IBM System Networking

Distributed Switch 5000V

Version 1.1, VMware Edition

User Guide

First Edition (June 2013)

Contents

Preface

This *User Guide* describes how to configure and use the IBM System Networking Distributed Switch 5000V (5000V) version 1.1 software to provide virtual network switching within a VMware-enhanced datacenter.

Who Should Use This Guide

This guide is intended for network installers and administrators engaged in configuring and maintaining a complex network. The administrator should be familiar with general Ethernet concepts and Layer 2 switching. They should also be familiar with the required VMware vCenter, vSphere, and ESX products and virtualization concepts.

What You'll Find in This Guide

This guide will help you plan, implement, and administer 5000V software. Where possible, each section provides feature overviews, usage examples, and configuration instructions. The following material is included:

Part 1: Getting Started

This material is intended to help those new to 5000V products with the basics of switch management. This part includes the following chapters:

- Chapter 1, "IBM DS 5000V Introduction," provides a conceptual overview of the 5000V solution.
- Chapter 2, "IBM DS 5000V Installation," describes how to install the IBM components of the 5000V solution. This chapter covers prerequisites and the installation and initial configuration of the 5000V controller and host modules.
- Chapter 3, "Host, VM and Port Management," describes how to connect hosts, virtual machines, and ports to the 5000V.

Part 2: Switch Features

- Chapter 4, "Securing Administration," describes methods for using Secure Shell for administration connections, and configuring end-user access control.
- Chapter 5, "VLANs," describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs. This chapter also describes Protocol-based VLANs, and Private VLANs.
- Chapter 6, "Authentication & Authorization Protocols," describes different secure administration for remote administrators. This includes using Remote Authentication Dial-in User Service (RADIUS), as well as TACACS+ and LDAP.
- Chapter 7, "Access Control Lists," describes how to use filters to permit or deny specific types of traffic, based on a variety of source, destination, and packet attributes.
- Chapter 8, "Quality of Service," discusses Quality of Service (QoS) features, including IP filtering using Access Control Lists (ACLs), Differentiated Services, and IEEE 802.1p priority values.
- Chapter 9, "Edge Virtual Bridging," describes using 802.1Qbg/Edge Virtual Bridging (EVB) to simplify virtual network resource management.

- Chapter 10, “Simple Network Management Protocol,” describes how to configure the switch for management through an SNMP client.
- Chapter 11, “sFlow, described how to use the embedded sFlow agent for sampling network traffic and providing continuous monitoring information to a central sFlow analyzer.
- Chapter 12, “Packet Monitoring,” discusses tools how copy selected port traffic to a monitor port for network analysis.

Part 3: CLI Reference

This section lists each command, together with the complete syntax and a functional description, from the built-in Command-Line Interface (CLI).

- Chapter 13, “CLI Basics,” describes how to connect to the switch and access the information and configuration commands. This chapter provides an overview of the command syntax, including command modes, global commands, and shortcuts.
- Chapter 14, “Information Commands,” shows how to view switch configuration parameters.
- Chapter 15, “Statistics Commands,” shows how to view switch performance statistics.
- Chapter 16, “Configuration Commands,” shows how to configure switch system parameters, ports, VLANs, Spanning Tree Protocol, SNMP, Port Mirroring, IP Routing, and more.
- Chapter 17, “Operations Commands,” shows how to use commands which affect switch performance immediately, but do not alter permanent switch configurations (such as temporarily disabling ports). The commands describe how to activate or deactivate optional software features.
- Chapter 18, “Boot Options,” describes the use of the primary and alternate switch images, how to load a new software image, and how to reset the software to factory defaults.
- Chapter 19, “Maintenance Commands,” shows how to generate and access a dump of critical switch state information, how to clear it, and how to clear part or all of the forwarding database.

Part 4: Appendices

- Appendix A, “Getting Help & Technical Assistance,” describes how to obtain product support.
- Appendix B, “Notices.”

Typographic Conventions

The following table describes the typographic styles used in this book.

Table 1. Typographic Conventions

Typeface or Symbol	Meaning	Example
ABC123	This italicized body type shows book titles, special terms, or words to be emphasized.	Read your *User's Guide* thoroughly.
`ABC123`	This plain, fixed-width type is used for names of commands, files, and directories used within the body of the text.	View the `readme.txt` file.
	It also depicts on-screen computer output and prompts.	`host#`
`ABC123`	This bold, fixed-width type appears in command examples. It depicts text that must be typed in exactly as shown.	`host#` **`show sys-info`**
< >	Angled brackets appear in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.	If the command syntax is: **`ping`** *<IP address>* You might enter: **`ping 192.32.10.12`**
`[ ]`	Square brackets depict optional elements within commands. These can be used or excluded as the situation demands. Do not type the brackets.	`host#` **`ls`** [**`-a`**]

{**`A`**|**`B`**}	Curled braces and vertical bars are used in command examples where there are multiple choices. Select only one of the listed options. Do not type the braces or bars.	If the command syntax is:
		`set` {**`left`**|**`right`**}
		You might enter:
		`set left`
		Or:
		`set right`

AaBbCc123	This bold type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the **Save** button.
A > B	This bold type with an angled right-bracket indicates nested menu items in a graphical interface.	**File > Save**

How to Get Help

IBM Support

If you need help, service, or technical assistance, visit our web site at the following address:

http://www.ibm.com/support

The warranty card received with your product provides details for contacting a customer support representative. If you are unable to locate this information, please contact your reseller. Before you call, prepare the following information:

- Serial number of the switch unit
- Software release version number
- Brief description of the problem and the steps you have already taken
- Technical support dump information (`# show tech-support`)

Part 1: Getting Started

Chapter 1. IBM DS 5000V Introduction

Solution Overview

The IBM System Networking Distributed Switch 5000V is a software-based network switching solution designed for use with the virtualized network resources in a VMware-enhanced datacenter.

Figure 1. The IBM DS 5000V in a VMware vCenter

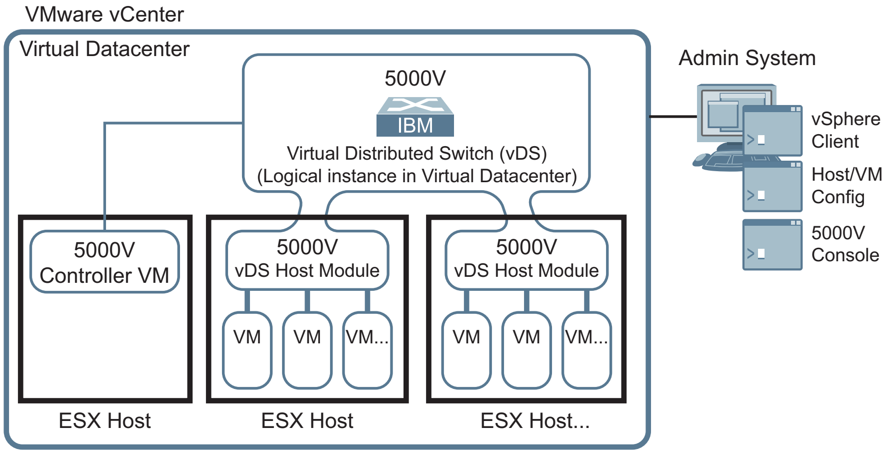

Using the VMware Virtual Distributed Switch (vDS) model, the IBM DS 5000V software switch modules are “distributed” to each participating VMware ESX host. Though each 5000V host module handles traffic for the local virtual machines (VMs), all distributed modules also work in unison as an aggregate virtual switching device. The 5000V solution can be roughly equated to an interconnected stack of independent switches which are unified and controlled by a single management plane.

The 5000V works with VMware vSphere and ESX 5.0 to provide an IBM Networking OS management plane, and advanced Layer 2 features in the control and data planes.

The management plane includes a built-in Command Line Interface (CLI) that runs on a VMware virtual machine. It is packaged as an Open Virtual Appliance (OVA) file.

The control/data plane is implemented by a software module that runs inside each participating ESX hypervisor. It is packaged as a vSphere Installation Bundle (VIB) file.

Using this VMware vDS model, the network administrator can define the 5000V at the datacenter level within the VMware vCenter. When ESX hosts in the data center join the 5000V, a virtual switch instance, or *portset,* is created on the host. Portsets inherit their properties from the global virtual switch. VMware vDS infrastructure synchronizes all the portsets and manages state migration during VMotion, the movement of virtual machines (VMs) within and among ESX hypervisors.

Note: The term "vDS" as used here refers to the infrastructure provided by VMware to achieve a distributed virtual switch implementation. VMware's vDS product has a different scope from the IBM DS 5000V and is expected to be managed by the server administrator, while the 5000V is generally managed by a network administrator.

Components

The IBM DS 5000V solution requires a VMware VCenter environment. It can be deployed as a stand-alone product within the VCenter, or as as a component of the IBM System Networking Software Defined Network for Virtual Environments (SDN-VE).

Note: This *User Guide* describes using the 5000V as a as a stand-alone vDS (without SDN-VE). For information on using the 5000V as part of the SDN-VE solution, refer to the *SDN-VE User Guide* instead.

The 5000V requires the following main components:

- VMware vCenter

  This is a VMware product that resides on a server within the datacenter. It provides a centralized tool for installing, managing and synchronizing vDS instances, hypervisors, and VMs on host servers throughout the datacenter.
- VMware vSphere Client

  This is a VMware product that resides on administrative client devices. It provides the server administrator or network administrator with rich, remote access to vCenter management tools.
- VMware ESX 5.0

  This is a VMware hypervisor product that resides on individual host servers within the datacenter. It provides the software infrastructure for installing, running, and managing VMs and vDSs on the hosts.
- IBM DS 5000V vDS Host Module

  This is an IBM product that resides in participating ESX hypervisors on host servers within the datacenter. It implements a vDS portset as defined in the VMware vDS API and acts a virtual network switch for the given host server. At its core, it forwards frames based on destination MAC addresses, controlling Layer 2 access to an from the associated VMs. It also provides advanced switching features such as VLANs, IGMP snooping, etc. The settings for each feature are configured by the network administrator through the 5000V controller.
- IBM DS 5000V Controller

  This is an IBM product than resides in a virtual machine within the datacenter. It works in conjunction with the VMware vCenter and ESX hypervisors to unify all 5000V vDS host modules into an aggregate superswitch. Through the VMware vSphere client, it provides a full CLI for switch configuration, operation, and the collection of switch information and statistics.

  All traffic to and from the controller is consolidated into single virtual NIC. This traffic includes the following:

  – Management traffic for applications like Telnet, SSH, SNMP, etc.

  – vSphere API traffic between the vSphere Client and the VMware vCenter.

  – Traffic between controller and the virtual switch elements on the ESX hosts.

Functional Overview

The installed 5000V controller creates a global 5000V vDS instance at the vCenter. The vCenter exposes a set of ports (a *portset*) to the 5000V controller which can then proceed to configure those ports.

The controller passes configuration settings to vCenter via vDS vendor-specific properties, using the vSphere API. The VMware vDS infrastructure then passes the settings to the individual ESX host portsets (instances of the 5000V vDS host modules) where the ports can be deployed (connected to VMs).

Ports can also be configured via policy profiles (similar to VMware port groups).

Operational Overview

The following steps describe the high-level operational tasks for installing and using the IBM DS 5000V.

For the network administrator:

1. Ensure that the required VMware software is installed and configured.
2. Download the 5000V controller OVA file from IBM.
3. Install and start the controller OVA file on a virtual machine in the datacenter.
4. Pair the 5000V controller with a VMware vCenter virtual datacenter.
5. Create a global 5000V vDS instance in the virtual datacenter.
6. Manage the global 5000V vDS just like a physical switch, using the controller's CLI. This includes provisioning, configuring, monitoring, and troubleshooting switch access/uplink ports and profiles in accordance with VM and host requirements, as well as your established network policies.

For the system administrator:

1. Download the 5000V vDS host module VIB file from IBM.
2. Install the vDS host module VIB on each VMware ESX 5.0 host where the 5000V will be deployed.
3. Join the ESX hosts to the global 5000V vDS instance created by the network administrator.
4. Designate one or more uplink ports (physical NICs) as necessary on each host.
5. Work with the network administrator to provision ports or vNIC profiles as needed.
6. Connect VMs to appropriate ports or port groups.

These topics will be covered in more detail in the other sections of this User Guide.

Advanced Features

The 5000V includes the following advanced Layer 2 features:

- VLAN
- Private VLAN
- Port Mirroring
- ERSPAN
- sFlow
- ACLs
- QoS
- LACP and Advanced Teaming
- SNMP
- RADIUS
- TACACS+
- Syslog
- EVB (802.1Qbg): VEPA, VDP, VSI Manager

Chapter 2. IBM DS 5000V Installation

Note: This User Guide describes using the 5000V as a as a stand-alone vDS in the VMware datacenter. For information on using the 5000V as part of the IBM Software Defined Network for Virtual Environments (SDN-VE), refer to the *SDN-VE User Guide* instead.

Prerequisites

The following tasks must be completed prior to IBM DS 5000V (5000V) installation:

- VMware vCenter Server must be installed and operational in your network (see the documentation provided with your vCenter product).
- The host servers which will be using the 5000V solution must be installed and operational, and include the following:
  - A valid VMware Enterprise Plus license must be installed on each host.
  - ESX 5.0 or 5.1 must be installed and operational.
  - There should be more than 1 host for vMotion.
  - Each host must have a minimum of one 1G or 10G physical NIC (three 1G or 10G NICs are recommended based on tested implementations). A separate 10G NIC is recommended for Management, VMKernel and virtual machine (VM) traffic.
  - All ESX hosts must have Layer 2/Layer 3 network connectivity.
- The controller: During the installation, a new VM will be deployed on an ESX host (any version) to hold the 5000V controller software. Requirements for the controller host include the following:
  - The host requires a minimum of 1 GB of physical RAM.
  - The host must have a minimum of one 1G or 10G physical NIC (three 1G or 10G NICs are recommended). A separate 10G NIC is recommended for Management, VMKernel and VM traffic.
  - The controller should have network connectivity to the target vCenter and all host servers which will hold a 5000V virtual switch element.
  - It is highly recommended that VMware High Availability and/or VMware Fault Tolerance features be configured to protect the 5000V against downtime or data loss.

The following 5000V software files are required:

- The IBM DS 5000V controller, an Open Virtual Appliance (OVA) file
- The IBM DS 5000V vDS host module, a vSphere Installation Bundle (VIB) file

Installing the 5000V Controller

The 5000V controller provides the core intelligence that unifies the operation of the individual 5000V vDS host modules that will be installed on the participating host servers. The 5000V controller software is contained in the IBM DS 5000V OVA file.

Follow these steps to install and start the 5000V controller:

1. Download the controller OVA file from IBM.
2. Place the OVA file on a system that has access to the VMware vSphere Client (such as an administrative laptop).

3. Specify an ESX host on which to deploy the controller.

   The controller host merely provides an environment in which the 5000V controller appliance will run. It is not required to participate as a vSD host and may be a different class of device than those where the vSD host modules will be installed. The primary requirement is for the controller host to have Layer 3 connectivity to the designated vCenter and the 5000V vDS host modules.

4. Launch the VMware vSphere Client and connect either to the vCenter that manages the host where the 5000V controller will be deployed or directly to the ESX host.

5. From the vSphere Client, select the target ESX host and choose **File > Deploy OVF Template** as shown below:

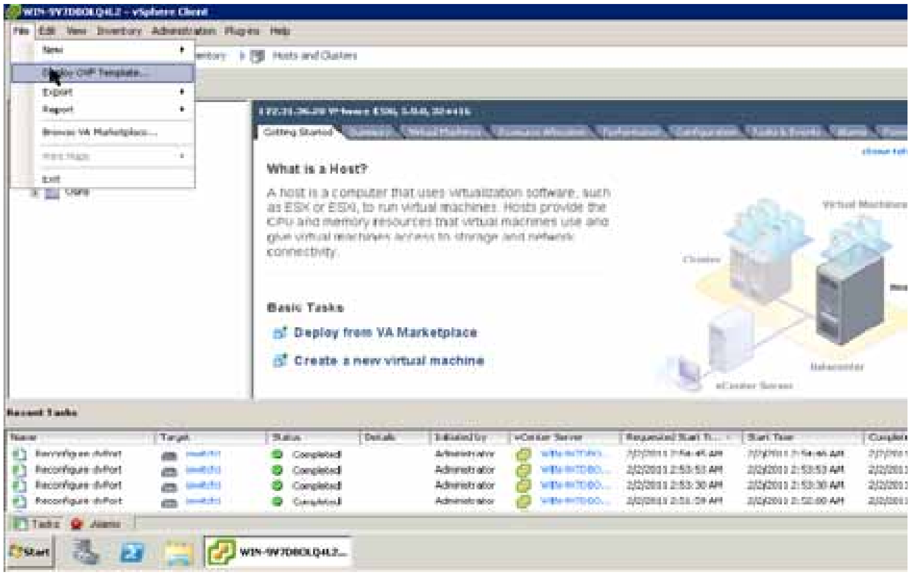

What is a Host?

A host is a computer that uses virtualization software, such as ESX or ESXi, to run virtual machines. Hosts provide the CPU and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

Basic Tasks

Deploy from VA Marketplace

Create a new virtual machine

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Ti...	Start Time	Complet
Reconfigure dvPort	[illegible]	Completed		Administrator	WIN-9V7DBO...	2/2/2011 2:54:45 AM	2/2/2011 2:54:45 AM	2/2/2011
Reconfigure dvPort	[illegible]	Completed		Administrator	WIN-9V7DBO...	2/2/2011 2:53:53 AM	2/2/2011 2:53:53 AM	2/2/2011
Reconfigure dvPort	[illegible]	Completed		Administrator	WIN-9V7DBO...	2/2/2011 2:53:30 AM	2/2/2011 2:53:30 AM	2/2/2011
Reconfigure dvPort	[illegible]	Completed		Administrator	WIN-9V7DBO...	2/2/2011 2:51:59 AM	2/2/2011 2:52:00 AM	2/2/2011

6. Select the location where the OVA file is stored and click **Next**.

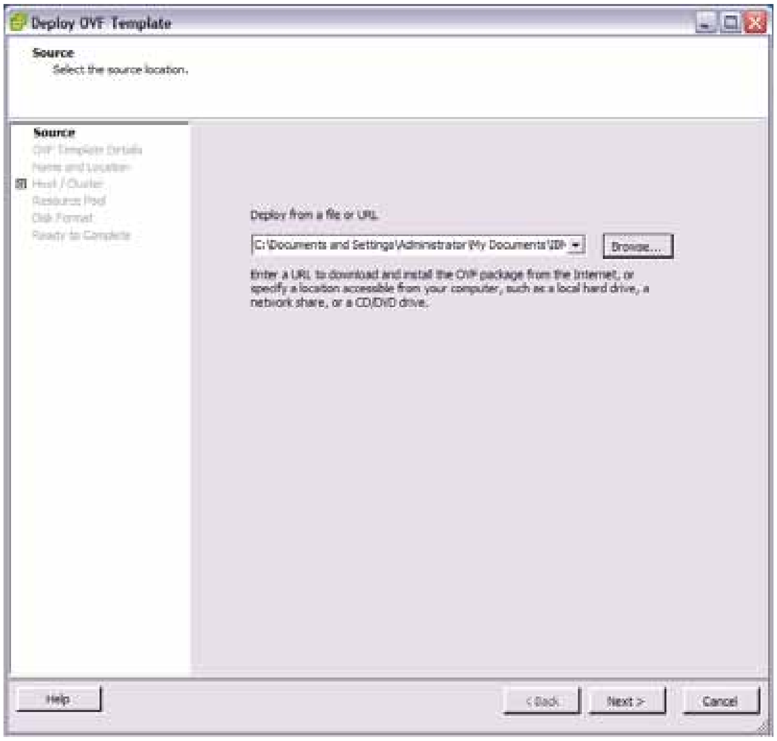

Source
Select the source location.

Deploy from a file or URL

C:\Documents and Settings\Administrator\My Documents\IBM

Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

7. Verify the OVA details and click **Next**.

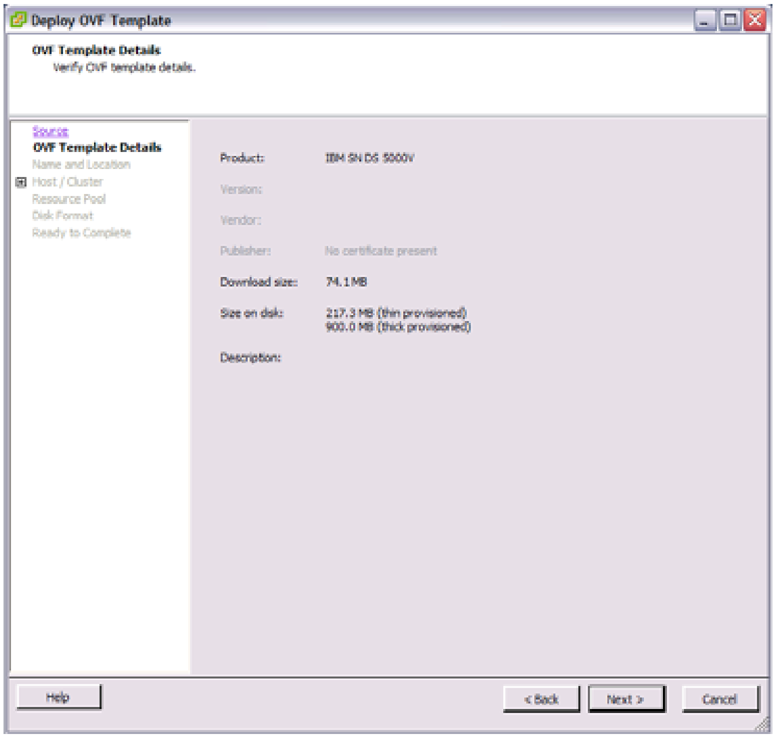

OVF Template Details
Verify OVF template details.

Product: IBM SN DS 5000V

Version:

Vendor:

Publisher: No certificate present

Download size: 74.1 MB

Size on disk: 217.3 MB (thin provisioned)
900.0 MB (thick provisioned)

Description:

8. Provide a name the 5000V controller and click Next.

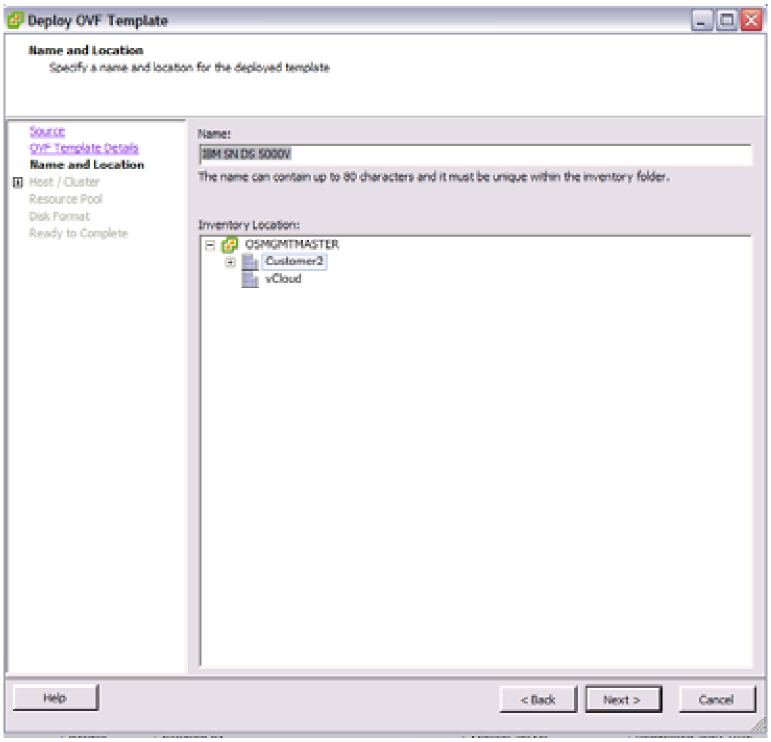

Deploy OVF Template

Name and Location

Specify a name and location for the deployed template

Name: IBM SN DS 5000V

The name can contain up to 80 characters and it must be unique within the inventory folder.

Inventory Location:

9. Specify the host or cluster on which to deploy the 5000V controller and click **Next**.

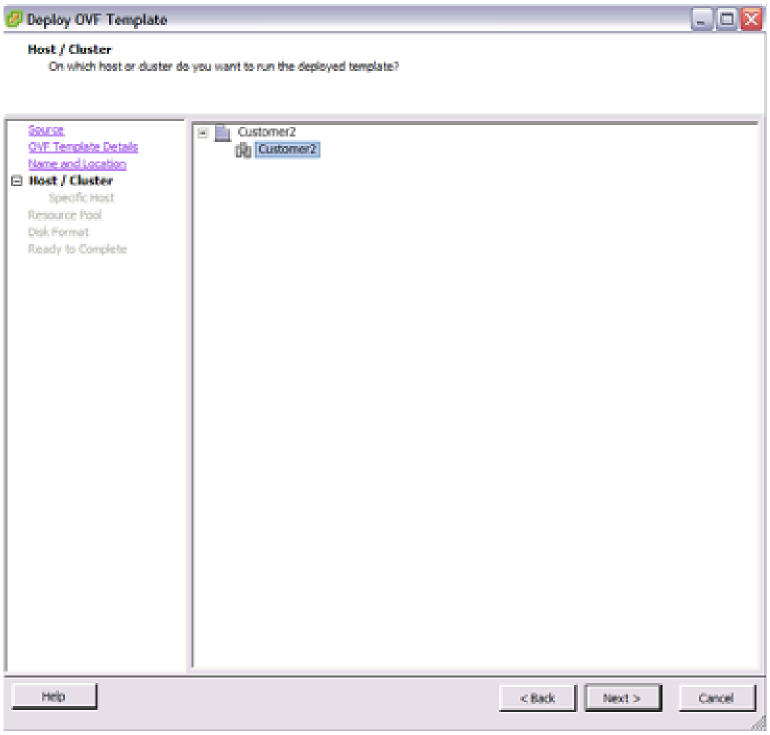

Deploy OVF Template

Host / Cluster

On which host or cluster do you want to run the deployed template?

10. Specify a location on the VM where 5000V controller files should be stored, and click **Next**.

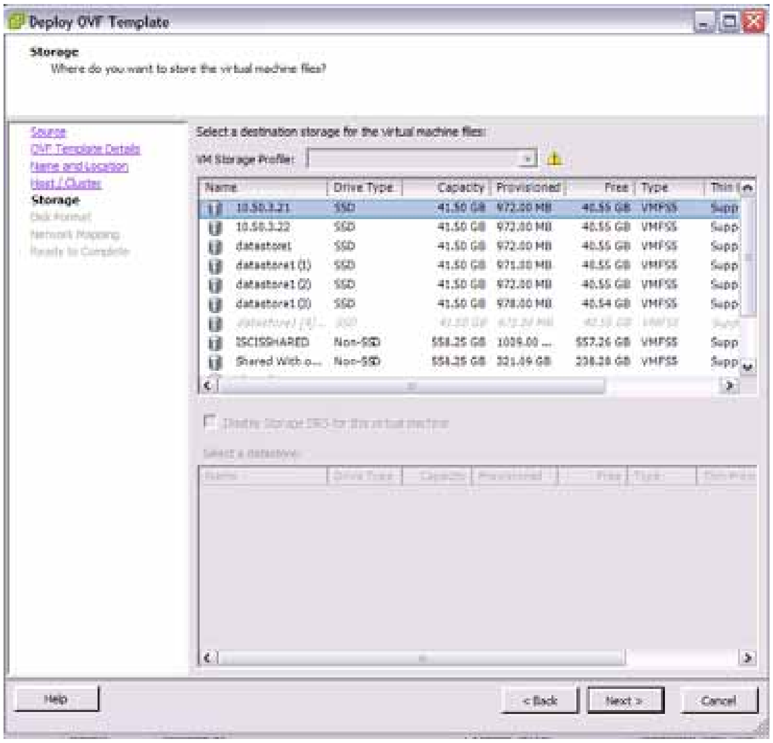

Storage

Where do you want to store the virtual machine files?

Select a destination storage for the virtual machine files:

VM Storage Profile:

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin
10.50.3.21	SSD	41.50 GB	972.00 MB	40.55 GB	VMFS5	Supp
10.50.3.22	SSD	41.50 GB	972.00 MB	40.55 GB	VMFS5	Supp
datastore1	SSD	41.50 GB	972.00 MB	40.55 GB	VMFS5	Supp
datastore1 (1)	SSD	41.50 GB	971.00 MB	40.55 GB	VMFS5	Supp
datastore1 (2)	SSD	41.50 GB	972.00 MB	40.55 GB	VMFS5	Supp
datastore1 (3)	SSD	41.50 GB	978.00 MB	40.54 GB	VMFS5	Supp
datastore1 (4)...	SSD	41.50 GB	972.00 MB	40.55 GB	VMFS5	Supp
ISCISSHARED	Non-SSD	558.25 GB	1029.00 ...	557.26 GB	VMFS5	Supp
Shared With o...	Non-SSD	558.25 GB	321.09 GB	238.28 GB	VMFS5	Supp

□ Disable Storage DRS for this virtual machine

Select a datastore:

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin Prov

11. Select a disk format and click **Next**. The recommended format is Thick Provisioned Lazy Zeroed.

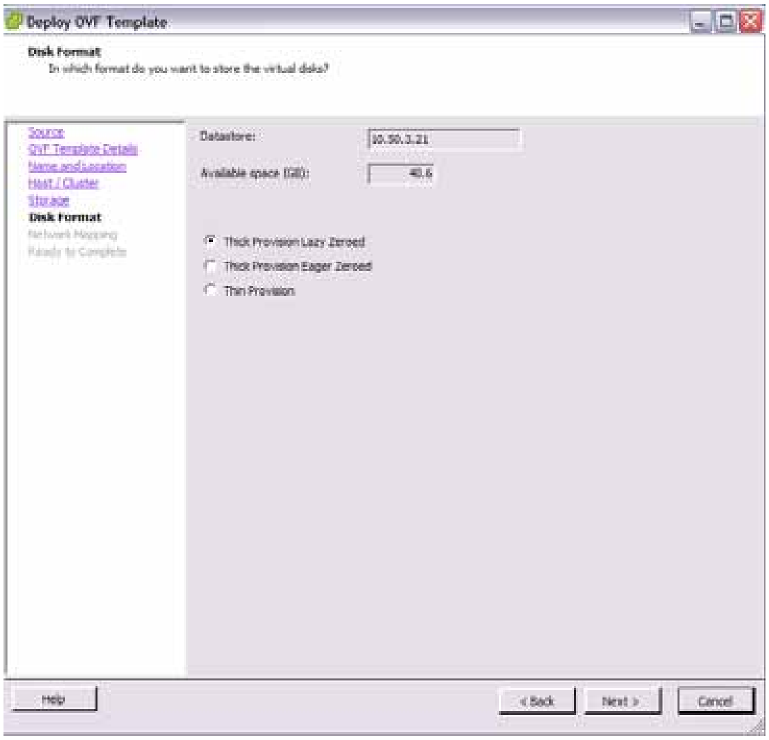

Disk Format

In which format do you want to store the virtual disks?

Datastore: 10.50.3.21

Available space (GB): 40.6

● Thick Provision Lazy Zeroed

○ Thick Provision Eager Zeroed

○ Thin Provision

12. Map the network for 5000V controller use and click **Next**.

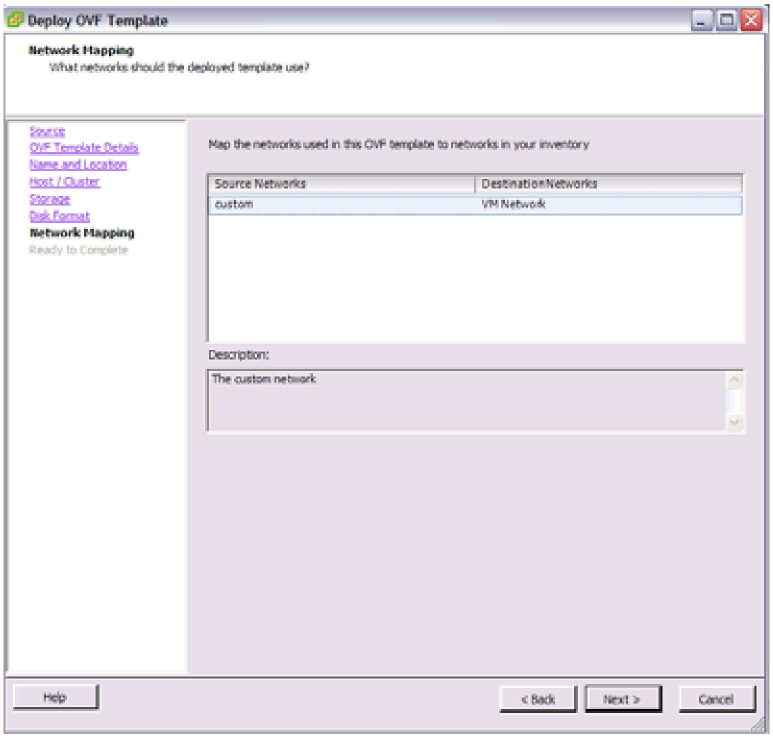

Network Mapping
What networks should the deployed template use?

Map the networks used in this OVF template to networks in your inventory

Source Networks	DestinationNetworks
custom	VM Network

Description:
The custom network

13. Verify the specified options, select the “Power on after deployment” option, and click **Next**.

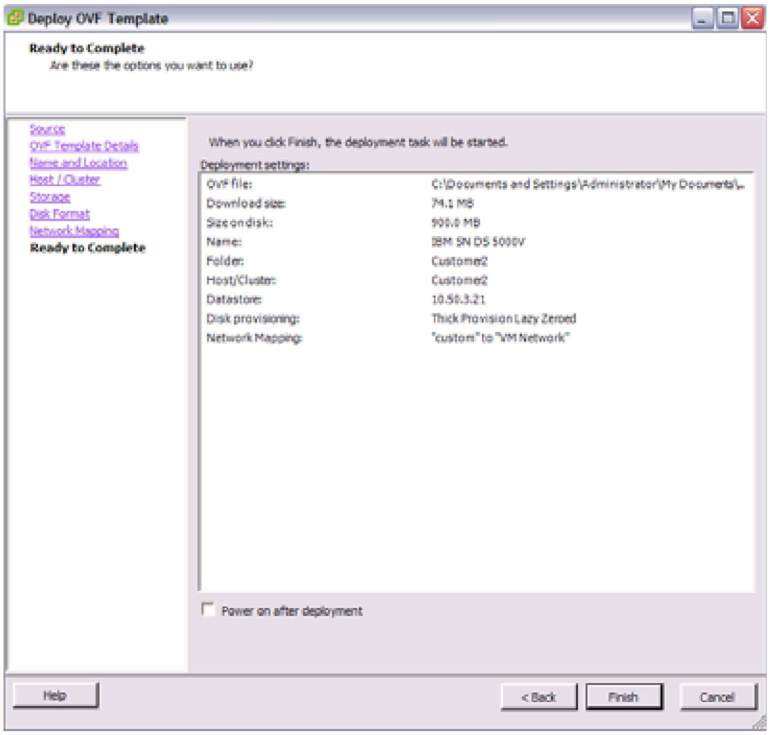

Ready to Complete
Are these the options you want to use?

When you click Finish, the deployment task will be started.

Deployment settings:

OVF file:	C:\Documents and Settings\Administrator\My Documents\...
Download size:	74.1 MB
Size on disk:	900.0 MB
Name:	IBM SN DS 5000V
Folder:	Customer2
Host/Cluster:	Customer2
Datastore:	10.50.3.21
Disk provisioning:	Thick Provision Lazy Zeroed
Network Mapping:	"custom" to "VM Network"

□ Power on after deployment

This will initiate the 5000V controller VM deployment:

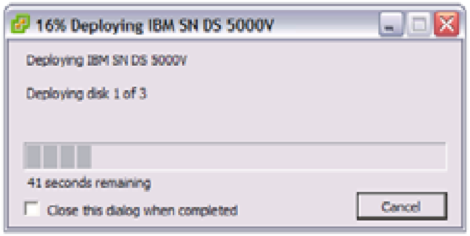

16% Deploying IBM SN DS 5000V

Deploying IBM SN DS 5000V

Deploying disk 1 of 3

41 seconds remaining

The 5000V controller VM will power on when deployment is complete, and the controller VM console will appear.

Initial Controller Setup

The 5000V must be manually configured by entering commands into the controller’s built-in Command-Line Interface (CLI). Initially, the CLI can be accessed only through the 5000V controller VM console on the vSphere Client. Later, if desired, the CLI can be accessed via remote Telnet or SSH connections.

Note: Configuration of the 5000V vDS must be performed solely from the 5000V CLI, and not through the vCenter interface (even in cases where the vCenter interface seems to allow it). However, host operations (such as adding ESX hosts and uplinks or assigning VM network interfaces to vDS ports or profiles) must be done through the vCenter interface.

Starting the 5000V Controller VM Console

When following the provided controller installation instructions (see Step 13 on page 26), the controller console automatically appears when the 5000V controller VM is powered on.

However, to manually access the controller console under other conditions, use the following procedure:

1. Log-in to the VMware vCenter via your vSphere Client.
2. Right-click on the 5000V controller VM and select the option to “Open Console.” Alternately, you can click on the Console icon.

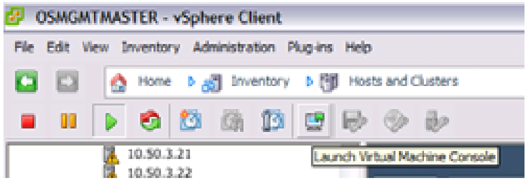

The VM console for the 5000V controller will appear.

Examine the License Agreement

The first time the 5000V controller is started, you will be prompted to read the Software Licence Agreement. When you select a language, the SLA will be displayed.

When you are finished examining the SLA, select **1** if you wish to accept the terms.

If you accept the SLA, the 5000V controller login prompt will appear.

Logging In to the Controller

CLI access is controlled through the use of a login name and password. Once you are connected to the 5000V controller, you are prompted to enter a login name and password. The default log-in user name and password are as follows:

Default user name:	`admin`
Default password:	`admin`

It is recommended that you change all default system password after initial configuration and as regularly as required under your network security policies.

Entering Configuration Mode

The 5000V controller uses a rich CLI command set with multiple command modes. For an overview of CLI modes and features, see "CLI Basics" on page 109. The remainder of this chapter will display all commands necessary for initial configuration, but only those command relevant to the specific configuration tasks will be called into example.

After logging in, perform the following commands to enter the CLI Global Configuration mode::

```
5000V> ena
5000V# configure terminal
5000V(config)#
```

The `ena` command initiates executive privilege mode, and the `configure terminal` command readies the controller for configuration.

Setting an IP Address

The 5000V controller must have IP connectivity to the VMware vCenter, as well as the hosts that will participate in the 5000V vDS.

By default, the 5000V controller is enabled for DHCP. If there is a DHCP server available in your network, the controller will automatically acquire its IP address.

However, if DHCP is not available in your network, configure a static IP address for the 5000V controller as follows:

```
5000V(config)# interface ip-mgmt address <IP address> [<mask>] [<gateway>]
5000V(config)# interface ip-mgmt gateway enable
```

Where *IP address* is the address of the controller in dotted-decimal notation, optionally followed by the network *mask* used for creating an address range, and finally, the *gateway* IP address that the controller should use for outbound traffic.

If preferred, use the following commands to independently configure each item:

```
5000V(config)# interface ip-mgmt address <IP address>
5000V(config)# interface ip-mgmt netmask <mask>
5000V(config)# interface ip-mgmt gateway <gateway IP address>
5000V(config)# interface ip-mgmt gateway enable
```

Enabling Remote Access

For system security, initial access to the 5000V controller is permitted only through the VM console. However, remote management can be performed via Telnet or SSH.

Telnet Access

A Telnet connection offers the convenience of accessing the switch from a workstation connected to the network. Telnet access provides the same options for user and administrator access as those available through the console port.

By default, Telnet access is disabled. Use the following commands (available on the console only) to enable Telnet access:

```
5000V(config)# access telnet enable
```

If the switch is configured with an IP address and gateway (either manually or via DHCP), you can use Telnet to access switch administration from any workstation connected to the management network.

To establish a Telnet connection with the switch, run the Telnet program on your workstation and issue the following Telnet command:

```
telnet <switch IP address>
```

You will then be prompted to log in.

Note: By default, application port 23 is used for Telnet traffic. To change the port, or to disable Telnet access once enabled, see the `access telnet` options in the CLI reference.

Secure Shell Access

Although a remote network administrator can manage the configuration of the 5000V via Telnet, this method does not provide a secure connection. The Secure Shell (SSH) protocol enables you to securely log into another device over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure.

