



Newer Cisco SBA for Government Guides Available

This guide is part of an older series of Cisco Smart Business Architecture for Government. To access the latest Cisco SBA for Government Guides, go to <http://www.cisco.com/go/govsba>

Cisco strives to update and enhance SBA guides on a regular basis. As we develop a new series of SBA guides, we test them together, as a complete system. To ensure the mutual compatibility of designs in Cisco SBA guides, you should use guides that belong to the same series.





Unified Computing Deployment Guide

● ● ● SBA FOR GOVERNMENT

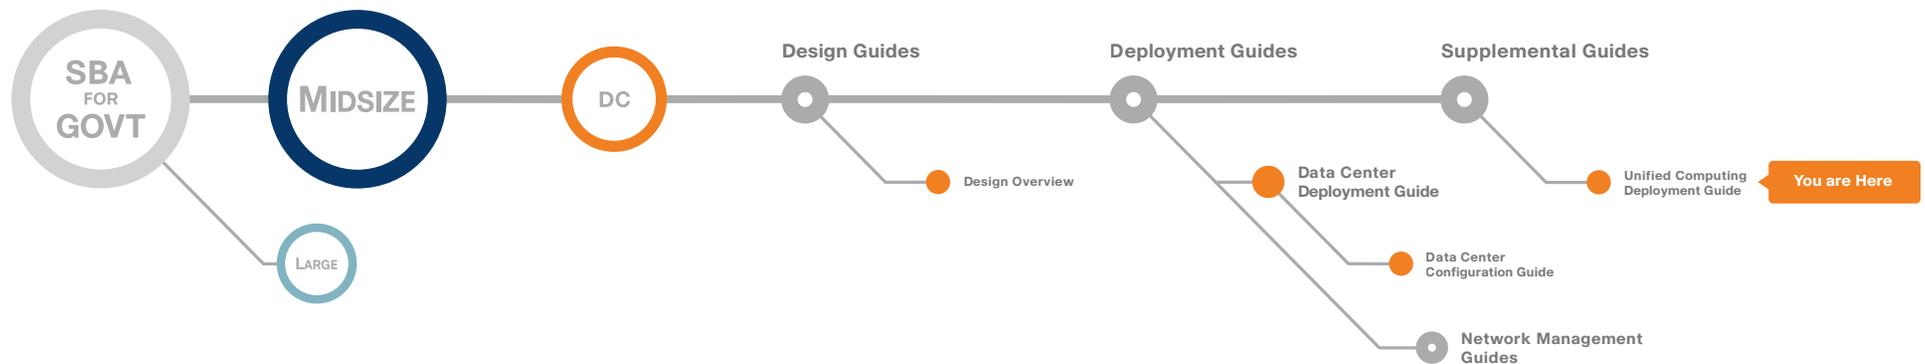
The Purpose of this Document

This guide is a concise reference on Cisco Unified Computing System (UCS) deployment. It includes an overview of some of the agency problems that Unified Computing can solve within your agency and the capabilities it brings to bear to solve them. It also provides step-by-step configuration instructions for the basic initial setup of Cisco UCS. Cisco UCS Manager GUI service profile examples are provided for basic server configuration, and for boot-from-LAN (PXE Boot) and boot-from-SAN setups.

Who Should Read This Guide

This guide is intended for the reader with any or all of the following:

- Responsibility for selection or implementation of server hardware
- Up to 250 physical or virtualized servers
- CCNA® certification or equivalent experience



The reader may be looking for any or all of the following:

- To reduce the complexity of managing application servers
- To increase application availability and reduce downtime
- To reduce the time required to deploy new servers, and upgrades
- To simplify cabling and more efficiently utilize space in equipment racks
- To prepare server hardware to support server virtualization
- To adopt centralized storage to more efficiently manage their storage environment
- To expand existing application servers to address growth
- To rely upon the assurance of a tested solution

Related Documents

Before reading this guide

- Borderless Networks Foundation Design Overview
- Borderless Networks Deployment Guide
- Borderless Networks Configuration Files Guide

Optional documents

- Borderless Networks Technology-Specific Guides
- Data Center Design Guide
- Data Center Deployment

Table of Contents

Introduction	1	Advanced Configurations	45
Guiding Principles	1	Working with Service Profile Templates.....	45
The Purpose of this Guide	1	Service Profiles using Multiple vNICs and Trunking	47
Agency Overview	2	Using a Virtual Interface Card.....	50
Technical Overview	4	Virtual Machine Integration	50
Network Infrastructure Systems.....	4	Enabling VLAN Trunking on vNICs	51
Computing Systems	5	Service Profiles Using Multiple vHBAs.....	52
Storage Systems	7	Appendix	53
Server Virtualization Software	8	Appendix A: Configuration Values Matrix	53
Deploying the SBA Unified Computing Architecture	9	Appendix B: Equipment List.....	55
Ethernet Network Infrastructure.....	9	Appendix C: SBA for Midsize Agencies Document System	56
Fibre Channel Network Infrastructure	10		
UCS Blade Server System	12		
Cisco UCS Rack Mount Servers	41		

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.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental. Cisco Unified Communications SRND (Based on Cisco Unified Communications Manager 7.x)

© 2010 Cisco Systems, Inc. All rights reserved.

Introduction

The Cisco® Smart Business Architecture (SBA) for Government is a comprehensive design for networks with up to 1000 users. This out-of-the-box design is simple, fast, affordable, scalable, and flexible.

The Cisco SBA for Midsize Agencies incorporates LAN, WAN, wireless, security, WAN optimization, and unified communication technologies tested together as a solution. This solution-level approach simplifies the system integration normally associated with multiple technologies, allowing you to select the modules that solve your agency's problems rather than worrying about the technical details.

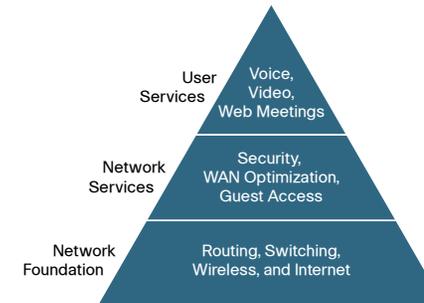
We have designed the Cisco SBA to be easy to configure, deploy, and manage. This architecture:

- Provides a solid network foundation
- Makes deployment fast and easy
- Accelerates ability to easily deploy additional services
- Avoids the need for re-engineering of the core network

By deploying the Cisco SBA, your agency can gain:

- A standardized design, tested and supported by Cisco
- Optimized architecture for midsize agencies with up to 1000 users and up to 20 branches
- Flexible architecture to help ensure easy migration as the agency grows
- Seamless support for quick deployment of wired and wireless network access for data, voice, teleworker, and wireless guest
- Security and high availability for agency information resources, servers, and Internet-facing applications
- Improved WAN performance and cost reduction through the use of WAN optimization
- Simplified deployment and operation by IT workers with CCNA® certification or equivalent experience
- Cisco enterprise-class reliability in products designed for midsize agencies

Figure 1. SBA for Government Model



Guiding Principles

We divided the deployment process into modules according to the following principles:

- **Ease of use:** A top requirement of Cisco SBA was to develop a design that could be deployed with the minimal amount of configuration and day-two management.
- **Cost-effective:** Another critical requirement as we selected products was to meet the budget guidelines for midsize agencies.
- **Flexibility and scalability:** As the agency grows, so too must its infrastructure. Products selected must have the ability to grow or be repurposed within the architecture.
- **Reuse:** We strived, when possible, to reuse the same products throughout the various modules to minimize the number of products required for spares.

The Cisco SBA can be broken down into the following three primary, modular yet interdependent components for the midsize agency.

- **Network Foundation:** A network that supports the architecture
- **Network Services:** Features that operate in the background to improve and enable the user experience without direct user awareness
- **User Services:** Applications with which a user interacts directly

The Purpose of this Guide

This Unified Computing Deployment Guide introduces the Cisco solutions for both Cisco UCS Blade Server systems and Cisco UCS C-Series rack mount systems.

It explains the requirements that were considered when building the Cisco SBA design and introduces each of the products that were selected.

Agency Overview

As a midsize agency begins to grow, the number of servers required to handle the information processing tasks of the agency grows as well. Using the full capabilities of the investment in server resources can help an agency add new applications while controlling costs as they move from a small server room environment into a mid-sized data center. Server virtualization has become a common approach to allow an agency to access the untapped processing capacity available in processor technology. Streamlining the management of server hardware and its interaction with networking and storage equipment is another important component of using this investment in an efficient manner.

Scaling a data center with conventional servers, networking equipment, and storage resources can pose a significant challenge to a growing agency. Multiple hardware platforms and technologies must be integrated to deliver the expected levels of performance and availability to application end users. These components in the data center also need to be managed and maintained, typically with a diverse set of management tools with different interfaces and approaches. In larger agencies, often multiple teams of people are involved in managing applications, servers, storage, and networking. In a midsize agency, the lines between these tasks are blurred and often a single, smaller team, or even one individual, may need to handle many of these tasks in a single day.

Consistent with the SBA approach, Cisco offers a simplified reference model for managing a small server room as it grows into a full-fledged data center. This model benefits from the ease of use offered by the Cisco UCS. Cisco UCS provides a single graphical management tool for the provisioning and management of servers, network interfaces, storage interfaces, and their immediately attached network components. Cisco UCS treats all of these components as a cohesive system, which simplifies these complex interactions and allows a midsize agency to deploy the same efficient technologies as larger agencies, without a dramatic learning curve.

This system integrates cleanly into the network foundation established in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide* and is designed to scale simply with the requirements of a growing agency.

This guide addresses many of the same agency issues encountered by growing agencies that are identified in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide* but it focuses on the server resources themselves and their interaction with network and storage systems. These common challenges include:

- Application Growth
- Increasing Data Storage Requirements
- Managing Processing Resources
- Availability and Business Continuance

Application Growth

As an application scales to support a larger number of users or as you deploy new applications, the number of servers required to meet the needs of the agency increases. The Cisco SBA Unified Computing model provides for rapid deployment of additional physical servers with common attributes through a simple graphical interface. Using Cisco UCS service profiles, the personality of an individual server is logically defined separately from any specific physical hardware, including boot characteristics, interface addresses, and even firmware versions. Service profiles can also be generated from a template, and may remain linked to the template to facilitate easier updates across multiple servers in the future.

Increasing Data Storage Requirements

As application requirements grow, the need for additional data storage capacity also increases. You can most efficiently manage the investment in additional storage capacity by moving to a centralized storage model. The SBA Unified Computing model decouples the computing functions of the server farm from the storage systems, which provides greater flexibility for system growth and migration. Basic local disk capacity is available on each server to facilitate local boot capability or to provide local caching capability to servers booted from the Ethernet IP network or Fibre Channel Storage Area Network (SAN).

Managing Processing Resources

As an agency grows, traditional servers may become dedicated to single applications to increase stability and simplify troubleshooting, but these servers do not operate at high levels of processor utilization for much of the day. Server virtualization technologies insert a hypervisor layer between the server operating systems and the hardware, allowing a single physical server to run multiple instances of different “guest” operating systems such as Microsoft Windows or Linux. This increases the utilization of the processors on the physical servers, which helps to optimize this costly resource.

The architecture of the SBA Unified Computing model is optimized to support the use of hypervisor-based systems or the direct installation of a base operating system such as Windows or Linux. The service profile structure of Cisco UCS, along with a centralized storage model, allows the easy portability of server definitions to different hardware with or without a hypervisor system in place. Built on the data center infrastructure foundation defined in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide*, the SBA Unified Computing architecture provides scalable connectivity options for not only Cisco UCS Blade Server chassis but also Cisco UCS C-Series Rack-Mount Servers, as well as connectivity options to support third-party servers.

Availability and Business Continuation

Midsized agencies rely on their investment in servers, storage, and networking technology to provide highly available access to critical electronic operational processes. We have taken many steps in all layers of the SBA to ensure this availability with the use of resilient network devices, links, and service models. The SBA Unified Computing model extends this resiliency to the servers themselves through the capabilities of Cisco UCS.

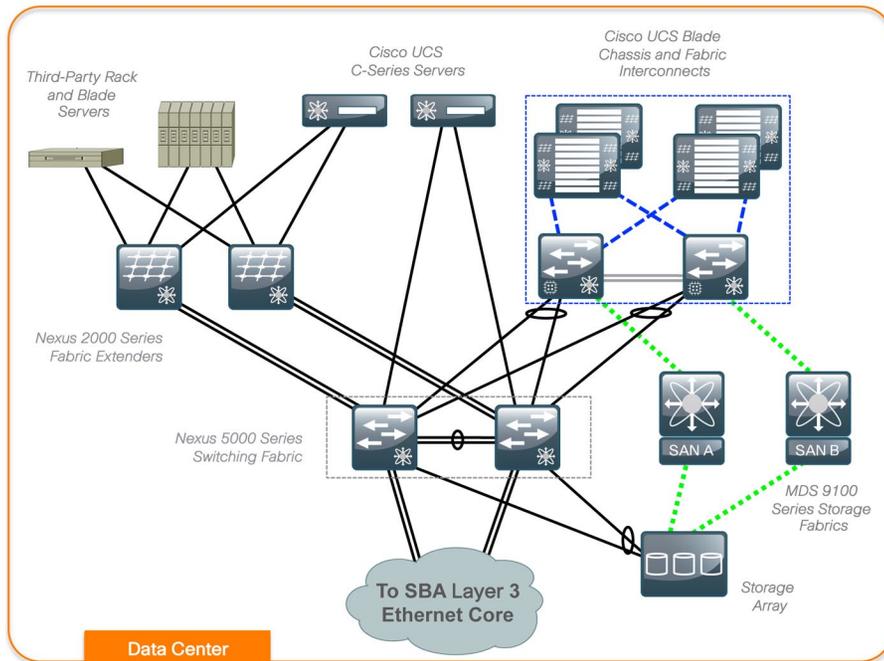
Cisco UCS uses service profiles to provide a consistent interface for managing all server resource requirements as a logical entity, independent of the specific hardware module that is used to provide the processing capacity. This service profile approach is applied consistently on both virtualized servers and “bare metal” servers, which do not run a hypervisor. This capability allows the entire personality of a given logical server to be ported easily to a different physical server module independent of any virtualization software when LAN or SAN boot are in use. This approach increases overall availability and dramatically reduces the time required to replace the function of an individual failed server module.

Notes

Technology Overview

The SBA Unified Computing reference design has been lab-validated in conjunction with the architecture defined in the *Cisco SBA Data Center for Midsize Agencies Deployment Guide*, available at: <http://www.cisco.com/go/sba>

Figure 2. The SBA Unified Computing Architecture



This architecture is flexible enough to be adapted to alternate Ethernet or Fibre Channel topologies, which can help you migrate from a legacy installed base of equipment towards a standardized reference design such as SBA. Figure 2 shows the data center components of this architecture and their interaction with the SBA headquarters core layer.

The SBA Unified Computing model is also adaptable to multiple ways of accessing centralized storage. Two alternatives for storage access are

included in the overall architecture. The simplest approach uses a pure Ethernet IP network to connect the servers to both their user community and the shared storage array. Communication between the servers and storage over IP can be accomplished using an Internet Small Computer System Interface (iSCSI), which is a block-oriented protocol encapsulated over IP, or traditional network-attached storage protocols such as Common Internet File System (CIFS) or Network File System (NFS). Servers can also be booted directly from the Local-Area Network (LAN), either for rapid OS deployment or for ongoing operations. LAN-based storage access follows the path through the Cisco Nexus 5000 Series Switching Fabric shown in Figure 2.

A more traditional but advanced alternative for providing shared storage access is using a separate SAN built using Fibre Channel switches such as the Cisco MDS 9100 Series. For resilient access, SANs are normally built with two distinct fabric switches that are not cross-connected. Currently, Fibre Channel offers the widest support for various disk-array platforms and also support for boot-from-SAN. This type of storage access follows the path through the Cisco MDS 9100 Series Storage Fabric Switches that are shown in Figure 2.

Many available shared storage systems offer multi-protocol access to the system, including iSCSI, Fibre Channel, CIFS, and NFS. The approaches using Ethernet and Fibre Channel are shown separately in this guide for clarity but you can combine them easily on the same system to meet the access requirements of a variety of server implementations. This flexibility also helps facilitate migration from legacy third-party server implementations onto Cisco UCS.

Network Infrastructure Systems

Ethernet

The *Cisco SBA Unified Computing Deployment Guide* is designed as an extension of the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide*. The basis of the SBA Unified Computing architecture is an Ethernet switch fabric that consists of two Cisco Nexus 5000 switches, as shown in Figure 2. This data center switching fabric provides Layer-2 Ethernet switching services to attached devices and, in turn, relies on the SBA Ethernet Core for Layer-3 switching services. The Cisco Nexus 5000-based switching fabric is shown in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide* in conjunction with a Cisco Catalyst 4507-R resilient switch, forming the network core. This SBA Unified Computing topology may also be extended to alternate Cisco switching platforms that provide Layer-3 services and 10-Gigabit Ethernet connectivity.

The two Cisco Nexus 5000 switches form the Ethernet Switch Fabric using Virtual Port Channel (vPC) technology. This feature provides loop-prevention services, and allows the two switches to appear as one logical Layer-2 switching instance to attached devices. In this way, the Spanning Tree Protocol (STP), which is a standard component of Layer-2 bridging, does not need to block any of the links in the topology to prevent bridging loops. Additional 1-Gigabit Ethernet switch port density may be added to the switch fabric using Cisco Nexus 2100 Series Fabric Extenders.

Fibre Channel (FC)

The optional Fibre Channel switching infrastructure in this topology consists of two Cisco MDS 9100 Series fabric switches. These two separate fabrics provide highly available connectivity between the centralized storage system and connected servers. The SBA Data Center for Midsize Agencies topology was validated using MDS 9124 and 9134 switches running 4Gbps Fibre Channel connections.

! **Tech Tip**

The Cisco MDS 9148 supports 8 Gbps Fibre Channel connectivity and has been validated with the topology shown in this guide. As of Cisco UCS Release 1.2(1d), the Cisco UCS 6100 Fabric Interconnects support 1, 2, 4, and 8 Gbps Fibre Channel connections.

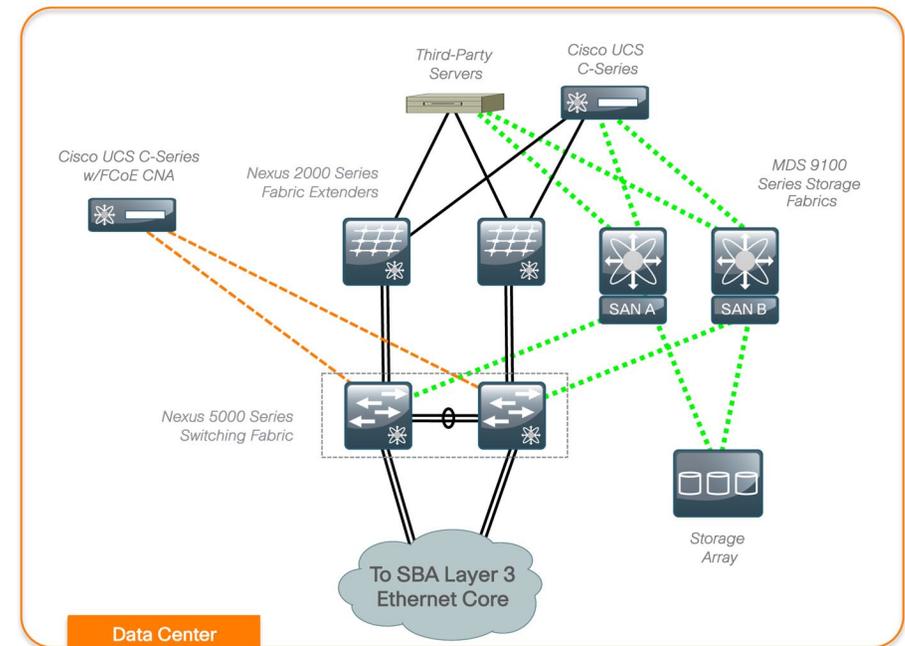
The Cisco UCS 6100 Series Fabric Interconnects also maintain separate Fibre Channel fabrics, so each fabric is attached to one of the Cisco MDS 9100 switches running either SAN A or SAN B as shown in Figure 2. When Fibre Channel is used for storage access from Cisco UCS Blade Servers, the system provides Virtual Host Bus Adaptors (vHBAs) to the service profiles to be presented to the host operating system.

On the Cisco UCS Fabric Interconnect, the Fibre Channel ports that connect to the Cisco MDS SAN operate in N-port Virtualization mode. Though there are multiple Fibre Channel ports on the fabric interconnects, Fiber Channel switching between these ports is not supported. All Fibre Channel switching happens upstream at the Cisco MDS switches running N-Port Identifier Virtualization (NPIV). NPIV allows multiple Fibre Channel port IDs to share a common physical port.

You can connect the Cisco UCS C-Series Rack-Mount Servers to the Fibre Channel SAN using dedicated Host Bus Adaptors (HBAs) that attach directly to the SAN switches. Alternately, you can use a Converged Network Adapter

(CNA), which allows Ethernet data and Fibre Channel over Ethernet (FCoE) storage traffic to share the same physical set of cabling. This Unified Wire approach allows these servers to connect directly to the Cisco Nexus 5000 Series switches for data traffic, as well as SAN A and SAN B highly available storage access, shown in Figure 3. The Cisco Nexus 5000 Ethernet switch fabric is responsible for splitting FCoE traffic off to the Fibre Channel attached storage array.

Figure 3. Cisco UCS C-Series Fibre Channel Connections



Computing Systems

The primary computing platforms targeted for the SBA Unified Computing reference architecture are Cisco UCS B-Series Blade Servers and Cisco UCS C-Series Rack-Mount Servers. The Cisco UCS Manager graphical interface provides ease of use that is consistent with the goals of the SBA. When deployed in conjunction with the SBA Data Center network foundation, the environment provides the flexibility to support the concurrent use of the Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack-Mount Servers, and third-party servers connected to 1 and 10-Gigabit Ethernet connections.

Tech Tip

Cisco UCS Manager functionality is provided through the fabric interconnects. As of Cisco UCS Manager release 1.2(1d), Cisco B-Series servers implemented in a Cisco UCS 5100 Series Blade Server Chassis are supported.

The Cisco UCS Blade Chassis is a blade-server style enclosure supporting compact, slide-in server modules, but architecturally is a significantly different approach from traditional blade server systems on the market. Most blade server systems essentially take the components that would have been in a standalone data center rack, such as a number of standardized rack-mount servers with a pair of redundant top-of-rack switches, and attempt to condense them into a single sheet-metal box. Some of these implementations even include localized storage arrays within the chassis. That approach achieves higher system density but retains most of the complexity of traditional rack systems in a smaller form factor. Also, the number of management interfaces and switching devices multiplies with each new chassis.

By extending a single low-latency network fabric directly into multiple enclosures, Cisco has removed the management complexity and cable-management issues associated with blade switching or pass-through module implementations common to blade servers. By consolidating storage traffic along this same fabric using lossless FCoE technology, Cisco UCS even further simplifies the topology by using the fabric interconnects as a common aggregation point for Ethernet data traffic and storage-specific Fibre Channel traffic. On top of this vastly simplified physical architecture, Cisco UCS Manager extends a single management interface across the physical server blades and all of their associated data and storage networking requirements.

Cisco UCS Blade Chassis System Components

The Cisco UCS Blade Chassis system has a unique architecture that integrates compute, data network access, and storage network access into a common set of components under a single-pane-of-glass management interface. The primary components included within this architecture are as follows:

- **Cisco UCS 6100 Series Fabric Interconnect:** The Cisco UCS Fabric Interconnects provide both network connectivity and management

capabilities to the other components in the system. The fabric interconnects are typically clustered together as a pair, providing resilient management access to the system as well as 10-Gigabit Ethernet, Fibre Channel, and FCoE capabilities.

- **Cisco UCS 2100 Series Fabric Extender:** The Cisco UCS 2100 Series Fabric Extenders, also referred to as I/O modules, are installed directly within the Cisco UCS 5100 Series Blade Server Chassis enclosure. These modules logically extend the fabric from the fabric interconnects into each of the enclosures for Ethernet, FCoE, and management purposes. The fabric extenders simplify cabling requirements from the blade servers installed within the system chassis.
- **Cisco UCS 5100 Series Blade Server Chassis:** The Cisco UCS 5100 Series Blade Server Chassis provides an enclosure to house up to eight half-width or four full-width blade servers, their associated fabric extenders, and four power supplies for system resiliency.

Tech Tip

As of Cisco UCS release 1.2(1d) up to eight Cisco UCS 5100 Series Blade Server Chassis may be connected to and managed as one system by a single pair of fabric interconnects.

- **Cisco UCS B-Series Blade Servers:** Cisco B-Series Blade Servers implement Intel Xeon 5500 and 5600 Series processors, and are available in both a half-width or full-width format. The Cisco UCS B200 M1 and M2 blade servers require a half-slot within the enclosure, providing high-density, high-performance computing resources in an easily managed system. The Cisco UCS B250 M1 and M2 Extended Memory Blade Servers provide up to 384 GB of memory on a single dual-socket server for memory-intensive processing such as extensive virtualization or workloads requiring large datasets.
- **Cisco UCS B-Series Network Adapters:** The Cisco UCS B-Series Blade Servers accept a variety of mezzanine adapter cards that allow the switching fabric to provide multiple interfaces to a server. These adapter cards fall into three categories:
 - Ethernet Adapters: The Cisco UCS 82598KR CI 10-GE Adapter can present up to two Ethernet interfaces to a server.
 - Converged Network Adapters: The Cisco UCS M71KR Converged

Network Adapters are available in two models, with chip sets from either Emulex or QLogic. These adapters can present up to two 10-Gbps Ethernet interfaces to a server, along with two 4-Gbps Fibre Channel interfaces.

- **Virtual Interface Cards:** The Cisco UCS M81KR Virtual Interface Card features new technology from Cisco, allowing additional network interfaces to be dynamically presented to the server. This adapter supports Cisco VN-Link technology in hardware, which allows each virtual adapter to appear as a separate Virtual Interface (VIF) on the fabric interconnects. The architecture of the Virtual Interface Card is capable of supporting up to 128 total interfaces split between vNICs and vHBAs. The specific number of interfaces currently supported is specific to the server and the installed operating system.

Cisco UCS Manager

Cisco UCS Manager is embedded software resident on the fabric interconnects, providing complete configuration and management capabilities for all of the components in the UCS system. This configuration information is replicated between the two fabric interconnects, providing a highly available solution for this critical function. The most common way to access Cisco UCS Manager for simple tasks is to use a Web browser to open the Java-based graphical user interface (GUI). For command-line or programmatic operations against the system, a command-line interface (CLI) and an XML API are also included with the system.

