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.

Collapsed Data Center and Campus Core Deployment Guide

Cisco Nexus 7000 Supplemental

SBA FOR GOVERNMENT

Revision: H2CY10

Using this Borderless Networks Guide

This is guide is a concise reference on consolidating the campus and data center core layers and is broken up into a few sections.

The Introduction outlines the use-cases that often drive the consolidation of the campus and data center cores

The Technology Overview section introduces the Cisco Nexus 7000 Series and the capabilities that uniquely position the platform as the device that should be utilized in a collapsed campus and data center core.

The Configuration Details section details the processes and procedures required to deploy the Nexus 7000 Series as a collapsed campus and data center core.

This document is for the reader who:

- Has read the *Cisco Smart Business Architecture (SBA) for Government Large Agencies—Borderless Networks LAN Deployment Guide*
- Has in total 2000 to 10,000 connected employees
- Has a campus LAN that supports up to 5000 connected employees co-located with a data center
- Wants to consolidate the campus and data center core devices
- Has high 10-Gigabit Ethernet density in the campus and data center core
- Has IT workers with a CCNA® certification or equivalent experience
- Wants the assurance of a tested solution

Who Should Read This Guide

This guide should be of interest to anyone in a large agency that wants to understand why they might want to utilize the Cisco Nexus 7000 Series when they are collapsing the campus and data center core.

The audience also includes information technology staff and technology resellers who want to understand more about the Cisco Nexus 7000 Series and how to deploy it in a collapsed campus and data center core environment.

This guide does assumes a CCNA-level technical background.

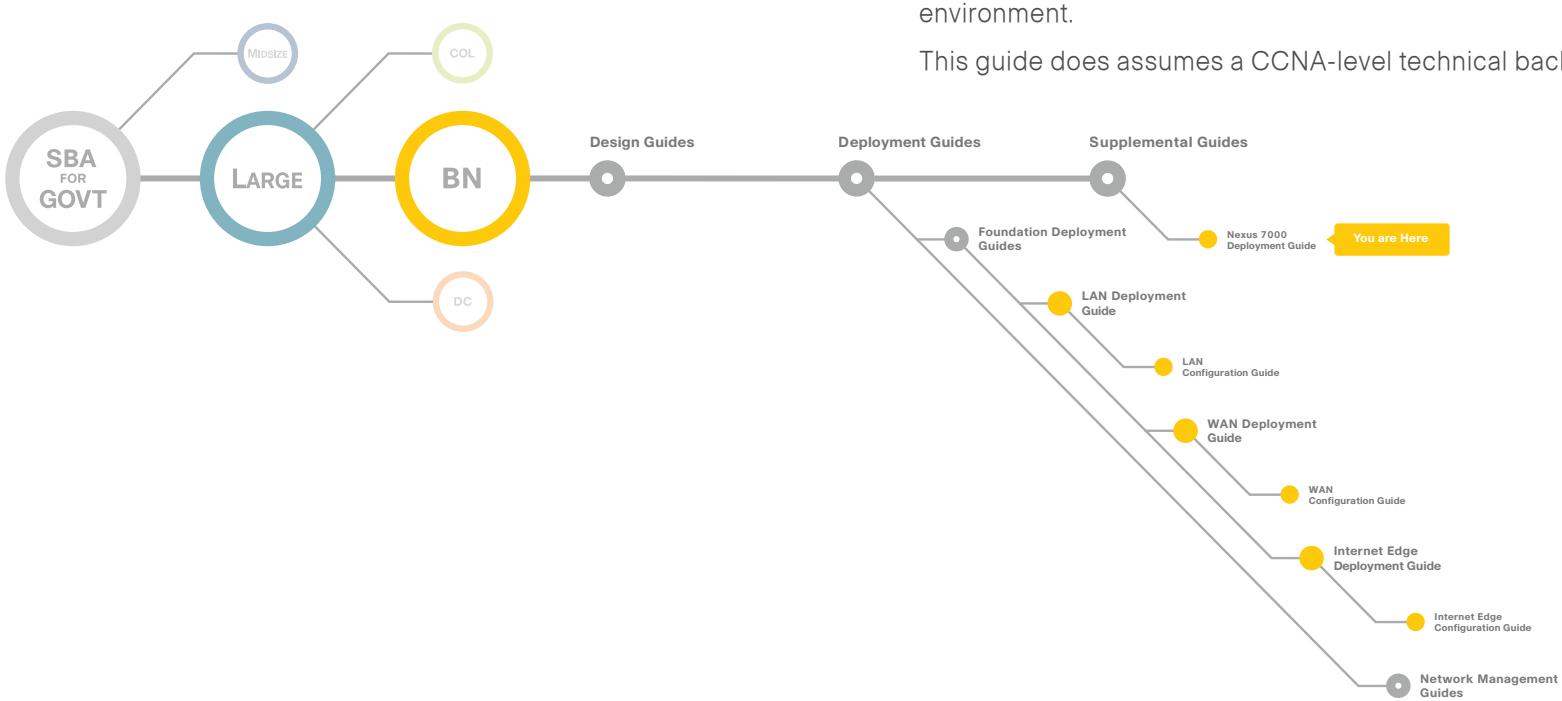

Table of Contents

Introduction 1

Agency Overview 2

Technology Overview 3

Configuration Details 6

Summary 11

For Additional Information 11

Appendix A: Large Agencies LAN Deployment Product List 12

Appendix B: SBA for Large Agencies Document System 13

Introduction

This guide is a companion document to the *Cisco SBA for Large Agencies—Borderless Networks Design Guide and LAN Deployment Guide*.

