

Data Acquisition

Best Practices Guide



This Documentation, which includes embedded help systems and electronically distributed materials (hereinafter referred to as the "Documentation"), is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2015 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at <http://ca.com/docs>.

Contents

Chapter 1: Introduction	7
Chapter 2: General Best Practices	9
Identify Mission-Critical Servers and Applications	9
Connect the Server to the Switches that Carry Application Traffic.....	9
Consider the Application Architecture	10
Mirror Ports Work Better than Network Taps	10
Do Not Oversubscribe the Output Capacity of the Mirror Port	10
Avoid Packet Duplication	11
Limit the Traffic Sent to Collection Devices	12
Use a Port with the Largest Buffer Size	13
Use Multi-Port Monitor and Data Aggregation Tools	13
Chapter 3: Best Practices for Application Delivery Analysis	15
Chapter 4: Best Practices for GigaStor	17
Chapter 5: Best Practices for Multi-Port Monitor	19
Chapter 6: Best Practices for UC Monitor	21
Chapter 7: What is Port Mirroring or SPAN?	25
Mirroring with Trunk Ports.....	26
Port Mirroring at Access-Layer Switches.....	27
Avoiding Duplication with Multi-Tier Applications	28
Impact of SPAN on Switch Performance	29
Verify Session Status	29
IOS Configuration for Port Mirroring.....	29
IOS Configuration for Ingress Only	30
What to Mirror in a Multi-tier Application	30
Chapter 8: What is VSPAN?	33
Packet Duplication.....	34
VSPAN with Trunk Ports	35

VSPAN with Layer 3 Switches	37
Example: VLAN as a Source	38

Chapter 9: What is a VACL? 39

What is an ACL?	39
What is an Access Map?	41
Configure an Access Map	42
Using VACLs to Capture Traffic	43
Capturing All Traffic on a VLAN	43
Filtering Captured Traffic	45
Blocking Traffic for Application Delivery Analysis	46
Configuring Multiple Ports to Capture Data	47

Chapter 10: What is Remote SPAN? 49

IOS Commands for RSPAN with VACL	50
--	----

Index 53

Chapter 1: Introduction

Data acquisition is a crucial component of effective network management. The most sophisticated monitoring applications can only report on the data they have. If that data is acquired in a manner that allows duplicated or asymmetric network traffic, the resulting performance and health metrics can be skewed.

CA Application Delivery Analysis, CA Multi-Port Monitor, and CA UC Monitor measure and analyze specific types of network traffic. Sending traffic that these applications cannot monitor only adds load to the server. In extreme cases, packets are discarded. NI GigaStor records all traffic that is sent to it. Duplicate packets severely reduce the amount of time that NI GigaStor can retain data. Duplicate packets make packet-capture analysis difficult.

This document explains the recommended methods for acquiring data from switches. Following these methods can help you prevent issues such as duplicate packets, asymmetric traffic, and oversubscription of switch ports. These best practices help ensure the accuracy of metrics calculated from captured data and reduce unnecessary load on network monitoring tools.

Chapter 2: General Best Practices

Best practices for port SPAN (port mirroring) and VSPAN configuration focus on methods that limit unnecessary, non-critical, and duplicated traffic. This section identifies general best practices appropriate for the following CA applications:

- Application Delivery Analysis
- GigaStor
- Multi-Port Monitor
- UC Monitor

Identify Mission-Critical Servers and Applications

Identify the mission-critical applications and servers that you want to monitor. Then, target the switches where you want to configure port mirroring.

Review the current VLAN configuration for each switch. If servers of interest are all on the same VLAN, include the VLAN in mirror configuration, or [VSPAN](#) (see page 33) in a Cisco environment. Consider mirroring multiple VLANs to capture all the traffic you want to monitor. For each VLAN, consider how many hosts are included and the resulting collection and capture load. Also be aware of the possibility of packet duplication.

Connect the Server to the Switches that Carry Application Traffic

Select a rack with cable-ready access to all switches that carry data to and from the larger enterprise network. Access-layer switches are the best candidates because they typically send fewer duplicate packets to the collection devices.

Configure a port on each access switch as a mirror output (destination) port. Verify that the traffic of interest is forwarded to the capture card on the server.

We recommend that you collect traffic at the same switch where the monitored servers are connected. This practice lets traffic be seen at the same time that it is transmitted to or from the server, resulting in more accurate data.

Consider the Application Architecture

When you mirror data from servers that support a [multi-tier application](#) (see page 30) architecture, collection devices send duplicate packets to the monitoring applications. The servers in a multi-tier architecture send data back and forth among themselves. When a server with mirror port sees a packet in both transmit and receive directions, the packet is sent to the mirror port. Generally, include only front-end servers in the port mirroring configuration. Mirroring the middle-tier servers sends duplicate packets because both transmit and receive packets are mirrored.