The Cisco UCS Manager GUI provides role-based access control (RBAC) to allow multiple levels of users granular administrative rights to system objects. Users can be restricted to certain portions of the system based on locale, which corresponds to an optional organizational structure that can be created within the system. Users can also be classified based on their access levels or areas of expertise, such as “Storage Administrator,” “Server Equipment Administrator,” or “Read-Only”. RBAC allows the comprehensive capabilities of Cisco UCS Manager GUI to be properly shared across multiple individuals or organizations within your agency in a flexible, secure manner.

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack-Mount Servers balance simplicity, performance, and density for production-level virtualization, Web infrastructure, and data center workloads. Cisco UCS C-Series servers extend Unified Computing innovations and benefits to the rack-mount server form factor. The Cisco UCS C-Series servers also implement Intel Xeon processor technology and are available in M1 (Xeon 5500) and M2 (Xeon 5600) models. The Cisco UCS C250 M1 and M2 servers also implement Cisco Extended Memory

Technology for demanding virtualization and large dataset workloads.

Third-Party Computing Systems

Third-party rack server and blade server systems may also be connected to the SBA Unified Computing topology with the available 10 Gigabit Ethernet interfaces on the Cisco Nexus 5000 Series switches, or interfaces on the Cisco Nexus 2000 Series Fabric Extenders that support 1-Gbps and 10-Gbps Ethernet connectivity, depending on the model selected. A previously installed base of running servers may be easily integrated into the environment to support existing applications and facilitate smooth migration to servers that support the Cisco Unified Computing System features.

Storage Systems

Centralized Storage Benefits

As application requirements grow, the need for additional data storage capacity also increases. This can initially cause issues when storage requirements for a given server increase beyond the physical capacity of the server hardware platform in use. As the agency grows, the investment in this additional storage capacity is most efficiently managed by moving to a centralized storage model. A centralized storage system uses Storage Area Network (SAN) technology to provide disk capacity across multiple applications and servers.

A dedicated storage system provides multiple benefits beyond raw disk capacity. SAN storage can increase the reliability of disk storage, which improves application availability. Storage systems allow increased capacity to be provided to a given server over the SAN without needing to physically attach new devices to the server itself. More sophisticated backup and data replication technologies are available in SAN storage, which helps protect the agency against data loss and application outages. This guide builds upon the design provided in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide*, which allows easy integration of centralized storage into the server farm with a choice of multiple storage-networking technology options.

Storage Access using iSCSI

Small Computer Systems Interface (SCSI) is a traditional storage protocol that was commonly used on a bus architecture directly cabled to an individual server system. Internet SCSI, or iSCSI, is an extension of that protocol over an IP network that is commonly implemented over the same Ethernet infrastructure used for client access to application servers. This approach allows growing customers to take advantage of the capabilities of centralized storage systems without investing in a separate dedicated switching infrastructure to build the SAN.

Storage Access using Fibre Channel

Fibre Channel SANs also implement a variant of the SCSI protocol, but over a dedicated, switched, fiber-based network. The switching infrastructure offers a variety of services that allow storage systems (targets) and application servers (initiators) to communicate securely over the shared switched infrastructure. Fibre Channel networks are commonly built in a redundant fashion, with separate A-side and B-side switch fabrics that allow servers to maintain continuous access to storage regardless of the failure of an individual fibre channel switch or interface.

Storage Access using NAS Protocols

Network Attached Storage (NAS) protocols access shared storage using the native protocols of traditional file-server operating systems. The most common implementations of NAS technology leverage Microsoft Common Internet File System (CIFS) or Network File System (NFS) which originated on Unix systems and has been implemented as a standard in many other operating systems. Client workstations may also access NFS shares directly as they would a typical file server system. Recently, NFS has also begun to be commonly implemented to extend server-to-server communication for centralized storage of databases and other application data.

Server Virtualization Software

Server virtualization technologies allow a single physical server to run multiple virtual instances of a guest operating system, creating virtual machines (VMs). Running multiple virtual machines on server hardware helps to increase processor utilization levels, while still allowing each VM to be viewed as independent from a security, configuration, and troubleshooting perspective.

Server Virtualization and UCS

Cisco Unified Communication System server platforms provide unique advantages that complement the implementation of server virtualization technologies. The Cisco UCS B-Series Blade Servers with Cisco UCS Manager allow the personality of a server instance to be easily ported to different physical hardware, similar to porting a virtual machine to a different host. Cisco UCS Manager provides the capability to directly integrate to the hypervisor system for dynamic network interface allocation to virtual machines. This is currently supported with VMware ESX 4.0 Update 1. The Cisco Extended Memory Technology allows individual servers to scale to large numbers of virtual machines, reducing support and licensing costs.

Hypervisor Options

Cisco UCS servers have been certified with hypervisor systems including VMware ESX, Microsoft Hyper-V, and Citrix Xen.

Please contact your Cisco Systems or authorized partner sales representative to verify the specifics of your implementation requirements with shipping hardware and software versions.

Deploying the SBA Unified Computing Architecture

The following sections provide detailed, step-by-step instructions to configure the basic elements of the SBA Unified Computing model. Common best-practices configurations are shown to allow a new user to quickly configure a new system for basic operations. This is also a flexible configuration, so additional information is provided, including pointers to more detailed documentation that is useful for more advanced system configurations.

Ethernet Network Infrastructure

The Ethernet network infrastructure for the SBA Unified Computing topology is based on the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide*. Cisco UCS C-Series Rack-Mount Servers may be connected to this infrastructure using available interfaces on the Cisco Nexus 5000 Series Switches or through the Cisco Nexus 2000 Series Fabric Extenders. Switching access or trunk port modes may be configured according to the settings appropriate for the installed operating system.

The Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis operate in conjunction with the Cisco UCS 6100 Series Fabric Interconnects to appear as a group of end-node servers to the data center Ethernet switching fabric. As of Cisco UCS Release 1.2(1b), the Cisco UCS 6100 Series Fabric Interconnects support 10-Gigabit Ethernet uplink connection into the switching fabric only. In the SBA Unified Computing architecture, the fabric interconnects are connected directly to the Cisco Nexus 5000 Series Ethernet switching fabric running Virtual Port Channel (vPC). vPC allows the two Cisco Nexus 5000 Series switches to appear as a single Layer-2 switching instance to attached devices with two or more interfaces configured as a port-channel. Cisco UCS 6100 Series Fabric Interconnects implement port-channel capability that uses the Link Aggregation Control Protocol (LACP).

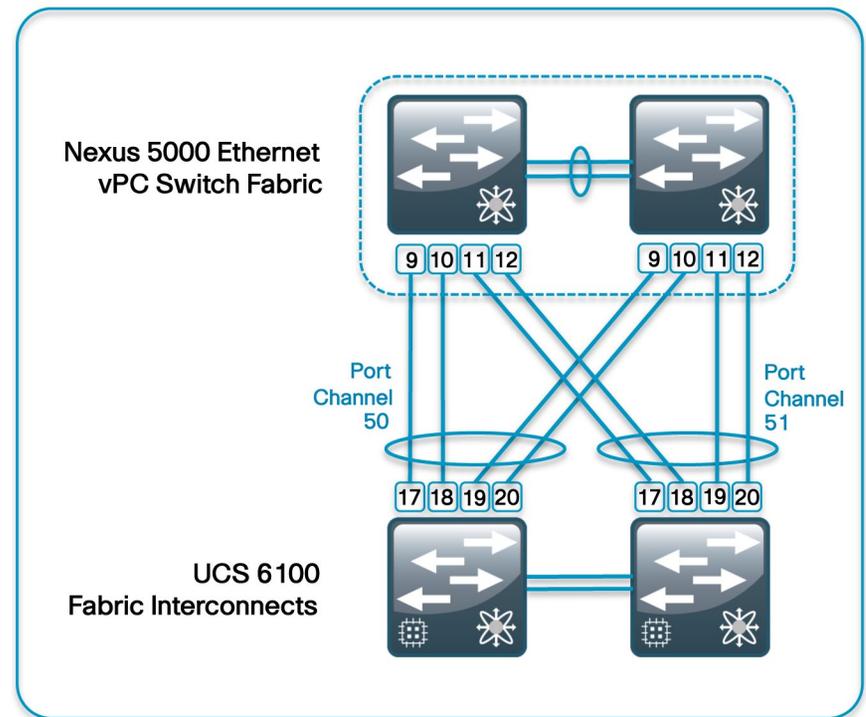
Configuration examples in this guide show the use of a port-channel with four physical 10-Gigabit Ethernet ports from each Cisco UCS 6100 Series Fabric Interconnect to the Cisco Nexus 5000 vPC pair. These interfaces

are numbered Ethernet 1/9 through 1/12 on each Cisco Nexus 5000 Series switch, and ports 17 through 20 on each fabric interconnect in the example configurations. The port channel from each fabric interconnect spans the two physical Cisco Nexus 5000 switches for resilient connectivity, as shown in Figure 4. You can use interface numbers specific to your implementation to achieve the same cabling approach.

Tech Tip

This example shows the use of integrated ports in the validation network for Ethernet Uplink connections. Expansion module Ethernet ports may also be used as uplink ports. Expansion module ports cannot be used as server ports.

Figure 4. UCS Fabric Interconnect Ethernet Detail



Process

Configuring the Port Channels

1. Configure Port Channels on Cisco Nexus 5000 Switches

For the best combination of throughput and resiliency, you should configure the links between the fabric interconnects and the Ethernet switching fabric as port channels.

Procedure 1 Config Port Channels on Nexus 5K Switches

Step 1: To configure the port channels on the Cisco Nexus 5000 switches, begin by creating the vPC port-channel interfaces on both switches.

```
interface port-channel50
  switchport mode trunk
  vpc 50
  spanning-tree port type edge trunk

interface port-channel51
  switchport mode trunk
  vpc 51
  spanning-tree port type edge trunk
```

Tech Tip

Setting the Spanning-Tree port type to “edge trunk” is appropriate for the default Fabric Interconnect configuration of End Host Mode. If the Fabric Interconnect is configured in switched mode, the Nexus 5000 port type should be left as “normal” for standard spanning-tree protocol loop prevention.

Step 2: After you have created the port channels, assign physical interfaces to the port channels as follows on the two Cisco Nexus 5000 switches, configured in vPC mode:

```
interface Ethernet1/9
  switchport mode trunk
  channel-group 50 mode active

interface Ethernet1/10
  switchport mode trunk
  channel-group 50 mode active

interface Ethernet1/11
  switchport mode trunk
  channel-group 51 mode active

interface Ethernet1/12
  switchport mode trunk
  channel-group 51 mode active
```

The port-channel interfaces will not become active until you complete the corresponding configuration on the Cisco UCS 6100 Series Chassis, which is covered in the Define the Ethernet Uplink Ports procedure in the Getting Started with UCS Manager Process.

Fibre Channel Network Infrastructure

Complete the following process to prepare a fibre channel SAN to support the system. Configuration instructions provided in this guide are based on the foundation of the Fibre Channel infrastructure in the SBA Data Center for Midsize Agencies topology.

Process

Configuring the Fibre Channel Network Infrastructure

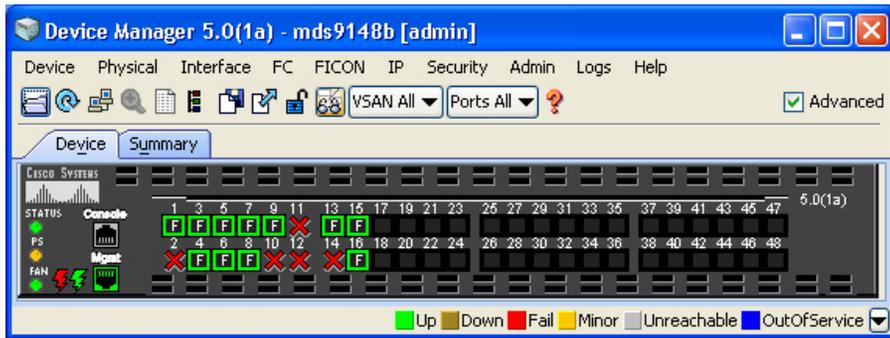
1. Prepare the Cisco MDS 9100s

Procedure 1 Prepare the Cisco MDS 9100

To prepare the Cisco MDS 9100 Series Fibre Channel switch fabrics for connecting the fabric interconnect, NPIV must be activated.

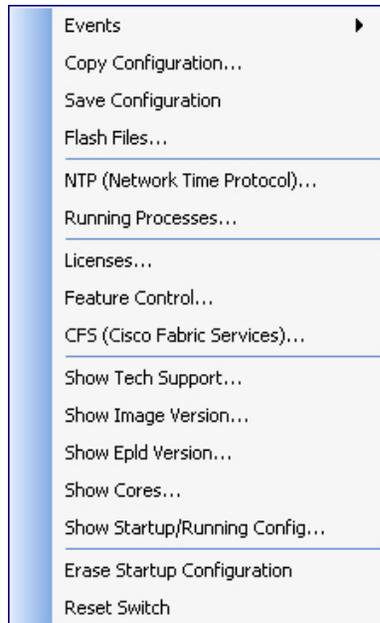
Step 1: Open Device Manager to the MDS as shown in Figure 5.

Figure 5. Device Manager



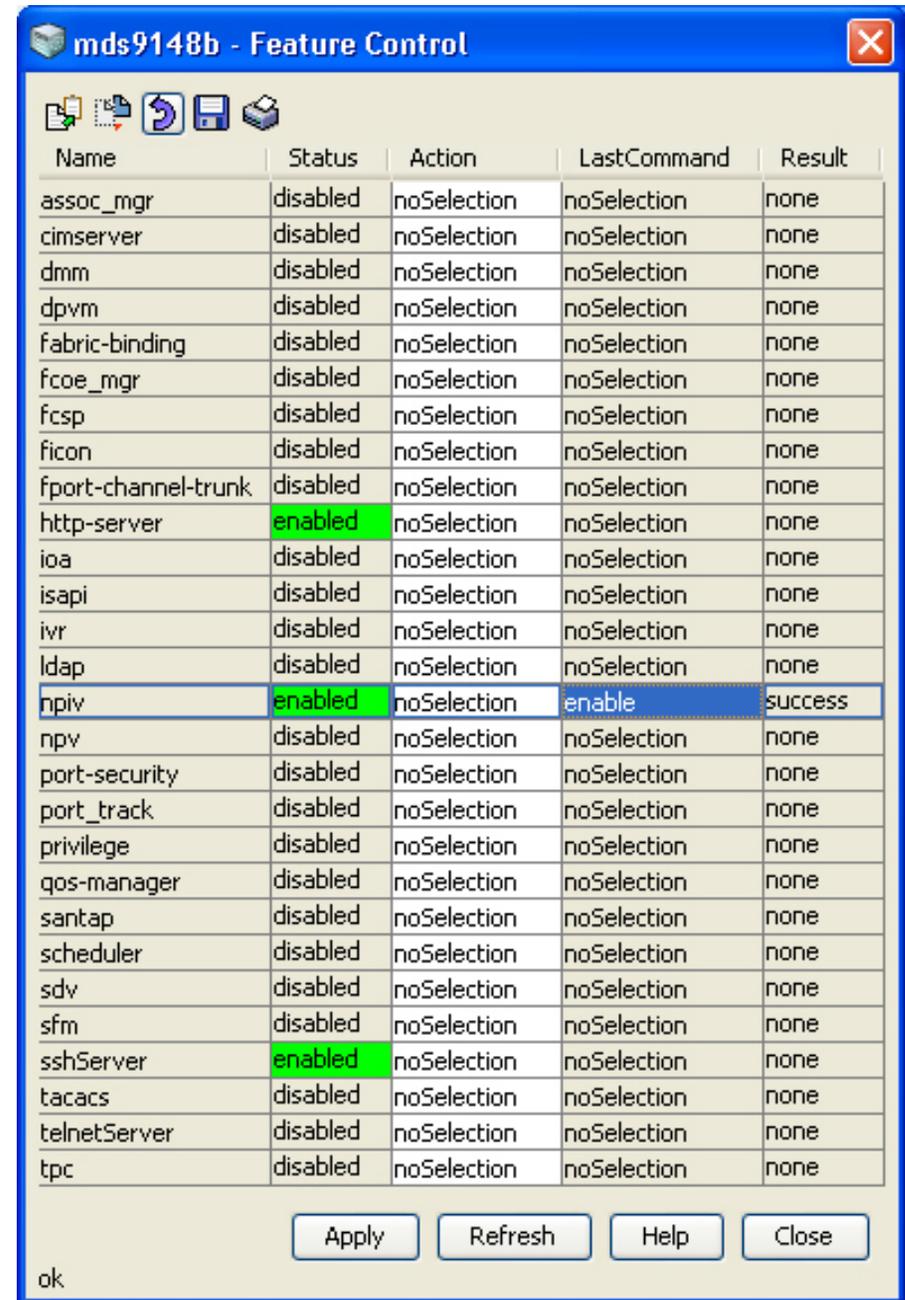
Step 2: From the Admin menu, shown in Figure 6, select Feature Control.

Figure 6. MDS Admin Menu



Step 3: Next to NPIV, select Enable as shown in Figure 7, and click Apply.

Figure 7. Feature Control



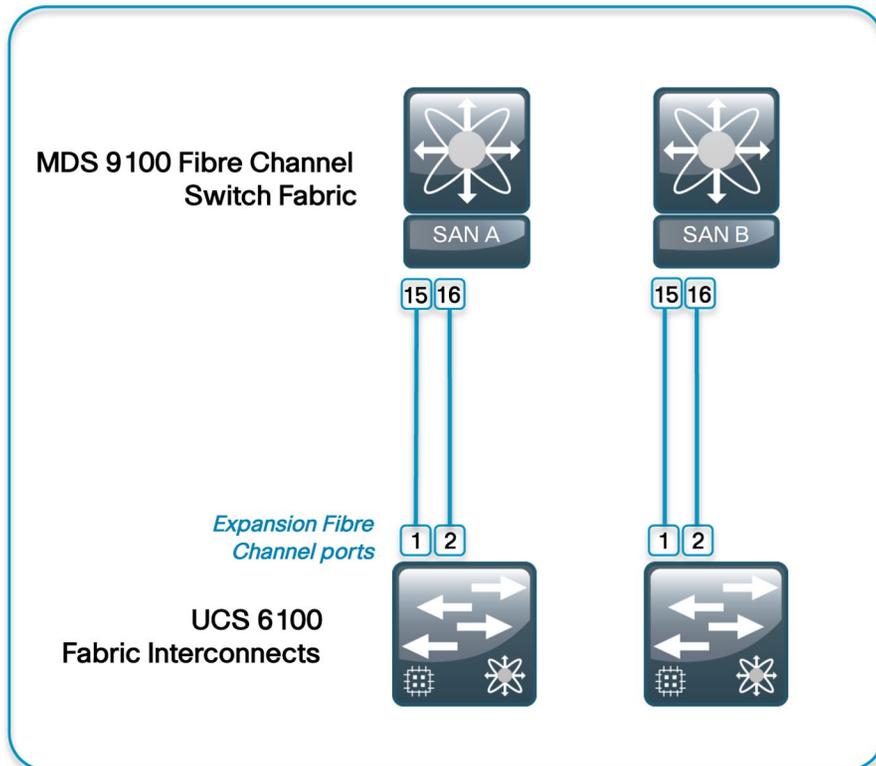
By default, the MDS ports are set to automatically negotiate speed. With NPIV enabled, you must assign a VSAN to the interface that connects to the fabric interconnect.

To assign a VSAN to the interface for the fabric interconnects, follow the instructions on activating an interface in the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide*. Cisco MDS configuration will need to be done in both fabrics A and B that connect to the fabric interconnects. Figure 8 illustrates the physical connectivity between the fabric interconnects and the fibre channel switch fabrics.

Tech Tip

The Fibre Channel interface will not show a status of up until further configuration of the fabric interconnects Fibre Channel interfaces.

Figure 8. UCS Fabric Interconnect Fibre Channel Detail



UCS Blade Server System

The Cisco UCS Blade Server system is the heart of the SBA Unified Computing architecture. This section provides information on initial system setup and basic service profile configuration to prepare your first running server to boot on one of the blade server modules. Additional information is provided on setting up service profiles with multiple interfaces, boot-from-LAN configurations, and boot-from-SAN configurations.

Process

Completing the Initial System Setup

1. Complete Physical Setup and Ensure Connectivity
2. Complete Initial Fabric Interconnect Setup

Procedure 1

Complete Physical Setup/Ensure

The Cisco UCS Fabric Interconnect acts as the concentration point for all cabling to and from the UCS Blade Chassis.

Step 1: Connect the two fabric interconnects together using the integrated ports labeled L1/L2. These ports are used for replication of cluster information between the two fabric interconnects, not the forwarding of data traffic.

Step 2: Attach the Management Ethernet ports from each fabric interconnect to a management network or appropriate Ethernet segment where they can be accessed for overall administration of the system.

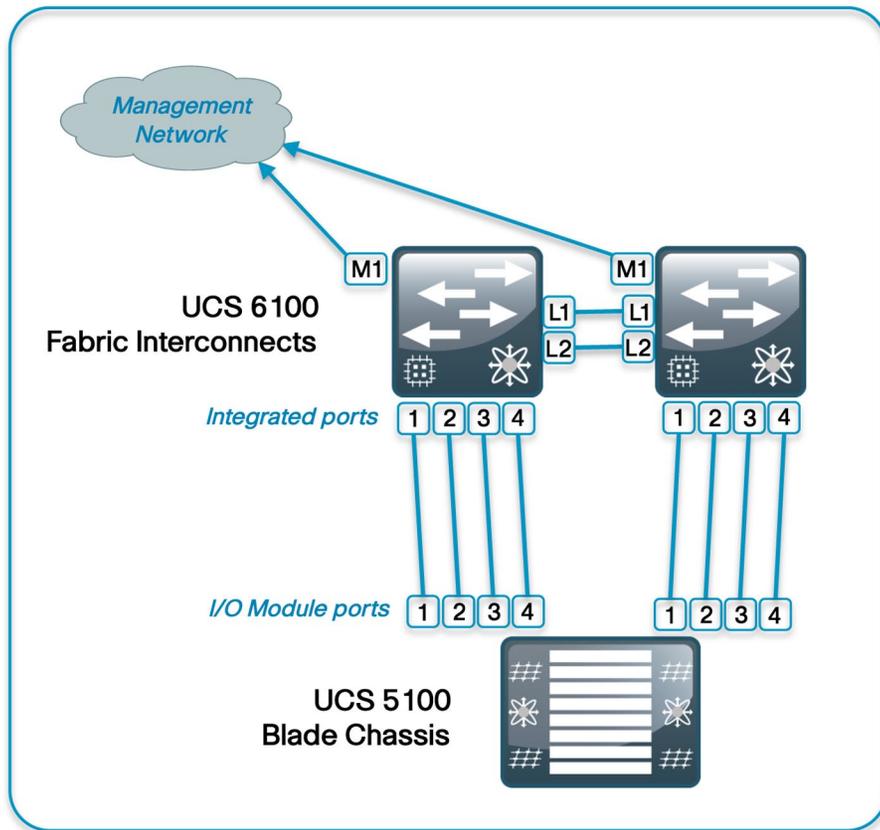
Step 3: Populate each blade chassis with two fabric extenders (I/O modules) to provide connectivity back to the fabric interconnects.

Step 4: Cable one I/O module to the first fabric interconnect. Cable the other I/O module to the second fabric interconnect. After you have configured the fabric interconnects, they will be designated as “A” and “B” fabrics.

You can connect the I/O modules to the fabric interconnects by using one, two, or four cables per module. For system resiliency and throughput we recommend a minimum of two connections per I/O module.

Figure 9 shows an example of completed connections for a single chassis system.

Figure 9. UCS Blade Chassis and Fabric Connections



Tech Tip

Ports 1 through 4 on the fabric interconnects are shown as an example. Additional blade chassis may be connected via their integrated I/O modules into any of the ports on the fabric interconnect. Each I/O module contains only four ports, labeled 1 through 4.

Procedure 2

Complete Initial Fabric Interconnect Setup

You can easily accomplish the initial configuration of the fabric interconnects through the Basic System Configuration Dialog that launches when you power on a new or un-configured unit.

Step 1: Connect a terminal to the console port of the first system to be configured and press Enter.

Step 2: In the Basic System Configuration Dialog that follows, enter console, setup, and yes, and then establish a password for the admin account.

```
---- Basic System Configuration Dialog ----
```

```
This setup utility will guide you through the basic
configuration of the system. Only minimal configuration
including IP connectivity to the Fabric interconnect and its
clustering mode is performed through these steps.
Type Ctrl-C at any time to abort configuration and reboot
system. To back track or make modifications to already entered
values, complete input till end of section and answer no when
prompted to apply configuration.
```

```
Enter the configuration method. (console/gui) ? console
```

```
Enter the setup mode; setup newly or restore from backup.
```

```
(setup/restore) ? setup
```

```
You have chosen to setup a new Fabric interconnect. Continue?
```

```
(y/n) : y
```

```
Enter the password for "admin": xxxxxxxx
```

```
Confirm the password for "admin": xxxxxxxx
```

Step 3: Next you are prompted to create a new cluster or add to an existing cluster. The Cisco UCS cluster consists of two fabric interconnects, with all associated configuration replicated between the two for all devices in the system. Enter "yes" to create a new cluster.

```
Do you want to create a new cluster on this Fabric interconnect
(select 'no' for standalone setup or if you want this Fabric
interconnect to be added to an existing cluster)? (yes/no)
[n] : yes
```

Step 4: Each fabric interconnect has a unique physical IP address. There is a shared cluster IP address that is used to access Cisco UCS Manager after the system initialization is completed. The fabric interconnects are assigned one of two unique fabric IDs for both Ethernet and Fibre Channel networking. Choose fabric A for the first fabric interconnect that you are setting up.

```
Enter the switch fabric (A/B) [] : a
```

Step 5: The system name is shared across both fabrics, so "-a" or "-b" is

automatically appended to the name that you specify in the Basic System Configuration Dialog when you set up one of the units.

Enter the system name: **sba-ucs-10**

Step 6: Apply the following example as you respond to the prompts.

```
Physical Switch Mgmt0 IPv4 address : 192.168.28.51
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 192.168.28.1
Cluster IPv4 address : 192.168.28.50
Configure the DNS Server IPv4 address? (yes/no) [n]: n
Configure the default domain name? (yes/no) [n]: n
```

Step 7: The Basic System Configuration Dialog displays a summary of the configuration options that you chose. Verify the accuracy of the settings. Unless the settings require correction, enter “yes” to apply the configuration. The system assumes the new identity that you configured.