The Cisco SBA for Large Agencies is a prescriptive architecture that delivers an easy to use, flexible and scalable network with wired, wireless, security, WAN, and internet edge modules. It eliminates the challenges of integrating the various network components be using a standardized design that is reliable with comprehensive support offerings.

The Cisco SBA for Large Agencies—Borderless Networks is designed to address the common requirements of agencies with 2000 to 10,000 employees. Each agency is unique however and so are its requirements so we ensured that the Cisco SBA was built so that additional capabilities could be added on without redesigning the network.

One way the Cisco SBA accomplishes this extensibility is by braking down the architecture into three primary layers, Network Foundation, Network Services, and User Services, as illustrated in Figure 1.

A collapsed campus and data center core is part of the Network Foundation. The reliable network foundation provided by the Cisco SBA helps ensure a government agency can rely on the network and user services required to operate effectively and achieve the agency mission objectives.

To learn more about Cisco SBA, visit:

http://www.cisco.com/go/smartarchitecture or

http://www.cisco.com/go/partner/smartarchitecture

Figure 1. SBA Model

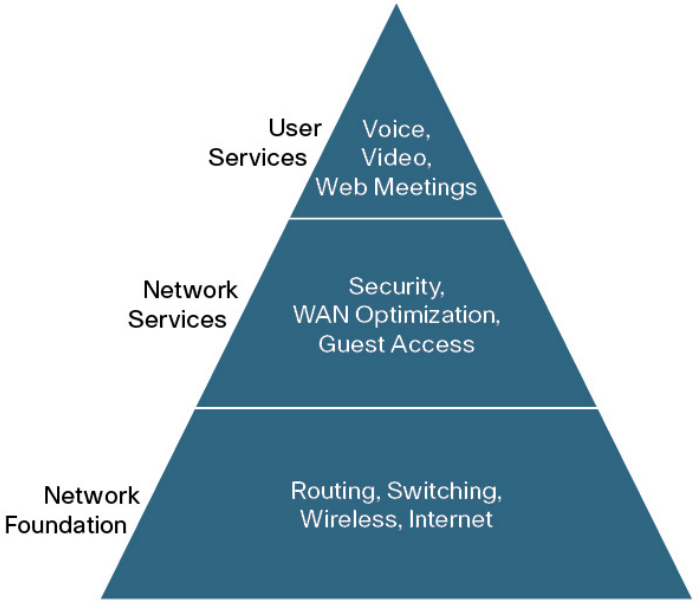

Agency Overview

As the LAN environment at a larger facility grows it often creates the need to use multiple LAN distribution layer blocks. As the number of required distribution layer blocks in a facility grows beyond two or three a solution is required to reduce the need and cost of fully meshing all interconnectivity while maintaining a design that provides a reliable infrastructure.

When an agency has a large primary site the agency's primary data center is often located there also. If the data center is large enough that it also includes multiple distribution layer blocks then the data center will also need a solution for interconnectivity.

The similarities between the needs of the campus and the data center present an opportunity to consolidate the needs into a single set of devices which can save both operational and equipment costs.

Notes

Technology Overview

Cisco has a rich tradition of building scalable and robust LAN and campus networks using the Cisco Catalyst switch line. The Catalyst 6500 series switch has a long history as the flagship of large campus aggregation and core layer networks. As of 2008, state-of-the-art data center networks started using Cisco Nexus family switching products. Specifically, data center aggregation (also referred to as distribution), and core layers are built using the Cisco Nexus 7000 Series switch.

The Cisco Nexus 7000 Series delivers high performance, port density, availability, and resiliency with a comprehensive feature set targeted for the core of data center and campus LAN networks. The Cisco Nexus 7000 Series switch effectively addresses data center core and aggregation requirements such as a high density 10-Gigabit Ethernet interface, robust Layer 3 protocols, and a zero-service disruption architecture. In environments where the core device also acts as a data center interconnect (DCI) platform, the Cisco Nexus 7000 Series switch supports Overlay Transport Virtualization (OTV), an industry solution that enables the extension of Layer 2 over Layer 3 networks, without the operational complexities of other interconnect solutions.

Outside of the data center, the Cisco Nexus 7000 can also be used in the core layer of the campus LAN. The main driver for utilizing the Cisco Nexus 7000 Series in the core of a campus LAN network is the ability to collapse the LAN and the data center core, which usually are on separate pairs of devices, onto a single pair of Cisco Nexus 7000 Series switches (Figure 2).

The Cisco Nexus 7000 Series switch is uniquely positioned as the platform to support a consolidated campus LAN and data center core because of its high 10-Gigabit Ethernet density, resiliency and high availability, and extensive support of virtualization.