Mirror Ports Work Better than Network Taps

However, you can connect a monitor port to a standard copper or fiber tap. You can also connect a monitor to an aggregating tap rather than a mirror port. For example, use a network tap if mirror ports are already used for another purpose, such as an intrusion detection system (IDS). Purchase a tap that sends the request and the response traffic over the same connection on the tap.

Purchase taps that support pass through on failure. If a tap fails without a pass through or fail-closed mechanism, data ceases to flow through the switch. The pass-through mechanism helps ensure that data stops flowing toward the monitoring application but passes through the switch ports.

Do Not Oversubscribe the Output Capacity of the Mirror Port

In high-traffic situations, you can limit the amount of traffic on the SPAN or mirror port. For example, set an Access Control List (ACL) on the mirror port to forward only traffic from key servers. With an [ACL](#) (see page 39), unnecessary traffic is discarded before it is sent out the mirror port. Cisco 4500 Series switches support the use of an ACL.

If you use an ACL, verify that all TCP traffic is forwarded to the monitor. Then add other protocols used by the critical applications you want to monitor. Specify the appropriate ports in the port mirroring statement.

Avoid situations in which a large-capacity switch sends data from all ports to one SPAN or mirror port. In these situations, data is lost.

Avoid Packet Duplication

The term [packet duplication](#) (see page 34) refers to reporting on the same traffic multiple times as it passes through interfaces on a switch. Several port mirroring configurations can result in duplication. The presence of duplicate packets can skew the metrics that are collected. Packet loss statistics are affected because duplicate packets are viewed as retransmissions.

As a best practice, configure mirror ports to minimize or eliminate duplicate packets. Another option for avoiding duplication is to mirror only packets traveling in the receive direction. This setup excludes traffic coming from clients into the VLAN.

When a VLAN sends all traffic to the mirror port, duplication occurs because traffic from all ports is forwarded to Application Delivery Analysis. In cases where packet duplication is likely, consider mirroring individual ports rather than whole VLANs. With this technique, only individual ports or interfaces are used as mirror sources. Only packets destined for selected servers are sent to the mirror port. Use the **show** command to see a list of all ports included in a VLAN.

Deduplication logic applies to all packets received on a given logical port. Therefore, if a duplicate packet from the same VLAN is received on a different logical port, it is not discarded. If you combine two physical ports into a single logical port definition, a duplicate is discarded in the following situations:

- If it arrives on a physical port within a few packets of the original packet on the other physical port
- If it arrives on a second switch.

Both packets are retained if the two physical ports are not combined into a logical port.

Note: The *CA Application Delivery Analysis Administrator Guide* provides instruction for deduplicating TCP packets on a CA Standard Monitor.

Limit the Traffic Sent to Collection Devices

When using applications that rely on specific traffic types, such as Application Delivery Analysis or GigaStor, filter the traffic to reduce volume as much as possible. Several technologies can help you limit the amount of data sent to monitoring applications.

VSPAN

A [VSPAN](#) (see page 33) is a SPAN port that uses a VLAN or multiple VLANs as the source. All the ports in the source VLANs are the source ports. If both ingress and egress are configured, packet duplication occurs each time packets are switched on the same VLAN. Use VSPANS to forward relevant traffic to the appropriate SPAN port and remove unnecessary packets. Otherwise, the captured VLAN traffic traverses multiple physical interfaces, which creates duplicate traffic.

Do not set up VSPAN sessions on your core switches. Instead, set up VSPAN sessions on your access-layer switches where packets are duplicated as they pass between switches at each layer.

VACL

A [VACL](#) (see page 39) is an Access Control List applied to a VLAN. All packets that enter the VLAN are verified against the rules in the list, such as packet type or destination. A VACL limits the amount of data sent over the SPAN port by denying certain types of data. VACLs are supported on Cisco 6500 Series switches.

A VACL filters unneeded traffic so that it is not sent to the SPAN port. A VACL allows you to filter by protocol.

Note: We recommend that multiple people review VACL configurations. The review helps to prevent misconfiguration that can result in dropped traffic. We also recommend that you test a VACL in a lab environment before applying it to a production environment.

RSPAN

[RSPAN](#) (see page 47) is an alternative method for monitoring multi-tier application traffic. RSPAN captures the traffic on one switch, mirrors it to a VLAN, and forwards the traffic to destination ports for analysis. Coupling this technology with Layer 2, 3, and 4 security provided by VACLs gives Application Delivery Analysis visibility into various applications without overloading the switch port. The RSPAN scenario spans the proper VLANs to the capture port while the VACL limits the captured traffic to prevent overloading of the switch port. You can use this technique to enable duplicate packet filtering to help ensure accurate results in Application Delivery Analysis.

Use a Port with the Largest Buffer Size

CA Application Delivery Analysis, CA Multi-Port Monitor, and CA UC Monitor perform passive monitoring of network traffic. Data is typically sourced from multiple Gigabit interfaces and then sent out of a single Gigabit interface. This many-to-one relationship means that it is easily possible to overrun the buffer on the destination interface of the switch. The resulting congestion causes the switch to discard packets. The monitoring application assumes the presence of packet loss and reports inaccurate volumes and rates.

