Chapter 2:

Cisco NX-OS Software Command-Line Interface Primer

Frequently used configuration commands for Layer 2 and Layer 3 topologies and their differences from Cisco Catalyst IOS Software

Contents

Introduction

This chapter is part of a larger document that provides guidelines for designing and deploying access and aggregation layers in the data center using Cisco Nexus and Catalyst products.

This document covers the main additions and differences in Cisco® NX-OS Software compared with Cisco Catalyst® IOS® Software. Use this document as a quick reference to the main commands for building and designing a data center Layer 2 and Layer 3 infrastructure with Cisco Nexus® products. This document is not intended as a replacement for the configuration guides available from the Cisco public website.

In addition, because features like virtual domain contexts, role-based access, checkpoint and rollback, fabric extenders, or port profiles are Cisco NX-OS Software-specific, they are explained here so that you can become familiar with the configuration-specific aspects.

For additional command-line interface (CLI) configuration information, visit these sites:

- For information on Cisco Nexus 7000 Series Switches: http://www.cisco.com/en/US/partner/products/ps9402/tsd_products_support_configure.html
- For information on Cisco Nexus 5000 Series Switches: http://www.cisco.com/en/US/partner/products/ps9670/tsd_products_support_series_home.html
- For information on Cisco Nexus 2000 Series Fabric Extenders: http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus2000/sw/configuration/guide/rel_4_0_1a/NX2000CLIConfig.html
- For information on Nexus 1000V Series Switches: http://www.cisco.com/en/US/partner/products/ps9902/tsd_products_support_series_home.html

Summary of Cisco NX-OS CLI Commands Used in Data Center Design

Most data center deployments use both Cisco NX-OS and Cisco Catalyst IOS Software. Although the underlying operating system characteristics differ, from a configuration point of view, the two operating systems are very similar. This section highlights how some key portions of the data center configuration can be performed with either operating system.

For more information about the two operating systems, visit: http://www.cisco.com/en/US/products/sw/iosswrel/products_ios_cisco_ios_software_category_home.html

Virtual Domain Context Concept

Cisco Nexus 7000 Series Switches running Cisco NX-OS Software have introduced the capability to divide a single physical switch into as many as four virtual switches, referred to as virtual device contexts (VDCs). Each VDC operates similarly to a standalone switch with a distinct configuration file, complement of physical ports, and separate instances of necessary control plane protocols such as routing protocols and spanning tree. This means that when you are operating a Cisco Nexus 7000 Series Switch, it is relevant to know which VDC you are operating.

There are four VDCs as of this writing, of which only one, VDC 1 (also called the default) has higher privileges than the others. VCD 1 can do the following:

- Create and destroy other VDCs.
- Allocate resources to VDCs (the main resources that you will allocate are physical interfaces).
- Set the dual-supervisor redundancy policy.
- Maintain software.
- Reboot the system.

In order to switch from the default VDC to the others, you can use the following command:

```
NXOS : #switchto vdc <name>
NXOS : #switchback
```

Notice that a user who starts a session with a non-default VDC (VDC > 1) cannot hop to other VDCs.

To allocate interfaces to a VDC, use the following commands:

```
NXOS : (config)# vdc <name>
NXOS : (config-vdc)# allocate interface ethernet <port/number>
NXOS : (config)# show vdc membership
```

From VDC 1, VDCs 2–4 can be restarted from the default VDC with the command:

```
NXOS: (config)# vdc <name> restart
```

This command shuts down local services running on that specific VDC (spanning tree, Hot Standby Router Protocol [HSRP], routing, and so on) and then brings them back up with their last saved configuration. Within the context of the VDC itself, the administrator can achieve the same result by issuing the **reload** command, which will reload the VDC, leaving the other VDCs unaffected (as long as this is a non-default VDC—that is, VDC > 1).

Currently, the VDC concept is implemented only on the Cisco Nexus 7000 Series.

The concept of VDC is most relevant at the aggregation layer of data center designs, while at the access layer the use of VLANs provides segmentation at the data plane layer. The reason for using VDCs at the aggregation layer is that if an attacker manages to get hold of the control plane of the router (through the default gateway), the attacker will not be able to hop into adjacent VDCs.