With higher density and the potential for combining multiple layers, the bar for availability and resiliency moves higher. From a high-availability standpoint, the Cisco Nexus 7000 Series is a very robust and resilient platform. In addition to the hardware redundancy that is built into every hardware component of the Cisco Nexus 7000 Series, its operating system, Cisco NX-OS, provides a unique, highly-resilient capability that supports non-stop packet forwarding, even in the event of a software upgrade with the hitless in-service software upgrade (ISSU) feature. Furthermore, Cisco NX-OS on the Cisco Nexus 7000 Series supports graceful restart, as well as non-stop routing (NSR) functionalities for routing protocols the operating system is capable of detecting and restarting a faulty process, without disrupting the other operations of the switch. The zero-service disruption architecture of the Cisco Nexus 7000 Series switch is paramount in a collapsed core environment, where infrastructure failure or maintenance windows can have a broad impact.

Figure 2. Collapsed Data Center and LAN Core

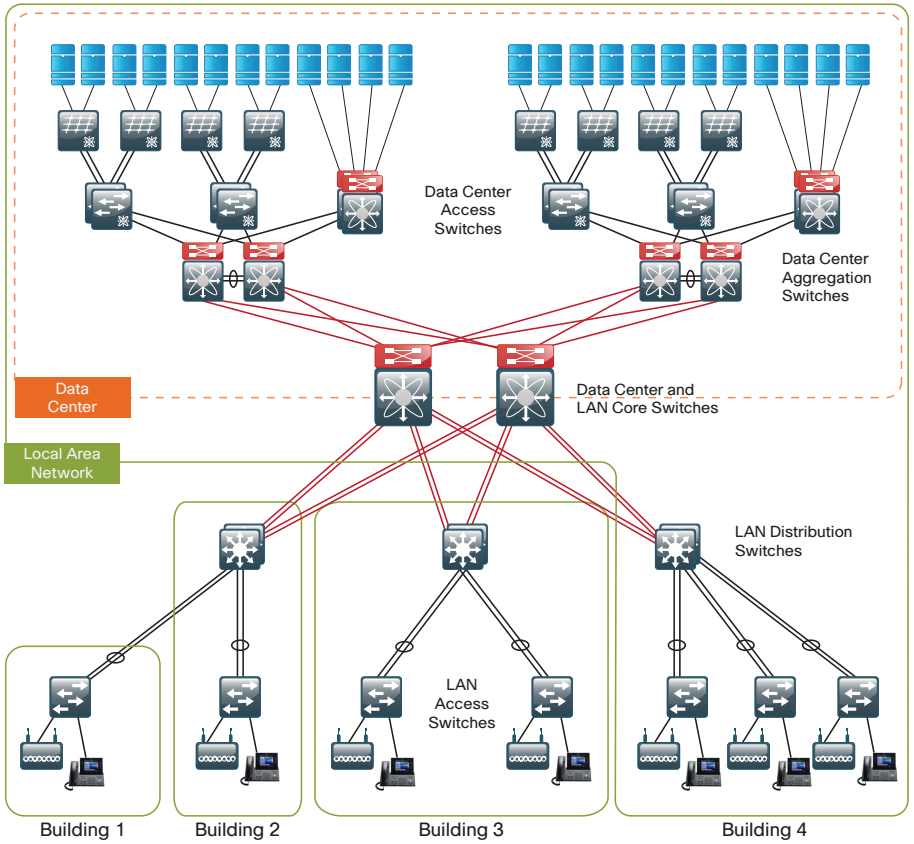

In a data center network with a 3-tier architecture (Figure 3), the Cisco Nexus 7000 Series switch is the state-of-the-art Layer 3 switch for the core and aggregation layers. Typically the aggregation layer is the Layer 2 or Layer 3 boundary.

Figure 3. 3-Tier Data Center Architecture

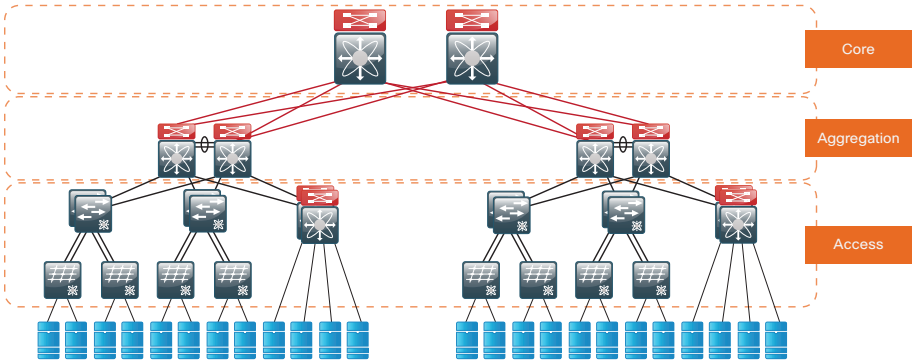

In a small or medium size data center network, the core and aggregation may collapse into one single layer that leverages the virtual device contexts capability for consolidation (Figure 4).

Virtual device contexts (VDCs) allow a Cisco Nexus 7000 Series switch to be carved up into logically separate devices. Each of those entities can be managed separately to allow operational independence and isolation from each other. This isolation can be useful for environments where layers of networks may be physically combined while logically staying independent, or where organizational units have to be separated from a management and configuration access standpoint (such as for compliance reasons).