Using SSH to Access the Switch

By default, the SSH feature is disabled. Use the following commands to enable SSH access:

```
5000V(config)# ssh enable
```

If the switch is configured with an IP address and gateway (either manually or via DHCP) and the SSH service is enabled, you can access the CLI using an SSH connection.

To establish an SSH connection with the switch, run the SSH program on your workstation by issuing the SSH command, followed by the switch IPv4 or IPv6 address:

```
ssh <switch IP address>
```

If SecurID authentication is required, use the following command:

```
ssh -1 ace <switch IP address>
```

You will then be prompted to enter a password.

Note: By default, application port 22 is used for SSH. Also, a 1 hour host key is used. For more information, see "Secure Shell and Secure Copy" on page 45.

Simple Network Management Protocol Access

The 5000V provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software, such as IBM Director or HP-OpenView.

Note: SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network.

To access the SNMP agent on the 5000V, the read and write community strings on the SNMP manager must be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands:

```
5000V(config)# snmp-server read-community <1-32 characters>
    -and-
5000V(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager must be able to reach any one of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch with the following commands:

```
5000V(config)# snmp-server trap-src-if <trap source IP interface>
5000V(config)# snmp-server host <IPv4 address> <trap host community string>
```

For more information on SNMP usage and configuration, see "Simple Network Management Protocol" on page 89.

Creating the Global vDS Instance

The VMware solution employs a hierarchical structure.

Figure 2. VMware vCenter Hierarchy

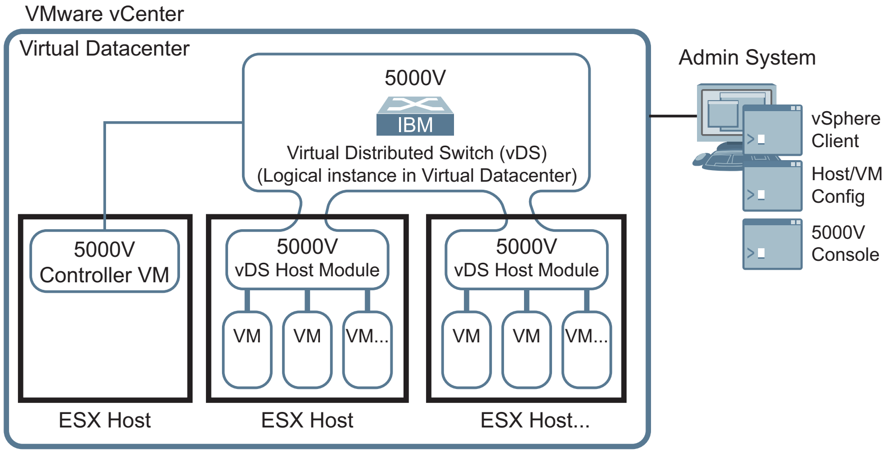

Within the vCenter interface, one or more virtual data centers can be defined. Each virtual data center is comprised of a logical set of ESX hosts, clusters, and virtual distributed switches (vDSs). Each ESX host can contain a vDS host module which serves the host itself, but these host modules also act as a part of the larger vDS operating at the virtual data center level.

The 5000V controller must be associated with a vDS for a particular virtual data center. The following CLI commands on the controller VM console are used to create the required association to the vCenter:

```
5000V(config)# iswitch vcenter <vCenter IP address> <user name>
```

The *vCenter IP address* represents the vCenter to which the 5000V will connect for configuration and *username* is the vCenter login name.

The system will then prompt you for the vCenter login password and its logical port number. By default, the vCenter operates on recommended TCP port number 443. However, if your vCenter communicates on a different port, enter the port number configured for the service.

Next, the 5000V controller must be associated with the vDS:

```
5000V(config)# iswitch vds <name> <datacenter>
```

When this configuration is complete, the 5000V vDS will appear at the vCenter in the **Home > Inventory > Networking** view:

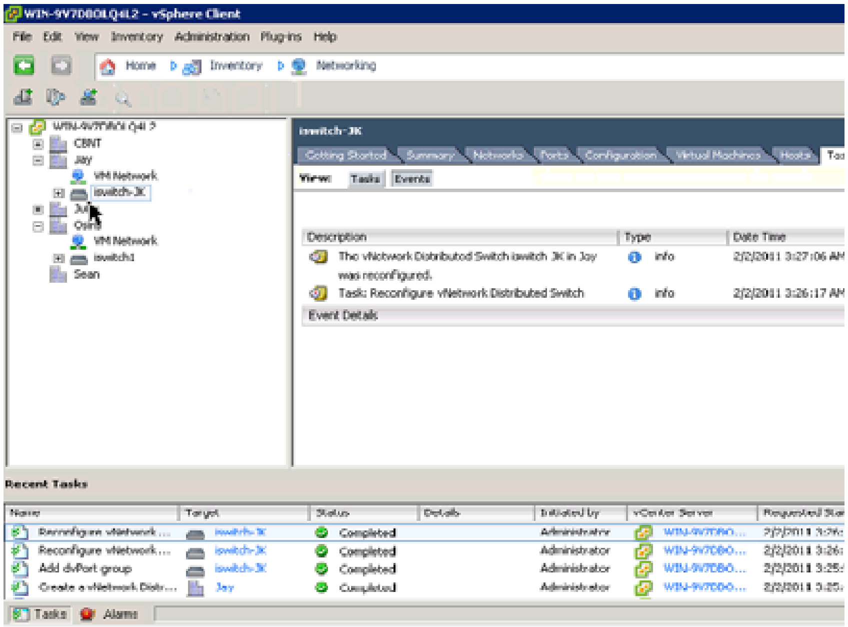

Description	Type	Date Time
The vNetwork Distributed Switch iswitch JK in Jay was reconfigured.	info	2/2/2011 3:27:06 AM
Task: Reconfigure vNetwork Distributed Switch	info	2/2/2011 3:26:17 AM

Event Details

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start
Reconfigure vNetwork ...	iswitch-JK	Completed		Administrator	WIN-9V7DBO...	2/2/2011 3:26:
Reconfigure vNetwork ...	iswitch-JK	Completed		Administrator	WIN-9V7DBO...	2/2/2011 3:26:
Add dvPort group	iswitch-JK	Completed		Administrator	WIN-9V7DBO...	2/2/2011 3:25:
Create a vNetwork Distr...	Jay	Completed		Administrator	WIN-9V7DBO...	2/2/2011 3:25:

Note: Once the controller is associated a vDS in the vCenter, whenever the IP address of the 5000V controller is changes (statically or via DHCP renewal), you must save the 5000V configuration and reload the controller in order to reestablish the required association.

Installing vDS Host Modules

The 5000V vDS host module can be installed in individual ESX hypervisors. There, it implements a vDS portset and acts a virtual network switch for the host. The software is available in two formats: as an IBM DS 5000V vSphere Installation Bundle (VIB) file or as an offline bundle (ZIP) file.

The software may be installed on one or more ESX hosts by using the VMware Update Manager (VUM) or on individual ESX hosts via the vCLI or ESXi shell (ESXCLI) on the host.

Using the ESXCLI

1. Ensure that SSH is enabled on the ESX host. This can be done via Troubleshooting Options in the ESX 5.0 console menu.
2. Copy the IBM DS 5000V VIB or ZIP file to the ESX host via Secure Copy (SCP).
3. Initiate the following ESXSCLI command on the host:

```
# esxcli software vib install -v=file://<path to the VIB file>
    -or-
# esxcli software vib install -d=file://<path to the ZIP file>
```

After the installation, the host will automatically reboot and the vDS host module will be activated, at which point the vDS host module can be joined to one or more global 5000V vDS instances in the virtual data center.

You can verify 5000V vDS host module installation with the following useful ESXCLI commands:

Table 2. ESXCLI Verification Commands

Command	Function
`esxcli software vib list`	List the VIBs present on the host.
`vmkload_mod -l`	Verify that the 5000V vDS host module is loaded.
`ps | grep 5000V`	Verify that the 5000V vDS agent process is running.

Using the VMware Update Manager

The 5000V offline bundle is a .zip file which can be installed on multiple ESX hosts using the VMware vSphere Update Manager (VUM) plug-in.

1. To load the 5000V offline bundle, access the VUM tool and select the Configuration tab.

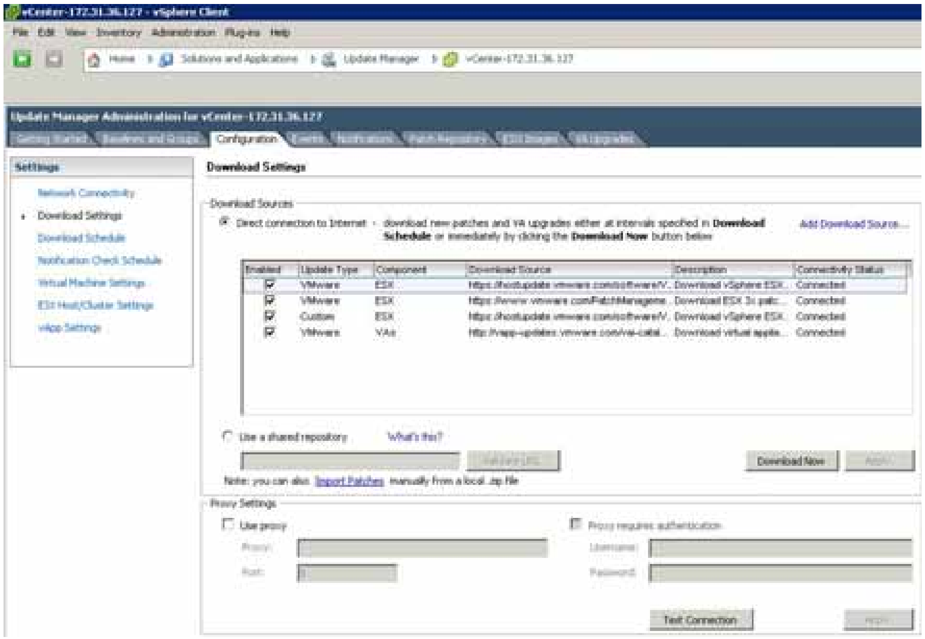

● Direct connection to Internet - download new patches and VA upgrades either at intervals specified in **Download Schedule** or immediately by clicking the **Download Now** button below

Enabled	Update Type	Component	Download Source	Description	Connectivity Status
☑	VMware	ESX	https://hostupdate.vmware.com/software/V...	Download vSphere ESX...	Connected
☑	VMware	ESX	https://www.vmware.com/PatchManageme...	Download ESX 3i patc...	Connected
☑	Custom	ESX	https://hostupdate.vmware.com/software/V...	Download vSphere ESX...	Connected
☑	VMware	VAs	http://vapp-updates.vmware.com/vai-catal...	Download virtual applia...	Connected

○ Use a shared repository

Note: you can also Import Patches manually from a local .zip file

□ Use proxy
Proxy:
Port:

□ Proxy requires authentication
Username:
Password:

2. Click on the “Import Patches” link in the note at the bottom of the Download Sources section of the page. The full note reads “Note: you can also Import Patches manually from a local .zip file.”

3. In the Import Patches window, select the offline bundle .zip file and click **Next**.

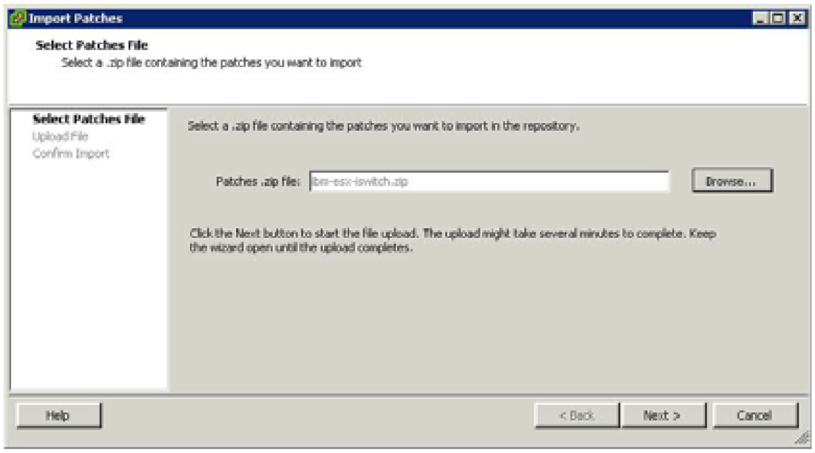

Select Patches File

Select a .zip file containing the patches you want to import

Select a .zip file containing the patches you want to import in the repository.

Patches .zip file: ibm-esx-iswitch.zip

Click the Next button to start the file upload. The upload might take several minutes to complete. Keep the wizard open until the upload completes.

Once the file has been processed, the 5000V vDS Host Module will appear in the Patch Repository tab as follows:

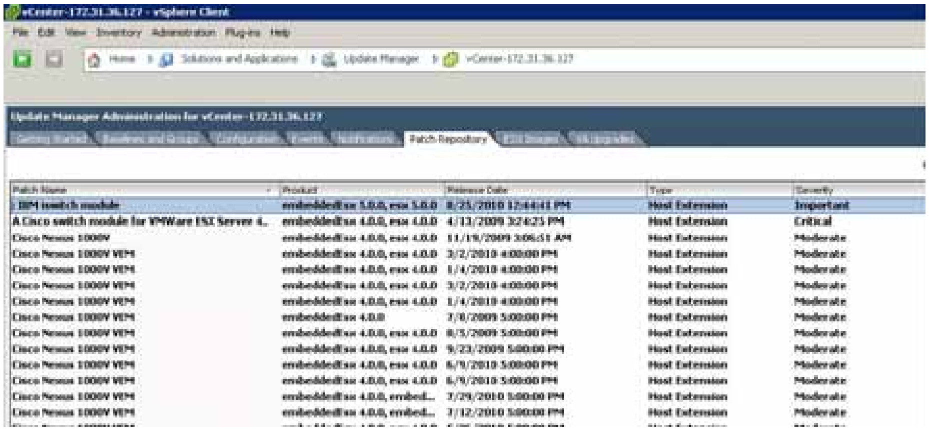

Patch Name	Product	Release Date	Type	Severity
IBM iswitch module	embeddedEsx 5.0.0, esx 5.0.0	8/25/2010 12:44:41 PM	Host Extension	Important
A Cisco switch module for VMWare ESX Server 4...	embeddedEsx 4.0.0, esx 4.0.0	4/13/2009 3:24:25 PM	Host Extension	Critical
Cisco Nexus 1000V	embeddedEsx 4.0.0, esx 4.0.0	11/19/2009 3:06:51 AM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	3/2/2010 4:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	1/4/2010 4:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	3/2/2010 4:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	1/4/2010 4:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0	7/8/2009 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	8/5/2009 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	9/23/2009 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	6/9/2010 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, esx 4.0.0	6/9/2010 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, embed...	7/29/2010 5:00:00 PM	Host Extension	Moderate
Cisco Nexus 1000V VEM	embeddedEsx 4.0.0, embed...	7/12/2010 5:00:00 PM	Host Extension	Moderate

Chapter 3. Host, VM and Port Management

Connecting vDS Host Modules & Uplink Ports

Within a virtual data center, a global 5000V vDS is comprised of a controller and its associated vDS Host Modules. To connect the various 5000V vDS Host Modules to to the global vDS instance representing the 5000V controller, use the following process:

1. Log-in to the VMware vCenter via your vSphere Client.
2. Right-click on the global vDS instance within target virtual data center and select the “Add Host” option.

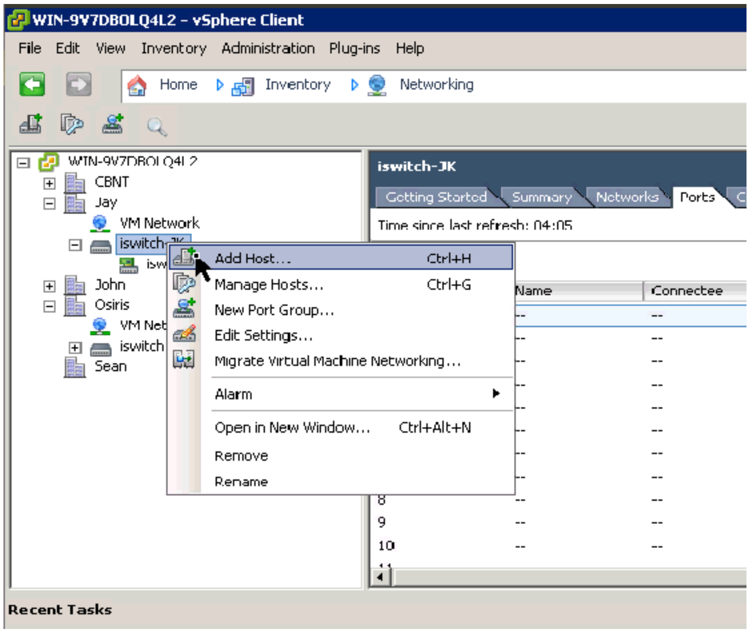

3. Select the hosts you wish to add. At this time, you can also specify any uplink ports (physical NICs) for the 5000V. Click **Next** when selections are complete.

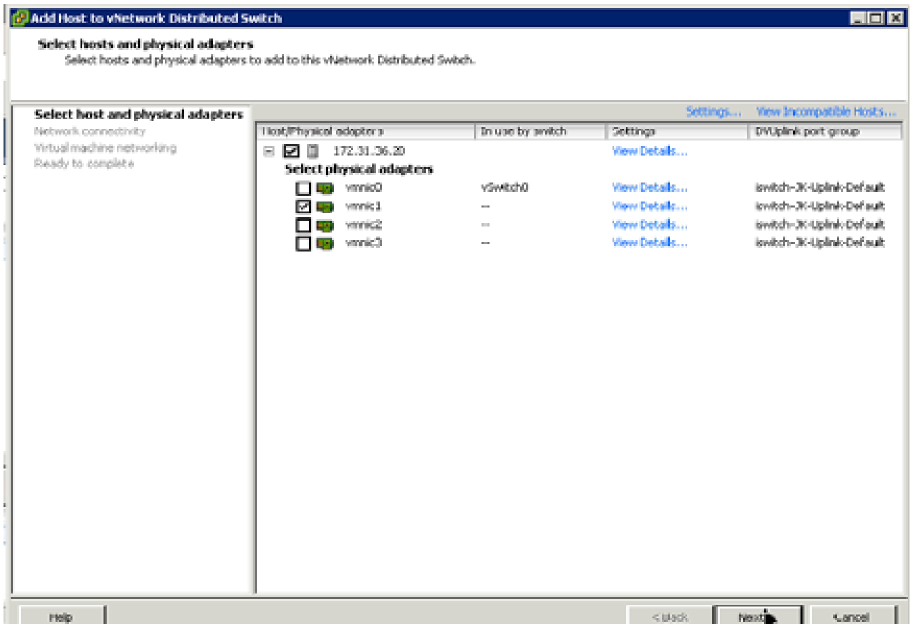

Add Host to vNetwork Distributed Switch

Select hosts and physical adapters

Select hosts and physical adapters to add to this vNetwork Distributed Switch.

Host/Physical adapters	In use by switch	Settings	DVUplink port group
☑ 172.31.36.20		View Details...	
Select physical adapters			
□ vmnic0	vSwitch0	View Details...	iswitch-JK-Uplink-Default
☑ vmnic1	--	View Details...	iswitch-JK-Uplink-Default
□ vmnic2	--	View Details...	iswitch-JK-Uplink-Default
□ vmnic3	--	View Details...	iswitch-JK-Uplink-Default

Select host and physical adapters
Network connectivity
Virtual machine networking
Ready to complete
Settings...
View Incompatible Hosts...
Help
< Back
Next >
Cancel

Connecting Virtual Machines

VMs located on the hosts associated with the global 5000V vDS can be connected either to a specific 5000V port or to a vNIC profile (port group). Both types of connections are performed from the VM properties window in the vSphere Client.

1. Right-click on the target VM and selet the “Properties” option.
2. Select the Hardware tab.

3. Specify the port or port group:

   – To connect a VM to a port group:

     Use the standard settings in the Network Connection section of the page, and select the port group from the “Network label” field:

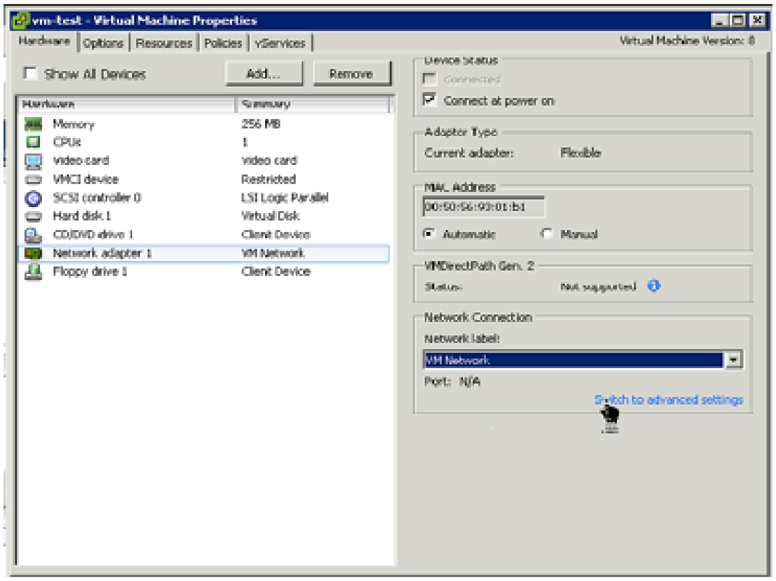

vm-test - Virtual Machine Properties

Hardware	Summary
Memory	256 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
CD/DVD drive 1	Client Device
Network adapter 1	VM Network
Floppy drive 1	Client Device

Device Status
☐ Connected
☑ Connect at power on

Adapter Type
Current adapter: Flexible

MAC Address
00:50:56:93:01:b1
● Automatic ○ Manual

VMDirectPath Gen. 2
Status: Not supported

Network Connection
Network label:
VM Network
Port: N/A

   – To connect a VM to a stand-alone port:

     Switch to the advanced settings (at the bottom of the Network Connection section). Specify the global vDS instance to which the VM will connect, and specify the virtual switch port for the connection:

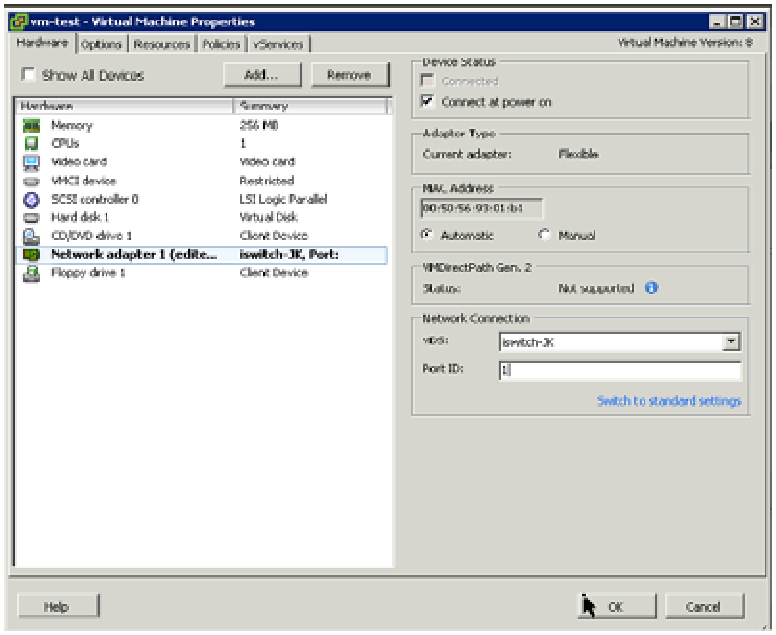

vm-test - Virtual Machine Properties

Hardware	Summary
Memory	256 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
CD/DVD drive 1	Client Device
Network adapter 1 (edite...	**iswitch-JK, Port:**
Floppy drive 1	Client Device

Device Status
☐ Connected
☑ Connect at power on

Adapter Type
Current adapter: Flexible

MAC Address
00:50:56:93:01:b1
● Automatic ○ Manual

VMDirectPath Gen. 2
Status: Not supported

Network Connection
vDS: iswitch-JK
Port ID: 1

4. Click **OK** when the selection is complete.

Port Management

The 5000V can be thought of as a massive switch comprised of access ports (those connected to virtual Ethernet NICs on virtual machines) and uplink ports (those connected to physical NICs, generally for the purpose of connection to the greater network.

Initially all the ports in the 5000V are defined as "undeployed." However, as hosts and VMs are connected, uplink ports get deployed as connections to pNICs which are attached to the vDS host module in an ESX host. The access ports get deployed as connections to virtual machine network adapters.

Each access port has a unique port number in the 5000V controller interface, and another in the vCenter interface. Configuration operations performed on ports via the controller CLI should use the 5000V controller port number for the port, while operations performed on ports via the vCenter (such as attaching a port to a VM) should use the VMware vDS Port number.

Uplinks, on the other hand, do not have a port number at the 5000V controller, but instead, the controller provisions one or more "uplink profiles" for use in each ESX host.

The system administrator chooses an appropriate uplink profile for each ESX host added to the global 5000V vDS. All uplinks connected to the same vDS within an ESX host are treated as a single aggregated link. Packets are never switched directly between the uplink ports of a given vDS. Because of this, the vDS ports can never introduce a loop in the network. Therefore STP is not required or supported on the 5000V.

Access Port Management

There are two types of access ports: stand-alone ports and those that belong to a vNIC profile (port group). Standalone ports are independent ports that must be configured individually. Ports that belong to a vNIC profile inherit all the properties of the profile as a group, requiring little or no individual configuration.

The choice of type depends on your network environment. For example, since there can be a maximum of 256 vNIC profiles on a single VMware vDS, if there is a need to define more than 256 unique vNIC profiles in a given environment, stand-alone ports would have to be used for some or all port definitions. In such environments, network and system administrators might be able to coordinate the required port settings by using specific ports rather than profiles. However, in other environments, hundreds of VMs might be needed, each with the exact same settings. In that case, it would be cumbersome to use stand-alone ports instead of profiles.

Access Port Numbers

When the global 5000V vDS instance is first created in the virtual data center, 100 stand-alone access ports are created with default properties for immediate use. The number of stand-alone ports can be increased up to a maximum of 4,000.

To increase the number of standalone access ports, use the following configuration command on the controller CLI:

```
5000V(config)# iswitch addports <amount> [<base>]
```

where *amount* is a multiple of 10, and `base` is an optional starting port number, if greater control over the port numbers is desired.

Configuring Stand-alone Ports

To configure individual ports, use the commands in the interface port mode:

```
5000V(config)# interface port <number>
5000V(config-if)# ?
```

where number is the target port number of the access port. The question mark (?) will display the various commands available for port configuration.

When finished, use the `exit` command to leave the interface port mode.

Configuring vNIC Profiles

To create and configure vNIC profiles, use the following base command:

```
5000V(config)# iswitch vnicprof <profile name>
5000V(config-vnic-profile)# ?
```

The question mark (?) will display the various commands available for port configuration.

Note that vNIC profiles support only a subset of features supported in stand-alone port configuration. This subset was chosen based on the most likely set of parameters that multiple VMs are likely to share. For example, VLANs are a part of vNIC profile configuration, but port mirroring is not. However, once a vNIC profile is created, you could enter the interface port menu for a specific port in the vNIC profile port and enable non-profile properties. For example, the user can configure port mirroring on an individual port that belongs to a vNIC profile.

Access Port Mapping

Due to the way that the VMware vCenter manages vDS port numbers, the port numbers used in the 5000V controller CLI do not necessarily match the port numbers displayed in the vCenter tools. To help manage any discrepancy, port mapping information is displayed at both the vCenter and in the 5000V controller CLI:

- At the vCenter, the port view of the vDS also displays the 5000V port number for each port.
- In the 5000V controller, the `show iswitch ports` command also displays vCenter port information.

Uplink Port Management

Uplink Profiles

Uplinks ports are not numbered or configured like the access ports on the 5000V. Instead, all of the uplinks connected to the same vDS host module within an ESX host are treated as a single link aggregation (LAG). To configure the uplinks, the administrator creates and applies uplink profiles. The uplink profiles contain settings such as the LAG type and MTU size.

A default uplink profile is defined when the global 5000V vDS is created in the virtual datacenter. By default, the aggregation model is asymmetric and the MTU is 1,500.

All uplink ports associated with the same ESX host (via the same vDS host module) must use the same uplink profile. Any profiles created at the 5000V controller are also available at the vCenter.

However, though the vCenter interface allows the system administrator to select a different uplink profile for each physical NIC (uplink port) when adding the NIC to the vDS, the 5000V only supports a single uplink profile per vDS host module. If you add a physical NIC and set a different uplink profile than the one currently active on the ESX host's vDS host module, the newly added physical NIC will be disabled (placed in a blocked state in the vCenter).

Uplink profiles are defined in the Uplink Profile mode:

```
5000V(config)# iswitch uprof <uplink profile name>
5000V(config-uprof)# ?
```

See "iSwitch vNIC Profile Configuration" on page 214 and "iSwitch Uplink Profile Configuration" on page 216 for profile options such as VLAN, tagging, PVID, 802.1 priority, ACLs, hash, and MTU.

LAG Modes

The LAG mode defines how the uplinks are aggregated. The modes supported in the uplink profile are: asymmetric, static and LACP.

Asymmetric LAG Mode

This mode allows all uplinks connected to the same vDS host module within an ESX host to form a logical LAG that is transparent to physical switches connected to the physical NICs. In this mode, the uplinks can connect to different physical switches, and their connected switch ports should not be configured to perform any type of link aggregation. ISLs (inter-switch links) should be configured properly to ensure all uplink-connected switch ports are in the same Layer 2 broadcast domain.

In this mode, the 5000V ensures that packets with a given source MAC address always use the same uplink port to avoid FDB problems on the connected switches. Duplicated and reflected packets will not be delivered to VMs.

Within the Uplink Profiles configuration mode, asymmetric LAG can be specified using the following CLI command:

```
5000V(config-uprof)# lagmode asymmetric
```

Note: This is the default mode and supports use of the EVB feature. See "Edge Virtual Bridging" on page 83 for more information.

Static LAG Mode

All uplinks ports connected to the same vDS host module within an ESX host will form a static 802.3ad-compatible link aggregation group. The physical switch ports connected the uplink ports' physical NICs must also be configured to form a corresponding static 802.3ad-compatible link aggregation group.

Within the Uplink Profiles configuration mode, static LAG can be specified using the following CLI command:

```
5000V(config-uprof)# lagmode static
```

To return to the global configuration mode, use the `exit` command.

Note: EVB is not supported for uplinks using static LAG mode.

LACP Mode

In this mode, 802.3ad LACP will run on each of the uplink. At most one LACP LAG will be formed for a vds within an ESX Host. When the uplinks connected to different partners, only the uplink(s) connected to the partner on which the LACP negotiation is firstly completed, will be LACP selected; others will be LACP unselected. When the aggregator is released, uplink(s) connected to another partner can be elected to use the aggregator.

Within the Uplink Profiles configuration mode, LACP LAG parameters can be specified using the following CLI commands:

```
5000V(config-uprof)# lagmode lacp
5000V(config-uprof)# lacp mode {active|passive}
5000V(config-uprof)# lacp system-priority <0-65535>
5000V(config-uprof)# lacp timeout {long|short}
```

Note: EVB is not supported for uplinks using LACP mode.

LACP Hash

The hash defines the load balancing method used within the LAG. This LAG load balancing can be based on MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The hash methods supported are as follows:

- `VPORT`: Transmit based on the hash value of VDS port ID which originates the packet. With this hash method, packets originated from the same access port are forwarded over the same uplink port.
- `SMAC`: Transmit based on the hash value of the source MAC address. With this hash method, packets with the same source MAC address are forwarded over the same uplink port.
- `DMAC`: Transmit based on the hash value of the destination MAC address.
- `DSMAC`: Transmit based on the hash value of both the destination and source MAC address.
- `SIP`: Transmit based on the hash value of the source IP address.
- `DIP`: Transmit based on the hash value of the destination IP address.
- `DSIP`: Transmit based on the hash value of both the destination and source IP address.

Within the Uplink Profiles configuration mode, asymmetric LAG can be specified using the following CLI command:

```
5000V(config-uprof)# laghash {sip|dip|dsip|smac|dmac|dsmac|vport}
```

Note: When using asymmetric LAG mode, the hash is restricted to use only the `VPORT` or `SMAC` method to ensure that a packet with a given source MAC address always uses the same uplink.

Advanced Teaming

Each VM can have its own bias to use a subset of uplinks in an LAG. This can be accomplished by specifying a *designated uplink* for the access port to which the VM is associated. The designated uplink is configured in port alias format, such as dvUplink1, dvUplink2, etc. When the designated uplink is set to `DROP`, the VM is not permitted to use the uplink.

Packets that originate from an access port and are destined for the physical network will use one of the valid uplinks in the specified set. By default (if no designated uplink is configured) one linke will be selected from all uplinks in the LAG.

The designated uplink can be configured in port configuration mode and vNIC profile configuration mode as follows:

```
5000V(config)# iswitch vnicprof <vNIC profile name>
5000V(config-vprof)# designated-uplinks {<port alias list>|DROP}
5000V(config-vprof)# exit
```

where port alias list is a set of port alias names (such as dvsuplink1), with individual aliases separated with a comma, or ranges separated with a dash. For example:

```
5000V(config-vprof)# designated-uplinks dvsuplink1,dvsuplink4
    -or-
5000V(config-vprof)# designated-uplinks dvsuplink1-dvsuplink3
```

Port Information Commands

The following controller CLI commands are useful for collecting information about port status and configuration.

Table 3. CLI Port Information Commands

Command	Function
`show iswitch hosts`	Show information for all hosts.
`show iswitch ports`	Show information for all ports.
`show interface port` *<#>* **`info`**	Show information for a specific access port.
`show iswitch uplink host` *<IP address>*	Show information for all uplink ports associated with a specific host.
`show iswitch lacp host` *<IP address>*	Show LACP information for the uplink ports associated with a specific host.

Part 2: Switch Features

Chapter 4. Securing Administration

Secure switch management is needed for environments that perform significant management functions across the Internet. Common functions for secured management are described in the following sections:

- "Secure Shell and Secure Copy" on page 45
- "End User Access Control" on page 50

Note: SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network.

Secure Shell and Secure Copy

Because using Telnet does not provide a secure connection for managing a 5000V, Secure Shell (SSH) and Secure Copy (SCP) features have been included for 5000V management. SSH and SCP use secure tunnels to encrypt and secure messages between a remote administrator and the switch.

SSH is a protocol that enables remote administrators to log securely into the 5000V over a network to execute management commands.

SCP is typically used to copy files securely from one machine to another. SCP uses SSH for encryption of data on the network. On a 5000V, SCP is used to download and upload the switch configuration via secure channels.

Although SSH and SCP are disabled by default, enabling and using these features provides the following benefits:

- Identifying the administrator using Name/Password
- Authentication of remote administrators
- Authorization of remote administrators
- Determining the permitted actions and customizing service for individual administrators
- Encryption of management messages
- Encrypting messages between the remote administrator and switch
- Secure copy support

Software Defined Network for Virtual Envirmonments implements the SSH version 2.0 standard and is confirmed to work with SSH version 2.0-compliant clients such as the following:

- OpenSSH_5.4p1 for Linux
- Secure CRT Version 5.0.2 (build 1021)
- Putty SSH release 0.60

Configuring SSH/SCP Features on the Switch

SSH and SCP features are disabled by default. To change the SSH/SCP settings, using the following procedures.

To Enable or Disable the SSH Feature

Begin a Telnet session from the console port and enter the following commands:

```
5000V(config)# [no] ssh enable
```

To Enable or Disable SCP Apply and Save

Enter the following commands from the switch CLI to enable the SCP `putcfg_apply` and `putcfg_apply_save` commands:

```
5000V(config)# [no] ssh scp-enable
```

Configuring the SCP Administrator Password

To configure the SCP-only administrator password, enter the following command (the default password is `admin`):

```
5000V(config)# [no] ssh scp-password
Changing SCP-only Administrator password; validation required...
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```

Using SSH and SCP Client Commands

This section shows the format for using some client commands. The following examples use 205.178.15.157 as the IP address of a sample switch.

To Log In to the Switch

Syntax:

```
>> ssh [-4|-6] <switch IP address>
    -or-
>> ssh [-4|-6] <login name>@<switch IP address>
```

Note: The `-4` option (the default) specifies that an IPv4 switch address will be used. The `-6` option specifies IPv6.

Example**:**

```
>> ssh scpadmin@205.178.15.157
```

To Copy the Switch Configuration File to the SCP Host

Syntax**:**

```
>> scp [-4|-6] <username>@<switch IP address>:getcfg <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getcfg ad4.cfg
```

To Load a Switch Configuration File from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg
```

To Apply and Save the Configuration

When loading a configuration file to the switch, the `apply` and `save` commands are still required for the configuration commands to take effect. The `apply` and `save` commands may be entered manually on the switch, or by using SCP commands.

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg_apply
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg_apply_save
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg_apply
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg_apply_save
```

- The CLI `diff` command is automatically executed at the end of `putcfg` to notify the remote client of the difference between the new and the current configurations.
- `putcfg_apply` runs the `apply` command after the `putcfg` is done.
- `putcfg_apply_save` saves the new configuration to the flash after `putcfg_apply` is done.
- The `putcfg_apply` and `putcfg_apply_save` commands are provided because extra `apply` and `save` commands are usually required after a `putcfg`; however, an SCP session is not in an interactive mode.

To Copy the Switch Image and Boot Files to the SCP Host

Syntax**:**

```
>> scp [-4|-6] <username>@<switch IP address>:getimg1 <local filename>
>> scp [-4|-6] <username>@<switch IP address>:getimg2 <local filename>
>> scp [-4|-6] <username>@<switch IP address>:getboot <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getimg1 6.1.0_os.img
```

To Load Switch Configuration Files from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putimg1
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putimg2
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putboot
```

Example:

```
>> scp 6.1.0_os.img scpadmin@205.178.15.157:putimg1
```

SSH and SCP Encryption of Management Messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication:Client RSA authenticates the switch at the beginning of every connection
- Key Exchange: RSA
- Encryption:3DES-CBC, DES
- User Authentication:Local password authentication, RADIUS, SecurID (via RADIUS or TACACS+ for SSH only—does not apply to SCP)

Generating RSA Host Key for SSH Access

To support the SSH host feature, an RSA host key is required. The host key is 1024 bits and is used to identify the 5000V.

To configure RSA host key, first connect to the 5000V through the console port (commands are not available via external Telnet connection), and enter the following command to generate it manually.

```
5000V(config)# ssh generate-host-key
```

When the switch reboots, it will retrieve the host key from the FLASH memory.

Note: The switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is performing key generation at that time. Also, key generation will fail if an SSH/SCP client is logging in at that time.

SSH/SCP Integration with Radius Authentication

SSH/SCP is integrated with RADIUS authentication. After the RADIUS server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS servers for authentication. The redirection is transparent to the SSH clients.

SSH/SCP Integration with TACACS+ Authentication

SSH/SCP is integrated with TACACS+ authentication. After the TACACS+ server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

SecurID Support

SSH/SCP can also work with SecurID, a token card-based authentication method. The use of SecurID requires the interactive mode during login, which is not provided by the SSH connection.

Note: There is no SNMP or Browser-Based Interface (BBI) support for SecurID because the SecurID server, ACE, is a one-time password authentication and requires an interactive session.

Using SecurID with SSH

Using SecurID with SSH involves the following tasks.

- To log in using SSH, use a special username, “ace,” to bypass the SSH authentication.
- After an SSH connection is established, you are prompted to enter the username and password (the SecurID authentication is being performed now).
- Provide your username and the token in your SecurID card as a regular Telnet user.

Using SecurID with SCP

Using SecurID with SCP can be accomplished in two ways:

- Using a RADIUS server to store an administrator password.

  You can configure a regular administrator with a fixed password in the RADIUS server if it can be supported. A regular administrator with a fixed password in the RADIUS server can perform both SSH and SCP with no additional authentication required.
- Using an SCP-only administrator password.

  Set the SCP-only administrator password (**`ssh scp-password`**) to bypass checking SecurID.

  An SCP-only administrator’s password is typically used when SecurID is not used. For example, it can be used in an automation program (in which the tokens of SecurID are not available) to back up (download) the switch configurations each day.

Note: The SCP-only administrator’s password must be different from the regular administrator’s password. If the two passwords are the same, the administrator using that password will not be allowed to log in as an SSH user because the switch will recognize him as the SCP-only administrator. The switch will only allow the administrator access to SCP commands.

End User Access Control

IBM DS 5000V allows an administrator to define end user accounts that permit end users to perform operation tasks via the switch CLI commands. Once end user accounts are configured and enabled, the switch requires username/password authentication.

For example, an administrator can assign a user, who can then log into the switch and perform operational commands (effective only until the next switch reboot).

Considerations for Configuring End User Accounts

Note the following considerations when you configure end user accounts:

- A maximum of 10 user IDs are supported on the switch.
- 5000V supports end user support for console, Telnet, BBI, and SSHv2 access to the switch.
- If RADIUS authentication is used, the user password on the Radius server will override the user password on the 5000V. Also note that the password change command only modifies only the user password on the switch and has no effect on the user password on the Radius server. Radius authentication and user password cannot be used concurrently to access the switch.
- Passwords for end users can be up to 128 characters in length for TACACS, RADIUS, Telnet, SSH, Console, and Web access.

User Access Control

The end-user access control commands allow you to configure end-user accounts.

Setting up User IDs

Up to 10 user IDs can be configured. Use the following commands to define any user name and set the user password at the resulting prompts:

```
5000V(config)# access user 1 name <1-8 characters>
5000V(config)# access user 1 password

Changing user1 password; validation required:
Enter current admin password: <current administrator password>
Enter new user1 password: <new user password>
Re-enter new user1 password: <new user password>
New user1 password accepted.
```

Defining a User’s Access Level

The end user is by default assigned to the user access level (also known as class of service, or COS). COS for all user accounts have global access to all resources except for User COS, which has access to view only resources that the user owns. For more information, see Table 5 on page 65.

To change the user’s level, select one of the following options:

```
5000V(config)# access user 1 level {user|operator|administrator}
```

Validating a User’s Configuration

```
5000V# show access user uid 1
```

Enabling or Disabling a User

An end user account must be enabled before the switch recognizes and permits login under the account. Once enabled, the switch requires any user to enter both username and password.

