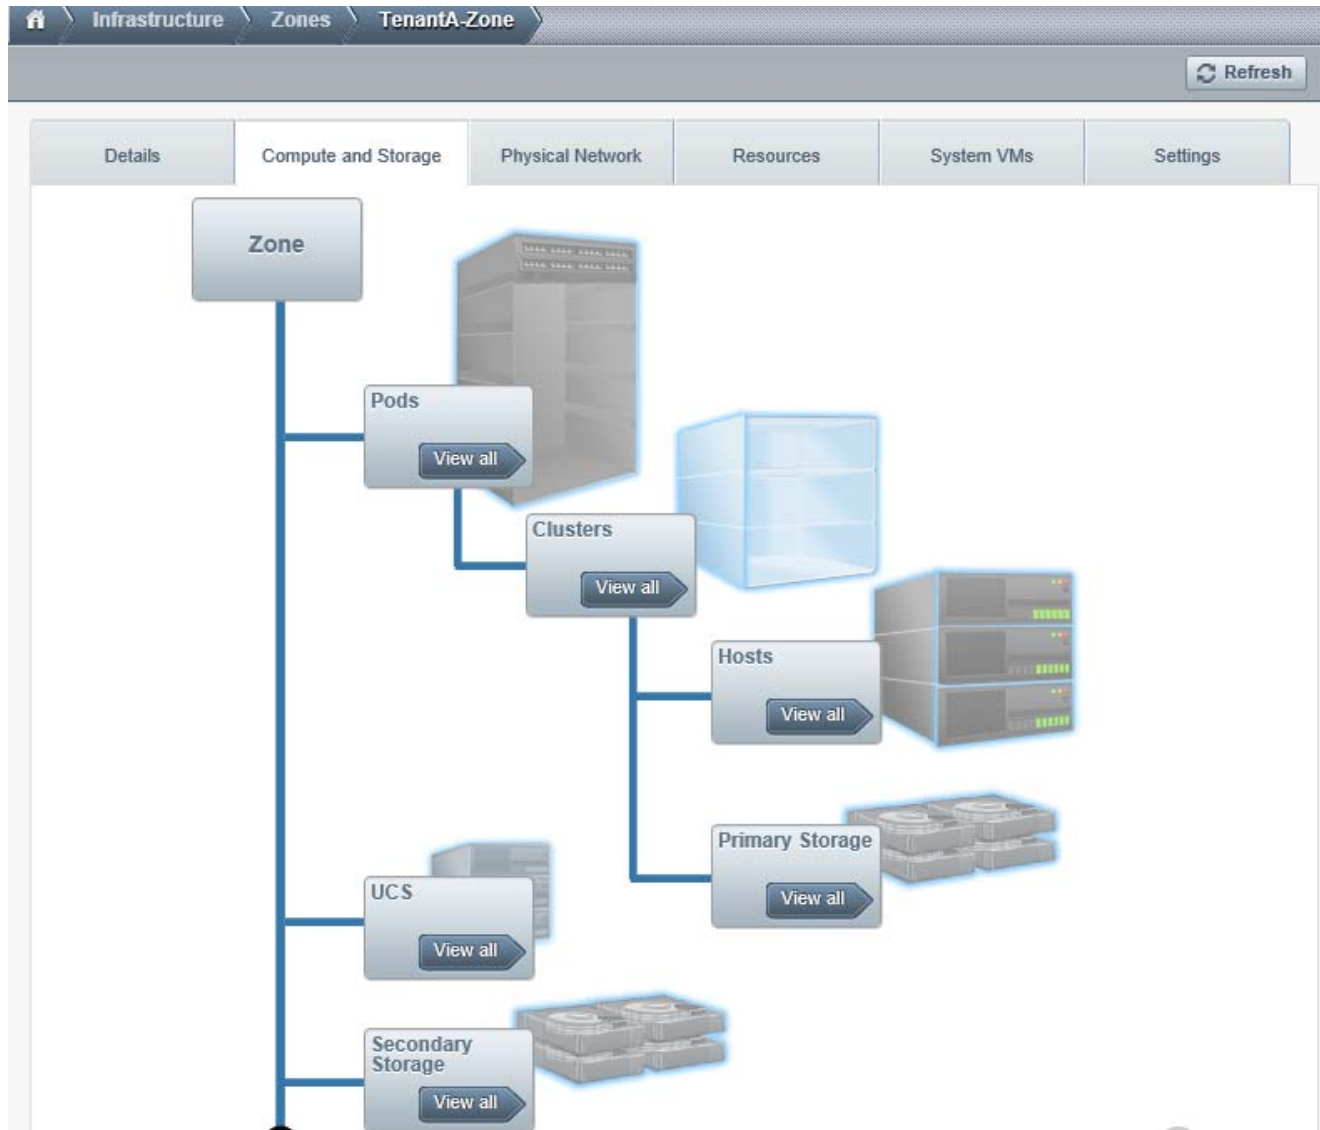
**Figure 303 Cisco UCS Manager Communication Services HTTP configuration details**



Login to CloudPlatform with User credentials to configure Cisco UCS Manager Plugin on the zone TenantA-BareMetal-Zone:

1. Click **Login**.
2. Click the **Infrastructure** tab on the left pane.
3. Click the Zones **View All**.
4. Click **<TenatA-BareMetal-Zone> Zone**.
5. Click **Compute and Storage**.
6. Click **UCS View All**.

Figure 304 Adding Cisco UCS Manager details



7. Click **Add Cisco UCS Manager**.
8. Enter Name <Cisco UCS Manager-Cloud>.
9. Enter URL: <10.65.121.14> UCS Cluster IP Address.



**Note** Enter only IP address.

10. Enter UserName <admin>.
11. Enter Password <xxxxx> .
12. Click **OK**.

**Figure 305**      *Displaying the Cisco UCS Manager Plugin Configuration*

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

| Field      | Value        |
|------------|--------------|
| Name       | UCSM-Cloud   |
| * IP       | 10.65.121.14 |
| * Username | admin        |
| * Password | ••••••••     |

13. Click **Cisco UCS Manager-Cloud**.

**Figure 306**      *Access Cisco UCS Manager Plugin*

| Name       | URL          |
|------------|--------------|
| UCSM-Cloud | 10.65.121.14 |

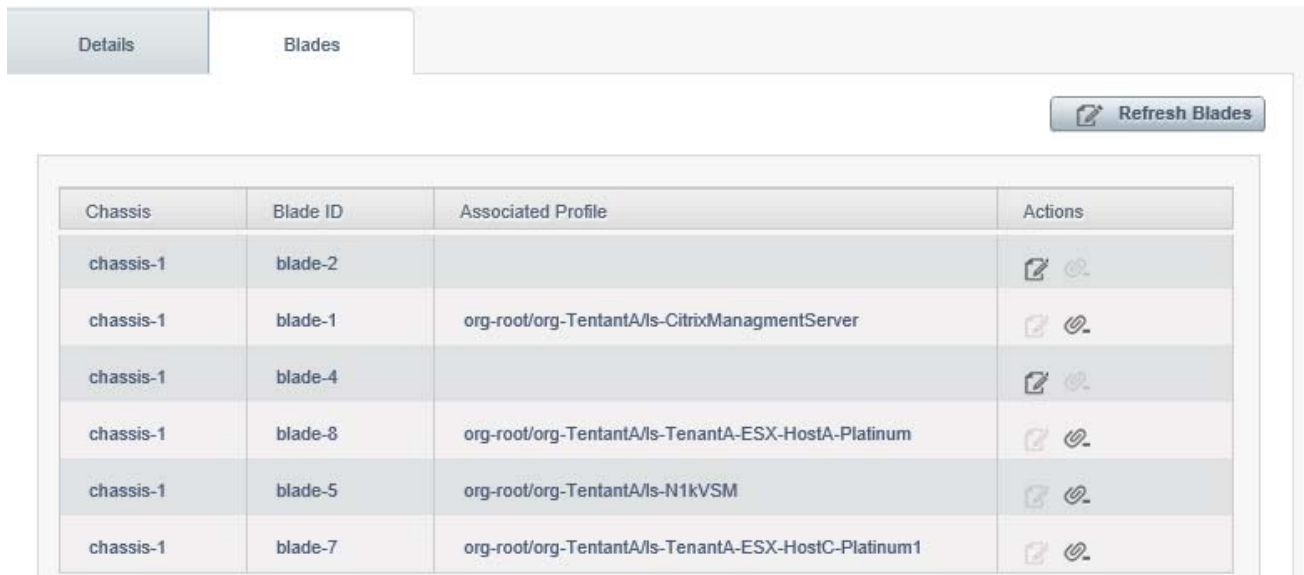
14. Click **Blades**.
15. The current CloudPlatform 4.2.1 displays all Blades which are associated with Service Profile or available for association with Service Profile to provision Baremetal host. CloudPlatform 4.2.1 allows associating freely available Blade by creating Service Profile derived using Service Profile Template created on Cisco UCS Manager.



**Note**

In current CloudPlatform 4.2.1 version the Cisco UCS Manager Plugin allows blade association with Service Profile derived by Service Profile Template created only in Root Org on Cisco UCS Manager, so make sure you have created Service Profile Templates under Root Org for Cisco UCS Manager Plugin to expose in CloudPlatform Manager and it is advised to Refresh Blades before starting Baremetal provisioning.

**Cisco UCS Manager displaying physical Blades available inventory**



- 16. On Blade-2 Under Action, click on Associate Profile to Blade.**

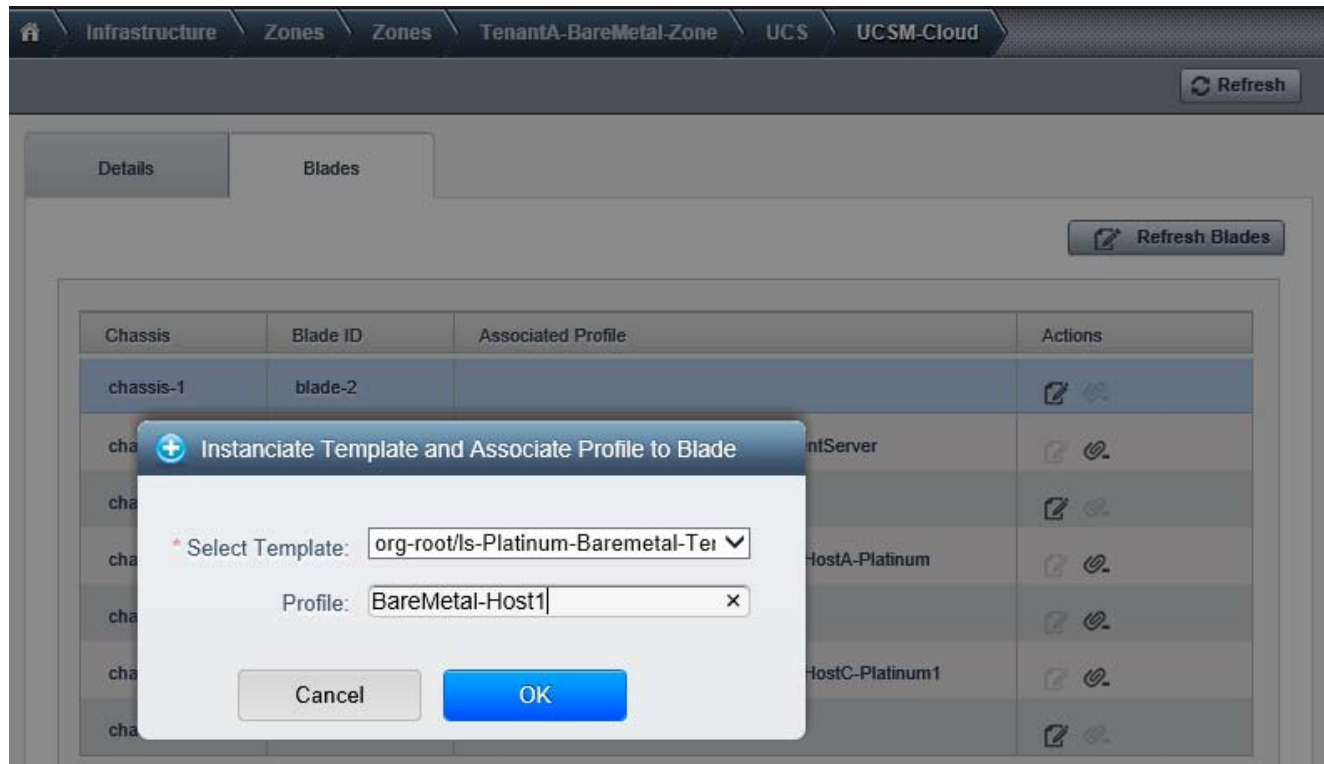
**Select blade -2 to associate with Service Profile in Cisco UCS Manager**



17. Select **Platinum-Baremetal-Template** in Select Template list box.
18. Enter **BareMetal-Host1** in Profile Name field.
19. Click **OK**.



**Figure 309**      **Select Platinum-Baremetal-Template to create BareMetal-Host1 Service Profile to associate blade-2**



20. After Service profile BareMetal-Host1 is associated with Blade 2, [Figure 310](#) shows Cisco UCS Manager and CloudPlatform mapping configuration.

**Figure 310** Cisco UCS Manager and CloudPlatform displaying BareMetal-Host1 Service Profile associated with blade-2

The figure consists of two screenshots. The top screenshot is a table from Cisco UCS Manager showing the association between chassis, blade IDs, and service profiles. The bottom screenshot is the CloudPlatform interface showing the details of the BareMetal-Host1 service profile.

| Chassis   | Blade ID | Assoc                                               | Actions |
|-----------|----------|-----------------------------------------------------|---------|
| chassis-1 | blade-2  | org-root/ls-BareMetal-Host1                         |         |
| chassis-1 | blade-1  | org-root/org-TenantA/ls-CitrixManagementServer      |         |
| chassis-1 | blade-4  |                                                     |         |
| chassis-1 | blade-8  | org-root/org-TenantA/ls-TenantA-ESX-HostA-Platinum  |         |
| chassis-1 | blade-5  | org-root/org-TenantA/ls-N1kVSM                      |         |
| chassis-1 | blade-7  | org-root/org-TenantA/ls-TenantA-ESX-HostC-Platinum1 |         |

The bottom screenshot shows the CloudPlatform interface with the following details:

- General Tab:**
  - Fault Summary:** 0 Critical, 4 Warning, 0 Error, 3 Info.
  - Status:** Overall Status: Up **Ok**.
  - Status Details:**
    - Desired Power State: Up
    - Assoc State: Associated
    - Assigned State: Assigned
  - Note:** The "Desired Power State" is the Power State of the server set via UCSM. It may be therefore different from the actual value. For the actual server power state click the "Server Details" Tab.
- Properties Tab:**
  - Name:** BareMetal-Host1
  - User Label:** [Empty field]
  - Description:** [Empty field]
  - Owner:** Local
  - UUID:** Hardware Default
  - UUID Pool:** [Empty field]
  - UUID Pool Instance:** org-root/uuid-pool-Platinum-Compute-UUID
  - Associated Server:** sys/chassis-1/blade-2
  - Service Profile Template:** [Empty field]
  - Template Instance:** [Empty field]
  - Assigned Server or Server Pool:** [Dropdown menu]
  - Management IP Address:** [Dropdown menu]
  - Maintenance Policy:** [Dropdown menu]