Figure 4. Collapsed Core/Aggregation Data Center Architecture

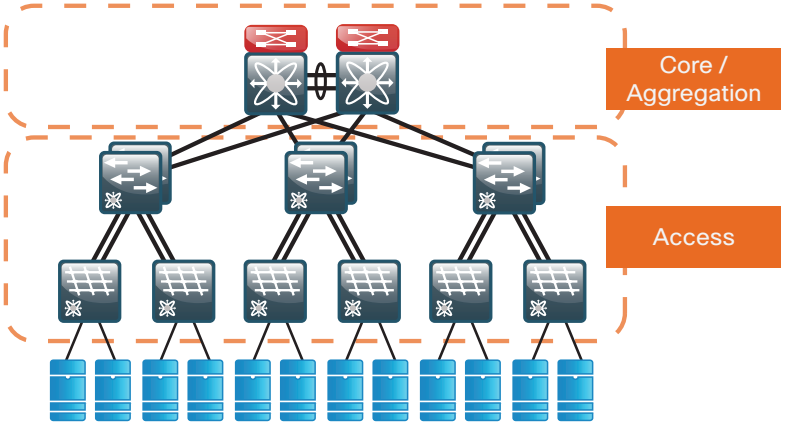

For example, a data center core and aggregation can fall into one pair of switches. To still follow the hierarchical network model as well as to respect operational boundaries, device virtualization can be used via a virtual device context (VDC), which is the first control plane virtualization in a network device. A single switch can be carved up into a maximum of four logical switches (that is, four VDCs), where resources such as interfaces, memory, and so on are explicitly allocated to the respective VDC. Besides offering consolidation and operational benefits, the configuration of VDCs offers flexible separation and distribution of resources. The hardware and software isolation per VDC not only delineates administrative contexts, but also allows scaling the overall systems. The virtualization capabilities of VDCs are key to deploy a collapsed core on a single Cisco Nexus 7000 Series switch with multiple logical cores.

The aggregation layer usually acts as a Layer 2 or Layer 3 boundary. Along with routing technologies, the Cisco Nexus 7000 Series switch offers flexible switching functionalities, such as virtual Port Channels (vPC), rapid spanning tree, and Port Channels, that are all handled in a stateful fashion to guarantee a very high level of system reliability.

Core Layer Architecture

The core layer of the LAN is a critical part of the scalable network, yet by design, is one of the simplest. Like the distribution layer, the core layer aggregates connectivity, but for multiple distribution layers instead of access layers. As networks grow beyond three distribution layers in a single location, a core layer should be used to optimize the design.

Beyond the simple aggregation of connectivity, the core layer serves to reduce the number of paths between distribution layers which in turn lowers the time required to converge the network after a failure. By upgrading bandwidth between a distribution layer and the core, multiple distribution layer blocks can benefit from the increase versus the need to upgrade the bandwidth to every other device in a design without a core. The core layer is especially relevant to designs where the data center resources might be collocated with the LAN.

In large modular and scalable LAN designs, a core layer is used to aggregate multiple user connectivity distribution layer blocks. In designs with a collocated data center, the core provides high speed fanout connectivity to the rest of the network. The core layer also serves as the interconnect for the Wide Area Network (WAN) and Internet Edge distribution layer blocks. Because of this central point of connectivity for all data flows, the core is part of the backbone IP routing address space and is designed to be highly resilient to protect from component-, power-, or operational-induced outages.

The core layer in the SBA design is based on two physically and logically separate switches. Connectivity to and from the core should be Layer 3 only. No VLANs should span the core to drive increased resiliency and stability. Since the core does not need to provide the same services or boundaries that the distribution layer does, the two-box design does not significantly increase the complexity of the solution. The core layer should not contain highly-complex or high-touch services that require constant care and tuning, to avoid downtime required by complex configuration changes, increased software upgrades for new services, or links that toggle up/down as part of normal operations like user endpoint connectivity.

The core is built on dual switches to provide a completely separate control plane housed on each switch, that provides redundant logic, line cards, hardware, and power for the backbone operation. Each distribution layer block, router, or other appliance connecting to the core should be dual-homed with an EtherChannel or link to each core switch. This dual homed approach provides Equal Cost Multiple Path (ECMP) load sharing of IP traffic across links for traffic traversing the core, and fast failover based on either EtherChannel or ECMP alternate routes without waiting for routing protocol topology changes to propagate the network.

The core is designed to be high speed and provides for connectivity ranging from Gigabit Ethernet, Gigabit EtherChannel, 10 Gigabit Ethernet, and up to 10 Gigabit EtherChannel. The core can provide non-blocking bandwidth based on design and configuration. EtherChannel links homed to a switch should be spread across line cards when possible.

The core switches can be provisioned with dual supervisors for Stateful Switchover (SSO) operation to protect the core bandwidth in the event a control plane hardware or software failure occurs. The core switches are Nonstop Forwarding (NSF) aware to provide enhanced resilience for any dual supervisor connected devices and NSF capable if provisioned with dual supervisors per switch.

Attaching the Data Center Aggregation Layer

The links from the Core layer facing the data center aggregation or distribution layer are configured in the same way as the links facing campus distribution layers. More specifically these are routed layer 3 interfaces with point-to-point subnets and a routing configuration as outlined below.

The Data Center aggregation layer configuration follows a similar model as the campus distribution layer. Integration of services such as firewalling, intrusion prevention and load-balancing are typically added. Design aspects and sample configurations for the aggregation layer using the Cisco Nexus 7000 Series switch are further described in the design guides referenced at the end of this document.

