



Citrix XenServer® 7.1 vSwitch Controller User Guide

Published August 2017
1.0 Edition



Citrix XenServer ® 7.1 vSwitch Controller User Guide

Copyright © 2017 Citrix Systems, Inc. All Rights Reserved.
Version: 7.1.1

Citrix, Inc.
851 West Cypress Creek Road
Fort Lauderdale, FL 33309
United States of America

Disclaimers

This document is furnished "AS IS." Citrix, Inc. disclaims all warranties regarding the contents of this document, including, but not limited to, implied warranties of merchantability and fitness for any particular purpose. This document may contain technical or other inaccuracies or typographical errors. Citrix, Inc. reserves the right to revise the information in this document at any time without notice. This document and the software described in this document constitute confidential information of Citrix, Inc. and its licensors, and are furnished under a license from Citrix, Inc.

Citrix Systems, Inc., the Citrix logo, Citrix XenServer and Citrix XenCenter, are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.

Trademarks

Citrix®
XenServer®
XenCenter®

This product contains an embodiment of the following patent pending intellectual property of Citrix Systems, Inc.:

1. United States Non-Provisional Utility Patent Application Serial Number 11/487,945, filed on July 17, 2006, and entitled "Using Writable Page Tables for Memory Address Translation in a Hypervisor Environment".
2. United States Non-Provisional Utility Patent Application Serial Number 11/879,338, filed on July 17, 2007, and entitled "Tracking Current Time on Multiprocessor Hosts and Virtual Machines".



Contents

1. Introduction	1
1.1. vSwitch and Controller for XenServer	1
2. Getting Started	2
2.1. Deploying the vSwitch Controller Virtual Appliance	2
2.2. Accessing the vSwitch Controller Command Line Interface	3
2.3. Accessing the vSwitch Controller Graphical User Interface	3
2.3.1. Accessing the vSwitch Controller GUI Remotely	3
2.4. Configuring the vSwitch Controller IP Address	3
2.5. Adding Resource Pools	4
2.6. Configuring High Availability	4
3. vSwitch Management	5
3.1. Interface Overview	5
3.1.1. Top Panel	5
3.1.2. Status Bar	5
3.1.3. Top Icons	6
3.1.4. Side Panel	6
3.1.5. Using the Resource Tree	6
3.1.5.1. Color-Coded Icons	7
3.1.6. Main Panel Data Area	7
3.2. Using the Dashboard to Monitor Network Activity	8
3.2.1. Server Statistics	8
3.2.2. Network Statistics	8
3.2.3. Recent Network Events	8
3.2.4. Recent Administrative Events	8
3.2.5. Throughput, Flows, and Bit Rate Graphs	9
4. Virtual Network Visibility & Control	10
4.1. Viewing Status	10
4.1.1. Global Level	10

4.1.2. Resource Pool Level	10
4.1.2.1. Fail safe mode	11
4.1.3. Server Level	11
4.1.4. Network Level	12
4.1.5. Virtual Machine (VM) Level	12
4.1.6. Virtual Interface (VIF) Level	13
4.1.7. Viewing Flow Statistics	13
4.2. Managing Address Groups	15
4.3. Managing Virtual Machine Groups	15
4.4. DVS Policy Configuration Hierarchy	16
4.5. Setting Up Access Control Policies	16
4.5.1. Global Access Control List (ACL) Rules	17
4.5.2. Resource Pool Access Control List (ACL) Rules	17
4.5.3. Network Access Control List (ACL) Rules	17
4.5.4. VM Access Control List (ACL) Rules	17
4.5.5. VIF Access Control List (ACL) Rules	17
4.5.6. Access Control List (ACL) Rule Enforcement Order	18
4.5.7. Defining Access Control List (ACL) Rules	18
4.6. Setting Up Port Configuration Policies	20
4.6.1. Configuring QoS. Configuring QoS	21
4.6.2. Configuring RSPAN	22
4.6.2.1. Identify your RSPAN VLAN	22
4.6.2.2. Configure the Physical Network with the Target VLAN	22
4.6.2.3. Configure vSwitch Controller with the Target VLAN	22
4.6.2.4. Modify port configuration to enable RSPAN for a set of VIFs	23
4.6.2.5. Configuring MAC Address Spoof Checking	23
4.6.2.6. Save Changes	23
5. vSwitch Controller Administration & Maintenance	24
5.1. Configuring IP Address Settings	24
5.2. Configuring the Controller Hostname	25
5.3. Collecting Information for Trouble Reports	25



5.4. Restarting the vSwitch Controller Software	25
5.5. Managing Administrative Accounts	25
5.6. Managing Configuration Snapshots	26
5.7. Adding Network Time Protocol (NTP) Servers	26
5.8. Exporting Syslog Files	27
6. Troubleshooting vSwitch Controller Issues	28
6.1. Resource Tree Node Status	28
6.2. Troubleshooting Access Policy Issues	29
6.3. Creating a Trouble Report	29
6.4. Controller Error Messages	30
7. Command Line Interface	31
7.1. CLI Commands	31
7.1.1. To terminate the current CLI session	31
7.1.2. To halt the vSwitch Controller	31
7.1.3. To get information on commands	31
7.1.4. To upgrade or downgrade the existing version of the Controller	31
7.1.5. To ping a specified remote system	31
7.1.6. To restart the Controller	31
7.1.7. To restart the Controller daemon	32
7.1.8. To set the hostname of the controller appliance	32
7.1.9. To set the IP address of the Controller management interface via DHCP	32
7.1.10. To set a static IP address for the Controller management interface	32
7.1.11. To display the current Controller hostname	32
7.1.12. To display a summary of the current configuration and status of the management interface	32
7.1.13. To display configuration values for the management interface	32
7.1.14. To display the current default gateway for the Controller	32
7.1.15. To display the current DNS configuration for the Controller	33
7.1.16. To display the current IP address of the Controller management interface	33
7.1.17. To display the current netmask of the Controller management interface	33
7.1.18. To display the software version of the Controller	33



Chapter 1. Introduction

The XenServer platform is a server virtualization platform for server and client operating systems that virtualizes each physical host on which it is installed, enabling a single physical machine to run multiple virtual machines (VMs) simultaneously.

XenServer allows you to combine multiple XenServer hosts into a *resource pool*, using industry-standard shared storage architectures and Citrix resource clustering technology. Resource pooling extends the basic single-server notion of virtualization to multiple servers, with VMs able to run on any server in the pool and even move between different servers in the pool using a technology called *live migration*. Each resource pool includes a master server, which stores configuration for all physical hosts and VMs in the pool.

XenCenter is a Windows-based management application that allows IT managers to create XenServer resource pools and to manage them and their resources from a single point of control. XenCenter provides a graphical interface to perform many of the same VM, storage, and clustering configuration operations that can be performed using the “xe” utility on the XenServer command line.

1.1. vSwitch and Controller for XenServer

The vSwitch brings visibility, security, and control to XenServer virtualized network environments. It consists of a virtualization-aware switch (the *vSwitch*) running on each XenServer and the *vSwitch Controller*, a centralized server that manages and coordinates the behavior of each individual vSwitch to provide the appearance of a single vSwitch.

The vSwitch Controller supports fine-grained security policies to control the flow of traffic sent to and from a VM and provides detailed visibility into the behavior and performance of all traffic sent in the virtual network environment. A vSwitch greatly simplifies IT administration within virtualized networking environments, as all VM configuration and statistics remain bound to the VM even if it migrates from one physical host in the resource pool to another.



Chapter 2. Getting Started

This chapter describes how to get started using the vSwitch Controller. Refer to the Release Notes for instructions on enabling the DVS vSwitch on the XenServers of a resource pool. The information in this chapter assumes that you have at least one XenServer resource pool configured in XenCenter and that you have sufficient capacity within that pool to deploy the vSwitch Controller virtual appliance VM. The requirements for controller deployment are described in the next section.

Setting up the vSwitch Controller involves the following tasks:

1. [Deploying the vSwitch Controller Virtual Appliance](#)
2. [Accessing the vSwitch Controller](#)
3. [Configuring the Controller IP Address](#)
4. [Adding Resource Pools](#)
5. [Configuring High-Availability \(optional\)](#)

Note:

This version of vSwitch Controller supports the following XenServer versions: 7.1, 7.0, 6.5.0, 6.2.0, 6.1.0, and 6.0.

2.1. Deploying the vSwitch Controller Virtual Appliance

The XenServer that runs the vSwitch Controller must meet the following minimum requirements:

- 2 CPUs
- 2GB DRAM
- 16GB Disk

The minimum allowed VM configuration for the vSwitch Controller appliance and the default configuration on import is:

- 2 vCPUs
- 2GB DRAM
- 16GB Disk

This configuration will support deployments up to 16 XenServers and 256 Virtual Interfaces (vifs) connected to the vSwitch Controller. For larger deployments (up to the maximum supported limit of 64 XenServers and 1024 vifs), the VM configuration should be modified to:

- 4 vCPUs
- 4GB DRAM
- 16GB Disk

The vSwitch Controller VM may run within a resource pool that it manages. Generally, this configuration runs as if the vSwitch Controller VM was running separately. However, it may take slightly longer (up to 2 minutes) to connect all the vSwitches in the event of a Controller migration or restart. This is due to differences in how the individual vSwitches route control connections.