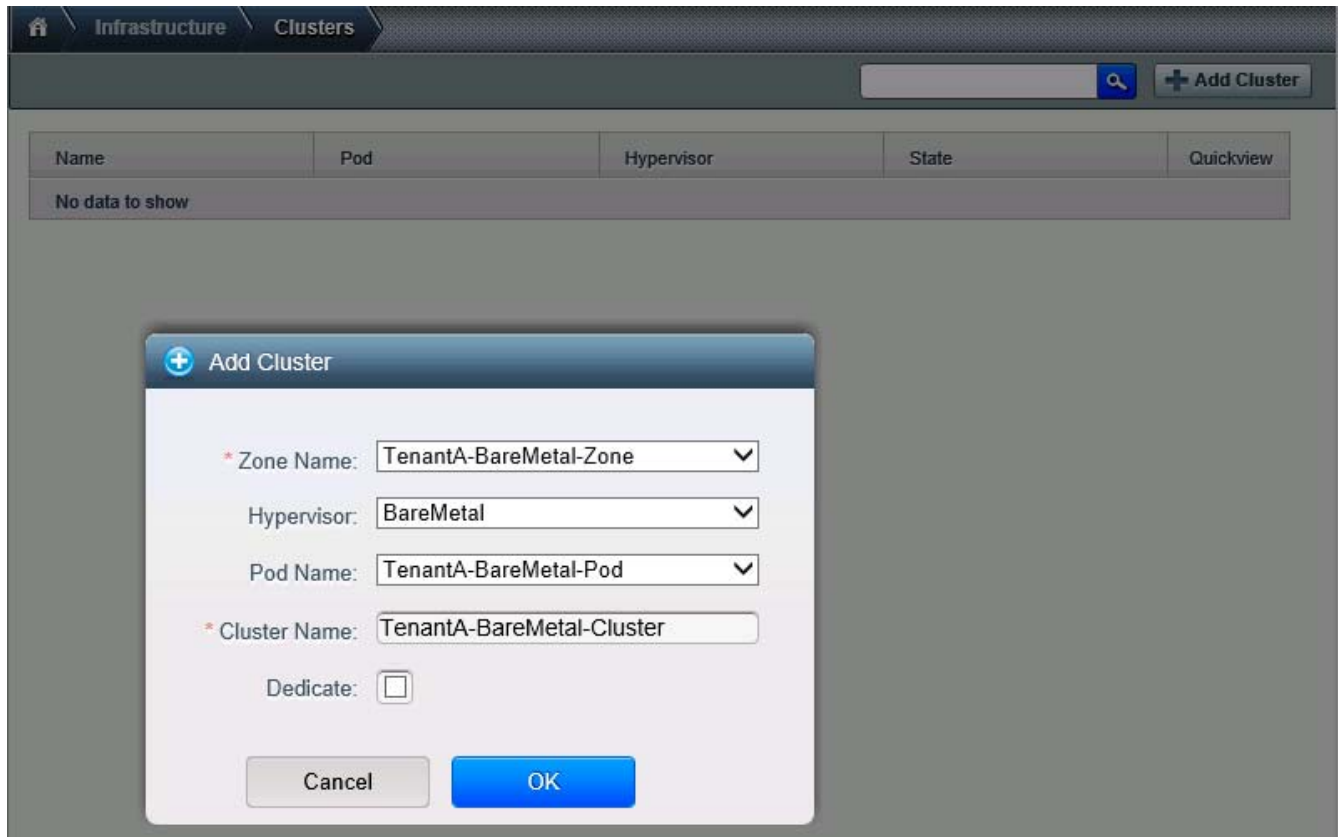
## Adding Cluster

Cluster provides a way to group hosts. Here we to be precise, we have cluster with set of Baremetal hosts part of cluster. The hosts in a cluster all have identical hardware, run the same Baremetal operating system, are on the same subnet, and access the same shared primary storage.

A cluster is the third-largest organizational unit within a CloudPlatform deployment. Clusters are contained within pods, and pods are contained within zones. Size of the cluster is limited by the underlying hypervisor, although the CloudPlatform recommends less in most cases. A cluster consists of one or more hosts and one or more primary storage servers.

To add cluster in zone using Citrix CloudPlatform application dedicated for TenantA multi-tenant to host cloud services, follow these steps:

1. Select <TenantA-BareMetal-Zone> in Zone Name list box.
2. Select <BareMetal> in Hypervisor list box.
3. Select <TenantA-BareMetal-Pod> in Pod Name list box.
4. Enter <TenantA-BareMetal-Cluster> in Cluster Name field.
5. Click **Next**.

**Figure 311**      **Adding the Cluster to the Zone**

## Adding Hosts

A host provides the computing resources that run the Baremetal Operating System. Each host has Baremetal linux based OS software installed on it to manage Applications. The Host here is an RHEL 6.3 Based server. The host is the smallest organizational unit within a CloudPlatform deployment. Hosts are contained within clusters, clusters are contained within pods, and pods are contained within zones.

To add host in cluster using Citrix CloudPlatform application dedicated for TenantA-BareMetal multi-tenant to host cloud services, follow these steps:

1. Select <TenantA-BareMetal-Zone > in Zone Name list box.
2. Select <TenantA-BareMetal-Pod > in Pod Name list box.
3. Enter <TenantA-BareMetal-Cluster > in Cluster Name field.
4. Enter <UCS-IPMI-IP-Address> 10.65.121.57 in Host Name field.



### Note

The IP Address of Host Name is the CIMC IP Address assigned to physical blade by Cisco UCS Manager which is been added as a Baremetal host to the zone.

5. Enter <admin> in Username field.

**Note**

The admin Username is created in IMPI Policy in Cisco UCS Manager which is applied to Service Profile associated to blade part of Baremetal host.

6. Enter < XXXX > in Password field.

**Note**

The password is created in IMPI Policy in Cisco UCS Manager which is applied to Service Profile associated to blade part of Baremetal host.

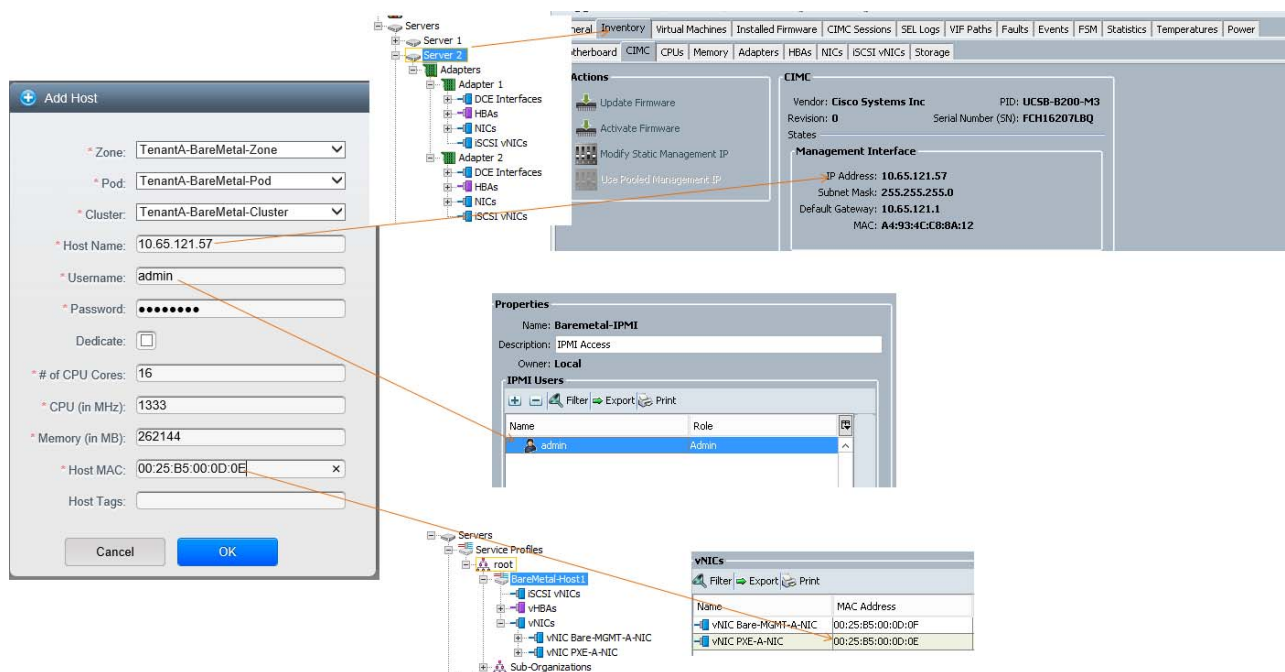
7. Uncheck **Dedicate** check box.
8. Enter <Physical\_Blade\_CPU\_Core> 16 in CPU Cores field.
9. Enter <Physical\_Blade\_CPU\_Mhz > 1333 in CPU (in MHz) field.
10. Enter <Physical\_Blade\_Memory > 262144 in Memory (MB) field.
11. Enter <00:25:b5:00:0D:0E > in Host MAC field.

**Note**

The MAC Address < 00:25:b5:00:0D:0E > is assigned to Static vNIC in Service Profile which handles PXE Boot operations.

12. Click OK.

**Figure 312** Adding the Host to the Zone



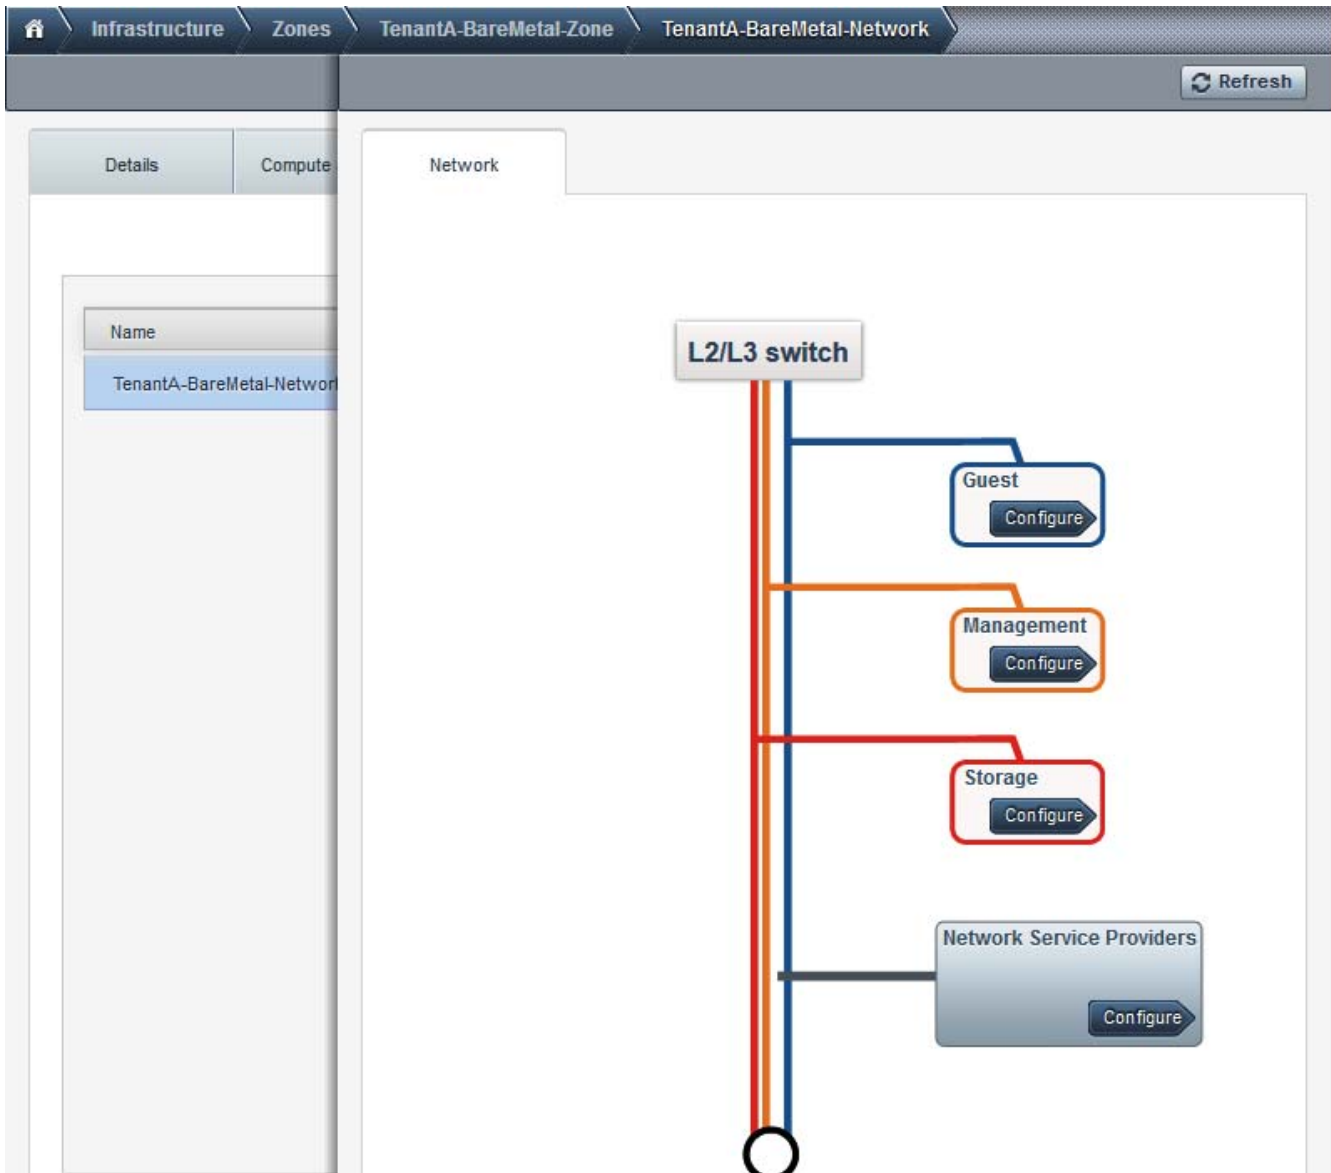
## Adding PXE and DHCP Servers

CloudPlatform 4.2.1 requires external PXE and DHCP servers to access to support Baremetal host PXE installation of Operating System on physical Server. The PXE and DHCP setup and configurations can be referred in Cloud Deployment section.

Login to CloudPlatform with User credentials to create PXE and DHCP servers on the zone TenantA-BareMetal-Zone:

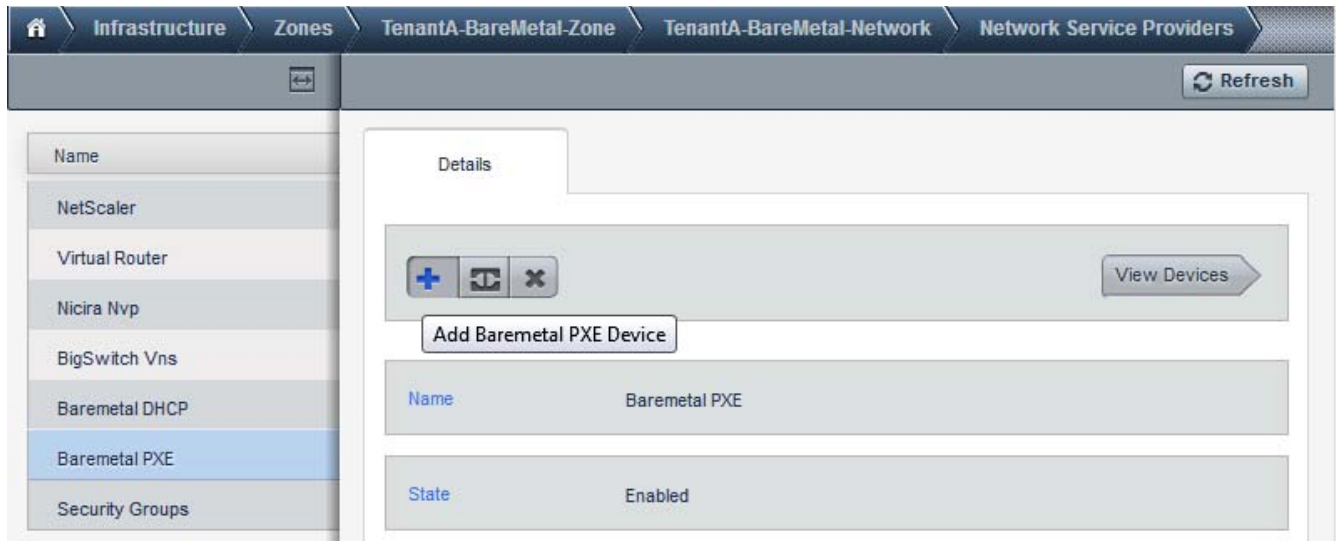
1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Infrastructure** tab.
4. Click **Zones View all**.
5. Click **TenantA-BareMetal-Zone**.
6. Click **Physical Network**.
7. Click **TenantA-BareMetal-Network**.
8. Click **Configure Network Service Providers**.