Configuration Details

This chapter describes the configuration for the core layer using the Cisco Nexus 7000 Series. Only core relevant features are described. The core layer is typically a Layer 3 configured device; therefore, configuration details of Cisco Nexus 7000 Series Layer 2 features, such as spanning tree, VLAN Trunking Protocol (VTP), and virtual Port Channel (vPC), are intentionally omitted. For those scenarios that require having spanning tree as a safeguard, Rapid Spanning Tree Protocol (RSTP) is enabled by default on the Cisco Nexus 7000 Series switch and does not require any additional configuration.

The core layer design is based on dual switches; therefore, programming of the core devices is symmetrical for simplicity, except for IP addressing and services that must be unique in the network. Because we cover the global configuration options in detail in the Access Layer section, the commands are listed here for easy reference.

Process

1. Platform Configuration
2. LAN Switch Global Configuration
3. Core Switch Global Configuration
4. Connectivity to LAN and Data Center Aggregation Switch Configuration

Procedure 1 Platform Configuration

Some platforms require a one-time initial configuration prior to configuring the features and services of the switch. If you don't have a platform listed in the following steps they can be skipped.

Step 1: Software License Prerequisites

The Cisco Nexus 7000 Series offers a simplified software management mechanism based on software licenses. These licenses are enforceable on a chassis basis and enable a full suite of functionalities. As the core layer is characterized by a Layer 3 configuration, the Cisco Nexus 7000 Series switch requires the Enterprise LAN license, which enables routing functionalities.

The VDC configuration is not included in this addendum. However, for those deployments where VDC is recommended, the Advanced LAN license needs to be installed. Other baseline features (such as QoS and Layer 2) do not require any additional licenses.

Step 2: Feature Commands

Due to the modular nature of Cisco NX-OS, processes are only started when a feature is enabled. As a result, commands and command chains only show up once the feature has been enabled. For licensed features, the feature **feature-name** command can only be used once the appropriate license is installed. For trial and testing purposes, there is a grace period of 120 days during which all features can be tried within the system. To enable this grace period, use the **license grace-period** command.

Procedure 2 LAN Switch Global Configuration

Within this design, there are features and services that are common across all LAN switches regardless of the type of platform or role in the network. These are system settings that simplify and secure the management of the solution.

This procedure provides examples for some of those settings. The actual settings and values will depend on your current network configuration.

Common Network Services Used in the Deployment Examples

Domain Name:	cisco.local
Multicast RP:	10.4.60.254
Authentication Control System:	10.4.200.15
Network Time Protocol Server:	10.4.200.17

Step 1: Configure the Device Hostname

Configure the device hostname and domain name to make it easy to identify the device.

```
hostname [hostname]
ip domain-name cisco.local
```

Step 2: Configure Device Management Protocols

Secure Shell (SSH) is an application and a protocol that provide a secure replacement to RSH and Telnet. Secure management access is enabled through the use of the SSH protocol. The protocol is encrypted for privacy.

SSH is enabled by default and does not require any specific configuration. Cisco NX-OS can store SSH authentication keys and therefore simplify the login process while still using strong authentication. Non-secure protocols like Telnet are disabled by default.

Simple Network Management Protocol (SNMP) is enabled to allow the network infrastructure devices to be managed by a network management system (NMS). SNMPv2c is configured both for a read-only and a read-write community string. The configuration of the NMS is outside the scope of this document.

The basic SNMP configuration is as follows:

```
snmp-server community cisco group network-operator
snmp-server community cisco123 group network-admin
```

Step 3: Configure Secure User Authentication Configuration

Authentication, Authorization and Accounting (AAA) is enabled for access control. All management access to the network infrastructure devices is controlled with AAA.

The AAA server used in this architecture is the Cisco Authentication Control System. Configuration of ACS is discussed in the ACS Deployment Supplement.

RADIUS is the primary protocol used to authenticate management logins on the infrastructure devices to the AAA server. A local AAA user database is also defined each network infrastructure device to provide a fallback authentication source in case the centralized RADIUS server is unavailable.

To enable RADIUS authentication and specify a RADIUS server, use the following commands:

```
username admin password c1sco123 role network-admin
aaa authentication login default group radius
radius-server host 10.4.200.15 key 7 "SecretKey"
```

Step 4: Setup Synchronized Clock and Timezone for Management

The Network Time Protocol (NTP) is designed to synchronize a network of devices. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two devices to within a millisecond of one another.

Network devices should be programmed to synchronize to a local NTP server in the network. The local NTP server typically references a more accurate clock feed from an outside source. Configuring console messages, logs, and debug output on switches, routers, and other devices in the network to provide time stamps on output allows cross-referencing of events in a network.

To configure NTP in Cisco NX-OS, use the following feature command and specify the NTP server. The example also shows how to enable the correct time zone and daylight savings settings.

```
ntp server 10.4.200.17
clock timezone PST -8 0
clock summer-time PDT 1 Sunday March 02:00 1 Sunday November
02:00 60
!
logging timestamp milliseconds
```

Step 5: Configure Device Resiliency Features

UniDirectional Link Detection (UDLD) is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When UDLD detects a unidirectional link, it disables the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree loops, black holes, and non-deterministic forwarding. In addition, UDLD enables faster link failure detection and quick reconvergence of port trunks, especially with fiber, which can be susceptible to unidirectional failures.