To install the vSwitch Controller, import the supplied virtual appliance VM image into a XenServer resource pool. During import, attach the single VIF of the imported VM to a network through which the XenServer or XenServer pool to be controlled by the VM is reachable. Refer to the XenServer documentation for more information.

After the VM has been imported, start it to begin the process of configuring the DVS.

2.2. Accessing the vSwitch Controller Command Line Interface

You can access the vSwitch Controller command line interface (CLI) from within XenCenter or remotely using an SSH client. When the vSwitch Controller VM first boots, the text console within XenCenter will display a message indicating the IP address that can be used to access the controller remotely. If the VM did not receive an IP address, the text console will indicate that an address must be assigned through the CLI. In either case, the text console will display a login prompt to log into the CLI locally in the XenCenter console. Full documentation of the available CLI commands is included in [Chapter 7](#).

2.3. Accessing the vSwitch Controller Graphical User Interface

You can access the vSwitch Controller graphical user interface (GUI) remotely using a web browser. When the vSwitch Controller VM boots, the text console within XenCenter will display a message indicating the IP address that can be used to access the GUI remotely. If the VM did not receive an IP address, the GUI can not be used locally or remotely until one is assigned. The text console will provide instructions on setting the IP address locally in the command line interface. Once the controller VM has the IP address, the GUI can be accessed locally within the XenCenter console by following the steps in the next section.

Note:

Since VNC is disabled, vSwitch Controller GUI can be accessed only from a web browser.

2.3.1. Accessing the vSwitch Controller GUI Remotely

To access the vSwitch Controller interface remotely:

1. Open a browser and enter the following URL, where *server* is the IP address or host name of the interface of the controller VM: `https://server:8080/`
2. Enter your user name and password, and click **Login**. The default user name and password are **admin** and **admin**.

Note:

By default, the vSwitch Controller webserver uses a self-signed certificate, which will cause many browsers to show a security error when connecting to the GUI. You can safely ignore the error and install the certificate into your browser.

The following browsers are supported: Firefox 3.x, Safari 4.x, Internet Explorer 7 and 8. Other modern browsers of similar capability (such as Opera or Google Chrome) are not supported, but may work as well. Internet Explorer 9 addresses known IE memory and resource leak issues; however it has not received full testing.

When you log in for the first time, the system prompts you to change the default admin password. It is important that you create a strong admin password to protect the security of your virtualized infrastructure.

2.4. Configuring the vSwitch Controller IP Address

When the vSwitch Controller is started for the first time, it attempts to obtain an IP address using DHCP; however, we recommend that you assign a static IP address. If DHCP is configured, resource pools cannot be set to Fail-Safe mode

To assign a static IP address:

1. Access the vSwitch Controller interface locally, as described in the previous section.
2. Click the **Settings** tab and then **IP Configuration** in the side panel. The current settings are shown.
3. Click **Modify Configuration**, specify the new IP address information, and click **Make Changes**.



Note:

If DHCP is configured, resource pools cannot be set to Fail-Safe Mode.

2.5. Adding Resource Pools

Adding a resource pool allows the vSwitch Controller to automatically begin managing all XenServer hosts in that pool.

To add a resource pool:

1. Under **Visibility & Control**, open the **Status** tab and choose **All Resource Pools** in the resource tree (side panel) to open the Status page for all resource pools.
2. Click **Add Resource Pool**. An error message is displayed if you do not have the correct license to add an additional resource pool.
3. Enter the IP address or DNS name of the master XenServer in the **Pool Master Server (DNS/IP)** field.
4. Enter the username and password for administrative access to the server.

The user must have full management capabilities in the resource pool. The vSwitch Controller will not be able to properly manage the pool if the account has restricted capabilities.

Typically, this will be the user named "root" but could be a different name if the RBAC features of the XenServer platform are in use.

5. Select the **Steal** check box only if you want to override any existing vSwitch Controller configuration that was previously set for this resource pool.
6. Click **Connect**.

The vSwitch Controller will use the provided username and password to communicate with the pool master server using the XAPI protocol. When communications are established, the new resource pool is added to the resource tree, along with all of the associated resources. If the vSwitch Controller VM is unable to communicate with the pool master, it displays an error message describing the failure.

Note:

For the vSwitch Controller and the XenServer resource pool to communicate with each other, the XenServer resource pool must be using the **Backwards compatibility mode** (which is the default). You can specify this setting on the **Pool Properties** page in XenCenter. For more information, see XenCenter Help.

2.6. Configuring High Availability

To ensure that XenServers can always reach an active vSwitch Controller, we recommend the use of Citrix High Availability for the vSwitch Controller VM. Refer to the *XenServer Administrator's Guide* for instructions on enabling high availability. Because continuous operation of the vSwitch Controller is critical to the operation of networking for all virtual machines, the vSwitch Controller VM restart-priority should be set to 1 and ha-always-run should be set to true.


Chapter 3. vSwitch Management

The vSwitch Controller GUI allows you to view status and flow statistics for elements within the virtual network, set up VM access control, QoS, and traffic mirroring policies, and modify configuration of the vSwitch Controller virtual appliance.

3.1. Interface Overview

The vSwitch Controller GUI is divided into the three different panels, as shown in the next figure.

Figure 3.1.



3.1.1. Top Panel

The top panel is always visible when using the GUI and includes a status bar and a set of main navigation icons.

3.1.2. Status Bar

The gray status bar at the top of the vSwitch Controller window contains the following information and functions (left to right):

- Version: Current vSwitch Controller version.
- Online Help: Click to display or close an online help area near the top of the controller window.
- Logout: Click to log out of the vSwitch Controller GUI.
- User: Displays the user name of the user that is currently logged in.
- Refresh icon: Click to manually update the information on the page.
- Play/Pause: Click to toggle whether the GUI should automatically refresh data on the screen using background updates. In play mode, the data that is shown refreshes automatically every 15 seconds. In pause mode, most data is not updated; however, a few elements are updated, notably the resource tree. The status bar



background behind the buttons turns orange and a “Data Updates Paused” indicator appears in the status bar when in pause mode.

3.1.3. Top Icons

Click the top icons to access the major functional areas within the GUI.

- Dashboard: View summary statistics and information about network and administrative events. See [Monitoring Network Status with the Dashboard](#).
- Visibility and Control: View network status and statistics or configure access control, QoS and traffic mirroring policies for virtual networks. See [Viewing and Controlling the Network](#).
- Settings: Perform vSwitch Controller maintenance and administrative functions. See [Administering and Maintaining the vSwitch Controller](#).

3.1.4. Side Panel

The side panel is available only in the Visibility and Control and Settings section.

For the Visibility and Control section, the side panel contains a resource tree that you can use to browse network elements within the virtual network environment. Similar to the resource tree in XenCenter, elements are organized hierarchically and provide an easy way to browse elements within the system. To expand a section of the resource tree, click the side-facing arrow next to the node text. An expanded node is marked with a down-facing arrow, which you can click to collapse.

When you select an element from the resource tree, the main panel displays status and configuration data for that node in the tree. For example, if you select a VM from the resource tree and choose **Status** in the **Visibility and Control** section, the main panel displays status information about the selected VM.

The resource tree includes a search function. To filter the contents based on a search string, enter text in the search field, and press **Enter**. Click the **X** symbol to clear the search. Searches support wildcards (* for one or more characters and ? for a single character). If wildcards are not used, the system performs a substring search as if a * wildcard were entered at the start and end of the search string. For example, the search “Lab” finds all items with “Lab” in the name, such as “Laboratory-1” and “New-Lab-5.”

For the Settings section, the side panel contains icons to select which area of vSwitch Controller configuration the user would like to view or modify

3.1.5. Using the Resource Tree

At the highest level, the resource tree displays the following items:

- All Resource Pools: List of all the available resource pools. This is the top-level resource for exploring all XenServers, Networks, VMs, and VIFs that are part of each resource pool.
- Address Groups: Named sets of IP addresses and subnet ranges to be used to limit the application of a rule in the access control section or to limit the scope of a query in the **Flow Statistics** section.
- VM Groups: Named sets of VMs to be used to simplify viewing the status and flow statistics of a particular collection of VMs.

When you expand a resource pool in the resource tree, the following items are displayed:

- Pool-wide networks: This list includes all networks in the resource pool and is similar to the list in the Network tab of XenCenter. You can expand the list to show the individual networks, expand a network to show the VMs on that network, and expand a VM to show its VIFs on that network.
- XenServers. This list is similar to the server hierarchy in XenCenter. You can expand the list to show all of the servers in the pool and expand a single server entry to show the networks, VMs, and VIFs associated with the server. The Server Networks listing is similar to what you see if you click a server in XenCenter and choose the Network tab.