At the access layer, this threat does not exist, provided that the management of the access layer device is out of band—that is, it uses, for example, the mgmt0 interface of a Cisco Nexus 5000 Series that may be connected to the default VDC of the Cisco Nexus 7000 Series. This is not to say that there will never be VDCs as an option for access layer devices, but of all devices, the ones that benefit the most are the ones at the aggregation layer.

Role-Based Access

Role-based access lets you specify the actions a user can perform on a given Cisco Nexus system. In the case of a Cisco Nexus 7000 Series, there are by default four different user roles:

- **network-admin:** Read-write privileges for the default VDC, so this user has higher privileges than all others. The network-admin can jump from VDC 1 into any of the other VDCs, as well as create and destroy VDCs.
- **network-operator:** Read privileges for the default VDC.
- **vdc-admin:** Read-write privileges for a VDC. This role exists only within a given VDC.
- **vdc-operator:** Read access to a VDC.

The difference between vdc-admin and network-admin is not significant on devices that do not implement VDCs (or in other words, on devices that implement only VDC1): the Cisco Nexus 5000 Series and Cisco Nexus1000V Series Switches.

For more information about role-based access on the Cisco Nexus 7000 Series, visit: http://www.cisco.com/en/US/partner/docs/switches/datacenter/sw/4_1/nx-os/security/configuration/guide/sec_rbac.html

For more information about role-based access on the Cisco Nexus 5000 Series, visit: http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/sec_rbac.html

Roles can be propagated across a Cisco Nexus infrastructure by using the Cisco Fabric Service Protocol over IP (CFSoIP) to define role-based access as a client to the CFS protocol.

The Cisco Nexus 5000 Series is a device that offers both SAN and LAN capabilities. To help manage access to the same device by a SAN administrator and a LAN administrator, you can use role-based access to define the actions that either administrator can perform.

Checkpoint and Rollback

Currently the only Cisco Nexus device that implements configuration checkpoint and rollback is the Cisco Nexus 7000 Series.

More information about checkpoint and rollback can be found here:
http://www.cisco.com/en/US/partner/docs/switches/datacenter/sw/4_0/nx-os/system_management/configuration/guide/sm_rollback.html

Port Profiles: Concept and Configuration

The port profiles concept applies to the Cisco Nexus 1000V Series and the Cisco Nexus 7000 Series.

A port profile is a configuration construct that simplifies the concurrent configuration of multiple switch interfaces. Port profiles allow repetitive commands to be typed only once and automatically applied to all interfaces that are members of the same profile.

Port profiles can be of different kinds, and the two that are of interest for the purpose of this solution guide are **type Ethernet** for regular interfaces and **type interface-vlan** for switch virtual interfaces (SVIs):

```
port-profile [type <ethernet | interface-vlan>] <name>
```

The port-profile-to-interface association is configured as follows:

```
interface etha/b
          inherit port-profile <name>
```

If the interface a/b is configured with switch-port commands and if a conflict arises, the interface-specific configurations take precedence over the port profile configuration.

Nexus1000V-Specific Commands

On the Cisco Nexus 1000V Series, port profiles allow some configurations that are very specific to virtualized server environments. These configurations include:

- Specifying the string name for the VMware port group associated with the port profile (**vmware port-group** command)
- Specifying whether a port profile is used for virtual Ethernet interfaces or for uplink purposes through the server network interface cards (NICs) (**capability uplink** command)
- Pushing the configuration to VMware vCenter (**state enable** command)
- Defining the VLANs that are used for communication between the Virtual Supervisor Module and Virtual Ethernet Module (**system vlans** command)

An example of uplink port profile configuration is:

```
n1000v(config)# port-profile SysProfile
n1000v(config-port-prof)# capability uplink
n1000v(config-port-prof)# vmware port-group
n1000v(config-port-prof)# switchport mode trunk
n1000v(config-port-prof)# switchport trunk allowed vlan 101-110
n1000v(config-port-prof)# no shutdown
n1000v(config-port-prof)# system vlan 103,104
n1000v(config-port-prof)# state enabled
```

For more information, visit:
http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus1000/sw/4_0/port_profile/configuration/guide/n1000v_port_profile.html

Dual-Supervisor Configuration

When deploying a system made of two supervisors, one supervisor is normally active and the other is at standby, at least from a management plane point of view.