In normal mode, if the link state of the port is determined to be unidirectional, then the port continues to forward traffic normally, but the port is marked as undetermined. The port cycles through the regular spanning tree protocol states and continues to forward traffic. In aggressive mode, the port enters the errdisable state and effectively shuts down. To recover from errdisable, you have to shut down and restart the port by issuing the shut and no shut commands. UDLD does not function any differently for either mode. The same messages are sent and the same messages are expected to be received. The modes only differ in the way that UDLD reacts to a unidirectional link failure.

To enable UDLD, use the following command:

```
feature udld
```

Procedure 3 Core Layer Switch Global Configuration

Step 1: Configure an In-Band Management Interface

The loopback interface is a logical interface that is always reachable as long as the device is powered on and any IP interface is reachable to the network. Because of this capability the loopback address is the best way to manage the switch in-band. Layer 3 process and features are also bound to the Loopback interface to ensure processes resiliency.

The loopback address is commonly a host address with a 32-bit address mask. Allocate the loopback address from the core IP address block .We have included the **ip pim sparse-mode** command that will be explained further in step 3.

```
interface loopback 1
 ip address [ip address]/32
 ip pim sparse-mode
```

Once the IP loopback interface is created, SNMP can be associated with that interface address.

```
snmp-server source-interface trap loopback1
```

Step 2: IP Unicast Routing Configuration

You must enter the feature **eigrp** command to enable this feature. Only after the feature is enabled, do all the EIGRP configuration commands appear.

Cisco NX-OS routing configuration follows an interface-centric model. Instead of adding networks to be advertised via network statements, EIGRP is enabled on a per-interface-basis. Each Layer 3 interface that carries a network that may be advertised via EIGRP requires the ip router eigrp statement. If a network may be advertised, but no adjacency may be formed, the ip passive-interface eigrp command should be used.

In this configuration the only parameter configured under the EIGRP process (router eigrp 100) is the router-ID. The loopback 1 IP address is used for the EIGRP router ID.

```
feature eigrp
router eigrp 100
  router-id [ip address of loopback 1]
interface loopback1
  ip router eigrp 100
```

Step 3: IP Multicast Routing Configuration

To enable Anycast RP operation, the first step is to configure a second loopback interface on each of the core switches. The key is that this second loopback interface will have the same IP address on both core switches and will use a host address mask (32 bits). All routers then point to this common IP address on loopback 2 for the RP in their configuration. We configured the RP address from the core IP address space.

```
interface Loopback2
 ip address 10.4.60.252/32
 ip pim sparse-mode
 ip router eigrp 100
```

The final step for the Anycast RP configuration is to enable Multicast Source Discovery Protocol (MSDP) to run between the two core RP switches. To enable MSDP, you must use unique addresses at each end of the link; therefore, we use the loopback 1 addresses of each core router to configure the MSDP session (Figure 5). Enter the feature msdp command first to enable the feature.

On core switch #1:

```
feature msdp
ip msdp originator-id loopback1
ip msdp peer 10.4.60.253 connect-source loopback1
! The IP address for the listed above is the core switch #2
loopback
```

Figure 5. MSDP Overview

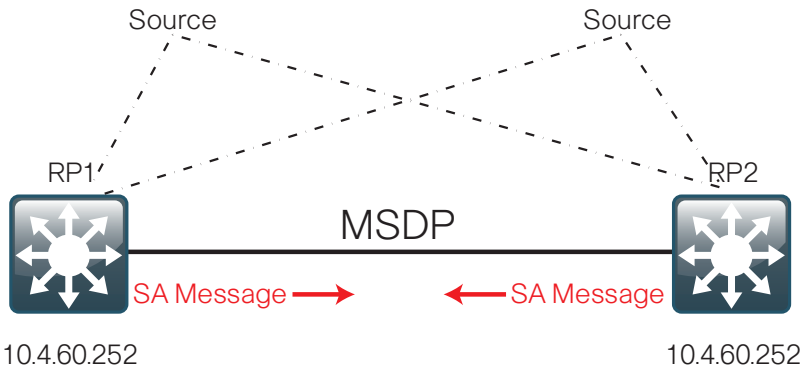

On core switch #2:

```
feature msdp
ip msdp originator-id loopback1
ip msdp peer 10.4.60.254 connect-source loopback1
! The IP address for the listed above is the core switch #1
loopback
```

The MSDP configuration is complete and convergence around a failed RP is now as fast as the unicast routing protocol (EIGRP) convergence. You can see the MSDP protocol session activate later as you enable the routing links between the core switches and the distribution layer blocks that establish Layer 3 connectivity. Use the **show ip msdp summary** command for verification.

Every Layer 3 switch and router must be configured with the address of the IP multicast RP, including the core switches that are serving as the RP. Configure the core switches to point their RP address to the loopback 2 IP address as well as every other remote router and Layer 3 switch. Use the **rp-address** command in conjunction with a group-list to limit the network size that the RP is responsible for. Also note that you need to enter the feature **pim** command to expose all the CLI commands for Protocol Independent Multicast (PIM).

```
feature pim
ip pim rp-address [rp address] group-list [multicast network]/
[mask]
```

All Layer 3 interfaces in the network must be enabled for sparse mode multicast operation.

```
ip pim sparse-mode
```

In the event that you add a core layer to your existing network and the RP is currently configured on a distribution layer, you may want to move the RP to the core. With Anycast RP, you can move the RP to a new location by programming the RP address on the loopback 2 interfaces at the new location, and then enabling and establishing IP multicast and MSDP peering. All remote routers should still point to the same RP address, which simplifies the move and reduces disruption to the IP multicast environment.