- All VMs: This list shows all VMs in the resource pool, whether or not they are configured for a single server. You can expand the list to show the individual VMs, and expand a VM to show its VIFs.

Right-click context menus on nodes are available on most nodes to provide a simple way of adding, modifying, and deleting items in the resource tree.

3.1.5.1. Color-Coded Icons

Color-coded icons in the resource tree show the state of tree nodes under the top-level “All Resource Pools” node. Similar to XenCenter, these color codes are based on data retrieved via XAPI from each pool master. When a node state changes, the icon is updated as follows:

- Green: A green icon indicates that the resource is active on the network and properly managed by the vSwitch Controller.
- Red: For a Resource Pool node, the red indicates that a XAPI connection could not be established to the pool master. If the Resource Pool node is green, a red icon for any node below it indicates the element is not currently active on the network (it is powered off or disconnected).
- Orange: An orange icon indicates that the node, or one of its descendants, is not properly connected or managed. The status page for the associated resource will display an error message describing the problem.

The color codes on the tree menu items are also displayed on the Status page for the node. Refer to [Troubleshooting vSwitch Controller Issues](#) for detailed information on the color codes and status information.

3.1.6. Main Panel Data Area

The main panel data area contains status information, statistics, and configuration settings.

- Dashboard: There is no side menu and the main panel data area takes up the full area below the top panel. The dashboard main panel provides an overview of all virtual networks managed by the vSwitch Controller.
- Visibility and Control: The main panel takes up the right side of the window below the top panel and includes tabs at the top that correspond to the following major visibility and control functions:
 - Status: View detailed status information for the selected resource tree node.
 - Flow Statistics: View a graph and data on network activity for the selected node.
 - Access Control: Set up access control policies for the selected node.
 - Port Configuration: Set up quality of service (QoS) and traffic mirroring policies for the selected node.
- Settings: The main panel takes up the right side of the window below the top panel. The setting main panel displays details for viewing or configuring vSwitch Controller settings based on the subsection selected in the side panel.

Within the Visibility and Control section, the type of data displayed in the main panel changes to reflect the hierarchy level as well as the specific item that you selected in the side panel.

For example, if you select a resource pool in the side panel and click the Access Control tab, the main panel displays the following:

- The global access control security policy
- The policy for the selected resource pool

If you select a virtual interface (VIF) from the side panel and click the Access Control tab, the main panel displays:

- The global access control security policy
- The policy for the resource pool that contains the VIF
- The policy for the VM that contains the VIF



- The policy for the selected VIF

3.2. Using the Dashboard to Monitor Network Activity

The dashboard presents summary statistics and information about events within the virtual network environment. To display the dashboard, click the **Dashboard** icon at the top of the vSwitch Controller interface.

The dashboard is divided into the areas described in this section. The information is automatically updated every few seconds.

3.2.1. Server Statistics

This section presents the following general information about the vSwitch Controller.

- Up Time: Length of time since the vSwitch Controller was last started.
- CPU Load: Current percent of CPU utilization for the vSwitch Controller virtual appliance.

3.2.2. Network Statistics

This section shows an inventory of network elements (resource pools, XenServers, networks, and VMs) For each of the following categories:

- Managed: Number of elements of this type that are in a running state according to XAPI and currently managed by the vSwitch Controller.
- Active: Number of elements of this type that are in a running state according to XAPI. Includes managed and unmanaged elements.
- Total: Number of elements of this type (active or not) that are known to exist via XAPI.

When the system is configured and operating correctly, the managed and active counts are the same. The total count is always equal to or greater than the managed and active count, because components that are powered off are not shown as being managed by the controller.

3.2.3. Recent Network Events

This section lists the most recent events that have occurred within the managed virtual networks since the vSwitch Controller was last restarted. Use the scroll bar on the right to scroll through the list. The most recent event is listed first. Over time, older events are deleted from the list.

The following information is reported for each network event:

- Priority: Relative importance of the event.
- Date/Time: Date and time that the event occurred.
- Event: Description of the event. You can click on hyperlinks in an event description to access the corresponding Visibility and Control Status pages of network elements mentioned in the event.

Network events can be exported to a syslog server for a more permanent record. Refer to [Exporting Syslog Files](#).

3.2.4. Recent Administrative Events

This section lists events that have occurred within the vSwitch Controller, often as a result of an administrator changing configuration within the GUI. Use the scroll bar on the right to scroll through the list. The most recent event is listed first. Over time, older events are deleted from the list.

The following information is reported for each administrative event:

- Priority: Relative importance of the event.



- Date/Time: Date and time that the event occurred.
- Event: Description of the event. You can click on hyperlinks in a event description to access the Visibility and Control Status pages of network elements mentioned in the event.

Network events can be exported to a syslog server for a more permanent record. Refer to [Exporting Syslog Files](#).

3.2.5. Throughput, Flows, and Bit Rate Graphs

These graphs display information about the behavior of the most active VMs and protocols.

The graphs display the following information:

- Aggregate Throughput (bits/sec) for the last hour
- Aggregate Packet Rate (packets/sec) for the last hour
- Aggregate Connection Rate (flows/sec) for the last hour

Chapter 4. Virtual Network Visibility & Control

The Visibility and Control section allows you to monitor network behavior and configure network policy. To access the pages, click the **Visibility and Control** icon at the top of the vSwitch Controller interface.

4.1. Viewing Status

The **Status** tab provides detailed information in table form about the node that is selected in the resource tree. The type of information that is presented varies according to the selected node. Most individual table entries are links that you can click to display the status page that applies to that entry.

All byte counts and error counts continue to accumulate even if a XenServer node is restarted or a VM restarts or migrates. The color codes follow the same rules as the color codes in the side panel. See [Color-Coded Icons](#).

4.1.1. Global Level

At the global level, the Status page presents a table listing all resources pools with the following information:

- Resource pool: Name of the resource pool.
- # Servers: Number of servers in the pool.
- # Networks: Number of networks in the pool.
- # VMs: Number of VMs in the pool.
- Status: Color-coded icon that shows the current pool status.

Clicking on the gear icon on the right side of a row provides options for modifying the resource pool.

On this page you can also specify available target VLANs for port configuration policies. See [Setting Up Port Configuration Policies](#).

4.1.2. Resource Pool Level

For a selected resource pool, the Status page presents the following information:

- Status: oded icon that shows the current pool status.
- Pool Master: IP address or DNS name of the master server in the pool.
- Pool-Wide Networks: Number of networks in the pool.
- XenServers: Number of servers in the pool.
- All VMs: Number of VMs in the pool.
- Server list: List of servers in the pool, including server name, number of networks, number of VMs, and status.

In addition to displaying status information, you can configure how Netflow data is forwarded by all XenServers in the pool. Select the following check boxes as appropriate, and click **Save Netflow Configuration**:

- vSwitch Controller (selected by default): Forwards Netflow information to the vSwitch Controller for use by the Flow Statistics section of the GUI. If you deselect this check box, the Netflow data is not sent to the vSwitch Controller and the Flow Statistics pages do not show data.
- External Netflow Controller: Allows you to forward Netflow data to an external third party Netflow collector. Enter the IP address of the external collector.

4.1.2.1. Fail safe mode

The Fail Mode section allows you to configure how a vSwitch in the resource pool enforces access control (ACL) rules when it is unable to connect with its configured vSwitch Controller. It is important to maintain a high level of vSwitch Controller availability to avoid data loss. During times of unavailability, the following fail modes apply:

- Fail-open: all traffic is allowed, previously defined ACLs no longer apply until the vSwitch is able to reconnect with the vSwitch Controller.
- Fail-safe: existing ACLs continue to apply.

Under normal operation, the vSwitch maintains connections to its configured vSwitch Controller to exchange network management and status information. If the vSwitch Controller becomes unavailable, for example due to network disruption or Controller restart, the vSwitch waits up to an inactivity timeout during which network traffic is dropped. After the inactivity timeout, the vSwitch enters into the configured fail mode.

In fail-safe mode, existing ACLs continue to apply after the vSwitch loses connectivity to its configured vSwitch Controller. Traffic that does not match existing ACLs are denied. Note that all ACLs (at any level of the policy hierarchy presented by the Controller) are enforced as sets of rules on VIFs in the vSwitch. As a result, new VIFs, or existing VIFs that are unplugged then re-plugged, that appear in fail-safe mode while the Controller is unavailable will not be able to communicate until the Controller becomes available again, even if higher-level ACL policy rules (Global, per-resource pool, per-network or per- VM) that allow communication are present on existing VIFs. Furthermore, the vSwitch Controller may define ACLs based on IP addresses it has learned. In fail-safe mode, packets sent by a VM using an IP address the Controller has not associated with the VM before it became unavailable are denied. For example, an existing VM that uses a new IP address will not be able to communicate until the Controller is reachable again. Other examples where traffic is denied while in fail-safe mode include:

- Newly plugged VIFs
- A new VM
- A migrated VM (e.g. XenMotion or Workload Balancing)
- VMs on hosts added to a pool
- Applications that act like a router

One additional behavior to note is, if the vSwitch is restarted in fail-safe mode and the controller is still unavailable after the vSwitch has started, all ACLs are lost which means all traffic is denied. The vSwitch stays in fail-safe mode until connectivity with the Controller is re-established and ACLs are pushed down to the vSwitch by the Controller.