Following configurations will be applied:

```
Switch Fabric=A
System Name=eng-cbc-ucs
Physical Switch Mgmt0 IP Address=192.168.28.51
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=192.168.28.1
Cluster Enabled=yes
Cluster IP Address=192.168.28.50
```

```
Apply and save the configuration (select 'no' if you want to
re-enter)? (yes/no): yes
Applying configuration. Please wait.
Configuration file - Ok
```

Step 8: After the system is reset, you can add the second fabric interconnect to the cluster. Because you have already defined the cluster, you only need to acknowledge the prompts to add the second fabric interconnect to the cluster and set a unique IP address.

```
Enter the configuration method. (console/gui) ? console
Installer has detected the presence of a peer Fabric
interconnect. This Fabric interconnect will be added to the
cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect...done
Peer Fabric interconnect Mgmt0 IP Address: 192.168.28.51
Peer Fabric interconnect Mgmt0 IP Netmask: 255.255.255.0
Cluster IP address: 192.168.28.50
Physical Switch Mgmt0 IPv4 address : 192.168.28.52
Apply and save the configuration (select 'no' if you want to
```

```
re-enter)? (yes/no): yes
Applying configuration. Please wait.
Configuration file - Ok
```

From this point forward, this document primarily shows the use of the GUI for management of the system; however, you should be familiar with the console in case you need very low-bandwidth remote access or a separate mode of access for administrative tasks such as code upgrades or system troubleshooting.

Process

Getting Started with UCS Manager

1. Launch UCSM
2. Discover System Components
3. Define Ethernet Uplink Ports
4. Define Fibre Channel Uplink Ports
5. Add a Management IP Address Pool

Cisco UCS Manager (UCSM) is the management service for all of the components in a Cisco UCS instance. Cisco UCS Manager runs on the fabric interconnects, and keeps configuration data synchronized between the resilient pair. The primary access method covered here for using Cisco UCS Manager is the GUI client, which is Java based and is launched from a Web browser, using Java Web Start.

Any computer that you want to use to run the Cisco UCS Manager client must meet or exceed the following minimum system requirements:

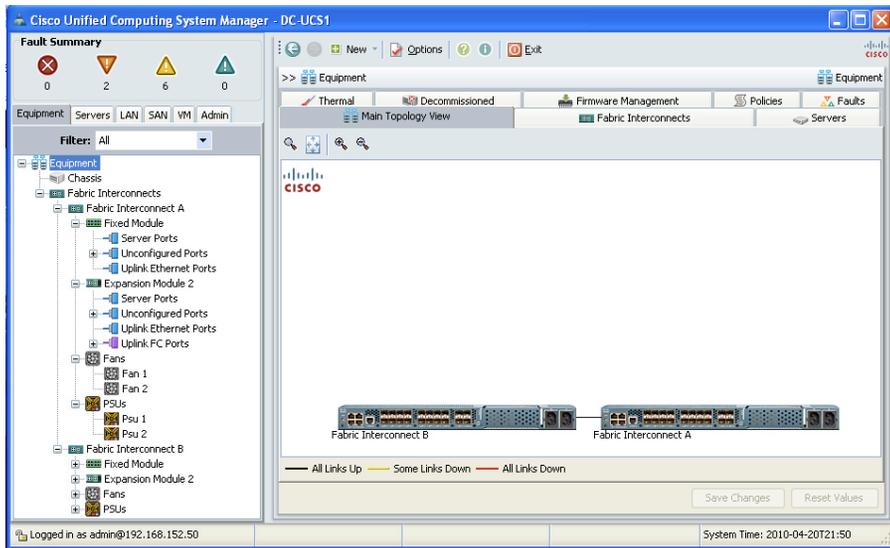
- The computer must have Sun JRE 1.6 or later installed.
- The computer must have one of the following Web browsers installed:
 - Microsoft Internet Explorer 6.0 or higher
 - Mozilla Firefox 3.0 or higher
- The computer must run one of the following operating systems:
 - Microsoft Windows XP
 - Microsoft Windows Vista
- Red Hat Enterprise Linux 5.0 or higher

Procedure 1

Launch UCS Manager

Step 1: Using a browser, access the cluster IP address that you assigned during initial setup, and choose **Launch** to download the UCSM Java application. Authenticate with the configured username and password and view the initial screen, as shown in Figure 10.

Figure 10. UCS Manager Initial Screen



The Cisco UCS Manager GUI consists of a navigation pane on the left side of the screen and a work pane on the right side of the screen. The navigation pane allows you to browse through containers and objects and to drill down easily through layers of the system management. In addition, the following tabs appear across the top of the navigation pane:

- **Equipment:** Inventory of hardware components and hardware-specific configuration
- **Servers:** Service profile configuration and related components such as policies and pools
- **LAN:** LAN-specific configuration for Ethernet and IP networking capabilities
- **SAN:** SAN-specific configuration for Fibre Channel networking capabilities

- **VM:** Configuration specific to linking to external server virtualization software, currently supported for VMware.
- **Admin:** User management tasks, fault management and troubleshooting.

The tabs displayed in the navigation pane are always present as you move through the system, and in conjunction with the tree structure shown within the pane itself, are the primary mechanisms for navigating the system.

After you choose a section of the GUI in the navigation pane, information and configuration options appear in the work pane on the right side of the screen. In the work pane, tabs divide information into categories. The work pane tabs that appear vary according to the context chosen in the navigation pane.

Procedure 2

Discover System Components

On a newly installed system, one of your first tasks in the Cisco UCS Manager GUI is to define which ports on the fabrics are attached to the I/O modules in each chassis. This allows Cisco UCS Manager to discover the attached system components and build a view of the entire system. These ports are referred to as server ports.

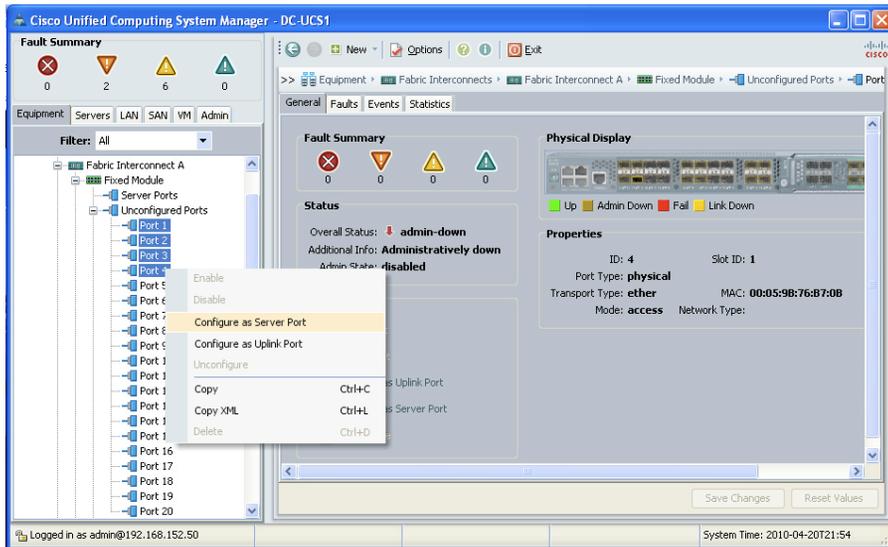
Step 1: In the navigation pane, choose the **Equipment** tab and then expand **Fabric Interconnects > Fabric Interconnect A > Fixed Module > Unconfigured Ports**.

Objects are displayed representing each of the physical ports on the base fabric interconnect system.

Step 2: Select the desired port by clicking the port object. Optionally, choose several sequential ports by clicking a second port while holding the shift key.

Step 3: Right-click the selected port or group of ports and choose **Configure as Server Port** from the pop-up menu as shown in Figure 11.

Figure 11. Configure Server Ports



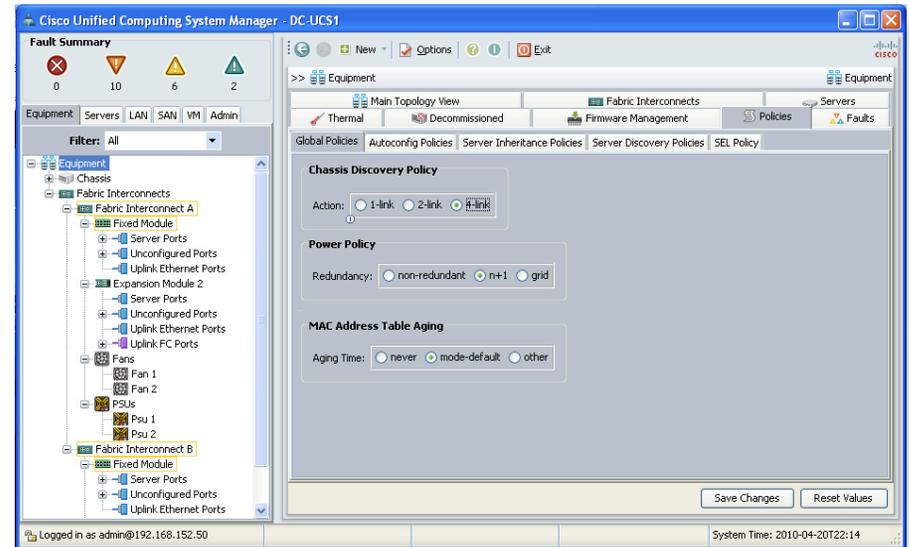
Step 4: Acknowledge this operation.

Step 5: In a similar manner, expand the tree to **Fabric Interconnect B**, and apply the corresponding configuration for the resilient links from Fabric B.

Step 6: Next, configure the number of physical links used to connect each chassis to each fabric. Click the **Equipment** tab in the navigation pane, and then choose the **Policies** tab in the work pane. Within the **Policies** tab, another set of tabs appears.

Step 7: By default, the **Global Policies** tab displays the Chassis Discovery Policy. This may be set at 1, 2, or 4 links per fabric, as shown in Figure 12. Choose the appropriate number of links for your configuration, and then click **Save Changes** at the bottom of the work pane.

Figure 12. Set Chassis Discovery Policy



Tech Tip

If the system gets out of synchronization for any reason during the chassis discovery process, you can clear up most issues by acknowledging the chassis. Right-click the chassis in the navigation pane and choose “Acknowledge Chassis”.

After Cisco UCS Manager has discovered each of the chassis attached to your system, you can use the Equipment tab in the navigation pane to verify that each chassis, I/O module, and server is properly reflected in the display.

Procedure 3 Define Ethernet Uplink Ports

In the SBA Unified Computing reference design, Ethernet uplink ports connect the fabric interconnects to the Cisco Nexus 5000 switches via 10-Gigabit Ethernet links. In alternate designs, these links may be attached to any 10-Gigabit Ethernet switch that provides access to the core of the network. These links carry IP-based client-server traffic, server-to-server traffic between IP subnets, and Ethernet-based storage access such as

iSCSI or NAS traffic. Ports from either the base fabric interconnect or expansion modules may be used as uplink ports.

Step 1: In the **Equipment** tab of the navigation pane, locate the ports that are physically connected to the upstream switches. These ports should initially be listed as unconfigured ports in the tree view.

Reference the Ethernet connectivity detail shown in Figure 4 or the specific ports selected for your implementation.

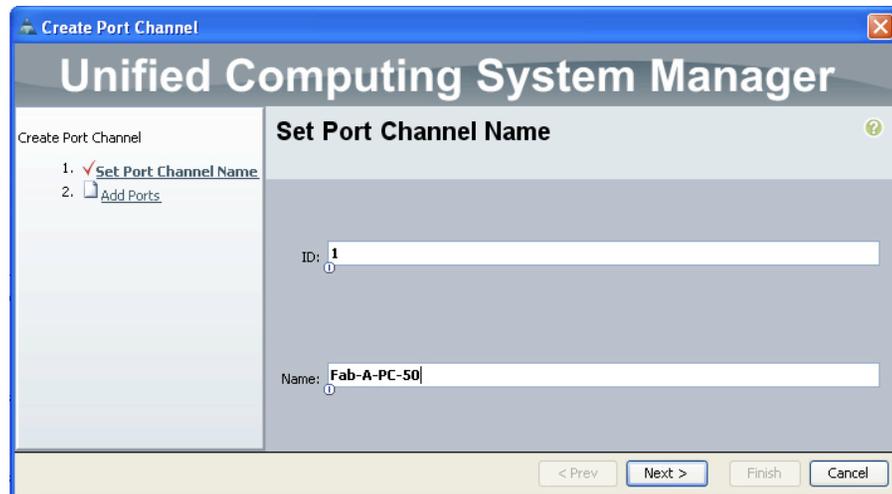
Step 2: For each port, right-click and choose **Configure as Uplink Port**.

Step 3: If you implemented port-channel configuration in the upstream switches such as the Cisco Nexus 5000 Series switches in this example, corresponding port-channel configuration is required for the Ethernet uplink ports in the Cisco UCS Manager GUI. Choose the **LAN** tab in the navigation pane, and expand **LAN > LAN Cloud > Fabric A**, and select the **Port Channels** container.

Step 4: Click the **Add** button (marked with a green plus sign) to create a new port-channel definition.

Step 5: Provide an ID and name for the new port channel, as shown in Figure 13.

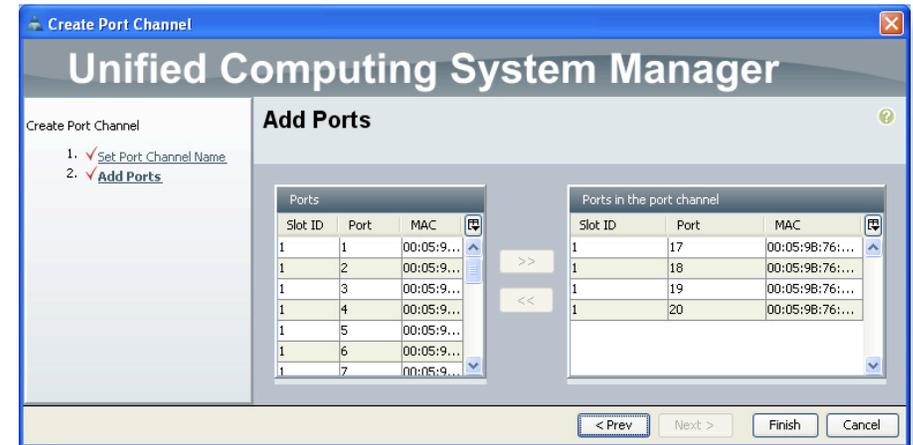
Figure 13. Set Port Channel Name



Step 6: Click **Next** and then select the Ethernet ports in use as uplinks from the list on the left of the screen as shown in Figure 14.

Step 7: Use the arrow button to add them to the list of ports in the channel.

Figure 14. Add Ports to the Port Channel



Tech Tip

Pay close attention to the Slot ID column when you select the ports to add to the channel. Integrated ports will be listed with a slot ID of 1. If using an expansion module, scroll down to find ports listed with a slot ID of 2.

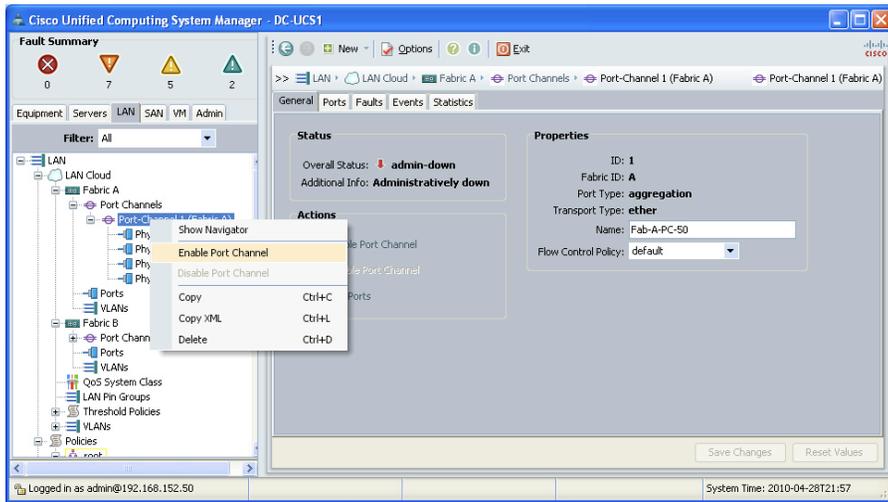
Step 8: Click **Finish** to complete creation of the Ethernet uplink port channel for Fabric A.

Step 9: Repeat this procedure to create a port channel for Fabric B, using a unique port-channel ID value.

Step 10: After you have created the port channels, you must enable them before they will become active. In the navigation pane, expand **Port Channels**.

Step 11: For each port channel, select and then right-click the port channel name, and choose **Enable Port Channel** as shown in Figure 15.

Figure 15. Enable the Port Channels



Tech Tip

Port channel IDs are locally significant to each device; therefore, as shown, the ID used to identify a port channel to the fabric interconnect does not have to match the ID used for the channels on the Cisco Nexus 5000 configuration. In some cases, it may be beneficial for operational support to use consistent numbering for representation of these channels.

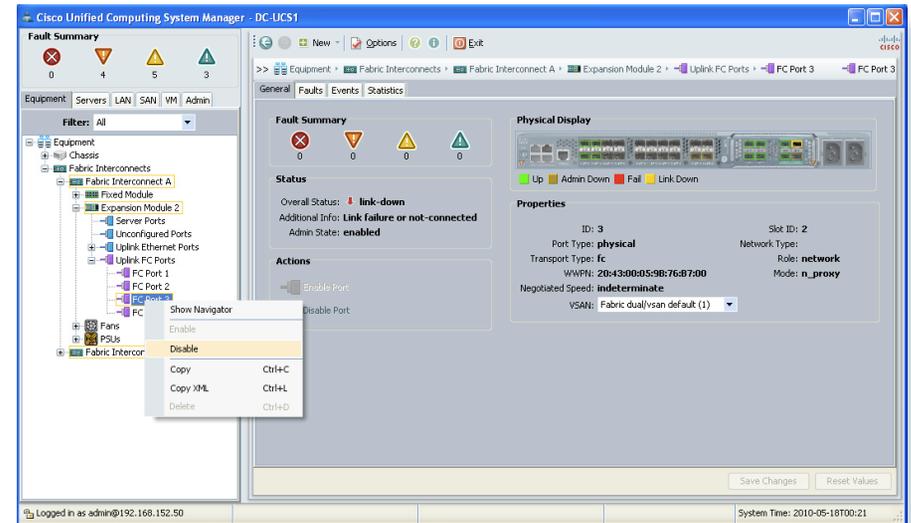
Procedure 4 Define Fibre Channel Uplink Ports

Fibre Channel uplink ports connect the fabric interconnects to the SAN, which in the SBA Unified Computing reference topology is built using Cisco MDS 9100 Series Switches. These links carry native Fibre Channel traffic to the SAN, and must be provisioned on expansion modules because there are no Fibre Channel ports on the base fabric interconnects. Because these ports can only be used as uplinks and not server ports, they automatically appear in the Uplink FC Ports folder in the navigation pane underneath the fabric interconnect expansion modules. The Fibre Channel ports are enabled by default.

Step 1: Connect the desired ports to the SAN and disable the unused ports to clear any system alerts tied to the unused ports.

Step 2: To disable a port, right-click the port name in the navigation pane and choose **Disable**, as shown in Figure 16 or select **Disable Port** in the **Actions** area of the **General** tab in the work pane.

Figure 16. Disable Unused Fibre Channel Ports



Step 3: You must create a VSAN and assign it to the Fibre Channel port to activate the port and transmit traffic. This VSAN should match the VSAN configured on the corresponding SAN fabric. Choose the **SAN** tab in the navigation pane and then expand **SAN > SAN Cloud > VSANs**.

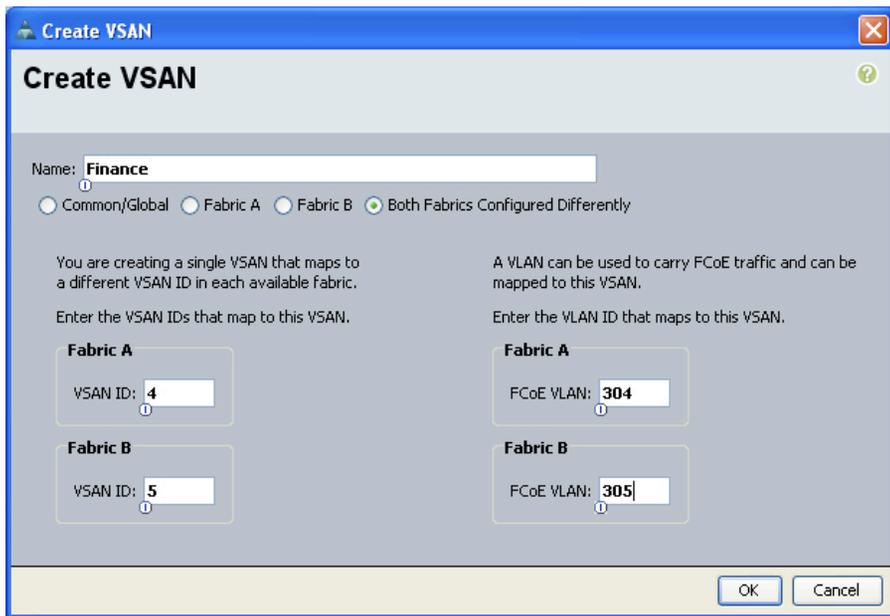
Step 4: Right-click the VSAN container and choose **Create VSAN**.

Step 5: Enter a description for the VSAN and select **Both Fabrics Configured Differently** as shown in Figure 17.

Step 6: Enter the VSANs corresponding to the primary and secondary VSANs configured in your SAN fabrics.

Step 7: For each fabric, enter the VLAN that the Fibre Channel traffic should use from the chassis to the fabric interconnects. In Figure 17, VSAN 4 on Fabric A corresponds to VLAN 304 on the fabric interconnect.

Figure 17. Create the VSANs

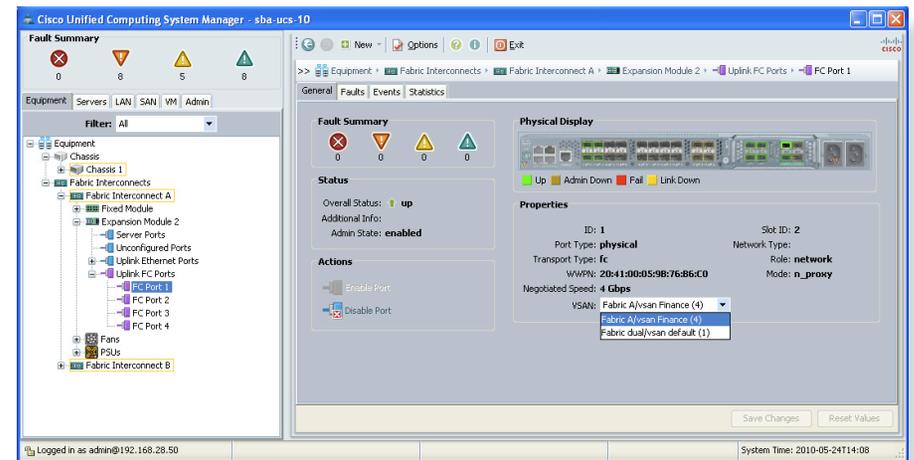


Step 8: When you have configured the VSAN IDs in this section, click OK. A window shows the successful creation of the VSAN.

Step 9: Now that you have created the VSAN, it can be assigned to the Fibre Channel uplinks. On the **Equipment** tab in the navigation pane, expand **Fabric Interconnects > Fabric Interconnect A > Expansion Module 2** and select the Fiber Channel uplink port.

Step 10: On the right side of the pane, from the VSAN drop-down list, select the VSAN that you created for the SAN Fabric A as shown in Figure 18 and click **Save Changes**.

Figure 18. Select VSANs for Fibre Channel Uplinks



Step 11: Repeat this procedure for the VSAN for SAN Fabric B on Fabric Interconnect B.

Tech Tip

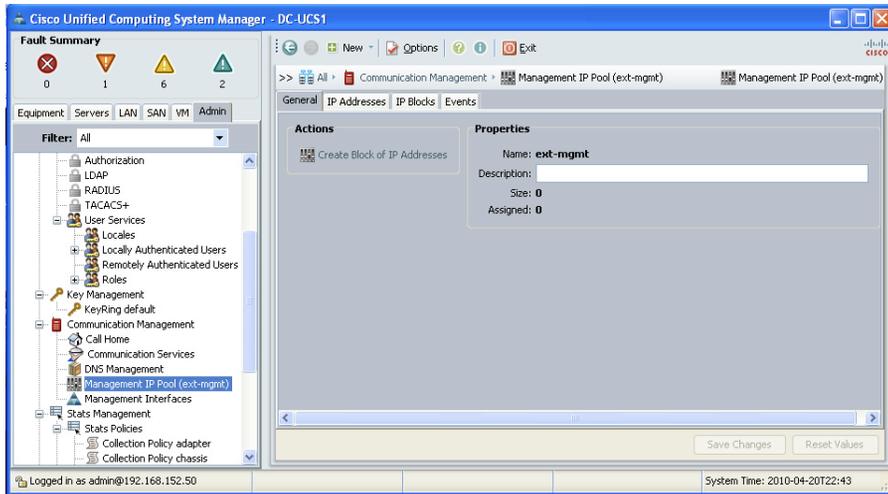
If you will access all of your storage strictly over Ethernet by using iSCSI or NAS protocols, it is not necessary to define or attach Fibre Channel uplinks and you can skip the Fibre Channel Uplinks procedure.

Procedure 5 Add a Management IP Address Pool

The Cisco UCS Manager GUI provides a launching point to direct Keyboard-Video-Mouse (KVM) access to control each of the blades servers within the system. To facilitate this remote management access, you must allocate a pool of IP addresses to the blade servers within the system. These addresses are used by the Cisco UCS KVM Console application to communicate with the individual blade servers. You must allocate this pool of addresses from the same IP subnet as the addresses assigned to the management interfaces of the fabric interconnects, because a common default gateway is used for their communication.