```
5000V(config)# [no] access user 1 enable
```

Listing Current Users

The following command displays defined user accounts and whether or not each user is currently logged into the switch.

```
5000V# show access user

Usernames:
  user     - Enabled - offline
  oper     - Disabled - offline
  admin    - Always Enabled - online 1 session

Current User ID table:
1: name jane   , ena, cos user    , password valid, online 1 session
2: name john   , ena, cos user    , password valid, online 2 sessions
```

Logging into an End User Account

Once an end user account is configured and enabled, the user can login to the switch using the username/password combination. The level of switch access is determined by the COS established for the end user account.

Chapter 5. VLANs

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs commonly are used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments. The following topics are discussed in this chapter:

- "VLANs and Port VLAN ID Numbers" on page 54
- "VLAN Tagging" on page 54
- "VLAN Topologies and Design Considerations" on page 58
  This section discusses how you can connect users and segments to a host that supports many logical segments or subnets by using the flexibility of the multiple VLAN system.
- "Private VLANs" on page 61

Note: VLANs can be configured from the Command Line Interface (see "VLAN Configuration" in the *Command Reference*).

VLANs Overview

Setting up virtual LANs (VLANs) is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

The IBM DS 5000V (5000V) supports jumbo frames with a Maximum Transmission Unit (MTU) of 9,216 bytes. Within each frame, 18 bytes are reserved for the Ethernet header and CRC trailer. The remaining space in the frame (up to 9,198 bytes) comprise the packet, which includes the payload of up to 9,000 bytes and any additional overhead, such as 802.1q or VLAN tags. On the access ports, jumbo frame support is automatic: it is enabled by default, requires no manual configuration, and cannot be manually disabled. However, on the uplink ports, the default MTU is 1,500 bytes, though this may be configured in the uplink profiles (see "iSwitch Uplink Profile Configuration" on page 216).

VLANs and Port VLAN ID Numbers

VLAN Numbers

The 5000V supports up to 4094 VLANs per switch, numbered 1 through 4094. VLAN 1 is the default VLAN for all 5000V ports.

Use the following command to view VLAN information:

```
5000V# show vlan

VLAN              Name                 Status          Ports
----  ------------------------------  ------  ------------------------
1     VLAN 1                           ena      6-4000
2     VLAN 2                           ena      1-5
```

PVID Numbers

Each port in the switch has a configurable default VLAN number, known as its *PVID*. By default, the PVID for all non-management ports is set to 1, which correlates to the default VLAN ID. The PVID for each port can be configured to any VLAN number between 1 and 4094.

Use the following command to view PVIDs:

```
5000V# show interface information

Alias  Port   Tag  Edge  Lrn  Fld  PVID  NAME             VLAN(s)
-----  ----   ---  ----  ---  ---  ----  -------------    --------

1      1      n          d    d    2     1                2
2      2      n          d    d    2     2                2
3      3      n          d    d    2     3                2
4      4      n          d    d    2     4                2
5      5      n          d    d    2     5                2
6      6      n          d    d    1     6                1
...
# = PVID is tagged.
```

Use the following command to set the port PVID:

```
5000V(config)# interface port <port number>
5000V(config-if)# pvid <PVID number>
```

Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN *tagging* enabled (see "VLAN Tagging" on page 54).

VLAN Tagging

The IBM DS 5000V software supports 802.1Q VLAN *tagging,* providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header of a packet, allowing each port to belong to multiple VLANs. When you add a port to multiple VLANs, you must enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID =3) assigns all untagged frames received on this port to VLAN 3. Any untagged frames received by the switch are classified with the PVID of the receiving port.
- Tagged frame—a frame that carries VLAN tagging information in the header. This VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the switch through a port that is configured as a tagged port.
- Untagged frame— a frame that does not carry any VLAN tagging information in the frame header.
- Untagged member—a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).

Note: If a 802.1Q tagged frame is received by a port that has VLAN-tagging disabled and the port VLAN ID (PVID) is different than the VLAN ID of the packet, then the frame is dropped at the ingress port.

Figure 3. Default VLAN settings

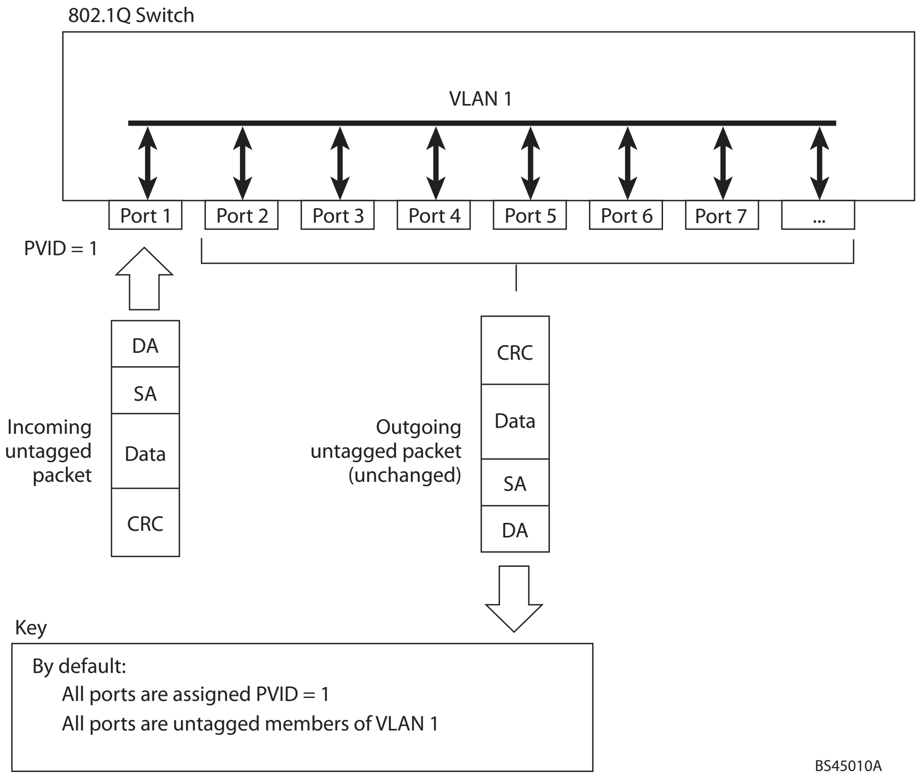

Note: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your switch model.

When a VLAN is configured, ports are added as members of the VLAN, and the ports are defined as either *tagged* or *untagged* (see Figure 4 through Figure 7).

The default configuration settings for the 5000V has all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in Figure 3, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID =1).

Figure 4 through Figure 7 illustrate generic examples of VLAN tagging. In Figure 4, untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Note: The port assignments in the following figures are not meant to match the 5000V.

Figure 4. Port-based VLAN assignment

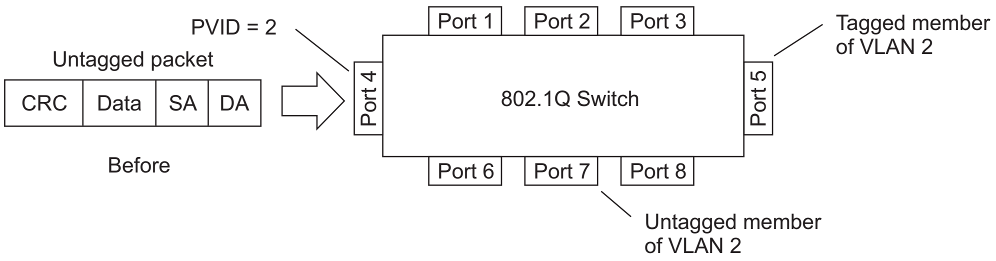

As shown in Figure 5, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 5. 802.1Q tagging (after port-based VLAN assignment)

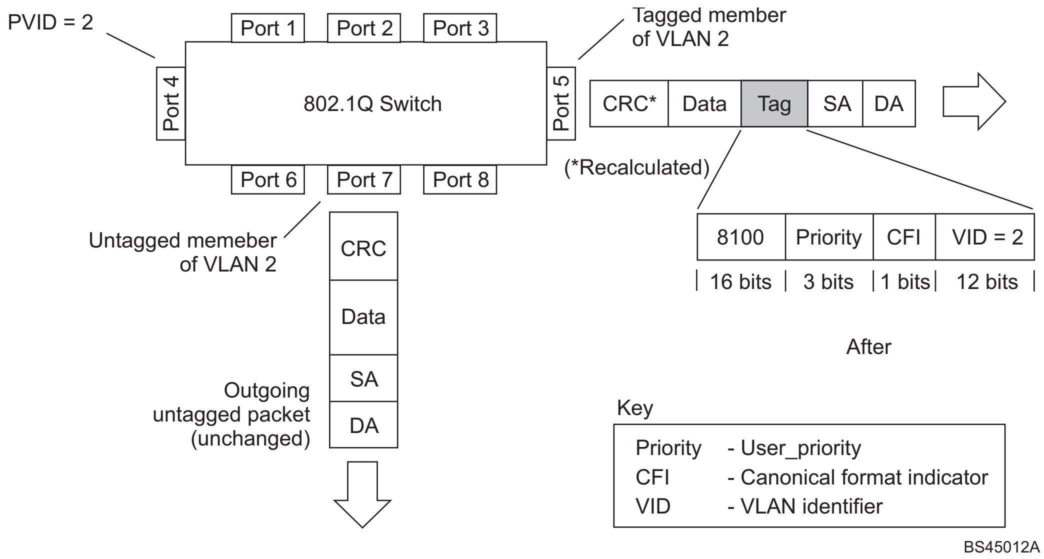

In Figure 6, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 6. 802.1Q tag assignment

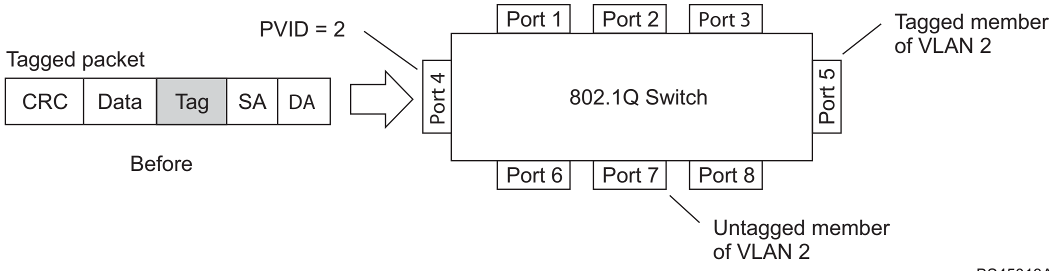

As shown in Figure 7, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 7. 802.1Q tagging (after 802.1Q tag assignment)

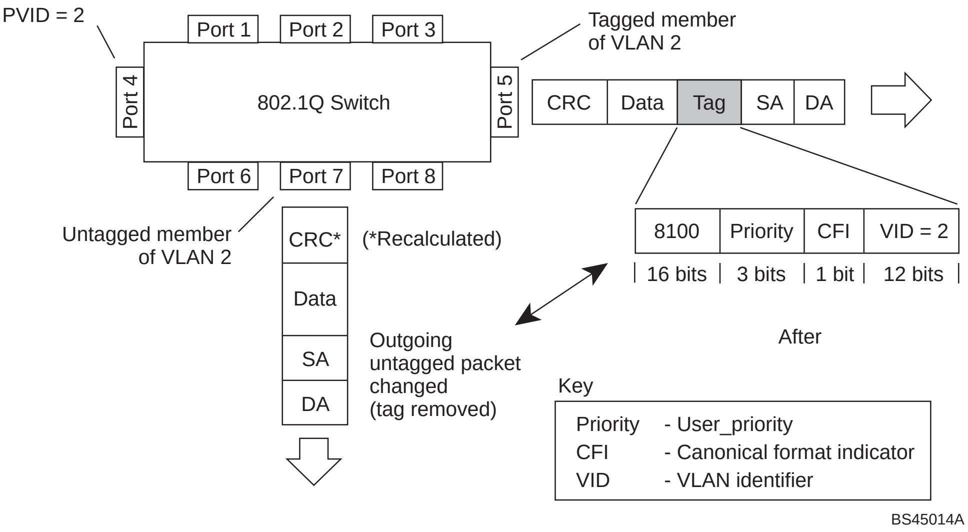

VLAN Topologies and Design Considerations

Note the following when working with VLAN topologies:

- By default, the 5000V software is configured so that tagging is disabled on all ports.
- By default, the 5000V software is configured so that all data ports are members of VLAN 1.
- All ports involved in both trunking and port mirroring must have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed. For more information trunk groups, see "Uplink Port Management" on page 39 and "Packet Monitoring" on page 103.

Multiple VLANs with Tagging Adapters

Figure 8 illustrates a network topology described in Note: on page 60 and the configuration example on page page 60.

Figure 8. Multiple VLANs with VLAN-Tagged Gigabit Adapters

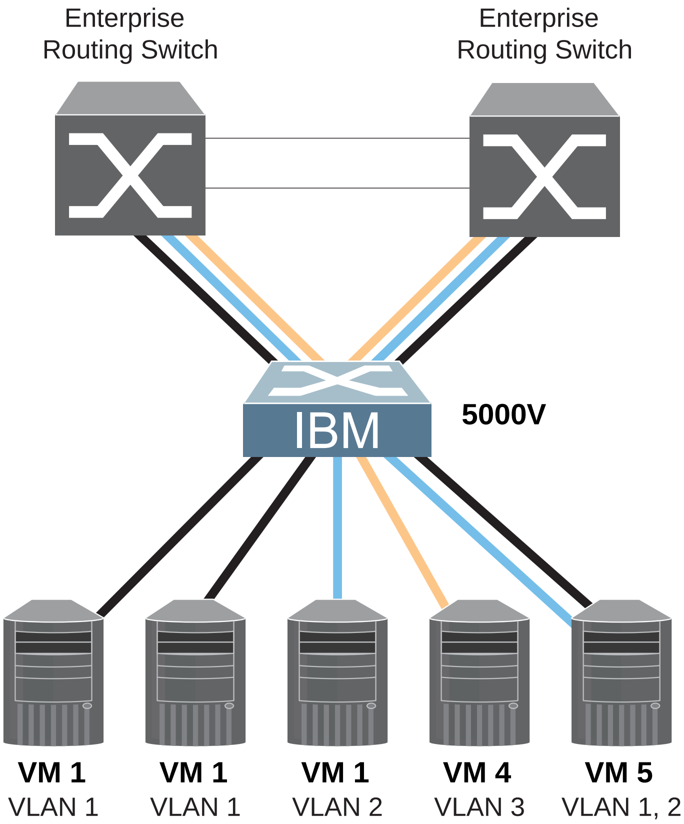

The features of this VLAN are described in the following table.

Table 4. Multiple VLANs Example

Component	Description
5000V switch	This switch is configured with three VLANs that represent three different IP subnets. Five ports are connected downstream to VMs. Two ports are connected upstream to routing switches. Uplink ports are members of all three VLANs, with VLAN tagging enabled.
VM 1	This VM is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging is disabled.
VM 2	This VM is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging is disabled.
VM 3	This VM belongs to VLAN 2, and it is logically in the same IP subnet as VM 5. The associated switch port has tagging disabled.

Table 4. Multiple VLANs Example (continued)

Component	Description
VM 4	A member of VLAN 3, this VM can communicate only with other VMs via a router. The associated switch port has tagging disabled.
VM 5	A member of VLAN 1 and VLAN 2, this VM can communicate only with VM 1, VM 2, and VM 3. The associated switch port has tagging enabled.
Enterprise Routing switches	These switches must have all three VLANs (VLAN 1, 2, 3) configured. They can communicate with VM 1, VM 2, and VM 5 via VLAN 1. They can communicate with VM 3 and VM 5 via VLAN 2. They can communicate with VM 4 via VLAN 3. Tagging on switch ports is enabled.

Note: VLAN tagging is required only on ports that are connected to other switches or on ports that connect to tag-capable end-stations, such as VMs with VLAN-tagging enabled.

VLAN Configuration Example

Use the following procedure to configure the example network shown in Figure 8 on page 59.

1. Enable VLAN tagging on access ports that support multiple VLANs.

```
5000V(config)# interface port 5
5000V(config-if)# tagging
5000V(config-if)# exit
```

2. Enable tagging on uplink ports that support multiple VLANs.

```
5000V(config)# interface port 19
5000V(config-if)# tagging
5000V(config-if)# exit
5000V(config)# interface port 20
5000V(config-if)# tagging
5000V(config-if)# exit
```

3. Configure the VLANs and their member ports.

```
5000V(config)# vlan 2
5000V(config-vlan)# enable
5000V(config-vlan)# member 3
5000V(config-vlan)# member 5
5000V(config-vlan)# member 19
5000V(config-vlan)# member 20
5000V(config-vlan)# exit
5000V(config)# vlan 3
5000V(config-vlan)# enable
5000V(config-vlan)# member 4,19,20
5000V(config-vlan)# exit
```

By default, all ports are members of VLAN 1, so configure only those ports that belong to other VLANs.

Private VLANs

Private VLANs provide Layer 2 isolation between the ports within the same broadcast domain. Private VLANs can control traffic within a VLAN domain, and provide port-based security for host servers.

Use Private VLANs to partition a VLAN domain into sub-domains. Each sub-domain is comprised of one primary VLAN and one or more secondary VLANs, as follows:

- Primary VLAN—carries unidirectional traffic downstream from promiscuous ports. Each Private VLAN configuration has only one primary VLAN. All ports in the Private VLAN are members of the primary VLAN.
- Secondary VLAN—Secondary VLANs are internal to a private VLAN domain, and are defined as follows:
  - Isolated VLAN—carries unidirectional traffic upstream from the host servers toward ports in the primary VLAN and the gateway. Each Private VLAN configuration can contain only one isolated VLAN.
  - Community VLAN—carries upstream traffic from ports in the community VLAN to other ports in the same community, and to ports in the primary VLAN and the gateway. Each Private VLAN configuration can contain multiple community VLANs.

After you define the primary VLAN and one or more secondary VLANs, you map the secondary VLAN(s) to the primary VLAN.

Private VLAN Ports

Private VLAN ports are defined as follows:

- Promiscuous—A promiscuous port is a port that belongs to the primary VLAN. The promiscuous port can communicate with all the interfaces, including ports in the secondary VLANs (Isolated VLAN and Community VLANs). Each promiscuous port can belong to only one Private VLAN.
- Isolated—An isolated port is a host port that belongs to an isolated VLAN. Each isolated port has complete layer 2 separation from other ports within the same private VLAN (including other isolated ports), except for the promiscuous ports.
  - Traffic sent to an isolated port is blocked by the Private VLAN, except the traffic from promiscuous ports.
  - Traffic received from an isolated port is forwarded only to promiscuous ports.
- Community—A community port is a host port that belongs to a community VLAN. Community ports can communicate with other ports in the same community VLAN, and with promiscuous ports. These interfaces are isolated at layer 2 from all other interfaces in other communities and from isolated ports within the Private VLAN.

Configuration Guidelines

The following guidelines apply when configuring Private VLANs:

- The default VLAN 1 cannot be a Private VLAN.
- Each secondary port's (isolated port and community ports) PVID must match its corresponding secondary VLAN ID.
- Ports within a secondary VLAN cannot be members of other VLANs.
- Each uplink ports profile support all VLANs. Restricting VLANs on uplink profiles is not allowed.

- Access ports can belong to a single Private VLAN (Primary or Secondary) only.
- The VLAN selection for access ports within a port profile must be done through the Access Port Profiles settings. All ports in a port profile will have the same VLAN settings.

Configuration Example

Follow this procedure to configure a Private VLAN.

1. Select a VLAN and define the Private VLAN type as primary.

```
5000V(config)# vlan 200
5000V(config-vlan)# enable
5000V(config-vlan)# member 201-220
5000V(config-vlan)# private-vlan type primary
5000V(config-vlan)# private-vlan enable
5000V(config-vlan)# exit
```

2. Configure a secondary VLAN and map it to the primary VLAN.

```
5000V(config)# vlan 201
5000V(config-vlan)# enable
5000V(config-vlan)# member 3
5000V(config-vlan)# member 4
5000V(config-vlan)# private-vlan type isolated
5000V(config-vlan)# private-vlan map 200
5000V(config-vlan)# private-vlan enable
5000V(config-vlan)# exit
```

3. Verify the configuration.

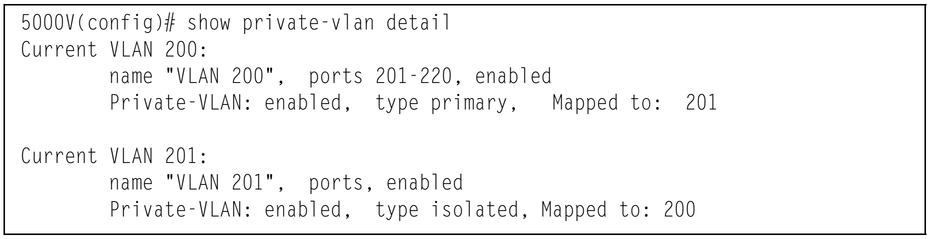

```
5000V(config)# show private-vlan detail
Current VLAN 200:
        name "VLAN 200",  ports 201-220, enabled
        Private-VLAN: enabled,  type primary,   Mapped to:  201

Current VLAN 201:
        name "VLAN 201",  ports, enabled
        Private-VLAN: enabled,  type isolated, Mapped to: 200
```

Chapter 6. Authentication & Authorization Protocols

Secure switch management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured IPv4 management and device access:

- “RADIUS Authentication and Authorization” on page 63
- “TACACS+ Authentication” on page 66

Note: Software Defined Network for Virtual Envirmonments 1.1 does not support IPv6 for RADIUS or TACACS+

RADIUS Authentication and Authorization

IBM DS 5000V supports the RADIUS (Remote Authentication Dial-in User Service) method to authenticate and authorize remote administrators for managing the switch. This method is based on a client/server model. The Remote Access Server (RAS)—the switch—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes UDP over IP (based on RFC 2138 and 2866)
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

The 5000V—acting as the RADIUS client—communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

How RADIUS Authentication Works

The RADIUS authentication process follows these steps:

1. A remote administrator connects to the switch and provides a user name and password.
2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the switch to grant or deny administrative access.

Configuring RADIUS on the Switch

Use the following procedure to configure Radius authentication on your switch.

1. Configure the IPv4 addresses of the Primary and Secondary RADIUS servers, and enable RADIUS authentication.

```
5000V(config)# radius-server primary-host 10.10.1.1
5000V(config)# radius-server secondary-host 10.10.1.2
5000V(config)# radius-server enable
```

Note: You can use a configured loopback address as the source address so the RADIUS server accepts requests only from the expected loopback address block. Use the following command to specify the loopback interface:
`5000V(config)# ip radius source-interface loopback <1-5>`

2. Configure the RADIUS secret.

```
5000V(config)# radius-server primary-host 10.10.1.1 key
      <1-32 character secret>
5000V(config)# radius-server secondary-host 10.10.1.2 key
      <1-32 character secret>
```

3. If desired, you may change the default UDP port number used to listen to RADIUS.

   The well-known port for RADIUS is 1812.

```
5000V(config)# radius-server port <UDP port number>
```

4. Configure the number retry attempts for contacting the RADIUS server, and the timeout period.

```
5000V(config)# radius-server retransmit 3
5000V(config)# radius-server timeout 5
```

RADIUS Authentication Features in IBM DS 5000V

5000V supports the following RADIUS authentication features:

- Supports RADIUS client on the switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows RADIUS secret password up to 32 bytes and less than 16 octets.
- Supports *secondary authentication server* so that when the primary authentication server is unreachable, the switch can send client authentication requests to the secondary authentication server. Use the following command to show the currently active RADIUS authentication server:

```
5000V# show radius-server
```

- Supports user-configurable RADIUS server retry and time-out values:
  - Time-out value = 1-10 seconds
  - Retries = 1-3

  The switch will time out if it does not receive a response from the RADIUS server in 1-3 retries. The switch will also automatically retry connecting to the RADIUS server before it declares the server down.
- Supports user-configurable RADIUS application port. The default is 1812/UDP-based on RFC 2138. Port 1645 is also supported.
- Supports user-configurable RADIUS application port. The default is UDP port 1645. UDP port 1812, based on RFC 2138, is also supported.
- Allows network administrator to define privileges for one or more specific users to access the switch at the RADIUS user database.

Switch User Accounts

The user accounts listed in Table 5 can be defined in the RADIUS server dictionary file.

Table 5. User Access Levels

User Account	Description and Tasks Performed	Password
User	The User has no direct responsibility for switch management. They can view all switch status information and statistics but cannot make any configuration changes to the switch.	`user`
Operator	The Operator manages all functions of the switch. The Operator can reset ports, except the management port.	`oper`
Administrator	The super-user Administrator has complete access to all commands, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	`admin`

RADIUS Attributes for IBM DS 5000V User Privileges

When the user logs in, the switch authenticates his/her level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the remote user is successfully authenticated by the authentication server, the switch will verify the *privileges* of the remote user and authorize the appropriate access. The administrator has an option to allow *secure backdoor* access via Telnet/SSH/BBI. Secure backdoor provides switch access when the RADIUS servers cannot be reached. You always can access the switch via the console port, by using `noradius` and the administrator password, whether secure backdoor is enabled or not.

Note: To obtain the RADIUS backdoor password for your 5000V, contact Technical Support.

All user privileges, other than those assigned to the Administrator, have to be defined in the RADIUS dictionary. RADIUS attribute 6 which is built into all RADIUS servers defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The following RADIUS attributes are defined for 5000V user privileges levels:

Table 6. IBM DS 5000V-proprietary Attributes for RADIUS

User Name/Access	User-Service-Type	Value
User	*Vendor-supplied*	255
Operator	*Vendor-supplied*	252
Admin	*Vendor-supplied*	6

TACACS+ Authentication

5000V supports authentication and authorization with networks using the Cisco Systems TACACS+ protocol. The 5000V functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the 5000V either through a data port or management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP-based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization and accounting.

How TACACS+ Authentication Works

TACACS+ works much in the same way as RADIUS authentication as described on page 63.

1. Remote administrator connects to the switch and provides user name and password.
2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using TACACS+ protocol, the authentication server instructs the switch to grant or deny administrative access.

During a session, if additional authorization checking is needed, the switch checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

TACACS+ Authentication Features in IBM DS 5000V

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. 5000V supports ASCII inbound login to the device. PAP, CHAP and ARAP login methods, TACACS+ change password requests, and one-time password authentication are not supported.

Authorization

Authorization is the action of determining a user’s privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization levels and 5000V management access levels is shown in Table 7. The authorization levels must be defined on the TACACS+ server.

Table 7. Default TACACS+ Authorization Levels

5000V User Access Level	TACACS+ level
`user`	0
`oper`	3
`admin`	6

Alternate mapping between TACACS+ authorization levels and 5000V management access levels is shown in Table 8. Use the following command to set the alternate TACACS+ authorization levels.

```
5000V(config)# tacacs-server privilege-mapping
```

Table 8. Alternate TACACS+ Authorization Levels

5000V User Access Level	TACACS+ level
`user`	0 - 1
`oper`	6 - 8
`admin`	14 - 15

If the remote user is successfully authenticated by the authentication server, the switch verifies the *privileges* of the remote user and authorizes the appropriate access. The administrator has an option to allow *secure backdoor* access via Telnet/SSH. Secure backdoor provides switch access when the TACACS+ servers cannot be reached. You always can access the switch via the console port, by using `notacacs` and the administrator password, whether secure backdoor is enabled or not.

Note: To obtain the TACACS+ backdoor password for your 5000V, contact Technical Support.

Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization is not performed via TACACS+, there are no TACACS+ accounting messages sent out.

You can use TACACS+ to record and track software login access, configuration changes, and interactive commands.

The 5000V supports the following TACACS+ accounting attributes:

- protocol (console/Telnet/SSH/HTTP/HTTPS)
- start_time
- stop_time
- elapsed_time
- disc_cause

Note: When using the Browser-Based Interface, the TACACS+ Accounting Stop records are sent only if the **Logout** button on the browser is clicked.

Command Authorization and Logging

When TACACS+ Command Authorization is enabled, 5000V configuration commands are sent to the TACACS+ server for authorization. Use the following command to enable TACACS+ Command Authorization:

```
5000V(config)# tacacs-server command-authorization
```

When TACACS+ Command Logging is enabled, 5000V configuration commands are logged on the TACACS+ server. Use the following command to enable TACACS+ Command Logging:

```
5000V(config)# tacacs-server command-logging
```

The following examples illustrate the format of 5000V commands sent to the TACACS+ server:

```
authorization request, cmd=shell, cmd-arg=interface ip
accounting request, cmd=shell, cmd-arg=interface ip
authorization request, cmd=shell, cmd-arg=enable
accounting request, cmd=shell, cmd-arg=enable
```

Configuring TACACS+ Authentication on the Switch

1. Configure the IPv4 addresses of the Primary and Secondary TACACS+ servers, and enable TACACS authentication. Specify the interface port (optional).

```
5000V(config)# tacacs-server primary-host 10.10.1.1
5000V(config)# tacacs-server primary-host mgtb-port
5000V(config)# tacacs-server secondary-host 10.10.1.2
5000V(config)# tacacs-server secondary-host data-port
5000V(config)# tacacs-server enable
```

Note: You can use a configured loopback address as the source address so the TACACS+ server accepts requests only from the expected loopback address block. Use the following command to specify the loopback interface:
`5000V(config)#` **`ip tacacs source-interface loopback`** *<1-5>*

2. Configure the TACACS+ secret and second secret.

```
5000V(config)# tacacs-server primary-host 10.10.1.1 key
<1-32 character secret>
5000V(config)# tacacs-server secondary-host 10.10.1.2 key
    <1-32 character secret>
```

3. If desired, you may change the default TCP port number used to listen to TACACS+.

   The well-known port for TACACS+ is 49.

```
5000V(config)# tacacs-server port <TCP port number>
```

4. Configure the number of retry attempts, and the timeout period.

```
5000V(config)# tacacs-server retransmit 3
5000V(config)# tacacs-server timeout 5
```

Chapter 7. Access Control Lists

Access Control Lists (ACLs) are filters that permit or deny traffic for security purposes. They can also be used with QoS to classify and segment traffic to provide different levels of service to different traffic types. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

IBM System Networking Distributed Switch 5000V 1.1 supports the following ACLs:

- IPv4 ACLs

  Up to 127 ACLs are supported for networks that use IPv4 addressing. IPv4 ACLs are configured using the following CLI command path:

  ```
  5000V(config)# access-list ip <128 - 254> {extended|standard}
  ```

- MAC ACLs

  Up to 127 ACLs are supported for filtering packets based on the MAC address. MAC ACLs are configured using the following CLI command path:

  ```
  5000V(config)# access-list mac extended <1-127>
  ```

Summary of Packet Classifiers

ACLs allow you to classify packets according to a variety of content in the packet header (such as the source address, destination address, protocol type, and others). Once classified, packet flows can be identified for more processing.

- IPv4 header options (for IPv4 ACLs only)
  - Source IPv4 address and subnet mask
  - Destination IPv4 address and subnet mask
  - IP protocol number or name as shown in Table 9:

Table 9. Well-Known Protocol Types

Number	Protocol Name
1	icmp
4	ipv4
6	tcp
17	udp
89	ospf
103	pim

MAC ACLs allow you to classify packets based on the following packet attributes:

- Ethernet header options (for MAC ACLs only)
  - Source MAC address
  - Destination MAC address
  - VLAN number
  - Ethernet type (ARP, IP, RARP)
  - Ethernet Priority (the IEEE 802.1p Priority)

- TCP/UDP header options (for extended ACLs only)
  - TCP/UDP application source port and mask as shown in Table 10
  - TCP/UDP application destination port as shown in Table 10

Table 10. Well-Known Application Ports

Port	TCP/UDP Application	Port	TCP/UDP Application	Port	TCP/UDP Application
20	ftp-data	79	finger	179	bgp
21	ftp	80	http	194	irc
22	ssh	109	pop2	220	imap3
23	telnet	110	pop3	389	ldap
25	smtp	111	sunrpc	443	https
37	time	119	nntp	520	rip
42	name	123	ntp	554	rtsp
43	whois	143	imap	1645/1812	Radius
53	domain	144	news	1813	Radius Accounting
69	tftp	161	snmp	1985	hsrp
70	gopher	162	snmptrap		

  - TCP/UDP flag value as shown in Table 11

Table 11. Well-Known TCP flag values

Flag	Value
URG	0x0020
ACK	0x0010
PSH	0x0008
RST	0x0004
SYN	0x0002
FIN	0x0001

Summary of ACL Actions

Once classified using ACLs, the identified packet flows can be processed differently. For each ACL, an *action* can be assigned. The action determines how the switch treats packets that match the classifiers assigned to the ACL. 5000V ACL packet actions include Pass or Drop.

Assigning Individual ACLs to a Port

Once you configure an ACL, you must assign the ACL to the appropriate ports. Each port can accept multiple ACLs, and each ACL can be applied for multiple ports. ACLs can be assigned individually.

To assign an individual IP ACL to a port, use the following IP Interface Mode commands:

```
5000V(config)# interface port <port>
5000V(config-if)# ip access-group <IPv4 ACL number> in
```

To assign an individual MAC ACL to a port, use the following IP Interface Mode commands:

```
5000V(config)# interface port <port>
5000V(config-if)# mac access-group <MAC ACL number> in
```

When multiple ACLs are assigned to a port, higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs. ACL order of precedence is discussed in the next section.

ACL Order of Precedence

When multiple ACLs are assigned to a port, they are evaluated in numeric sequence, based on the ACL number. Lower-numbered ACLs take precedence over higher-numbered ACLs. For example, ACL 1 (if assigned to the port) is evaluated first and has top priority. MAC ACLs take precedence over IP ACLs.

If multiple ACLs match the port traffic, only the one with the lowest ACL number is applied. The others are ignored.

If no assigned ACL matches the port traffic, the default action (permit) is applied.

ACL Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the 5000V by configuring a QoS meter (if desired) and assigning ACLs to ports.

Note: When you add ACLs to a port, make sure they are ordered correctly in terms of precedence.

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- **In-Profile**If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile**If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kbps (in multiples of 64 Mbps). All traffic within this Committed Rate is In-Profile. Additionally, you can set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level that traffic receives.
- Change the 802.1p priority of a packet.

Viewing ACL Statistics

ACL statistics display how many packets have matched each ACL. Use ACL statistics to check filter performance or to debug the ACL filter configuration.

You must enable statistics for each IP or MAC ACL that you wish to monitor:

```
5000V(config)# access-list ip <128 - 254> {extended|standard} statistics

5000V(config)# access-list mac extended <1-127> statistics
```

ACL Configuration Examples

ACL Example 1

Use this configuration to block traffic to a specific host. All traffic that ingresses on port 1 is denied if it is destined for the host at IP address 100.10.1.1

1. Configure an Access Control List.

```
5000V(config)# access-list ip 128 standard
5000V(config-std-nacl)# deny any host 100.10.1.1
```

2. Add ACL 128 to port 1.

```
5000V(config)# interface port 1
5000V(config-if)# ip access-group 128 in
5000V(config-if)# exit
```

ACL Example 2

Use this configuration to block traffic from a network destined for a specific host address. All traffic that ingresses in port 2 with source IP from class 100.10.1.0/24 and destination IP 200.20.2.2 is denied.

1. Configure an Access Control List.

```
5000V(config)# access-list ip 130 standard
5000V(config-std-nacl)# deny 100.10.1.0 255.255.255.0 host 200.20.2.2
```

2. Add ACL 130 to port 2.

```
5000V(config)# interface port 2
5000V(config-if)# ip access-group 130 in
5000V(config-if)# exit
```

ACL Example 3

Use this configuration to deny all ARP packets that ingress a port.

1. Configure an Access Control List.

```
5000V(config)# access-list mac extended 1
5000V(config-ext-macl)# deny any any arp
```

2. Add ACL 1 to port 2.

```
5000V(config)# interface port 2
5000V(config-if)# mac access-group 1 in
5000V(config-if)# exit
```

Chapter 8. Quality of Service

Quality of Service (QoS) features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to factors such as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate QoS level.

The following topics are discussed in this section:

QoS Overview

QoS helps you allocate guaranteed bandwidth to the critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

Software Defined Network for Virtual Envirmonments QoS functionality includes packet classification and traffic conditioning. Packets are classified based on the content in the packet header. Traffic conditioning includes metering, policing, and/or re-marking to ensure that traffic entering the IBM DS 5000V domain conforms to the traffic classifier rules, traffic profiles, metering, marking, and discarding rules that are applied to the traffic.

Figure 9 shows the basic functionality of the traffic classifier and conditioner.

Figure 9. Logical View of a Packet Classifier and Traffic Conditioner

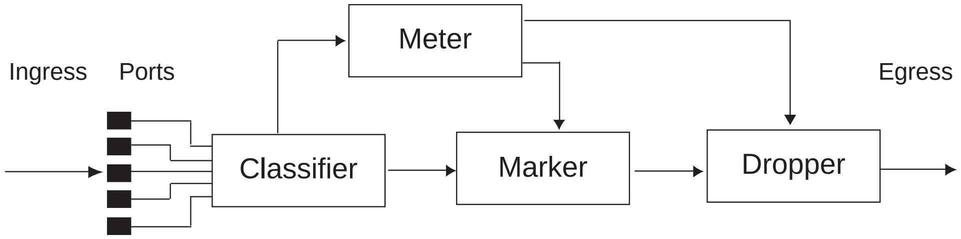

The basic QoS model works as follows:

- Classify traffic:
  – Read DiffServ Code Point (DSCP) value.
  – Read 802.1p priority value.
  – Match ACL filter parameters.

- Perform actions:
  – Define bandwidth and burst parameters
  – Select actions to perform on in-profile and out-of-profile traffic
  – Deny packets
  – Permit packets
  – Mark DSCP or 802.1p Priority

Using CoS/ACL/DSCP Filters

Using Class of Service (CoS), ACLs, and DSCP values, you can classify and segment traffic to provide different levels of service to different traffic types. The three classifiers work as follows:

- CoS: Classifies packets based on the 802.1p priority value.
- DSCP: Classifies packets based on the DSCP value.
- ACLs: Classify packets based on a combination of one or more header fields such as source address, destination address, protocol ID, and port number.

Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

Software Defined Network for Virtual Envirmonments 1.1 supports up to 254 ACLs. For ACL details, see "Access Control Lists" on page 71.

Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the 5000V by configuring a QoS meter (if desired). Actions are performed based on the configured meter.

Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile, which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic, as follows:

- **In-Profile**–If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile**–If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kbps. All traffic within this Committed Rate is In-Profile. Additionally, you set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter and perform actions on metered traffic, such as packet re-marking.

Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level traffic receives.
- Change the 802.1p priority of a packet.

The following example includes steps to configure a meter and out-of-profile DSCP remarking:

1. Create an ACL.

```
5000V(config)# access-list mac extended 1
5000V(config-ext-macl)# permit any any user-priority 5
5000V(config-ext-macl)# exit
```

2. Create a class map.

```
5000V(config)# class-map 1
5000V(config-cmap)# match access-group mac 1
5000V(config-cmap)# exit
```

3. Create a policy map.