Warning:

Removing a resource pool from vSwitch Controller management while in fail-safe mode may result in the vSwitch losing network connectivity and forcing an emergency reset situation. To prevent this, a resource pool should only be removed while its status is green.

You can also specify available target VLANs for port configuration policies on this page. See [Setting Up Port Configuration Policies](#).

4.1.3. Server Level

For a selected server, the Status page presents the following information:

- Server Status: Color-coded icon that shows the current server status.
- Server Networks: Number of networks in the resource pool.
- MAC Address: MAC address of the server management interface.
- IP Address: IP address of the server management interface.
- vSwitch Version: Build and version number of the vSwitch running on this XenServer.



- **Server Networks:** List of all networks associated with the server, including the number of VMs on the server using that network, associated physical interface, VLAN, number of bytes transmitted and received, number of errors, and status.
- **Server VMs:** List of all VMs associated with the server, and for each VIF on the VM, list of the MAC address, network, IP address, total bytes transmitted and received since the VM was booted, and status.

On this page you can also specify available target VLANs for port configuration policies. See [Setting Up Port Configuration Policies](#).

4.1.4. Network Level

The Status tab for pool-wide networks lists summary information about each network in the resource pool. The Status tab for an individual network lists information about the network itself and includes hyperlinked tables of information about the physical interfaces and VM interfaces currently connected to the network.

The status icon is green if the network is active and properly managed by the vSwitch Controller, red if it has no connected interfaces, and orange if there is an error condition described by the associated text.

For pool-wide networks, the following information is displayed:

- **Network name:** Specific network.
- **VMs:** Number of VMs associated with the network.
- **XenServer:** Server for the network.
- **Physical Interface:** Server interface for the network.
- **Transmit (Tx) and receive (Rx) packets:** Aggregated counters across all VIFs on the specified network.
- **Errors:** Aggregated counters across all VIFs on the specified network.
- **Status:** Color-coded icon that shows the current network.

For a selected network, the following information is presented:

- **Network Status:** Color-coded icon that shows the current network.
- **VMs:** Number of VMs associated with the network.
- **Physical interfaces:** List of physical interfaces, including VLAN, number of bytes transmitted and received, errors, and status.
- **Switching XenServer (present on cross-server private networks only):** Specifies the current active switching host for the network.

A Cross-server private network enables communication between VMs in the same resource pool, without need for any additional configuration of the physical network and regardless of whether the VMs are running on the same host. This is accomplished by having a "switching host" establish GRE tunnels (in a star topology) to each of the other hosts (which have an active VM running on the private network) in the pool.

If a switching host becomes unavailable or is deleted, a new switching host is automatically selected and new GRE tunnels are configured. See the *XenServer Administrator's Guide* for more information on cross-server private networks.

- **VM interfaces:** List of VMs, including MAC address, IP address, number of bytes transmitted and received, and status.

On this page you can also specify available target VLANs for port configuration policies. See [Setting Up Port Configuration Policies](#).

4.1.5. Virtual Machine (VM) Level

The following information is displayed for all VMs:



- VM name: Name of the specific VM.
- MAC address: MAC address assigned to the VM.
- Network name: Network to which the VM is assigned.
- Detected IP address: IP address(es) assigned to the VM.
- Transmit (Tx) and receive (Rx) packets: Aggregated counters across all VIFs on the specified VM.
- Errors: Aggregated counters across all VIFs on the specified VM.

For a selected VM, the Status page displays the following information:

- Status: Color-coded icon that displays the current VM status.
- Resource Pool: Resource pool to which the VM belongs.
- Server Name: Name of the server to which the VM is currently assigned. This is blank if the VM is not running and is not tied to a specific server.
- VM Group Membership: List of administrative groups to which the VM is assigned.
- VM interfaces: List of the VIFs on the VM, including MAC address, network name, detected IP address, transmit and receive byte, packet, and error counts, and status.
- Network Events: List of network events involving the VM, including priority, date/time, and description.

4.1.6. Virtual Interface (VIF) Level

For a selected VIF, the Status page presents the following information:

- Status: Color-coded icon that shows the current VIF status.
- Resource Pool: Resource pool to which the VIF belongs.
- Network: Network to which the VIF belongs.
- VM Name: VM to which the VIF belongs.
- MAC Address: MAC address of the VIF.
- IP Address: IP address of the VIF.
- Transmit and Receive bytes, packets, and errors: Traffic counts for the VIF.
- Switch Port ACL Statistics: Unlike transmit and receive counts, the ACL hit counts are instantaneous statistics read from the ACL rule statistics of the current vSwitch. Therefore, policy changes and VM actions, such as suspension, shut down, or migration will cause these statistics to reset.

The vSwitch ACL statistics require an IP address to be identified on the network and able to collect statistics for IP-based protocols. If you find that there are no counts on IP-based rules, verify that an IP address is displayed in the IP address field.

4.1.7. Viewing Flow Statistics

By default, the vSwitch on each managed XenServer sends Netflow data to the vSwitch Controller, which uses this data to generate Flow Statistics tables and charts. Netflow records are generated for all IPv4 flows after five seconds of inactivity or 60 seconds of total activity.

The data rate of a flow is represented as the total traffic of the flow averaged across the duration of the flow. For example, if a flow lasts 10 seconds with 900KB sent in the first second and 10KB sent in each of the nine remaining seconds, the resulting data is plotted as if the rate were 100KB/second for the entire flow period.

Due to Netflow's use of UDP datagrams to transport NetFlow records between a switch and a collector (e.g., the vSwitch Controller), there is usually no way for the collector to know why a NetFlow record was not received, and dropped records may result in nondeterministic data with Flow Statistics tables or charts. For example, assume


that a network generating 10 flows per second has a single 1GB file transfer that lasts 10 seconds. A total of 202 flows are generated (100 hping stimuli, 100 hping responses, 1 file transfer stimulus, and 1 file transfer response). If 50 percent of the UDP datagrams carrying NetFlow records are dropped, there is a 50/50 probability that the collector will report either 1GB of data, or 2KB.

Because Netflow records are generated by each vSwitch in a resource pool, sources and destinations that are running on different XenServers result in two records, doubling the statistics counts.

We recommend disabling flow visibility in deployments of more than 100 VMs to avoid overloading the vSwitch Controller virtual appliance and the network used to send NetFlow records.

The **Flow Statistics** tab displays a graph and associated table to show flows for the selected node.

Use the drop-down lists at the top of the page to specify the following:



- Direction: Bidirectional, Inward, Outbound
- Units: Bytes, Bits, Packets, Flows
- The top or bottom items (highest or lowest values) of one of the following groupings:
 - VMs: VMs residing within the resource pool as sources/destinations for traffic
 - IP Addresses: IP addresses as source or destination for traffic
 - Protocols: IP protocol traffic such as ICMP, TCP, and UDP

Note:

Ethernet layer protocols (such as ARP) are not displayed due to the limitations in the Netflow protocol used to generate results.

- Application: “application”-level protocol traffic, identified by TCP/UDP port or ICMP type/code
- Traffic (by type): VMs, IP Address, Protocols, Applications (shown by protocol type and port number, this can allow you to infer the service)
- Time interval.

The table below the graph displays some or all of the following information, depending upon the type of item selected in the drop-down list:

- VM
- IP
- Inbound bytes
- Inbound data rate (Kbit/s)
- Outbound bytes
- Outbound data rate (Kbit/s)
- Total bytes
- Total data rate (bps)

If NetFlow is not being forwarded to the vSwitch Controller, a warning blue status text will be displayed under the **Flow Statistics** tab: "one or more selected pools is not configured to forward NetFlow records to vSwitch Controller".

To re-configure forwarding, click the blue status text to see a list of resource pools. Select the resource pool desired from the list to navigate to the pool status page. From the status page, you can configure NetFlow data forwarding.



4.2. Managing Address Groups

You can set up address groups to specify the IP addresses to use as the source or destination for ACLs and for reporting of flow statistics.

To add an address group:

1. Under **Visibility & Control**, select **Address Groups** in the resource tree (side panel) to open the Status page for all address groups.
2. Click **Create Group**.
3. Enter the name to identify the group, and an optional description.
4. Click **Create Group**. The new group is added to the list of address groups.
5. Select the new group in the resource tree to open its **Status** page.
6. Click the **Add Members** button.
7. In the pop-up window, specify one or more IP addresses or subnets (comma separated). Example: 192.168.12.5, 192.168.1.0/24
8. Click **Add**. Continue to add additional networks as needed. Each set of addresses is added as a node under the network in the Address Groups list.