Cisco NX-OS Software offers the following commands to check the supervisor redundancy status:

```
NXOS : #show system redundancy status
```

In Cisco NX-OS, you can trigger a supervisor switch-over by issuing the command:

```
NXOS : #system switchover
```

In Cisco Catalyst IOS Software, when two supervisors are present in the same chassis (which is currently not applicable to virtual switching system (VSS)-mode deployments), redundant supervisor configurations are handled through the following commands:

```
Catalyst IOS : (config)# redundancy
Catalyst IOS : (config)#  mode sso
Catalyst IOS : (config)#  auto-sync running-config
```

Cisco Catalyst IOS Software offers the following commands to check the supervisor redundancy status:

```
Catalyst IOS : (config)# show redundancy state
```

In Cisco Catalyst IOS Software, you can trigger a supervisor switchover using the following command:

Catalyst IOS : (config)# redundancy force-switchover

Upgrade Procedures

This document is not an operational guide, so it doesn't cover in detail software upgrades and troubleshooting. On the other hand, for a proper design it is important to have a basic idea of which devices support In-Service Software Upgrade (ISSU) and how it works. It is also important to understand upgrades in a virtualized server environment.

For Cisco Nexus 7000 Series software upgrade procedures:
http://www.cisco.com/en/US/partner/products/ps9402/prod_installation_guides_list.html

For Cisco Nexus 1000V Series software upgrades procedures:
http://www.cisco.com/en/US/partner/products/ps9902/prod_installation_guides_list.html

ISSU

The ISSU concept is applicable to the Cisco Nexus 7000 Series platform. In a Cisco NX-OS system, you can perform in-service software upgrades as described in the documentation at this link:
http://www.cisco.com/en/US/partner/docs/switches/datacenter/sw/4_1/nx-os/upgrade/guide/nx-os_upgrade.html

In a Cisco Catalyst IOS system, you can upgrade the system in an ISSU fashion using the Enhanced Fast Supervisor Upgrade (eFSU) as follows:
http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/issu_efsu.html#wp1073900

If the system is configured for VSS, the upgrade procedure follows the guidelines described in this document:
http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/vss.html#wp1169328

Configuration of Fabric Extender Ports

The fabric extender concept applies to the Cisco Nexus 5000 Series family of products. Fabric extenders are not an independent manageable entity; the Cisco Nexus 5000 Series manages them through in-band connectivity. Currently, the Cisco Nexus 2148T Fabric Extender implements the fabric extender architecture.

Each fabric extender is identified by a number in the range 100–199. You configure the number using the Cisco Nexus 5000 Series CLI. All you need to do is configure the 10 Gigabit Ethernet ports (any 10 Gigabit Ethernet port of the Cisco Nexus 5000 Series) that are going to be connected to a fabric extender module. You do so by specifying the fabric extender number under the interface configuration mode, as follows:

```
nexus5k(config)#interface ethernet <a/b>
nexus5k(config-if)#switchport mode fex-fabric
nexus5k(config-if)#fex associate <fex-id>
```

The preceding configuration indicates that when the fabric extender module is wired to port 1/1 through a twinax or optical cable, it assumes the identity number indicated in the **fex associate** statement.

The other fundamental configuration that is fabric-extender-specific is **pinning**. Pinning is covered in the chapter on Cisco Nexus 5000 Series design, Chapter 6. The configuration commands are as follows:

```
fex <fex-id>
  pinning max-links <1-4>
  description <description>
```

Gigabit and 10 Gigabit Interfaces Configuration

In Cisco NX-OS Software, interfaces are named as EthernetA/B regardless of whether their speed is Gigabit Ethernet or 10 Gigabit Ethernet.

```
NXOS: interface Ethernet1/1
Catalyst IOS : interface GigabitEthernet1/1 or interface TenGigabitEthernet1/1
```