**Figure 313** DHCP and PXE Server configuration Details



9. Click **Baremetal PXE** tab.
10. Click + **ADD** icon to add Baremetal PXE Device.

**Figure 314** Add Baremetal PXE server configuration Details



11. Enter < BaremetalAgent\_PXE\_Server\_IP\_Address > http://20.1.1.5 in URL field.



**Note** IP Address 20.1.1.5 is assigned to server which has CloudPlatform 4.2.1 Baremetal agent and PXE services configured.


12. Enter < root > in Username field.  
 13. Enter < XXXXX > in Password field.  
 14. Enter < /export > in Tftp root directory.



**Note** /export path is where PXE and Baremetal OS image files are stored and NFS exported for more details refer to Cloud Deployment section.

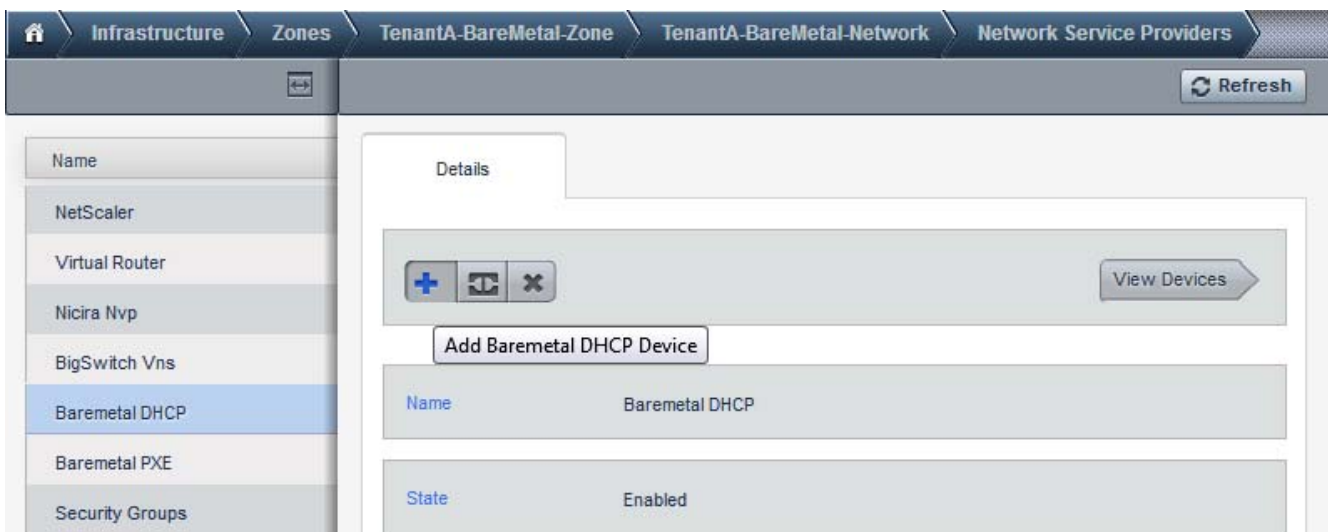
15. Click **OK**.

**Figure 315** Add Baremetal PXE server and Path configuration Details



16. Click **Baremetal DHCP** tab.
17. Click + **ADD** icon to add Baremetal DHCP Device.

**Figure 316** Add Baremetal DHCP server configuration Details



18. Enter <CloudPlatform\_Management\_IP\_Address> http://10.65.121.70 in URL field.






**Note** IP Address 10.65.121.70 is assigned to server which has DHCP services configured for more details refer to Cloud Deployment section.

19. Enter < **root** > in Username field.
20. Enter < **XXXXXX** > in Password field.
21. Click **OK**.

**Figure 317** Baremetal DHCP Server Configuration Details



## Baremetal Host Deployment

This section explains how to register Baremetal host template with Linux Operating System PXE Image and create Baremetal instance with Compute and Network offerings.

Bare metal hosts can run any of the following operating systems CentOS 5.5, CentOS 6.2, CentOS 6.3, Fedora 17, RedHat 6.1 – 6.3 and Ubuntu12.04

In this study we will install RHEL 6.3 operating system on newly created Baremetal instance BareMetal-Host1 using PXE method on FCoE NetApp Cluster Target.

Login to CloudPlatform with User credentials to create ISO Image TenantA Zone:

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Template** tab.

4. Select **Templates** in Select view list box.
5. Click **Register template**.
6. Enter **BareMetal-RHEL6-3-Template** in the Name field.
7. Enter **RHEL6-3 Image** in the Description field.
8. Enter URL Path  
ks=http://20.1.1.5/RHEL63/RHEL63.ks;kernel=20.1.1.5:/var/www/html/RHEL63/vmlinuz;initrd=20.1.1.5:/var/www/html/RHEL63/initrd.img

**Note**

Make sure you provide correct path for RHEL6-3 kickstart, kernel and initrd files. The directory /export and /var/www/html/RHEL63 where these files are stored is been NFS exported. For more details refer Cloud Deployment section.

9. Select **TenantA-BareMetal-Zone** in Zone list box.
10. Select **BareMetal** in Hypervisor list box.
11. Select **BareMetal** in Format list box.
12. Select **Red Hat Enterprise Linux 6.3 (64 bit)** in OS Type list box.
13. Check **Extractable** check box.
14. Check **Public** check box.
15. Check **Featured** check box.
16. Click **OK**.

**Figure 318**      *Displaying the Baremetal Template Details*

The screenshot shows a web-based interface for managing templates. A modal dialog titled "Register template" is displayed in the foreground. The background shows a table with columns: Name, Zone, Hypervisor, Order, and Quickview. The table currently contains no data, indicated by "No data to".

The "Register template" dialog contains the following fields and options:

- Name:** BareMetal-RHEL6-3-Template
- Description:** RHEL6-3 Image
- URL:** ks=http://20.1.1.5/RHEL63/RHEL63
- Zone:** TenantA-BareMetal-Zone
- Hypervisor:** BareMetal
- Format:** BareMetal
- OS Type:** Red Hat Enterprise Linux 6.3 (64-b
- Extractable:** ☒
- Password Enabled:** ☐
- Dynamically Scalable:** ☐
- Public:** ☒
- Featured:** ☒
- Routing:** ☐

At the bottom of the dialog are two buttons: "Cancel" and "OK".

17. Click the **Instances** tab.
18. Click **Add Instance**.
19. Select **TenantA-BareMetal-Zone** in Select a Zone list box.
20. Select **Template** radio button.
21. Click **Next**.

**Figure 319**      **Selecting Template**

**Add Instance**

1 Setup   2 Select a template   3 Compute offering   4 Data Disk Offering   5 Affinity   6 Network   7 Review

**Select a zone**  
A zone typically corresponds to a single datacenter. Multiple zones help make the cloud more reliable by providing physical isolation and redundancy.

TenantA-BareMetal-Zone

**Select ISO or template**

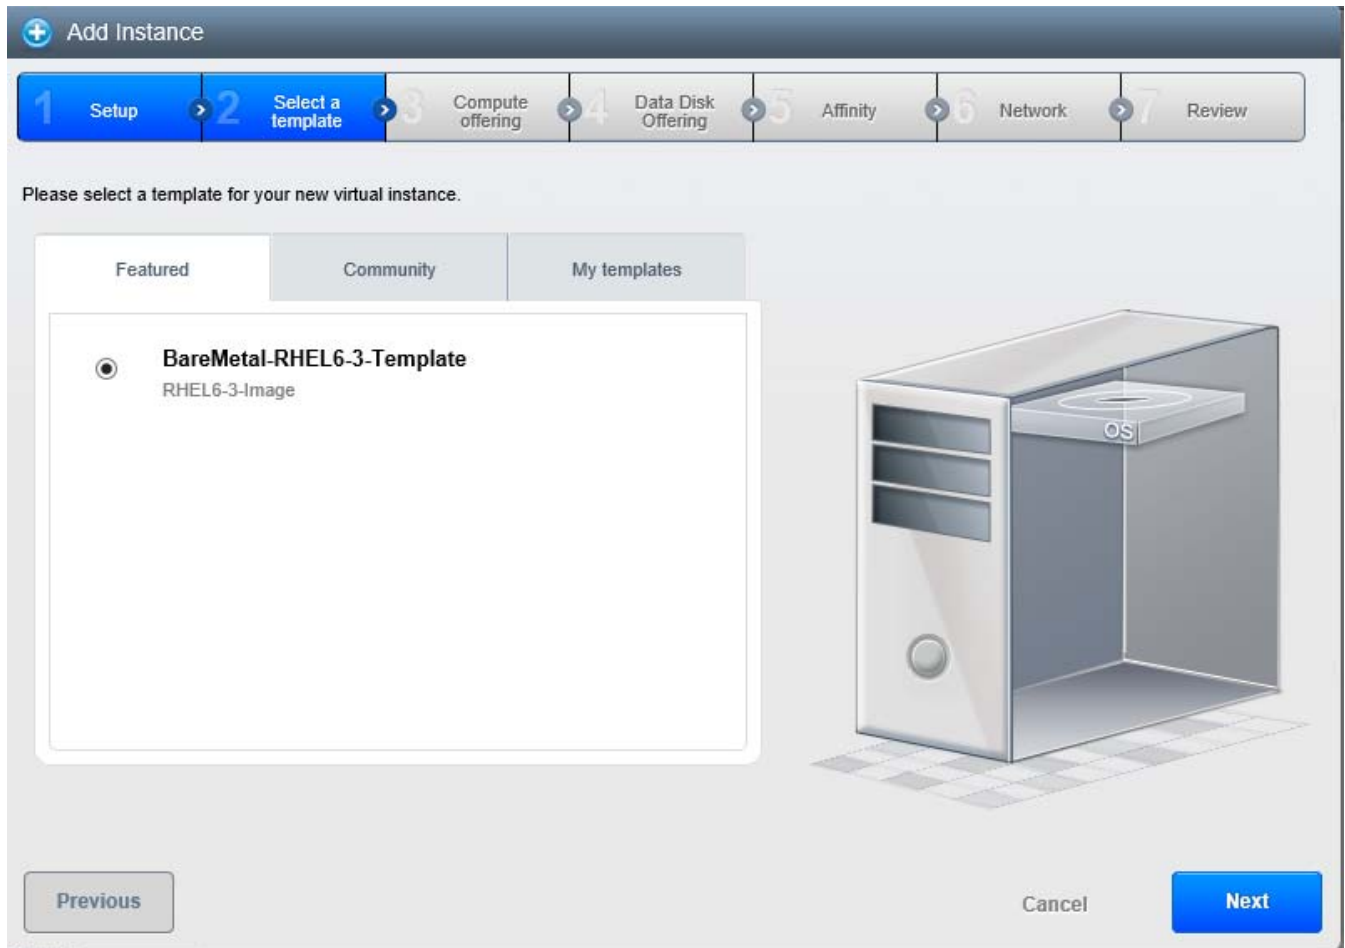
☒ **Template**   OS image that can be used to boot VMs

☐ **ISO**   Disc image containing data or bootable media for OS

Cancel   **Next**

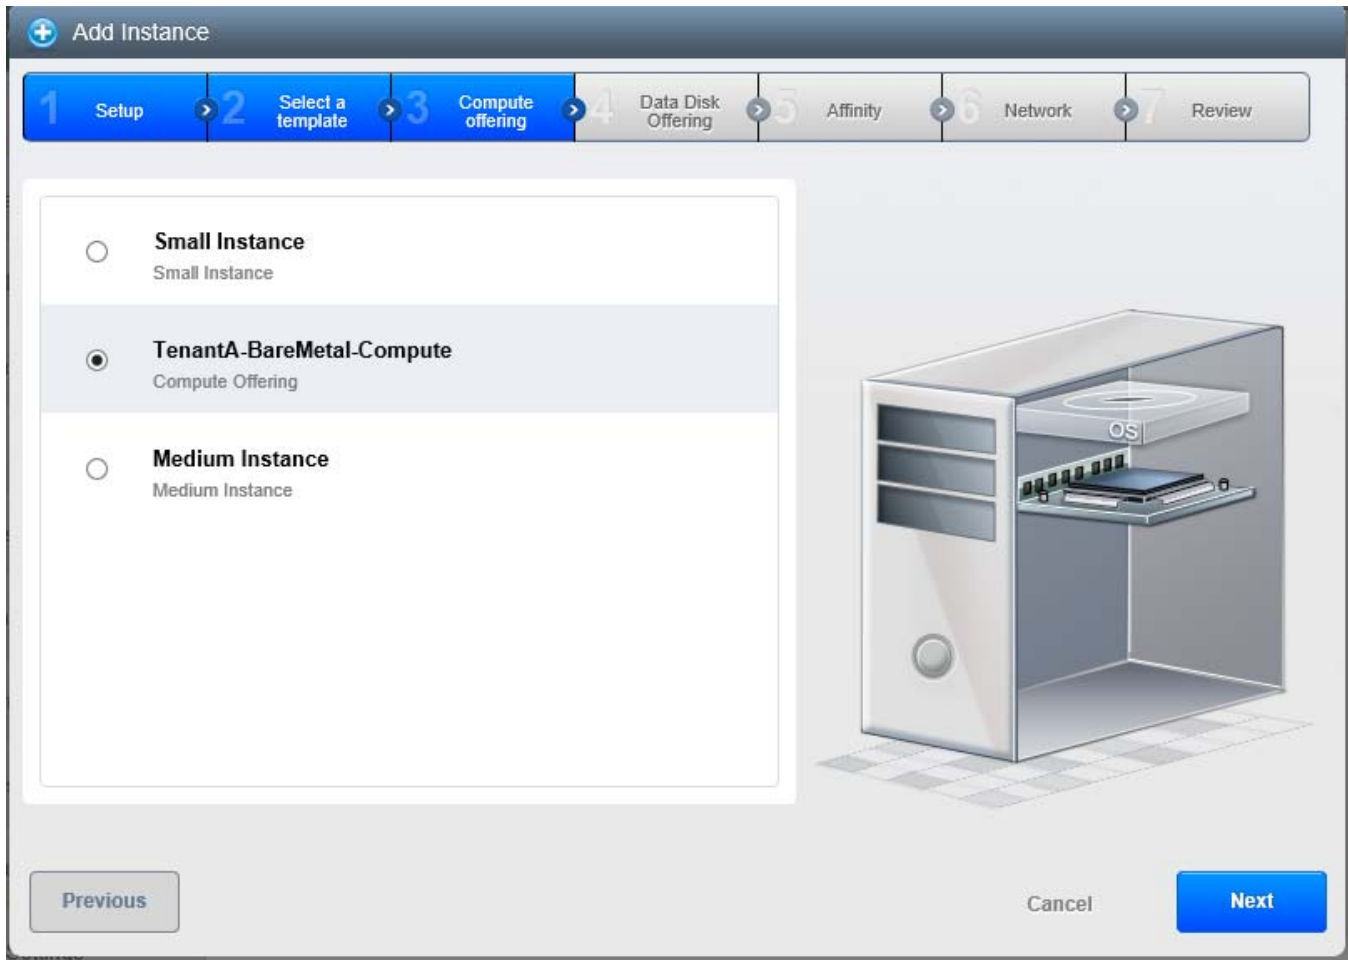
22. Enter **BareMetal-RHEL6-3-Template** in the Name field.
23. Enter **RHEL6-3 Image** in the Description field.
24. Select **BareMetal-RHEL6-3Template** in template.
25. Click **Next**.