```
5000V(config)# policy-map 1
5000V(config-pmap)# class 1
```

4. Configure the meter with a committed rate and set out-of-profile action to re-mark DSCP value based on the global DSCP mapping. The DSCP remarking value is global for all ACLs that have DSCP re-marking enabled.

```
5000V(config-pmap-c)# police cir 10000000 be 1536 conform
                     set-dscp-transmit 40         (Assign new DSCP value)
```

Metering and Re-Marking Configuration Examples

Example 1

The following example includes the basic steps for re-marking the priority value for packets with the destination MAC address 00:50:56:B3:76:EC.

1. Define a MAC ACL.

```
5000V(config)# access-list mac extended 1
5000V(config-ext-macl)# permit any host 00:50:56:B3:76:EC
5000V(config-ext-macl)# exit
```

2. Create a Class Map.

```
5000V(config)# class-map match-any 1
5000V(config-cmap)# match access-group mac 1
5000V(config-cmap)# exit
```

3. Create a Policy Map and assign the Class Map to it.

```
5000V(config)# policy-map 1
5000V(config-pmap)# class 1
5000V(config-pmap-c)# set cos 3
5000V(config-pmap-c)# exit
5000V(config-pmap)# exit
```

Example 2

The following example includes steps to configure metering. Packets are transmitted if the data rate is within the configured meter. If there is a violation, priority of packets that have a destination MAC address 00:50:56:B3:76:EC is changed to 7.

1. Define a MAC ACL.

```
5000V(config)# access-list mac extended 1
5000V(config-ext-macl)# permit any host 00:50:56:B3:76:EC
5000V(config-ext-macl)# exit
```

2. Create a Class Map.

```
5000V(config)# class-map match-any 1
5000V(config-cmap)# match access-group mac 1
5000V(config-cmap)# exit
```

3. Create a Policy Map and assign the Class Map to it.

```
5000V(config)# policy-map 1
5000V(config-pmap)# class 1
5000V(config-pmap-c)# police cir 256 conform transmit violate
                      set-cos-transmit 7
5000V(config-pmap-c)# exit
5000V(config-pmap)# exit
```

Using 802.1p Priority to Provide QoS

The 5000V provides QoS functions based on the priority bits in a packet’s VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1Q VLAN header.) The 802.1p bits, if present in the packet, specify the priority to be given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority value.

The IEEE 802.1p standard uses eight levels of priority (0-7). Priority 7 is assigned to highest priority network traffic, such as OSPF or RIP routing table updates, priorities 5-6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a “best effort” traffic prioritization, and this is the default when traffic priority has not been configured on your network. The switch can filter packets based on the 802.1p values.

Figure 10. Layer 2 802.1q/802.1p VLAN tagged packet

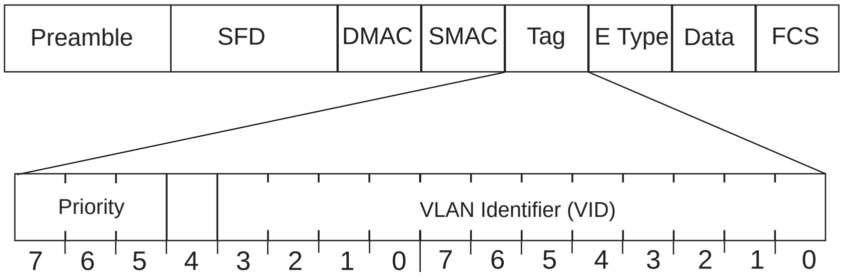

Ingress packets receive a priority value, as follows:

- **Tagged packets**—switch reads the 802.1p priority in the VLAN tag.
- **Untagged packets**—switch tags the packet and assigns an 802.1p priority value, based on the port’s default 802.1p priority.

To configure a port’s default 802.1p priority value, use the following commands.

```
5000V(config)# interface port 1
5000V(config-if)# dot1p <802.1p value (0-7)>
5000V(config-if)# exit
```

Chapter 9. Edge Virtual Bridging

The 802.1Qbg/Edge Virtual Bridging (EVB) is an emerging IEEE standard for allowing networks to become virtual machine (VM)-aware. EVB bridges the gap between physical and virtual network resources. The IEEE 802.1Qbg simplifies network management by providing a standards-based protocol that defines how virtual Ethernet bridges exchange configuration information. In EVB environments, virtual NIC (vNIC) configuration information is available to EVB devices. This information is generally not available to an 802.1Q bridge.

The EVB features on the 5000V are compliant with the IEEE 802.1Qbg Authors Group Draft 0.2. For a list of documents on this feature, see: http://www.ieee802.org/1/pages/802.1bg.html.

The 5000V implementation of EVB supports the following protocols:

- Virtual Ethernet Bridging (VEB) and Virtual Ethernet Port Aggregator (VEPA): VEB and VEPA are mechanisms for switching between VMs on the same hypervisor. VEB enables switching with the server, either in the software (virtual switch), or in the hardware (using single root I/O virtualization capable NICs). VEPA requires the edge switch to support “Reflective Relay”— an operation where the switch forwards a frame back to the port on which it arrived if the destination MAC address is on the same port.
- Edge Control Protocol (ECP): ECP provides reliable delivery of service dispatcher units (SDUs) between the station and bridge, and between the port extender and bridge.
- Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP): VDP allows hypervisors to advertise VSIs to the physical network. This protocol also allows centralized configuration of network policies that will persist with the VM, independent of its location.
- EVB Type-Length-Value (TLV): EVB TLV is a component of Link Layer Discovery protocol (LLDP)-based TLV used to discover and configure VEPA, ECP, and VDP.

EVB Operations Overview

The 5000V includes a pre-standards VSI Type Database (VSIDB) implemented through the IBM System Networking Distributed Switch 5000V. The VSIDB is the central repository for defining sets of network policies that apply to VM network ports. You can configure only one VSIDB.

The 5000V includes the following features that enable EVB for VM environments:

- VSI Type database: The 5000V controller includes an embedded VSIDB that can be used as the central repository for VM network policies.
- VDP support: The 5000V enables VDP support for a hypervisor such that the hypervisor can proactively exchange information about VM networking requirements with a physical switch.
- VEPA mode: The 5000V can be enabled for VEPA mode either at an individual VM port level or at a port group level. This provides the flexibility to use it both as an host-internal distributed switch for certain VMs as well as a VEPA for others. The 5000V supports VDP for VMs connected to VEPA-enabled ports.

Note: While the 5000V enables EVB support at the hypervisor level, you also need a physical switch layer that supports EVB capabilities. The VSIDB exports the database to the physical switch when it receives a request. The hypervisor sends a VSI ASSOCIATE, which contains the VSI type ID, to the physical switch port after the VM is started. The physical switch updates its configuration based on the requested VSI type. The physical switch configures the per-VM bandwidth using the VM policies.

EVB Configuration

Prerequisites

Before configuring EVB, ensure you have completed the following tasks:

- Install the 5000V controller appliance on the host. The controller provides the management interface for the associated distributed switch.

  Note: You do not need to install the virtual appliance on a host that will have uplinks to the 5000V.
- Configure the controller. This includes changing the administrator password, setting the IP address, enabling remote management (SSH/telnet), installation of licenses, and configuring vCenter credentials in the controller.
- Create an instance of the 5000V in the target vSphere datacenter.
- Install the 5000V vSphere Installation Bundle (VIB).
- Assign the hosts and uplinks to the 5000V.

Note: For uplinks, EVB is supported only in uplink profiles configured for asymmetric mode.

Configuration Tasks

Following is an overview of the primary components required for the configuration of EVB:

- vib: A 5000V component that comes packaged as a vSphere installation bundle and requires installation on each ESX host.
- evbprof: An EVB profile is configured on the physical switch to enable features such as Reflective Relay as well as to enable the VDP protocol on the physical switch ports connecting to the hypervisor.
- vnicprof: A 5000V port group that allows network policies and parameters to be set on a group of virtual ports. This includes enabling VEPA mode within the hypervisor as well as specifying VSI type ID at the port group level.

For configuring EVB, you need the perform the following sequence of tasks:

On the 5000V:

1. Create network policy definitions in the VSI database to cater to VM networking requirements.
2. Configure VLAN-centric port groups and/or standalone ports with a VSI type ID and native VLAN ID. Enable VEPA mode for the port group or standalone ports.

On the physical switch:

3. Configure the physical switch to:
   a. Query the VSI type database using the IP address and document path.
   b. Enable LLDP support.
   c. Create and assign EVB profiles to enable reflective relay and VDP support on the physical ports connected to the 5000V uplink adapters.

On the 5000V:

4. Connect VM network ports to standalone ports or port groups on the 5000V that have the appropriate VLAN and VSI type ID for each VM.

Following values are used as examples in the configuration steps:

- Type ID: 11
- Version: 1
- Profile Name: profile1
- VLAN: 1100
- vNIC Profile Name: vnicprof1

Task 1: Create Network Policy Definitions

Use the following commands to create a network profile (VSI type ID):

```
5000V(config)# vsiman
5000V(config-vsiman)# typeid 11 version 1
5000V(typeid-version)# name profile1
5000V(typeid-version)# vlans 1100
5000V(typeid-version)# exit
5000V(config-vsiman)# exit
```

Note: You can assign additional parameters such as ACLs and QoS, and extend ACLs to mark packet priority (see “VSI Type ID Configuration” on page 218).

Task 2: Configure Port Groups and Enable VEPA

Use the following commands to create VLAN-centric port groups:

1. Create a VLAN.

```
5000V(config)# vlan 1100
   Feb  1 2012 21:47:44 5000V:VLAN-INFO: INFO  mgmt: VLAN number 1100
   with name
   "VLAN 1100" created
5000V(config-vlan)# enable
5000V(config-vlan)# exit
```

Note: The 5000V will tag all outgoing traffic with the pvid of the source port/port group that the traffic originated from. For this reason the pvid configured for each port group must match a VLAN specified in the corresponding VSI type definition. The physical switch uses this information to validate the association between a VM and its requested VSI type ID. If the VLAN tag differs from the list in the VSI type definition, the physical switch will deny association for that VM and the virtual NIC port will show as being in a “blocked” state.

2. Create port groups and enable VEPA.

```
5000V(config)# iswitch vnicprof vnicprof1
    Ports Allocated: 101-120
5000V(config-vprof)# pvid 1100
    Nov 4 2011 06:29:10 5000V:VLAN-INFO: Ports 101-120 are UNTAGGED
    Ports and their PVIDs are changed to 1100
5000V(config-vprof)# vsitype 11 version 1
5000V(config-vprof)# vepa
5000V(config-vprof)# exit
```

Task 3: Configure Physical Switch

Note: The following configuration steps are for a physical switch and are not for the 5000V. These have been provided for informational purpose. An IBM RackSwitch G8264 is used as an example.

This section includes the steps to configure EVB based on the following values:

– Profile number: 1

– Port number: 1

– Retry interval: 8000 milliseconds

– VSI Database:

  - Manager IP: 172.31.37.187
  - Port: 80

1. Create an EVB profile.

```
>> Main# /cfg/virt/evb/profile 1                     (Enter number from 1-16)
```

2. Enable Reflective Relay.

```
>> EVB Profile 1# rr enable
```

3. Enable VSI discovery.

```
>> EVB Profile 1# vsidisc enable
```

4. Add EVB profile to port.

```
>> EVB Profile 1# /cfg/port 1/evbprof 1        (Enter EVB profile ID (0-16))
```

Note: On some switches, this port should be a server port (`>>Main # /cfg/sys/srvports/add` *<port number>*)

5. Configure ECP retransmission interval.

```
>> Port 1# /cfg/l2/ecp/retrans 8000
                        (Enter retransmission interval in milliseconds (100-9000)
```

6. Set VSI database information.

```
>> Edge Control Protocol Configuration# /cfg/virt/evb/vsidb 1
>> VSI Type DB 1# managrip 172.31.37.187  (Set VSI database Manager IP)
>> VSI Type DB 1# port 80                 (Set VSI database Manager port)
>> VSI Type DB 1# docpath “vsidb”         (Set VSI database document path)
>> VSI Type DB 1# alltypes “all.xml”      (Set VSI database document path)
>> VSI Type DB 1# interval 30            (Set update interval in seconds)
```

7. Enable LLDP.

```
>> VSI Type DB 1# /cfg/l2/lldp/on                   (Turn on LLDP)
```

8. Disable VMready.

```
>> LLDP Configuration# /cfg/virt/disvmr            (Disable VMready)
```

Task 4: Connect VMs to VEPA-enabled ports

You can now connect virtual machines to the VM-aware network.

Limitations

- If both ACL and egress bandwidth metering are enabled, traffic will first be matched with the ACL and will not be limited by bandwidth metering.
- ACLs based on a source MAC or VLAN must match the source MAC and VLAN of the VM. If not, the policy will be ignored and you will see the following warning message:

```
"vm: VSI Type ID 100 Associated mac 00:50:56:b6:c0:ff on port 6,
ignore 1 mismatched ACL"
```

- A new policy cannot be applied immediately after the following three actions take place:
  - A VM gets associated.
  - The VSI database in the XML server is updated.
  - The local VSI database is updated.

  You need to re-associate the VM in order to apply the policy.

Unsupported features

The following features are not supported:

- S-channel and Channel Discovery and Configuration Protocol (CDCP)
- LAG/VLAG
- VMReady
- VNIC
- FCoE
- Stacking

Chapter 10. Simple Network Management Protocol

Software Defined Network for Virtual Envirmonments provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software, such as IBM Director or HP-OpenView.

Note: SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network.

SNMP Version 1 & Version 2

To access the SNMP agent on the 5000V, the read and write community strings on the SNMP manager must be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands on the CLI:

```
5000V(config)# snmp-server read-community <1-32 characters>

   -and-

5000V(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager must be able to reach the management interface or any one of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch with the following command:

```
5000V(config)# snmp-server trap-src-if <trap source IP interface>
5000V(config)# snmp-server host <IPv4 address> <trap host community string>
```

Note: You can use a loopback interface to set the source IP address for SNMP traps. Use the following command to apply a configured loopback interface:
`5000V(config)# snmp-server trap-source loopback <1-5>`

SNMP Version 3

SNMP version 3 (SNMPv3) is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMPv3 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators and encryption to protect against threats such as masquerade, modification of information, message stream modification and disclosure.

SNMPv3 allows clients to query the MIBs securely.

SNMPv3 configuration is managed using the following command path menu:

```
5000V(config)# snmp-server ?
```

For more information on SNMP MIBs see the command reference material starting at “System SNMP Configuration” on page 180.

Default Configuration

IBM DS 5000V has two SNMPv3 users by default. Both of the following users have access to all the MIBs supported by the switch:

- User 1 name is `adminmd5` (password `adminmd5`). Authentication used is MD5.
- User 2 name is `adminsha` (password `adminsha`). Authentication used is SHA.

Up to 16 SNMP users can be configured on the switch. To modify an SNMP user, enter the following commands:

```
5000V(config)# snmp-server user <1-16> name <1-32 characters>
```

Users can be configured to use the authentication/privacy options. The 5000V support two authentication algorithms: MD5 and SHA, as specified in the following command:

```
5000V(config)# snmp-server user <1-16> authentication-protocol
    {md5|sha} authentication-password
   -or-
5000V(config)# snmp-server user <1-16> authentication-protocol none
```

User Configuration Example

1. To configure a user with name “admin,” authentication type MD5, and authentication password of “admin,” privacy option DES with privacy password of “admin,” use the following CLI commands.

```
5000V(config)# snmp-server user 5 name admin
5000V(config)# snmp-server user 5 authentication-protocol md5
    authentication-password
Changing authentication password; validation required:
Enter current admin password:          <admin. password>
Enter new authentication password:     <auth. password>
Re-enter new authentication password:  <auth. password>
New authentication password accepted.

5000V(config)# snmp-server user 5 privacy-protocol des
    privacy-password
Changing privacy password; validation required:
Enter current admin password:          <admin. password>
Enter new privacy password:            <privacy password>
Re-enter new privacy password:         <privacy password>
New privacy password accepted.
```

2. Configure a user access group, along with the views the group may access. Use the access table to configure the group’s access level.

```
5000V(config)# snmp-server access 5 name admingrp
5000V(config)# snmp-server access 5 level authpriv
5000V(config)# snmp-server access 5 read-view iso
5000V(config)# snmp-server access 5 write-view iso
5000V(config)# snmp-server access 5 notify-view iso
```

   Because the read view, write view, and notify view are all set to “iso,” the user type has access to all private and public MIBs.

3. Assign the user to the user group. Use the group table to link the user to a particular access group.

```
5000V(config)# snmp-server group 5 user-name admin
5000V(config)# snmp-server group 5 group-name admingrp
```

Configuring SNMP Trap Hosts

SNMPv1 Trap Host

1. Configure a user with no authentication and password.

```
5000V(config)# snmp-server user 10 name “v1trap”
```

2. Configure an access group and group table entries for the user. Use the following menu to specify which traps can be received by the user:

```
5000V(config)# snmp-server access <access group>
```

In the following example the user will receive the traps sent by the switch. First, assign an access group to view SNMPv1 traps:

```
5000V(config)# snmp-server access 10 name v1trap
5000V(config)# snmp-server access 10 security snmpv1
5000V(config)# snmp-server access 10 notify-view iso
```

Next, assign the user to the access group:

```
5000V(config)# snmp-server group 10 security snmpv1
5000V(config)# snmp-server group 10 user-name v1trap
5000V(config)# snmp-server group 10 group-name v1trap
```

3. Configure an entry in the notify table.

```
5000V(config)# snmp-server notify 10 name v1trap
5000V(config)# snmp-server notify 10 tag v1trap
```

4. Specify the IPv4 address and other trap parameters in the `targetAddr` and `targetParam` tables. Use the following commands to specify the user name associated with the targetParam table:

```
5000V(config)# snmp-server target-address 10 name v1trap address
    10.70.70.190
5000V(config)# snmp-server target-address 10 parameters-name v1param
5000V(config)# snmp-server target-address 10 taglist v1param
5000V(config)# snmp-server target-parameters 10 name v1param
5000V(config)# snmp-server target-parameters 10 user-name v1only
5000V(config)# snmp-server target-parameters 10 message snmpv1
```

Note: 5000V 1.1 supports only IPv4 addresses for SNMP trap hosts.

5. Use the community table to specify which community string is used in the trap.

```
5000V(config)# snmp-server community 10 index v1trap
5000V(config)# snmp-server community 10 name public
5000V(config)# snmp-server community 10 user-name v1trap
```

SNMPv2 Trap Host Configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model, use `snmpv2` instead of `snmpv1`.

```
5000V(config)# snmp-server user 10 name v2trap

5000V(config)# snmp-server group 10 security snmpv2
5000V(config)# snmp-server group 10 user-name v2trap
5000V(config)# snmp-server group 10 group-name v2trap
5000V(config)# snmp-server access 10 name v2trap
5000V(config)# snmp-server access 10 security snmpv2
5000V(config)# snmp-server access 10 notify-view iso

5000V(config)# snmp-server notify 10 name v2trap
5000V(config)# snmp-server notify 10 tag v2trap

5000V(config)# snmp-server target-address 10 name v2trap
     address 100.10.2.1
5000V(config)# snmp-server target-address 10 taglist v2trap
5000V(config)# snmp-server target-address 10 parameters-name
     v2param
5000V(config)# snmp-server target-parameters 10 name v2param
5000V(config)# snmp-server target-parameters 10 message snmpv2c
5000V(config)# snmp-server target-parameters 10 user-name v2trap
5000V(config)# snmp-server target-parameters 10 security snmpv2

5000V(config)# snmp-server community 10 index v2trap
5000V(config)# snmp-server community 10 user-name v2trap
```

Note: 5000V 1.1 supports only IPv4 addresses for SNMP trap hosts.

SNMPv3 Trap Host Configuration

To configure a user for SNMPv3 traps, you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

This is configured in the access table using the following commands:

```
5000V(config)# snmp-server access <1-32> level
5000V(config)# snmp-server target-parameters <1-16>
```

Configure the user in the user table accordingly.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user `v3trap` with authentication only:

```
5000V(config)# snmp-server user 11 name v3trap
5000V(config)# snmp-server user 11 authentication-protocol md5
    authentication-password
Changing authentication password; validation required:
Enter current admin password:         <admin. password>
Enter new authentication password:    <auth. password>
Re-enter new authentication password: <auth. password>
New authentication password accepted.
5000V(config)# snmp-server access 11 notify-view iso
5000V(config)# snmp-server access 11 level authnopriv
5000V(config)# snmp-server group 11 user-name v3trap
5000V(config)# snmp-server group 11 tag v3trap
5000V(config)# snmp-server notify 11 name v3trap
5000V(config)# snmp-server notify 11 tag v3trap
5000V(config)# snmp-server target-address 11 name v3trap address
    47.81.25.66
5000V(config)# snmp-server target-address 11 taglist v3trap
5000V(config)# snmp-server target-address 11 parameters-name v3param
5000V(config)# snmp-server target-parameters 11 name v3param
5000V(config)# snmp-server target-parameters 11 user-name v3trap
5000V(config)# snmp-server target-parameters 11 level authNoPriv
```

Note: 5000V 1.1 supports only IPv4 addresses for SNMP trap hosts.

Chapter 11. sFlow

The IBM DS 5000V supports sFlow version 5 technology for monitoring data networks. The embedded sFlow agent can be configured to provide continuous monitoring in the form of random packet sampling and time-based sampling of statistical counters for IPv4 traffic.

Note: IBM DS 5000V 1.1 does not support IPv6 for sFlow.

The 5000V is responsible only for forwarding sFlow information. One or more separate sFlow collectors (or analyzers) are required elsewhere on the network to interpret sFlow data.

The 5000V provides a global sampling engine and up to 31 additional sampling engines which can be customized to monitor specific ports and/or VLANs. Each sampling engine has independent sampling rates, counter poll intervals, and can be directed to different sFlow collectors.

Further, each ESX host associated with the 5000V vDS is an sFlow sub-agent, and each ESX host also maintains independent sample-rate counters.

Enabling sFlow

To enable the sFlow feature, use the following CLI configuration commands:

```
5000V(config)# sflow
5000V(config-sflow)# enable
5000V(config-sflow)# agent-ip <agent IP address>
```

The *agent IP address* represents the 5000V to the sFlow collectors and analyzers. Set the IP address of the controller management interface. You can find this address using the `show interface ip-mgmt` command.

Although enabled, actual sampling will not occur until packet sampling or statical counters sampling are configured as shown in the following sections.

Note: Communication between the 5000V and target sFlow collectors uses established sFlow service port 6343. sFlow operation requires that any VMware firewalls or security features permit UDP port 6343 traffic between the VMware vCenter and 5000V controller and vDS host modules. See "Firewall Considerations" on page 101 for more information.

Global Packet Sampling

When global sampling is configured, the 5000V sFlow engine samples all packets that traverse the 5000V vDS.

Packets are sampled only if they successfully egress the 5000V switching fabric, either via an access port attached to a VM or via an uplink port. Packets that are dropped by ACLs or other features will not be sampled.

When a packet sample is taken, 128 bytes are copied, UDP-encapsulated, and sent to a configured sFlow collector.

Configuring Global Packet Sampling

Global packet sampling configuration is performed by setting the sample-rate and the IP address of the sFlow collector as follows:

```
5000V(config-sflow)# sample-rate <packet period (1-65534)>
5000V(config-sflow)# collector <sFlow IP address>
```

The sFlow global sampling rate can be configured to occur once each 1 to 65534 packets. A sampling rate of 256 means that one sample will be taken for approximately every 256 packets received.

The sampling rate is statistical. It is possible to have slightly more or fewer samples sent to the collector for any specific group of packets (especially under low traffic conditions). The actual sample rate becomes most accurate over time, and under higher traffic flow.

Note: Although sFlow sampling is not generally a CPU-intensive operation, configuring extremely fast sampling rates on ports under heavy traffic loads can cause high CPU utilization on the controller or ESX hosts. Use larger rate values of 256 or more for ports that experience heavy traffic.

Disabling Global Network Sampling

To disable global packet sampling while leaving other sFlow features operational, negate the sample-rate as follows:

```
5000V(config-sflow)# no sample-rate
```

Network Sampling Limitations

When combined with other features, sFlow sampling the following behaviors are expected:

- Packets that are dropped by ACLs or other features will not be sampled.
- sFlow sampling will not occur on packets that are duplicated during the port mirroring process. If sFlow sampling is enabled on a port that is configured as a port monitor, the mirrored traffic will not be sampled. However, the original packet may be sampled if sFlow network sampling is enabled on its original (non-monitor) port or VLAN destination.

Statistical Counters

When global counters sampling is configured, the 5000V sends information regarding network statistical counters to an sFlow collector (or analyzer) at regular, configurable intervals.

Configuring Global Counters Sampling

Global counters sampling configuration is performed by setting the poll interval and the IP address of the sFlow collector as follows:

```
5000V(config-sflow)# counter-poll <interval in seconds (20-65534)>
5000V(config-sflow)# collector <sFlow IP address>
```

Note: If the sFlow collector IP address was previously configured for packet sampling, the `collector` command can be ignored.

When the configured polling interval has elapsed, the 5000V will report general port statistics and port Ethernet statistics to the sFlow collector. In addition, each sub-agent (vDS host module) will send its own statistical counters data. Each sub-agent maintains an independent sFlow engine. A packet traversing one ESX host will not impact the sFlow counters on another ESX host in the same vDS.

Disabling Global Counters Sampling

To disable global counters sampling while leaving other sFlow features operational, negate the polling interval as follows:

```
5000V(config-sflow)# no counter-poll
```

Custom Sampling Groups

The 5000V supports up to an additional 31 sFlow sampling engines. Each can be be customized to focus packet sampling and/or counters sampling on a single port or VLAN, or set of ports or VLANs. Each sampling engine is independent of the others. Sampling engines are configured using the sFlow `group` configuration mode.

Sampling groups are numbered grouped from 1 to 31. If any port or VLAN is assigned to multiple sampling groups, the sampling group with the lowest ID number will have priority.

Configuring Sampling Groups

Sampling group configuration is performed in the sFlow Group configuration mode as follows:

```
5000V(config-sflow)# group <group number (1-31)>
5000V(config-sflow-group)# sample-rate <packet period (1-65534)>
5000V(config-sflow-group)# counter-poll <interval in seconds (20-65534)>
5000V(config-sflow-group)# collector <sFlow IP address>
5000V(config-sflow-group)# add port <port list>
5000V(config-sflow-group)# add vlan <VLAN list>
```

The sample-rate, poll interval, and collector IP address work the same as with global sampling ("Global Packet Sampling" on page 95), but apply to only the ports and VLANs specified for the group.

The port and VLAN lists can specify a single port, or it can specify multiple ports using the standard range syntax: using no extra spaces, separate list items with a comma (**,**) or specify list ranges using the dash (**-**). For example:

```
5000V(config-sflow-group)# add port 1-10,20
```

This specifies the inclusion of ports 1 through 10, and also port 20.

Adding ports and VLANs to the group is cumulative. The add commands can be used multiple times with different settings to add different items to the group.

To remove previously added ports or VLANs from the group, use the appropriate `del` configuration option:

```
5000V(config-sflow-group)# del port <port list>
5000V(config-sflow-group)# del vlan <VLAN list>
```

Use the `exit` command to leave the sFlow Group configuration mode.

Enabling or Disabling All Custom Sampling Groups

Custom sampling groups require the sFlow feature to be enabled (see "Enabling sFlow" on page 95). When sFlow is disabled, custom sampling groups are inactive. However, when sFlow is enabled, custom sampling groups are independent of the global sampling engine. Custom sampling groups can be used even when global packet or global counters sampling are disabled.

Enabling or Disabling Individual Groups

The collector IP address is required for sampling. To disable both packet sampling and counters sampling simultaneously, you can negate the collector IP address in the sFlow Group configuration mode:

```
5000V(config-sflow-group)# no collector
```

Enabling or Disabling Individual Group Functions

Each custom sampling group can be independently configured packet sampling, counters sampling, or both. The sample-rate is required only if packet sampling is desired. The polling interval is required only if counters sampling is desired. To disable either or both sampling functions, use the appropriate negation command in the sFlow Group configuration mode:

```
5000V(config-sflow-group)# no sample-rate
5000V(config-sflow-group)# no counter-poll
```

Order of Precedence

Each packet will be considered for sampling or counting no more than once. There is no duplication between the global sFlow engine or any of the customer sFlow sampling groups. Whether a packet is considered by the global engine or by one of the custom groups is based on the following priorities:

1. Custom group port matching.

   If the packet egress port matches a port assigned to a custom sampling group, the packet will participate only in that group's sampling process.

   If the egress port is included in more than one group, the group with the lowest ID is given priority. For example, group 1 has a higher priority than group 10. If the port belongs to both groups, sampling or counting of that packet will be processed according to group 1 sample-rate and counted only toward group 1 statistics.

2. Custom group VLAN matching.

   Upon switch egress, if the packet’s VLAN matches a VLAN assigned to a custom sampling group, the packet will participate only in that group’s sampling process.

   And just as with port matching, if the packet’s VLAN is included in more than one custom sampling group, the group with the lowest ID is given priority.

3. Global sFlow engine

The packet will be processed for sampling and/or statistical counting only at the level where the first match is found. If a packet matches one level but is subsequently not selected to be forwarded to the sFlow collector (based on the group’s sample-rate), it will not be considered for packet sampling in any other sFlow engine. Similarly, at that same match level, if the packet is not selected for statistical counting (such as when no polling interval is configured), it will not be counted at any other group or level.

sFlow Configuration Information

To obtain information about the current state of sFlow configuration, us the following global show command (shown with sample output):

```
5000V# show sflow
sFlow sampling is globally enabled
        Global sampling rate:           1 in 250 packets
        Global counter polling rate:    not configured
        Collector: 172.31.46.40
        Agent IP:  172.31.38.155

 Group 10: sample rate is 100, poll is 60, collector is 10.100.200.150
        ports: 10 20
        vlans: 10
5000V(config-sflow-group)# no counter-poll
```

sFlow Configuration Example

In the following example, a customer sampling group is configured. Only packets that egress port 10 or 20 (regardless of VLAN), or on VLAN 10 (regardless of port) are considered. For packets that match those criteria, 1 in 200 packets are sampled, and statistical counters are collected every two minutes.

1. Enable sFlow:

   ```
   5000V(config)# sflow
   5000V(config-sflow)# enable
   ```

2. Set the IP address the switch will use to identify itself to the sFlow collector:

   ```
   5000V(config-sflow)# agent-ip 10.100.200.10
   ```

3. Enter the sFlow Group configuration mode:

   ```
   5000V(config-sflow)# group 10
   ```

4. Set the sample-rate and polling interval

```
5000V(config-sflow-group-10)# sample-rate 200
5000V(config-sflow-group-10)# counter-poll 120
```

5. Specify the IP address of the sFlow collector or analyzer:

```
5000V(config-sflow-group-10)# collector 10.100.200.150
```

6. Add the appropriate ports and VLANs to the group:

```
5000V(config-sflow-group-10)# add port 10
5000V(config-sflow-group-10)# add port 20
5000V(config-sflow-group-10)# add vlan 10
```

7. Check the configuration information and exit the sub-modes.

```
5000V(config-sflow-group-10)# show sflow
5000V(config-sflow-group-10)# exit
5000V(config-sflow)# exit
5000V(config)#
```

Firewall Considerations

In order for the sFlow feature to work correctly, UDP traffic to destination port 6343 on the target sFlow collectors must be permitted through any VMware firewall or security features that are configured in the network.

For convenience, an example configuration is provided in the following section. However, the example may not apply in all environments. It is recommended that you defer to the appropriate VMware documented for reconfiguring security and firewall options on the ESX hypervisor.

The VMware documentation to configure the firewall can be found at the following URL:

> http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-security-guide.pdf

Example ESX hypervisor configuration

Perform the following tasks on the ESX hypervisor (login via SSH).

1. Create the `file /etc/vmware/firewall/ibm5000V.xml` with the following content:

```
<ConfigRoot>
 <service>
  <id>ibm5000VsFlow</id>
   <rule id='0000'>
    <direction>outbound</direction>
    <protocol>udp</protocol>
    <porttype>dst</porttype>
    <port>6343</port>
   </rule>
   <enabled>true</enabled>
   <required>false</required>
 </service>
</ConfigRoot>
```

2. At the ESXCLI, perform a network firewall refresh .

```
# esxcli network firewall refresh
```

3. Verify that the `ibm5000VsFlow` entry is now in the list and enabled (as shown by the final line of the sample output):

```
# esxcli network firewall ruleset list

...
remoteSerialPort      false
ibm5000VsFlow          true
```

4. Disable `allow-all` for the ruleset (unless you want to enable the firewall so that sFlow packets can go to any destination):

```
# esxcli network firewall ruleset set --allowed-all false
--ruleset-id="ibm5000VsFlow"
```

5. Enable the ruleset for the target sFlow collector’s subnet or IP address. In this example, the collector is in 172.30.0.0/24:

   ```
   # esxcli network firewall ruleset allowedip add
   --ip-address=172.30.0.0/24 --ruleset-id="ibm5000VsFlow"
   ```

   Replace 172.30.0.0/24 with an address relevant to your deployment.

   Repeat this command for each additional IP address or subnet you wish to add.

6. Verify that the ruleset is properly applied:

   ```
   # esxcli network firewall ruleset allowedip list

   ...
   remoteSerialPort    All
   ibm5000VsFlow       172.30.0.0/24
   ```

7. Configure sFlow on the 5000V as described in the other sections of this chapter.

Once full configuration is complete, sFlow packets should traverse the ESX firewall and arrive at the sFlow collector. Please consult the VMware documentation for best practices on reconfiguring that product’s firewall.

Chapter 12. Packet Monitoring

Port Mirroring

The IBM DS 5000V port mirroring feature allows you to mirror (copy) the packets of a target port, and forward them to a monitoring port. Port mirroring functions for all layer 2 and layer 3 traffic on a port. This feature can be used as a troubleshooting tool or to enhance the security of your network. For example, an IDS server or other traffic sniffer device or analyzer can be connected to the monitoring port to detect intruders attacking the network.

The 5000V supports a “many to one” mirroring model. As shown in Figure 11, selected traffic for ports 1 and 2 is being monitored by port 3. In the example, both ingress traffic and egress traffic on port 2 are copied and forwarded to the monitor. However, port 1 mirroring is configured so that only ingress traffic is copied and forwarded to the monitor. A device attached to port 3 can analyze the resulting mirrored traffic.

Figure 11. Mirroring Ports

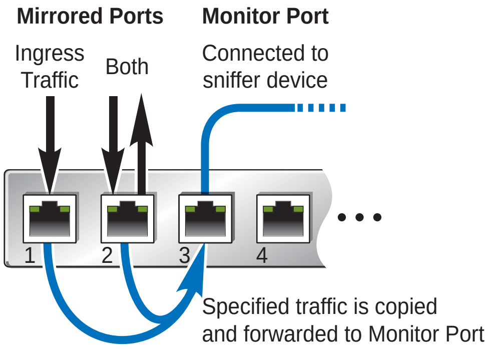
Mirrored Ports
Monitor Port
Ingress Traffic
Both
Connected to sniffer device
1
2
3
4
Specified traffic is copied and forwarded to Monitor Port

Mirroring Restrictions

The 5000V supports multiple monitor ports. Each monitor port can receive mirrored traffic from any number of target ports.

Each monitor port must be on the same ESX host as its assigned mirror ports.

The 5000V does not support “one to many” or “many to many” mirroring models where traffic from a specific port traffic is copied to multiple monitor ports. For example, port 1 traffic cannot be monitored by both port 3 and 4 at the same time, nor can port 2 ingress traffic be monitored by a different port than its egress traffic.

Ingress and egress traffic is duplicated and sent to the monitor port after processing. Any packets discarded while processing will not be mirrored.

Configuring Port Mirroring

The following procedure may be used to configure port mirroring for the example shown in Figure 11 on page 103:

1. Specify the monitoring port, the mirroring port(s), and the port-mirror direction.

```
5000V(config)# port-mirroring monitor-port 3 mirroring-port 1 in
5000V(config)# port-mirroring monitor-port 3 mirroring-port 2 both
```

2. Enable port mirroring.

```
5000V(config)# port-mirroring enable
```

3. View the current configuration.

```
5000V# show port-mirroring

Port Monitoring : Enabled

Monitoring Ports       Mirrored Ports
1                      none
2                      none
3                      (1, in) (2, both)
4                      none
5                      none
6                      none
7                      none
8                      none
9                      none
10                     none
...
```

ERSPAN

While port mirroring requires that the monitor and mirrored ports are on the same host (associated with the same vDS host module), ERSPAN allows the source and destination ports to be on different hosts.

The ERSPAN Configuration Mode

ERSPAN configuration is performed in the ERSPAN mode. To enter ERSPAN mode, use the following CLI command:

```
5000V(config)# erspan
```

Defining an ERSPAN source

ERSPAN sources are defined in the ERSPAN Source sub-mode. To enter the sub-mode, specify a the source ID:

```
5000V(config-erspan)# source <ID>
```

The following commands are available in the source sub-mode:

- **direction** {**in|out|both**}

  Specify the direction of the traffic to be monitored.

- **traffic** {**all|unicast|multicast|broadcast**}

  Specify the type of traffic to monitor.

- **mode** {**l2|l3**}

  Specify Layer 2 or Layer 3 mode.

- **vlan** *<VLAN number>*

  Specify the VLAN ID that is applied to the ERSPAN headers.

- **priority** *<0-6>*

  Specify the priority level of the ERSPAN packets.

- **add port** *<port list>*

  Specify access ports to be added to the source session.

- **add vlan** *<VLAN list>*

  Specify VLANs to be added to the source session.

- **del port** *<port list>*

  Specify access ports to be removed from the source session.

- **del vlan** *<VLAN list>*

  Specify VLANs to be removed from the source session

The following commands are also available in this sub-menu:

- **curr**

  Display current settings for this source.

- **exit**

  Return to the parent ERSPAN mode.

Defining ERSPAN Flows

Once one or more ERSPAN sources are defined, you can define ERSPAN flows using the following ERSPAN mode command:

```
5000V(config-erspan)# flow <flow ID> source <source ID> destination
   <IP address>
```

where *flow ID* is a unique value from 0 thru 1023 and the destination is the IP address of the ERSPAN collector.

Enabling ERSPAN

When the sources and flows are configured, enable ERSPAN with the ERSPAN mode command:

```
5000V(config-erspan)# enable
```

ERSPAN may be enabled at any time, but monitoring will only occur when valid sources and flows are defined, and when the associated access ports and VLANs are enabled and operations.

ERSPAN Information Commands

The following ERSPAN mode commands are available for reviewing ERSPAN settings:

- **`show`**

  Shows all ERSPAN settings

- **`show source`** [<*source ID*>]

  Show settings for a specific source (if the *source ID* option is used) or for all ERSPAN source sessions (if the *source ID* is not specified).

- **`show flow`** [<*flow ID*>]

  Show settings for a specific flow (if the *flow ID* option is used) or for all ERSPAN flows (if the *flow ID* is not specified).

Exit the ERSPAN Mode

From withing the ERSPAN mode, return to the global configuration mode using the `exit` command.

Part 3: CLI Reference

Chapter 13. CLI Basics

Your IBM DS 5000V is ready to perform basic switching functions after basic installation. Some of the more advanced features, however, require some administrative configuration before they can be used effectively.

The 5000V provides a Command-Line Interface (CLI) for collecting switch information and performing switch configuration. Using a basic terminal, the CLI allows you to view information and statistics about the switch, and to perform any necessary configuration.

This chapter explains how to use the CLI for the switch.

User Access Levels

To enable better switch management and user accountability, three levels or *classes* of user access have been implemented on the 5000V. Levels of access to CLI functions increase as needed to perform various switch management tasks. Conceptually, access classes are defined as follows:

- **User**

  Interaction with the switch is completely passive—nothing can be changed on the 5000V. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.

- **Operator**

  Operators can make temporary changes on the 5000V. These changes are lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any change an operator makes is undone by a reset of the switch, operators cannot severely impact switch operation.

- **Administrator**

  Administrators are the only ones that may make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the 5000V. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique surnames and passwords. Once you are connected to the switch via local Telnet, remote Telnet, or SSH, you are prompted to enter a password. The default user names/password for each access level are listed in the following table.

Note: It is recommended that you change default switch passwords after initial configuration and as regularly as required under your network security policies.

Table 12. User Access Levels

User Account	Password	Description and Tasks Performed
`user`	`user`	The User has no direct responsibility for switch management. He or she can view all switch status information and statistics, but cannot make any configuration changes to the switch.
`oper`	`oper`	The Operator manages all functions of the switch. The Operator can reset ports, except the management ports.
`admin`	`admin`	The superuser Administrator has complete access to all menus, information, and configuration commands on the 5000V, including the ability to change both the user and administrator passwords.