Step 4: Quality of Service (QoS)

The core layer QoS configuration honors the QoS values present in the traffic because classification and marking of traffic has already happened at the edge of the network. For traffic not hitting any congestion point when traversing through the core, the original QoS (DSCP) values are preserved and not altered. In case congestion occurs once traffic travels through the core, a reclassification or markdown can occur.

A Cisco Nexus 7000 Series swtich has QoS enabled by default. All interfaces are configured to be trusted, which means that the QoS values present in the IP packets are honored for queuing and scheduling. By default then, no QoS configuration is required in the core.

Procedure 4 Distribution Layer Aggregation Configuration

In this design, links in the core layer are configured as point-to-point Layer 3 routed links or Layer 3 routed EtherChannels.

If you are using the Cisco Catalyst 6500 VSS 1440 system in the distribution layer, we recommend that all peer-connected links are EtherChannel links. EtherChannel to the Catalyst 6500 VSS provides for optimal forwarding as a packet that is received on the switch will be forwarded out a link on that same switch in normal operation versus traversing the VSL link.

Other benefits of EtherChannel to any single physical or logical device are the ease of growing bandwidth without changing the topology and that a single link failure uses EtherChannel recovery versus using ECMP or a routing topology change to reroute the data flows for fastest recovery.

Since the core links are point-to-point routed links, use 30-bit IP address subnets and masks and do not use Switched Virtual Interfaces (SVI).

For core connected devices that do not require EtherChannel, configure routed interfaces with the IP address directly on the physical interface and do not use a switch virtual interface (SVI).

```
interface ethernet[number]
 logging event link-status
 ip address [ip address]/[mask]
 ip pim sparse-mode
 ip router eigrp 100
```

Creating the EtherChannel (sometimes referenced as a port channel) in Cisco NX-OS means adding the channel-group command under the respective channel member interfaces. Note that Cisco NX-OS labels all Ethernet interfaces as interface ethernet irrespective of the actual interface speed. (This style of labeling is different than IOS which uses interface GigabitEthernet, interfaceTenGigabitEthernet and so on.).

While the port-channel interface is automatically created, it does not carry any IP address, routing, and logging configuration yet. The port-channel interface number matches the number specified at the channel-group command. To get a LACP configuration started, make sure you first enter the respective feature command.

```
feature lacp
interface ethernet[interface 1]-[interface 2]
 logging event port link-status
 channel-group [number] mode active

interface port-channel [number]
 logging event port link-status
 ip address [ip address]/[mask]
 ip pim sparse-mode
 ip router eigrp 100
```

Make sure you enter the no shut command to bring up the port-channel interface and the channel member interfaces to an operational state.

Summary

The Cisco Cisco Nexus 7000 Series is a key platform for the data center and campus core deployments.

Built around the flexible and scalable Cisco NX-OS operating system, the Cisco Nexus 7000 Series combines hardware and software to deliver unprecedented high availability. Full hardware redundancy and sophisticated mechanisms for recovering from failure conditions make the Cisco Nexus 7000 Series a very robust, zero-downtime platform for core networks for both scenarios of collapsed data center and campus core or separate cores.

The VDC technology provides a variety of benefits ranging from simplified operations, hardware and software resources separation, virtualization, and consolidation. While relevant to all the core architectures, VDC adds great value in a collapsed core environment. In this scenario, each VDC can virtualize a core device, presenting the physical switch as multiple logical devices with their own unique set of VLANs, IP routing tables, VRFs, and physical interfaces.

Furthermore, scalable and efficient multicast capabilities, along with a flexible Quality of Service model, enable the Cisco Nexus 7000 Series to embrace the Cisco Medianet solution.

For Additional Information

Cisco NX-OS Product Page: http://www.cisco.com/go/nxos

Cisco NX-OS / IOS Comparison: http://docwiki.cisco.com/wiki/Cisco_Nexus_7000_NX-OS/IOS_Comparison_Tech_Notes

Cisco Nexus 7000 Series Product Page: http://www.cisco.com/go/nexus7000

Cisco Nexus 7000 Series Product Documentation: http://www.cisco.com/en/US/products/ps9402/tsd_products_support_series_home.html

Medianet on Cisco.com: http://www.cisco.com/go/medianet

Borderless Networks: http://www.cisco.com/go/borderless

Data Center Design and Configuration documents:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/DC-3_0_IPInfra.html

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/nx_7000_dc.html

Appendix A:
Large Agencies LAN Deployment Product List

Functional Area	Product	Part Numbers	Software Version
Core Layer	Nexus 7000	N7K-C7010 Nexus 7000 C7010 (10 Slot) Chassis N7K-SUP1 Nexus 7000 Supervisor module-1X N7K-M148GS-11 Nexus 7000 48-port GigE Mod (SFP) N7K-M132XP-12 Nexus 7000 32-port 10 GigE Module N7K-M108X2-12L Nexus 7000 8-port 10 GigE Ethernet Module (X2)	4.2(4)