**Figure 320**      **Selecting Template Image**



26. Select **TenantA-BareMetal-Compute** in Compute offering.
27. Click **Next**.

**Figure 321**      *Selecting Compute Offering*



28. Select **Medium** in Data Disk offering.
29. Click **Next**.

**Figure 322**      **Selecting Disk Offering**

**Add Instance**

1 Setup > 2 Select a template > 3 Compute offering > **4 Data Disk Offering** > 5 Affinity > 6 Network > 7 Review

☐ No thanks

☐ **Small**  
Small Disk, 5 GB

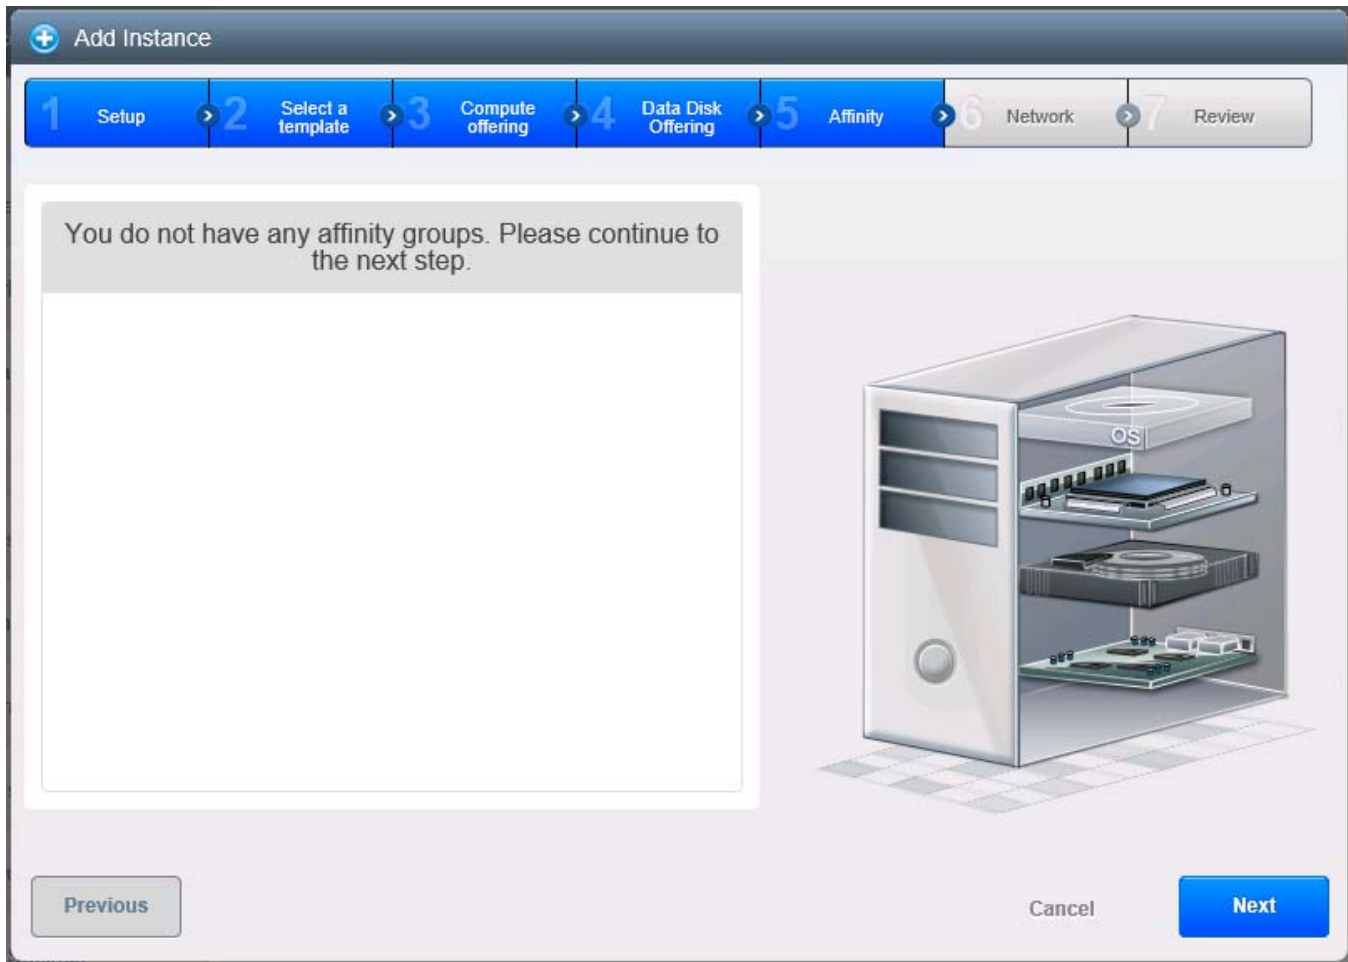
☒ **Medium**  
Medium Disk, 20 GB

☐ **Large**  
Large Disk, 100 GB

☐ **Custom**  
Custom Disk

Previous Cancel **Next**

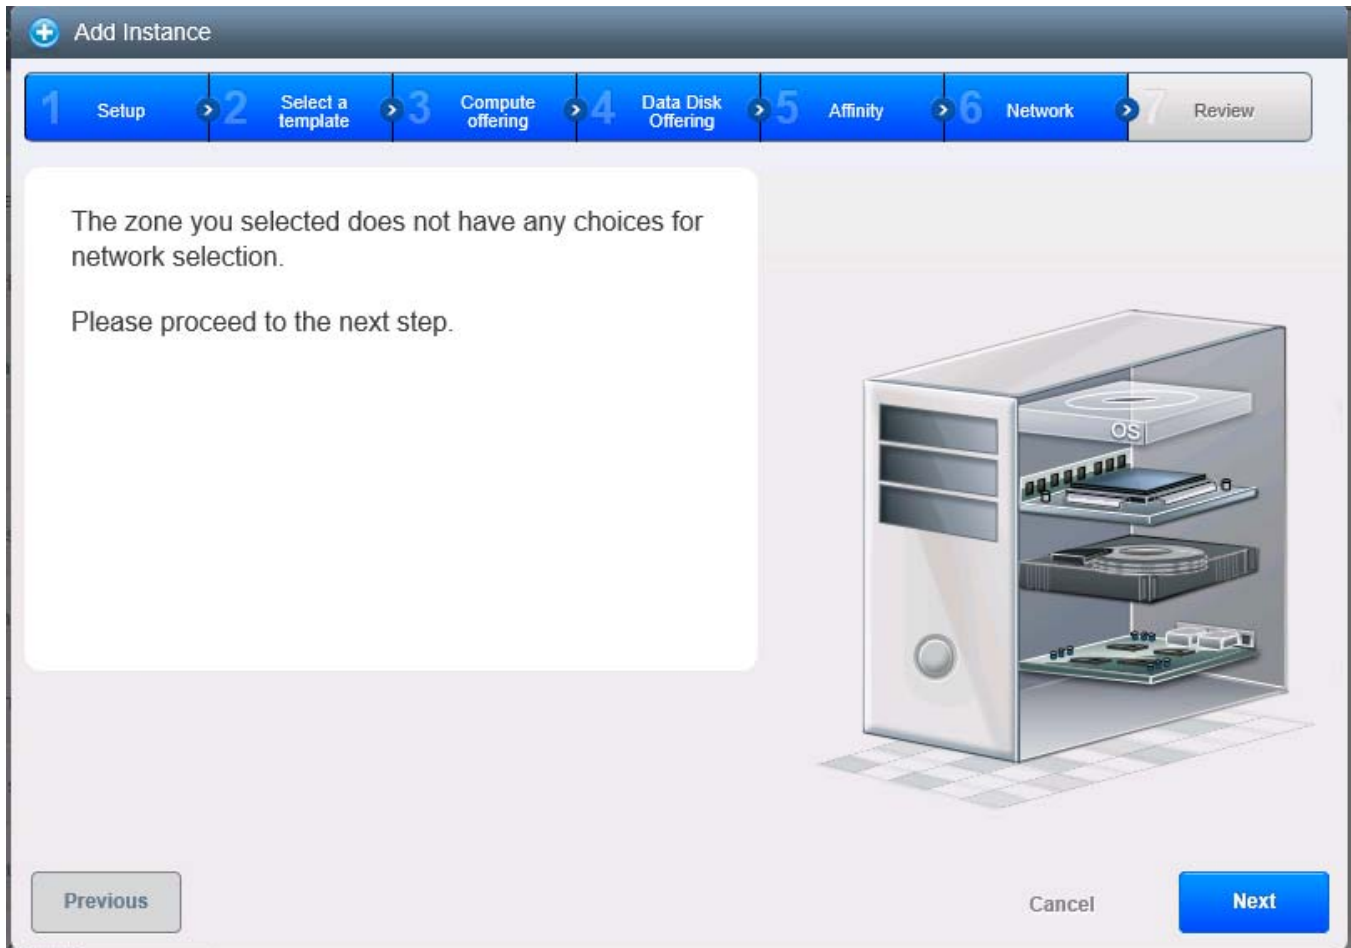
30. Click Next.

**Figure 323**      *No Affinity Rules Configured*

31. Click Next.

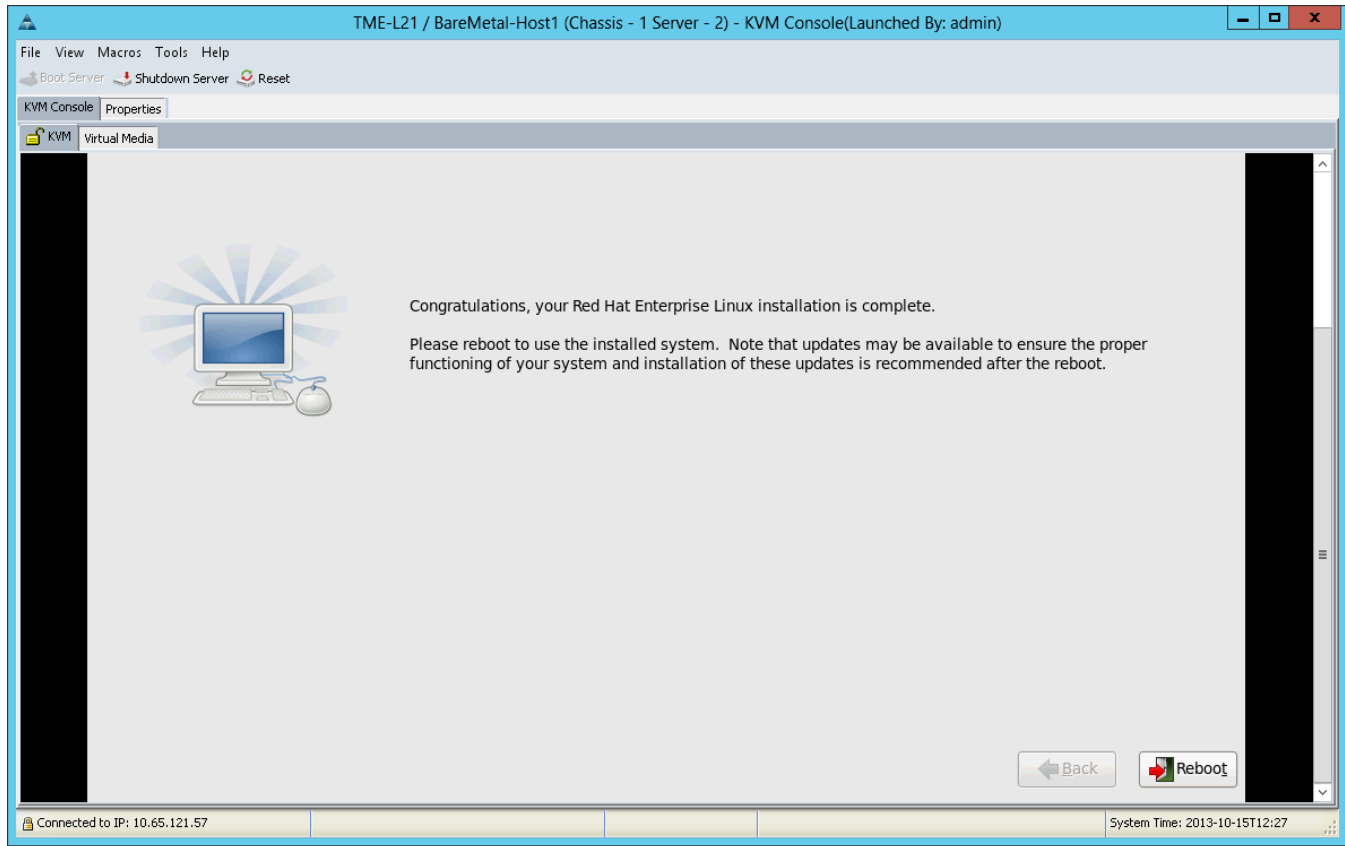


**Figure 324**      **No Network Configured**



32. Click **Reboot**.

33. Final screen after PXE Boot on host **BareMetal-Host1**.

**Figure 325**      **Instance Final Configuration setting**

## Baremetal Blade Deprovision

This section explains the steps taken for cleaning up Baremetal Instance as part of maintenance efforts by the cloud administrator which releases compute, network and storage resources back to cloud pool for reusability. As part of this process Cisco UCS Manager will disassociate Service Profile on Blade and delete Service Profile,