Note: With the exception of the “admin” user, access to each user level can be disabled by setting the password to an empty value.

CLI Command Modes

The CLI has three major command modes listed in order of increasing privileges, as follows:

- **User EXEC mode**

  This is the initial mode of access. By default, password checking is disabled for this mode, on console.

- **Privileged EXEC mode**

  This mode is accessed from User EXEC mode. This mode can be accessed using the following command: `enable`

- **Global Configuration mode**

  This mode allows you to make changes to the running configuration. If you save the configuration, the settings survive a reload of the 5000V. Several sub-modes can be accessed from the Global Configuration mode. For more details, see Table 13.

Each mode provides a specific set of commands. The command set of a higher-privilege mode is a superset of a lower-privilege mode—all lower-privilege mode commands are accessible when using a higher-privilege mode.

Table 13. lists the CLI command modes.

Table 13. CLI Command Modes

Command Mode/Prompt	**Command used to enter or exit**
User EXEC `5000V>`	Default mode, entered automatically on console Exit: `exit` or `logout`
Privileged EXEC `5000V#`	Enter Privileged EXEC mode, from User EXEC mode: `enable` Exit to User EXEC mode: `disable` Quit CLI: `exit` or `logout`
Global Configuration `5000V(config)#`	Enter Global Configuration mode, from Privileged EXEC mode: `configure terminal` Exit to Privileged EXEC: `end` or `exit`
Interface port `5000V(config-if)#`	Enter Port Configuration mode, from Global Configuration mode: `interface port` *<port number or alias>* Exit to Global Configuration mode: `exit` Exit to Privileged EXEC mode: `end`
VLAN `5000V(config-vlan)#`	Enter VLAN Configuration mode, from Global Configuration mode: `vlan` *<VLAN number>* Exit to Global Configuration mode: `exit` Exit to Privileged EXEC mode: `end`
ERSPAN `5000V(config-erspan)#`	Enter ERSPAN Configuration mode, from Global Configuration mode: `erspan` Exit to Global Configuration mode: `exit`
sFlow `5000V(config-sflow)#`	Enter sFlow Configuration mode, from Global Configuration mode: `sflow` Exit to Global Configuration mode: `exit`
Access Control List (ACL) IP Standard access list `5000V(config-std-nacl)#`	Enter ACL IP Standard Configuration mode, from Global Configuration mode: `access-list ip` *<128 - 254>* `standard` Exit to Global Configuration mode: `exit` Exit to Privileged EXEC mode: `end`

Table 13. CLI Command Modes (continued)

Command Mode/Prompt	**Command used to enter or exit**
ACL IP Extended access list `5000V(config-ext-nacl)#`	Enter ACL IP Extended Configuration mode, from Global Configuration mode: `access-list ip` *<128 - 254>* `extended` Exit to Global Configuration mode: `exit` Exit to Privileged EXEC mode: `end`
ACL MAC access list `5000V(config-ext-macl)#`	Enter ACL MAC Extended Configuration mode, from Global Configuration mode: `access-list mac extended` *<1 - 127>* Exit to Global Configuration mode: `exit` Exit to Privileged EXEC mode: `end`
iSwitch vNIC Profile `5000V(config-vprof)#`	Enter vNIC Profile Configuration mode, from Global Configuration mode: `iswitch vnicprof` *<vNIC Profile Name>* Exit to Global Configuration mode: `exit`
iSwitch Uplink Profile `5000V(config-uprof)#`	Enter vNIC Profile Configuration mode, from Global Configuration mode: `iswitch uprof` *<Uplink Profile Name>* Exit to Global Configuration mode: `exit`
Class Map `5000V(config-cmap)#`	Enter Class Map Configuration mode, from Global Configuration mode: `class-map` *<Class Map Number>* Exit to Global Configuration mode: `exit`

Global Commands

Some basic commands are recognized throughout the CLI command modes. These commands are useful for obtaining online help, navigating through the interface, and for saving configuration changes.

For help on a specific command, type the command, followed by `help`.

Table 14. Description of Global Commands

Command	**Action**
`?`	When entered at the prompt, this command lists the commands that are available in the current mode. When placed after a full or partial command, it provides more information about a specific command and its available parameters.
`exit`	Go up one level in the command mode structure. If already at the top level, exit from the command line interface and log out.

Table 14. Description of Global Commands

Command	**Action**
`copy running-config startup-config`	Write configuration changes to non-volatile flash memory.
`logout`	Exit from the command line interface and log out.
`ping`	Use this command to verify station-to-station connectivity across the network. The format is as follows: `ping` *<host name>* | *<IP address>* [*<number of packets (1-32)>*] [*<delay in seconds (1-100)*]
`prompting`	Enables command prompting.
`show history`	This command displays the last ten issued commands.
`show who`	Displays a list of users who are currently logged in.
`terminal-length`	Sets the length of the terminal. Valid range is 0-300 lines per screen.
`traceroute`	Use this command to identify the route used for station-to-station connectivity across the network. The format is as follows: `traceroute` {*<hostname>* | *<IP address>*} [*<max-hops (1-32)>*] [*<timeout in seconds (1-100)>*] Where *hostname/IP address* is the hostname or IP address of the target station, *max-hops* (optional) is the maximum distance to trace (1-32 devices), and *timeout in seconds* (optional) is the number of milliseconds to wait for the response. As with `ping`, the DNS parameters must be configured if specifying hostnames.

Negation

To turn off, disable, clear, delete, or otherwise undo a command or one or more of its options or parameters, use negation. This is accomplished by placing `no` at the start of the command. For example:

```
5000V(config)# ssh enable             Turns SSH access on.
5000V(config)# no ssh enable          Turns SSH access off.
```

Commands that permit negation are listed in the command reference section with [`no`] or `no` as the first option.

CLI Shortcuts

The following shortcuts allow you to enter commands quickly and easily.

List and Range Inputs

For VLAN and port commands that allow an individual item to be selected from within a numeric range, lists and ranges of items can now be specified. For example, the `show vlan` command permits the following options:

```
# show vlan 1,3,4094                          (view VLANs 1, 3, and 4094)
# show vlan 1-20                              (access VLANs 1 through 20)
# show vlan 1-5,90-99,4090-4094               (access multiple ranges)
# show vlan 1-5,19,20,4090-4094               (access a mix of lists and ranges)
```

The numbers in a range must be separated by a dash: *<start of range>* - *<end of range>*

Multiple ranges or list items are permitted using a comma: *<range or item 1>* , *<range or item 2>*

Do not use spaces within list and range specifications.

Ranges can also be used to apply the same command option to multiple items. For example, to access multiple ports with one command:

```
# interface port 1-4                          (Access ports 1 though 4)
```

Note: Port ranges accept only port numbers, not port aliases.

Command Abbreviation

Most commands can be abbreviated by entering the first characters which distinguish the command from the others in the same mode. For example, consider the following full command and a valid abbreviation:

```
5000V(config)# vlan 1
    or
5000V(config)# vl 1
```

Idle Timeout

By default, the switch will disconnect your Telnet session after five minutes of inactivity. This function is controlled by the following command, which can be set from 1 to 60 minutes:

`system idle` *<1-60>*

Command mode: Global Configuration

Chapter 14. Information Commands

You can view configuration information for the switch in both the user and administrator command modes. This chapter discusses how to use the command line interface to display switch information.

Table 15. Information Commands

Command Syntax and Usage
`show information-dump` Dumps all switch information available (10K or more, depending on your configuration). If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands. **Command mode:** All

System Information

The information provided by each command option is briefly described in Table 16, with pointers to where detailed information can be found.

Table 16. System Information Options

Command Syntax and Usage
`show sys-info` Displays system information, including: – System date and time – Switch model name and number – Switch name and location – Time of last boot – MAC address of the switch management processor – IP address of management interface – Software image file and version number – Configuration name – Log-in banner, if one is configured For details, see page 125. **Command mode:** All
`show system daylight` Displays daylight saving time setting. For details, see page 125. **Command mode:** All
`show system timezone` Displays the current time zone. For details, see page 125. **Command mode:** All

Table 16. System Information Options (continued)

Command Syntax and Usage
`show system acknowledgement` Displays the open source acknowledgements. For details, see page 126. **Command mode:** Privilege EXEC
`show system cpu` Displays system CPU utilization. For details, see page 127. **Command mode:** Privilege EXEC
`show system memory` Displays system memory. For details, see page 128. **Command mode:** Privilege EXEC
`show clock` Displays the current date and time. For details, see page 129. **Command mode:** All
`show logging [messages]` Displays most recent syslog messages. (Optional) Specify the messages command for view last 100 syslog messages. For details, see page 129. **Command mode:** All
`show access user` Displays configured user names and their status. For details, see page 129. **Command mode:** All

SNMPv3 System Information

SNMP version 3 (SNMPv3) is an extensible SNMP Framework that supplements the SNMPv2 framework by supporting the following:

- a new SNMP message format
- security for messages
- access control
- remote configuration of SNMP parameters

For more details on the SNMPv3 architecture please refer to RFC2271 to RFC2276.

Table 17. SNMPv3 Information Options

Command Syntax and Usage
`show snmp-server v3 user` Displays User Security Model (USM) table information. To view the table, see page 118. **Command mode:** All
`show snmp-server v3 view` Displays information about view, subtrees, mask and type of view. To view a sample, see page 118. **Command mode:** All
`show snmp-server v3 access` Displays View-based Access Control information. To view a sample, see page 119. **Command mode:** All
`show snmp-server v3 group` Displays information about the group, including the security model, user name, and group name. To view a sample, see page 120. **Command mode:** All
`show snmp-server v3 community` Displays information about the community table information. To view a sample, see page 121. **Command mode:** All
`show snmp-server v3 target-address [`*<1-16>*`]` Displays the Target Address table information. To view a sample, see page 121. **Command mode:** All
`show snmp-server v3 target-parameters [`*<1-16>*`]` Displays the Target parameters table information. To view a sample, see page 122. **Command mode:** All
`show snmp-server v3 notify [`*<1-16>*`]` Displays the Notify table information. To view a sample, see page 122. **Command mode:** All
`show snmp-server v3` Displays all the SNMPv3 information. To view a sample, see page 124. **Command mode:** All

SNMPv3 USM User Table Information

The User-based Security Model (USM) in SNMPv3 provides security services such as authentication and privacy of messages. This security model makes use of a defined set of user identities displayed in the USM user table. The following command displays SNMPv3 user information:

```
show snmp-server v3 user <1-16>
```

Command mode: All

The USM user table contains the following information:

- the user name
- a security name in the form of a string whose format is independent of the Security Model
- an authentication protocol, which is an indication that the messages sent on behalf of the user can be authenticated
- the privacy protocol

```
5000V# show snmp-server v3 user 1

Name adminmd5, auth md5, privacy des
```

Table 18. USM User Table Information Parameters

Field	Description
User Name	A string representing the user name you can use to access the switch.
Protocol	Whether messages sent from this user are protected from disclosure using a privacy protocol. IBM N/OS supports DES algorithm for privacy and two authentication algorithms: MD5 and HMAC-SHA.

SNMPv3 View Table Information

The user can control and restrict the access allowed to a group to only a subset of the management information in the management domain that the group can access within each context by specifying the group’s rights in terms of a particular MIB view for security reasons.

The following command displays the SNMPv3 View Table:

```
show snmp-server v3 view <1-128>
```

Command mode: Al

The View table contains the following information:

- the user name
- the SNMP version
- the type

```
5000V# show snmp-server v3 view 1
 name v1v2only,subtree 1
 type Included
```

Table 19. SNMPv3 View Table Information Parameters

Field	**Description**
View Name	Displays the name of the view.
Subtree	Displays the MIB subtree as an OID string. A view subtree is the set of all MIB object instances which have a common Object Identifier prefix to their names.
Type	Displays whether a family of `view subtrees` is included or excluded from the MIB view.

SNMPv3 Access Table Information

The access control subsystem provides authorization services.

The vacmAccessTable maps a group name, security information, a context, and a message type, which could be the read or write type of operation or notification into a MIB view.

The View-based Access Control Model defines a set of services that an application can use for checking access rights of a group. This group's access rights are determined by a read-view, a write-view and a notify-view. The read-view represents the set of object instances authorized for the group while reading the objects. The write-view represents the set of object instances authorized for the group when writing objects. The notify-view represents the set of object instances authorized for the group when sending a notification.

The following command displays SNMPv3 access information:

```
show snmp-server v3 access <1-32>
```

Command mode: All

```
group name admingrp model usm
   level AuthPriv,
   read view iso, writeview iso, notify view iso
```

Table 20. SNMPv3 Access Table Information

Field	**Description**
Group Name	Displays the name of group.
Model	Displays the security model used, for example, SNMPv1, or SNMPv2 or USM.
Level	Displays the minimum level of security required to gain rights of access. For example, `noAuthNoPriv`, `authNoPriv`, or `authPriv`.
Read View	Displays the MIB view to which this entry authorizes the read access.

Table 20. SNMPv3 Access Table Information (continued)

Field	Description
Write View	Displays the MIB view to which this entry authorizes the write access.
Notify View	Displays the Notify view to which this entry authorizes the notify access.

SNMPv3 Group Table Information

A group is a combination of security model and security name that defines the access rights assigned to all the security names belonging to that group. The group is identified by a group name.

The following command displays SNMPv3 group information:

```
show snmp-server v3 group <1 - 16>
```

Command mode: All

```
5000V# show snmp-server v3 group 1

 model usm, user Name adminmd5, group Name admingrp
```

Table 21. SNMPv3 Group Table Information Parameters

Field	Description
Sec Model	Displays the security model used, which is any one of: USM, SNMPv1, SNMPv2, and SNMPv3.
User Name	Displays the name for the group.
Group Name	Displays the access name of the group.

SNMPv3 Community Table Information

The following command displays the SNMPv3 community table information stored in the SNMP engine:

```
show snmp-server v3 community
```

Command mode: All

```
5000V(config)# show snmp-server v3 community

Index      Name       User Name            Tag
---------- ---------- -------------------- ----------
```

Table 22. SNMPv3 Community Table Information Parameters

Field	Description
Index	Displays the unique index value of a row in this table
Name	Displays the community string, which represents the configuration.
User Name	Displays the User Security Model (USM) user name.
Tag	Displays the community tag. This tag specifies a set of transport endpoints from which a command responder application accepts management requests and to which a command responder application sends an SNMP trap.

SNMPv3 Target Address Table Information

The following command displays SNMPv3 target address information stored in the SNMP engine:

```
show snmp-server v3 target-address
```

Command mode: All

```
5000V(config)# show snmp-server v3 target-address
Name         Tranport Addr    Port   Taglist    Params
--------   ---------------   -----   --------  ---------
trapr      10.10.53.200      162    traptag    traptag
```

Table 23. SNMPv3 Target Address Table Information Parameters

Field	Description
Name	Displays the locally arbitrary, but unique identifier associated with this `snmpTargetAddrEntry`.
Transport Addr	Displays the transport addresses.
Port	Displays target port.

Table 23. SNMPv3 Target Address Table Information Parameters (continued)

Field	**Description**
Taglist	This column contains a list of tag values which are used to select target addresses for a particular SNMP message.
Params	The value of this object identifies an entry in the `snmpTargetParamsTable`. The identified entry contains SNMP parameters to be used when generating messages to be sent to this transport address.

SNMPv3 Target Parameters Table Information

The following command displays SNMPv3 target parameters information:

```
show snmp-server v3 target-parameters
```

Command mode: All

```
5000V(config)# show snmp-server v3 target-parameters
Name                MP Model   User Name        Sec Model  Sec Level
 ----------------- --------- ----------------- --------- ---------
 traptag             snmpv1                      snmpv1    noAuthNoPriv
```

Table 24. SNMPv3 Target Parameters Table Information

Field	**Description**
Name	Displays the locally arbitrary, but unique identifier associated with this `snmpTargeParamsEntry`.
MP Model	Displays the Message Processing Model used when generating SNMP messages using this entry.
User Name	Displays the `securityName`, which identifies the entry on whose behalf SNMP messages will be generated using this entry.
Sec Model	Displays the security model used when generating SNMP messages using this entry. The system may choose to return an `inconsistentValue` error if an attempt is made to set this variable to a value for a security model which the system does not support.
Sec Level	Displays the level of security used when generating SNMP messages using this entry.

SNMPv3 Notify Table Information

The following command displays the SNMPv3 Notify Table:

```
show snmp-server v3 notify
```

Command mode: All

```
5000V(config)# show snmp-server v3 notify

Name                             Tag
------------------------------   --------------------
trapr                            traptag
```

Table 25. SNMPv3 Notify Table Information

Field	Description
Name	The locally arbitrary, but unique identifier associated with this `snmpNotifyEntry`.
Tag	This represents a single tag value which is used to select entries in the `snmpTargetAddrTable`. Any entry in the `snmpTargetAddrTable` that contains a tag value equal to the value of this entry, is selected. If this entry contains a value of zero length, no entries are selected.

SNMPv3 Dump Information

The following command displays SNMPv3 information:

```
show snmp-server v3
```

Command mode: All

```
5000V# show snmp-server v3

EngineId: 80.00.08.1c.04.46.53

usmUser Table:
User Name                       Protocol
---------------------------- ------------------------------
adminmd5                      HMAC_MD5   DES PRIVACY
adminsha                      HMAC_SHA   DES PRIVACY
v1v2only                      No Auth    NO PRIVACY

vacmAccess Table:
Group Name  Model    Level        ReadV      WriteV     Notify
---------- ------- ------------ ----------- ---------- ----------
v1v2grp    snmpv1   noAuthNoPriv  iso         iso         v1v2only
admingrp   usm      AuthPriv      iso         iso         iso


vacmViewTreeFamily Table:
View Name            Subtree                          Mask            Type
-------------------- -------------------------------- --------------  ---------
iso                  1                                                Included

v1v2only             1                                                Included

v1v2only             1.3.6.1.6.3.15                                   Excluded

v1v2only             1.3.6.1.6.3.16                                   Excluded

v1v2only             1.3.6.1.6.3.18                                   Excluded

vacmSecurityToGroup Table:
Sec Model  User Name                      Group Name
---------- ------------------------------ --------------------
snmpv1      v1v2only                        v1v2grp
usm         adminmd5                        admingrp
usm         adminsha                        admingrp


snmpCommunity Table:
Index      Name       User Name            Tag
---------- ---------- -------------------- ----------



snmpTargetAddr Table:
Name            Tranport Addr   Port   Taglist    Params
 --------------- --------------- -----  ---------- ---------------


snmpTargetParams Table:
Name                MP Model   User Name        Sec Model  Sec Level
 ------------------ --------- ------------------ --------- ---------


snmpNotify Table:
Name                             Tag
------------------------------   --------------------
```

System Information

General System Information

The following command displays system information:

```
show sys-info
```

Command mode: All

```
5000V# show sys-info

IBM System Networking Distributed Switch 5000V

System Information at
 Wed Jan 25 2012 02:22:22
 Switch has been up for 0 day, 19 hours, 14 minutes and 19 seconds

Last boot: power cycle

Management Port MAC Address:  00:50:56:af:79:d2
Management Port IP Address:  172.20.213.229
Software Version 1.0.0.1506, Boot Version 1.0.0.1506, active config block
```

System information includes:

- System date and time
- Switch model
- Time of last boot
- MAC address of the switch management processor
- Software image file and version number, and configuration name.
- IP address of the management interface
- Log-in banner, if one is configured

System Daylight Savings Time (DST) Information

The following command displays system information:

```
show system daylight
```

Command mode: All

```
5000V# show system daylight
DST is Disabled
```

System Time Zone Information

The following command displays system time zone information:

```
show system timezone
```

Command mode: All

```
5000V(config)# show system timezone
Time zone set to 64 - America/US/Pacific (UTC -8:00)
```

System Open Source Acknowledgement Information

The following command displays system acknowledgement information:

```
show system acknowledgement
```

Command mode: All

```
5000V# show system ack

NOTICES AND INFORMATION

IBM System Networking Distributed Switch 5000V 1.0
IBM System Networking Distributed Switch 5000V 1.0
IBM System Networking Distributed Switch 5000V 1.0

The IBM license agreement and any applicable information on the web
download page for IBM products refers Licensee to this file for details
concerning notices applicable to code included in the products listed
above ("the Program").

Notwithstanding the terms and conditions of any other agreement Licensee
may have with IBM or any of its related or affiliated entities
(collectively "IBM"), the third party code identified below is subject
to the terms and conditions of the IBM license agreement for the Program
and not the license terms that may be contained in the notices below.
The notices are provided for informational purposes.

Please note: This Notices file may identify information that is not used
by, or that was not shipped with, the Program as Licensee installed it.

IMPORTANT: IBM does not represent or warrant that the information in this
NOTICES file is accurate. Third party websites are independent of IBM and
IBM does not represent or warrant that the information on any third party
website referenced in this NOTICES file is accurate. IBM disclaims any
and all liability for errors and omissions or for any damages accruing
from the use of this NOTICES file or its contents, including without
limitation URLs or references to any third party websites.


The following code packages or programs contained in these packages contain
certain requisite notices which are reproduced below:

@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
PACKAGE: audit

ATTRIBUTION[844]:

This originates from X11R5 (mit/util/scripts/install.sh), which was
later released in X11R6 (xc/config/util/install.sh) with the following
copyright and license.

Copyright (C) 1994 X Consortium
--More--
```

System CPU Information

The following command displays system CPU information:

```
show system cpu
```

Command mode: All

```
5000V# show system cpu
-----------------------------------------------------------
CPU information:
Load Average (over the last 1 min):     0.28
Load Average (over the last 5 mins):    0.27
Load Average (over the last 15 mins):   0.31
Runnable tasks/Total processes:         1/70
PID of the most recent process:         174
```

System Memory Information

The following command displays system memory utilization information:

```
show system memory
```

Command mode: All

```
5000V# show system memory
-----------------------------------------------------------
Memory information:
MemTotal:        1035564 kB
MemFree:          644376 kB
Buffers:            4424 kB
Cached:            37844 kB
SwapCached:            0 kB
Active:           348328 kB
Inactive:          31128 kB
Active(anon):     337192 kB
Inactive(anon):        0 kB
Active(file):      11136 kB
Inactive(file):    31128 kB
Unevictable:           0 kB
Mlocked:               0 kB
HighTotal:        135112 kB
HighFree:          18288 kB
LowTotal:         900452 kB
LowFree:          626088 kB
SwapTotal:             0 kB
SwapFree:              0 kB
Dirty:                 0 kB
Writeback:             0 kB
AnonPages:        339236 kB
Mapped:             8616 kB
Shmem:                 4 kB
Slab:               4840 kB
SReclaimable:       1592 kB
SUnreclaim:         3248 kB
KernelStack:         544 kB
PageTables:          908 kB
NFS_Unstable:          0 kB
Bounce:                0 kB
WritebackTmp:          0 kB
CommitLimit:      517780 kB
Committed_AS:     416360 kB
VmallocTotal:     122880 kB
VmallocUsed:        2324 kB
VmallocChunk:     118208 kB
AnonHugePages:    215040 kB
DirectMap4k:        8184 kB
DirectMap2M:      905216 kB
```

System Clock Information

The following command displays system clock settings information:

```
show clock
```

Command mode: All

```
5000V# show clock
Current system time:
 Tue Jan 24 2012 18:59:51
```

System Logging Information

The following command displays system logging information:

```
show logging
```

Command mode: All

```
5000V# show logging
Current syslog configuration:
  host  : 0.0.0.0, severity : 7, facility : 0
  host2 : 0.0.0.0, severity2 : 7, facility2 : 0
  console logging  : enabled

syslogging for  are disabled
```

System Access Users Information

The following command displays information about configured users:

```
show access user
```

Command mode: All

```
5000V# show access user
Usernames:
  admin - Always Enabled   - online  1 session.
  user  - enabled          - offline
  oper  - disabled         - offline
```

iSwitch Information

The following commands display the iSwitch configuration information.

Table 26. iSwitch 2 Information Commands

Command Syntax and Usage
`show iswitch hosts` Displays iSwitch hosts information. For details, see page 130. **Command mode:** All except USER Executive
`show iswitch lacp host` *<IP address>* Displays dynamic trunk information of an iSwitch host. For details, see page 131. **Command mode:** All except USER Executive
`show iswitch uplinks host` *<IP address>* Displays uplink information of an iSwitch host. For details, see page 133. **Command mode:** All except USER Executive
`show iswitch portmap [`*<port number>*`]` Displays iSwitch port map information. (Optional) You can specify a port number. For details, see page 134. **Command mode:** All except USER Executive
`show iswitch myvsidb` Displays iSwitch VSI database information. For details, see page 134. **Command mode:** Global Configuration
`show iswitch ports` Displays iSwitch ports information. For details, see page 135. **Command mode:** All except USER Executive

iSwitch Hosts Information

iSwitch All Hosts Information

The following command displays iSwitch information for all hosts:

```
show iswitch hosts
```

Command mode: All

```
5000V# show iswitch hosts

Hosts Connection Information:
 IP Address        Age Time    Agent Version              DPM Version
--------------------------------------------------------------------------------
 172.20.213.112          0s    1.0.0.1429                 1.0.0.1429
 172.20.213.104          4s    1.0.0.1461                 1.0.0.1461
```

iSwitch LACP Host Information

The following command displays iSwitch information for LACP host:

`show iswitch lacp host` *<IP address>*

Command mode: All

```
5000V# show iswitch lacp host 172.20.213.112
vds port 136, ifindex 0
---------------------------------------------
port MAC address - 00:50:56:5a:54:10
Port enabled                 - TRUE
LACP enabled                 - FALSE
LACP admin enabled           - FALSE

Actor System Priority        - 32768
Actor System ID              - 00:50:56:5a:54:10
Actor Admin Key              - 0
Actor Oper Key               - 0
Actor Port Number            - 136
Actor Port Priority          - 32768

Partner Oper System Priority  - 0
Partner Oper System ID       - 00:00:00:00:00:00
Partner Oper Key             - 0
Partner Oper Port Number     - 0
Partner Oper Port Priority   - 0

Actor Oper Port state
  Activity:        Active  Timeout:      Long    Aggregation:   FALSE
  Synchronization: FALSE   Collecting:   FALSE   Distributing:  FALSE
  Defaulted:       TRUE    Expired:      FALSE


Partner Oper Port state
  Activity:        Passive Timeout:      Long    Aggregation:   FALSE
  Synchronization: FALSE   Collecting:   FALSE   Distributing:  FALSE
  Defaulted:       FALSE   Expired:      FALSE


Bound Aggregator ID          - 0
Selected Aggregator ID       - 65535
Attached Aggregator ID       - 65535
ready_n                      - FALSE
ntt                          - FALSE
selected                     - Unselected
port_moved                   - FALSE

Rx machine state             - LACP_RX_LACP_DISABLED_STATE
Mux machine state            - LACP_MUX_DETACHED_STATE
Periodic machine state       - LACP_PERIODIC_NO_STATE

Periodic transmit timer      - 0 s
Current while timer          - 0 s
Wait_while_timer             - 0 s

...continued
```

```
...continued
5000V# show iswitch lacp host 172.20.213.112

vds port 137, ifindex 1
----------------------------------------------
port MAC address - 00:50:56:5a:54:12
Port enabled                  - FALSE
LACP enabled                  - FALSE
LACP admin enabled            - FALSE

Actor System Priority         - 32768
Actor System ID               - 00:50:56:5a:54:10
Actor Admin Key               - 0
Actor Oper Key                - 0
Actor Port Number             - 137
Actor Port Priority           - 32768

Partner Oper System Priority  - 0
Partner Oper System ID        - 00:00:00:00:00:00
Partner Oper Key              - 0
Partner Oper Port Number      - 0
Partner Oper Port Priority    - 0

Actor Oper Port state
  Activity:       Active  Timeout:       Long    Aggregation:   FALSE
  Synchronization: FALSE   Collecting:    FALSE   Distributing:  FALSE
  Defaulted:       TRUE    Expired:       FALSE


Partner Oper Port state
  Activity:       Passive Timeout:       Long    Aggregation:   FALSE
  Synchronization: FALSE   Collecting:    FALSE   Distributing:  FALSE
  Defaulted:       FALSE   Expired:       FALSE


Bound Aggregator ID           - 0
Selected Aggregator ID        - 65535
Attached Aggregator ID        - 65535
ready_n                       - FALSE
ntt                           - FALSE
selected                      - Unselected
port_moved                    - FALSE

Rx machine state              - LACP_RX_PORT_DISABLE_STATE
Mux machine state             - LACP_MUX_DETACHED_STATE
Periodic machine state        - LACP_PERIODIC_NO_STATE

Periodic transmit timer       - 0 s
Current while timer           - 0 s
Wait_while_timer              - 0 s


Aggr Id: 0
----------------------------------------------
ref_count: 2
aggr_name: dvportgroup-35
actor_oper_key: 0
ready: 0
attaching ports [0] (vds port ID):
ports_bound_bitmap:  3 0 0 0
ports_not_ready_bitmap:  0 0 0 0
ports_waiting_attach_bitmap:  0 0 0 0
```

iSwitch Uplinks Host Information

The following command displays iSwitch information for uplinks host:

`show iswitch uplinks host` *<IP address>*

Command mode: All

```
5000V# show iswitch uplinks host 172.20.213.112

Uplink Team 0: dvportgroup-35
Uprofile name: QA_VDS-Uplink-Default
----------------------------------------
Team State: active
LAG MODE  : asymmetric
LAG HASH  : vport
Fowarding uplink ports (alias):
 0

Uplink Port Information:
 DVS      Alias      DEV           MAC            Link
 Port      Name      Name        Address          Status     Forwarding
--------------------------------------------------------------------------------
 136    DVSUplink0    vmnic2    00:50:56:5a:54:10    LINKUP      YES
 137    DVSUplink1    vmnic3    00:50:56:5a:54:12    LINKDOWN    NO

VDS port 0, alias i 1
Name: Type_62_112_Port62.eth0
--------------------------------------------------------------------------------
Teaming State: ISWITCH_TEAMING_PORT_DEFAULT_STATE
Designated uplinks (alias):
Dedicated uplink port (vds port): 136

VDS port 1, alias i 2
Name: Type_63_112_Port63.eth0
--------------------------------------------------------------------------------
Teaming State: ISWITCH_TEAMING_PORT_DEFAULT_STATE
Designated uplinks (alias):
Dedicated uplink port (vds port): 136

VDS port 2, alias i 3
Name: Type_64_112_Port64.eth0
--------------------------------------------------------------------------------
Teaming State: ISWITCH_TEAMING_PORT_DEFAULT_STATE
Designated uplinks (alias):
Dedicated uplink port (vds port): 136

VDS port 3, alias i 4
Name: Type_65_112_Port65.eth0
--------------------------------------------------------------------------------
Teaming State: ISWITCH_TEAMING_PORT_DEFAULT_STATE
Designated uplinks (alias):
Dedicated uplink port (vds port): 136

VDS port 4, alias i 5
Name: Type_66_112_Port66.eth0
--------------------------------------------------------------------------------
Teaming State: ISWITCH_TEAMING_PORT_DEFAULT_STATE
Designated uplinks (alias):
Dedicated uplink port (vds port): 136
```

iSwitch Port Map Information

The following command displays iSwitch port map information:

show iswitch portmap [<*port number*>]

Command mode: All

```
5000V# show iswitch portmap 1

Port mapping Information:

BladeOs Port ID    Vds Port ID    Profile name

=================|===============|=================|
     1                 0      QA_VDS-Standalone-Ports
```

iSwitch VSI Database Information

The following command displays iSwitch VSI database information:

```
show iswitch myvsidb
```

Command mode: Global Configuration

```
5000V(config)# show iswitch myvsidb
VSI Data Base Address: 127.0.0.1
VSI Data Base Port   : 80
VSI Data Base Path   : /vsitypes/
INDEX : 1
-----------------------------------
        Name           :
        Type ID        : 1
        Version        : 1
        Manager ID     : 0
        VLAN           : 2
        TxRate         : 0
        TxBurst        : 0
        RxRate         : 0
        RxBurst        : 0

        ACL Index      : 1
        ---------------------
            SRC MAC     : 00:00:00:00:00:00
            SRC MAC MASK: 00:00:00:00:00:00
            DST MAC     : 00:00:00:00:00:00
            DST MAC MASK: 00:00:00:00:00:00
            VLAN        : 2
            VLAN MASK   : 0 (0x0)
            Ether Type  : 0x0
            ACL Action  : permit
```

iSwitch Port Information

The following command displays iSwitch ports information:

```
show iswitch ports
```

Command mode: All

```
5000V(config)# show iswitch ports
Port vDs-Port Profile     Connectee      Host          Mac-Address      Status
=================================================================================
   1     0  STANDALONE   Type_62_11.. 172.20.213.112  00:50:56:a8:9f:6a Enabled
   2     1  STANDALONE   Type_63_11.. 172.20.213.112  00:50:56:a8:9f:6e Enabled
   3     2  STANDALONE   Type_64_11.. 172.20.213.112  00:50:56:a8:9f:6b Enabled
   4     3  STANDALONE   Type_65_11.. 172.20.213.112  00:50:56:a8:9f:6c Enabled
   5     4  STANDALONE   Type_66_11.. 172.20.213.112  00:50:56:a8:9f:6d Enabled
   6     5  STANDALONE
   7     6  STANDALONE
   8     7  STANDALONE
   9     8  STANDALONE
--More--
```

VLAN Information

Use the following commands to view VLAN information.

Table 27. VLAN Information Options

Command Syntax and Usage
`show vlan` Displays VLAN configuration information for all configured VLANs, including: – VLAN Number – VLAN Name – Status – Port membership of the VLAN For details, see page 136. **Command mode:** All
`show vlan` <*VLAN number or range*> Displays general VLAN information. The valid VLAN numbers are 1 - 4094. **Command mode:** All
`show private-vlan` <*VLAN number*> Displays Private VLAN information. **Command mode:** All

Table 27. VLAN Information Options (continued)

Command Syntax and Usage
`show private-vlan detail` Displays detailed Private VLAN information. **Command mode:** All
`show vlan information` Displays information about all VLANs, including: – VLAN number and name – Port membership – VLAN status (enabled or disabled) – Private VLAN status For details, see page 136. **Command mode:** All

All VLANs Information

The following commands display VLAN information for all VLANs:

```
show vlan
```

```
show vlan information
```

Command mode: All

```
5000V# show vlan

VLAN                Name                Status            Ports
----  -------------------------------  ------  ------------------------
1     VLAN 1                           ena      6-4000
2     VLAN 2                           ena      1-5
```

```
5000V# show vlan information

Current VLAN 1 :
        name "VLAN 1",  ports 6-4000, enabled


Current VLAN 2 :
        name "VLAN 2",  ports 1-5, enabled
```

This information display includes all configured VLANs and all member ports that have an active link state. Port membership is represented in slot/port format.

VLAN information includes:

- VLAN Number
- VLAN Name
- Status
- Port membership of the VLAN
- Private VLAN information (if available)

Single VLAN Information

The following command displays information for a single VLAN:

show vlan <*VLAN number or range*>

Command mode: All

```
5000V# show vlan 1

VLAN            Name              Status          Ports
----  --------------------------- ------  ---------------------
1     VLAN 1                      ena       6-4000
```

Private VLANs Information

The following commands display VLAN information for private VLANs:

show private-vlan <*VLAN number*>

show private-vlan detail

Command mode: All

```
5000V(config)# show private-vlan 200

Private-VLAN    Type      Mapped-To   Status      Ports
------------  ---------  ---------- ----------  ----------
200           primary       201        ena       201-220
```

```
5000V(config)# show private-vlan detail
Current VLAN 200:
        name "VLAN 200",  ports 201-220, enabled
        Private-VLAN: enabled,  type primary,   Mapped to:  201

Current VLAN 201:
        name "VLAN 201",  ports, enabled
        Private-VLAN: enabled,  type community, Mapped to: 200
```

Port Information

The following table includes commands to view port information:

Table 28. Port Information Commands

Command Syntax and Usage
`show interface port` <*port alias or number*> `information` Displays port information. For details, see page 138. **Command mode**: All
`show interface port` <*port alias or number*> `access-list` Displays access lists information for the port. For details, see page 138. **Command mode:** All

Port Information

```
5000V# show interface port 1 information
===============================================================
Blade OS port: 1
VDS port: 0
Port Name: 1
Allocated: Yes
Profile: STANDALONE
Connected: Yes
Connectee: Type_62_112_Port62
Host Name: 172.20.213.112
MAC Address: 00:50:56:a8:9f:6a
IP Address:
Link State: Up
Status: Enabled
ACL(s):
VLAN(s): 2
Tag: n
PVID: 2
Learning: Disabled
Flooding: Disabled
===============================================================
```

```
5000V# show interface port 1 access-list
Current interface port 1 ACL configuration:

    IP ACCESS-LIST CONFIGURED: NIL

    MAC ACCESS-LIST CONFIGURED:NIL
```

Layer 2 Information

The following table lists commands to view Layer 2 information.

Table 29. Layer 3 Information Commands

Command Syntax and Usage
`show layer2 information` Dumps all Layer 2 switch information available (10K or more, depending on your configuration). If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands. **Command mode:** All

```
5000V# show layer2 information
VLAN Information:

VLAN                Name                Status           Ports
----  --------------------------------  ------  -------------------------
1     VLAN 1                            ena      6-4000
2     VLAN 2                            ena      1-5


Forwarding database information:
```

IP Information

The following table lists commands to view Layer 3 information.

Table 30. Layer 3 Information Commands

Command Syntax and Usage
`show ip information` Displays all IP information. For details, see page 139. **Command mode**: All
`show ip route` Displays all routes configured on the switch. For details, see page 140. **Command mode:** All
`show ip arp` Displays Address Resolution Protocol (ARP) information. For details, see page 141. **Command mode:** All
`show interface ip` Displays management IP interface configuration. For details, see page 141. **Command mode**: All
`show ip dns` Displays DNS settings information. For details, see page 142. **Command mode**: All

Show IP Information

The following command displays all IP information:

```
show ip information
```

Command mode: All

```
5000V# show ip information
No configured interfaces exist

Default gateway information:

Current IP forwarding settings: OFF, noicmprd enabled
```

Show All IP Route Information

The following command displays IP route information:

```
show ip route
```

Command mode: All

```
5000V# show ip route
Status code: * - best
   Destination       Mask            Gateway       Type    Tag   Metr If
 --------------- --------------- --------------- ------- ------- ---- --
* 0.0.0.0         0.0.0.0         172.20.1.1      Indirect static      0
* 172.20.0.0      255.255.0.0     172.20.213.229 direct   fixed       0
* 172.20.213.229  255.255.255.255 172.20.213.229 local    addr    0   0
* 172.20.255.255  255.255.255.255 172.20.255.255 bcast    bcast   0   0
```

The following table describes the `Type` parameters.

Table 31. IP Routing Type Parameters

Parameter	**Description**
indirect	The next hop to the host or subnet destination will be forwarded through a router at the `Gateway` address.
direct	Packets will be delivered to a destination host or subnet attached to the switch.
local	Indicates a route to one of the switch’s IP interfaces.
bcast	Indicates a broadcast route.
martian	The destination belongs to a host or subnet which is filtered out. Packets to this destination are discarded.
mcast	Indicates a multicast route.

The following table describes the `Tag` parameters.

Table 32. IP Routing Tag Parameters

Parameter	**Description**
fixed	The address belongs to a host or subnet attached to the switch.
static	The address is a static route which has been configured on the IBM DS 5000V.
addr	The address belongs to one of the switch’s IP interfaces.
bcast	Indicates a broadcast address.
martian	The address belongs to a filtered group.
mcast	Indicates a multicast address.

Show All ARP Entry Information

The following command displays ARP information:

```
show ip arp
```

Command mode: All

```
5000V# show ip arp

Current ARP configuration:
 rearp 5

Current static ARP:
ip                mac                 interface   vlan
--------------    -----------------   ---------   ----



IP Address        Flags   Age(min)   Hardware Address    Interface   Vlan
--------------    -----   --------   -----------------   ---------   ----
172.20.1.1                0          fc:cf:62:40:69:00               1
172.20.27.220             0          00:21:5e:43:75:ee               1
172.20.113.120            1          e4:1f:13:97:66:ce               1
172.20.213.104            1          e4:1f:13:18:c0:b8               1
172.20.213.112            1          e4:1f:13:18:bd:d4               1
172.20.213.229    P       -          00:50:56:af:79:d2               1
```

The `Port` field shows the target port of the ARP entry.

The `Flags` field is interpreted as follows:

Table 33. ARP Flag Parameters

Flag	Description
P	Permanent entry created for switch IP interface.
R	Indirect route entry.
U	Unresolved ARP entry. The MAC address has not been learned.