The new address group is now available for ACL policies and flow statistics.

You can remove an existing address group by clicking the **Remove** link in the row of the **All Address Groups** for that address group.

You can also update the name or description of an address group:

1. Select the new group in the resource tree to open its **Status** page.
2. Click the **Modify Group** button.
3. In the dialog that opens, change the name and description.
4. Click the **Modify Group** button to save the changes.

4.3. Managing Virtual Machine Groups

A VM group is a set of VMs that you identify as a group for viewing status and flow statistics. Each VM in a VM group must already be in a resource pool. The groups are otherwise independent of resource pools and servers.

To add a VM group:

1. Under **Visibility & Control**, select **VM Groups** in the resource tree (side panel) to open the Status page for all VM groups.
2. Click the **Create Group** button.
3. Enter the name to identify the group, and an optional description.
4. Click **Create Group**. The new group is added to the list of VM groups.
5. Select the new group in the resource tree to open its **Status** page.
6. Click **Add Member**.
7. In the pop-up window, select the VM from the drop-down list.
8. Click **Add**. Continue to add additional VMs as needed. Each VM is added as a sub-node under the group in the VM Groups list.

The following right-click options are available for each VM group:

- Add VM to group: Add a new group member.

- **Modify Name/Description:** Change the name or description.
- **Remove Group:** Delete the group.

4.4. DVS Policy Configuration Hierarchy

The Access Control and Port Configuration tabs within Visibility & Control provide a way to configure access control, QoS, and traffic mirroring policies within the virtual network environment. While all policies are applied at the VIF level, vSwitch Controller exposes a hierarchical policy model that supports declaring default policies across a collection of VIFs (e.g., a resource pool) while also providing a way to override this default policy by creating fine-grained exceptions when needed (e.g., exempting a particular VM from the default resource pool policy).

Similar to the hierarchy used in the resource tree, the policy hierarchy has the following levels:

- **Global (most general level):** Includes all VIFs in all resource pools.
- **Resource pools:** All VIFs in a particular resource pool.
- **Networks:** All VIFs attached to a particular network.
- **VMs:** All VIFs attached to a particular VM
- **VIFs (most specific level):** A single VIF.

Note:

XenServers are not included in the policy hierarchy, since policies must apply regardless of what XenServer in a resource pool is currently running a VM.

4.5. Setting Up Access Control Policies

Choose the **Access Control** tab to set up policies that allow or deny VM traffic based on packet attributes.

An ACL policy consists of a set of rules, each of which includes the following:


- **Action:** Indication of whether traffic matching the rule should be permitted (Allow) or dropped (Deny).
- **Protocol:** Network protocol to which the rule applies. You can apply the rule to all protocols (Any), choose from an existing protocol list, or specify a new protocol.
- **Direction:** Direction of traffic to which the rule applies. The text of the rules is meant to be read from left to right, so “to” means traffic outbound from the VM, while “from” means traffic inbound to the VM.
- **Remote Addresses:** Indicates whether the rule is limited to traffic to/from a particular set of remote IP addresses.

Management of ACL policies closely follows the resource tree hierarchy. You can specify policies at any supported level of the hierarchy. At each level, rules are organized as follows:

- **Mandatory rules:** These are evaluated before any child policy rules. The only rules that take precedence over them are mandatory rules of parent (less specific) policies. Mandatory rules are used to specify rules that cannot be overridden by child (more specific) policies.
- **Child rules:** The child policy placeholder indicates the location in the rule order at which rules in child policies will be evaluated. It divides the mandatory rules from the default rules.
- **Default rules:** These are evaluated last, after all mandatory rules and all child policy default rules. They only take precedence over default rules of parent policies. They are used to specify behavior that should only be applied if a more specific child policy does not specify conflicting behavior.

The next figure shows the **Access Control** tab for a VIF.

Figure 4.1.



4.5.1. Global Access Control List (ACL) Rules

To set up global ACL rules, click **All Resource Pools** in the resource tree. The page lists all of the ACL rules that are defined at the global level.

4.5.2. Resource Pool Access Control List (ACL) Rules

To set up ACL rules for a resource pool, select the resource pool in the resource tree.

The page shows an expandable bar for global policy, and an expanded area for resource pool rules. If you click the **Expand All** button, you can see how the resource pool rules are embedded in the global policy framework.

4.5.3. Network Access Control List (ACL) Rules

To set up ACL rules at the network level, click the network in the resource tree.

The page shows an expandable bar for global rules, an expandable bar for the resource pool to which the network belongs, and an expanded area for network rules. If you click the **Expand All** button, you can see how the network policies are embedded in the resource policy framework, and, in turn, in the global policy framework.

4.5.4. VM Access Control List (ACL) Rules

To set up policies at the VM level, click the VM in the resource tree.

The page shows an expandable bar for global rules, expandable bars for the resource pool and network to which the VM belongs, and an expanded area for VM rules. If you click the **Expand All** button, you can see how the VM rules are embedded in the network, resource pool, and global framework.

If a VM contains VIFs on multiple networks, a “Change Network” link will appear on the right side of the example bar for the network, allowing you to view the rules for each network level policy that might apply to a VIF on that VM.

4.5.5. VIF Access Control List (ACL) Rules

To set up policies at the VIF level, click the VIF in the resource tree. Because policies are packaged and applied only at the VIF level, you must display the VIF pages to see the full policy context.



The page shows expandable bars for global rules, expandable bars for the resource pool, network, and VM to which the VIF belongs, and an expanded area for VIF rules. If you click the **Expand All** button, you can see how the VIF rules are embedded in the VM, network, resource pool, and global framework.

4.5.6. Access Control List (ACL) Rule Enforcement Order

While ACLs can be defined at different levels of the policy configuration hierarchy, ACLs are enforced on a per-VIF basis. For actual enforcement, the hierarchy is combined in the order described in this section and applied to each VIF. To see the currently-applied rules on a VIF along with the associated statistics, select the VIF in the resource tree and view the ACL list in the Status tab.

The enforcement order is as follows:

1. Mandatory rules at the global level
2. Mandatory rules for the resource pool containing the VIF
3. Mandatory rules for the network containing the VIF
4. Mandatory rules for the VM containing the VIF
5. Rules for the VIF containing the VIF
6. Default rules for the VM containing the VIF
7. Default rules for the network containing the VIF
8. Default rules for the resource pool containing the VIF
9. Default rules for the global containing the VIF

The first rule that matches is executed, and no further rules are evaluated.

Note:

When a vSwitch Controller is unavailable, the resource pool will enforce access control rules based on the configured fail mode. See the section called "Resource Pool Level" under "Viewing Status" for more details about a resource pool's fail mode.

4.5.7. Defining Access Control List (ACL) Rules

To define a new ACL rule, use the resource tree to choose the node at the appropriate level in the policy configuration hierarchy. At each level, you can add rules for that level and higher levels. For example, if you select a resource pool, you can add rules for that resource pool and global rules.

If you choose a resource tree node that does not correspond to a level in the policy configuration hierarchy (such as a XenServer), a message is displayed with links to choose another levels.

New rules can be added in the following ways:

- To add a new mandatory rule, click the gear icon in the header bar for the level, and choose **Add New Mandatory ACL**.
- To add a new default rule, click the gear icon in the header bar for the level, and choose **Add New Default ACL**.
- To add a new rule above an existing rule entry, click the gear icon for the entry, and choose **Add New ACL Above**.
- To add a new rule below an existing rule entry, click the gear icon for the entry, and choose **Add New ACL Below**.

The new rule is added to the page with the following default settings:

- Action: Allow
- Protocol: Any
- Direction: To/From
- Remote Addresses: Any

- Description: None

To change a particular field within a rule, click the link representing the current field value and apply changes as described in the following table. When you apply a change, the rule is updated to show the values.