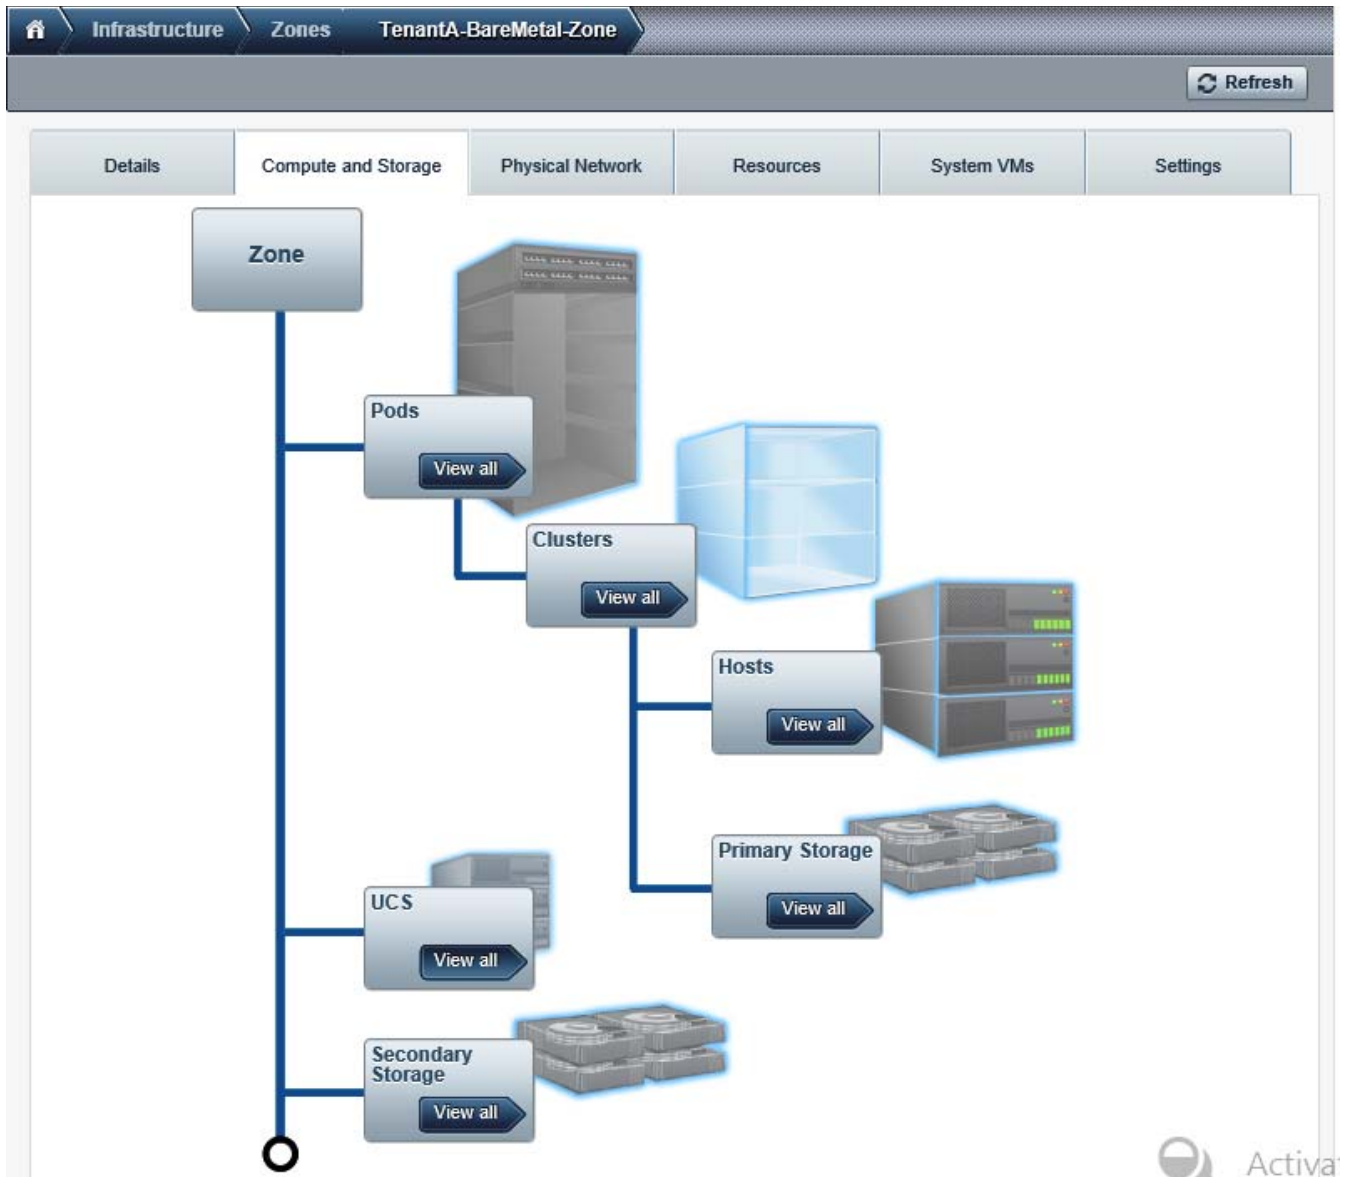
### Note

In this study we will disassociate BareMetal-Host1 using CloudPlatform.

Login to CloudPlatform with User credentials to delete BareMetal-Host1 Instance on the zone TenantA-BareMetal-Zone:

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Infrastructure** tab.
4. Click the **Zones** and click on **TenantA-BareMetal-Zone**.
5. Click **UCS ViewAll**.

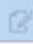



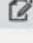


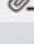
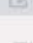

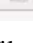
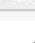
Figure 326 Accessing Cisco UCS Plugin



6. Click **TME-L21** and click in **Blades**.
7. Click **Disassociate Profile from Blade** icon on blade-2 of chassis-1.

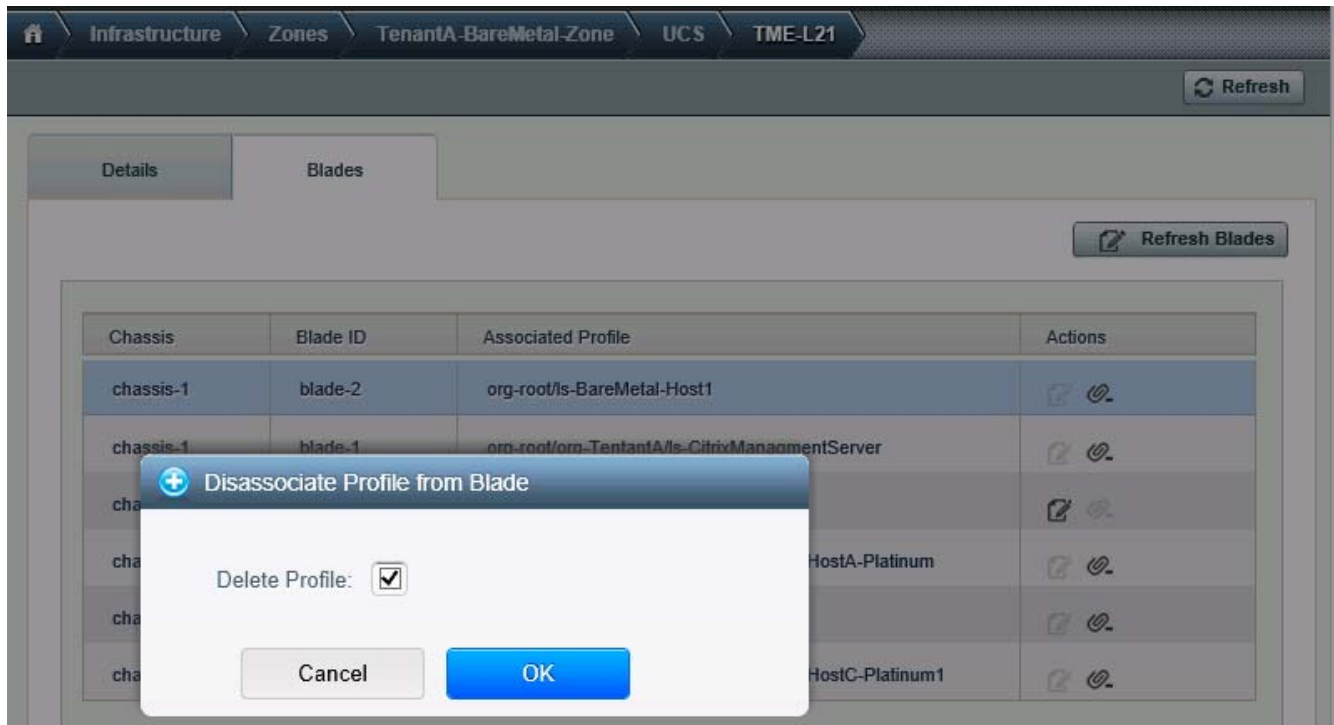
**Figure 327**      *Selecting blade-2 on chassis-1 to Disassociate from BareMetal-Host1 Service Profile in Cisco UCS Plugin*

The screenshot shows the Cisco UCS Manager interface. The breadcrumb navigation at the top indicates the path: Infrastructure > Zones > TenantA-BareMetal-Zone > UCS > TME-L21. A 'Refresh' button is located in the top right. Below the navigation, there are two tabs: 'Details' and 'Blades'. The 'Blades' tab is active. A 'Refresh Blades' button is located above the table. The table has four columns: 'Chassis', 'Blade ID', 'Associated Profile', and 'Disassociate Profile from Blade'. The first row is highlighted in blue, showing 'chassis-1', 'blade-2', and 'org-root/ls-BareMetal-Host1'. The other rows show different blades and their associated profiles.

| Chassis   | Blade ID | Associated Profile                                  | Disassociate Profile from Blade                                                                                                                                           |
|-----------|----------|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| chassis-1 | blade-2  | org-root/ls-BareMetal-Host1                         |     |
| chassis-1 | blade-1  | org-root/org-TenantA/ls-CitrixManagmentServer       |     |
| chassis-1 | blade-4  |                                                     |     |
| chassis-1 | blade-8  | org-root/org-TenantA/ls-TenantA-ESX-HostA-Platinum  |     |
| chassis-1 | blade-5  | org-root/org-TenantA/ls-N1kVSM                      |     |
| chassis-1 | blade-7  | org-root/org-TenantA/ls-TenantA-ESX-HostC-Platinum1 |   |

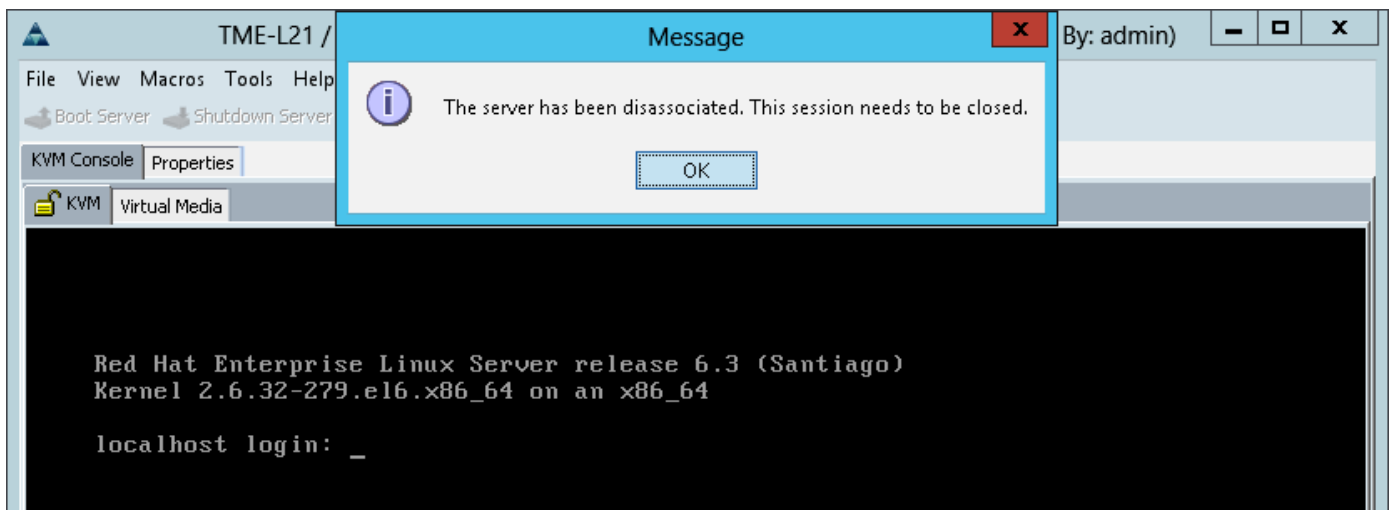
8. Enable **Delete Profile** check box (This option will delete Service profile on Cisco UCS Manager).
9. Click **OK**.

**Figure 328** Enabling Delete Profile check box to delete Service Profile BareMetal-Host1



10. Service Profile BareMetal-Host1 is disassociated with blade-2 and deleted in Cisco UCS Manager.

**Figure 329** Service Profile BareMetal-Host1 disassociated and deleted on blade-2



## Baremetal Instance Life Cycle Management

This section explains the steps taken for providing Baremetal Instance life cycle management as part of maintenance efforts by the cloud administrator which releases compute, network and storage resources back to cloud pool for reusability.

The CloudPlatform 4.2.1.1 invokes Cisco UCS Manager API Calls using IPMI interface to start, stop, reboot and destroy RHEL Operating System running on Cisco UCS Blade Server, part of Baremetal host.

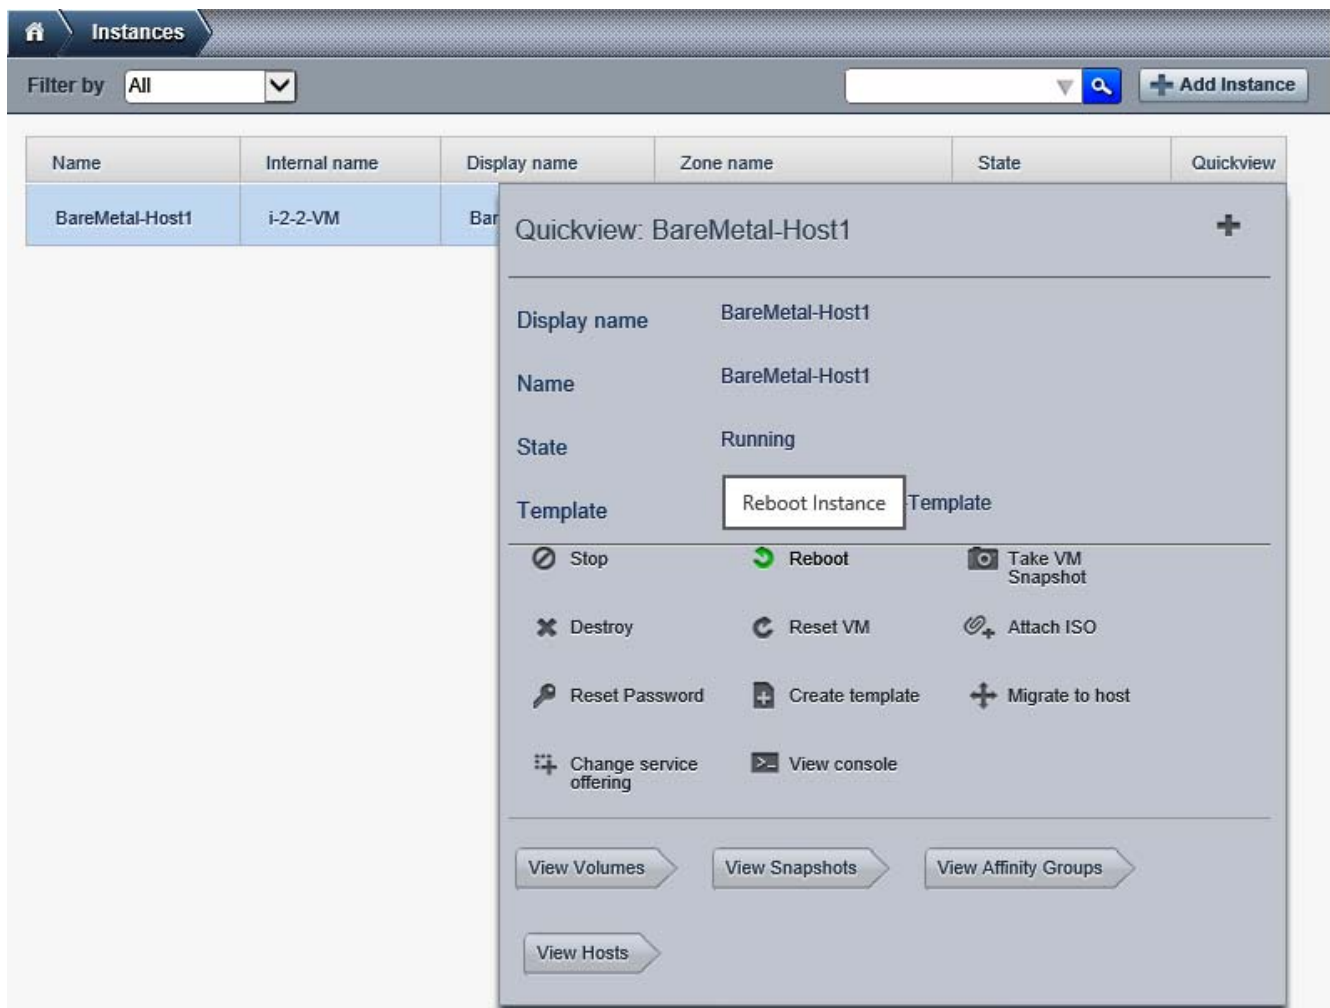
## Reboot Baremetal Instance

In this study we reboot BareMetal-Host1 Instance using CloudPlatform.