Step 1: Choose the **Admin** tab from the navigation pane, and expand **All > Communication Management**, and select **Management IP Pool**,

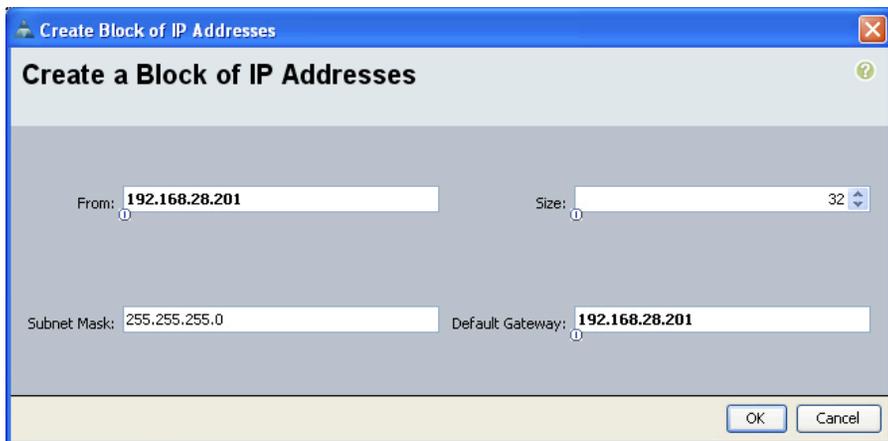
Figure 19. Add Management IP Pool



Step 2: In the work pane, under **Actions**, click **Create Block of Addresses**. As shown in Figure 19.

Step 3: Allocate a contiguous block of IP addresses by specifying the starting address in the **From** field, and then use the corresponding fields to specify the **Size** of the block, the **Subnet Mask** and the **Default Gateway**.

Figure 20. Create Management IP Block



Step 4: Click **OK** to finish creating the block of addresses.

Tech Tip

After you complete the initial setup, ensure that the system firmware is updated to the most current version or to the version recommended for your installation. Detailed information on upgrading firmware is available at: http://www.cisco.com/en/US/products/ps10281/prod_installation_guides_list.html.

Process

Creating an Initial Service Profile for Local Boot

1. Access the Service Profile Wizard
2. Create UUIDs
3. Configure Storage
4. Complete Networking Configuration
5. Define the Server Boot Order Policy
6. Assign the Server

One of the core concepts of Cisco UCS Manager is the service profile. A service profile defines all characteristics that are normally presented by a physical server to a host operating system or a hypervisor, including presence of network interfaces and their addresses, host adapters and their addresses, boot order, disk configuration, and firmware versions. The profile can be assigned to one or more physical server blades within the chassis. In this way, what is traditionally thought of as the personality of a given server or host is tied to the service profile rather than to the physical server blade where the profile is running. This is particularly true if network-based or SAN-based boot is configured for the profile. If local-boot is configured for the profile, the boot images installed on the local hard drives of the physical blade do tie the identity of the service profile to a given physical server blade.

There are multiple supporting objects within the Cisco UCS Manager GUI to streamline the creation of a service profile. These objects contain items such as pools of MAC addresses for Ethernet, World Wide Port Names (WWPNs) for Fibre Channel, disk configurations, VLANs, VSANs, etc. These objects are stored by the system so that they may be referenced by multiple service profiles, so that you do not need to redefine them as each new profile is created. Walking through the process of creating an initial service profile on a brand-new system may be used as a system-initialization exercise, providing a structured path through the process and the option to create these reusable configuration objects along the way to be referenced by additional service profiles in the future.

This process provides an example of how to create a basic service profile for initial installation and boot of a host operating system or a hypervisor. Throughout this process, we explore the creation of reusable system objects to facilitate faster creation of additional profiles with similar characteristics. For simplicity, configuration of a basic boot policy using local mirrored disks is shown. Later sections in this guide show options for network-based or SAN-based boot. The configuration of this initial profile creates the base system setup upon which additional, more advanced profiles can be built.

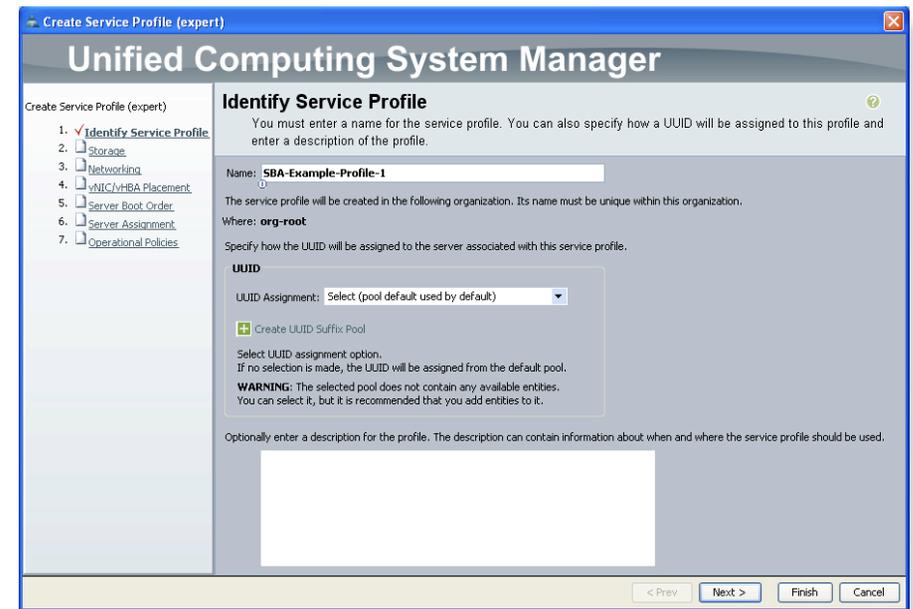
Procedure 1 Access the Service Profile Wizard

Step 1: Select the **Servers** tab in the navigation pane, expand the containers underneath **Service Profiles**, and select the **Root** container.

Step 2: On the **General** tab within the work pane, click on **Create Service Profile (expert)**. An initial pop-up window displays, as shown in Figure 21. Assign a name to the service profile. The following sections will walk you through a wizard-based process for defining these attributes of the first service profile:

- Identification/UUIDs
- Storage
- Networking
- vNIC/vHBA Placement
- Server Boot Order
- Server Assignment
- Operational Policies

Figure 21. Initial Service Profile (Expert) Screen



Procedure 2 Create UUIDs

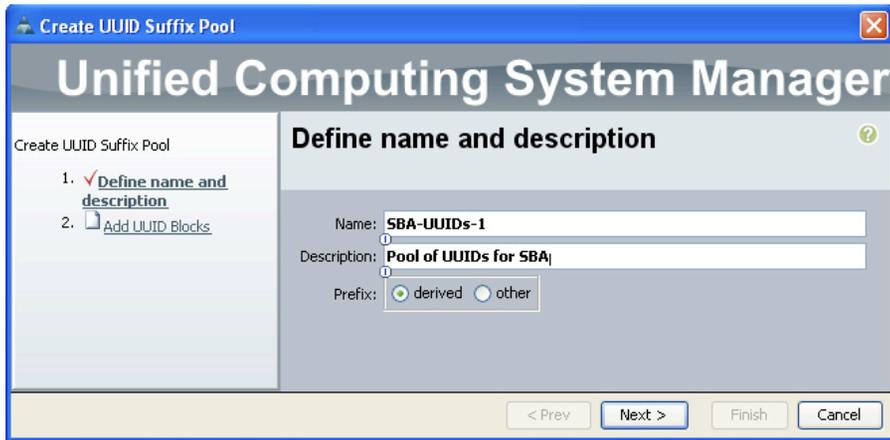
Step 1: Click **Create UUID Suffix Pool** to add a Universally Unique Identifier (UUID) suffix pool to the system.

Tech Tip

A UUID suffix pool is a collection of SMBIOS UUIDs that are available to be assigned to servers. The first number of digits that constitute the prefix of the UUID are fixed. The remaining digits, the UUID suffix, are variable. A UUID suffix pool avoids conflicts by ensuring that these variable values are unique for each server associated with a service profile that uses that particular pool.

Step 2: Assign a name and description to the pool, as shown in Figure 22. To use a UUID prefix which is derived from hardware, leave **derived** selected in the Prefix box.

Figure 22. Set UUID Pool Name and Description



Step 3: Click **Next** to proceed to the step of defining the UUID Blocks to be included in the pool, and click **Add**.

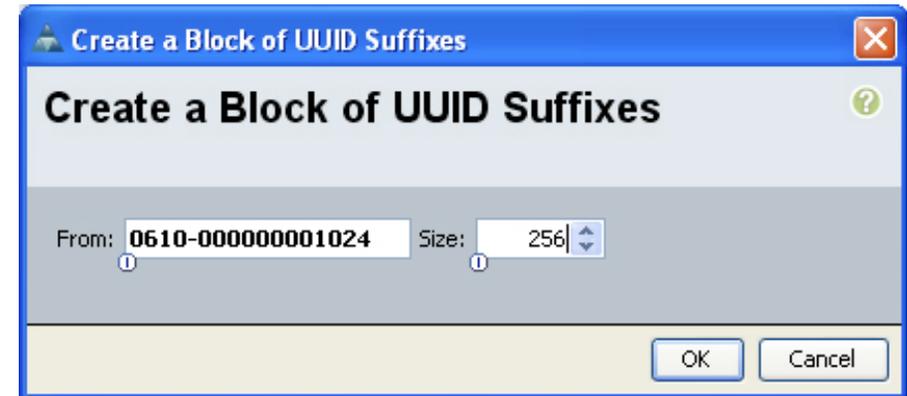
Step 4: In the **From** field, enter a unique, randomized base value as a starting point as shown in Figure 23.

Tech Tip

UUID generation tools compliant with RFC 4122 are available on the Internet. For an example, see <http://www.famkruihof.net/uuid/uuidgen>.

Step 5: Set the **Size** field to exceed the number of servers or service profiles that you will require to use the same pool. If future expansion is required, you can add multiple UUID suffix blocks to the same pool. For a base system startup, a simple, small pool is sufficient.

Figure 23. Create UUID Suffix Blocks



Step 6: After you enter the starting point and size of the suffix block, click **OK** to create the suffix block and **Finish** to complete the creation of the pool.

Step 7: Then click **Next** to proceed to the storage section.

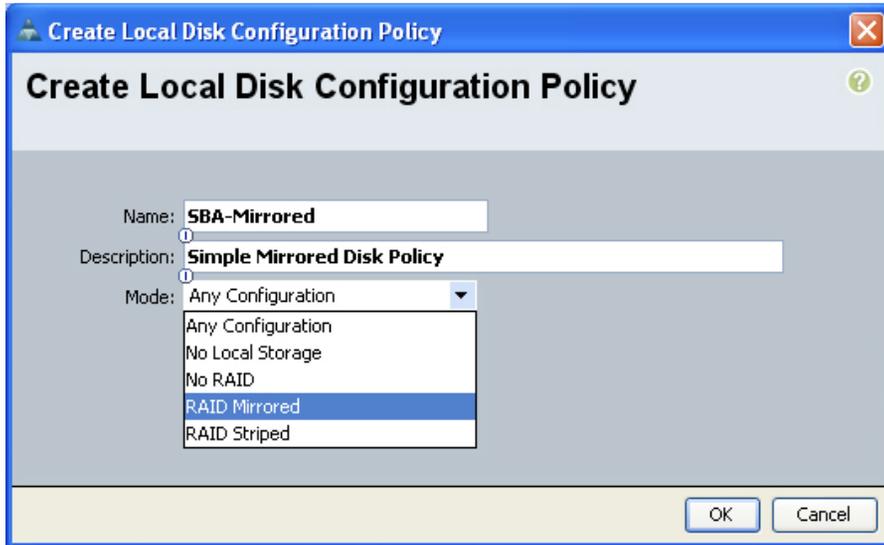
Procedure 3 Configure Storage

The local disk configuration policy allows the service profile to define how the block storage is structured on the local disks installed in each UCS blade server. A common configuration is to have two locally installed disks in each blade, set up for mirrored storage. To speed configuration of profiles and ensure consistency, a local disk configuration policy may be created and saved as a reusable object.

Step 1: Click **Create Local Disk Configuration Policy**.

Step 2: Provide a name and description for the policy, and choose **RAID Mirrored** from the **Mode** drop-down list as shown in Figure 24.

Figure 24. Create Local Disk Configuration Policy



Step 3: Click **OK** to create the policy, acknowledge the creation, and then choose the name of the newly created disk policy from the **Local Storage** drop-down list on the storage screen.

Step 4: The next item on the storage screen is scrub policy. This controls the action of the system on the data stored on the disks of blades being disassociated from a profile. For purposes of this example, leave scrub policy at default. Also for this example, we will not show Virtual Host Bus Adapter (vHBA) setup, which is specific to a Fibre Channel SAN configuration.

Step 5: To proceed with creating a basic service profile for local boot, choose **No vHBAs** next to the **How would you like to configure SAN connectivity?** prompt.

Step 6: Choose **Next** to proceed to the screen for Networking configuration.

Tech Tip

See the *Creating Service Profiles for SAN Boot* process for detailed information on enabling a service profile to access Fibre Channel attached storage over a SAN.

Procedure 4

Complete Networking Configuration

The Networking configuration screen allows you to define virtual network interface cards (vNICs) that the system will present to the installed operating system in the same way that a standalone physical server presents hardware NICs installed in a PCI bus. The type of mezzanine card installed in the blade server affects how many vNICs may be defined in the profile and presented to the server operating system. Leave the Dynamic vNIC Connection Policy drop-down list at its default setting for this example.

Tech Tip

Dynamic vNICs only apply to configurations that use the Cisco UCS M81KR Virtual Interface Card. Such configurations are discussed in the *Service Profiles Using Multiple vNICs* section under *Advanced Configurations*.

Step 1: Click **Expert** next to the **How would you like to configure LAN connectivity?** prompt. The expert mode allows you to walk through the creation of a MAC address pool, instead of using the default MAC pool which will not contain any address blocks on a new system.

Step 2: Click **Add** at the bottom of the screen to define the first vNIC in the profile.

Step 3: First assign a name to the profile. For the example configuration, we will use **eth0** as the interface name, representing Ethernet 0 as shown in Figure 25.

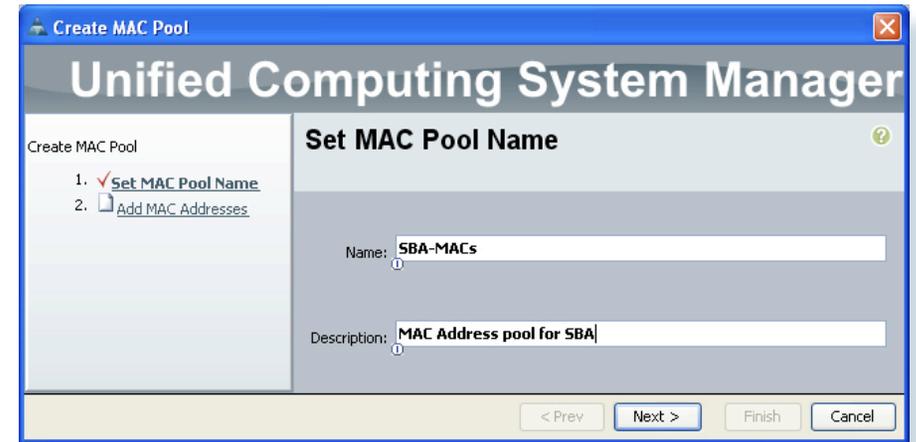
Figure 25. Create vNIC



Step 4: Next, click **Create MAC Pool** to add a pool of MAC addresses to be used by vNIC interfaces in service profiles. Using a pool of MAC addresses instead of hardware-based MAC addresses allows a service profile to retain the same MAC address for its network interfaces, even when it is assigned to a new blade server in the system.

Step 5: Set a name and description for the MAC pool, as shown in Figure 26.

Figure 26. Naming the MAC Address Pool



Step 6: Click **Next** to continue with adding MAC addresses to the pool.

Step 7: Click **Add** in the bottom of the window to add a block of MAC addresses.

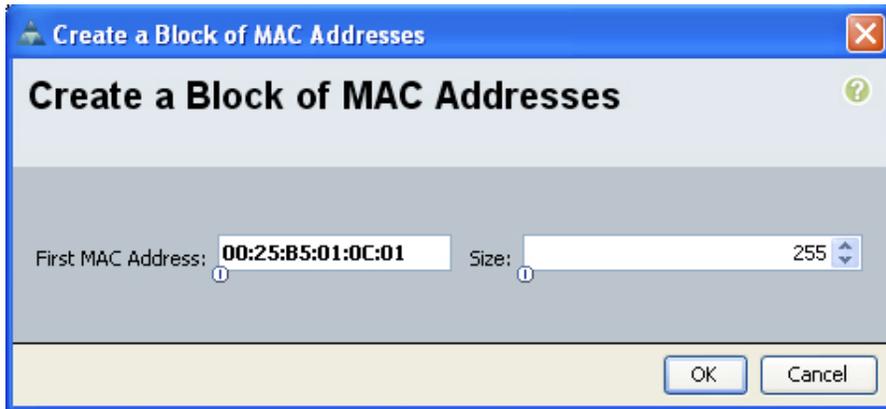
Step 8: The dialog box for creating a block of MAC addresses allows you to define the starting address and the number of addresses to include in the block. Create a block of addresses large enough to allocate one address to each vNIC that will exist in the system.

 **Tech Tip**

Consider using multiple MAC address blocks with specific numbering conventions relevant to your implementation to assist with troubleshooting.

Step 9: Enter a starting address for the MAC address block, and a quantity of addresses to allocate as shown in Figure 27.

Figure 27. Create A MAC Address Block



Step 10: Click **OK** to add the new block into the MAC address pool, and then click **Finish** and **OK** to acknowledge creation of the pool.

Step 11: Use the **MAC Address Assignment** drop-down list to select the name of the MAC address pool that you created.

The next section of the Create vNIC screen allows you to define the vNIC traffic path through the fabric interconnects and what VLANs are present on the vNIC. The Cisco UCS system has the capability to present multiple vNICs to the installed operating system and pass the traffic from a specific vNIC to either fabric interconnect A or B. In addition, a fabric-failover capability is available on specific NIC hardware to allow the system to continue forwarding traffic through the other fabric interconnect if the primary selection has failed. For this basic service profile example, select fabric A as the primary traffic path and enable failover.

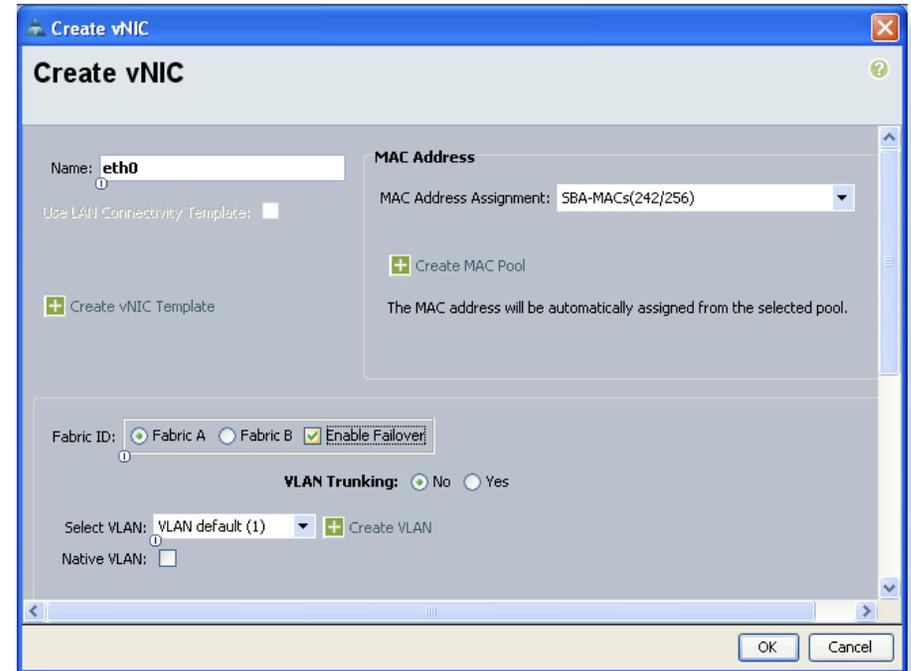
Tech Tip

Fabric failover is appropriate for configurations with a single host operating system installed directly on the blade server. For a virtualized environment, we recommend presenting multiple vNICs to the hypervisor and allowing the hypervisor system to manage the failover of traffic in the event of a loss of connection to one of the fabrics.

See the Service Profiles Using Multiple vNICs discussion under Advanced Configurations for more information on configurations with multiple vNICs.

Step 12: Next to **Fabric ID**, choose **Fabric A** and select **Enable Failover** as shown in Figure 28.

Figure 28. vNIC Fabric Selection

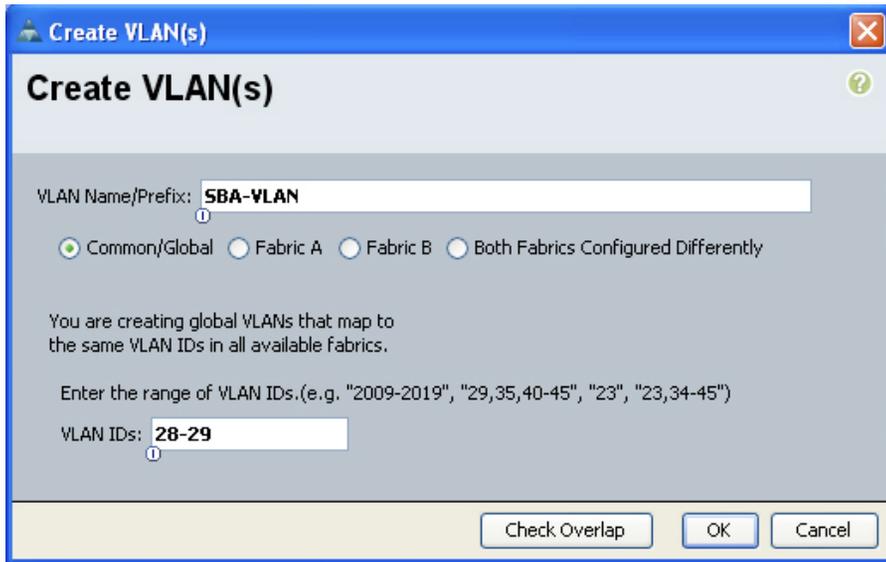


Step 14: The Cisco UCS system also allows vNICs to connect to the hosts as 802.1q VLAN trunks. In this basic example, we are placing this vNIC on a single VLAN in the Ethernet switching fabric, therefore, next to **VLAN Trunking**, leave **No** selected.

Step 15: To receive traffic from the server vNICs, you must define each VLAN needed in the Cisco UCS system. Click **Create VLAN** to identify the new VLAN number to the system.

Step 16: The Create VLAN(s) window allows you to create multiple VLANs. The number of the VLANs created is appended to the Name/Prefix entered. For example: The entries shown in Figure 29 would result in VLANs called SBA-VLAN28 through SBA-VLAN29. Enter the desired name and group of VLANs to create and click **OK**.

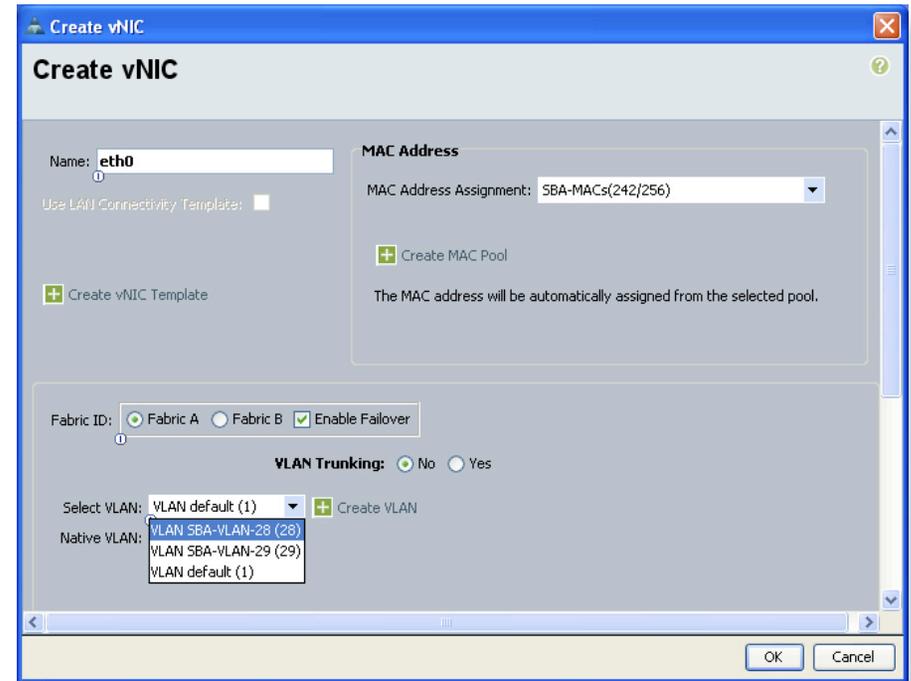
Figure 29. Create VLANs Window



Step 17: From the **Create vNIC** main screen, you can now choose the newly created VLAN from the **Select VLAN** drop-down list as shown in Figure 30.

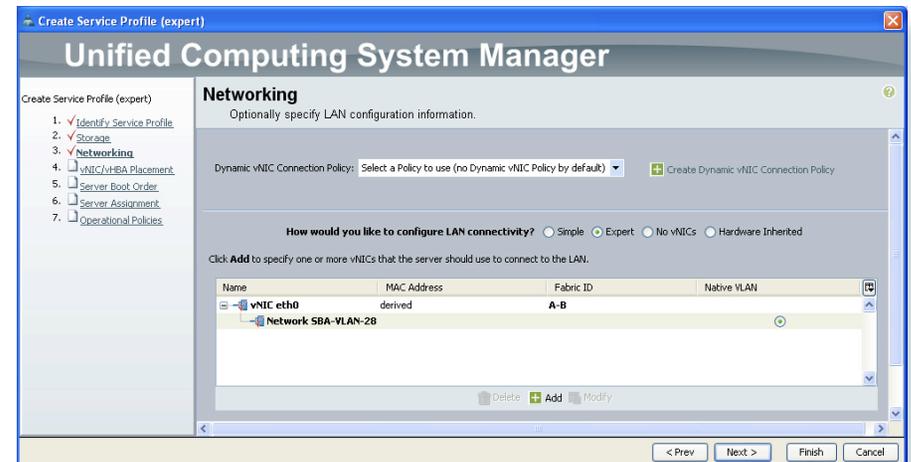
Step 18: When running a single VLAN from a host that will be transmitting traffic without 802.1Q VLAN tags, select **Native VLAN** to ensure that the untagged traffic can be properly forwarded by the fabric interconnects.

Figure 30. Select VLAN for the vNIC



Step 19: Leave the remainder of the fields on the **Create vNIC** screen at the default settings and click **OK**. The next screen shows the resulting created vNIC, its fabric association, and the VLANs on which it forwards traffic.