Item	Description
Action	Click the link and choose Change Action to Deny or Change Action to Allow .
Protocol	<p>Click and choose one of these options:</p> <ul style="list-style-type: none"> • Choose Match Any Protocol to apply the rule to all protocols. • Choose Use an Existing Protocol to specify a protocol. Select the protocol from the drop-down list, and click Use Protocol. • Choose Use a New Protocol to specify custom protocol characteristics. Specify the following information in the pop-up window, and click Save & Use: <ul style="list-style-type: none"> • Ethertype: Select IP or enter another Ether type. • IP Protocol: Select one of the listed protocols, or enter another. • Destination Port (TCP/UDP only): Enter a port number or specify Any. • Source Port (TCP/UDP only): Enter a port number or specify Any. When defining an application that uses a well-known server port (e.g., HTTP uses port 80), it is best to define that well-known port as the destination port and leave the source port as Any. • ICMP Type (ICMP only): Choose Any or enter a specific ICMP type Protocol (ICMP) type. • ICMP Code (ICMP only): Choose Any or enter a specific ICMP code. • Match reply traffic: Indicate whether return traffic will automatically be allowed as part of the rule. For example, if the rule is to allow UDP destination port 7777 traffic from the VM to a specified remote address, and Match reply traffic is selected, then UDP traffic is also allowed from source port 7777 of the remote address to the VM. This option should be enabled for any UDP protocol that requires bidirectional communication (the option is always enabled for TCP). • One-time Use vs. Multiple Uses: Select whether you want to use this protocol only for the current rule or add it to the list of protocols that can be selected in the drop-down protocol menu. • Choose View/Modify Current Protocol to modify characteristics for an already defined protocol.
Direction	Choose whether the rule will apply from or to the specified remote addresses, or both.
Remote Addresses	<p>To specify the remote addresses:</p> <ol style="list-style-type: none"> 1. Click the Any link to open a pop-up window that lists the available address groups. 2. Select one or more address groups and use the arrows to move them to the Selected column. 3. Use the All buttons to select or deselect all of the groups. 4. To specify an IP address or subnet that is not part of an existing address group, enter the address or subnet (x.x.x.x or x.x.x.x/n), and click Add. Repeat to add additional addresses. 5. Click Done.

Item	Description
Description	To add a text description of the rule: <ol style="list-style-type: none"> 1. Click the Description button. 2. Click the entry (<None> if there is no current description). A text entry area is displayed. Enter the text and press Enter.
Rule Details	Click the Rule Details button to display a brief summary of the rule.

You must click **Save Policy Changes** to apply the new rules. When you do so, the changes take effect immediately within the virtual network environment. If you have not already saved the rules, you can click **Undo Changes** to reverse the changes you have named.

When you change an ACL, all background updates for the vSwitch Controller GUI are paused. If another administrator is modifying the policy simultaneously and commits changes before you, you must refresh the page to retrieve the new policy from the server and then reenter the changes.

You can change order of rules in a level by clicking the gear icon for the rule and choosing **Move Up** or **Move Down**. You cannot move a rule between levels in the hierarchy. To remove a rule, click the gear icon and choose **Delete**. Click the **Description** button to display the ACL description. or the **Rule** button to display the ACL rule that you constructed.

ACL rules should always be interpreted from the point of view of the virtual interface of the VM, even if configured at higher levels of the policy hierarchy. This is particularly important when thinking about the meaning of the **Remote Addresses** field in the rules.

For example, if a VM within a resource pool has the IP address 10.1.1.1, it might be expected that a rule on that resource pool specifying "deny all protocols to IP 10.1.1.1" would prevent any traffic from reaching the VM. This will be the case for all other VMs in the resource pool because each VM will enforce the rule when the VM transmits. However, machines that are external to the resource pool *will* be able to communicate with the VM with IP address 10.1.1.1. This is because no rules control the transmit behavior of the external machines. It is also because the VIF of the VM with IP address 10.1.1.1 has a rule that drops transmit traffic with that address but not *receive traffic* with that address.

If the policy behavior is unexpected, it can be helpful to view the **Status** tab for the virtual interface on which the entire configured set of rules from all policy levels is visualized.

4.6. Setting Up Port Configuration Policies

Use the **Port Configuration** tab to configure policies that apply to the VIF ports. The following policy types are supported:

- QoS: Quality of service (QoS) policies control the maximum transmit rate for a VM connected to a DVS port.
- Traffic Mirroring: Remote Switched Port Analyzer (RSPAN) policies support mirroring traffic sent or received on a VIF to a VLAN in order to support traffic monitoring applications.
- Disable MAC address spoof check: MAC address spoof check policies control whether MAC address enforcement is performed on traffic outbound from a VIF. If the vSwitch Controller detects a packet with an unknown MAC address from a VIF, it drops the packet and all subsequent traffic from the VIF. MAC address spoof check policies are on by default and should be disabled on VIFs running software like Network Load Balancing on Microsoft Windows servers.


Warning:

Enabling RSPAN without correct configuration of your physical and virtual network can cause a serious network outage. Read the instructions in [Configuring RSPAN](#) carefully before enabling this feature.

You can configure QoS and Traffic Mirroring port policies at the global, resource pool, network, VM, and VIF levels. When you select a node in the resource tree and choose the **Port Configuration** tab, the configured value for each parent level in the hierarchy is shown, but only the configuration at the selected policy level can be changed. For example, if you select a VM, the **Port Configuration** tab shows the values configured at the global, resource pool, and network levels, and lets you change the value at the VM level.

QoS and Traffic Mirroring configurations at a given level override the configurations at the higher levels. If a configuration is overridden, then the **Port Configuration** tab shows the higher level configuration crossed out. For example, the next figure shows a QoS configuration at the network level that overrides the configuration at the resource pool level.

Figure 4.2.



To configure port policies, choose the node in the resource tree and choose the **Port Configuration** tab. If you choose a node that does not support port configuration policies (such as a XenServer), a message is displayed with links to nodes that do support port configuration.

4.6. Configuring QoS. Configuring QoS

For QoS policies, choose from the following options:

- Inherit QoS policy from parent (default): Applies the policy from the higher (i.e., less specific) hierarchy level. This option does not exist at the global level.
- Disable inherited QoS policy: Ignores any policies that are set at higher (i.e., less specific) levels such that all VIFs included in this policy level have no QoS configuration.
- Apply a QoS limit: Select a rate limit (with units), and a burst size (with units). Traffic to all VIFs included in this policy level is limited to the specified rate, with individual bursts limited to the specified number of packets.

Warning:

Setting the burst size to be too small relative to the rate limit can prevent a VIF from even being able to send enough traffic to reach the rate limit, especially with protocols that perform congestion control such as TCP.

At minimum, the burst rate must be larger than the Maximum Transmission Unit (MTU) of the local network.



Setting QoS to an inappropriately low burst rate (for example, 1 KB) on any interface which the vSwitch Controller sits may result in losing all communication with the vSwitch Controller and forcing an emergency reset situation.

To prevent any inherited enforcement from taking place, the QoS policy at the VM level should be disabled

Click **Save Port Configuration Changes** to implement the changes, or click **Undo Changes** to remove any unsaved changes. The policy takes effect immediately after saving.

4.6.2. Configuring RSPAN

Warning:

Configuring RSPAN when the server is connected to a switch that does not understand VLANs or is not properly configured to support the RSPAN VLAN can lead to traffic duplication and network outages. Review the documentation and configuration of your physical switches before enabling the RSPAN feature, especially at higher levels of the hierarchy where multiple physical switches may be involved.

Enabling RSPAN requires a series of steps, outlined below:

4.6.2.1. Identify your RSPAN VLAN

When RSPAN is enabled on a VIF, the vSwitch for that VIF will make a copy of each packet sent to or from that VIF and transmit the copy of that packet tagged with VLAN value called the target VLAN. An administrator would then place a host performing monitoring on the switch port that is configured to use the target VLAN. If the monitoring host interface uses promiscuous mode, it can see all traffic sent to and from the VIFs configured to use RSPAN.

4.6.2.2. Configure the Physical Network with the Target VLAN

It is critical to correctly configure the physical network to be aware of the RSPAN traffic to avoid network outages. RSPAN should only be enabled if the physical switching infrastructure connecting all RSPAN-enabled VIFs can be configured to disable learning on the target VLAN (see the documentation from your switch manufacturer).

Additionally, traffic sent on the target VLAN must be forwarded from each of the vSwitches to the monitoring hosts. If your physical infrastructure includes many switches in a hierarchy, this requires trunking the target VLAN between the different switches (see the documentation from your switch manufacturer).

4.6.2.3. Configure vSwitch Controller with the Target VLAN

You must tell the vSwitch Controller about each target VLAN before using that VLAN ID for RSPAN port configuration. You can specify available target VLAN IDs at the resource pool, network, or server level. Target VLANs that are added at a level of the hierarchy are available when configuring RSPAN port configuration at that level and all lower levels of the hierarchy. The correct level to specify a target VLAN depends on how widely you have configured your physical infrastructure to be aware of that target VLAN.

To specify available target VLANs:

1. Under **Visibility & Control**, open the **Status** tab for all resource pools, a specific resource pool, a specific server, or a specific network.
2. In the **RSPAN Target VLAN IDs** area, click + and enter the VLAN ID.
3. Repeat to add additional VLAN IDs.
4. Click **Save Target VLAN Change**.

The VLANs are now available for selection on the Port Configuration tab, as described in this section.



4.6.2.4. Modify port configuration to enable RSPAN for a set of VIFs

To configure RSPAN policies within the **Port Configuration** tab, select the appropriate node in the resource tree and choose from the following options:

- Inherit RSPAN policy from parent (default): Applies the policy from the next higher (i.e., less specific) hierarchy level.
- Disable inherited RSPAN policy: Ignores any policies that are set at higher (i.e., less specific) levels such that all VIFs included in this policy level have no RSPAN configuration.
- RSPAN traffic on VLAN: Choose a VLAN from the list of target VLANs. The only target VLANs that will appear in the list are those configured for policy levels containing the currently selected node.