Interface Information

The following command displays interface information:

```
show interface ip-mgmt
```

Command mode: All

```
5000V# show interface ip-mgmt
Management IP Interface Configurations
IP Address  : 172.20.213.229 MANUAL
Subnet Mask : 255.255.0.0
Gateway Address  : 172.20.1.1,enabled
```

DNS Information

The following command displays DNS settings information:

```
show ip dns
```

Command mode: All

```
5000V# show ip dns
Current DNS settings
--------------------
DNS domain name
DNS primary server address none
DNS secondary server address none
```

ERSPAN Information

The following table includes commands to view the ERSPAN information.

Table 34. ERSPAN Information Commands

Command Syntax and Usage
`show erspan` Displays ERSPAN settings information. For details, see page 143. **Command mode**: All
`show erspan flow [`*<flow ID>*`]` Displays ERSPAN flow settings information. For details, see page 143. **Command mode:** All
`show erspan source [`*<source ID>*`]` Displays ERSPAN source settings information. For details, see page 144. **Command mode:** All
`curr` Displays current ERSPAN settings information. For details, see page 143. **Command mode:** ERSPAN Configuration

ERSPAN ALL Settings Information

The following commands display ERSPAN-related information:

`show erspan`

Command mode: All

```
5000V# show erspan
erspan
        source  1
                direction in
                traffic   unicast
                mode      L2
                vlan      1
                priority  0
                add port
                add vlan
        !
!
```

`curr`

Command mode: ERSPAN configuration

```
5000V(config-erspan)# curr
erspan
        source  1
                direction in
                traffic   unicast
                mode      L2
                vlan      1
                priority  0
                add port
                add vlan
        !
!
```

ERSPAN Flow Settings Information

The following command displays ERSPAN flow settings information:

`show erspan flow`

Command mode: All

```
5000V# show erspan flow
flow 1 1 11.1.1.1
flow 2 src2 1.1.1.12
```

`show erspan flow` *<flow ID>*

Command mode: All

```
5000V# show erspan flow 2
flow 2 src2 1.1.1.12
```

ERSPAN Source Settings Information

The following command displays ERSPAN source settings information:

`show erspan source`

Command mode: All

```
5000V# show erspan source
        source  1
                direction in
                traffic   unicast
                mode      L2
                vlan      1
                priority  0
                add port
                add vlan
        !
```

`show erspan source` *<source ID>*

Command mode: All

```
5000V# show erspan source 2
        source  2
                direction out
                traffic   multicast
                mode      L2
                vlan      2
                priority  0
                add port 10
                add vlan 2
        !
```

ERSPAN Source Information

The following table includes commands to view the ERSPAN source information.

Table 35. ERSPAN Source Information Commands

Command Syntax and Usage
`curr` Displays current ERSPAN source configuration information. **Command mode**: ERSPAN source configuration

```
5000V(erspan-source-1)# curr
        source  1
                direction in
                traffic   unicast
                mode      L2
                vlan      1
                priority  0
                add port
                add vlan
        !
```

Chapter 15. Statistics Commands

You can use the Statistics Commands to view switch performance statistics in both the user and administrator command modes. This chapter discusses how to use the command line interface to display switch statistics.

Table 36. Statistics Commands

Command Syntax and Usage
`show snmp-server counters` **Command mode:** All Displays SNMP statistics. See page 164 for sample output.
`show ntp` Displays Network Time Protocol (NTP) Statistics. **Command mode:** All See page 168 for a sample output and a description of NTP Statistics.
`show ntp verbose` Displays detailed Network Time Protocol (NTP) Statistics. **Command mode:** All See page 168 for a sample output and a description of NTP Statistics.
`show counters` Dumps all switch statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command. **Command mode:** All For details, see page 169.
`show system packet` Displays packet statistics. For details, see page 152. **Command mode:** All
`clear ntp` Clears NTP statistics. **Command mode:** Privilege EXEC
`clear ntp primary-server` Clears primary server NTP statistics. **Command mode:** Privilege EXEC

Table 36. Statistics Commands

Command Syntax and Usage
`clear ntp secondary-server` Clears secondary server NTP statistics. **Command mode:** Privilege EXEC
`clear nvram` Clears contents of NVRAM. **Command mode:** Privilege EXEC

Access Control Statistics

The following commands display ACL statistics.

Table 37. Port Statistics Commands

Command Syntax and Usage
`show access-list counters` Displays ACL statistics. See page 147 for sample output. **Command mode:** All
`show access-list` *<1 - 254>* `counters` Displays statistics for the specified ACL . See page 147 for sample output. **Command mode:** All
`show access-list ip counters` Displays statistics for all IP ACLs. See page 147 for sample output. **Command mode:** All
`show access-list mac counters` Displays statistics for all MAC ACLs. See page 147 for sample output. **Command mode:** All
`show access-list ip` *<128 - 254>* `counters` Displays statistics for the specified IP ACL. See page 147 for sample output. **Command mode:** All
`show access-list mac` *<1 - 127>* `counters` Displays statistics for the specified MAC ACL. See page 147 for sample output. **Command mode:** All
`clear access-list [`*<1 - 254>*`|ip` *<128 - 254>*`|mac` *<1 - 127>*`] counters` Clears statistics for the ACL you specify. **Command mode:** All

Table 37. Port Statistics Commands (continued)

Command Syntax and Usage
`clear access-list counters` Clears all ACL statistics. **Command mode:** All
`clear access-list {ip|mac} counters` Clears statistics for all IP or MAC ACLs. **Command mode:** All

All Access Control Counters

The following command displays statistics for all ACLs. If you need statistics for any one ACL, use the command: `show access-list` *<1 - 254>* `counters`.

```
show access-list counters
```

Command mode: All

IP Access Control Counters

The following command displays statistics for all IP ACLs. If you need statistics for any one IP ACL, use the command: `show access-list ip` *<128 - 254>* `counters`.

```
show access-list ip counters
```

Command mode: All

MAC Access Control Counters

The following command displays statistics for all MAC ACLs. If you need statistics for any one MAC ACL, use the command: `show access-list mac` *<1-127>* `counters`.

```
show access-list mac counters
```

Command mode: All

Port Statistics

These commands display traffic statistics on a port-by-port basis. Traffic statistics include SNMP Management Information Base (MIB) objects.

Table 38. Port Statistics Commands

Command Syntax and Usage
`show interface port` *<port alias or number>* `interface-counters` Displays interface statistics for the port. See page 149 for sample output. **Command mode:** All
`show interface port` *<port alias or number>* `interface-rate` Displays per-second interface statistics for the port. **Command mode:** All
`show interface port` *<port alias or number>* `link-counters` Displays link statistics for the port. See page 151 for sample output. **Command mode:** All
`show interface port` *<port alias or number>* `mp-counters` Displays interface CPU statistics. See page 152 for sample output. **Command mode:** All
`clear interface port` *<port alias or number>* `counters` Clears all statistics for the port. **Command mode:** Privileged EXEC

Interface Statistics

Use the following command to display the interface statistics of the selected port:

`show interface port` *<port alias or number>* `interface-counters`

Command mode: All.

```
5000V# show interface port 1 interface-counters
------------------------------------------------------
Interface statistics for port 1
                    ifHCIn Counters    ifHCOut Counters
 Octets:                    6816838                6816950
 UcastPkts:                   70343                  70344
 BroadcastPkts:                   1                      2
 MulticastPkts:                   0                      0
 Discards:                        0                  14769
 Errors:                          0                      0

Ingress Discard reasons for port 1
 VLAN Discards:                            0
 Empty Egress Portmap:                     0
 Filter Discards:                          0
 Policy Discards:                          0
 Non-Forwarding State:                     0
 IBP/CBP Discards:                         0
```

Table 39. Interface Statistics of a Port

Statistics	Description
ifHCInOctets	The total number of octets received on the interface, including framing characters.
ifHCInUcastPkts	The number of packets, delivered by this sub-layer to a higher sub- layer, which were not addressed to a multicast or broadcast address at this sub-layer.
ifHCInBroadcastPkts	The number of packets, delivered by this sub-layer to a higher sub- layer, which were addressed to a broadcast address at this sub-layer.
ifHCInMulticastPkts	The total number of packets that higher-level protocols requested to be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
ifHCInDiscards	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being delivered to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.

Table 39. Interface Statistics of a Port (continued)

Statistics	**Description**
ifHCInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
ifHCOutOctets	The total number of octets transmitted out of the interface, including framing characters.
ifHCOutUcastPkts	The total number of packets that higher-level protocols requested to be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
ifHCOutBroadcastPkts	The total number of packets that higher-level protocols requested to be transmitted, and which were addressed toa broadcast address at this sub-layer, including those that were discarded or not sent. This object is a 64-bit version of `ifOutBroadcastPkts`.
ifHCOutMulticastPkts	The total number of packets that higher-level protocols requested to be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses. This object is a 64-bit version of `ifOutMulticastPkts`.
ifHCOutDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifHCOutErrors	For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
VLAN Discards	Discarded because the packet was tagged with a VLAN to which this port is not a member.
Filter Discards	Dropped by the Content Aware Engine (user-configured filter).
Policy Discards	Dropped due to policy setting. For example, due to a user-configured static entry.
Non-Forwarding State	Discarded because the ingress port is not in the forwarding state.

Table 39. Interface Statistics of a Port (continued)

Statistics	Description
IBP/CBP Discards	Discarded because of Ingress Back Pressure (flow control), or because the Common Buffer Pool is full (for example, insufficient packet buffering).
Empty Egress Portmap	Dropped due to an egress port bitmap of zero condition (no ports in the egress mask). This counter increments whenever the switching decision found that there was no port to send out.

Interface Rate Statistics

Use the following command to display the per-second interface statistics of the selected port:

`show interface port` *<port alias or number>* `interface-rate`

Command mode: All

```
5000V# show interface port 1 interface-rate

 in Cfa show interfaces rate
------------------------------------------------------
Interface statistics for port 1
                         ifHCIn Rate    ifHCOut Rate
 Bits:                          3136                      3136
 UcastPkts:                        4                         4
 BroadcastPkts:                    0                         0
 MulticastPkts:                    0                         0
 Discards:                         0                         0
 Errors:                           0                         0
```

See Table 39 on page 149 for a description of the statistics displayed.

Link Statistics

Use the following command to display the link statistics of the selected port:

`show interface port` *<port alias or number>* `link-counters`

Command mode: All

```
5000V# show interface port 1 link-counters
------------------------------------------------------------------
Link statistics for port:1
linkStateChange:2
```

Table 40. Link Statistics of a Port

Statistics	Description
linkStateChange	The total number of link state changes.

CPU Statistics

Use the following command to display the CPU statistics of the selected port:

`show interface port` *<port alias or number>* `mp-counters`

Command mode: All

```
5000V# show interface port 1 mp-counters
------------------------------------------------------
Interface MP Statistics for port 1
 MP Pkts Received:                 0
 MP Mcast Received:                0
```

Table 41. CPU Statistics of a Port

Statistics	Description
MP Pkts Received	Number of packets received by the management processor.
MP Mcast Received	Number of multicast packets received by the management processor.

System Statistics

Use the following command to display packet statistics:

`show system packet`

Command mode: All

```
5000V# show system packet
---------------------------------------------------------------------
Packet counts:
allocs:          119723 frees:           119723
hi-watermark:       126 failures:             0
```

Table 42. System Packet Statistics

Statistics	Description
allocs	Total number of packet allocations from the packet buffer pool by the TCP/IP protocol stack.
frees	Total number of times the packet buffers are freed (released) to the packet buffer pool by the TCP/IP protocol stack.
hi-watermark	The highest number of packet allocation with size less than 128 bytes from the packet buffer pool by the TCP/IP protocol stack.
failures	Total number of packet allocation failures from the packet buffer pool by the TCP/IP protocol stack.

iSwitch Statistics

Use the following commands to view iSwitch statistics.

Table 43. Layer 3 Statistics Commands

Command Syntax and Usage
`show iswitch uplinks interface-counters host` *<IP address>* Displays uplink statistics of a iSwitch host interface. **Command mode:** All except USER Executive
`clear iswitch uplinks counters [host` *<IP address>*`]` Clears uplink iSwitch uplink counters. (Optional) You can clear the uplink counters of a particular host. See for sample output. **Command mode:** All except USER Executive

```
5000V# show iswitch uplinks interface-counters host 172.20.213.112
------------------------------------------------------
Interface statistics for port DVS port 136
                    ifHCIn Counters    ifHCOut Counters
 Octets:                    65471455                  13695346
 UcastPkts:                   141486                    141318
 BroadcastPkts:               526357                        16
 MulticastPkts:               101004                         0
 Discards:                    517331                         0
 Errors:                           0                         0

Ingress Discard reasons for port DVS port 136
 VLAN Discards:                        515597
 Empty Egress Portmap:                   3915
 Filter Discards:                           0
 Policy Discards:                           0
 Non-Forwarding State:                      0
 IBP/CBP Discards:                          0
------------------------------------------------------
Interface statistics for port DVS port 137
                    ifHCIn Counters    ifHCOut Counters
 Octets:                     1094367                       378
 UcastPkts:                        0                         0
 BroadcastPkts:                12454                         9
 MulticastPkts:                  179                         0
 Discards:                      5449                         0
 Errors:                           0                         0

Ingress Discard reasons for port DVS port 137
 VLAN Discards:                          3894
 Empty Egress Portmap:                     85
 Filter Discards:                           0
 Policy Discards:                           0
 Non-Forwarding State:                      0
 IBP/CBP Discards:                          0
```

See Table 39 on page 149 for a description of the statistics displayed.

IP Statistics

Table 44. Layer 3 Statistics Commands

Command Syntax and Usage
`show ip counters` Displays IP statistics. See page 155 for sample output. **Command mode:** All
`show ip dns counters` Displays Domain Name System (DNS) statistics. See page 156 for sample output. **Command mode:** All
`show ip icmp counters` Displays ICMP statistics. See page 156 for sample output. **Command mode:** All
`show ip tcp counters` Displays TCP statistics. See page 158 for sample output. **Command mode:** All
`show ip udp counters` Displays UDP statistics. See page 159 for sample output. **Command mode:** All
`clear ip counters` Clears IP statistics. **Command mode:** Privileged EXEC
`clear ip dns counters` Clears Domain Name System (DNS) statistics. **Command mode:** Privileged EXEC
`clear ip dhcp counters` Clears DHCP statistics. **Command mode:** Privileged EXEC

IP Counters

The following command displays IP statistics:

```
show ip counters
```

Command mode: All

```
5000V# show ip counters

-------------------------------------------------------------------
IP statistics:
ipInReceives:     37462 ipInHdrErrors:    37462
ipInAddrErrors:       0 ipForwDatagrams:      1
ipInUnknownProtos:   26 ipInDiscards:         0
ipInDelivers:     37408 ipOutRequests:    17305
ipOutDiscards:        2 ipOutNoRoutes:        4
ipReasmReqds:         0 ipReasmOKs:           0
ipReasmFails:         0 ipFragOKs:            0
ipFragFails:          0 ipFragCreates:        0
ipRoutingDiscards:    0 ipDefaultTTL:        64
ipReasmTimeout:      30
```

Table 45. IP Statistics

Statistics	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity (the switch). This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities which are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ipInUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in `ipForwDatagrams` if any such packets met this (discretionary) discard criterion.
ipReasmReqds	
ipReasmFails	
ipFragFails	

Table 45. IP Statistics

Statistics	Description
ipRoutingDiscards	
ipReasmTimeout	

DNS Statistics

The following command displays Domain Name System statistics.

```
show ip dns counters
```

Command mode: All

```
5000V# show ip dns counters

---------------------------------------------------------------
DNS statistics:
dnsInRequests:     0   dnsOutRequests:    0
dnsBadRequests:    0
```

Table 46. DNS Statistics

Statistics	Description
dnsInRequests	The total number of DNS response packets that have been received.
dnsOutRequests	The total number of DNS response packets that have been transmitted.
dnsBadRequests	The total number of DNS request packets received that were dropped.

ICMP Statistics

The following command displays ICMP statistics:

```
show ip icmp counters
```

Command mode: All

```
5000V# show ip icmp counters

---------------------------------------------------------------------
ICMP statistics:
icmpInMsgs:                   7 icmpInErrors:                 0
icmpInDestUnreachs:           7 icmpInTimeExcds:              0
icmpInParmProbs:              0 icmpInSrcQuenchs:             0
icmpInRedirects:              0 icmpInEchos:                  0
icmpInEchoReps:               0 icmpInTimestamps:             0
icmpInTimestampReps:          0 icmpInAddrMasks:              0
icmpInAddrMaskReps:           0 icmpOutMsgs:                 16
icmpOutErrors:                0 icmpOutDestUnreachs:         16
icmpOutTimeExcds:             0 icmpOutParmProbs:             0
icmpOutSrcQuenchs:            0 icmpOutRedirects:             0
icmpOutEchos:                 0 icmpOutEchoReps:              0
icmpOutTimestamps:            0 icmpOutTimestampReps:         0
icmpOutAddrMasks:             0 icmpOutAddrMaskReps:          0
```

Table 47. ICMP Statistics

Statistic	**Description**
icmpInMsgs	The total number of ICMP messages which the entity (the switch) received. Note that this counter includes all those counted by `icmpInErrors`.
icmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
icmpInParmProbs	The number of ICMP Parameter Problem messages received.
icmpInRedirects	The number of ICMP Redirect messages received.
icmpInEchoReps	The number of ICMP Echo Reply messages received.
icmpInTimestampReps	The number of ICMP Timestamp `Reply` messages received.
icmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
icmpOutErrors	The number of ICMP messages which this entity (the switch) did not send due to problems discovered within ICMP such as a lack of buffer. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of errors that contribute to this counter's value.
icmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
icmpOutSrcQuenchs	The number of ICMP Source Quench (buffer almost full, stop sending data) messages sent.
icmpOutEchos	The number of ICMP Echo (request) messages sent.
icmpOutEchoReps	The number of ICMP Echo Reply messages sent.
icmpOutTimestamps	The number of ICMP Timestamp (request) messages sent.
icmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.

TCP Statistics

The following command displays TCP statistics:

```
show ip tcp counters
```

Command mode: All

```
5000V# show ip tcp counters

--------------------------------------------------------------------
TCP statistics:
tcpRtoAlgorithm:      4 tcpRtoMin:           0
tcpRtoMax:            0 tcpMaxConn:        500
tcpActiveOpens:       0 tcpPassiveOpens:     4
tcpAttemptFails:      0 tcpEstabResets:      0
tcpInSegs:        29381 tcpOutSegs:      18038
tcpRetransSegs:       0 tcpInErrs:          15
tcpCurrEstab:         3 tcpCurrConn:         6
tcpOutRsts:          15
```

Table 48. TCP Statistics

Statistic	Description
tcpRtoAlgorithm	The algorithm used to determine the `timeout` value used for retransmitting unacknowledged octets.
tcpRtoMax	The maximum value permitted by a TCP implementation for the retransmission `timeout`, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission `timeout`. In particular, when the `timeout` algorithm is rsre(3), an object of this type has the semantics of the UBOUND quantity described in RFC 793.
tcpActiveOpens	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
tcpAttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
tcpInSegs	The total number of segments received, including those received in error. This count includes segments received on currently established connections.
tcpRetransSegs	The total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
tcpCurrEstab	The total number of outstanding memory allocations from heap by TCP protocol stack.
tcpOutRsts	The number of TCP segments sent containing the RST flag.

UDP Statistics

The following command displays UDP statistics:

```
show ip udp counters
```

Command mode: All

```
5000V# show ip udp counters

-------------------------------------------------------------------
UDP statistics:
udpInDatagrams:    9539 udpOutDatagrams:      1
udpInErrors:       9536 udpNoPorts:        9536
```

Table 49. UDP Statistics

Statistic	Description
udpInDatagrams	The total number of UDP datagrams delivered to the switch.
udpOutDatagrams	The total number of UDP datagrams sent from this entity (the switch).
udpInErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
udpNoPorts	The total number of received UDP datagrams for which there was no application at the destination port.

Management Processor Statistics

Use the following commands to view management processor statistics.

Table 50. Management Processor Statistics Commands

Command Syntax and Usage
`show mp packet` Displays packet statistics, to check for leads and load. To view a sample output and a description of the stats, see page 161. **Command mode:** All
`show mp tcp-block` Displays all TCP control blocks that are in use. To view a sample output and a description of the stats, see page 161. **Command mode:** All
`show mp udp-block` Displays all UDP control blocks that are in use. To view a sample output, see page 162. **Command mode:** All
`show mp cpu` Displays CPU utilization statistics. To view a sample output, see page 163. **Command mode:** All

MP Packet Statistics

The following command displays MP packet statistics:

`show mp packet`

Command mode: All

```
 5000V# show mp packet
---------------------------------------------------------------------
Packet counts:
allocs:          136166 frees:          136166
hi-watermark:       126 failures:            0
```

Table 51. MP Packet Statistics

Statistics	**Description**
allocs	Total number of packet allocations from the packet buffer pool by the TCP/IP protocol stack.
frees	Total number of times the packet buffers are freed (released) to the packet buffer pool by the TCP/IP protocol stack.
hi-watermark	The highest number of packet allocation with size less than 128 bytes from the packet buffer pool by the TCP/IP protocol stack.
failures	Total number of packet allocation failures from the packet buffer pool by the TCP/IP protocol stack.

MP TCP Statistics

The following command displays MP TCP statistics:

`show mp tcp-block`

Command mode: All

```
5000V# show mp tcp-block
-----------------------------------------------------------
All TCP allocated control blocks:
  0.0.0.0          0   <=> 0.0.0.0         22      LISTEN
  0.0.0.0          0   <=> 0.0.0.0         23      LISTEN
  0.0.0.0          0   <=> 0.0.0.0        902      LISTEN
  10.10.53.160 56836   <=> 172.20.213.229    23 ESTABLISHED
  172.20.213.104 60565   <=> 172.20.213.229   902 ESTABLISHED
  172.20.213.112 62043   <=> 172.20.213.229   902 ESTABLISHED
```

Table 52. MP Specified TCP Statistics

Statistics	**Description**
`10.10.53.160`	Destination IP address
`56836`	Destination port
`172.20.213.229`	Source IP

Table 52. MP Specified TCP Statistics

Statistics	Description
`23`	Source port
`Established`	State

MP UDP Statistics

The following command displays MP UDP statistics:

```
show mp udp-block
```

Command mode: All

```
5000V# show mp udp-block
-----------------------------------------------------------
All UDP allocated control blocks:
    68:    LISTEN
   123:    LISTEN
   161:    LISTEN
  1812:    LISTEN
  1813:    LISTEN
  6123:    LISTEN
```

MP CPU Statistics

The following commands display the MP CPU utilization statistics:

```
show mp cpu
```

Command mode: All

```
 5000V# show mp cpu
-----------------------------------------------------------
CPU information:
Load Average (over the last 1 min):     0.54
Load Average (over the last 5 mins):    0.45
Load Average (over the last 15 mins):   0.40
Runnable tasks/Total processes:         1/68
PID of the most recent process:         154
-----------------------------------------------------------
Memory information:
MemTotal:        1035564 kB
MemFree:          651480 kB
Buffers:            4436 kB
Cached:            34528 kB
SwapCached:            0 kB
Active:           343724 kB
Inactive:          31140 kB
Active(anon):     335904 kB
Inactive(anon):        0 kB
Active(file):       7820 kB
Inactive(file):    31140 kB
Unevictable:           0 kB
Mlocked:               0 kB
HighTotal:        135112 kB
HighFree:          17872 kB
LowTotal:         900452 kB
LowFree:          633608 kB
SwapTotal:             0 kB
SwapFree:              0 kB
Dirty:                 0 kB
Writeback:             0 kB
AnonPages:        335900 kB
Mapped:             8592 kB
Shmem:                 4 kB
Slab:               4740 kB
SReclaimable:       1580 kB
SUnreclaim:         3160 kB
KernelStack:         544 kB
PageTables:          900 kB
NFS_Unstable:          0 kB
Bounce:                0 kB
WritebackTmp:          0 kB
CommitLimit:      517780 kB
Committed_AS:     416492 kB
VmallocTotal:     122880 kB
VmallocUsed:        2324 kB
VmallocChunk:     118208 kB
AnonHugePages:    210944 kB
DirectMap4k:        8184 kB
DirectMap2M:      905216 kB
```

SNMP Statistics

The following command displays SNMP statistics:

```
show snmp-server counters
```

Command mode: All

```
5000V# show snmp-server counters
SNMP statistics:
-----------------------------------------------------------------
snmpInPkts:                    0    snmpInBadVersions:                0
snmpInBadC'tyNames:            0    snmpInBadC'tyUses:                0
snmpInASNParseErrs:            0    snmpEnableAuthTraps:              2
snmpOutPkts:                   0    snmpInBadTypes:                   0
snmpInTooBigs:                 0    snmpInNoSuchNames                 0
snmpInBadValues                0    snmpInReadOnlys                   0
snmpInGenErrs                  0    snmpInTotalReqVars                0
snmpInTotalSetVars             0    snmpInGetRequests                 0
snmpInGetNexts                 0    snmpInSetRequests                 0
snmpInGetResponses             0    snmpInTraps                    4004
snmpOutTooBigs                 0    snmpOutNoSuchNames                0
snmpOutBadValues               0    snmpOutReadOnlys                  0
snmpOutGenErrs                 0    snmpOutGetRequests                0
snmpOutGetNexts                0    snmpOutSetRequests                0
snmpOutGetResponses            0    snmpOutTraps                      0
snmpSilentDrops                0    snmpProxyDrops                    0
```

Table 53. SNMP Statistics

Statistic	Description
snmpInPkts	The total number of Messages delivered to the SNMP entity from the transport service.
snmpInBadVersions	The total number of SNMP Messages, which were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
snmpInBadC'tyNames	The total number of SNMP Messages delivered to the SNMP entity which used an SNMP community name not known to the said entity (the switch).
snmpInBadC'tyUses	The total number of SNMP Messages delivered to the SNMP protocol entity which represented an SNMP operation which was not allowed by the SNMP community named in the Message.

Table 53. SNMP Statistics (continued)

Statistic	**Description**
snmpInASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding SNMP Messages received. **Note:** OSI's method of specifying abstract objects is called ASN.1 (Abstract Syntax Notation One, defined in X.208), and one set of rules for representing such objects as strings of ones and zeros is called the BER (Basic Encoding Rules, defined in X.209). ASN.1 is a flexible notation that allows one to define a variety of data types, from simple types such as integers and bit strings to structured types such as sets and sequences. BER describes how to represent or encode values of each ASN.1 type as a string of eight-bit octets.
snmpEnableAuthTraps	An object to enable or disable the authentication traps generated by this entity (the switch).
snmpOutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
snmpInBadTypes	The total number of SNMP Messages which failed ASN parsing.
snmpInTooBigs	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is *too big.*
snmpInNoSuchNames	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is `noSuchName`.
snmpInBadValues	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is `badValue`.
snmpInReadOnlys	The total number of valid SNMP Protocol Data Units (PDUs), which were delivered to the SNMP protocol entity and for which the value of the error-status field is `read-Only'. It should be noted that it is a protocol error to generate an SNMP PDU, which contains the value `read-Only' in the error-status field. As such, this object is provided as a means of detecting incorrect implementations of the SNMP.
snmpInGenErrs	The total number of SNMP Protocol Data Units (PDUs), which were delivered to the SNMP protocol entity and for which the value of the error-status field is `genErr`.

Table 53. SNMP Statistics (continued)

Statistic	**Description**
snmpInTotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as a result of receiving valid SNMP Get-Request and Get-Next Protocol Data Units (PDUs).
snmpInTotalSetVars	The total number of MIB objects, which have been altered successfully by the SNMP protocol entity as a result of receiving valid SNMP Set-Request Protocol Data Units (PDUs).
snmpInGetRequests	The total number of SNMP Get-Request Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInGetNexts	The total number of SNMP Get-Next Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInSetRequests	The total number of SNMP Set-Request Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInGetResponses	The total number of SNMP Get-Response Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInTraps	The total number of SNMP Trap Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpOutTooBigs	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is *too big*.
snmpOutNoSuchNames	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status is `noSuchName`.
snmpOutBadValues	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is `badValue.`
snmpOutReadOnlys	Not in use.
snmpOutGenErrs	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is `genErr`.
snmpOutGetRequests	The total number of SNMP Get-Request Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.

Table 53. SNMP Statistics (continued)

Statistic	**Description**
snmpOutGetNexts	The total number of SNMP Get-Next Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
snmpOutSetRequests	The total number of SNMP Set-Request Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
snmpOutGetResponses	The total number of SNMP Get-Response Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
snmpOutTraps	The total number of SNMP Trap Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
snmpSilentDrops	The total number of `GetRequest`-PDUs, `GetNextRequest`-PDUs, `GetBulkRequest`-PDUs, `SetRequest`-PDUs, and `InformRequest`-PDUs delivered to the SNMPv2 entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
snmpProxyDrops	The total number of `GetRequest`-PDUs, `GetNextRequest`-PDUs, `GetBulkRequest`-PDUs, `SetRequest`-PDUs, and `InformRequest`-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the message to a proxy target failed in a manner such that no Response-PDU could be returned.

NTP Statistics

IBM N/OS uses NTP (Network Timing Protocol) version 3 to synchronize the switch’s internal clock with an atomic time calibrated NTP server. With NTP enabled, the switch can accurately update its internal clock to be consistent with other devices on the network and generates accurate syslogs.

The following command displays NTP statistics:

```
show ntp
```

Command mode: All

Table 54. NTP Statistics

Field	**Description**
Primary Server	• **Requests Sent:** The total number of NTP requests the switch sent to the primary NTP server to synchronize time. • **Responses Received:** The total number of NTP responses received from the primary NTP server. • **Updates:** The total number of times the switch updated its time based on the NTP responses received from the primary NTP server.
Secondary Server	• **Requests Sent:** The total number of NTP requests the switch sent to the secondary NTP server to synchronize time. • **Responses Received:** The total number of NTP responses received from the secondary NTP server. • **Updates:** The total number of times the switch updated its time based on the NTP responses received from the secondary NTP server.
Last update based on response from primary server	Last update of time on the switch based on either primary or secondary NTP response received.
Last update time	The time stamp showing the time when the switch was last updated.
Current system time	The switch system time when the following command was issued: `show ntp counters`

Statistics Dump

The following command dumps switch statistics:

```
show counters
```

Use the dump command to dump all switch statistics (40K or more, depending on your configuration). This data can be used to tune or debug switch performance.

If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command.

Chapter 16. Configuration Commands

This chapter discusses how to use the Command Line Interface (CLI) for making, viewing, and saving switch configuration changes. Many of the commands, although not new, display more or different information than in the previous version. Important differences are called out in the text.

Table 55. General Configuration Commands

Command Syntax and Usage
`show running-config` Dumps current configuration to a script file. For details, see page 228. You can specify additional optional parameters for more specific information. **Command mode:** All
`show config-not-restored` Displays unsaved configuration that was not restored when the switch was last reloaded. **Command mode:** All
`show startup-config` Dumps the startup configuration stored in Flash. **Command mode:** All
`show active-config` Dumps the active configuration stored in Flash. **Command mode:** All
`show backup-config` Dumps the backup configuration stored in Flash. **Command mode:** All
`copy running-config backup-config` Copy the current (running) configuration from switch memory to the `backup-config` partition. For details, see page 228. **Command mode:** Privileged EXEC
`copy running-config startup-config` Copy the current (running) configuration from switch memory to the `startup-config` partition. **Command mode:** Privileged EXEC
`copy running-config active-config` Copy the current (running) configuration from switch memory to Flash. **Command mode:** Privileged EXEC
`copy running-config {tftp|scp}` Backs up current configuration to a file on the selected TFTP/SCP server. **Command mode:** Privileged EXEC

Table 55. General Configuration Commands

Command Syntax and Usage
`copy {tftp|scp} active-config` Restores current configuration from a TFTP/SCP server. **Command mode:** Privileged EXEC
`copy {tftp|scp} backup-config` Restores backup configuration from a TFTP/SCP server. **Command mode:** Privileged EXEC
`copy active-config {tftp|scp}` Copies current configuration to a TFTP/SCP server. **Command mode:** Privileged EXEC
`copy backup-config {tftp|scp}` Copies backup configuration to a TFTP/SCP server. **Command mode:** Privileged EXEC

Viewing and Saving Changes

As you use the configuration commands to set switch parameters, the changes you make take effect immediately (you do not need to apply them).

All changes are stored in a draft configuration state until the switch reboots, or until you save them to permanent memory.

If the 5000V controller is rebooted without current configuration changes being explicitly saved, a notification will be displayed during the next boot cycle and you will be prompted to select whether to restore the draft configuration or continue with last saved configuration.

You can display unsaved and non-restored commands with the following executive mode command:

```
5000V# show config-not-restored
```

To save the current draft configuration settings to permanent memory, enter the following command:

```
5000V# copy running-config startup-config
```

When you save configuration changes, the draft changes are saved to the *active* configuration block and the draft configuration is cleared. For instructions on selecting the configuration to run at the next system reset, see “Selecting a Configuration Block” on page 235.

It is also recommended to save a copy of the configuration file in a remote TFTP or SCP server in case the configuration must later be recovered.

Note: Some operations can override the settings of the Configuration commands. Therefore, settings you view using the Configuration commands (for example, port status) might differ from run-time information that you view using the Information commands. The Information commands display current run-time information of switch parameters.

System Configuration

These commands provide configuration of switch management parameters such as user and administrator privilege mode passwords, Web-based management settings, and management access lists.

Table 56. System Configuration Options

Command Syntax and Usage
`system date` *<yyyy> <mm> <dd>* Prompts the user for the system date. The date retains its value when the switch is reset. **Command mode:** Global configuration
`system time` *<hh>*`:`*<mm>*`:`*<ss>* Configures the system time using a 24-hour clock format. The time retains its value when the switch is reset. **Command mode:** Global configuration
`[no] system timezone` Configures the time zone where the switch resides. Select your location (continent, country, region) from the menu. Once a location is selected, the switch updates the time to reflect local changes to Daylight Savings Time, etc. **Command mode:** Global configuration
`system timezone` *<time zone ID>* Configures the time zone where the switch resides. Enter the time zone ID for your location (continent, country, region). Once the ID is entered, the switch updates the time to reflect local changes to Daylight Savings Time, etc. **Command mode:** Global configuration
`system timezone` *<Continent|Country|City>* Configures the time zone where the switch resides. Enter your Continent or Country or City. Once a region is entered, the switch updates the time to reflect local changes to Daylight Savings Time, etc. **Command mode:** Global configuration
`[no] system daylight` Disables or enables daylight savings time in the system clock. When enabled, the switch will add an extra hour to the system clock so that it is consistent with the local clock. By default, this option is disabled. **Command mode:** Global configuration

Table 56. System Configuration Options (continued)

Command Syntax and Usage
`system idle` *<1-60>* Sets the idle timeout for CLI sessions, from 1 to 60 minutes. The default is 5 minutes. **Command mode:** Global configuration
`[no] system notice[1-5]` *<maximum 1024 character multi-line login notice>* `<'.'` *to end>* Displays login notice immediately before the “Enter password:” prompt. This notice can contain up to 1024 characters and new lines. **Command mode:** Global configuration
`[no] banner` *<1-80 characters>* Configures a login banner of up to 80 characters. When a user or administrator logs into the switch, the login banner is displayed. It is also displayed as part of the output from the `show sys-info` command. **Command mode:** Global configuration
`[no] hostname` *<character string>* Enables or disables displaying of the host name (system administrator’s name) in the Command Line Interface (CLI). **Command mode:** Global configuration
`show system` Displays the current system parameters. **Command mode:** All

System Host Log Configuration

Use the following commands to configure logging on hosts.

Table 57. Host Log Configuration Options

Command Syntax and Usage
`[no] logging host` *<1-2>* `address` *<IP address>* Sets the IP address of the first or second syslog host. **Command mode:** Global configuration
`logging host` *<1-2>* `severity` *<0-7>* This option sets the severity level of the first or second syslog host displayed. The default is 7, which means log all severity levels. **Command mode:** Global configuration
`logging host` *<1-2>* `facility` *<0-7>* This option sets the facility level of the first or second syslog host displayed. The default is 0. **Command mode:** Global configuration
`logging console` Enables delivering syslog messages to the console. It is enabled by default. **Command mode:** Global configuration
`no logging console` Disables delivering syslog messages to the console. When necessary, disabling `console` ensures the switch is not affected by syslog messages. It is enabled by default. **Command mode:** Global configuration
`[no] logging log [`*<feature>*`]` Displays a list of features for which syslog messages can be generated. You can choose to enable/disable specific features (such as `vlans`), or enable/disable syslog on all available features. **Command mode:** Global configuration
`show logging` Displays the current syslog settings. **Command mode:** All

SSH Server Configuration

For the IBM DS 5000V, these commands enable Secure Shell access from any SSH client.

Table 58. SSH Server Configuration Options

Command Syntax and Usage
`ssh generate-host-key` Generate the RSA host key. **Command mode:** Global configuration
`ssh generate-server-key` Generate the RSA server key. **Command mode:** Global configuration
`ssh port` *<SSH Access port number>* Sets the SSH access port number. **Command mode:** Global configuration
`ssh enable` Enables the SSH server. **Command mode:** Global configuration
`no ssh enable` Disables the SSH server. **Command mode:** Global configuration
`ssh interval` *<0-24>* Set the interval for generation RSA server key in seconds. **Command mode:** Global configuration
`show ssh` Displays the current SSH server configuration. **Command mode:** All

RADIUS Server Configuration

Use the following commands to configure RADIUS server.

Table 59. RADIUS Server Configuration Options

Command Syntax and Usage
`[no] radius-server primary-host` *<IP address>* Sets the primary RADIUS server address. **Command mode:** Global configuration
`[no] radius-server secondary-host` *<IP address>* Sets the secondary RADIUS server address. **Command mode:** Global configuration

Table 59. RADIUS Server Configuration Options (continued)

Command Syntax and Usage
`radius-server primary-host` *<IP address>* `key` *<1-32 characters>* This is the primary shared secret between the switch and the RADIUS server(s). **Command mode:** Global configuration
`radius-server secondary-host` *<IP address>* `key` *<1-32 characters>* This is the secondary shared secret between the switch and the RADIUS server(s). **Command mode:** Global configuration
`[default] radius-server port` *<UDP port number>* Enter the number of the UDP port to be configured, between 1500 - 3000. The default is 1645. **Command mode:** Global configuration
`radius-server retransmit` *<1-3>* Sets the number of failed authentication requests before switching to a different RADIUS server. The default is 3 requests. **Command mode:** Global configuration
`[no] radius-server secure-backdoor` Enables or disables the RADIUS back door using secure password for telnet/SSH/HTTP/HTTPS. This command does not apply when backdoor (`telnet`) is enabled. **Command mode:** Global configuration
`radius-server timeout` *<1-10>* Sets the amount of time, in seconds, before a RADIUS server authentication attempt is considered to have failed. The default time is three seconds. **Command mode:** Global configuration
`radius-server enable` Enables the RADIUS server. **Command mode:** Global configuration
`no radius-server enable` Disables the RADIUS server. **Command mode:** Global configuration
`show radius-server` Displays the current RADIUS server parameters. **Command mode:** All

TACACS+ Server Configuration

TACACS (Terminal Access Controller Access Control system) is an authentication protocol that allows a remote access server to forward a user's logon password to an authentication server to determine whether access can be allowed to a given system. TACACS is not an encryption protocol, and therefore less secure than TACACS+ and Remote Authentication Dial-In User Service (RADIUS) protocols. Both TACACS and TACACS+ are described in RFC 1492.

TACACS+ protocol is more reliable than RADIUS, as TACACS+ uses the Transmission Control Protocol (TCP) whereas RADIUS uses the User Datagram Protocol (UDP). Also, RADIUS combines authentication and authorization in a user profile, whereas TACACS+ separates the two operations.

TACACS+ offers the following advantages over RADIUS as the authentication device:

- TACACS+ is TCP-based, so it facilitates connection-oriented traffic.
- It supports full-packet encryption, as opposed to password-only in authentication requests.
- It supports de-coupled authentication, authorization, and accounting.