Figure 31. Example Created vNIC



Step 20: Verify the information displayed regarding the new vNIC as shown in Figure 31. Click **Next** to continue the service profile creation process.

Step 21: Click **Next** on the **vNIC/vHBA Placement** screen to let the system perform the placement of the virtual interfaces on the physical interfaces that exist on the blade servers to which this profile will be associated.

Procedure 5 Define the Server Boot Order Policy

The server boot order policy allows you to control the priority of different boot devices to which the server will have access. A basic configuration is to boot from removable media first, such as an attached CD/DVD drive, and then from the internal disk. More advanced configurations allow boot from LAN or boot from SAN. Having a preconfigured policy as a reusable object promotes consistency between similar service profile configurations.

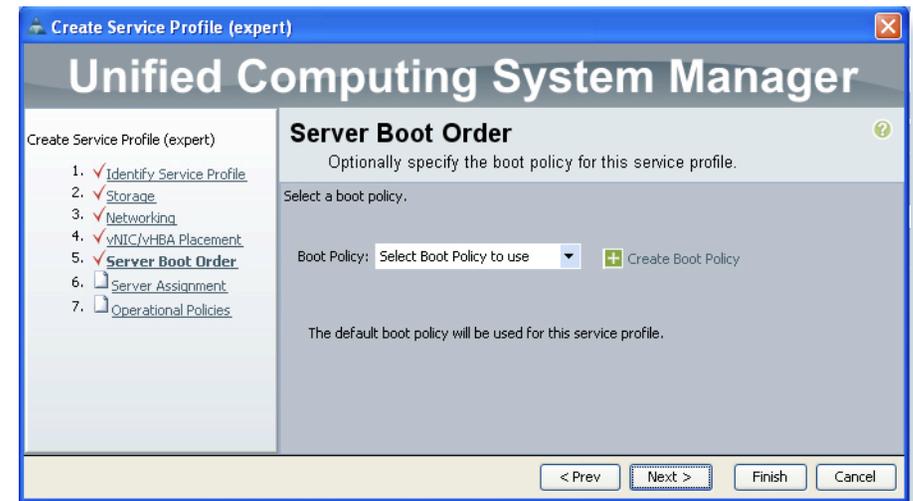
Step 1: On the **Server Boot Order** screen, click **Create Boot Policy**. This launches the **Create Boot Policy** screen, which allows you to assign various boot sources in order to a named policy. The lower left side of this screen has three containers for boot sources: **Local Devices**, **vNICs** and **vHBAs**.

Step 2: Click the down arrows on the **Local Devices** container to display the choices.

Step 3: Click **Add CD-ROM** first, to add a removable media source to the list.

Step 4: Click **Add Local Disk** to add the locally installed drives on the blade server itself as the next boot source.

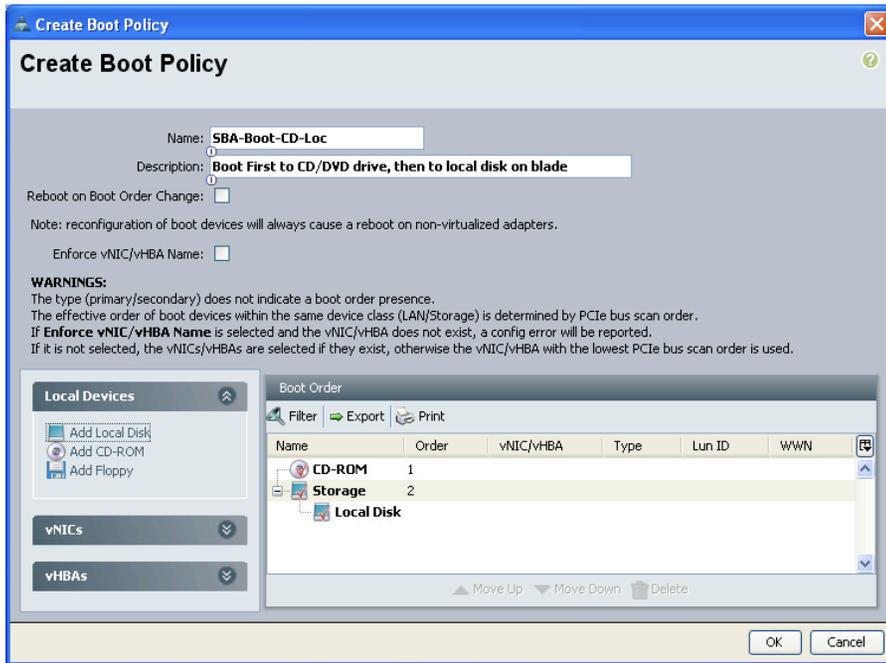
Figure 32. Create Boot Policy Screen



Step 5: The order of the devices in the list is displayed as a number in the **Order** column of the table. Assign a name and description to the policy in the spaces provided, verify the choices, and click **OK** to create the named boot policy, as shown in Figure 32.

Step 6: In the **Server Boot Order** screen, use the **Boot Policy** drop-down list to select the name of the policy just created to be applied to this profile, as shown in Figure 33.

Figure 33. Server Boot Policy Selection



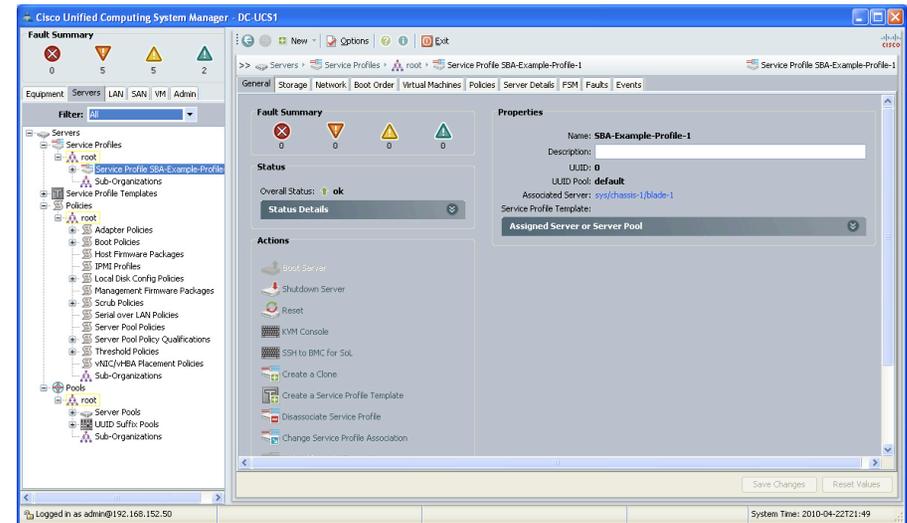
Step 7: Click **Next** to proceed with the service profile creation.

Procedure 6 Assign the Server

Cisco UCS has the ability to assign a service profile directly to a specific server, pre-provision an unused chassis slot, assign the profile to a pool of servers, or assign the profile to a physical blade server later.

Step 1: To simplify creation of this basic initial service profile, choose **Assign Later** from the **Server Assignment** drop-down list, as shown in Figure 34.

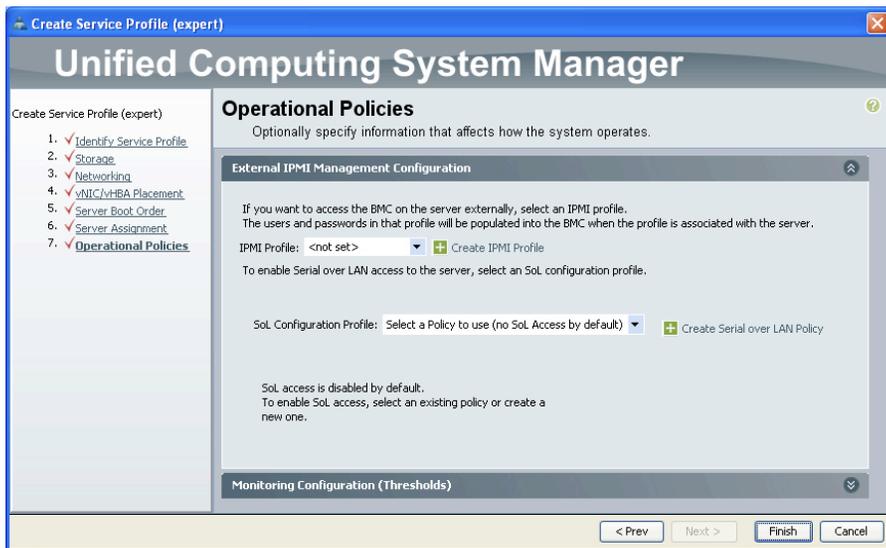
Figure 34. Server Assignment



Step 2: Leave the power state set up to ensure that a physical blade server will be powered on when the service profile is eventually applied. Click **Next** to proceed.

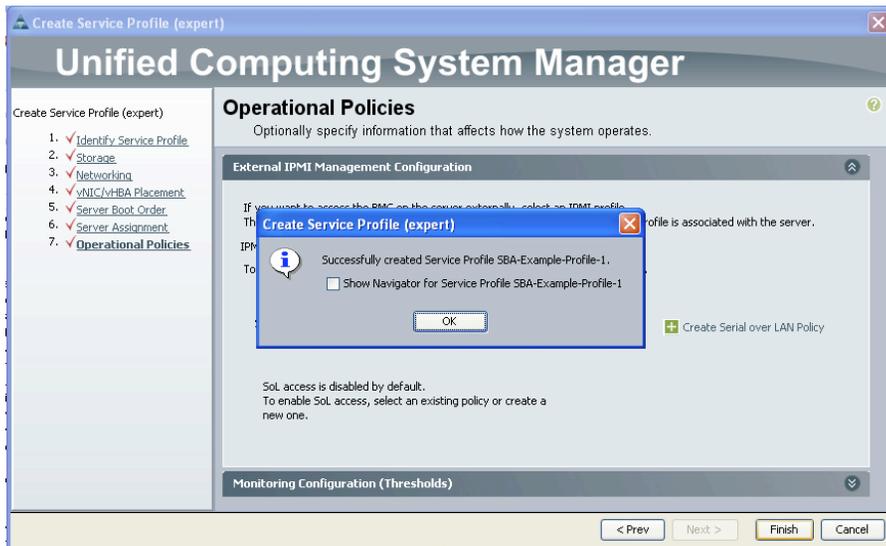
Step 3: The final screen in the service profile creation expert wizard allows configuration of access by Intelligent Platform Management Interface (IPMI) clients, and Serial over LAN (SoL) access as shown in Figure 35. Detail on these tools is out of the scope of this guide. For more information, please refer to Cisco UCS product guides at http://www.cisco.com/en/US/partner/products/ps10281/tsd_products_support_series_home.html.

Figure 35. Operational Policies Screen 7



Step 4: Click **Finish** to complete creation of the initial service profile on the system. The system displays a message indicating successful completion of the **Create Service Profile (expert)** wizard as shown in Figure 36.

Figure 36. Create Service Profile Success



Process

Applying Service Profiles to Physical Servers

1. View Your Profile
2. Associate Profile with Server
3. Complete Basic OS Installation

You've now completed the basic initial service profile creation process. This process shows you how to apply this profile to a physical blade server, and install a base operating system on the locally installed disks.

Procedure 1 View Your Profile

After you complete the service profile creation wizard, view the resulting profile in the Cisco UCS Manager GUI.

Step 1: In the navigation pane, choose the **Servers** tab, expand **Service Profiles**, and select the working service profile.

The work pane shows multiple tabs that roughly correspond to the sections of service profile configuration that you walked through in the **Create Service Profile (expert)** wizard. On the **General** tab in the work pane is an **Actions** area with a list of links for performing various tasks associated with the profile, as shown in Figure 37.

Figure 37. View Service Profile in UCS Manager

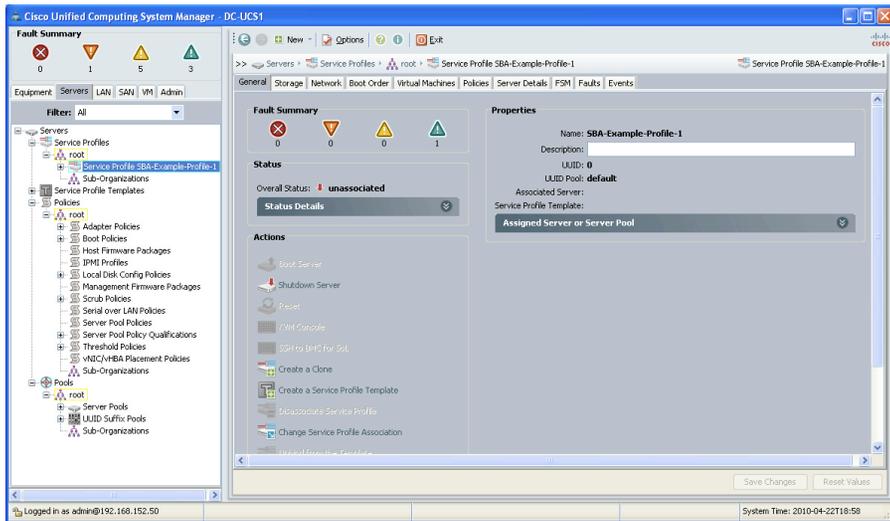
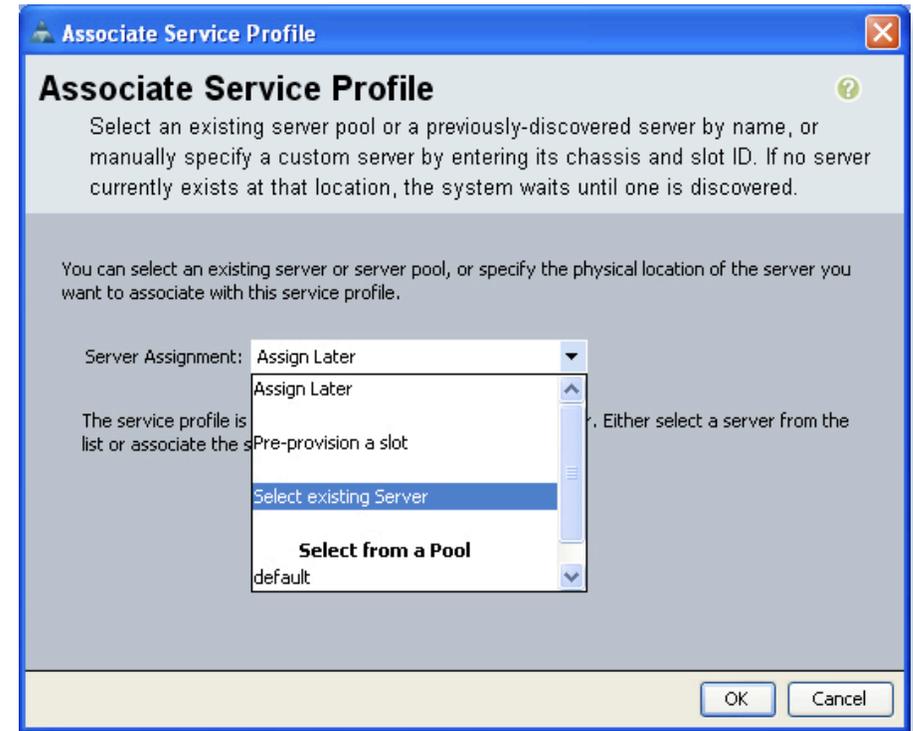


Figure 38. Associate Existing Server Selection



Step 3: In the **Select** column, choose a server, as shown in Figure 39.

Procedure 2 Associate Profile with Server

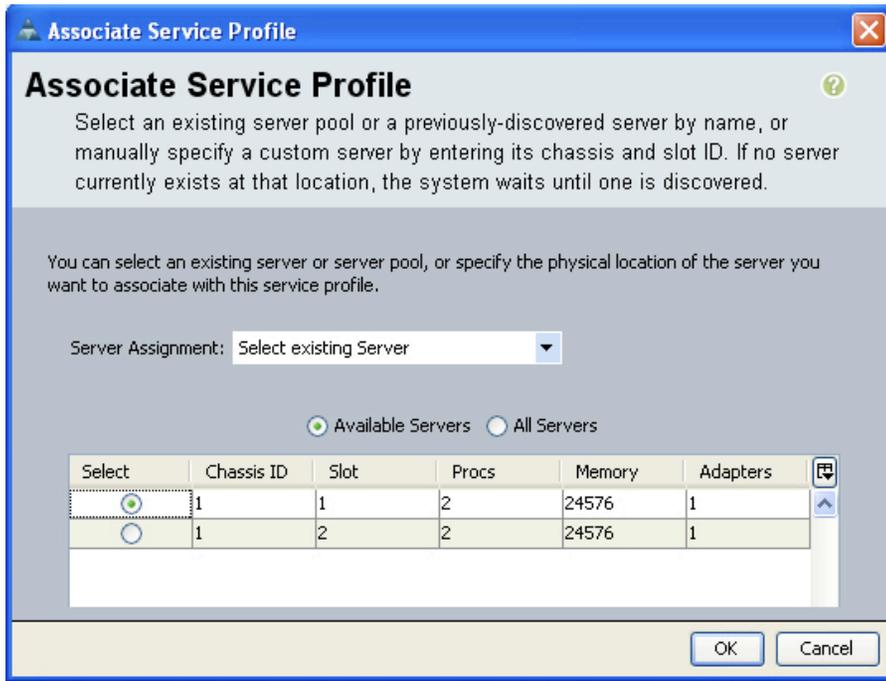
After you finish viewing the service profile and verifying the settings, you are ready to associate the profile with a physical blade server:

Step 1: Click **Change Service Profile Association** to launch the Associate Service Profile screen.

Step 2: From the **Server Assignment** drop-down list, choose **Select existing Server** as shown in Figure 38.

By default, the **Associate Service Profile** screen shows available servers only, that is, servers that are not already associated to another service profile. The available servers are displayed in a table sorted by chassis ID, slot number, and characteristics.

Figure 39. Available Server Selection



Step 4: Click OK to initiate the process of associating the service profile to the selected server.

Step 5: To track the progress of the association, choose the service profile by name under the **Servers** tab in the navigation pane and view either **Overall Status** on the **General** tab or the progress of the specific operation on the **FSM** tab.

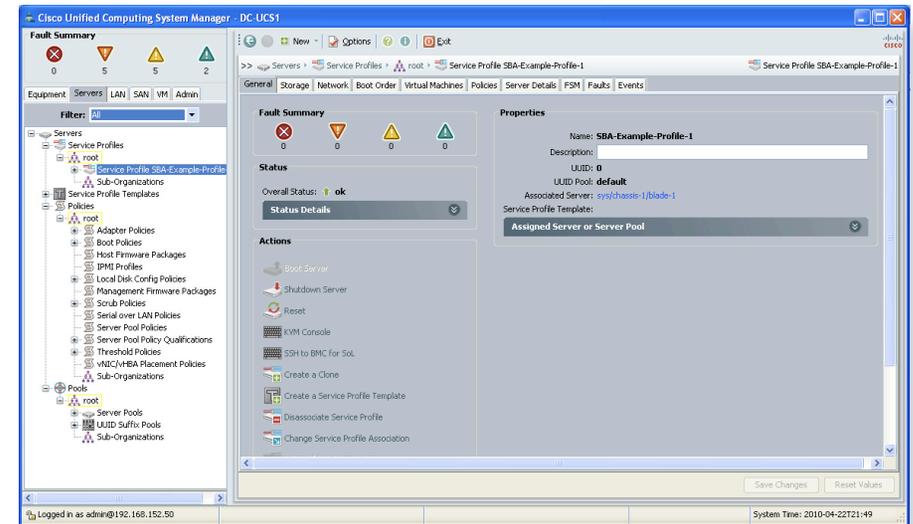
Procedure 3 Complete Basic OS Installation

To complete a simple operating system installation with traditional non-scripted approach, CD or DVD media must be available during the server boot process, and defined first in the server boot order as shown in our example service profile creation.

Step 1: To view the state of the server as it boots, select the service profile name in the **Servers** tab of the navigation pane, and view the general tab in the work pane as shown in Figure 41.

Step 2: Access the KVM Console for the server by clicking **KVM Console** in the **Actions** area of the work pane.

Figure 40. Associated Service Profile General Tab

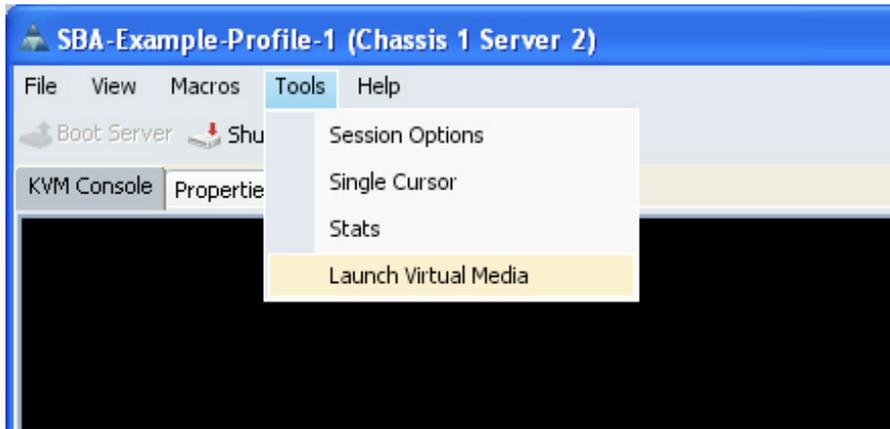


Step 3: Removable CD/DVD media may be presented to the system in two primary ways. The first approach is to use the USB connection provided by the console port on the front of each blade server to connect an external drive. This approach will provide the fastest file access for initial operating system install.

Tech Tip

An alternate approach is to use the virtual media capability of the KVM Console utility. In the KVM Console window, choose **Tools > Launch Virtual Media** as shown in Figure 41. The **Virtual Media Session** window opens.

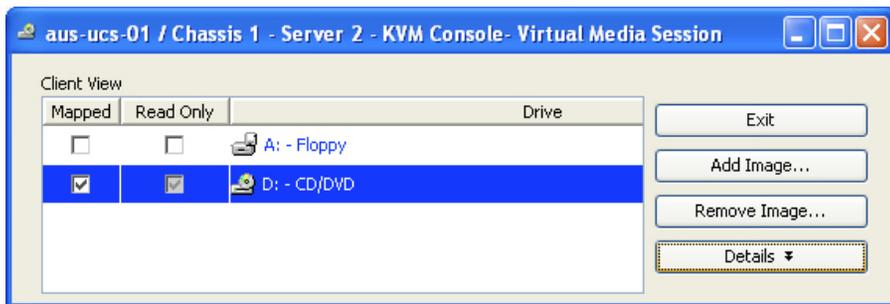
Figure 41. KVM Launch Virtual Media



The Virtual Media feature allows the server being viewed within the KVM Console to access the removable media drives of the computer running the Cisco UCS Manager client. The mapped drive will be seen as a locally attached device for purposes of boot policy, for easy installation of operating systems without needing to manually load media locally to the blade server. Map ISO disk images present on the computer running the Cisco UCS Manager client as virtual media by clicking **Add Image**.

Step 4: To map the local drive to the active blade server being viewed, next to the drive of your choice, select **Mapped**, as shown in Figure 42. After it is mapped, the blade server can boot to the virtual media session just as it would to a USB-attached CD/DVD player attached to the physical console port.

Figure 42. Virtual Media Session Mapping



Tech Tip

When using virtual media, the media does need to travel across the network from the UCSM client machine through the fabric interconnects to the server. Install times may be slightly longer with this approach, depending on the speed and latency of the connection between the computer running the client and the blade server.

Step 5: After the blade server associated to the service profile has booted from the provided install media, the installation process proceeds in the same way as a typical standalone rack-mount server. You can use the KVM Console interface for any ongoing interaction required to complete the installation.

Tech Tip

An alternative approach for operating system rollout is to use a PXE-boot server for LAN boot of the installation media. This approach is covered in the next process, Creating Service Profiles for LAN Boot.

Process

Creating Service Profiles for LAN Boot

1. Create and Activate PXE Boot

Booting over Network (LAN or SAN) is an important step in moving towards stateless computing in which there is no static binding between a physical server and the operating system. The operating system is decoupled from the physical server.

Pre-Execution Environment (PXE) boot allows for LAN booting of operating systems without local disk or SAN required. PXE is based on industry standard protocols TCP/IP, DHCP, and TFTP. PXE allows booting from a network by copying a boot image file from a server. A NIC that supports PXE is required. All network adapters in Cisco UCS support PXE boot. In addition to configuring the Cisco UCS boot profile for PXE boot, you will need to configure the PXE server with the MAC address of the Cisco UCS server that you are applying the PXE boot profile to.

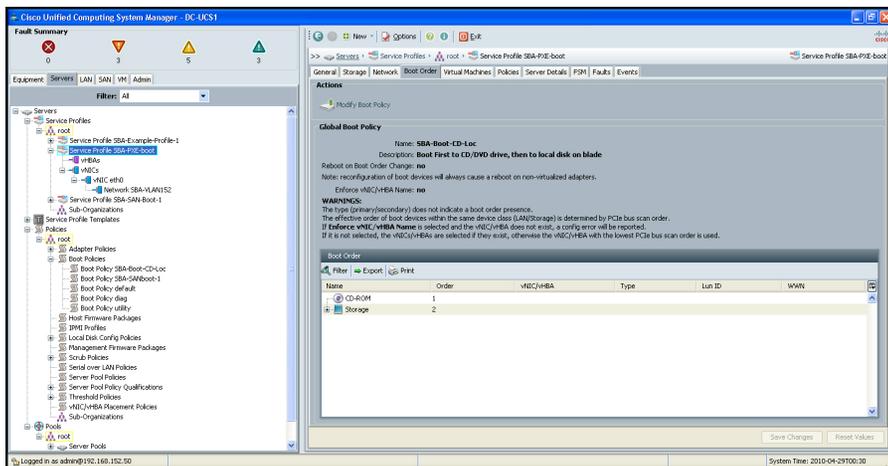
Procedure 1 Create and Activate PXE Boot

For a Cisco UCS server to support PXE boot, a boot policy specifically for PXE boot is needed. With an existing service policy in place, complete this procedure to create and activate PXE boot.

Step 1: On the **Server** tab in the navigation pane, select the service profile that you want to configure to support PXE boot, as shown in Figure 43.