4.6.2.5. Configuring MAC Address Spoof Checking

To disable MAC address enforcement, select the MAC address spoof checking check box. Enforcement can only be configured on a per VIF basis and does not inherit or override parent configurations.

4.6.2.6. Save Changes

Click Save Port Configuration Changes to implement the changes, or click **Undo Changes** to remove any unsaved changes. The policy takes effect immediately after saving.

Chapter 5. vSwitch Controller Administration & Maintenance

Use the **Settings** pages to perform administration and maintenance functions on the vSwitch Controller. To access the Settings pages, click the **Settings** icon in the top panel of the vSwitch Controller window.

5.1. Configuring IP Address Settings

Use the **IP Configuration** page to verify and configure the IP address of the vSwitch Controller. When the vSwitch Controller is started for the first time, it obtains an IP address through DHCP; however, we recommend that you assign a static IP address. If DHCP is configured, resource pools cannot be set to Fail-Safe mode.

To view and configure the controller IP address:

1. Under **Settings**, choose **IP Configuration** to display the current configuration.
2. To modify the configuration, click **Modify Configuration**.
3. Select **Manual Configuration** to assign a static IP address.
4. Enter the new IP address, netmask, gateway IP address, and, optionally, one or two DNS server IP address(es).

Note:

At least one DNS server IP address must be specified to enable name resolution on the Controller.

5. Click **Make Changes** to implement the changes.

Warning:

If, after changing the IP address of the vSwitch Controller, you see an error message (displaying **Pool Managed By [old IP Address]**) in the **Status** column of the resource pool(s) that the vSwitch Controller manages, you will need to instruct the Controller to begin managing the pool(s) again.

In the **All Resource Pools** tab, click the gear icon next to the **Status** column of the resource pool(s). Select **Steal Pool**.

How to Upgrade the vSwitch SSL Certificate

By default, the vSwitch Controller virtual appliance uses a self-signed SSL certificate for connections with the vSwitch running on each XenServer. You can get a certificate authority to provide you with a signed certificate for your vSwitch connections. Follow the instructions of the certificate authority you plan to use for generating the public/private key pair to be signed and submit it to the authority. After you obtain the signed certificate from the authority, follow the steps in this section.

1. Under **Settings**, click **Server and Certificate Maintenance**.
2. Click **Update OVS Certificate**.
3. Browse to select the SSL/TLS certificate file.
4. After uploading the file, click **Update Certificate**.

To view information about the vSwitch SSL security certificate or determine when it expires:

1. Under **Settings**, click **Server and Certificate Maintenance**.
2. Click **View OVS Certificate**.

After updating the vSwitch SSL certificate, as new resource pools are added for management the vSwitch of each XenServer in the new resource pool automatically downloads and starts using the updated SSL certificate.



However, the SSL certificate on vSwitches running on existing pools under management need to have their SSL certificates updated manually.

To update the vSwitch SSL Certificate on a XenServer

1. On the XenServer host, copy the SSL certificate to `/etc/openvswitch/vswitchd.cacert`
2. Restart the XenServer host.

5.2. Configuring the Controller Hostname

To verify and configure the Controller hostname and DNS domain, use the **IP Configuration** page. By default, the controller hostname is "dvsc", and the DNS domain name is unassigned.

To change the hostname or domain:

1. Under **Settings**, chose **IP Configuration** to display the current configuration.

Click **Modify Host Settings**.

2. Enter the desired hostname and domain name into the appropriate fields.

The value of the domain name is used for both the domain name of the host and the domain to search for unqualified host names.

3. Click **Make Changes** to save changes, or choose **Cancel**.

5.3. Collecting Information for Trouble Reports

To collect information to supply for trouble reports, click **Server and Certificate Maintenance** under **Settings**, and then click **Collect & Zip All Logs** to add all relevant vSwitch Controller logs to a zip file for download. When the zip operation is complete, click the **here** link in the pop-up window to download the dump.tar.gz file. After downloading, click **Close** to close the pop-up window.

5.4. Restarting the vSwitch Controller Software

To restart the vSwitch Controller software, click **Server and Certificate Maintenance** under **Settings**, and then click **Restart Network Controller**. When the restart is complete, the login page opens.

5.5. Managing Administrative Accounts

Multiple user accounts can be used to provide certain users with limited privileges when accessing the GUI. Additionally, since entries in the Administrative Events log contain the name of the user who performed the action, having multiple users can help determine who made a recent configuration change.

To add user accounts for access to the vSwitch Controller and to change user passwords:

1. Under **Settings**, choose **Administrative Accounts**.
2. Click **Create Account**.
3. Enter a user name and password, and reenter the password to confirm. Specify any of the following user privilege levels:
 - Superuser: All privileges.
 - Read-write: All privileges, except for the ability to modify other user accounts and restore snapshots.
 - Read-Only: Can see most information in the GUI but cannot modify anything in the vSwitch Controller except the user's own password.
4. Click **Add User**.



To change a user password, click the **Password** link for the user. Enter and confirm a new password, and click **Change Password**.

To remove a user, click the **Remove** link for the user. You cannot remove the admin user.

5.6. Managing Configuration Snapshots

Snapshots provide a mechanism to save the current vSwitch Controller configuration so that you may restore to that exact configuration at a later point. For example, it might be useful to snapshot the system prior to making major configuration changes. By default, the system automatically creates an automatic snapshot every 12 hours.

Click **Configuration Snapshots** under **Settings** to view the list of configuration backups and restore from backup. The page lists all recent backups, with the most recent listed first. Automatic backups are taken twice per day and each time the vSwitch Controller is restarted. When restoring from a backup, the current IP configuration of the vSwitch Controller is not updated. To change the vSwitch Controller IP address, see [Section 5.1, “Configuring IP Address Settings”](#).

To restore the configuration from a backup, click the gear icon for the snapshot and choose **Restore to Snapshot**. When asked if you want to continue, click **Yes, Restore**.

To create a backup on demand, click **Create New Snapshot**. You can enter an optional description to identify the snapshot. Click **Create Snapshot**. The new backup is added to the top of the list.

To download a snapshot to store on another system, click the gear icon for the snapshot and choose **Download**. Follow the instructions in the popup windows to save the snapshot file.

To upload a previously-saved snapshot to the controller, click **Upload Snapshot**. Browse to select the snapshot file, and click **Upload Snapshot**. The uploaded snapshot is added to the list on the Configuration Snapshots page.

To delete a snapshot, click the gear icon for the snapshot and choose **Delete Snapshot**. When asked if you want to continue, click **Delete Snapshot**.

The snapshot table also includes information on the software version and compatibility. Compatibility indicates whether the data in the snapshot is compatible with the current software version. It displays a green indicator if it is compatible and a red indicator if it is not. To revert to an incompatible snapshot, you must first change the software to a compatible version, as listed in the Software Version column.

By default, the system creates a configuration snapshot every 12 hours. These snapshots are listed with a description label of “Automatic periodic snapshot”. In addition, configuration snapshots are created each time the vSwitch Controller is restarted. These snapshots are listed with a description label of “Startup snapshot”. System initiated snapshots are automatically deleted if more than 30 days old. When manually creating a new snapshot, enter a unique description label so it is not mistaken as a system initiated snapshot and deleted after 30 days. If a system initiated snapshot needs to be preserved beyond 30 days, it can be downloaded and re-uploaded using a unique description label.

5.7. Adding Network Time Protocol (NTP) Servers

The time setting on the vSwitch Controller virtual appliance is managed by a connection to external Network Time Protocol (NTP) servers. The controller comes with default servers already configured. Because these may not be optimal for your environment, it is best to replace them with a local NTP server according to the following instructions.

To add an NTP server:

1. Under **Settings**, choose **Time & NTP**.
2. Click **Add Server**.
3. Enter the IP address of the server, and click **Add**.
4. Add additional servers as needed.



To remove an NTP server, click the **Remove** link.

5.8. Exporting Syslog Files

Use the Syslog page to add servers to receive remote syslog messages, which consist of administrative and network event messages generated by the system. The most recent syslog entries are also displayed on the dashboard.

To add syslog servers:

1. Under **Settings**, choose **Syslog**.
2. Click **Add Server Address**.
3. Enter the IP address of the server, and click **Add**.
4. Add additional servers as needed.

To remove a server, click the **Remove** link.

Chapter 6. Troubleshooting vSwitch Controller Issues

This chapter contains information to help in troubleshooting vSwitch Controller issues.

6.1. Resource Tree Node Status

The following table describes the status icons for each resource type. These appear in the resource tree and on the Status page for the item.