Table 60. TACACS+ Server Configuration Options

Command Syntax and Usage
`[no] tacacs-server primary-host` *<IP address>* Defines the primary TACACS+ server address. **Command mode:** Global configuration
`[no] tacacs-server secondary-host` *<IP address>* Defines the secondary TACACS+ server address. **Command mode:** Global configuration
`[no] tacacs-server primary-host` *<IP address>* `key` *<1-32 characters>* This is the primary shared secret between the switch and the TACACS+ server(s). **Command mode:** Global configuration
`[no] tacacs-server secondary-host` *<IP address>* `key` *<1-32 characters>* This is the secondary shared secret between the switch and the TACACS+ server(s). **Command mode:** Global configuration
`[default] tacacs-server port` *<TCP port number>* Enter the number of the TCP port to be configured, between 1 and 65000. The default is 49. **Command mode:** Global configuration

Table 60. TACACS+ Server Configuration Options (continued)

Command Syntax and Usage
`[no] tacacs-server secure-backdoor` Enables or disables TACACS+ secure back door access through Telnet, SSH/SCP, or HTTP/HTTPS only when the TACACS+ servers are not responding. This feature is recommended to permit access to the switch when the TACACS+ servers become unresponsive. If no back door is enabled, the only way to gain access when TACACS+ servers are unresponsive is to use the back door via the console port. The default is `disabled`. **Command mode:** Global configuration
`[no] tacacs-server privilege-mapping` Enables or disables TACACS+ privilege-level mapping. The default value is `disabled`. **Command mode:** Global configuration
`[no] tacacs-server command-authorization` Enables or disables TACACS+ command authorization. **Command mode:** Global configuration
`[no] tacacs-server command-logging` Enables or disables TACACS+ command logging. **Command mode:** Global configuration
`[no] tacacs-server enable` Enables or disables the TACACS+ server. By default, the server is disabled. **Command mode:** Global configuration
`[no] preemption enable` Enables or disables the TACACS+ server preemption. **Command mode:** Global configuration
`tacacs-server retransmit` *<1 - 3>* Set the TACACS+ server retry interval. **Command mode:** Global configuration
`tacacs-server timeout` *<4 - 15>* Set the TACACS+ server timeout interval. **Command mode:** Global configuration
`show tacacs-server` Displays current TACACS+ configuration parameters. **Command mode:** All

NTP Server Configuration

These commands allow you to synchronize the switch clock to a Network Time Protocol (NTP) server. By default, this option is disabled.

Table 61. NTP Server Configuration Options

Command Syntax and Usage
`[no] ntp primary-server` *<IP address>* Prompts for the hostname or IP addresses of the primary NTP server to which you want to synchronize the switch clock. **Command mode:** Global configuration
`[no] ntp secondary-server` *<IP address>* Prompts for the hostname or IP addresses of the secondary NTP server to which you want to synchronize the switch clock. **Command mode:** Global configuration
`[no] ntp interval` *<1-10080>* Specifies the interval, that is, how often, in minutes, to re-synchronize the switch clock with the NTP server. **Command mode:** Global configuration
`ntp enable` Enables the NTP synchronization service. **Command mode:** Global configuration
`no ntp enable` Disables the NTP synchronization service. **Command mode:** Global configuration
`show ntp` Displays the current NTP service settings. **Command mode:** All

System SNMP Configuration

IBM N/OS supports SNMP-based network management. In SNMP model of network management, a management station (client/manager) accesses a set of variables known as MIBs (Management Information Base) provided by the managed device (agent). If you are running an SNMP network management station on your network, you can manage the switch using the following standard SNMP MIBs:

- MIB II (RFC 1213)
- Ethernet MIB (RFC 1643)
- Bridge MIB (RFC 1493)

An SNMP agent is a software process on the managed device that listens on UDP port 161 for SNMP messages. Each SNMP message sent to the agent contains a list of management objects to retrieve or to modify.

SNMP parameters that can be modified include:

- System name
- System location
- System contact
- Use of the SNMP system authentication trap function
- Read community string
- Write community string
- Trap community strings

Table 62. System SNMP Options

Command Syntax and Usage
`snmp-server name` *<1-64 characters>* Configures the name for the system. The name can have a maximum of 64 characters. **Command mode:** Global configuration
`[no] snmp-server location` *<1-64 characters>* Configures the name of the system location. The location can have a maximum of 64 characters. **Command mode:** Global configuration
`[no] snmp-server contact` *<1-64 characters>* Configures the name of the system contact. The contact can have a maximum of 64 characters. **Command mode:** Global configuration
`[no] snmp-server read-community` *<1-32 characters>* Configures the SNMP read community string. The read community string controls SNMP “get” access to the switch. It can have a maximum of 32 characters. The default read community string is *public*. **Command mode:** Global configuration
`snmp-server write-community` *<1-32 characters>* Configures the SNMP write community string. The write community string controls SNMP “set” and “get” access to the switch. It can have a maximum of 32 characters. The default write community string is *private*. **Command mode:** Global configuration
`[no] snmp-server authentication-trap` Enables or disables the use of the system authentication trap facility. The default setting is `disabled`. **Command mode:** Global configuration

Table 62. System SNMP Options (continued)

Command Syntax and Usage
`[no] snmp-server link-trap` Enables or disables the sending of SNMP link up and link down traps. The default setting is `enabled`. **Command mode:** Global configuration
`show snmp-server` Displays the current SNMP configuration. **Command mode:** All

SNMPv3 Configuration

SNMP version 3 (SNMPv3) is an extensible SNMP Framework that supplements the SNMPv2 Framework by supporting the following:

- a new SNMP message format
- security for messages
- access control
- remote configuration of SNMP parameters

For more details on the SNMPv3 architecture please refer to RFC3411 to RFC3418.

Table 63. SNMPv3 Configuration Options

Command Syntax and Usage
`snmp-server user` *<1-16>* This command allows you to create a user security model (USM) entry for an authorized user. You can also configure this entry through SNMP. **Command mode:** Global configuration To view command options, see page 184.
`snmp-server view` *<1-128>* This command allows you to create different MIB views. **Command mode:** Global configuration To view command options, see page 185.
`snmp-server access` *<1-32>* This command allows you to specify access rights. The View-based Access Control Model defines a set of services that an application can use for checking access rights of the user. You need access control when you have to process retrieval or modification request from an SNMP entity. **Command mode:** Global configuration To view command options, see page 186.

Table 63. SNMPv3 Configuration Options (continued)

`snmp-server group <1-16>` A group maps the user name to the access group names and their access rights needed to access SNMP management objects. A group defines the access rights assigned to all names that belong to a particular group. **Command mode:** Global configuration To view command options, see page 187.
`snmp-server community <1-16>` The community table contains objects for mapping community strings and version-independent SNMP message parameters. **Command mode:** Global configuration To view command options, see page 188.
`snmp-server target-address <1-16>` This command allows you to configure destination information, consisting of a transport domain and a transport address. This is also termed as transport endpoint. The SNMP MIB provides a mechanism for performing source address validation on incoming requests, and for selecting community strings based on target addresses for outgoing notifications. **Command mode:** Global configuration To view command options, see page 189.
`snmp-server target-parameters <1-16>` This command allows you to configure SNMP parameters, consisting of message processing model, security model, security level, and security name information. There may be multiple transport endpoints associated with a particular set of SNMP parameters, or a particular transport endpoint may be associated with several sets of SNMP parameters. **Command mode:** Global configuration To view command options, see page 190.
`snmp-server notify <1-16>` A notification application typically monitors a system for particular events or conditions, and generates Notification-Class messages based on these events or conditions. **Command mode:** Global configuration To view command options, see page 191.
`snmp-server version {v1v2v3|v3only}` This command allows you to enable or disable the access to SNMP versions 1, 2 or 3. This command is enabled by default. **Command mode:** Global configuration
`show snmp-server v3` Displays the current SNMPv3 configuration. **Command mode:** All

User Security Model Configuration

You can make use of a defined set of user identities using this Security Model. An SNMP engine must have the knowledge of applicable attributes of a user.

These commands help you create a user security model entry for an authorized user. You need to provide a security name to create the USM entry.

Table 64. User Security Model Configuration Options

Command Syntax and Usage
`snmp-server user` *<1-16>* `name` *<1-32 characters>* This command allows you to configure a string that represents the name of the user. This is the login name that you need in order to access the switch. **Command mode:** Global configuration
`snmp-server user` *<1-16>* `authentication-protocol {md5|sha|none} authentication-password` *<password value>* This command allows you to configure the authentication protocol and password. The authentication protocol can be HMAC-MD5-96 or HMAC-SHA-96, or none. The default algorithm is `none`. When you configure an authentication algorithm, you must provide a password, otherwise you will get an error message during validation. This command allows you to create or change your password for authentication. **Command mode:** Global configuration
`snmp-server user` *<1-16>* `privacy-protocol {des|none} privacy-password` *<password value>* This command allows you to configure the type of privacy protocol and the privacy password. The privacy protocol protects messages from disclosure. The options are `des` (CBC-DES Symmetric Encryption Protocol) or `none`. If you specify `des` as the privacy protocol, then make sure that you have selected one of the authentication protocols (MD5 or HMAC-SHA-96). If you select `none` as the authentication protocol, you will get an error message. You can create or change the privacy password. **Command mode:** Global configuration
`no snmp-server user` *<1-16>* Deletes the USM user entries. **Command mode:** Global configuration
`show snmp-server v3 user` *<1-16>* Displays the USM user entries. **Command mode:** All

SNMPv3 View Configuration

Note that the first five default `vacmViewTreeFamily` entries cannot be removed, and their names cannot be changed.

Table 65. SNMPv3 View Configuration Options

Command Syntax and Usage
`snmp-server view` *<1-128>* `name` *<1-32 characters>* This command defines the name for a family of view subtrees. **Command mode:** Global configuration
`snmp-server view` *<1-128>* `tree` *<1-64 characters>* This command defines MIB tree, which when combined with the corresponding mask defines a family of view subtrees. **Command mode:** Global configuration
`[no] snmp-server view` *<1-128>* `mask` *<1-32 characters>* This command defines the bit mask, which in combination with the corresponding tree defines a family of view subtrees. **Command mode:** Global configuration
`snmp-server view` *<1-128>* `type {included|excluded}` This command indicates whether the corresponding instances of `vacmViewTreeFamilySubtree` and `vacmViewTreeFamilyMask` define a family of view subtrees, which is included in or excluded from the MIB view. **Command mode:** Global configuration
`no snmp-server view` *<1-128>* Deletes the `vacmViewTreeFamily` group entry. **Command mode:** Global configuration
`show snmp-server v3 view` *<1-128>* Displays the current `vacmViewTreeFamily` configuration. **Command mode:** All

View-based Access Control Model Configuration

The view-based Access Control Model defines a set of services that an application can use for checking access rights of the user. Access control is needed when the user has to process SNMP retrieval or modification request from an SNMP entity.

Table 66. View-based Access Control Model Options

Command Syntax and Usage
`snmp-server access` *<1-32>* `name` *<1-32 characters>* Defines the name of the group. **Command mode:** Global configuration
`snmp-server access` *<1-32>* `security {usm|snmpv1|snmpv2}` Allows you to select the security model to be used. **Command mode:** Global configuration
`snmp-server access` *<1-32>* `level {noauthnopriv|authnopriv|authpriv}` Defines the minimum level of security required to gain access rights. The level `noAuthNoPriv` means that the SNMP message will be sent without authentication and without using a privacy protocol. The level `authNoPriv` means that the SNMP message will be sent with authentication but without using a privacy protocol. The `authPriv` means that the SNMP message will be sent both with authentication and using a privacy protocol. **Command mode:** Global configuration
`snmp-server access` *<1-32>* `read-view` *<1-32 characters>* Defines a read view name that allows you read access to a particular MIB view. If the value is empty or if there is no active MIB view having this value then no access is granted. **Command mode:** Global configuration
`snmp-server access` *<1-32>* `write-view` *<1-32 characters>* Defines a write view name that allows you write access to the MIB view. If the value is empty or if there is no active MIB view having this value then no access is granted. **Command mode:** Global configuration
`snmp-server access` *<1-32>* `notify-view` *<1-32 characters>* Defines a notify view name that allows you notify access to the MIB view. **Command mode:** Global configuration
`no snmp-server access` *<1-32>* Deletes the View-based Access Control entry. **Command mode:** Global configuration
`show snmp-server v3 access` *<1-32>* Displays the View-based Access Control configuration. **Command mode:** All

SNMPv3 Group Configuration

The Group Access Control Model defines a set of services that an application can use for checking access rights of the group.

Table 67. SNMPv3 Group Configuration Options

Command Syntax and Usage
`snmp-server group` *<1-16>* `security {usm|snmpv1|snmpv2}` Defines the security model. **Command mode:** Global configuration
`snmp-server group` *<1-16>* `user-name` *<1-32 characters>* Sets the user name as defined in the following command on page 184: `snmp-server user` *<1-16>* `name` *<1-32 characters>* **Command mode:** Global configuration
`snmp-server group` *<1-16>* `group-name` *<1-32 characters>* The name for the access group as defined in the following command: `snmp-server access` *<1-32>* `name` *<1-32 characters>* on page 184. **Command mode:** Global configuration
`no snmp-server group` *<1-16>* Deletes the `vacmSecurityToGroup` entry. **Command mode:** Global configuration
`show snmp-server v3 group` *<1-16>* Displays the current `vacmSecurityToGroup` configuration. **Command mode:** All

SNMPv3 Community Table Configuration

These commands are used for configuring the community table entry. The configured entry is stored in the community table list in the SNMP engine. This table is used to configure community strings in the Local Configuration Datastore (LCD) of SNMP engine.

Table 68. SNMPv3 Community Table Configuration Options

Command Syntax and Usage
`snmp-server community` *<1-16>* `index` *<1-32 characters>* Allows you to configure the unique index value of a row in this table. **Command string:** Global configuration
`snmp-server community` *<1-16>* `name` *<1-32 characters>* Defines the user name as defined in the following command on page 184: `snmp-server user` *<1-16>* `name` *<1-32 characters>* **Command string:** Global configuration
`snmp-server community` *<1-16>* `user-name` *<1-32 characters>* Defines a readable string that represents the corresponding value of an SNMP community name in a security model. **Command mode:** Global configuration
`snmp-server community` *<1-16>* `tag` *<1-255 characters>* Allows you to configure a tag. This tag specifies a set of transport endpoints to which a command responder application sends an SNMP trap. **Command mode:** Global configuration
`no snmp-server community` *<1-16>* Deletes the community table entry. **Command mode:** Global configuration
`show snmp-server v3 community` *<1-16>* Displays the community table configuration. **Command mode:** All

SNMPv3 Target Address Table Configuration

These commands are used to configure the target transport entry. The configured entry is stored in the target address table list in the SNMP engine. This table of transport addresses is used in the generation of SNMP messages.

Table 69. Target Address Table Configuration Options

Command Syntax and Usage
`snmp-server target-address` *<1-16>* `address` *<IP address>* `name` *<1-32 characters>* Allows you to configure the locally arbitrary, but unique identifier, target address name associated with this entry. **Command mode:** Global configuration
`snmp-server target-address` *<1-16>* `name` *<1-32 characters>* `address` *<transport IP address>* Configures a transport IPv4 or IPv6 address that can be used in the generation of SNMP traps. IPv6 addresses are not displayed in the configuration, but they do receive traps. **Command mode:** Global configuration
`snmp-server target-address` *<1-16>* `port` *<port range>* Allows you to configure a transport address port that can be used in the generation of SNMP traps. **Command mode:** Global configuration
`snmp-server target-address` *<1-16>* `taglist` *<1-255 characters>* Allows you to configure a list of tags that are used to select target addresses for a particular operation. **Command mode:** Global configuration
`snmp-server target-address` *<1-16>* `parameters-name` *<1-32 characters>* Defines the name as defined in the following command on page 190: `snmp-server target-parameters` *<1-16>* `name` *<1-32 characters>* **Command mode:** Global configuration
`no snmp-server target-address` *<1-16>* Deletes the Target Address Table entry. **Command mode:** Global configuration
`show snmp-server v3 target-address` *<1-16>* Displays the current Target Address Table configuration. **Command mode:** All

SNMPv3 Target Parameters Table Configuration

You can configure the target parameters entry and store it in the target parameters table in the SNMP engine. This table contains parameters that are used to generate a message. The parameters include the message processing model (for example: SNMPv3, SNMPv2c, SNMPv1), the security model (for example: USM), the security name, and the security level (`noAuthnoPriv` , `authNoPriv`, or `authPriv`).

Table 70. Target Parameters Table Configuration Options

Command Syntax and Usage
`snmp-server target-parameters` *<1-16>* `name` *<1-32 characters>* Allows you to configure the locally arbitrary, but unique, identifier that is associated with this entry. **Command mode:** Global configuration
`snmp-server target-parameters` *<1-16>* `message {snmpv1|snmpv2c|snmpv3}` Allows you to configure the message processing model that is used to generate SNMP messages. **Command mode:** Global configuration
`snmp-server target-parameters` *<1-16>* `security {usm|snmpv1|snmpv2}` Allows you to select the security model to be used when generating the SNMP messages. **Command mode:** Global configuration
`snmp-server target-parameters` *<1-16>* `user-name` *<1-32 characters>* Defines the name that identifies the user in the USM table (page 184) on whose behalf the SNMP messages are generated using this entry. **Command mode:** Global configuration
`snmp-server target-parameters` *<1-16>* `level {noAuthNoPriv|authnopriv|authpriv}` Allows you to select the level of security to be used when generating the SNMP messages using this entry. The level `noAuthNoPriv` means that the SNMP message will be sent without authentication and without using a privacy protocol. The level `authNoPriv` means that the SNMP message will be sent with authentication but without using a privacy protocol. The `authPriv` means that the SNMP message will be sent both with authentication and using a privacy protocol. **Command mode:** Global configuration
`no snmp-server target-parameters` *<1-16>* Deletes the `targetParamsTable` entry. **Command mode:** Global configuration
`show snmp-server v3 target-parameters` *<1-16>* Displays the current `targetParamsTable` configuration. Command mode: All

SNMPv3 Notify Table Configuration

SNMPv3 uses Notification Originator to send out traps. A notification typically monitors a system for particular events or conditions, and generates Notification-Class messages based on these events or conditions.

Table 71. Notify Table Options

Command Syntax and Usage
`snmp-server notify` *<1-16>* `name` *<1-32 characters>* Defines a locally arbitrary, but unique, identifier associated with this SNMP notify entry. **Command mode:** Global configuration
`snmp-server notify` *<1-16>* `tag` *<1-255 characters>* Allows you to configure a tag that contains a tag value which is used to select entries in the Target Address Table. Any entry in the `snmpTargetAddrTable`, that matches the value of this tag, is selected. **Command mode:** Global configuration
`no snmp-server notify` *<1-16>* Deletes the notify table entry. **Command mode:** Global configuration
`show snmp-server v3 notify` *<1-16>* Displays the current notify table configuration. **Command mode:** All

System Access Configuration

Use the following commands to enable secure access to the switch.

Table 72. System Access Configuration Options

Command Syntax and Usage
`access user user-password` Sets the user (`user`) password. The user has no direct responsibility for switch management. The user view switch status information and statistics, but cannot make any configuration changes. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. **Note:** To disable the user account, set the password to null (no password). **Command Mode**: Global configuration
`access user operator-password` Sets the operator (`oper)`password. The operator manages all functions of the switch. The operator can view all switch information and statistics and can reset ports. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. **Note:** To disable the operator account, set the password to null (no password). The default setting is disabled (no password). **Command Mode**: Global configuration
`access user administrator-password` Sets the administrator (`admin`) password. The administrator has complete access to all menus, information, and configuration commands on the 5000V, including the ability to change both the user and administrator passwords. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. Access includes “`oper`” functions. **Note:** You cannot disable the administrator password. **Command Mode**: Global configuration
`[no] access snmp {read-only|read-write}` Disables or provides read-only/write-read SNMP access. **Command mode:** Global configuration
`no access snmp` Disables SNMP access control. **Command mode:** Global configuration
`[no] access telnet enable` Enables or disables Telnet access. This command is enabled by default. **Command mode:** Global configuration

Table 72. System Access Configuration Options (continued)

Command Syntax and Usage
`[default] access telnet port [`*<1-65535>*`]` Sets an optional Telnet server port number for cases where the server listens for Telnet sessions on a non-standard port. **Command mode:** Global configuration
`[default] access tftp-port [`*<1-65535>*`]` Sets the TFTP port for the switch. The default is port 69. **Command mode:** Global configuration
`[default] access scp server-port` *<1 - 65535>* Set the SCP server access port. **Command Mode**: Global configuration
`default access ssh port` Set SSH as the default access. **Command Mode**: Global configuration
`show access` Displays the current system access parameters. **Command mode:** All

User Access Control Configuration

The following table describes user-access control commands.

Passwords can be a maximum of 128 characters.

Table 73. User Access Control Configuration Options

Command Syntax and Usage
`access user eject` *<user name>* Ejects the specified user from the 5000V. **Command mode:** Global configuration
`access user eject console-user` Ejects the console user from the 5000V. **Command mode:** Global configuration
`access user user-password` Sets the user (`user`) password. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. **Command mode:** Global configuration

Table 73. User Access Control Configuration Options

Command Syntax and Usage
`access user operator-password` Sets the operator (`oper`)password. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. **Command mode:** Global configuration
`access user administrator-password` Sets the administrator (`admin`) password. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. Access includes "`oper`" functions. **Command mode:** Global configuration
`show access user` Displays the current user status. **Command mode:** All except User EXEC

System User ID Configuration

The following table describes user ID configuration commands.

Table 74. User ID Configuration Options

Command Syntax and Usage
access user <*1-10*> level {user|operator|administrator} Sets the Class-of-Service to define the user's authority level. IBM N/OS defines these levels as: User, Operator, and Administrator, with User being the most restricted level. **Command mode:** Global configuration
access user <*1-10*> name <*1-8 characters*> Defines the user name of maximum eight characters. **Command mode:** Global configuration
access user <*1-10*> password Sets the user (`user`) password. This command will prompt for required information: current admin password, new password (up to 128 characters) and confirmation of the new password. **Command mode:** Global configuration
access user <*1-10*> enable Enables the user ID. **Command mode:** Global configuration
no access user <*1-10*> enable Disables the user ID. **Command mode:** Global configuration

Table 74. User ID Configuration Options

Command Syntax and Usage
`no access user` *<1-10>* Deletes the user ID. **Command mode:** Global configuration
`show access user uid` *<1 - 10>* Displays the current user ID configuration. **Command mode:** All except User EXEC

Port Configuration

Use the Port Configuration commands to configure settings for interface ports.

Table 75. Port Configuration Options

Command Syntax and Usage
`interface port` *<port alias or number>* Enter Interface port mode. **Command mode:** Global configuration
`[no] interface port` *<port alias or number>* `learning` **Note:** Not supported in this release.
`dot1p` *<0-7>* Configures the port's 802.1p priority level. **Command mode:** Interface port
`pvid` *<1 - 4094>* Sets the default VLAN number which will be used to forward frames which are not VLAN tagged. The default number is 1 for non-management ports. **Command mode:** Interface port
`[no] name` *<1-64 characters>* Set or reset a name for the port. The assigned port name appears next to the port number on some information and statistics screens. The default is set to `None`. **Command mode:** Interface port
`[no] tagging` Disables or enables VLAN tagging for this port. The default setting is `disabled`. **Command mode:** Interface port
`[no] tag-pvid` Disables or enables VLAN tag persistence. When disabled, the VLAN tag is removed from packets whose VLAN tag matches the port PVID. The default setting is `disabled`. **Command mode:** Interface port
`[no] shutdown` Enables or disables the port. **Command mode:** Interface port
`[no] allvlans` **Note:** Not supported in this release.

Table 75. Port Configuration Options (continued)

Command Syntax and Usage
`[no] link-logging` Disables or enables syslog for link up/down. **Command mode:** Interface port
`[no] vepa` Disables or enables Virtual Ethernet Port Aggregator (VEPA) mode. **Command mode:** Interface port
`vsitype` *<1 - 16777215>* Set the VSI type ID. **Command mode:** Interface port
`vsitype` *<1 - 16777215>* `version` *<0 - 255>* Set the VSI type ID and version. **Command mode:** Interface port
`no vsitype` Disables VSI type setting. **Command mode:** Interface port
`[no] dps-disable` Enable or disable DPS setting. **Command mode:** Interface port
`[no] ops-disable` Enable or disable OPS setting. **Command mode:** Interface port
`designated-uplinks` *<port number(s)>* Assign ports as designated uplink ports. Enter a single port or multiple ports separated by a ‘,’ or a range separated by a ‘-’. **Command mode:** Interface port
`designated-uplinks drop` Un-assign all ports from being designated uplink ports. **Command mode:** Interface port
`no designated-uplinks` Disable designated uplinks setting on the interface. **Command mode:** Interface port
`show interface port` *<port alias or number>* `{access-list|information|interface-counters|interface-rate|link-counters|mp-counters}` Displays current port parameters. **Command mode:** All

Table 75. Port Configuration Options (continued)

Command Syntax and Usage
`show vsitypeid [<VSI Type ID>] [<version>]` Displays VSI Type ID information. (Optional) Specify the VSI Type ID value (0 - Clears D16777215) or the version (0 - 255). **Command mode:** Privileged EXEC

Port Link Configuration

Use these commands to set flow control for the port link.

Table 76. Port Link Configuration Options

Command Syntax and Usage
`[no] flowcontrol` **Note:** Not supported in this release.
`[no] auto` Turns auto-negotiation on or off. **Command mode:** Interface port

Port Access List Configuration

Use the following commands to apply an access list to a port.

Table 77. Access List Configuration Options

Command Syntax and Usage
`[no] ip access-group <128 - 254> in` Apply the IP ACL on inbound packets. **Command mode:** Interface port
`[no] mac access-group <1 - 127> in` Apply the MAC ACL on inbound packets. **Command mode:** Interface port

Port Mirroring

Port mirroring is disabled by default. Port Mirroring commands are used to configure, enable, and disable the monitor port. When enabled, network packets being sent and/or received on a target port are duplicated and sent to a monitor port. By attaching a network analyzer to the monitor port, you can collect detailed information about your network performance and usage.

The mirror and monitor ports must be on the same host. Use ERPAN for remote port mirroring. See “ERSPAN Configuration” on page 225.

Table 78. Port Mirroring Configuration Options

Command Syntax and Usage
`[no] port-mirroring enable` Enables or disables port mirroring. **Command mode:** Global configuration
`show port-mirroring` Displays current settings of the mirrored and monitoring ports. **Command mode:** All except User EXEC

Port-Mirroring Configuration

Use the following commands to configure port mirroring.

Table 79. Port-Based Port-Mirroring Configuration Options

Command Syntax and Usage
`port-mirroring monitor-port` *<port alias or number>* `mirroring-port` *<port alias or number>* `{in|out|both}` Adds the port to be mirrored. This command also allows you to enter the direction of the traffic. It is necessary to specify the direction because: If the source port of the frame matches the mirrored port and the mirrored direction is ingress or both (ingress and egress), the frame is sent to the monitoring port. If the destination port of the frame matches the mirrored port and the mirrored direction is egress or both, the frame is sent to the monitoring port. **Command mode:** Global configuration
`no port-mirroring monitor-port` *<port alias or number>* `mirroring-port` *<port alias or number>* Removes the mirrored port. **Command mode:** Global configuration
`no port-mirroring monitor-port` *<port alias or number>* Disables port mirroring. **Command mode:** Global configuration
`show port-mirroring` Displays the current settings of the monitoring port. **Command mode:** All except User EXEC

Access Control Configuration

Use these commands to create Access Control Lists. ACLs define matching criteria used for IP filtering and Quality of Service (QoS) functions.

For information about assigning ACLs to ports, see “Port Access List Configuration” on page 198.

Table 80. General ACL Configuration Options

Command Syntax and Usage
`[no] access-list ip` *<128 - 254>* `{extended|standard}` Configures an Access Control List. For the configuration options, see “Standard ACL IP Configuration” on page 201 and “Extended ACL IP Configuration” on page 201. **Command mode:** Global configuration
`[no] access-list ip` *<128 - 254>* `{extended|standard} statistics` Enables or disables statistics on the IP ACL. **Command mode:** Global configuration

Table 80. General ACL Configuration Options

Command Syntax and Usage
`[no] access-list mac extended` *<1-127>* Configures an Access Control List. For the configuration options, see "Extended ACL MAC Configuration" on page 207. **Command mode:** Global configuration
`[no] access-list mac extended` *<1-127>* `statistics` Enable or disable statistics for MAC ACL. **Command mode:** Global configuration
`show access-list [`*<1-254>*`|ip` *<128-254>*`|mac` *<1-127>*`]` Displays the current ACL parameters. **Command mode:** All

Standard ACL IP Configuration

The commands in Table 81 allow you to define filtering criteria for a standard Access Control List (ACL). Both the source and destination criteria are used for filtering packets.

Table 81. Standard ACL Configuration Options

Command Syntax and Usage
`{deny|permit}` *<source IP address> <subnet mask>* `{`*<destination IP address> <subnet mask>*`|any|host` *<host IP address>*`}` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets based on the source IP address. **Command mode:** Standard ACL configuration mode
`{deny|permit} any {`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Standard ACL configuration mode
`{deny|permit} host` *<source IP address>* `{`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action for packets that match the ACL definitions. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Standard ACL configuration mode

Extended ACL IP Configuration

These commands allow you to define filtering criteria for an extended ACL. Both the source and destination criteria are used for filtering packets.

Protocol Filtering Configuration

The following table includes ACL commands for filtering packets based on the protocol type.

Table 82. Protocol Filtering Configuration Options

Command Syntax and Usage
{deny|permit} *<protocol number>* *<source IP address>* {*<destination IP address>* *<subnet mask>*| any|host *<host IP address>*} Configures a filter action based on protocol type number 0 - 255. You can choose to permit (pass) or deny (drop) packets from a source IP address. **Command mode:** Extended ACL configuration mode
{deny|permit} *<protocol number>* any {*<destination IP address>* *<subnet mask>*| any|host *<host IP address>*} Configures a filter action based on protocol type number 0 - 255. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Extended ACL configuration mode
{deny|permit} *<protocol number>* host *<source IP address>* {*<destination IP address>* *<subnet mask>*| any|host *<host IP address>*} Configures a filter action based on protocol type number 0 - 255. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Extended ACL configuration mode

ICMP Filtering Configuration

The following table includes ACL commands for filtering packets based on ICMP protocol type. Both the source and destination criteria are used for filtering packets.

Table 83. ICMP Filtering Configuration Options

Command Syntax and Usage
{deny|permit} icmp *<source IP address>* *<subnet mask>* {*<destination IP address>* *<subnet mask>*| any|host *<host IP address>*} [*<message type>*] [*<message code>*] Configures a filter action based on ICMP protocol. You can choose to permit (pass) or deny (drop) packets from a source IP address. You can specify the optional message type and code with a value in the range 0 - 255. **Command mode:** Extended ACL configuration mode

Table 83. ICMP Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} icmp any` {<*destination IP address*> <*subnet mask*>`| any|host` <*host IP address*>} [<*message type*>] [<*message code*>] Configures a filter action based on ICMP protocol. You can choose to permit (pass) or deny (drop) packets from any source. You can specify the optional message type and code with a value in the range 0 - 255. **Command mode:** Extended ACL configuration mode
`{deny|permit} icmp host` <*source IP address*> {<*destination IP address*> <*subnet mask*>`| any|host` <*host IP address*>} [<*message type*>] [<*message code*>] Configures a filter action based on ICMP protocol. You can choose to permit (pass) or deny (drop) packets from a particular host. You can specify the optional message type and code with a value in the range 0 - 255. **Command mode:** Extended ACL configuration mode

IP Filtering Configuration

The following table includes ACL commands for filtering packets based on IP protocol type. Both the source and destination criteria are used for filtering packets.

Table 84. IP Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} ip` <*source IP address*> <*subnet mask*> {<*destination IP address*> <*subnet mask*>`| any|host` <*host IP address*>} Configures a filter action based on IP protocol. You can choose to permit (pass) or deny (drop) packets from a source IP address. **Command mode:** Extended ACL configuration mode
`{deny|permit} ip any` {<*destination IP address*> <*subnet mask*>`| any|host` <*host IP address*>} Configures a filter action based on IP protocol. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Extended ACL configuration mode
`{deny|permit} ip host` <*source IP address*> {<*destination IP address*> <*subnet mask*>`| any|host` <*host IP address*>} Configures a filter action based on IP protocol. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Extended ACL configuration mode

OSPF Filtering Configuration

The following table includes ACL commands for filtering packets based on OSPF protocol. Both the source and destination criteria are used for filtering packets.

Table 85. OSPF Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} ospf` *<source IP address> <subnet mask>* `{`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on OSPF protocol. You can choose to permit (pass) or deny (drop) packets from a source IP address. **Command mode:** Extended ACL configuration mode
`{deny|permit} ospf any {`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on OSPF protocol. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Extended ACL configuration mode
`{deny|permit} ospf host` *<source IP address>* `{`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on OSPF protocol. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Extended ACL configuration mode

PIM Filtering Configuration

The following table includes ACL commands for filtering packets based on PIM. Both the source and destination criteria are used for filtering packets.

Table 86. PIM Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} pim` *<source IP address> <subnet mask>* `{`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on PIM. You can choose to permit (pass) or deny (drop) packets from a source IP address. **Command mode:** Extended ACL configuration mode
`{deny|permit} pim any {`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on PIM. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Extended ACL configuration mode
`{deny|permit} pim host` *<source IP address>* `{`*<destination IP address> <subnet mask>*`| any|host` *<host IP address>*`}` Configures a filter action based on PIM protocol. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Extended ACL configuration mode

TCP Filtering Configuration

The following table includes ACL commands for filtering packets based on TCP. Both the source and destination criteria are used for filtering packets. For each of these commands, you can specify the following optional source and destination parameters:

Source Parameters:

- `eq` *<0 - 65535>*: Specify a port. Packets are allowed if the source port matches this specified port.
- `gt` *<0 - 65534>*: Specify a port. Packets are allowed if the source port is greater than this specified port.
- `lt` *<1 - 65535>*: Specify a port. Packets are allowed if the source port is less than this specified port.
- `range` *<0 - 65535> <0 - 65535>*: Specify a range of ports. Enter a starting and ending port number. Packets are allowed if the source port is in this specified range of ports.

Destination Parameters:

- `ack`: Check TCP ACK bit against the packet.
- `eq` *<0 - 65535>*: Specify a port. Packets are allowed if the destination port matches this specified port.
- `fin`: Check TCP FIN bit against the packet.
- `gt` *<0 - 65534>*: Specify a port. Packets are allowed if the destination port is greater than this specified port.
- `lt` *<1 - 65535>*: Specify a port. Packets are allowed if the destination port is less than this specified port.
- `psh`: Check TCP PSH bit against the packet.
- `range` *<0 - 65535> <0 - 65535>*: Specify a range of ports. Enter a starting and ending port number. Packets are allowed if the destination port is in this specified range of ports.
- `rst`: Check TCP RST bit against the packet.
- `syn`: Check TCP SYN bit against the packet.
- `urg`: Check TCP URG bit against the packet.

Table 87. TCP Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} tcp` *<source IP address> <subnet mask>* `{`*<destination IP address> <subnet mask>*`|any|host` *<host IP address>*`}` Configures a filter action based on TCP. You can choose to permit (pass) or deny (drop) packets from a source IP address. **Command mode:** Extended ACL configuration mode

Table 87. TCP Filtering Configuration Options

Command Syntax and Usage
`{deny|permit} tcp any {`*<destination IP address> <subnet mask>*`|any|host` *<host IP address>*`}` Configures a filter action based on TCP. You can choose to permit (pass) or deny (drop) packets from any source. **Command mode:** Extended ACL configuration mode
`{deny|permit} tcp host` *<host IP address>* `{`*<destination IP address> <subnet mask>*`|any|host` *<host IP address>*`}` Configures a filter action based on TCP. You can choose to permit (pass) or deny (drop) packets from a particular host. **Command mode:** Extended ACL configuration mode

UDP Filtering Configuration

The following table includes ACL commands for filtering packets based on UDP. Both the source and destination criteria are used for filtering packets. For each of these commands, you can specify the following optional source and destination parameters:

Source Parameters:

- `eq` *<0 - 65535>*: Specify a port. Packets are allowed if the source port matches this specified port.
- `gt` *<0 - 65534>*: Specify a port. Packets are allowed if the source port is greater than this specified port.
- `lt` *<1 - 65535>*: Specify a port. Packets are allowed if the source port is less than this specified port.
- `range` *<0 - 65535> <0 - 65535>*: Specify a range of ports. Enter a starting and ending port number. Packets are allowed if the source port is in this specified range of ports.

Destination Parameters:

- `eq` *<0 - 65535>*: Specify a port. Packets are allowed if the destination port matches this specified port.
- `gt` *<0 - 65534>*: Specify a port. Packets are allowed if the destination port if greater than this specified port.
- `lt` *<1 - 65535>*: Specify a port. Packets are allowed if the destination port if less than this specified port.
- `range` *<0 - 65535> <0 - 65535>*: Specify a range of ports. Enter a starting and ending port number. Packets are allowed if the destination port is in this specified range of ports

.

Table 88. UDP Filtering Configuration Options

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td>{deny|permit} udp <i>&lt;source IP address> <subnet mask&gt;</i> {<i>&lt;destination IP address> <subnet mask&gt;</i>|any|host <i>&lt;host IP address&gt;</i>}
Configures a filter action based on UDP. You can choose to permit (pass) or deny (drop) packets from a source IP address.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
<tr><td>{deny|permit} udp any {<i>&lt;destination IP address> <subnet mask&gt;</i>|any|host <i>&lt;host IP address&gt;</i>}
Configures a filter action based on UDP. You can choose to permit (pass) or deny (drop) packets from any source.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
<tr><td>{deny|permit} udp host <i>&lt;host IP address&gt;</i> {<i>&lt;destination IP address> <subnet mask&gt;</i>|any|host <i>&lt;host IP address&gt;</i>}
Configures a filter action based on UDP. You can choose to permit (pass) or deny (drop) packets from a particular host.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
</table>

Extended ACL MAC Configuration

These commands allow you to define filtering criteria for an extended MAC-based ACL. Both the source and destination criteria are used for filtering packets.

Table 89. MAC ACL Configuration Options

<table>
<tr><th>Command Syntax and Usage</th></tr>
<tr><td>{deny|permit} {any|host <i>&lt;source MAC address&gt;</i>} {any|host <i>&lt;destination MAC address&gt;</i>}
You can choose to permit (pass) or deny (drop) packets from any source or a particular host.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
<tr><td>{deny|permit} {any|host <i>&lt;source MAC address&gt;</i>} {any|host <i>&lt;destination MAC address&gt;</i>} [<i>&lt;user-defined protocol&gt;</i>] [user-priority <i>&lt;0 - 7&gt;</i>] [vlan <i>&lt;1 - 4094&gt;</i>]
You can choose to permit (pass) or deny (drop) packets from any source or a particular host based on a user-defined protocol. This protocol value can be in the range 1536 - 65535. You can also specify the optional filtering parameters based on user-priority and VLAN.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
<tr><td>{deny|permit} {any|host <i>&lt;source MAC address&gt;</i>} {any|host <i>&lt;destination MAC address&gt;</i>} [arp|ipv4|rarp|] [user-priority <i>&lt;0 - 7&gt;</i>] [vlan <i>&lt;1 - 4094&gt;</i>]
You can choose to permit (pass) or deny (drop) packets from any source or a particular host based the ARP or IPv4 or RARP protocols. You can also specify the optional filtering parameters based on user-priority and VLAN.
<b>Command mode:</b> Extended ACL configuration mode</td></tr>
</table>

Table 89. MAC ACL Configuration Options

Command Syntax and Usage
`{deny|permit} {any|host` <*source MAC address*>`} {any|host` <*destination MAC address*>`} [user-priority` <*0 - 7*>`]` You can choose to permit (pass) or deny (drop) packets from any source or a particular host. You can also specify the optional filtering parameter based on the user-priority. **Command mode:** Extended ACL configuration mode
`{deny|permit} {any|host` <*source MAC address*>`} {any|host` <*destination MAC address*>`} [vlan` <*1 - 4094*>`] [user-priority` <*0 - 7*>`]` You can choose to permit (pass) or deny (drop) packets from any source or a particular host. You can also specify the optional filtering parameter based on the user-priority. **Command mode:** Extended ACL configuration mode

QoS Configuration

Use the following commands to configure quality of service (QoS) parameters. .

Table 90. QoS Configuration Options

Command Syntax and Usage
`class-map [match-all|match-any]` *<1 - 32>* Create or modify a class map that defines a Class of Service (CoS). You can configure the following optional parameters: • `match-all`: Applies the class map to a packet if all of the criteria specified in the class map matches the packet. • `match-any`: Applies the class map to a packet if any of the criteria specified in the class map matches the packet. This is the default action. **Command mode:** Global configuration
`[no] policy-map` *<1 - 64>* Configure a policy map. Use the command to create the policy map and enter the policy map configuration mode. **Command mode:** Global configuration
`show policy-map` *<1 - 64>* Displays policy map information. **Command mode:** Global configuration
`show class-map` *<1 - 32>* Displays class map information. **Command mode:** Global configuration
`show running-config qos` Displays all QoS configuration information. **Command mode:** Global configuration

Class Map Configuration

The following commands in the class map configuration mode allow you to define a match criteria and apply it to the class map.