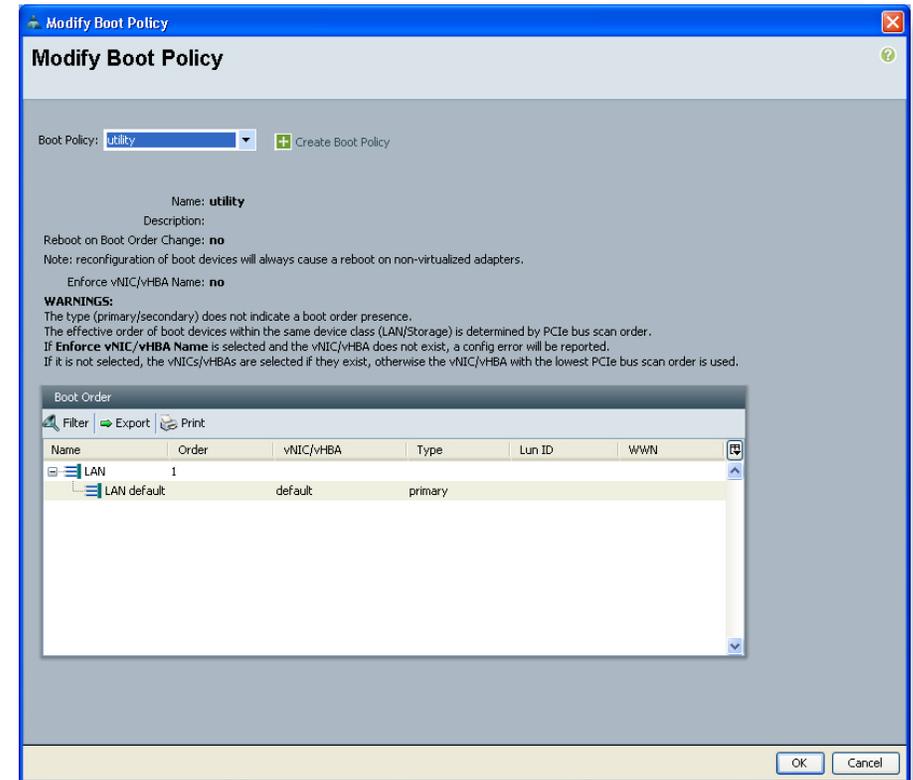
Step 2: In the working pane, click the **Boot Order** tab and select **Modify Boot Policy**. The **Boot Policy** window will display, as shown in Figure 44.

Figure 43. Modify Boot Order



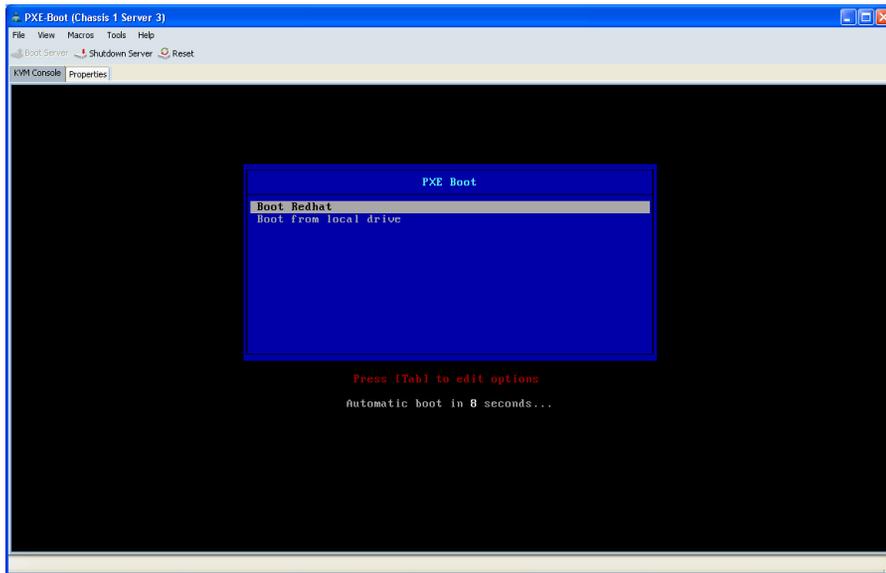
Step 3: To select a boot policy that is configured for PXE boot. A predefined boot policy called utility is already configured to allow PXE boot. Select **Utility** under the **Boot Policy** pull-down list.

Figure 44. Modify Boot Policy



Step 4: After the service profile with the utility boot policy is associated with a Cisco UCS server, when you boot the server, you can launch the KVM console see the system DHCP request and access the PXE boot menu on the server. An example of a KVM Console screen to a successfully PXE-booted server is shown in Figure 45. Your screen will be different, based on the capabilities configured in your PXE server.

Figure 45. PXE Boot Menu from PXE Server



Process

Creating A Service Profile for SAN Boot

1. Create a WWNN Pool
2. Add a vHBA to the Service Profile
3. Provide Access to the Disk Array Through SAN Fabric
4. Assign the WWNN Pool
5. Modify the Boot Policy

Booting service profiles directly from a Fibre Channel SAN can provide key advantages for ensuring server and application availability. With all operating system files and application data specific to the server stored on the SAN, your agency benefits from SAN disk redundancy and backup practices. This approach works in conjunction with the hardware independence provided by Cisco UCS-specific constructs such as shared pools of Ethernet and Fibre Channel addressing. Together, these attributes provide

the ability to move a service profile between server blades within the system programmatically, with no physical intervention required. This concept is known as stateless computing.

This process illustrates how to create a new service profile for SAN boot based on the existing example profile already created in the Creating an Initial Service Profile for Local Boot process.

Procedure 1 Create a WWNN Pool

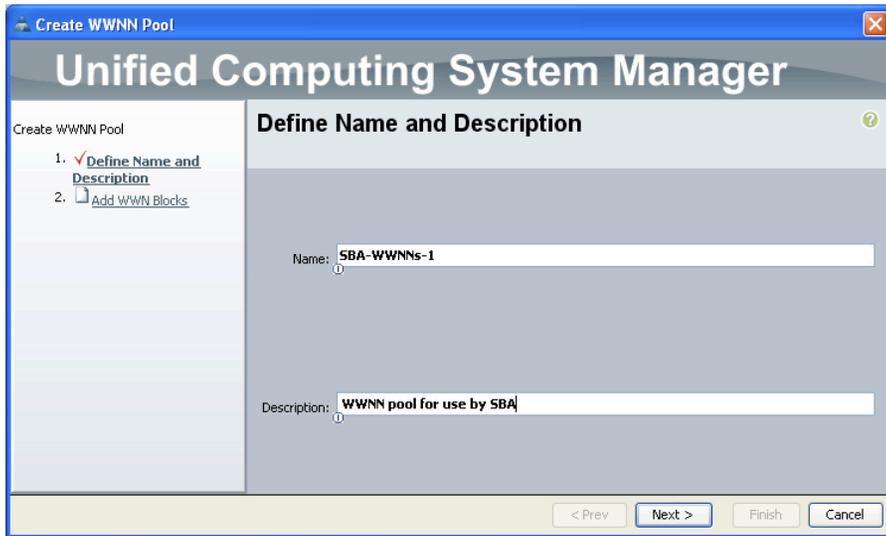
First, you must provision a pool of World Wide Node Names (WWNNs) in the system. The WWNNs in the pool will be assigned to service profiles that need to access the Fibre Channel SAN. One WWNN is assigned from the pool to each service profile. Each WWNN corresponds to the identity of a Fibre Channel end-node.

Step 1: Choose the **SAN** tab in the navigation pane, and expand **SAN > Pools > Root**.

Step 2: Select **WWNN Pools** in the navigation pane, and then in the work pane, click **Add**.

Step 3: Provide a name and description for the new pool as shown in Figure 46, and click **Next**.

Figure 46. Add a WWNN Pool

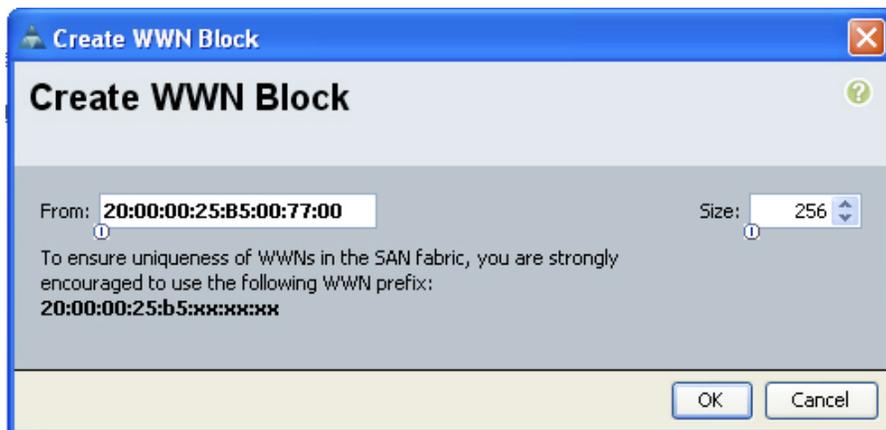


Step 4: The next step is to create a block of WWNN addresses by defining a starting point and the quantity of addresses in the block. In the **From** field, the system provides a prefix to help ensure uniqueness of the WWNN values on the SAN. Assign the last three segments of the base WWNN value in colon-delimited notation, as needed for your system.

Step 5: In the **Size** field, specify the number of node names required, as shown in Figure 47.

Step 6: Click OK.

Figure 47. Create WWNN Block



Step 7: Click **Finish** to complete creation of the new WWNN pool. This pool will be referenced by the service profile created for SAN boot in the following section.

For more information on WWN, WWNN, and WWP, see the Cisco UCS Manager GUI Configuration Guide at: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/GUI_Config_Guide_chapter20.html.

Procedure 2 Add a vHBA to the Service Profile

Cloning the example service profile previously created allows you to reuse the storage and boot order configuration of an existing profile to create a new profile with very little additional work.

Step 1: On the **Servers** tab in the navigation pane, right-click the profile name and choose **Create a Clone**.

Step 2: Assign a name that clearly identifies the profile as a SAN boot server instance.

Step 3: After you have created the new profile, click the name of the new profile in the navigation pane and select the **Storage** tab in the work pane. This tab is initially blank because we did not add any storage adapters in the example profile.

Step 4: Click **Add** at the bottom of the screen to create a vHBA for Fibre Channel storage access in this service profile. This will launch the Create vHBA screen, as shown in Figure 48. Assign a name to the vHBA; the example configuration uses the name **fc0**.

Figure 48. Add vHBA for Storage Access

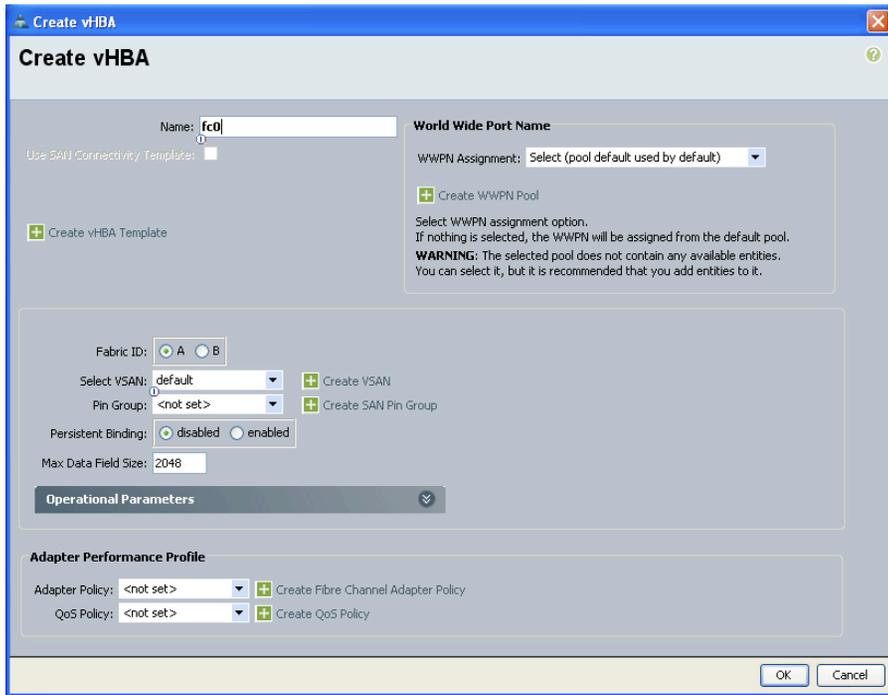
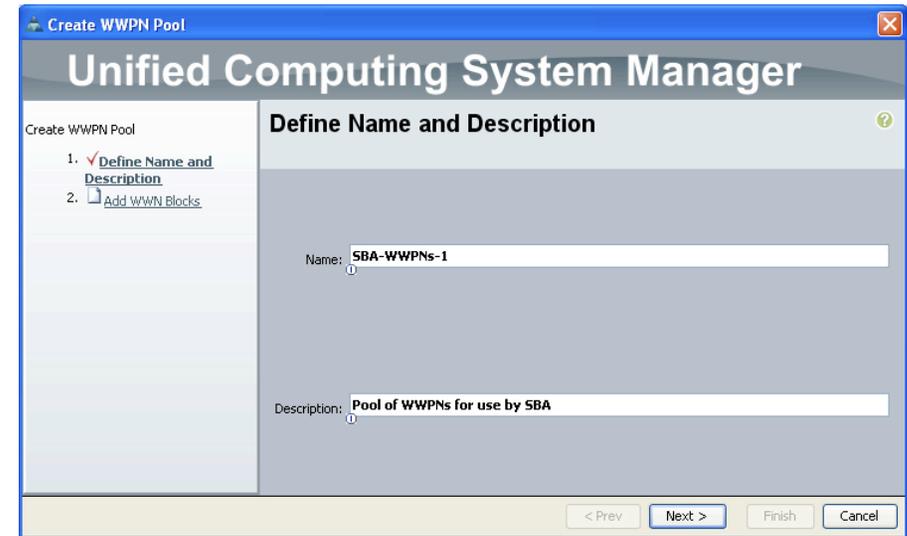


Figure 49. Create WWPN Pool



Step 3: The next screen allows assignment of WWPN blocks to the system. Click **Add** to create a new WWPN block.

Step 4: In the **From** field, the system provides a prefix value to help ensure uniqueness of the WWPN values on the SAN. Assign the last three segments of the base WWPN value in colon-delimited notation, as needed for your system.

Step 5: In the **Size** field, set the number of port names required in the pool, as shown in Figure 50.

Step 6: Click **OK**.

Procedure 3 Provide Access to the Disk Array

Both Fibre Channel port and node addressing assignments are required in order for Cisco UCS to provide access through the SAN fabric to the disk array. Using WWPNs and WWNNs that are independent from the physical hardware allows you to assign the service profile to any server in the Cisco UCS system and assume the correct server identity and SAN access privileges. Similar to the WWNN pool that you created, you must provision a pool of WWPNs for the system to assign port addresses consistently when you add new service profiles to the system.

Step 1: On the **Create vHBA** screen, click **Create WWPN Pool** to provision a new WWPN pool on the system.

Step 2: Assign a name and description to the pool as shown in Figure 49 and click **Next**.

Figure 50. Creating a WWPN Block



Step 7: Click Finish to complete creation of the new WWPN pool.

Step 8: In the Create vHBA window, choose the new WWPN pool from the WWPN Assignment drop-down list.

Step 9: Because this is the first vHBA that you have added to this profile, leave the Fabric ID setting as Fabric A. If you add a second vHBA to the profile for SAN fault-tolerance, you should assign the second vHBA to Fabric B.

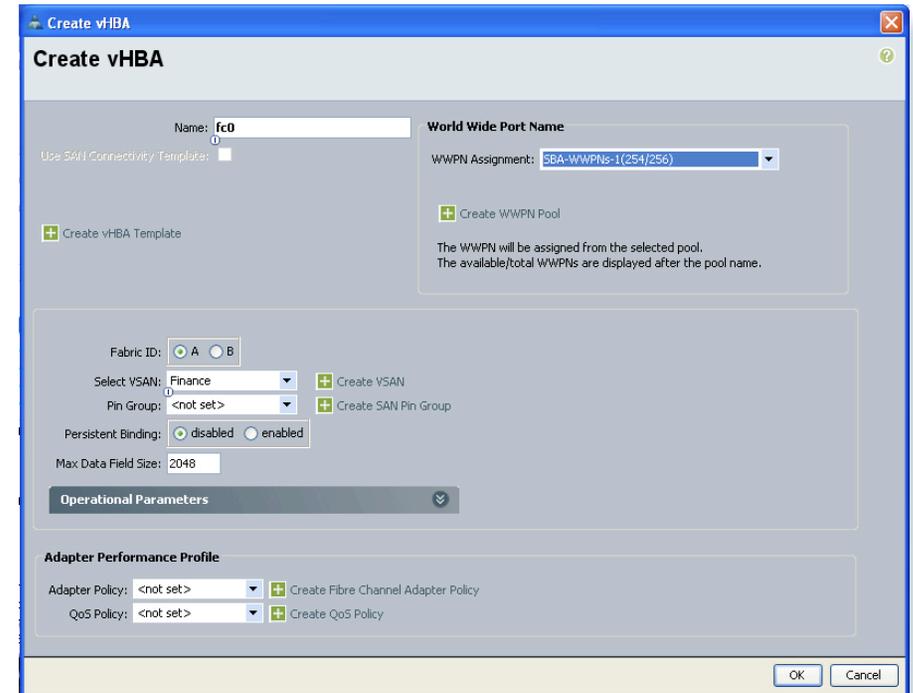
Next, you must configure this vHBA to communicate on a specific Virtual SAN (vSAN). vSANs allow multiple logical SAN networks to share a common physical Fibre Channel infrastructure. The desired vSAN numbering must match what you have created on your Cisco MDS 9100 Series storage-switching network, and what was previously created in the Define Fibre Channel Uplink Ports procedure within the Getting Started with UCS Manager process.

Step 10: In the Create vHBA window, use the Select vSAN drop-down list to assign the vSAN to the vHBA.

Step 11: Leave the remaining fields in the Create vHBA window at their default settings.

Step 12: Click OK to complete creation of the vHBA, as shown in Figure 51.

Figure 51. Complete vHBA Creation



Step 13: After you have completed the vHBA configuration, click **Save Changes** at the bottom of the **Storage** tab in the work pane to ensure that the changes are applied to the service profile.

When you save and apply changes to the service profile, the system assigns a specific WWPN from the pool to the new vHBA. This WWPN is the Fibre Channel initiator value that must be assigned to the correct SAN zone to access the desired boot LUN assigned by the SAN administrator. The storage system must also be properly configured in its LUN masking to expose the desired LUNs to this specific WWPN.

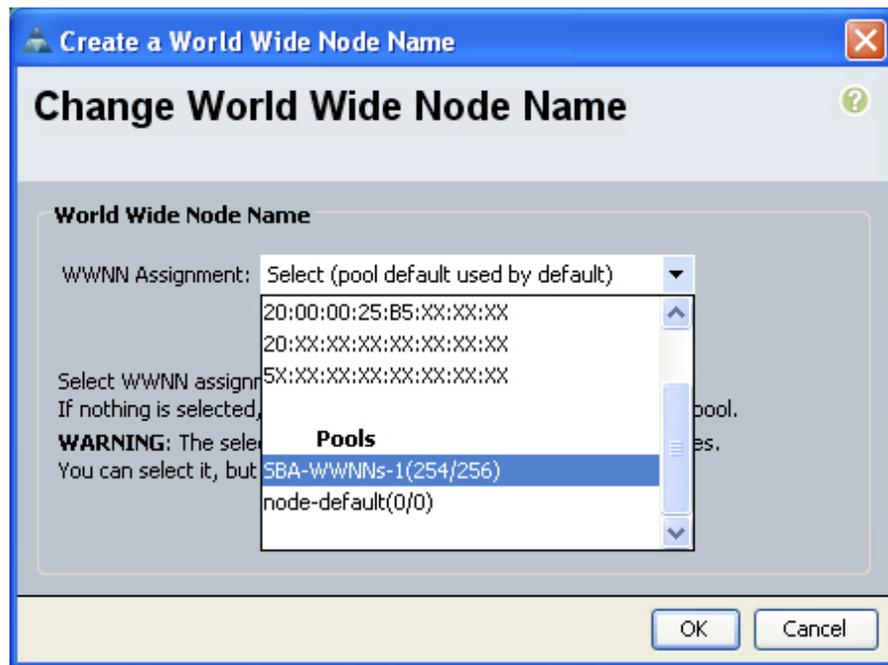
Procedure 4 Assign the WWNN Pool

Now you have defined a vHBA in the service profile to provide access to the Fibre Channel SAN, you must assign the WWNN pool to the profile.

Step 1: In the **Storage** tab of the work pane, in the **Actions** area, click **Change World Wide Node Name**.

Step 2: From the **WWNN Assignment** drop-down list, choose the WWNN pool name that you created, as shown in Figure 52.

Figure 52. Assign WWNN Pool to Service Profile



Step 3: Click **OK** to complete assignment of the pool to the profile. On the **Storage** tab, the World Wide Node Name field now reflects the WWNN assigned to the profile from the pool and the name of the WWNN pool that you specified.

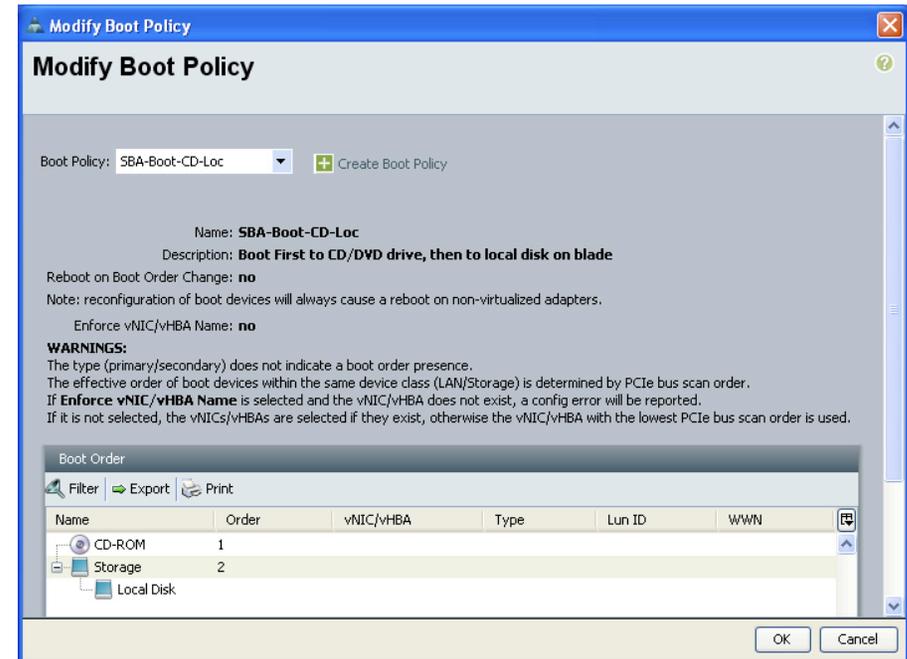
Procedure 5 Modify the Boot Policy

Now update the profile to use the new vHBA storage adapter and the addresses assigned to it to access a new boot LUN over the SAN. The base profile that we created uses a simple boot policy with a boot order that starts with any attached removable media, such as CD or DVD, and then default to the local (internal) disk drives on the server.

Step 1: Choose the **Boot Order** tab in the work pane, and then in the **Actions** area, click **Modify Boot Policy** to change the boot policy to use the SAN connections. You will create a new boot policy that is specific to accessing the target WWPN of the storage system that houses the server boot LUN.

Step 2: Click **Create Boot Policy** on the **Modify Boot Policy** screen as shown in Figure 53.

Figure 53. Modify Boot Policy Screen



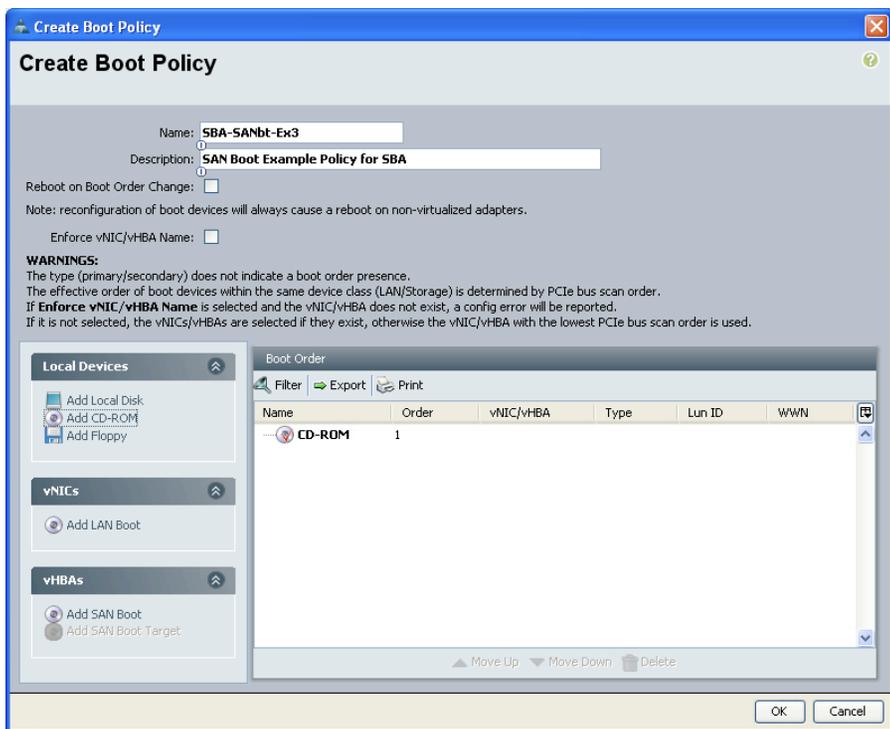
Tech Tip

When you create a boot policy that targets the WWPN of the storage system, the boot policy may be reused across multiple service profiles. Many storage systems can present a different LUN as a boot LUN or LUN 0 to different initiators, based on the initiator WWPN address. Referencing a common boot policy promotes configuration consistency across similar service profiles.

Step 3: On the **Create Boot Policy** screen, assign a name and description to the new boot policy as shown in Figure 54. The choices along the left of the window allow you to add different devices into the boot sequence.

Step 4: Click **Local Devices**, and then click **Add CD ROM** to allow the profile to first boot to a physically attached removable drive or a KVM Console Virtual Media drive, if present.