Items/Status Icons	Description
VIFs	
Red	Associated virtual machine (VM) is shut down or unreachable.
Green	Virtual interface (VIF) is currently up and being managed.
Orange	VM is running but the XenServer on which the VIF resides is not connected to the vSwitch Controller.
VMs	
Red	VM is shut down or unreachable.
Green	VM is in running state and VIF's are being managed.
Orange	VM is running but the XenServer on which the VM resides is not correctly connected to the vSwitch Controller (depends on the collective state of the respective VIFs).
Server Networks	
Red	XenServer is shut down or unreachable or no VMs have VIFs that are associated with the network.
Green	XenServer is correctly connected to the vSwitch Controller.
Orange	XenServer is not correctly configured to connect to the vSwitch Controller (depends on the collective state of the associated physical interfaces and VIFs).
XenServers	
Red	XenServer is shut down or unreachable.
Green	XenServer is correctly connected to the vSwitch Controller.
Orange	XenServer is not configured to connect to the vSwitch Controller (depends on the collective state of the associated physical interfaces and VIFs).
Pool-Wide Networks	
Red	Master XenServer is shut down or unreachable.
Green	Master XenServer is configured to connect to the vSwitch Controller and the connection is up and working.

Items/Status Icons	Description
Orange	Master XenServer is not configured to connect to the vSwitch Controller (depends on the collective state of the associated physical interfaces and VIFs).
Resource Pools	
Red	Master XenServer is shut down or unreachable.
Green	Master XenServer is configured to connect to the vSwitch Controller and the connection is up and working.
Orange	Master XenServer is not configured to connect to the vSwitch Controller (depends on the collective state of the associated physical interfaces and VIFs).

6.2. Troubleshooting Access Policy Issues

The following suggestions may help in troubleshooting when access control policies are not operating properly:

1. Select the **Status** page for the VIF of a VM that should be affected by the policy. View the hit counts for each rule while you generate traffic that is not being handled correctly by the policy. Identify the rule that the traffic is actually hitting instead of the rule you expected it to be hitting. If the policy for this VIF does not already have a default rule that will match all traffic, for the purposes of debugging, add a rule that will match all traffic as the lowest priority default rule at the global level (Note: This rule can have either an allow or deny action, depending on your desired network behavior while debugging. Remove this rule after debugging).
2. If the traffic is hitting a rule of lower priority than the one you expected, carefully check the rule matching criteria. Is the direction of the traffic correctly specified? Are the remote hosts properly identified? Is the protocol correctly defined? For example, could the protocol be specifying a UDP port instead of a TCP port or vice versa?
3. If the traffic is hitting a rule of higher priority than the one you expected, you must resolve the conflict between this rule and the rule you expected the traffic to hit. You can resolve conflicts by redefining rules to be more/less granular (e.g., scoping a rule to only apply to a particular set of remote IP addresses) or by changing the relative priorities of the two rules.
4. If the VM has multiple VIFs, verify that it is transmitting/receiving the traffic on the VIF to which the policy is applied. If appropriate, use RSPAN to mirror traffic from the VIF to a network analyzer to ensure the traffic that should match the rule that is present.

Note:

When a vSwitch Controller is unavailable, the resource pool will enforce access control rules based on the configured fail mode. See the section called “Resource Pool Level” under “Viewing Status” for more details about a resource pool's fail mode.

6.3. Creating a Trouble Report

To address issues efficiently, you will need to collect information from the XenServer and vSwitch Controller that are involved in the issue as soon as possible after the issue occurs and submit the information along with your trouble report.

- Include a Server Status report for each XenServer that is involved in the issue. Refer to the *XenServer Administrator's Guide* for instructions on generating Server Status reports.
- Include a log bundle from the vSwitch Controller by clicking **Collect and Zip All Logs** in the Server & Certificate Maintenance Settings page. Refer to [Collecting Information for Trouble Reports](#).

6.4. Controller Error Messages

The following table describes error messages.

Message	Description
Connecting to Pool	<p>Displayed when a new pool is added and vSwitch Controller has not yet successfully connected to the pool master.</p> <p>OR</p> <p>Displayed when the vSwitch Controller restarts and it has not yet successfully connected to the pool master. If a successful connection is not established in 30 seconds, this message will be replaced 'Pool Connection Failed'</p>
Network control channels disconnected	XenServer is not correctly connected to the vSwitch Controller.
Missing Pool Address	No DNS name or IP address is available for the pool.
Pool Connection Failed	There is a network problem between the controller and the pool master, a failure in DNS name resolution, an invalid DNS name or pool master IP address, or the pool master is down or misconfigured.
Unsupported Pool Version	The DNS name or IP address configured to the pool does not resolve to a compatible version of XenServer.
Duplicate Pool: Pool Disabled	The pool reports the same XAPI UUID as another pool already in the vSwitch Controller database.
Pool Authentication Failure	vSwitch Controller was unable to authenticate to the pool master using the username and password provided.
Pool Identity Changed	The pool has been reinstalled and does not match the state of the matching pool in vNetManager.
Pool Synchronization Error	An unsupported operation was seen while using XAPI to communicate with the pool master.
Unknown Error	Cause of the error is not known.

Chapter 7. Command Line Interface

This chapter describes the vSwitch Controller CLI commands. You can access the CLI locally from the text console of the Controller VM in XenCenter. To access the CLI remotely, use an SSH client application and connect to the controller VM hostname or IP address on port 22 (the default ssh port).

During a CLI session you can get help with CLI commands in either of the following ways:

- Type **help** and then press **Enter**.
- Enter part of a command followed by a space and question mark (?), and then press **Enter**.

7.1. CLI Commands

This section lists the available CLI commands. The interface supports completion of the command argument when you press the **Tab** key. Generally, you can abbreviate commands to the shortest, unique string at each level to reduce typing. You can access the command history within the current session is available by pressing the **Arrow** keys.

7.1.1. To terminate the current CLI session

Run the command: `exit`

7.1.2. To halt the vSwitch Controller

Run the command: `halt controller`

This command halts the vSwitch Controller appliance by gracefully shutting down the Controller.

7.1.3. To get information on commands

Run the command: `help`

7.1.4. To upgrade or downgrade the existing version of the Controller

Run the command: `install controller software-update <scp-format-remote-filename>`

This command secure copies (scp) a controller update file from the specified remote location and installs that version in place of the existing version.

This command can be used to install software versions that are both upgrades and downgrades. Upgrades will automatically migrate configuration to the new version. Downgrades will revert to the most recent compatible configuration snapshot or an empty configuration if no compatible snapshot exists.

7.1.5. To ping a specified remote system

Run the command: `ping <name-or-IP-address> [<count>]`

This command sends ICMP echo requests to the remote system identified by `<name-or-IP-address>` and waits for replies. If no count is specified, requests will be sent once per second until interrupted with Ctrl-C. If a count is specified, that number of pings will be sent.

7.1.6. To restart the Controller

Run the command: `restart controller appliance`

This command shuts down and restarts the entire controller appliance.



This command is primarily for troubleshooting and should not generally be required. Generally, the `halt` command should be used to power off the controller appliance.

7.1.7. To restart the Controller daemon

Run the command: `restart controller daemon`

This command shuts down and restarts the processes that implement the controller functions.

This command is primarily for troubleshooting and should not generally be required.

7.1.8. To set the hostname of the controller appliance

Run the command: `set controller hostname <hostname>`

This command sets the hostname of the controller appliance.

If the provided hostname contains one or more period (".") character(s), the hostname of the appliance will be set to the string *before* the first period; the domain name of the appliance will be set to the string *after* the first period.

7.1.9. To set the IP address of the Controller management interface via DHCP

Run the command: `set controller management-interface config dhcp`

This command sets the Controller management interface IP address using DHCP. If DHCP is configured, resource pools cannot be set to Fail-Safe mode.

This command takes effect when executed, so remote access to the CLI may be lost if the address changes.

7.1.10. To set a static IP address for the Controller management interface

Run the command: `set controller management-interface config static <IP-address> <netmask> <gateway-IP> [<dns-server-IP>] [<dns-server-IP2> <dns-search>]]`

This command sets a static IP address for the Controller management interface. The DNS configuration information is optional. The ability to specify a DNS search path requires the specification of two DNS servers.

This command takes effect when executed so remote access to the CLI may be lost if the address changes.

7.1.11. To display the current Controller hostname

Run the command: `show controller hostname`

7.1.12. To display a summary of the current configuration and status of the management interface

Run the command: `show controller management-interface`

7.1.13. To display configuration values for the management interface

Run the command: `show controller management-interface config`

7.1.14. To display the current default gateway for the Controller

Run the command: `show controller management-interface default-gateway`



7.1.15. To display the current DNS configuration for the Controller

Run the command: `show controller management-interface dns-server`

7.1.16. To display the current IP address of the Controller management interface

Run the command: `show controller management-interface ip-address`

7.1.17. To display the current netmask of the Controller management interface

Run the command: `show controller management-interface netmask`

7.1.18. To display the software version of the Controller

Run the command: `show controller version`