Table 91. QoS Class Map Configuration Options

Command Syntax and Usage
`[no] match access-group ip` *<128 - 254>* Specify the IP ACL that needs to be matched. **Command mode:** Class map configuration
`[no] match access-group mac` *<1 - 127>* Specify the MAC ACL that needs to be matched. **Command mode:** Class map configuration

Table 91. QoS Class Map Configuration Options (continued)

Command Syntax and Usage
`[no] match cos` *<cos number(s) or range>* Specify the CoS that needs to be matched. You can specify a single CoS value or a list of CoS values. **Command mode:** Class map configuration
`[no] match dscp` *<dscp number(s) or range>* Specify the Differentiated Services Code Point (DSCP) that needs to be matched. You can specify a single DSCP value or a list of DSCP values. **Command mode:** Class map configuration

Policy Map Configuration

Use the policy map configuration mode to assign a class map to a policy map.

Table 92. QoS Policy Map Configuration Options

Command Syntax and Usage
`[no] class` *<1 - 32>* Add a Quality of Service (QoS) class to the policy map. See Policy Map Class Configuration for more options on policy map class configuration. **Command mode:** Policy map configuration

Policy Map Class Configuration

In the policy map class configuration mode, you can configure the metering and remarking options, and assign CoS values, DSCP values, and priorities.

Following table lists the `set` command options available in the policy map class configuration mode.

Table 93. QoS Policy Map Class Configuration Set Options

Command Syntax and Usage
`[no] set cos` *<0 - 7>* Set IEEE 802.1q CoS class attributes to a packet. **Command mode:** Policy map class configuration
`[no] set ip dscp` *<0 - 63>* Assign IP DSCP values to a packet. **Command mode:** Policy map class configuration
`[no] set ip precedence` *<0 - 7>* Assign precedence value to a packet. **Command mode:** Policy map class configuration

Following table lists the `police` command options available in the policy map class configuration mode.

Table 94. QoS Policy Map Class Configuration Police Options

Command Syntax and Usage
[no] police cir <*cir-value*> {bc <*bc-value*>|be <*be-value*>} [be <*be-value*>] conform {set-cos-transmit <*cos-value*>|set-dscp-transmit <*dscp-value*>|transmit} [exceed <*exceed-options*> [violate <*violate-options*>]] Configure committed information rate (CIR), Burst Commit (BC), and Extended Burst (BE). If you configure the BC after configuring the CIR, you must also configure the BE. However, if you configure the BE after configuring the CIR, you need not configure the BC. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration
[no] police cir <*cir-value*> bc <*bc-value*> pir <*pir-value*> be <*be-value*> conform {set-cos-transmit <*cos-value*>|set-dscp-transmit <*dscp-value*>|transmit} [exceed <*exceed-options*> [violate <*violate-options*>]] Configure CIR, BC, Peak Information Rate (PIR), and BE. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration
[no] police cir <*cir-value*> bc <*bc-value*> pir <*pir-value*> conform {set-cos-transmit <*cos-value*>|set-dscp-transmit <*dscp-value*>|transmit} [exceed <*exceed-options*> [violate <*violate-options*>]] Configure CIR, BC, and PIR. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration
[no] police cir <*cir-value*> conform {set-cos-transmit <*cos-value*>|set-dscp-transmit <*dscp-value*>|transmit} [exceed <*exceed-options*> [violate <*violate-options*>]] Configure CIR. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration

Table 94. QoS Policy Map Class Configuration Police Options (continued)

Command Syntax and Usage
`[no] police cir <`*cir-value*`> pir <`*pir-value*`> conform {set-cos-transmit <`*cos-value*`>|set-dscp-transmit <`*dscp-value*`>|transmit} [exceed <`*exceed-options*`> [violate <`*violate-options*`>]]` Configure CIR and PIR. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration
`[no] police cir <`*cir-value*`> pir <`*pir-value*`> be <`*be-value*`> conform {set-cos-transmit <`*cos-value*`>|set-dscp-transmit <`*dscp-value*`>|transmit} [exceed <`*exceed-options*`> [violate <`*violate-options*`>]]` Configure CIR, PIR, and BE. Set the action to be performed if packets conform the configured rates. (Optional) Configure the action for packets that exceed or violate the configured rates. **Command mode:** Policy map class configuration

Syntax Description

`police`	Policies traffic based on a token bucket algorithm.
`cir`	Sets the Committed Information Rate (CIR). CIR is the minimum bandwidth you allocate to the network and is configured in bits per second (bps).
cir-value	1 - 1000000 bps
`bc`	Sets the Burst Commit (BC). BC is an allowance bandwidth that you allocate in case traffic exceeds the CIR.
bc-value	1536 - 16776960 bps
`pir`	Sets the Peak Information Rate (PIR). PIR is the allowed bandwidth in excess of the CIR. The PIR allows an overhead for the throughput when traffic exceeds the CIR.
pir-value	1 - 1000000 bps
`be`	Sets the Extended Burst (BE). BE is the bandwidth you allocate in cases where traffic exceed the burst commit. When you configure an extended burst, packets are dropped less aggressively.
be-value	1536 - 16776960 bps
`conform`	Sets the action to be performed on packets that are within the configured meter.
`set-cos-transmit`	Assigns a CoS value and transmits the packet.
cos-value	0 - 7
`set-dscp-transmit`	Assigns a DSCP value and transmits the packet.

dscp-value	0 - 63
`transmit`	Processes the packet.
`exceed`	(Optional) Specifies the action to be performed on packets that exceed the configured meter. Configure any one exceed option.
exceed-options	• `drop`: Drops the packet. • `set-cos-transmit` *<0-7>*: Assigns a CoS value and transmits the packet. • `set-dscp-transmit` *<0-63>*: Assigns a DSCP value and transmits the packet.
`violate`	(Optional) Sets the actions to be performed when a configured meter is violated. Configure any one violate option.
violate-options	• `drop`: Drops the packet. • `set-cos-transmit` *<0-7>*: Assigns a CoS value and transmits the packet. • `set-dscp-transmit` *<0-63>*: Assigns a DSCP value and transmits the packet.

iSwitch Configuration

The following table lists the iSwitch configuration commands.

Table 95. iSwitch Configuration Options

Command Syntax and Usage
`iswitch delports` *<range of ports>* Removes ports from standalone profile. Enter a range separated by a ‘-’. **Command mode:** Global configuration
`iswitch myvsidb` *<IP address>* `[path` *<VSI URI path>*`] [port` *<1 - 65535>*`]` Set VSI database IP address. Enter optional parameters: VSI path or port. **Command mode:** Global configuration
`[no] iswitch vcenter` *<vCenter IP address> <vCenter user name>* Configure VMware vCenter. **Command mode:** Global configuration
`[no] iswitch vds` *<vDS name> <Datacenter name>* Create a virtual distributed switch (vDS). **Command mode:** Global configuration
`[no] iswitch vnicprof` *<vNIC profile name>* `[numports` *<10 - 200>*`]` Configure a vNIC profile. (Optional) Specify the number of ports. The default is 20 ports. **Command mode:** Global configuration

Table 95. iSwitch Configuration Options

Command Syntax and Usage
`[no] iswitch uprof` *<uplink profile name>* Configure an uplink profile. **Command mode:** Global configuration
`iswitch addports` *<10 - 200>* `[`*<base port number>*`]` Add standalone ports. (Optional) Specify the base port number **Command mode:** Global configuration
`iswitch uports` *<1- 32>* Configure number of uplink ports. **Command mode:** Global configuration

iSwitch vNIC Profile Configuration

Use the following commands to configure vNIC profile.

Table 96. iSwitch vNIC Profile Configuration Options

Command Syntax and Usage
`addports` *<10 - 200>* `[`*<base port number>*`]` Specify the number of ports to add. Default number of ports added is 20. (Optional) Specify a base port number. **Command mode:** vNIC Profile configuration
`delports` *<range of ports>* Remove ports from a standalone profile. Enter ranges of 10 separated by a ‘,’. **Command mode:** vNIC Profile configuration
`pvid` *<port PVID>* Enter a numeric PVID for the port. Default PVID is 1. **Command mode:** vNIC Profile configuration
`tagpvid {0|1}` Enable or disable PVID tagging. Enter 1 to enable PVID tagging or 0 to disable PVID tagging. Default value is 0. **Command mode:** vNIC Profile configuration
`dot1p` *<0 - 7>* Assign a default 802.1p priority. Default priority is 0. **Command mode:** vNIC Profile configuration
`vlanlist` *<VLAN number(s)>* Assign VLANs to be enabled. Enter a single VLAN number or a range separated by a colon. **Command mode:** vNIC Profile configuration

Table 96. iSwitch vNIC Profile Configuration Options

Command Syntax and Usage
`addvlan` *<VLAN number(s)>* Assign VLANs to the vNIC profile. Enter a single VLAN number or a range separated by a colon. **Command mode:** vNIC Profile configuration
`remvlan` *<VLAN number(s)>* Remove VLANs assigned to the vNIC profile. Enter a single VLAN number or a range separated by a colon. **Command mode:** vNIC Profile configuration
`tagging {0|1}` Enable or disable egress VLAN tagging. Enter 1 to enable tagging or 0 to disable tagging. Default value is 0. **Command mode:** vNIC Profile configuration
`[no] vepa` Enable or disable VEPA on the vNIC profile. **Command mode:** vNIC Profile configuration
`vsitype` *<1 - 16777215>* Configure a VSI Type ID. **Command mode:** vNIC Profile configuration
`designated-uplinks` *<port number(s)>* Assign ports as designated uplink ports. Enter a single port or multiple ports separated by a ‘,’ or a range separated by a ‘-’. **Command mode:** vNIC Profile configuration
`designated-uplinks drop` Un-assign all ports from being designated uplink ports. **Command mode:** vNIC Profile configuration
`curr` Display current vNIC profile setting. **Command mode:** vNIC Profile configuration
`[no] service-policy {input|output}` *<policy-map number (1 - 64)>* Assign a QoS service policy to the in or out traffic of an interface. **Command mode:** vNIC Profile configuration
`[no] rate-limit {input|output}` *<1 - 1000000>* Configure rate limit in kilobits per second. **Command mode:** vNIC Profile configuration
`[no] access-list {mac` *<1 - 127>*`|ip` *<128 - 254>*`} in` Assign a MAC or IP ACL in the inward direction. **Command mode:** vNIC Profile configuration

Table 96. iSwitch vNIC Profile Configuration Options

Command Syntax and Usage
`no vsitype` Disable VSI Type setting on the vNIC profile. **Command mode:** vNIC Profile configuration
`no designated-uplinks` Disable designated uplinks setting on the vNIC profile. **Command mode:** vNIC Profile configuration

iSwitch Uplink Profile Configuration

Use the following commands to configure iSwitch Uplink profile.

Table 97. iSwitch Uplink Profile Configuration Options

Command Syntax and Usage
`lagmode {asymmetric|lacp|static}` Specify the link aggregation (LAG) mode for the Uplink profile. Default mode is asymmetric. **Command mode:** Uplink Profile configuration
`laghash {sip|smac|vport|dip|dsip|dmac|dsmac}` Specify the LAG hash basis. default is vport. Options include: • sip: source IP • smac: source MAC • vport: virtual port number • dip: destination IP • dsip: destination and source IP • dmac: destination MAC • dsmac: destination and source MAC **Command mode:**Uplink Profile configuration
`mtu` *<1500 - 9000>* Specify the maximum transmission frame size. Setting the MTU on the uplink port will set the MTU on every physical NIC attached to the uplink profile. The default MTU is 1500. **Command mode:** Uplink Profile configuration
`lacp mode {active|passive}` Configure the LACP mode. Default mode is active. **Command mode:** Uplink Profile configuration
`lacp system-priority` *<0 - 65535>* Configure the LACP system priority. Default priority is 32768. **Command mode:** Uplink Profile configuration

Table 97. iSwitch Uplink Profile Configuration Options

Command Syntax and Usage
`lacp timeout {short|long}` Configure LACP timeout. **Command mode:** Uplink Profile configuration
`curr` Display current Uplink profile setting. **Command mode:** Uplink Profile configuration

VSI Manager Configuration

Use the following command to configure the VSI Manager.

Table 98. VSI Manager Configuration Options

Command Syntax and Usage
`vsiman` Enter VSI Manager configuration mode. **Command mode:** Global configuration
`managerid` *<1 - 255>* Assign a VSI Manager ID. **Command mode:** VSI Manager configuration
`no managerid` Clear the VSI Manager ID. **Command mode:** VSI Manager configuration
`port` *<1 - 65534>* Assign a VSI service port. **Command mode:** VSI Manager configuration
`no port` Un-assign the VSI service port. **Command mode:** VSI Manager configuration
`[no] typeid` *<1 - 512>* `version` *<0 - 255>* Configure the type ID and version of the VSI Manager. **Command mode:** VSI Manager configuration

VSI Type ID Configuration

Use the following command to configure the VSI Type ID.

Table 99. VSI Type ID Configuration Options

Command Syntax and Usage
`acls` *<ACL number>* Assign ACLs by entering the ACL number. You can specify multiple ACLs separated by a ‘,’ or a range of ACLs separated by a ‘-’. **Command mode:** VSI Type ID configuration
`[no] action-setpriority acl` *<ACL number(s)>* `priority` *<priority level>* Set the ACL(s) priority. Enter ACL number or multiple ACL numbers separated by a comma (,). All packets matching ACL qualifiers will be marked with the specified packet priority. **Command mode:** VSI Type ID configuration
`name` *<VSI Type name>* Configure a name for the VSI Type. **Command mode:** VSI Type ID configuration
`qos rx-burst` *<32 - 4096>* Configure maximum burst size to receive in kilobits. Configure rate in multiples of 64. **Command mode:** VSI Type ID configuration
`qos rx-cbr` *<64 - 10000000>* Configure committed rate to receive in kilobits. Configure rate in multiples of 64. **Command mode:** VSI Type ID configuration
`qos tx-burst` *<32 - 4096>* Configure maximum burst size to transmit in kilobits. Configure rate in multiples of 64. **Command mode:** VSI Type ID configuration
`qos tx-cbr` *<64 - 10000000>* Configure committed rate to transmit in kilobits. Configure rate in multiples of 64. **Command mode:** VSI Type ID configuration
`vlans` *<VLAN number(s)>* Assign a single VLAN or multiple VLANs separated by a ‘,’ or assign a range of VLANs separated by a ‘-’. **Command mode:** VSI Type ID configuration
`no acls` Un-assign ACLs from the VSI type. **Command mode:** VSI Type ID configuration

Table 99. VSI Type ID Configuration Options

Command Syntax and Usage
`no qos {rx-burst|rx-cbr|tx-burst|tx-cbr}` Clear the QoS configurations assigned to the VSI type. **Command mode:** VSI Type ID configuration
`no vlans` Un-assign VLANs from the VSI type. **Command mode:** VSI Type ID configuration
`no name` Clear the name assigned to the VSI type. **Command mode:** VSI Type ID configuration

VLAN Configuration

These commands configure VLAN attributes, change the status of each VLAN, change the port membership of each VLAN, and delete VLANs.

By default, VLAN 1 is the only VLAN configured on the switch. All ports are members of VLAN 1 by default. Up to 2048 VLANs can be configured on the 5000V.

VLANs can be assigned any number between 1 and 4094.

Table 100. VLAN Configuration Options

Command Syntax and Usage
`vlan` *<VLAN number>* Enter VLAN configuration mode. **Command mode:** Global configuration
`name` *<1-32 characters>* Assigns a name to the VLAN or changes the existing name. The default VLAN name is the first one. **Command mode:** VLAN
`member` *<port alias or number>* Adds port(s) to the VLAN membership. You can add only ports 1 - 100 to a VLAN. Ports 101-4000 are vNIC profile ports and cannot be added to a VLAN. **Command mode:** VLAN
`no member` *<port alias or number>* Removes port(s) from this VLAN. **Command mode:** VLAN
`enable` Enables this VLAN. **Command mode:** VLAN

Table 100. VLAN Configuration Options (continued)

Command Syntax and Usage
`no enable` Disables this VLAN without removing it from the configuration. **Command mode:** VLAN
`no vlan` *<VLAN number>* Deletes this VLAN. **Command mode:** VLAN
`show vlan information` Displays the current VLAN configuration. **Command mode:** All

Note: All ports must belong to at least one VLAN. Any port which is removed from a VLAN and which is not a member of any other VLAN is automatically added to default VLAN 1. You cannot remove a port from VLAN 1 if the port has no membership in any other VLAN. Also, you cannot add a port to more than one VLAN unless the port has VLAN tagging turned `on`. Private VLANs cannot have VLAN 1 as their primary VLAN.

Private VLAN Configuration

Use the following commands to configure Private VLANs.

Table 101. Private VLAN Options

Command Syntax and Usage
`private-vlan type primary` Configures the VLAN type as a Primary VLAN. A Private VLAN must have only one primary VLAN. The primary VLAN carries unidirectional traffic to ports on the isolated VLAN or to community VLAN. **Command mode:** VLAN
`private-vlan type community` Configures the VLAN type as a community VLAN. Community VLANs carry upstream traffic from host ports. A Private VLAN may have multiple community VLANs. **Command mode:** VLAN
`private-vlan type isolated` Configures the VLAN type as an isolated VLAN. The isolated VLAN carries unidirectional traffic from host ports. A Private VLAN may have only one isolated VLAN. **Command mode:** VLAN
`private-vlan map` *<2 - 4094>* Set private VLAN mapping to a primary VLAN. **Command mode:** VLAN

Table 101. Private VLAN Options (continued)

Command Syntax and Usage
`private-vlan enable` Enables the private VLAN. **Command mode:** VLAN
`no private-vlan enable` Disables the Private VLAN. **Command mode:** VLAN
`show private-vlan [`*<2-4094>*`]` Displays current parameters for the selected Private VLAN(s). **Command mode:** VLAN

Management IP Interface Configuration

The following table describes the management IP interface Configuration commands. The following sections provide more detailed information and commands.

Table 102. Management IP Interface Configuration Commands

Command Syntax and Usage
`[no] interface ip-mgmt enable` Enable or disable the management IP Interface. **Command mode:** Global configuration
`no interface ip-mgmt` Delete the management IP Interface configuration. **Command mode:** Global configuration
`interface ip-mgmt address` *<IP address> <Subnet Mask> <Gateway IP address>* Set the IP address of the management IP Interface. **Command mode:** Global configuration
`interface ip-mgmt dhcp {enable|release|renew}` Configure DHCP. The configuration options include: – enable - Enables DHCP – release - releases a DHCP IP address – renew - renews the DHCP IP address **Command mode:** Global configuration
`no interface ip-mgmt dhcp enable` Disables DHCP. **Command mode:** Global configuration
`[no] interface ip-mgmt gateway enable` Enable or disable the management gateway. **Command mode:** Global configuration
`no interface ip-mgmt gateway` Delete the management gateway. **Command mode:** Global configuration
`interface ip-mgmt gateway` *<IP address>* Set the IP address of the management IP gateway. **Command mode:** Global configuration

Table 102. Management IP Interface Configuration Commands (continued)

Command Syntax and Usage
`interface ip-mgmt gateway` *<Subnet Mask>* Set the subnet mask of the management IP interface. **Command mode:** Global configuration
`show interface ip-mgmt` Displays management IP interface information. **Command mode:** Global configuration

ARP Configuration

Address Resolution Protocol (ARP) is the TCP/IP protocol that resides within the Internet layer. ARP resolves a physical address from an IP address. ARP queries machines on the local network for their physical addresses. ARP also maintains IP to physical address pairs in its cache memory. In any IP communication, the ARP cache is consulted to see if the IP address of the computer or the router is present in the ARP cache. Then the corresponding physical address is used to send a packet.

Table 103. ARP Configuration Options

Command Syntax and Usage
`ip arp rearp` *<2-120>* Defines re-ARP period, in minutes, for entries in the switch arp table. When ARP entries reach this value the switch will re-ARP for the address to attempt to refresh the ARP cache. The default value is 5 minutes. **Command mode:** Global configuration
`show ip arp` Displays the current ARP configurations. **Command mode:** All except User EXEC

ARP Static Configuration

Static ARP entries are permanent in the ARP cache and do not age out like the ARP entries that are learned dynamically. Static ARP entries enable the switch to reach the hosts without sending an ARP broadcast request to the network. Static ARPs are also useful to communicate with devices that do not respond to ARP requests. Static ARPs can also be configured on some gateways as a protection against malicious ARP Cache corruption and possible DOS attacks.

Table 104. ARP Static Configuration Options

Command Syntax and Usage
`ip arp` *<IP address> <MAC address> <IP interface number>* Adds a permanent ARP entry. **Command mode:** Global configuration
`no ip arp` *<IP address>* Deletes a permanent ARP entry. **Command mode:** Global configuration
`no ip arp all` Deletes all static ARP entries. **Command mode:** Global configuration
`no ip arp all` *<IP interface number>* Deletes all static ARP entries that use the interface. **Command mode:** Global configuration
`show ip arp` Displays current ARP configuration. **Command mode:** All

Domain Name System Configuration

The Domain Name System (DNS) commands are used for defining the primary and secondary DNS servers on your local network, and for setting the default domain name served by the switch services. DNS parameters must be configured prior to using hostname parameters with the `ping`, `traceroute`, and `tftp` commands.

Table 105. Domain Name Service Options

Command Syntax and Usage
`[no] ip dns primary-server` *<IP address>* Set the IP address for your primary DNS server, using dotted decimal notation. **Command mode:** Global configuration
`[no] ip dns secondary-server` *<IP address>* Set the IP address for your secondary DNS server, using dotted decimal notation. If the primary DNS server fails, the configured secondary will be used instead. **Command mode:** Global configuration

Table 105. Domain Name Service Options

Command Syntax and Usage
`[no] ip dns domain-name` *<string>* Sets the default domain name used by the switch. For example: `mycompany.com` **Command mode:** Global configuration
`ip dns resolve` *<String of up to 191 characters>* Set DNS resolution. **Command mode:** Global configuration
`show ip dns` Displays the current Domain Name System settings. **Command mode:** All except User EXEC

ERSPAN Configuration

Encapsulated Remote Switched Port Analyzer (ERSPAN) is used for remote port mirroring. ERSPAN uses a Generic routing Encapsulation (GRE) tunnel to carry traffic. An ERSPAN flow is an association between a source session and a destination IP. Table 106 lists the commands you can use to configure ERSPAN.

Table 106. ERSPAN Configuration Options

Command Syntax and Usage
`erspan` Enter ERSPAN configuration mode. **Command mode:** Global configuration
`no erspan` Delete ERSPAN configuration. **Command mode:** Global configuration
`[no] enable` Enable ERSPAN. **Command mode:** ERSPAN configuration
`exit` Exit from ERSPAN configuration mode. **Command mode:** ERSPAN configuration
`flow` *<1 - 1023> <ERSPAN source ID> <collector IP>* Create an ERSPAN flow. **Command mode:** ERSPAN configuration

Table 106. ERSPAN Configuration Options

Command Syntax and Usage
`[no] flow` *<1 - 1023>* `[`*<ERSPAN source ID> <collector IP>*`]` Delete an ERSPAN flow. (Optional) Specify the ERSPAN source ID. **Command mode:** ERSPAN configuration
`[no] source` *<ERSPAN source ID>* Create an ERSPAN source and enter the ERSPAN source configuration mode. Specify a string. **Command mode:** ERSPAN configuration

ERSPAN Source Configuration

Use the following commands to configure ERSPAN source.

Table 107. ERSPAN Source Configuration Options

Command Syntax and Usage
`add port` *<port number(s)>* Add ports to the ERSPAN source. You can add multiple ports separated by a ‘,’ or specify a range of ports separated by a ‘-’. **Command mode:** ERSPAN source configuration
`add vlan` *<VLAN number(s)>* Add VLAN to the ERSPAN source. You can add multiple VLANs separated by a ‘,’ or specify a range of VLANs separated by a ‘-’. **Command mode:** ERSPAN source configuration
`del port` *<port number(s)>* Add ports to the ERSPAN source. You can add multiple ports separated by a ‘,’ or specify a range of ports separated by a ‘-’. **Command mode:** ERSPAN source configuration
`del vlan` *<VLAN number(s)>* Add VLAN to the ERSPAN source. You can add multiple VLANs separated by a ‘,’ or specify a range of VLANs separated by a ‘-’. **Command mode:** ERSPAN source configuration
`direction {both|in|out}` Specify the direction of the traffic that needs to be captured. **Command mode:** ERSPAN source configuration
`mode {l2|l3}` Specify the mode that needs to be captured. **Note:** Only l2 mode is supported in this release. **Command mode:** ERSPAN source configuration

Table 107. ERSPAN Source Configuration Options

Command Syntax and Usage
`priority` *<0 - 7>* Specify the ERSPAN packet priority. **Command mode:** ERSPAN source configuration
`traffic {all|broadcast|unicast|multicast}` Specify the type of traffics that needs to be captured. **Command mode:** ERSPAN source configuration
`vlan` *<1 - 4094>* Specify the ERSPAN VLAN ID. **Command mode:** ERSPAN source configuration

sFlow Configuration

IBM N/OS supports sFlow version 5. sFlow is a sampling method used for monitoring high speed switched networks. Use these commands to configure the sFlow agent on the switch.

Table 108. sFlow Configuration Options

Command Syntax and Usage
`sflow` Enter sFlow configuration mode. **Command mode:** Global configuration
`no sflow` Disables the sFlow agent. **Command mode:** Global configuration
`agent-ip` *<IP address>* Configure the sFlow agent source IP. **Command mode:** sFlow configuration
`collector` *<IP address>* Configure the IP address of the global sFlow collector. **Command mode:** sFlow configuration
`counter-poll` *<20 - 65534>* Set counter poling interval in seconds. **Command mode:** sFlow configuration
`[no] group` *<1 - 31>* Create a unique sFlow instance for specific port(s) and/or VLAN(s). **Command mode:** sFlow configuration

Table 108. sFlow Configuration Options

Command Syntax and Usage
`[no] sample-rate` *<1-65534>* Set the sFlow sampling rate. One packet will be sampled for the sample size you specify. You must set the sample rate to enable sFlow engine. **Command mode:** sFlow configuration
`show sflow` Displays sFlow configuration parameters. **Command mode:** All

Configuration Dump

The dump program writes the current switch configuration to the terminal screen. To start the dump program, at the prompt, enter:

```
5000V(config)# show running-config
```

The configuration is displayed with parameters that have been changed from the default values. The screen display can be captured, edited, and placed in a script file, which can be used to configure other switches through a Telnet connection. When using Telnet to configure a new switch, paste the configuration commands from the script file at the command line prompt of the switch. The active configuration can also be saved or loaded via TFTP, as described on page 229.

Saving the Active Switch Configuration

When the `copy running-config` {`scp`|`tftp`} command is used, the switch's active configuration commands (as displayed using `show running-config`) will be uploaded to the specified script configuration file on the SCP/TFTP server. To start the switch configuration upload, at the prompt, enter:

```
5000V(config)# copy running-config tftp
```

The switch prompts you for the server address and filename.

Note: The output file is formatted with line-breaks but no carriage returns—the file cannot be viewed with editors that require carriage returns (such as Microsoft Notepad).

Note: If the SCP/ TFTP server is running SunOS or the Solaris operating system, the specified configuration file must exist prior to executing the `copy running-config` command and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current configuration data.

Restoring the Active Switch Configuration

When the `copy {scp|tftp} active-config` command is used, the active configuration will be replaced with the commands found in the specified configuration file. The file can contain a full switch configuration or a partial switch configuration.

To start the switch configuration download, at the prompt, enter:

```
5000V(config)# copy tftp active-config
```

The switch prompts you for the server address and filename.

You need to reload the switch for the restored configuration to take effect.

Chapter 17. Operations Commands

Operations commands generally affect switch performance immediately, but do not alter permanent switch configurations. For example, you can use Operations commands to immediately disable a port (without the need to apply or save the change), with the understanding that when the switch is reset, the port returns to its normally configured operation.

These commands enable you to alter switch operational characteristics without affecting switch configuration.

Table 109. General Operations Commands

Command Syntax and Usage
`password` *<1-128 characters>* Allows the user to change the password. You must enter the current password in use for validation. The switch prompts for a new password between 1-128 characters. **Command Mode**: Privileged EXEC
`console-log` Enables or disables session console logging. **Command Mode**: Privileged EXEC
`clear logging` Clears all Syslog messages. **Command Mode**: Privileged EXEC
`ntp send` Allows the user to send requests to the NTP server. **Command Mode**: Privileged EXEC

Operations-Level Port Commands

Operations-level port options are used for temporarily disabling or enabling a port, and for re-setting the port.

Table 110. Port Operations

Command Syntax and Usage
`no interface port` *<port number or alias>* `shutdown` Temporarily enables the port. The port will be returned to its configured operation mode when the switch is reset. **Command Mode**: Privileged EXEC
`interface port` *<port number or alias>* `shutdown` Temporarily disables the port. The port will be returned to its configured operation mode when the switch is reset. **Command Mode**: Privileged EXEC

Table 110. Port Operations

Command Syntax and Usage
`interface port` *<port number or alias>* `learning` Temporarily enables FDB learning on the port. **Command Mode**: Privileged EXEC
`no interface port` *<port number or alias>* `learning` Temporarily disables FDB learning on the port. **Command Mode**: Privileged EXEC

Chapter 18. Boot Options

To use the Boot Options commands, you must be logged in to the switch as the administrator. The Boot Options commands provide options for:

- Selecting a switch software image to be used when the switch is next reset
- Selecting a configuration block to be used when the switch is next reset
- Downloading or uploading a new software image to the switch via SCP/TFTP

The boot options are discussed in the following sections.

Scheduled Reboot of the Switch

This feature allows the switch administrator to schedule a reboot to occur at a particular time in future. This feature is particularly helpful if the user needs to perform switch upgrades during off-peak hours. You can set the reboot time, cancel a previously scheduled reboot, and check the time of the current reboot schedule.

Table 111. Scheduled Reboot Options

Command Syntax and Usage
`show boot` Displays the current switch reboot schedule. Includes information about the software version numbers and indicates whether or not the switch will be upgraded to a higher version after a reboot. **Command Mode**: All except User EXEC

```
5000V# show boot
Currently set to boot software image2, active config block
image1: version 1.0.0.1506, downloaded
image2: version 1.0.0.1506, downloaded
boot: version 1.0.0.1506
```

Updating the Switch Software Image

The switch software image is the executable code running on the IBM DS 5000V. As new versions of the image are released, you can upgrade the software.

Click on software updates. Use the following command to determine the current software version: `show boot`

Upgrading the software image on your switch requires the following:

- Loading the new image onto a SCP or TFTP server on your network
- Transferring the new image from the SCP or TFTP server to your switch
- Selecting the new software image to be loaded into switch memory the next time the switch is reset

Loading New Software to Your Switch

The switch can store up to two different software images, called `image1` and `image2`, as well as boot software, called `boot`. When you load new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

To load a new software image to your switch, you need the following:

- The image or boot software loaded on a TFTP/SCP server on your network
- The hostname or IP address of the TFTP/SCP server
- The name of the new software image or boot file

Note: The DNS parameters must be configured if specifying hostnames.

When the above requirements are met, use the following procedure to download the new software to your switch.

1. In Privileged EXEC mode, enter the following command:

```
5000V# copy {tftp|scp} {image1|image2}
```

2. Enter the hostname or IP address of the SCP or TFTP server.

```
Address or name of remote host: <IP address or hostname>
```

3. Enter the name of the new software file on the server.

```
Source file name: <filename>
```

   The exact form of the name will vary by server. However, the file location is normally relative to the SCP or TFTP directory (usually `tftpboot`).

4. Enter your username and password for the server, if applicable.

```
User name: {<username>|<Enter>}
```

5. The system prompts you to confirm your request.

Next, select a software image to run, as described in the following section.

Selecting a Software Image to Run

You can select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot.

1. In Global Configuration mode, enter:

```
5000V(config)# boot image {image1|image2}
```

2. Enter the name of the image you want the switch to use upon the next boot.

   The system informs you of which image set to be loaded at the next reset:

```
Next boot will use switch software image1 instead of image2.
```

Selecting a Configuration Block

When you make configuration changes to the IBM DS 5000V, you must save the changes so that they are retained beyond the next time the switch is reset. When you perform a save operation (`copy running-config startup-config`), your new configuration changes are placed in the *active* configuration block. The previous configuration is copied into the *backup* configuration block.

There is also a *factory* configuration block. This holds the default configuration set by the factory when your IBM DS 5000V was manufactured. Under certain circumstances, it may be desirable to reset the switch configuration to the default. This can be useful when a custom-configured IBM DS 5000V is moved to a network environment where it will be re-configured for a different purpose.

In Global Configuration mode, use the following command to set which configuration block you want the switch to load the next time it is reset:

```
5000V(config)# boot configuration-block {active|backup|factory}
```

Resetting the Switch

You can reset the switch to make your software image file and configuration block changes occur.

Note: Resetting the switch causes the Spanning Tree Group to restart. This process can be lengthy, depending on the topology of your network.

Enter the following command to reset (reload) the switch:

```
5000V# reload
```

You are prompted to confirm your request.

```
Reset will use software "image2" and the active config block.
>> Note that this will RESTART the Spanning Tree,
>> which will likely cause an interruption in network service.
Confirm reload (y/n) ?
```

Chapter 19. Maintenance Commands

The maintenance commands are used to manage dump information and forward database information. They include debugging commands to help with troubleshooting.

Dump information contains internal switch state data that is written to flash memory on the IBM DS 5000V after any one of the following occurs:

- The watchdog timer forces a switch reset. The purpose of the watchdog timer is to reboot the switch if the switch software freezes.
- The switch detects a hardware or software problem that requires a reboot.

To use the maintenance commands, you must be logged in to the switch as the administrator.

Table 112. General Maintenance Commands

Command Syntax and Usage
`show flash-dump` Displays dump information. **Command mode:** All
`copy flash-dump {tftp|scp}` Saves the system dump information via TFTP or SCP. For details, see page 238. **Command mode:** Privileged EXEC
`copy` *<STRING-scp>* Copy to a file on a remote host. Enter the complete SCP path including the file name. **Command mode:** Privileged EXEC
`copy` *<STRING-tftp>* Copy to a file on a TFTP server. Enter the complete TFTP path. **Command mode:** Privileged EXEC
`clear flash-dump` Clears dump information from flash memory. **Command mode:** Privileged EXEC
`show tech-support` Dumps all 5000V information, statistics, and configuration. You can log the output (`tsdmp`) into a file. **Command mode:** All except User EXEC
`copy tech-support {tftp|scp}` Redirects the technical support dump (tsdmp) to an external TFTP or SCP server. **Command mode:** Privileged EXEC

Table 112. General Maintenance Commands (continued)

Command Syntax and Usage
`copy tech-support` <*STRING-scp*> Copy the technical support dump to a file on a remote host. Enter the complete SCP path including the file name. **Command mode:** Privileged EXEC
`copy tech-support` <*STRING-tftp*> Copy the technical support dump to a file on a TFTP server. Enter the complete TFTP path including the file name. **Command mode:** Privileged EXEC
`copy support-archive {tftp|scp}` Copies archived support information to a TFTP or SCP server. **Command mode:** Privileged EXEC

ARP Cache Maintenance

Table 113 describes the ARP cache maintenance commands.

Table 113. Address Resolution Protocol Maintenance Options

Command Syntax and Usage
`show ip arp` Shows all ARP entries. **Command mode:** All
`clear ip arp-cache` Clears the entire ARP list from switch memory. **Command mode:** Privileged EXEC

Note: To display all or a portion of ARP entries currently held in the switch, you can also refer to “ARP Information” on page 141.

TFTP System Dump Put

Use these commands to `put` (save) the system dump to a TFTP server.

Note: If the TFTP server is running SunOS or the Solaris operating system, the specified `copy flash-dump tftp` file must exist *prior* to executing the `copy flash-dump tftp` command, and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current dump data.

To save dump information via TFTP, enter:

```
5000V# copy flash-dump tftp
```

You are prompted for the TFTP server IP address or hostname, and the *filename* of the target dump file.

Clearing Dump Information

To clear dump information from flash memory, enter:

```
5000V# clear flash-dump
```

The switch clears the dump region of flash memory and displays the following message:

```
FLASH dump region cleared.
```

If the flash dump region is already clear, the switch displays the following message:

```
FLASH dump region is already clear.
```

Unscheduled System Dumps

If there is an unscheduled system dump to flash memory, the following message is displayed when you log on to the switch:

```
Note: A system dump exists in FLASH. The dump was saved
      at 13:43:22 Wednesday January 30, 2011. Use show flash-dump
      uuencode to
      extract the dump for analysis and clear flash-dump to
      clear the FLASH region. The region must be cleared
      before another dump can be saved.
```

Part 4: Appendices

Appendix A. Getting Help & Technical Assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM *Documentation* CD that comes with your system.
- Go to the IBM support website at http://www.ibm.com/systems/support/ to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the documentation that is provided with your IBM product. The documentation that comes with IBM systems also describes the diagnostic tests that you can perform. Most systems, operating systems, and programs come with documentation that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

Using the documentation

Information about your IBM system and pre-installed software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, online documents, ReadMe files, and Help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to http://www.ibm.com/systems/support/ and follow the instructions. Also, some documents are available through the IBM Publications Center at http://www.ibm.com/shop/publications/order/.

Getting help and information on the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM systems, optional devices, services, and support. The address for IBM System x® and xSeries® information is http://www.ibm.com/systems/x/. The address for IBM BladeCenter information is http://www.ibm.com/systems/bladecenter/. The address for IBM IntelliStation® information is http://www.ibm.com/intellistation/.

You can find service information for IBM systems and optional devices at http://www.ibm.com/systems/support/.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with System x and x Series servers, BladeCenter products, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, see http://www.ibm.com/services/sl/products/.

For more information about Support Line and other IBM services, see http://www.ibm.com/services/, or see http://www.ibm.com/planetwide/ for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to http://www.ibm.com/partnerworld/ and click **Find Business Partners** on the right side of the page. For IBM support telephone numbers, see http://www.ibm.com/planetwide/. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

IBM Taiwan product service

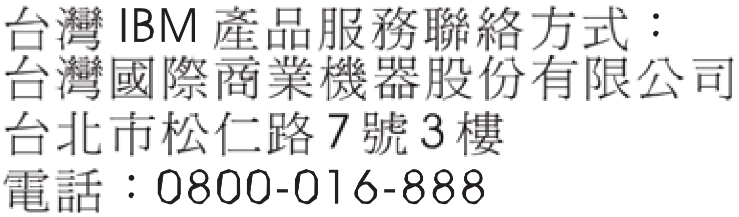

台灣 IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation
3F, No 7, Song Ren Rd.
Taipei, Taiwan
Telephone: 0800-016-888

Appendix B. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product, and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml.

Important Notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

Maximum memory might require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

Particulate contamination

Attention: Airborne particulates (including metal flakes or particles) and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might pose a risk to the devices that run the software described in this document. Risks that are posed by the presence of excessive particulate levels or concentrations of harmful gases include damage that might cause the device to malfunction or cease functioning altogether. This specification sets forth limits for particulates and gases that are intended to avoid such damage. The limits must not be viewed or used as definitive limits, because numerous other factors, such as temperature or moisture content of the air, can influence the impact of particulates or environmental corrosives and gaseous contaminant transfer. In the absence of specific limits that are set forth in this document, you must implement practices that maintain particulate and gas levels that are consistent with the protection of human health and safety. If IBM determines that the levels of particulates or gases in your environment have caused damage to the device, IBM may condition provision of repair or replacement of devices or parts on implementation of appropriate remedial measures to mitigate such environmental contamination. Implementation of such remedial measures is a customer responsibility.

Contaminant	**Limits**
Particulate	• The room air must be continuously filtered with 40% atmospheric dust spot efficiency (MERV 9) according to ASHRAE Standard 52.2[1]. • Air that enters a data center must be filtered to 99.97% efficiency or greater, using high-efficiency particulate air (HEPA) filters that meet MIL-STD-282. • The deliquescent relative humidity of the particulate contamination must be more than 60%[2]. • The room must be free of conductive contamination such as zinc whiskers.
Gaseous	• Copper: Class G1 as per ANSI/ISA 71.04-1985[3] • Silver: Corrosion rate of less than 300 Å in 30 days

[1] ASHRAE 52.2-2008 - *Method of Testing General Ventilation Air-Cleaning Devices for Removal Efficiency by Particle Size*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

[2] The deliquescent relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote ionic conduction.

[3] ANSI/ISA-71.04-1985. *Environmental conditions for process measurement and control systems: Airborne contaminants*. Instrument Society of America, Research Triangle Park, North Carolina, U.S.A.

Documentation format

The publications for this product are in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when you use the PDF files and want to request a web-based format or accessible PDF document for a publication, direct your mail to the following address:

Information Development
IBM Corporation
205/A0153039 E. Cornwallis Road
P.O. Box 12195
Research Triangle Park, North Carolina 27709-2195
U.S.A.

In the request, be sure to include the publication part number and title.

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Index