If you are configuring a 10 Gigabit Ethernet line card on the Cisco Nexus 7000 Series, you'll have to consider that 10 Gigabit Ethernet ports can be operated in **dedicated** or in **shared** mode. This means that you need to configure the operation mode under the Ethernet interface configuration as follows:

```
NXOS: (config-if)#rate-mode {dedicated | shared }
```

Link Layer Encryption Configuration

The Cisco Nexus 7000 Series supports link layer encryption (IEEE 802.1ae), referred to as Cisco Trusted Security: http://www.cisco.com/en/US/partner/docs/switches/datacenter/sw/4_1/nx-os/security/configuration/guide/sec_trustsec.html

To encrypt on a link between two Cisco Nexus 7000 Series Switches, use the following configuration:

```
interface Etherneta/b
  switchport
  switchport mode trunk
  switchport trunk allowed vlan <vlan list>
  cts manual
    sap pmk <key>
  channel-group <number> mode <on | active | passive >
no shutdown
```

If there is no key in the remote site, the link won't go up. **As you can see in the below CLI capture, the remote interface won't go up because the authentication failed.**

```
tc-nexus7k01# show int etha/b
Ethernet7/25 is down (Authorization pending)
  Hardware: 10000 Ethernet, address: 001b.54c1.65e8 (bia 001b.54c1.65e8)
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

If you add the key on a port that is already up, you'll need to do a "shut-no shut" operation to enable the SAP configuration:

```
tc-nexus7k02-VDC2# show cts interface ethernet a/b
CTS Information for Interface Ethernet7/25:
    CTS is enabled, mode:   CTS_MODE_MANUAL
    IFC state:              CTS_IFC_ST_CTS_OPEN_STATE
    Authentication Status:  CTS_AUTHC_SKIPPED_CONFIG
      Peer Identity:
      Peer is:              Not CTS Capable
      802.1X role:          CTS_ROLE_UNKNOWN
      Last Re-Authentication:
    Authorization Status:   CTS_AUTHZ_SKIPPED_CONFIG
      PEER SGT:             0
      Peer SGT assignment:  Not Trusted
    SAP Status:             CTS_SAP_SUCCESS
      Configured pairwise ciphers: GCM_ENCRYPT
      Replay protection: Enabled
      Replay protection mode: Strict
      Selected cipher: GCM_ENCRYPT
      Current receive SPI: sci:1b54c1a7940000 an:0
      Current transmit SPI: sci:1b54c165e80000 an:0
```

Spanning Tree Configuration

Most configuration commands are identical between Cisco IOS and Cisco NX-OS Software. The main difference is how to configure **portfast** and **bridge assurance**.

Portfast is configured as follows:

```
NXOS: (config-if)#spanning-tree port type edge  [trunk]
```

Bridge assurance is enabled globally by default and locally on the interface when the option **network** is selected. If the option selected is **normal**, bridge assurance doesn't run on the link. On a per Layer 2 link level, the normal option is the default:

```
Nexus(config)# spanning-tree bridge assurance
NXOS: (config-if)#spanning-tree port type { normal | network }
```

Link Aggregation Control Protocol Configuration

Port channel configurations are, for the most part, identical in Cisco NX-OS and Cisco Catalyst IOS Software. One difference is that in Cisco NX-OS, you must enable the Link Aggregation Control Protocol service, as follows:

```
NXOS: (config)#feature lacp
```

The configuration of the hashing algorithm follows the same syntax in Cisco NX-OS, but on the Cisco Nexus 7000 Series, it can be specified per module:

```
NXOS : (config)# port-channel load-balance ethernet { various options } [module]
```

The **show** command is slightly different between Cisco NX-OS and Cisco IOS Software:

```
NXOS : # show port-channel summary
Catalyst IOS: # show etherchannel summary
```

Commands Specific to the Cisco Nexus 1000V Series

On the Cisco Nexus 1000V Series Switches, the port channel configuration varies slightly from the regular port channel configuration, mostly because dual-homed VMware ESX servers require that the channel be split between two upstream switches.

The syntax is as follows:

```
channel-group auto mode {on | active | passive} [subgroup cdp]
```

If the network adapters on the VMware ESX host are connected to a single upstream switch, the port channel configuration is very simple:

```
n1000v(config)# port-profile uplink-profile
n1000v(config-port-prof)# channel-group auto
n1000v(config-port-prof)# capability uplink
```

If the network adapters on the ESX host are connected to two upstream switches that are virtual PortChanel (vPC)- or VSS-capable, the port channel configuration is a regular port-channel configuration as follows:

```
n1000v(config)# port-profile uplink-profile
n1000v(config-port-prof)# channel-group auto
n1000v(config-port-prof)# capability uplink
```

If the network adapters on the ESX host are connected to two upstream switches that are neither vPC- nor VSS-capable, and if these adapters can use Cisco Discovery Protocol, the configuration is as follows:

```
n1000v(config)# port-profile uplink-profile
n1000v(config-port-prof)# channel-group auto subgroup cdp
n1000v(config-port-prof)# capability uplink
```

If the network adapters on the ESX host are connected to two upstream switches that are not vPC or VSS-capable and do not use Cisco Discovery Protocol, the configuration is as follows:

```
n1000v(config)# port-profile uplink-profile
n1000v(config-port-prof)# channel-group auto subgroup cdp
n1000v(config-port-prof)# capability uplink

n1000v(config)# interface Ethernet <ESX ports going to switch 1>
n1000v(config-if-range)#  sub-group-id 0
n1000v(config-if-range)# interface Ethernet <ESX ports going to switch 2>
n1000v(config-if-range)#  sub-group-id 1
```

Notice that the **cdp** keyword is used even when the upstream devices don't run Cisco Discovery Protocol. The manual configuration of the **sub-group-id** ensures that the links are correctly split into two different port channels, one per upstream switch (group 0 to switch 1, group 1 to switch 2).

Virtual PortChannel Commands

The procedure for configuring virtual PortChannels on the Cisco Nexus 7000 or 5000 Series is summarized in this section.

1. Execute this command to enable the feature:

```
agg(config)# feature vpc
```

2. Create a vPC domain and give it a priority, as follows:

```
agg(config)# vpc domain 1
agg(config-vpc-domain)# role priority 100
```

3. Create a peer link for the communication between the vPC peers:

```
agg(config)# interface port-channel10
agg(config-if)#  vpc peer-link
agg(config-if)#  switchport trunk allowed vlan <all access vlans>
```

4. Create an out-of-band communication path to verify the health of the vPC peer in case the peer link is cut:

```
vpc domain 1
peer-keepalive destination 192.168.1.2 source 192.168.1.1 vrf vpc-keepalive
```

5. Make Layer 2 ports members of the vPC:

```
agg(config)#interface ethernet2/9
agg(config-if)# channel-group 51 mode <on | active | passive>

agg(config)#interface Port-channel 51
agg(config-if)# switchport
agg(config-if)# vpc 51
```

Private VLANs

In order to use private VLANs in Cisco NX-OS, you need to enable the feature:

```
NXOS : (config)#feature private-vlan
```

Assuming that VLAN 50 is the primary VLAN and that VLAN 51 is the secondary VLAN, and that this VLAN is configured as isolated, the configuration in Cisco NX-OS Software is as follows.

```
vlan 50
  private-vlan primary
  private-vlan association 51
vlan 51
private-vlan isolated
```

The SVI configuration for VLAN 50, where VLAN 51 needs to be remapped, follows this syntax:

```
interface Vlan50
  no shutdown
private-vlan mapping 51
```

For trunks carrying primary and secondary VLANs where you desire no remapping, the configuration is the usual trunk configuration:

```
interface Ethernet2/9
  description tc-nexus5k01 - Eth2/1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan …50-51…
no shutdown
```

Remember the difference between association and mapping:

- The keyword **association** is used in conjunction with isolated port configuration (that is, it tells the port, typically an access port, that the primary VLAN needs to be forwarded, or remapped to the secondary, and the secondary VLAN is to be kept as is).
- The keyword **mapping** is used in conjunction with the promiscuous ports configuration and it tells the switch that the secondary VLAN should be translated into the primary VLAN.

Here is an example of configuring an isolated access port:

```
nexus5000(config-if)#
switchport mode private-vlan host
switchport private-vlan association trunk 50 51
```

Here is an example of configuring an isolated trunk port:

```
nexus5000(config-if)#
switchport mode private-vlan trunk secondary
switchport private-vlan association trunk 50 51
```

Here is an example of configuring an uplink port that operates as a promiscuous trunk:

```
nexus5000(config-if)#
switchport mode private-vlan trunk promiscuous
switchport private-vlan mapping trunk 50 51
```

If the trunk port of a promiscuous trunk or if isolated trunks carry VLANs other than the private VLANs, you need to use a different command to make sure they are correctly forwarded (remember that the port is not in the switch port mode trunk, but in the switch port mode **private-vlan** trunk):

```
switchport private-vlan trunk allowed vlan <list of regular non-private VLANs>
```

Port Security

In a data center switching environment hosting virtual servers, the only place where it makes sense to configure port security is on the Cisco Nexus 1000V Series. In all other deployments, where the Cisco Nexus 1000V Series is not deployed, a moving virtual machine would otherwise trigger an alarm and the port would be error-disabled because the MAC move with no port link down would be perceived as a security violation. In addition to this, limiting the number of MAC addresses that a physical port is allowed to forward is impractical when the server is virtualized because new virtual machines may appear unexpectedly (as a result of a VMotion migration or of provisioning).

For this reason, we include the syntax for port security only on the Cisco Nexus 1000V Series Switch:

```
switchport port-security maximum max-val
 switchport port-security violation {shutdown | protect | restrict}
 switchport port-security aging time time-value
 switchport port-security aging type {absolute | inactivity}
 switchport port-security mac-address sticky
```

Layer 3 Interface VLAN Configuration

Interface VLAN configurations are typically identical in Cisco NX-OS and Cisco Catalyst IOS Software. One difference is that in Cisco NX-OS, you must enable the interface-VLAN service, as follows:

```
NXOS: (config)#feature interface-vlan
NXOS: (config)#interface vlan <number>
```

HSRP Configuration

Hot Standby Router Protocol (HSRP) configurations are generally identical in Cisco NX-OS and Cisco Catalyst IOS Software. One difference is that in Cisco NX-OS, you must enable the HSRP service, as follows:

```
NXOS: (config)#feature hsrp
```

Additionally, the HSRP configuration under the interface is slightly different than it is in Cisco Catalyst IOS Software:

```
NXOS: (config-if)# hsrp <group>
NXOS: (config-if-hsrp)#ip <IP address>
NXOS: (config-if-hsrp)#priority <priority>
NXOS: (config-if-hsrp)#[no] preempt

Catalyst IOS: (config-if)# standby <group> ip
Catalyst IOS: (config-if)# standby <group> priority <number>
Catalyst IOS: (config-if)# standby <group> preempt
```

The **show** commands differ as follows:

```
NXOS: #show hsrp brief
Catalyst IOS: #show standby brief
```

Routing Configuration

In Cisco NX-OS Software, the **network** command is not utilized; it has been replaced by the **ip router ospf <area>** command under the interface VLAN or the Layer 3 interface configuration, as follows:

```
NXOS: (config)#feature ospf

NXOS: config)# interface <VLAN>
NXOS: config-if)#ip router ospf <area>
```

Additionally, if you want to make an interface passive, enter the following:

```
NXOS: config-if)# ip ospf passive-interface

NXOS: router ospf 1
NXOS: auto-cost reference-bandwidth 1000000
```

Virtual Route Forwarding

Both Cisco Nexus and Cisco Catalyst platforms support the concept of virtual route forwarding (VRF).

Cisco Nexus platforms utilize VRF by default for management purposes. Two VRFs are defined by default:

- Default
- Management

In order to configure a VRF in Cisco NX-OS Software, do the following:

```
NXOS : (config)#vrf context <VRF NAME>
NXOS : (config)#interface Ethernet1/1
NXOS : (config-if)# vrf member <VRF NAME>
```

In order to configure a VRF in Cisco Catalyst IOS Software, do the following:

```
Catalyst IOS : (config)#ip vrf <VRF NAME>
Catalyst IOS : (config)#interface Ethernet1/1
Catalyst IOS : (config-if)# ip vrf forwarding <VRF NAME>
```

Management Port

Cisco NX-OS platforms offer a management interface. This port is clearly indicated on the front plate of Cisco Nexus devices and is configured through the following command:

```
NXOS: (config)#interface mgmt0
NXOS: (config-if)# ip address …
```

Routing for management traffic is configured starting from the **management vrf:**

```
NXOS: (config)#vrf context management
NXOS: (config-if)# ip route …
```

Quality of Service

The Cisco Nexus switches follow the Cisco Modular Quality of Service command-line interface (MQC). Quality-of-service (QoS) features and capabilities are hardware-specific, and as a result, several options are specific to each of the Cisco Nexus platforms.

MQC consists of three configuration steps:

1. Define match criteria with a class map.
2. Associate an action for each defined class with a policy map.
3. Apply the policy to the entire system or to an interface service policy.

Classification

A simple configuration example of **classification** is:

```
Ip access-list ACL-A
  Permit ip any 1.1.1.0/24

class-map type qos qosgroupA
  match access-group ACL-A

policy-type qos policy-classify
  class type qos qosgroupA
    set qos-group 3
  class type qos qosgroupB
    set qos-group 4
  […]
```

Classification can use any of the following options and is supported on the Cisco Nexus 1000V, 5000, and 7000 Series:

```
ACL-based (SMAC/DMAC, IP SA/DA, Protocol, L4 ports, L4 protocol fields)
CoS, IP prec, DSCP
Internal QoS values (e.g. QoS-group)
Protocol type (non-IP packets)
```

Marking

Marking on the Cisco Nexus products is configured with **network qos policy**, as follows:

```
class-map type network-qos class-app-1
  match qos-group 3
class-map type network-qos class-app-2
  match qos-group 4

policy-map type network-qos marking-apps
   class-map type qos class-app-1
     set cos 3
   class-map type qos class-app-1
     set cos 4
```

Marking on the Cisco Nexus 7000 and 1000V Series can rewrite class of service (CoS), IP precedence, or differentiated services code point (DSCP) values. Marking on the Cisco Nexus 5000 Series can rewrite the CoS.

Marking can define an internal value for further processing such as the **qos-group** (which was used in the example about classification).

The Cisco Nexus 5000 Series, a unified-fabric-capable device, also allows the following:

- Defining the maximum transmission unit (MTU) for a given class
- Specifying whether the class is drop-no drop

Queuing

On Cisco Nexus products, quality of service is configured with a **queue qos policy**. Queuing depends strictly on the hardware implementation. The Cisco Nexus 5000 Series follows the Enhanced Transmission draft standard IEEE 802.1Qaz, so the queuing specifies the bandwidth percentage.

An example of queuing configuration for the Cisco Nexus 5000 is as follows:

```
class-map type queueing class-app-1
  match qos-group 3
class-map type queuing class-app-2
  match qos-group 4

policy-map type queuing queue-apps
   class-map type queueing class-app-1
     bandwidth percent 30
   class-map type queueing class-app-1
     bandwidth percent 40
```

The possible queuing configuration parameters on the Cisco Nexus 7000 Series are:

```
bandwidth (weight), queue-limit, priority, random-detect, shape
```

The possible queuing configuration parameters on the Cisco Nexus 5000 Series are:

```
bandwidth percent, queue-limit, priority
```

Applying Policies

For the policy to apply to the traffic, it needs to be applied either at the system level or to an interface, as follows:

```
N5k(config)#system qos
N5k(config-sys-qos)#service-policy type qos input policy-classify
N5k(config-sys-qos)#service-policy type queueing queue-apps
```

Multicast

For information about multicast, see the following:

- Cisco Nexus 7000 Series Multicast configuration guide: http://www.cisco.com/en/US/partner/docs/switches/datacenter/sw/4_1/nx-os/multicast/configuration/guide/multicast_cli.html
- Cisco Nexus 5000 Series Multicast configuration guide: http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli_rel_4_0_1a/IGMPSnooping.html
- Cisco Nexus 1000V Series Multicast configuration guide: http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus1000/sw/4_0/layer2/configuration/guide/l2_5igmp_snoop.html

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.

Printed in USA C07-572833-00 01/10