Figure 54. Create Boot Policy Screen



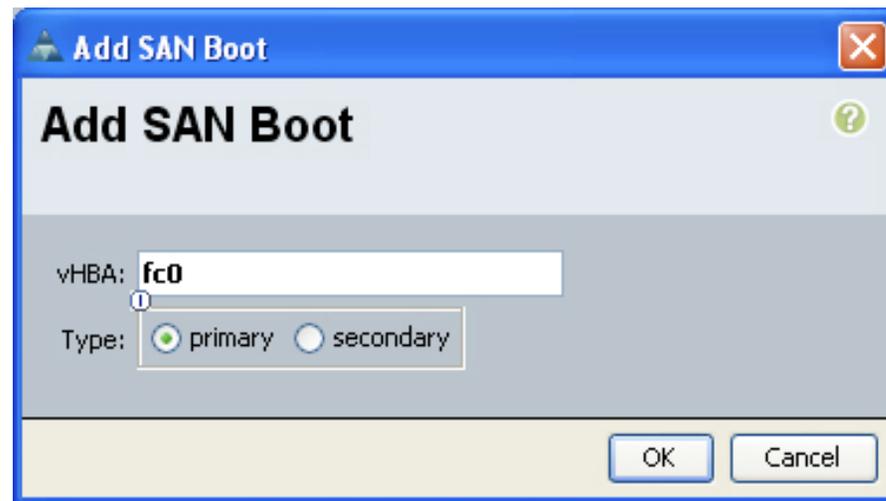
Step 5: To define the vHBA to be used for booting the Service Profile, click on vHBAs and choose **Add SAN Boot**.

Step 6: You must enter in the same name used for defining the vHBA to the system in the Add a vHBA to the Service Profile procedure in the Creating New Service Profile for SAN Boot process. The example name provided was **fc0**. Enter the correct name into the vHBA field of the Add SAN Boot window as shown in Figure 55. Click **OK**.

Tech Tip

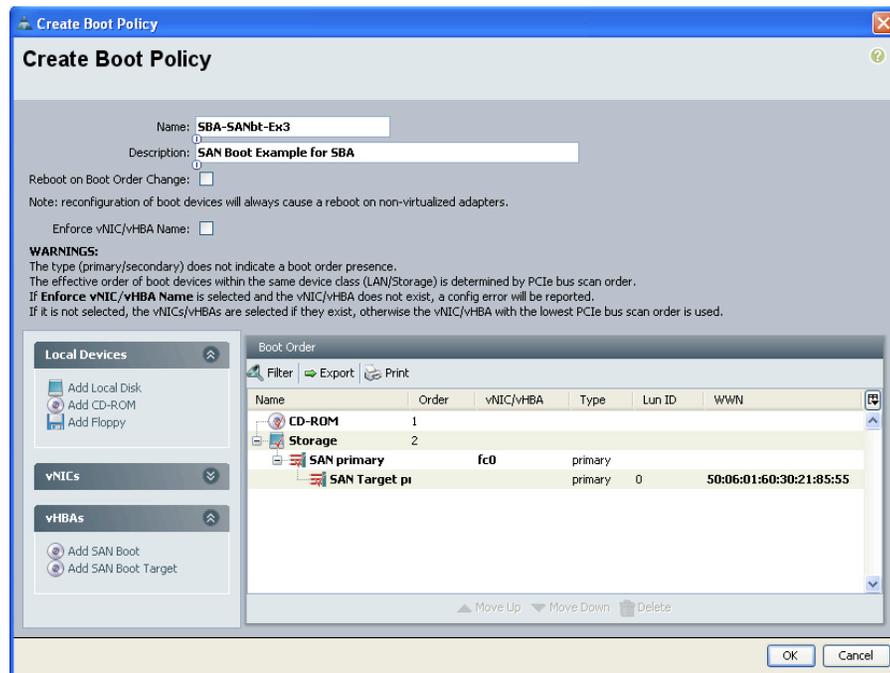
Since the boot policy references the vHBA by name, you must name interfaces consistently across service profiles that need to share a common boot policy definition.

Figure 55. Define vHBA for SAN Boot



Step 7: On the **Create Boot Policy** screen, choose **Add SAN Boot Target** to define the specific LUN number for the system to boot and the WWPN of the target storage system. Typically, the boot LUN is presented by the storage system as LUN 0 to the requesting initiator. With the proper SAN target WWPN provided, the boot policy should appear similar to Figure 56.

Figure 56. Example Boot Policy for SAN Boot



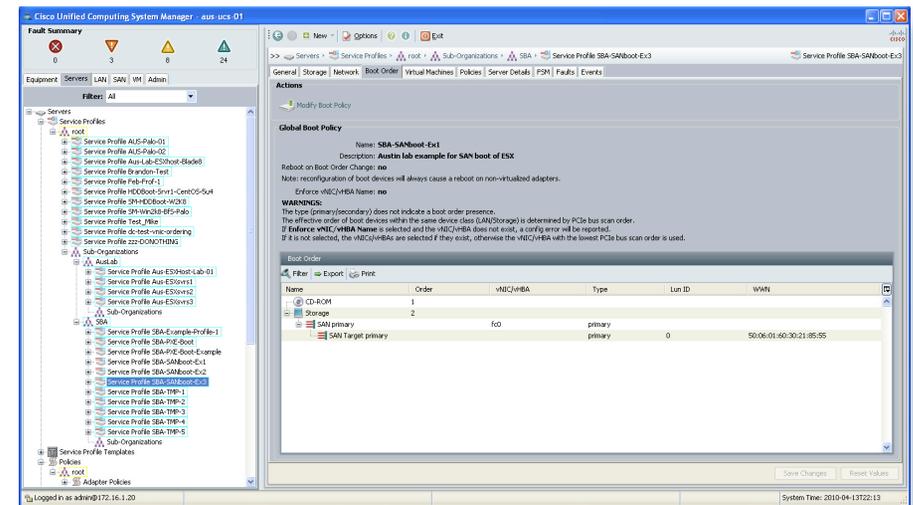
Step 8: Click OK to complete creation of the new boot policy.

Step 9: In the **Modify Boot Policy** window, use the **Boot Policy** drop-down list to select the new boot policy and assign it to the profile.

Step 10: After you select the new boot policy, the work pane shows the new boot order information, including the correct target LUN ID and WWPN number, as shown in Figure 57.

After you complete this step, you can apply the service profile to a server in Cisco UCS.

Figure 57. Updated Boot Policy Applied to Profile



After you apply the service profile applied to a server, you can boot the server, which needs to have the operating system installation media available by locally attached removable media or KVM Console Virtual Media. The installation begins as is typical for the given operating system. When you choose a target disk destination for the installation, ensure that the new LUN 0, accessible over the Fibre Channel SAN, is selected.

You can provision the SAN to expose multiple LUNs to a given initiator. For example: You can use separate LUNs to house operating system boot files and files that contain application specific data or database contents. In a hypervisor environment, a LUN specific to an individual profile is presented as a boot LUN. A larger LUN accessible to multiple initiators is used to house VM-specific files. In this way, multiple virtualized servers can access the VM files.

Tech Tip

Redundant access to the boot LUN may be configured at install on some operating systems, on others it must be added after the initial installation is complete. Please see Service Profiles using Multiple vHBAs in the Advanced Configurations Section for more information.

Cisco UCS Rack Mount Servers

Physical Setup and Connectivity

Cisco UCS rack servers ship with onboard 10/100/1000 Ethernet adapters and a Cisco Integrated Management Controller (CIMC). To get the most out of the rack servers and minimize cabling in the SBA Unified Computing architecture, the Cisco UCS C210 rack-mount server is connected to a unified fabric. The Cisco Nexus 5010 Series switch that connects the Cisco UCS 5100 Series Blade Server Chassis to the network can also be used to extend Fibre Channel traffic over 10-Gigabit Ethernet. The Cisco Nexus 5010 Series switch consolidates I/O onto one set of cables, eliminating redundant adapters, cables, and ports. A single card and set of cables connects servers to the Ethernet and Fibre Channel networks and also allows the use of a single cabling infrastructure within server racks.

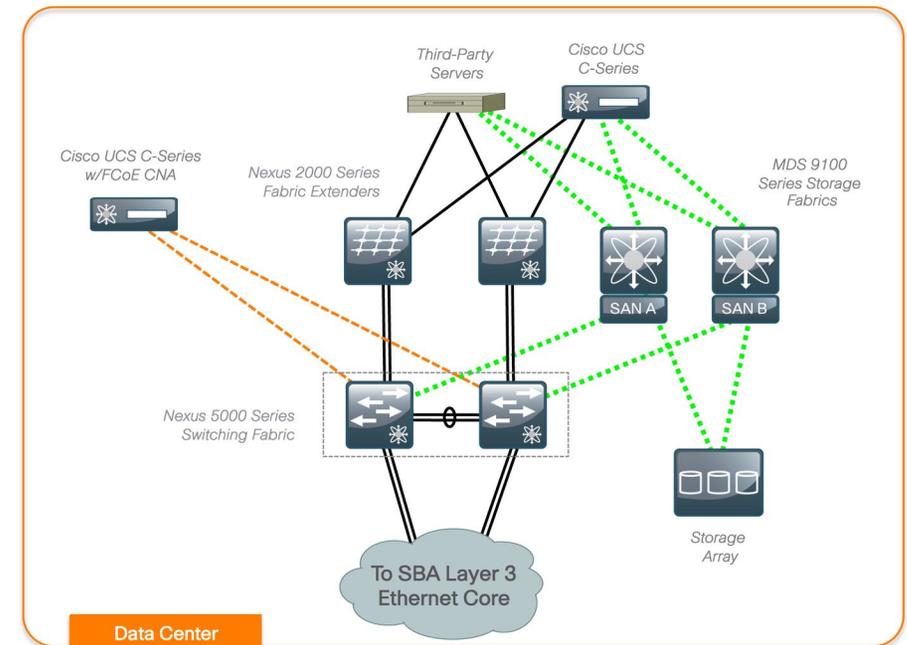
In the SBA Unified Computing architecture, the Cisco UCS C210 rack-mount server is configured with a dual-port CNA. Cabling the Cisco UCS C210 with a CNA limits the cables to three, one for each port on the CNA and the CIMC connection. A standard server without a CNA could have a few Ethernet connections or multiple Ethernet and Fibre Channel connections. Figure 58 shows a topology with mixed unified fabric and standard Ethernet and Fibre Channel connections.

The Cisco UCS C210 is connected to both Cisco Nexus 5000 Series switches from the CNA with twinax cabling. The CIMC management port connects to an Ethernet port on the fabric extender.

The Cisco Nexus 5010 switch has one available expansion slot which can be used to add Fibre Channel or to add 10-Gigabit Ethernet ports. The available options for Fibre Channel are:

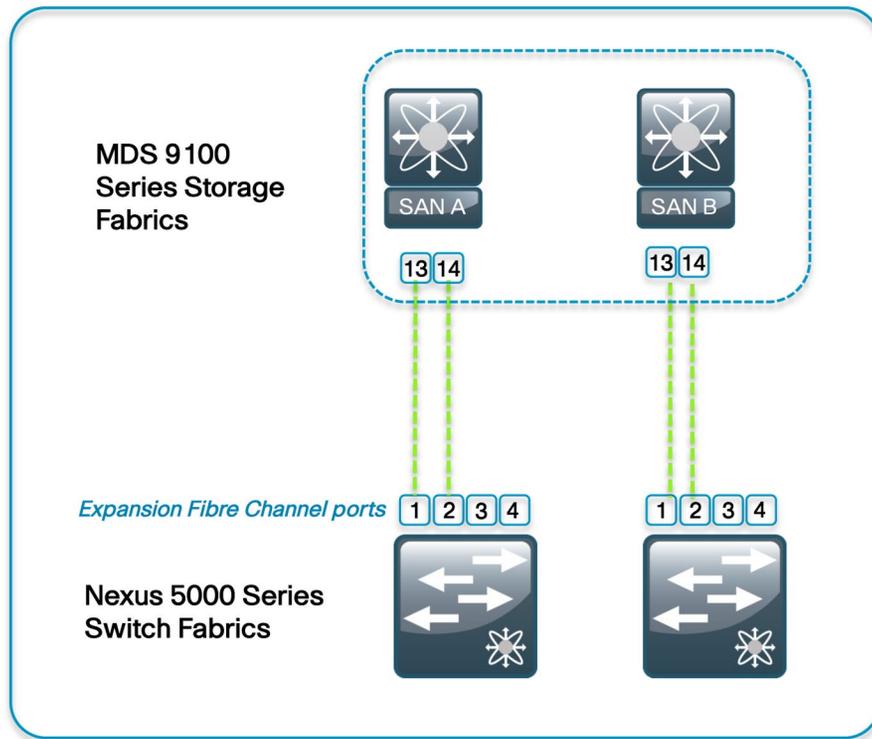
- 8-port 4-Gbs Fibre Channel card,
- Split 4-port 10-Gigabit Ethernet/4-port 4-Gbs Fibre Channel card
- 6-port 8-Gbs Fibre Channel card

Figure 58. Unified Fabric and Non-Unified Fabric Topologies



The Cisco Nexus 5010 with a Fibre Channel expansion card acts a standard Fiber Channel switch or it can act as a Fibre Channel switch in host mode. In N-port Virtualization (NPV) mode all traffic is managed at the upstream MDS. NPV allows for the storage to be configured the same as in the previous Cisco UCS chassis configuration. In the SBA Unified Computing architecture, all the Fiber Channel zoning and switching occurs upstream on the Cisco MDS 9148 switches. The Fibre Channel connectivity is configured as shown in Figure 59, with one connection from each Cisco Nexus 5000 Series switch to each SAN fabric.

Figure 59. SAN and Nexus 5010 Connectivity



The Cisco MDS 9100 Series switch was configured for NPIV in the Configuring the Fibre Channel Network Infrastructure process earlier in this guide to support the Cisco UCS fabrics. NPIV is required on the Cisco MDS 9100 Series switch when the Cisco Nexus 5000 Series switches are configured for N-port Virtualization mode.

Process

Configuring Nexus 5010

1. Configure Nexus 5010 for FCo.

The Cisco Nexus 5010 Series switch supports FCoE. To configure it for unified fabrics and support of the 2nd generation CNAs, you must configure

the following steps:

- FCoE functionality
- NPV (optional)
- Fibre Channel VSAN creation
- Fibre Channel Uplink
- VLAN association to fibre channel VSAN
- Creation of a virtual fibre channel interface
- Assign VSAN to virtual fibre channel interface
- Configure Ethernet port and trunking

Tech Tip

Configuration will be similar across both of the Cisco Nexus 5010 Series switches with the exception of the VSAN configured for SAN fabric A and for SAN fabric B. You must perform this procedure once for each of the two Cisco Nexus 5010 Series switches.

Procedure 1 Configure Nexus 5010 for FCoE

Step 1: First, enable FCoE on each Cisco Nexus 5010 Series switch.

```
feature fcoe
```

Tech Tip

A temporary 180-day license is activated when you enter the `feature fcoe` command. For long-term use, you must install a proper license. For more information, please see the Cisco NX-OS Licensing Guide at http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/licensing/guide/b_Cisco_NX-OS_Licensing_Guide.html.



Tech Tip

Enabling NPV erases the working configuration and reboots the switch. You must then reconfigure the switch over the console interface. The only information that remains is the admin username and password. Please understand the impact of this change on a production network device.

If you do not enable NPV, the Cisco Nexus 5000 Series switches are used as a switch.

All zoning and Fibre Channel configuration of the Cisco Nexus 5000 Series switches is similar to the Cisco MDS 9100 Series switch zoning and configuration in the Cisco SBA Data Center for Midsize Networks Deployment Guide.

Step 2: Enter the following at the command line.

```
feature npv
```

Step 3: Configure the VSAN on the Cisco Nexus 5000 Series switch and assign it to the interface connected to the MDS

```
vsan database
vsan 4 name finance
vsan 4 interface fc2/1
exit
```

Step 4: Configure and bring up the Fibre Channel port that is connected to the Cisco MDS 9100 Series switch.

```
interface fc 2/1
no shut
exit
```



Tech Tip

The port will need to be enabled on the MDS and have the correct VSAN assigned. Please refer to the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide* for more information on configuring the Cisco MDS 9100.

Step 5: Use the show interface brief command on the Cisco Nexus 5000 Series switch to view the operating mode of the interface. For example: In the output below, the operating mode is NP (proxy N-Port). Because the default port configuration on the Cisco MDS 9148 Series switch is auto and NPIV feature has been enabled previous in the Cisco UCS Fabric Configuration, the switch negotiates as an NP port.

```
DC-5010a# show interface brief
```

```
-----
Interface Vsan Admin Admin Status SFP Oper Oper Port
           Mode Trunk Mode          Mode Speed Channel
           Mode                                     (Gbps)
-----
fc2/1      4    NP    off  up    sw1  NP    4    --
```

Step 6: Check the Fibre Channel interface on the corresponding Cisco MDS 9148 switch.

With fibre channel configuration complete between the Cisco Nexus 5010 Series switch and the Cisco MDS 9148 Series switch, connectivity to the host can begin. On the Cisco Nexus 5010 Series switch, configure the Ethernet ports connected to the CNA in the host.

Step 7: Create a VLAN that will carry FCoE traffic to the host. In this example, VLAN 304 is mapped to VSAN 4. VLAN 304 carries all VSAN 4 traffic to the CNA over the trunk.

```
vlan 304
fcoe vsan 4
exit
```

Step 8: Create a Virtual Fibre Channel (vfc) interface for Fiber Channel traffic and bind it to the corresponding host Ethernet interface.

```
interface vfc1
bind interface Ethernet 1/3
no shutdown
exit
```

Step 9: Add the virtual Fibre Channel interface to the VSAN database

```
vsan database
vsan 4 interface vfc 1
exit
```

Step 10: Configure the Ethernet interface to operate in trunk mode.

Step 11: Also configure the interface with the FCoE VSAN and any data VLANs required by the host.

Step 12: Configure the spanning-tree port type as trunk edge.

```
interface Ethernet 1/3
switchport mode trunk
switchport trunk allowed vlan 152,304
spanning-tree port type edge trunk
no shut
```

Procedure

Verify FCoE Connectivity

Use the show interface command to verify the status of the virtual Fibre Channel interface. The interface should now be up as seen below if the host is properly configured to support the CNA. Host configuration is beyond the scope of this guide. Please see CNA documentation for specific host drivers and configurations.

Step 1: On the Cisco Nexus 5000 Series switches, use the show interface command to display the status of the virtual Fibre Channel interface.

```
show interface vfc1
vfc1 is up
  Bound interface is Ethernet1/3    Hardware is Virtual
  Fibre Channel
  Port WWN is 20:00:00:0d:ec:b4:7d:ff
  Admin port mode is F, trunk mode is off
  snmp link state traps are enabled
  Port mode is F, FCID is 0x050601
  Port vsan is 4
  1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    14 frames input, 1768 bytes
    0 discards, 0 errors
    14 frames output, 1620 bytes
    0 discards, 0 errors
  Interface last changed at Fri Apr 23 17:42:44 2010
```

Step 2: Use the show fcoe database command to display the FCoE addresses.

```
show fcoe database
-----
INTERFACE  FCID      PORT NAME                MAC ADDRESS
-----
vfc1       0x050601  21:00:00:c0:dd:11:28:89  00:c0:dd:11:28:89
```

Step 3: On the Cisco MDS 9000 Series switch, use the show flogi database command to view the addresses in the current Fiber Channel login database. The first line below is the Cisco Nexus 5010 Series switch. The second ID is the host on the vfc 1 interface.

```
mds9148a# show flogi database
-----
INTERFACE  VSAN  FCID      PORT NAME                NODE NAME
-----
fc1/13     4     0x050600  20:41:00:0d:ec:b4:7d:c0  20:04:00:0d:ec:b4:7d:c1
fc1/13     4     0x050601  21:00:00:c0:dd:11:28:89  20:00:00:c0:dd:11:28:89
```

Step 4: Use the show fcns data command to display the Fiber Channel name server database information, which differentiates the Cisco Nexus 5010 Series switch WWN from the actual host WWN. The switch appears as type NPV and the host as expected will show up as an initiator.

```
mds9148a# show fcns database
VSAN 4:
-----
FCID      TYPE  PWWN                VENDOR  FC4-TYPE:FEATURE
-----
0x050600  N     20:41:00:0d:ec:b4:7d:c0  (Cisco)  npv
0x050601  N     21:00:00:c0:dd:11:28:89  (Qlogic) scsi-fcp:init
```

Step 5: Follow the instructions in the Fibre Channel section for the *Cisco SBA for Midsize Agencies—Data Center Deployment Guide* to configure zoning and device aliases.



Tech Tip

Much of the configuration of the Cisco Nexus 5010 Series switch can also be done from within Device Manager; however, Device Manager cannot be used to configure VLANs or Ethernet Trunks on the Cisco Nexus 5010 Series switches.

Advanced Configurations

This section covers some additional configuration options for service profiles you may choose to use for production systems. Some basic configuration examples are provided, but additional features are also discussed that require referencing the product documentation or other sources to complete a full deployment.

Working with Service Profile Templates

Process

Working with Service Profile Templates

1. Create a Service Profile Template
2. Generate Service Profiles from a Template
3. Validate an Updating Template

Service profile templates allow you to rapidly create multiple service profiles that emulate servers with identical characteristics to be mapped to different physical blades during a server rollout. These templates may be configured as either initial templates, which are only used for profile creation, or updating templates that continue to be linked to the created profiles for ongoing maintenance procedures. The updating template feature is a powerful tool for managing updates to multiple servers with minimal administrative overhead.

Procedure 1 Create a Service Profile Template

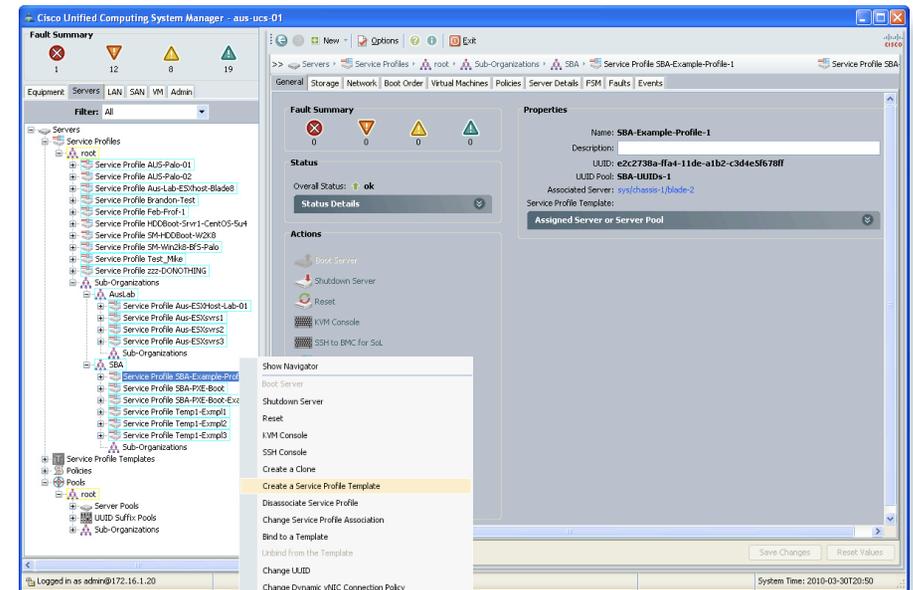
You can easily create an example service profile template from the initial service profile that you built in the Creating an Initial Service Profile for Local Boot process.

Step 1: In the **Servers** tab of the navigation pane, right-click the name of the base service profile, and choose **Create a Service Profile Template** as shown in Figure 60.

Tech Tip

Alternatively, you can clone a service profile. To do so, in the **Servers** tab of the navigation pane, right-click the profile name and choose **Create a Clone**.

Figure 60. Create Service Profile Template



Step 2: Provide a name for the template, place it under the **Root** organization and then choose the desired type of template. Our example shows the capabilities of the updating template feature, so choose **Updating Template** in the **Type** box as shown in Figure 61.

Step 3: Click **OK** to create the template with the same settings as the base service profile.

Figure 61. Template Name and Type



Step 4: After you have created the template, you can modify the attributes of the template by choosing the template name in the navigation pane and then using the tabbed details in the work pane.

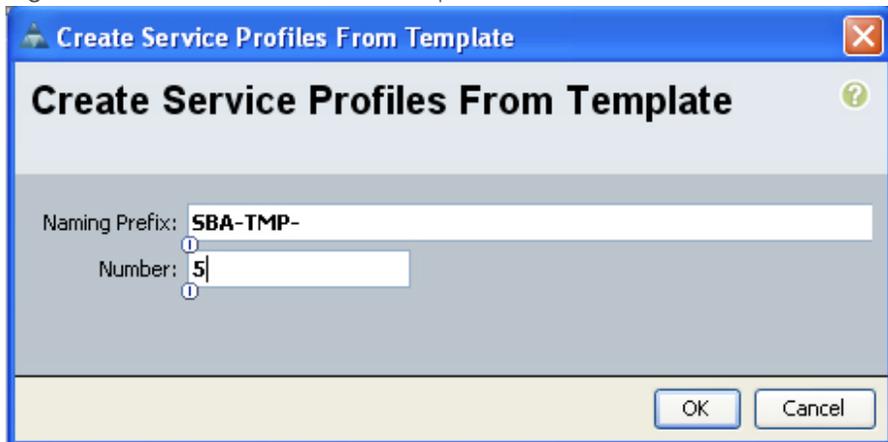
Procedure 2 Generate Service Profiles from a Template

To generate multiple service profiles from the created template:

Step 1: In the Actions area of the work pane, choose **Create Service Profiles from Template**.

Step 2: Provide a naming prefix for the profiles to be generated, and enter the quantity of profiles desired in the **Number** field as shown in Figure 62.

Figure 62. Create Profiles from Template



The template uses the naming prefix as the base name of the service profiles that it generates. For example, the naming prefix and number used in this example creates profiles SBA-TMP-1 through SBA-TMP-5.

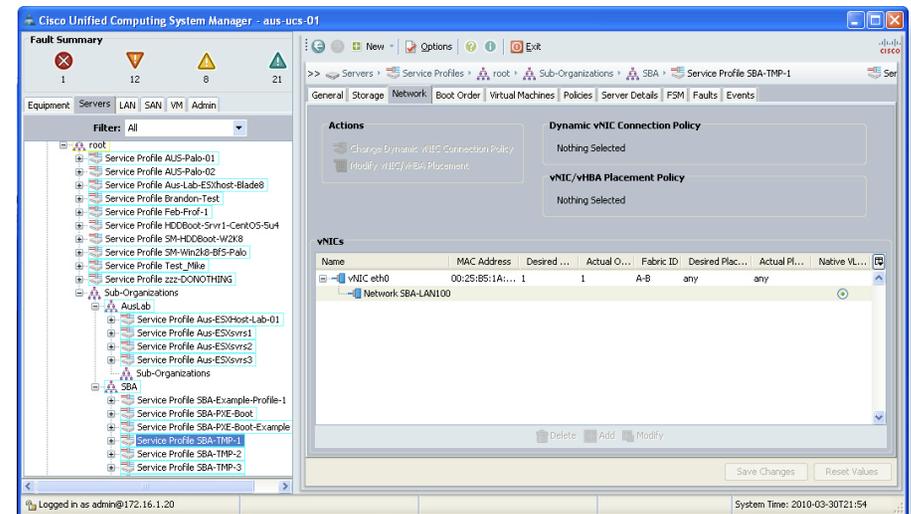
Procedure 3 Validate an Updating Template

Because this is an updating template, you can edit multiple service profiles in bulk by using the template. The template configuration created from our base service profile only includes a single vNIC for network access.