Login to CloudPlatform with User credentials to reboot BareMetal-Host1 Instance on the zone TenantA-BareMetal-Zone:

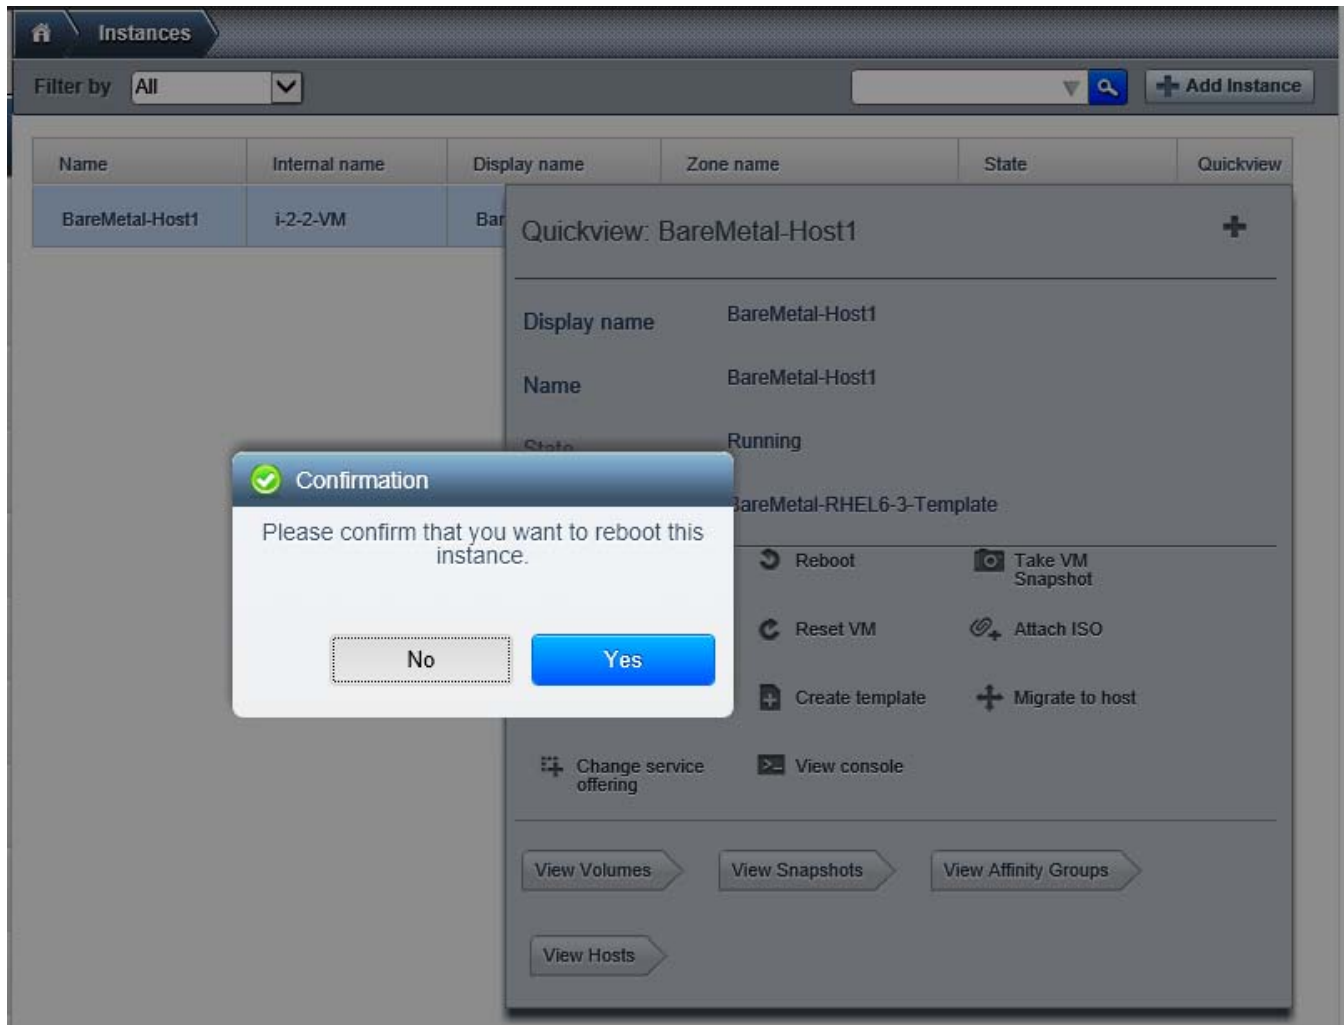
1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instances** tab.
4. Click **Quickview** icon on BareMetal-Host1.
5. Click **Reboot**.

**Figure 330 Reboot BareMetal-Host1 Instance**



6. Click **Yes** to confirm instance.

**Figure 331**      **Confirm Reboot BareMetal-Host1 Instance**



## Stop and Start Baremetal Instance



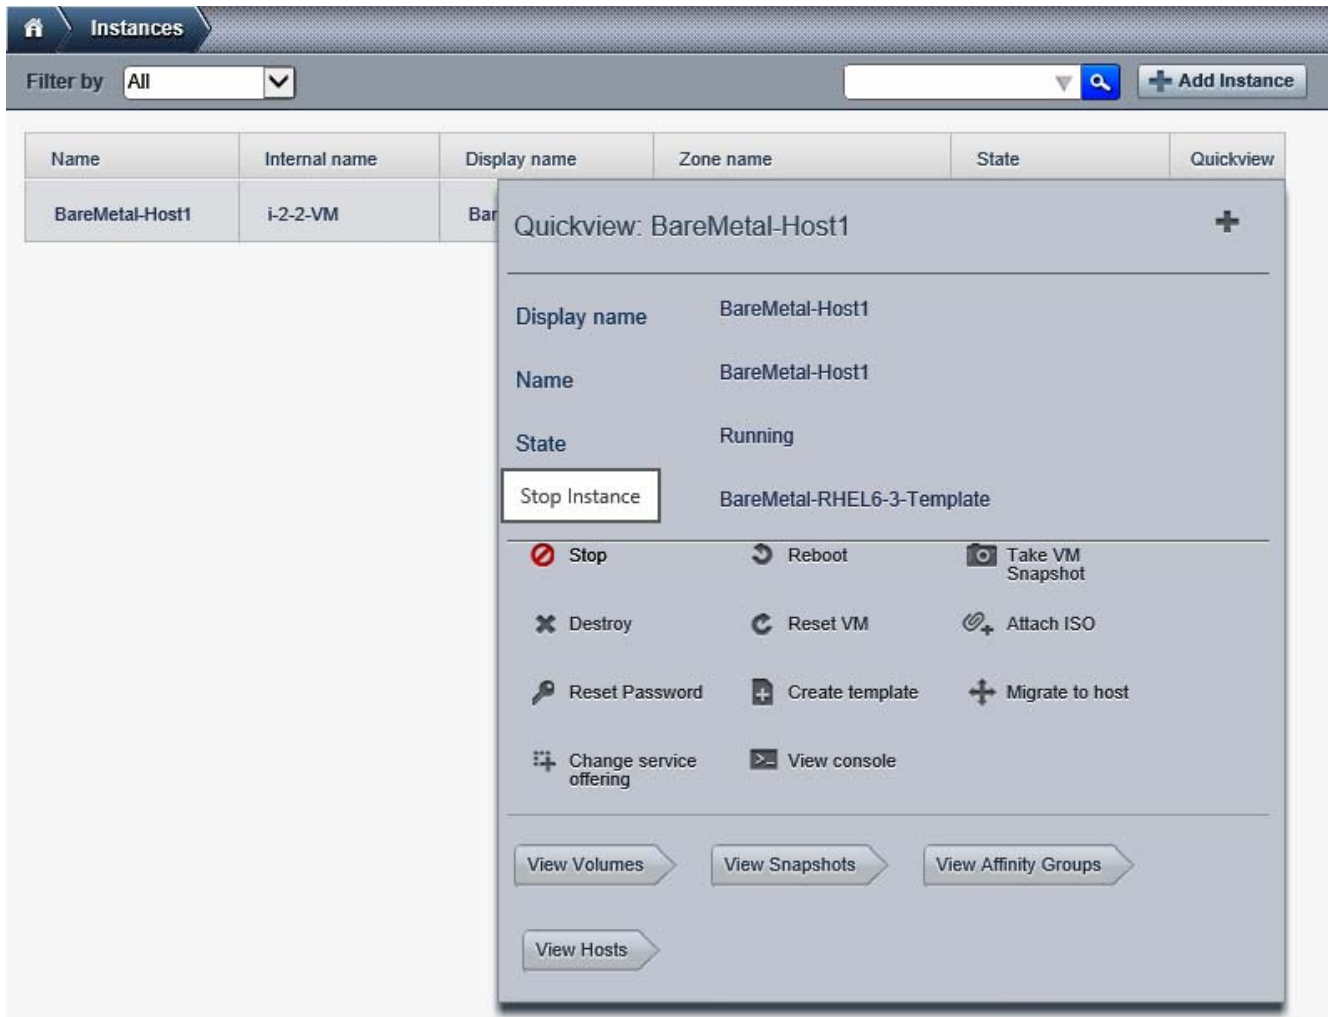
### Note

In this study we stop and start BareMetal-Host1 Instance using CloudPlatform.

Login to CloudPlatform with User credentials to reboot BareMetal-Host1 Instance on the zone TenantA-BareMetal-Zone:

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instances** tab.
4. Click **Quickview** icon on **BareMetal-Host1**.
5. Click **Stop**.

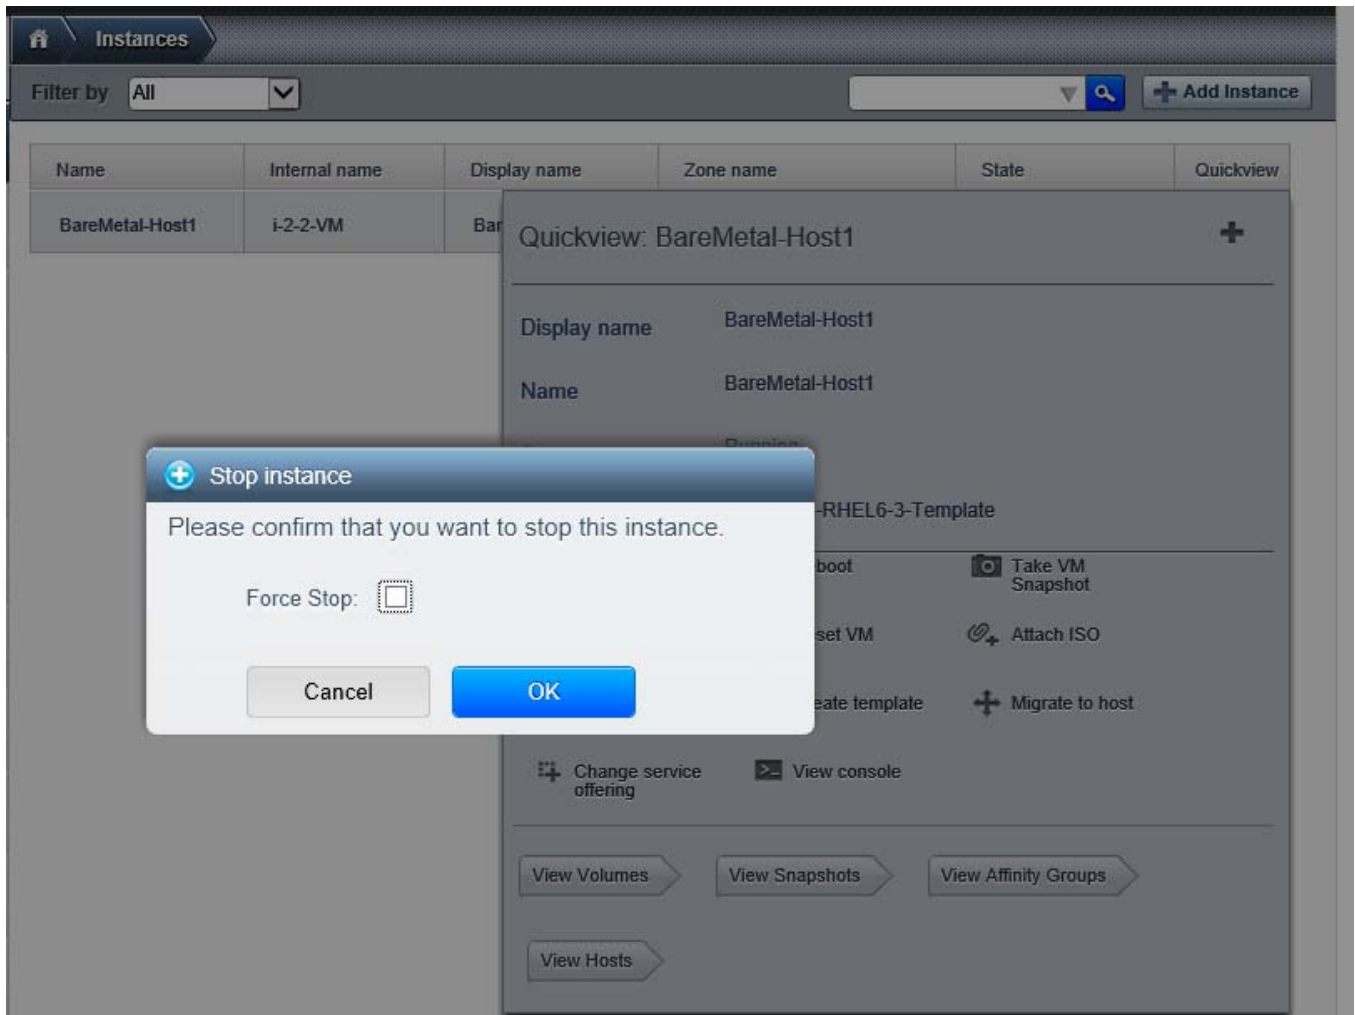
Figure 332 Stop BareMetal-Host1 Instance



6. Click **Yes** to confirm Stop instance.
7. Uncheck **Force Stop** instance.



**Figure 333**      **Confirm Stop BareMetal-Host1 Instance**



8. Stop instance status.

**Figure 334**      **Stop status of BareMetal-Host1 Instance**



9. Click the **Instances** tab.
10. Click **Quickview** icon on **BareMetal-Host1**.
11. Click **Start instance**

Figure 335 Start BareMetal-Host1 Instance

The screenshot shows the OpenStack dashboard 'Instances' page. A table lists instances, with 'BareMetal-Host1' selected. A 'Quickview' modal is open, displaying details for 'BareMetal-Host1'. The 'Start Instance' button is highlighted with a red box.