Appendix B:
SBA for Large Agencies Document System

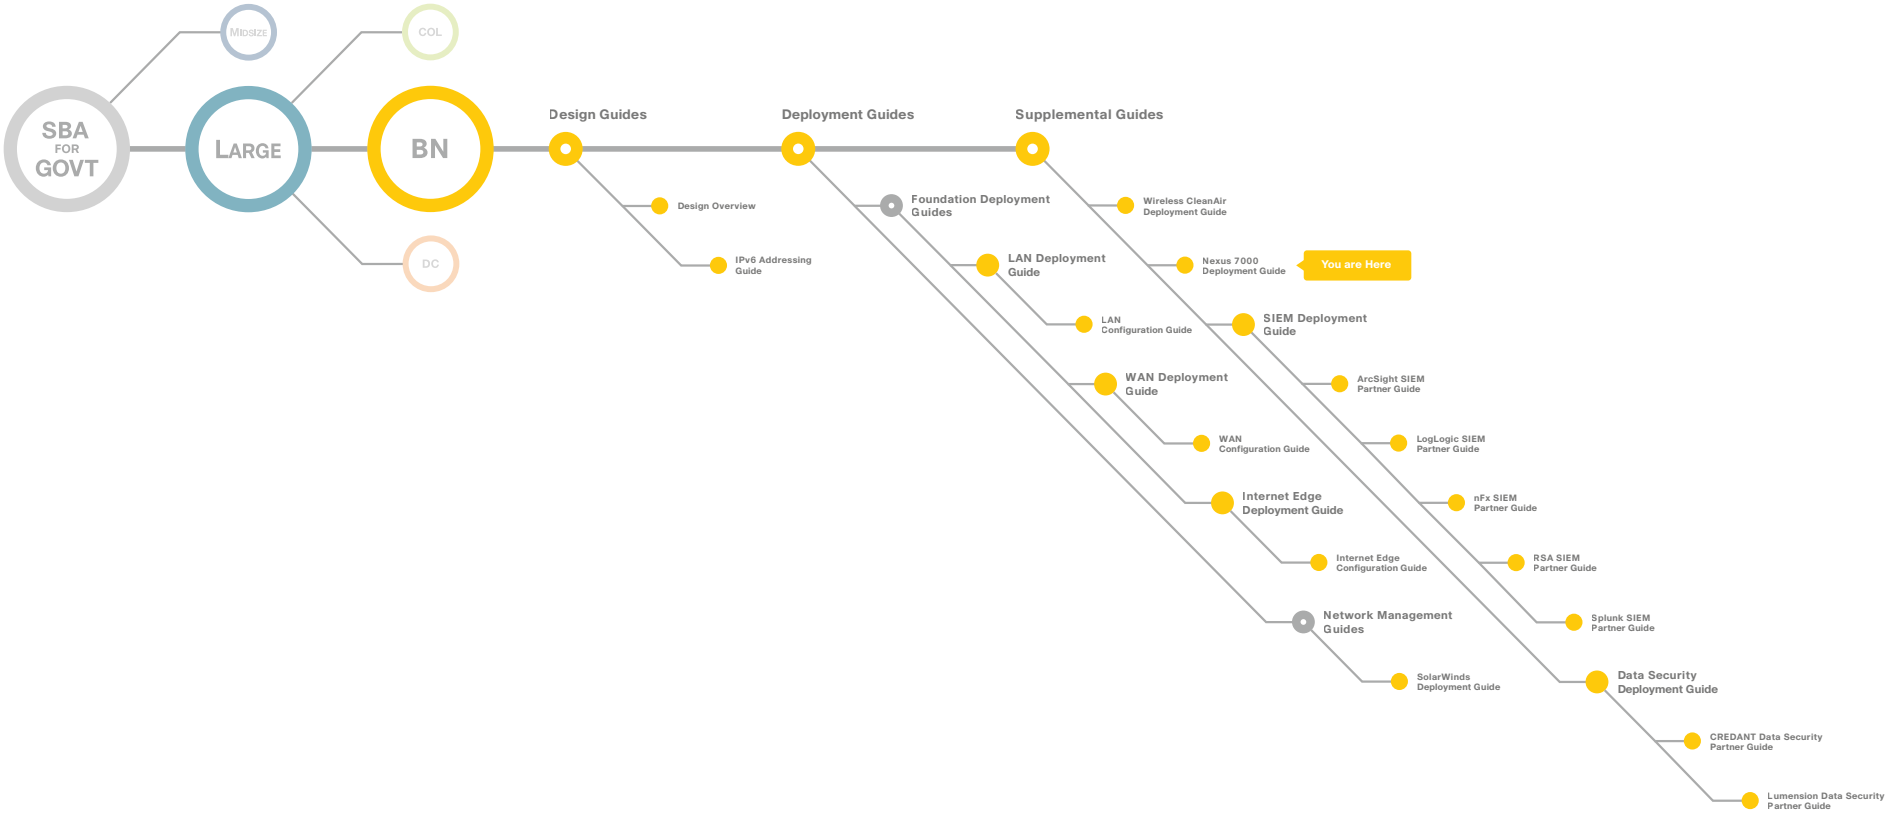

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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.**

C07-640796-00 02/11