Step 1: Select the profile name in the navigation pane

Step 2: View the **Network** tab in the work pane as shown in Figure 63.

Figure 63. Network Tab

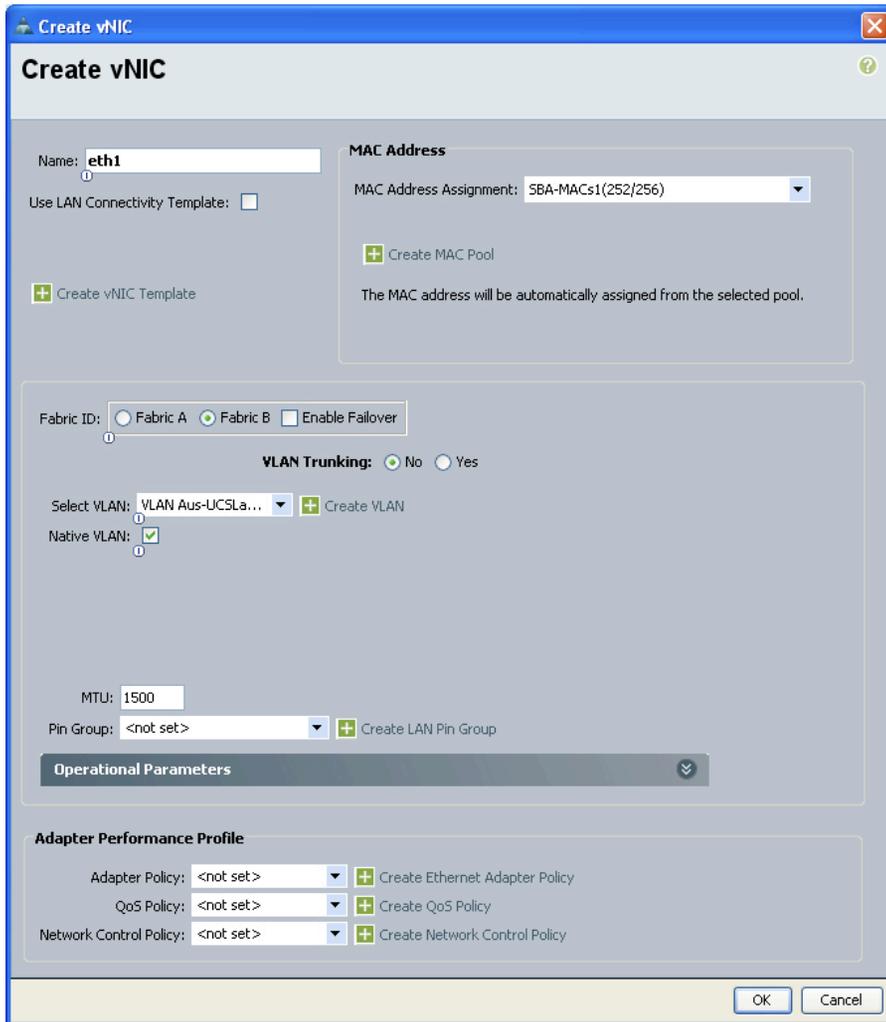


Step 3: Choose the updating template name in the navigation pane and click on the **Network** tab in the work pane for the template.

Step 4: Click **Add** on the bottom of the window to add a new vNIC to the template and also update its bound service profiles.

Step 5: Complete the configuration of the additional vNIC using the configuration settings as shown in Figure 64.

Figure 64. Add vNIC to Template

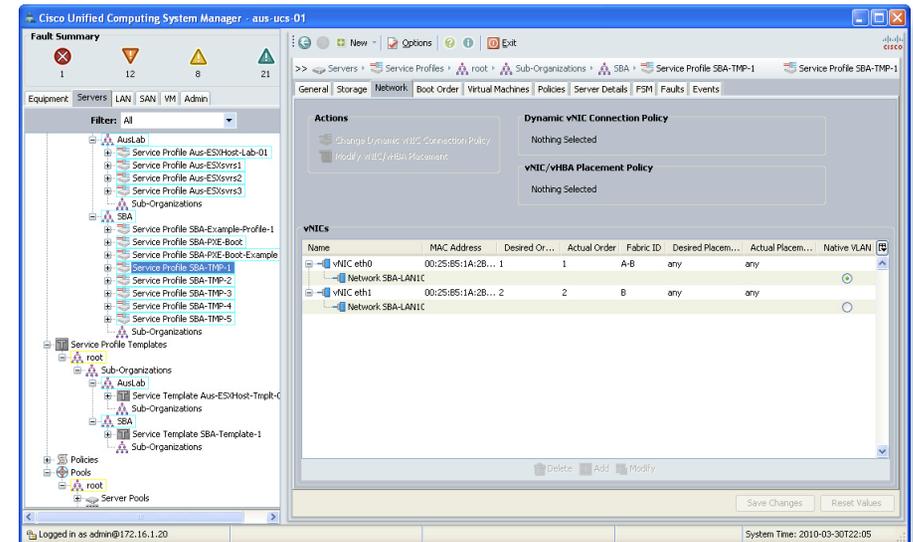


Step 6: Click OK after completing the configuration settings.

Step 7: To commit the changes to the template, click **Save Changes** at the bottom of the window.

Step 8: Verify in one of the service profiles that are bound to the updating template that the new vNIC has also been applied.

Figure 65. Service Profile Updated Through Template



As shown in Figure 65, the system has applied the same change to each service profile that is bound to the updating template.

Tech Tip

Service profiles generated from updating templates can also be later unbound from the template. After a profile is unbound, it is not affected by a change to the original template.

Service Profiles Using Multiple vNICs and Trunking

Service profiles allow the system to present one or more vNICs to each blade server. There are multiple approaches available to achieve network interface resiliency in the system. You can use a single vNIC with fabric failover, you can use NIC-teaming approaches in the installed server operating system, or you can take advantage of advanced features that link hypervisor capabilities into Cisco UCS Manager. Choosing between these options requires a clear understanding of your agency requirements, the capabilities of the installed operating system, and the functionality available in the specific network adapter hardware.

Using Fabric Failover

The Creating an Initial Service Profile for Local Boot process illustrated a basic single-vNIC configuration utilizing fabric failover.

Fabric A is the primary path for traffic when all system components are up and running. If Fabric Interconnect A or all of its associated uplinks fail, the alternate path through Fabric B can assume responsibility for forwarding traffic onto the Ethernet network. To accomplish this failover, Fabric B must transmit a Gratuitous Address Resolution Protocol (Gratuitous ARP or GARP) upstream to its connected Ethernet switches in order to update their forwarding tables with the new path to the attached servers.

This approach works cleanly for operating systems that are installed directly to the service profile and do not leverage a hypervisor.

Tech Tip

Fabric failover is a Cisco value-added feature that is specific to the Cisco UCS M71KR and M81KR interface adapters. Support for additional interface adapters from other manufacturers is planned for a future release of Cisco UCS. If you plan to use a configuration that uses this feature, ensure that the mezzanine adapter that you choose for your deployment supports this capability.

Using Dual vNICs for Failover

The Cisco UCS 82598KR-CI and M71KR network adapters are capable of presenting a maximum of two vNICs to a service profile. If a hypervisor system is running on the service profile, the capability of Cisco UCS fabric failover is limited to the hypervisor instance itself. As of Cisco UCS release 1.2(1d), the resilient fabric cannot provide a GARP for the MAC addresses of the guest-OS virtual machines (VMs) running under the hypervisor. The VMs carry their own MAC addresses and are serviced by the virtual switching instance running within the hypervisor. To provide resilient network connectivity for the VMs, dual interfaces must be configured in the service profile. The inherent NIC-teaming capability of the hypervisor can be used to provide a resilient network path through these two interfaces.

To extend our example service profile to contain a dual-vNIC capability, change the configuration of the existing vNIC to disable fabric failover.

Step 1: On the **Servers** tab in the navigation pane, select the service profile name and then choose the **Network** tab in the work pane.

Step 2: Select the vNIC that you previously created and click **Modify** at the bottom of the screen.

Step 3: Uncheck **Enable Failover**, as shown in Figure 66.

Figure 66. Disable Fabric Failover on vNIC

The screenshot shows the 'Modify vNIC' configuration window. The 'Name' is 'eth0'. Under 'MAC Address', the 'MAC Address Assignment' is 'SBA-MACs(243/256)'. In the 'Fabric ID' section, 'Fabric A' is selected, and the 'Enable Failover' checkbox is checked. The 'VLAN Trunking' section has 'No' selected. The 'Operational Parameters' section is collapsed. The 'Adapter Performance Profile' section has 'Adapter Policy', 'QoS Policy', and 'Network Control Policy' all set to '<not set>'. There are 'OK' and 'Cancel' buttons at the bottom right.

Step 4: Click **OK** to confirm the change.

Step 5: To add a second vNIC to the profile, click **Add** at the bottom of the **Network** tab in the work pane.

Step 6: Assign a name and MAC address pool to the vNIC, and choose the appropriate VLAN configuration to be consistent with the selections for the first vNIC.

Step 7: In the **Fabric ID** area for this vNIC, choose **Fabric B** to provide a resilient path for network traffic, and ensure that **Enable Failover** is unchecked as shown in Figure 67.

Figure 67. Adding a Second vNIC

The screenshot shows the 'Create vNIC' dialog box with the following configuration details:

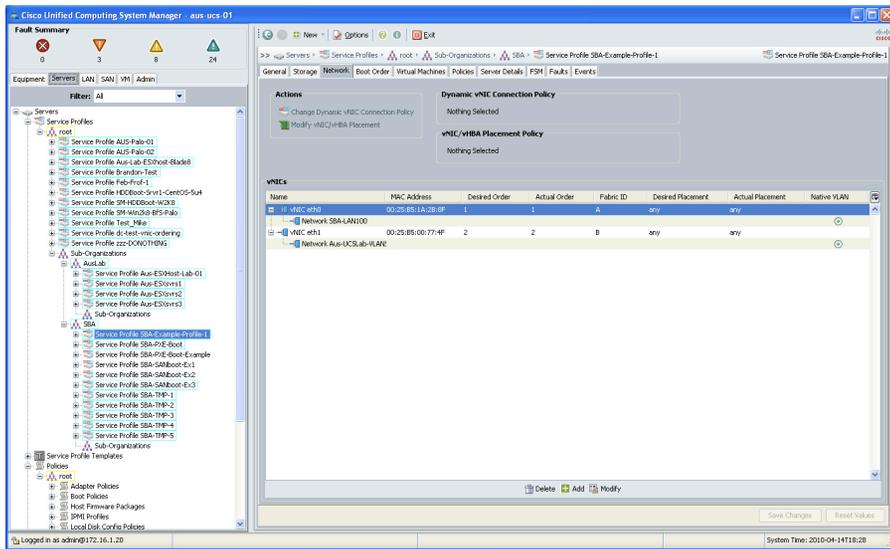
- Name:** eth1
- MAC Address:** MAC Address Assignment: SBA-MACs(249/256)
- Fabric ID:** Fabric A (unselected), Fabric B (selected), Enable Failover (unchecked)
- VLAN Trunking:** No (selected), Yes (unchecked)
- Select VLAN:** VLAN SBA-LAN10... (selected)
- Native VLAN:** (unchecked)
- MTU:** 1500
- Pin Group:** <not set>
- Adapter Performance Profile:** Adapter Policy: <not set>, QoS Policy: <not set>, Network Control Policy: <not set>

Step 8: Click **OK**. Then click **Save Changes** at the bottom of the **Network** tab in the work pane to ensure that changes are committed to the service profile.

The network configuration should show one interface assigned to Fabric A, and the second interface assigned to Fabric B.

Network Tab with Dual vNICs

Figure AC-9



After you have created the dual-vNIC service profile, you can install and/or boot the hypervisor system and it will recognize the availability of two network interfaces to the system. Use the configuration specific to the operating system or hypervisor in use to enable NIC-teaming capabilities for network interface resiliency

Using a Virtual Interface Card

The Cisco UCS M81KR Virtual Interface Card provides the capability to assign more than two vNICs to a given service profile. The specific number of interfaces is dependent on the number of uplinks in use between the chassis and the fabric interconnects and on the operating system in use on the blade server. These interfaces may be used for multiple purposes by either a directly installed bare-metal operating system or a hypervisor-based system.

Tech Tip

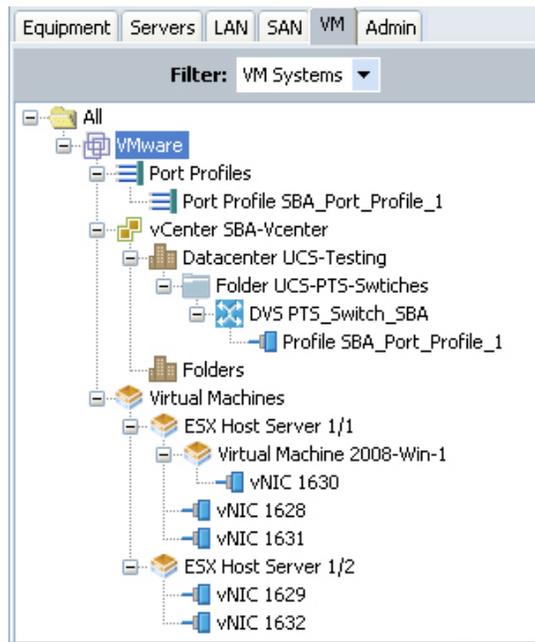
Service Profiles with more than two static vNICs can only be applied to the Cisco UCS M81KR Virtual Interface Card. Applying such a service profile to servers with other mezzanine cards will result in a config failure error in UCS Manager due to an insufficient number of available vNICs.

Virtual Machine Integration

Cisco UCS Manager supports direct integration with hypervisor management systems to facilitate dynamic allocation of vNIC resources on the Cisco UCS M81KR Virtual Interface Card to virtual machines. This capability is available for VMware ESX 4.0 Update 1 as of Cisco UCS release 1.2(1b). This feature is an implementation of the Cisco VN-Link technology, similar to the Cisco Nexus 1000V Distributed Virtual Switch.

Cisco UCS Manager provides extensions which can be exported into a file, which is then installed as a Plug-In into vCenter Server. This linkage allows the management of networking in vSphere using Virtual Distributed Switch configuration that is aware of the port profiles created in Cisco UCS Manager. Figure 68 shows the contents of the VM tab in a Cisco UCS system that has been linked to a vCenter Server. The port profiles and the virtual switching constructs are created in Cisco UCS Manager. The virtual machine information is learned through the link to vCenter Server.

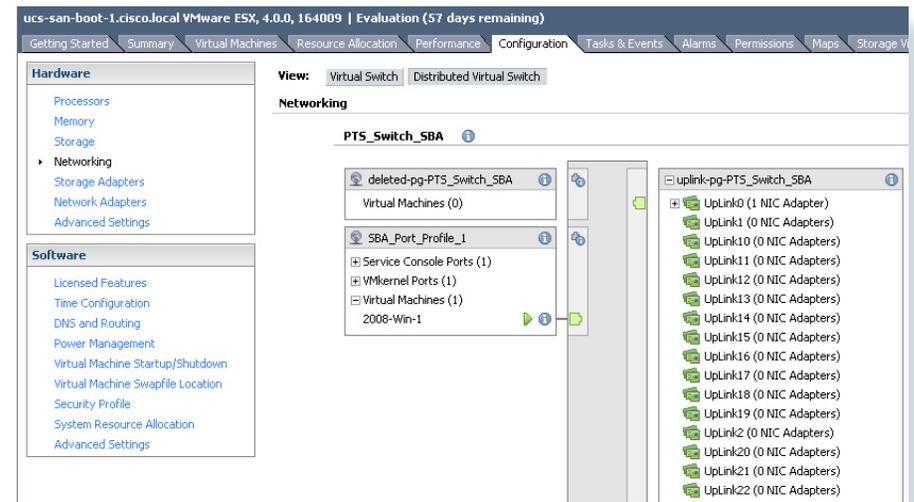
Figure 68. UCS Manager VM Tab Detail



You can view the corresponding vCenter information in a vSphere host by accessing the Configuration tab, clicking Networking in the Hardware area, and then clicking Distributed Virtual Switch, as shown in Figure 69. The system defines a pool of available uplinks which are dynamically assigned to virtual machines as required.

This configuration allows all port-specific information such as traffic statistics and QoS configuration to move with a virtual machine that is moved between hosts using vMotion.

Figure 69. vSphere Distributed Virtual Switch Linked to UCS Manager



Tech Tip

For more information on UCS Manager and VMware vSphere see: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.2.1/UCSM_GUI_Configuration_Guide_1_2_1_chapter28.html.

Enabling VLAN Trunking on vNICs

You can also easily configure vNICs to support trunking of multiple VLANs to a blade server that is using IEEE standard 802.1Q headers. Cisco UCS Manager supports a native VLAN configuration selection that directs the interface on how to handle untagged traffic.

The use of multiple VLANs is common in a server virtualization environment to handle traffic destined to different virtual machines and administrative traffic destined to a hypervisor service console or kernel interface. In some cases, using a Cisco UCS M81KR Virtual Interface Card may allow enough interfaces to be assigned to the system to preclude the need for trunking multiple VLANs on a single interface.

Service Profiles Using Multiple vHBAs

When you set up a server profile with a Fibre Channel SAN, two resilient storage fabrics, A and B, will usually be available to the server profile.

One caveat of booting from SAN is that, normally, both SAN fabrics are configured for a Server Profile. If you are installing a fresh operating system on a LUN, it can be simpler to make one vHBA available until you complete the install of the host operating system. Then with host-specific or storage-provider-specific multi-pathing software installed and properly configured, bring up the second vHBA to fabric B.



Tech Tip

For example: Windows requires a single HBA during install until multipath drivers are installed: <http://www.microsoft.com/downloads/details.aspx?FamilyID=f4095fae-553d-4700-aafa-1cce38b5618f&displaylang=en>

Other operating systems have different requirements. Please refer to your specific operating system documentation for handling redundant SAN connections.

The setup for the second vHBA is identical to the first vHBA with two exceptions. Normally VSANs are not identical between the fabrics. In the example used earlier, VSAN 4 is in fabric A and VSAN 5 is in fabric B. The other exception is the fabric selected. Fabric B will be selected so all SAN traffic will flow to the second Cisco UCS 6100 Series Fabric Interconnect and off to the SAN fabric B.

Multiple Fibre Channel ports from the Cisco UCS 6100 Series Fabric Interconnect can connect to an upstream Cisco MDS 9100 Series Multilayer Fabric Switch. Managing which server profiles use which uplink can be configured automatically or statically (using manual Pin Groups).



Tech Tip

Please reference the Cisco UCS documentation, Configuring SAN Pin Groups section at: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/GUI_Config_Guide_chapter19.html

Appendix A

Configuration Values Matrix

The following matrix can be used as a worksheet when following the basic configuration examples in this guide.

Item	Example Value	Configured Value
Ethernet Infrastructure		
Port-Channel 6100A	50, VPC ID 50	
Port-Channel 6100B	51, VPC ID 51	
Uplink Ports 6100A	1/17-20	
Uplink Ports 6100B	1/17-20	
Server Ports 6100A	1/1-4	
Server Ports 6100B	1/1-4	
Fibre Channel Infrastructure		
VSAN Fabric A	4 (Finance)	
VSAN Fabric B	5 (Finance)	
Fibre Channel Uplink Ports 6100A	FC 2/1-2	
Fibre Channel Uplink Ports 6100B	FC 2/1-2	
Initial Fabric Interconnect Setup		
Fabric A Physical IP Address / Netmask	192.168.28.51/255.255.255.0	
Default Gateway	192.168.28.1	
Cluster IP Address	192.168.28.50	
DNS (optional)	192.168.28.10	
Fabric B Physical IP Address Netmask	192.168.28.51/255.255.255.0	
Management IPs		
KVM IP Address Pool	192.168.28.201 (Size 32)	
KVM Subnet mask	255.255.255.0	
KVM Default Gateway	192.168.28.1	

Item	Example Value	Configured Value
Initial Service Profile		
UUID Pool	0610:000000001024 (Size 256)	
MAC Address Block	00:25:B5:01:0c:01 (Size 256)	
VLANs	28-29	
Target WWN	50:0a:09:81:89:0a:df:b1	
SAN Boot		
WWNN Block	20:00:00:25:B5:00:00:00 (Size 64)	
WWPN Block	20:00:00:25:B5:00:77:00 (Size 256)	
Service Profile Interfaces		
LAN Interface	eth0	
HBA Interface	fc0	
C-Series FCoE		
VSAN	4	
FCoE VSAN	304	

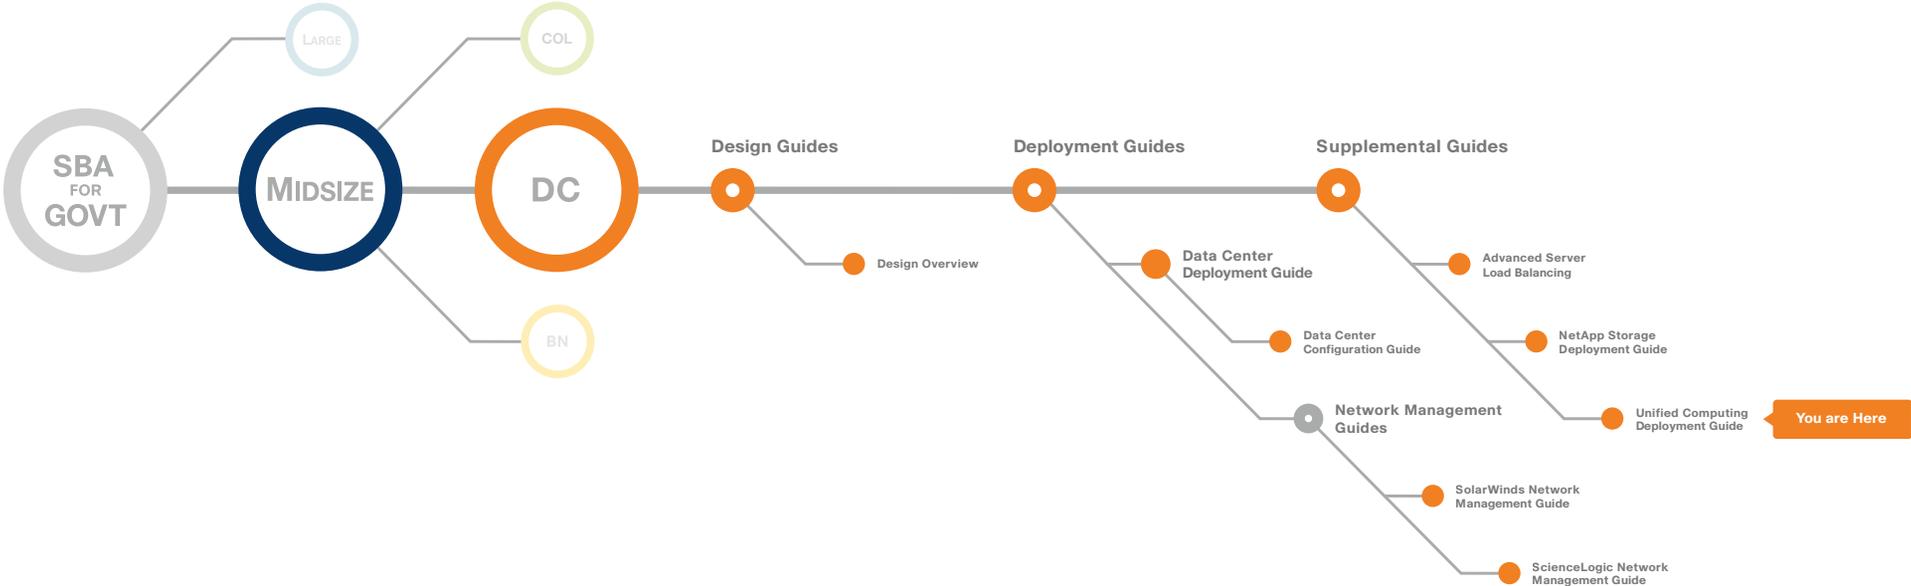
Appendix B

Equipment List

This table provides a listing of equipment hardware and software versions used in a lab validation of this design. Note that representative part numbers and descriptions are provided for the primary components of the topology. For an orderable configuration, please work with your Cisco representative.

Functional Area	Product	Part Numbers	Software Version
UCS Fabric Interconnects	UCS 6120XP 20-port Fabric Interconnect	N10-S6100	Cisco UCS release version 1.2(1d)
Fibre Channel Module	6-port 8Gb FC/Expansion module/ UCS 6100 Series	N10-E0060	Cisco UCS release version 1.2(1d)
UCS Blade Chassis	UCS 5108 Blade Server Chassis	N20-C6508	Cisco UCS release version 1.2(1d)
UCS I/O Module	UCS 2104XP Fabric Extender	N20-I6584	Cisco UCS release version 1.2(1d)
Blade Server, Half-slot	UCS B200 M1 Blade Server	N20-B6620-1	Cisco UCS release version 1.2(1d)
Blade Server, Full-slot	UCS B250 M1 Blade Server	N20-B6620-2	Cisco UCS release version 1.2(1d)
Blade Server Interface Mezzanine Adapters	UCS M81KR Virtual Interface Card	N20-AC0002	Cisco UCS release version 1.2(1d)
UCS Rack Mount Server, 1RU	UCS C200 M1 Server	R200-1120402	Cisco UCS release version 1.2(1d)
UCS Rack Mount Server, 2RU	UCS C210 M1 Srvr	R210-2121605	Cisco UCS release version 1.2(1d)
Data Center Switching Fabric	Nexus 5010	N5K-C5010P-BF	NXOS 4.1(3)
Data Center Fabric Extenders	Nexus 2148T	N2K-C2148T-1GE	NXOS 4.1(3)
Fibre Channel SAN Switches	MDS 9148	DS-C9148D-8G16P-K9	5.0(1a)
Storage Array	NetApp FAS3100 Series	FAS3140	7.3.2
Converged Network Adapter - Rackmount	QLogic 8152	QLE8152	

Appendix C: SBA for Midsize Agencies Document System





SMART BUSINESS ARCHITECTURE



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

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 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)

C07-641161-00 12/10