| Name            | Internal name | Display name    | Zone name | State   | Quickview |
|-----------------|---------------|-----------------|-----------|---------|-----------|
| BareMetal-Host1 | i-2-2-VM      | BareMetal-Host1 |           | Stopped | +         |

Quickview: BareMetal-Host1

Display name: BareMetal-Host1

Name: BareMetal-Host1

State: Stopped

Start Instance: BareMetal-RHEL6-3-Template

Start Instance

Take VM Snapshot

Destroy

Reset VM

Change affinity

Attach ISO

Reset Password

Create template

Migrate to storage

Change service offering

Assign Instance to Another Account

View Volumes

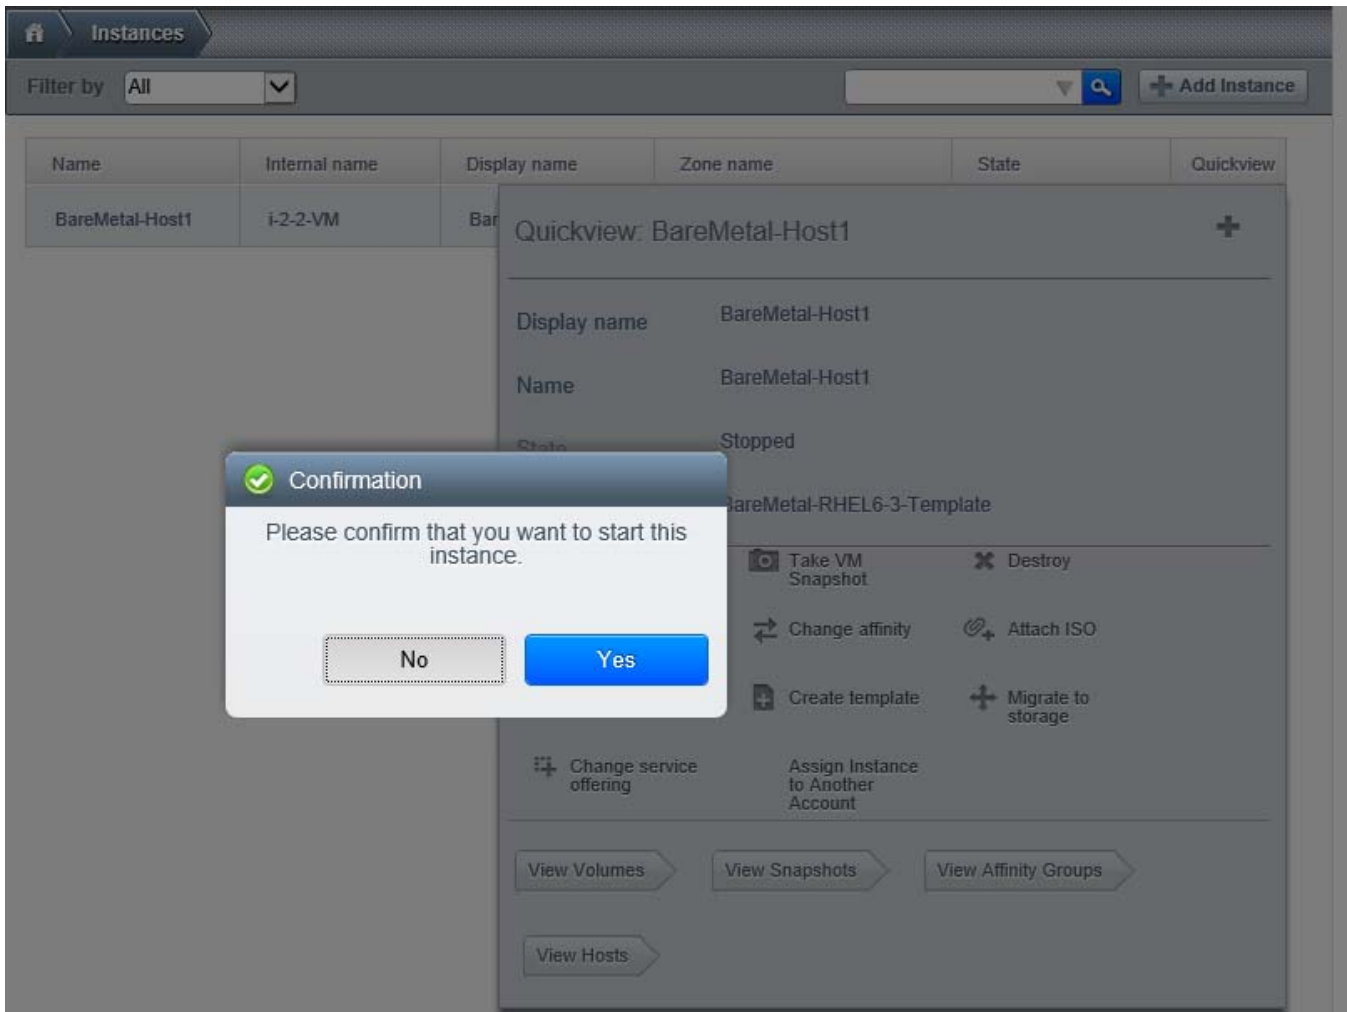
View Snapshots

View Affinity Groups

View Hosts

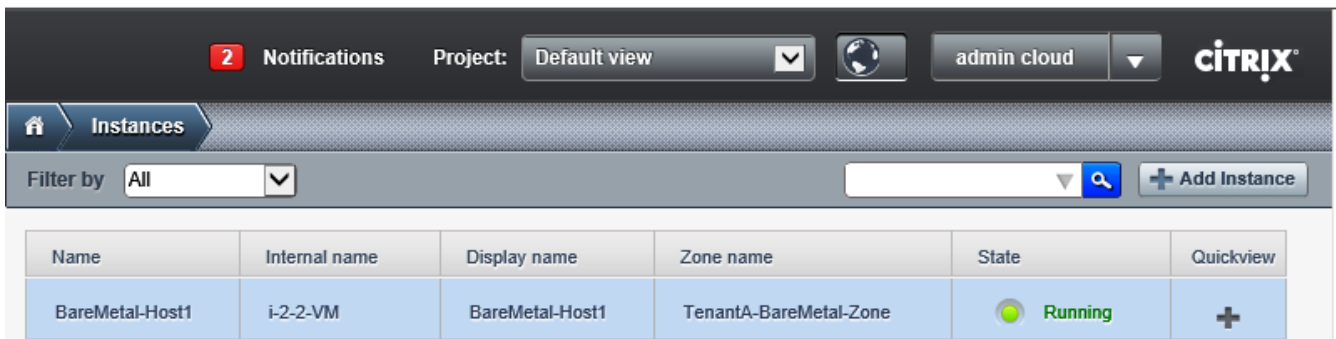
12. Click **Yes** to confirm start instance.

**Figure 336**      **Start BareMetal-Host1 Instance Confirmation**



13. Running BareMetal Status after start instance.

**Figure 337**      **BareMetal-Host1 running status after start instance**



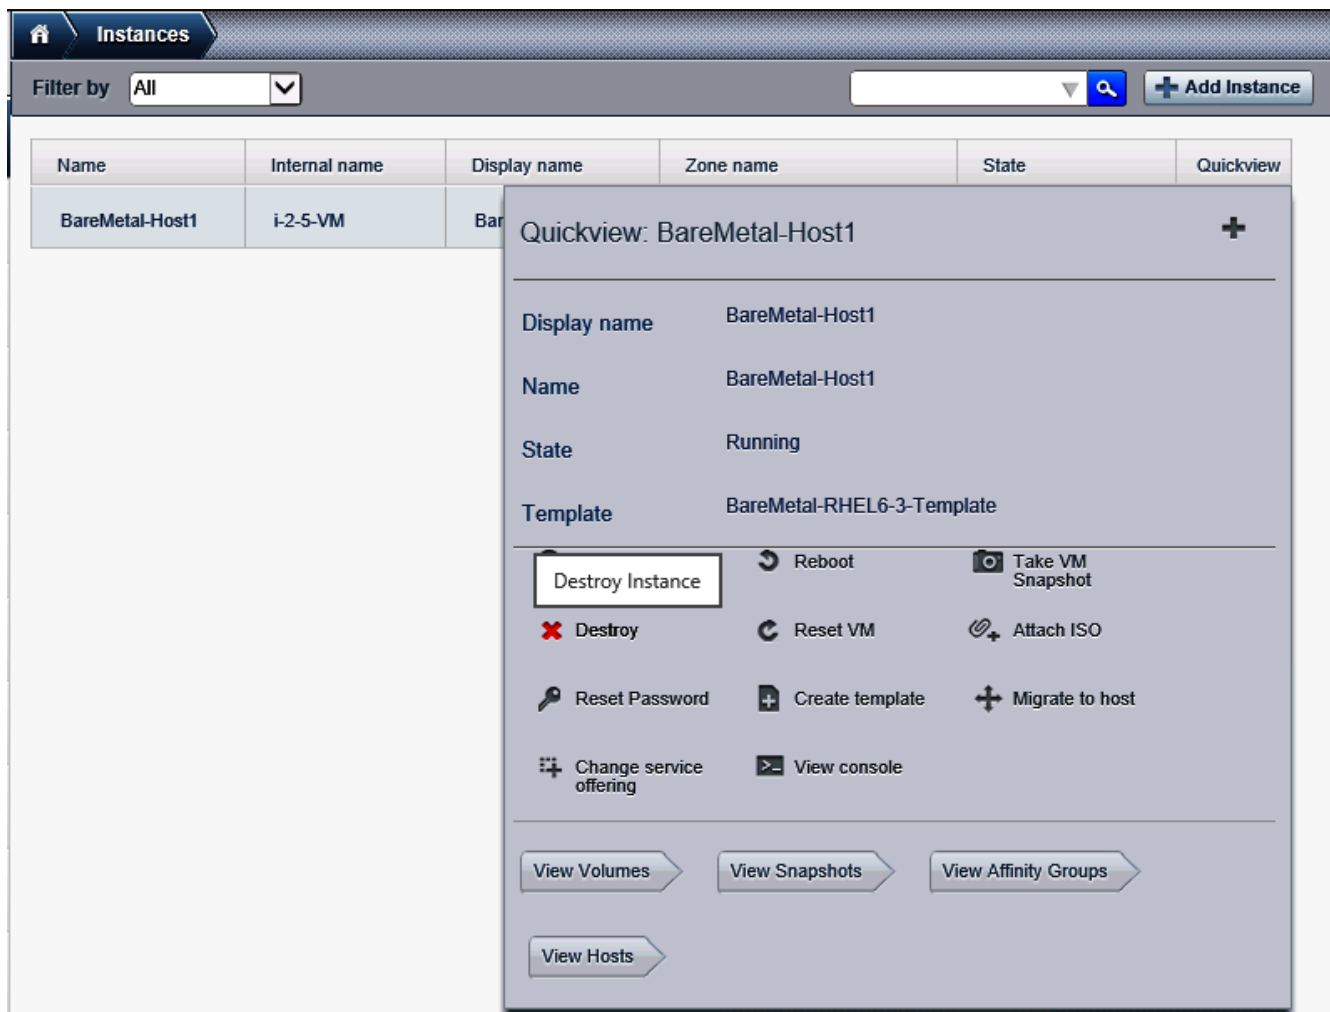
## Destroy and Restore Baremetal Instance

In this study we destroy and restore BareMetal-Host1 Instance using CloudPlatform.

Login to CloudPlatform with User credentials to destroy BareMetal-Host1 instance on the zone TenantA-BareMetal-Zone:

1. Provide User Name <root> and Password <XXXXXX> and Domain.
2. Click **Login**.
3. Click the **Instances** tab.
4. Click **Quickview** icon on **BareMetal-Host1**.
5. Click **Destory**.

**Figure 338**      *Destroy BareMetal-Host1 Instance*



6. Click **OK** to confirm destroy.
7. Uncheck **Expunge** check box

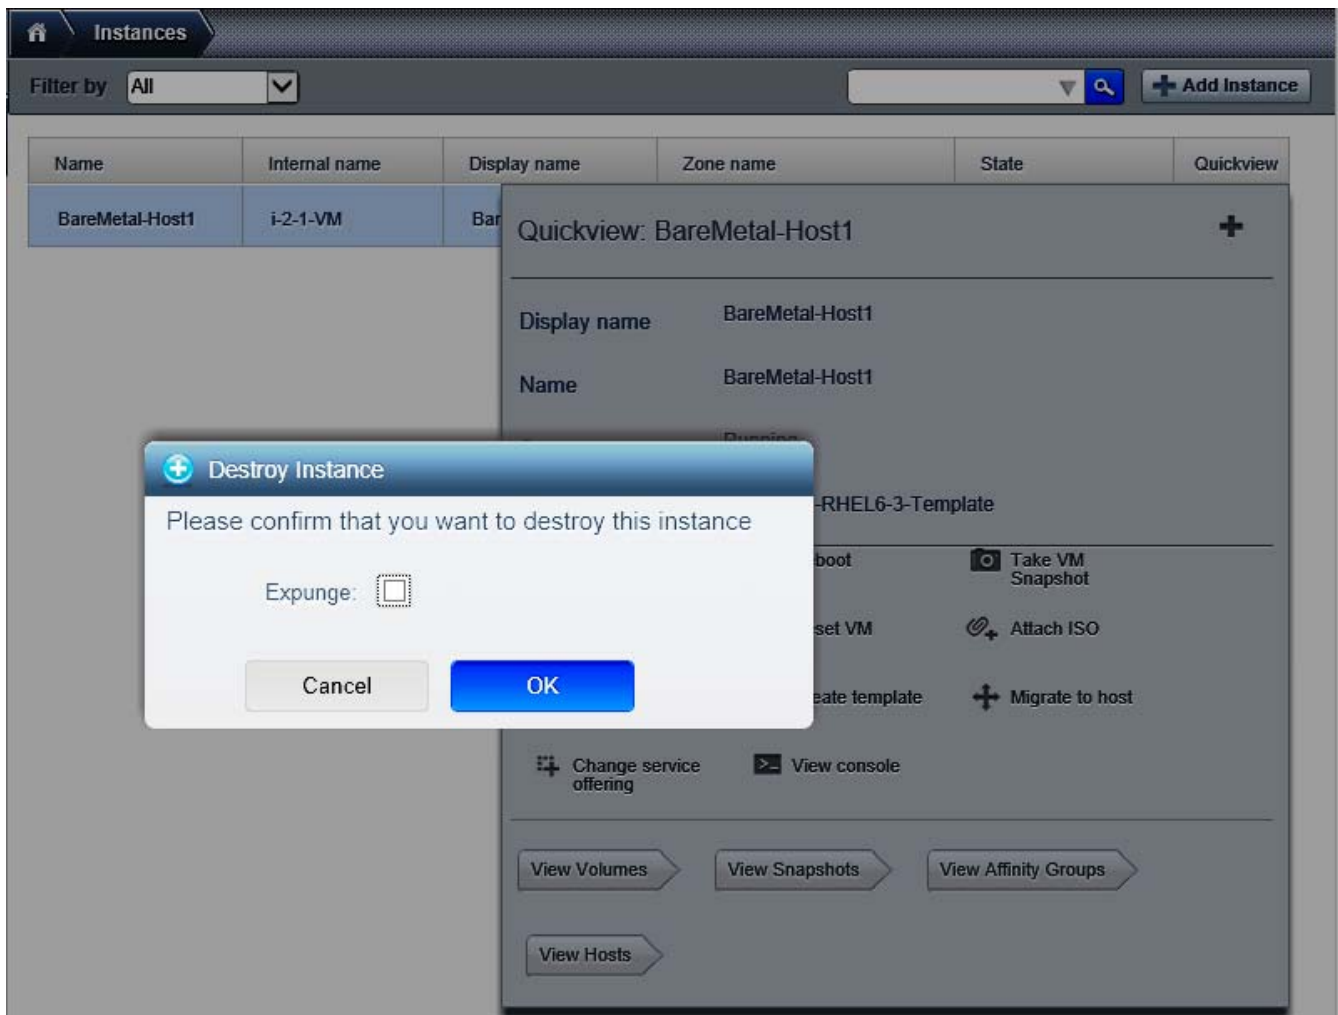


**Note** If **Expunge** is unchecked, the BareMetal-Host1 instance can be restored later.



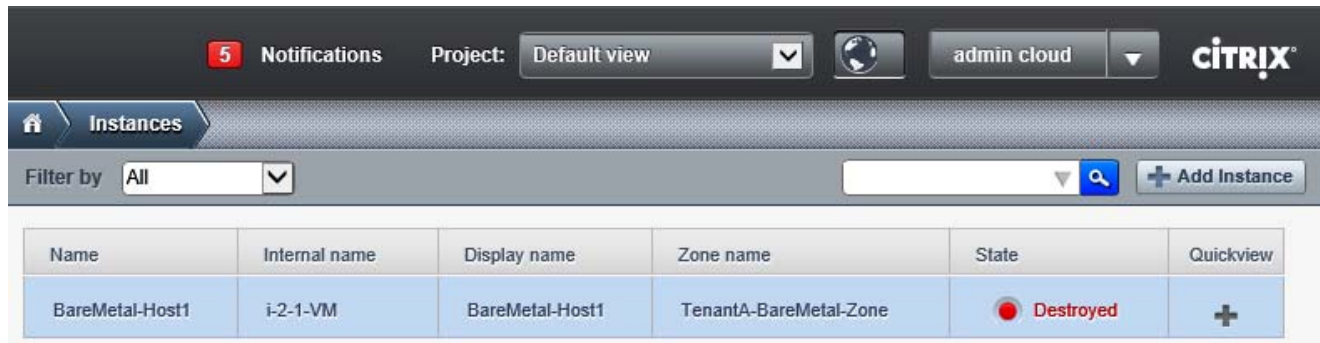
**Note** If **Expunge** is checked, the BareMetal-Host1 instance is deleted and cannot be restored.

**Figure 339** Confirming Destroying BareMetal-Host1 Instance



8. Instance BareMetal-Host1 destroyed.

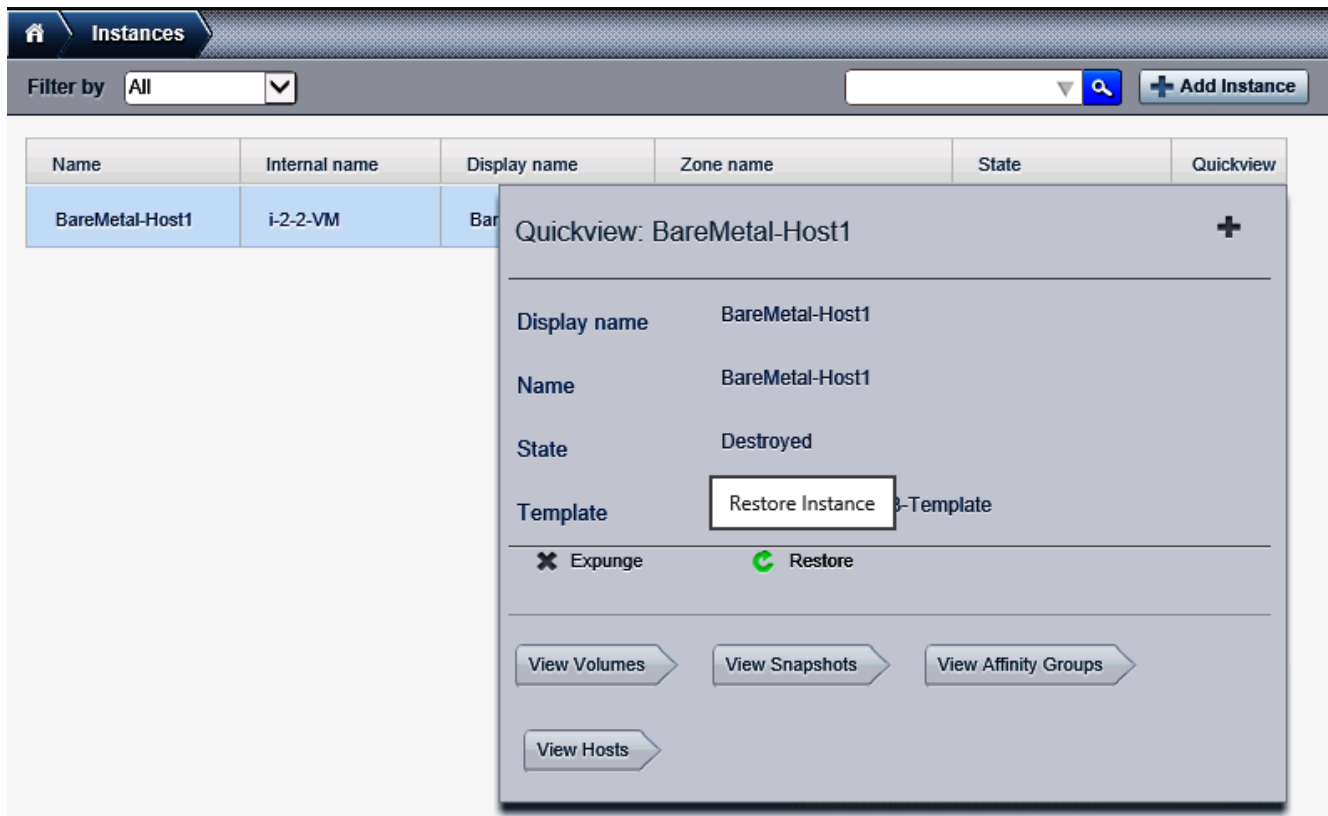
**Figure 340**      *Destroyed BareMetal-Host1 Instance*



## Restore

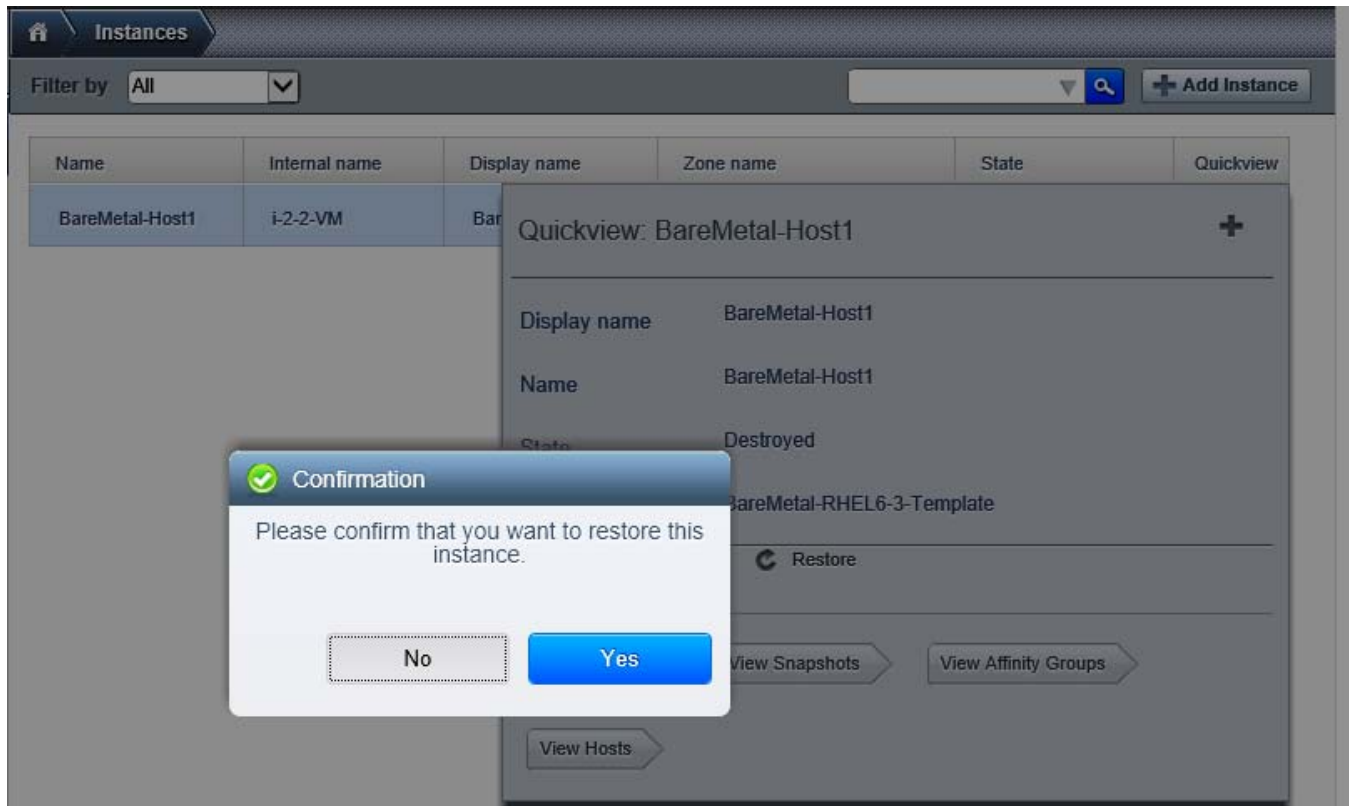
1. Click the **Instances** tab.
2. Click **Quickview** icon on BareMetal-Host1.
3. Click **Restore**.

**Figure 341**      *Restore BareMetal-Host1 Instance*



4. Click **Yes** to confirm restore on BareMetal-Host1.

**Figure 342**      *Restore BareMetal-Host1 Instance Confirmation*



5. Click stopped instance **BareMetal-Host1**.

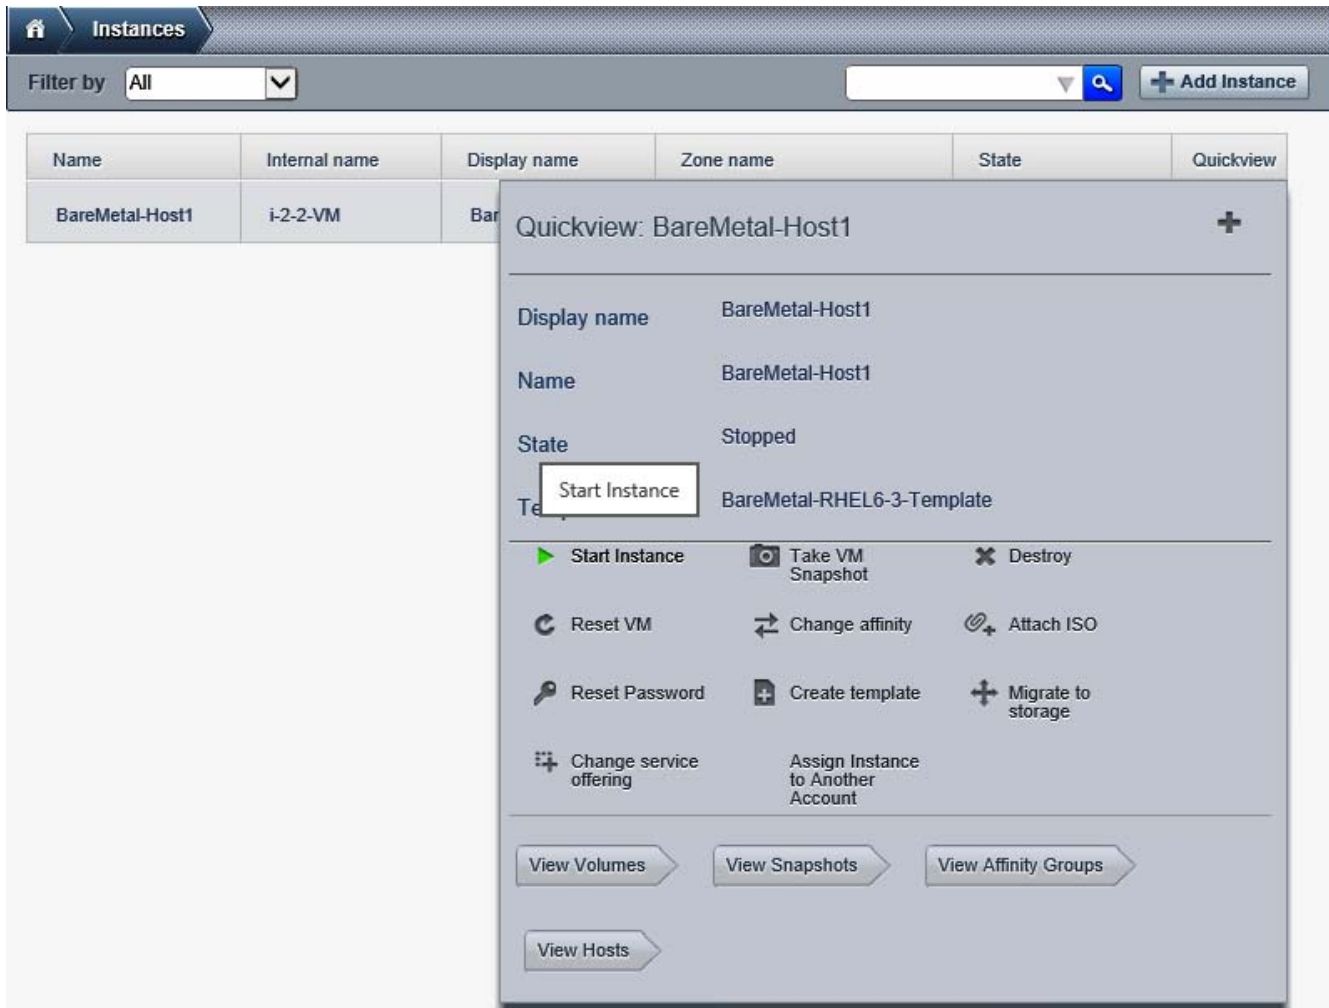
**Figure 343**      *Stopped BareMetal-Host1 Instance*



6. Click stopped instance **BareMetal-Host1**.
7. Click **Quickview** icon.
8. Click **Start Instance**.



Figure 344 Start BareMetal-Host1 Instance



9. Running BareMetal Status after restore instance.

Figure 345 BareMetal-Host1 running status after Restore instance





## Conclusion

Cisco and Citrix are both committed to providing superior cloud solutions to a global clientele. Together, products from both companies enable the flexible and agile delivery of cloud-based services with simplified physical and virtual infrastructure management. Many established enterprises have selected the combination of Cisco and Citrix offerings after analyzing best-in-class cloud products and have built robust and scalable enterprise clouds based on Cisco UCS and Citrix CloudPlatform. The reference architecture in this document further simplifies the deployment of these products by providing best practices for infrastructure and software configurations.

## References

Documents listed here provide additional information relevant to implementing Citrix Private Cloud with NetApp Storage System on Cisco UCS B-Series Servers.

- [Cisco Nexus QoS Switch Configuration Guide: Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide](#)
- [Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 5.0\(3\)N2\(1\)](#)
- [Citrix CloudPlatform Installation and configuration Guides](#)
- [Cisco UCS System Hardware and Software Interoperability Matrix](#)
- [NetApp Storage Deployment Guide](#)