We recommend exporting mirrored data to a port on the Ethernet module with the largest buffer size per port. You can obtain a list of Cisco 6500 modules and the buffer depth per port on each module from the [Cisco website](#). Use this list and the **show module** command to determine the best locations from which to export traffic. The increased buffer depth decreases the likelihood of packet loss at each switch port, which helps ensure that each packet is counted.

Use Multi-Port Monitor and Data Aggregation Tools

In many cases, you may need to gather data from multiple points within the network. The easiest way to combine and send this data to Application Delivery Analysis is to use Multi-Port Monitor. Multi-Port Monitor lets you aggregate 4x1 Gbps, 8x1 Gbps, or 2x10 Gbps streams of traffic.

You can create logical ports to define whether interfaces are treated as separate streams of traffic or combined into one stream. Multi-Port Monitor offers hardware filtering to remove unnecessary traffic and packet slicing to reduce disk usage.

Data aggregation tools from other vendors let you combine traffic feeds, which you can filter and slice before sending to Application Delivery Analysis or other monitoring applications.

Chapter 3: Best Practices for Application Delivery Analysis

Application Delivery Analysis monitors specific types of traffic. Filter the traffic to reduce volume as much as possible.

- Access-layer switches carry the application (TCP) data that Application Delivery Analysis monitors.
- Use port [SPAN](#) (see page 25) and source only from the ports that are directly connected to servers of interest.
- Use [VACLs](#) (see page 39) to allow only TCP traffic. If you are unsure which ports your applications run on, use server IP addresses. After Application Delivery Analysis identifies the application ports, modify the VACL to allow only the ports that Application Delivery Analysis monitors.
- Apply destination filtering or capture port filtering to specify which VLANs exit a destination or capture port. Use this type of filtering on a distributed Application Delivery Analysis system to send VLANs to different collection devices with a single session. One SPAN session can have multiple destination interfaces.
- For monitoring VLANs in a virtual environment:
 - If you have separate VLANs for front-end servers (traffic coming from outside the ESX) and back-end servers (internal traffic inside the ESX), mirror the front-end server traffic to a physical monitoring device.
 - If you do not have separate VLANs for front-end and back-end servers, mirror all the traffic in the ESX and mirror the external traffic coming to the ESX to a physical collector, then “pin” the front-end servers to the physical monitoring device.
- When you use two core switches to load-balance a server, assign the server to the two Application Delivery Analysis monitors mirrored off each switch. Application Delivery Analysis combines the metrics from both monitors when reporting the traffic. For information about assigning a server to more than one monitor feed, see the *CA Application Delivery Analysis Administrator Guide* or online help.
- Do not feed traffic that is captured on either side of a firewall to the same Application Delivery Analysis monitor. The firewall may disrupt the order of the TCP sequence numbers, making it impossible for Application Delivery Analysis to correlate associated sequences.

More information:

[Configuring Multiple Ports to Capture Data](#) (see page 47)

Chapter 4: Best Practices for GigaStor

When using GigaStor to capture data for long-term forensic analysis, consider the following:

- Capture all traffic of interest on the devices being monitored. In general, use [port mirroring](#) (see page 25) sessions because IP VACLs filter out traffic that is not Layer 3, such as STP.
- Configure destination interfaces as trunk ports to export VLAN headers so that VLAN statistics can be calculated. Configure the interface as a trunk port before configuring it as a destination interface. Shut down the port before configuring it to prevent a spanning tree recalculation. Use the **vlan allowed** command on the trunk port to filter VLANs that are allowed to leave the destination interface.
- Reduce packet loss at the destination interface. Use the **switchport trunk allowed vlan** command to specify which VLANs are sent out of a destination interface. With this command, you can split the SPAN across multiple destination interfaces and reduce contention. Because GigaStor has up to eight 1-Gigabit interfaces or two 10-Gigabit interfaces, it is well-suited for using multiple destination interfaces.
- If backup traffic consumes too much storage space on the GigaStor, filter the GigaStor active instance to prevent it from writing that data to disk.

Chapter 5: Best Practices for Multi-Port Monitor

Mirror the switch ports, where traffic travels to and from the monitored servers, to the ports where Multi-Port Monitor is connected.

Exclude irrelevant traffic

Application Delivery Analysis measures and analyzes only TCP network traffic. Therefore, sending additional traffic through the mirror port adds unnecessary load to the capture card on Multi-Port Monitor. In extreme cases, the unneeded data can cause packet loss. However, Multi-Port Monitor analyzes traffic composition and performance metrics from all active protocols on the network. Consider these valuable metrics, which are complementary to the TCP metrics of Application Delivery Analysis, when deciding which traffic is irrelevant.

Minimize or eliminate duplicate packets

Multi-Port Monitor provides a packet deduplication setting that applies to the capture card and is enabled by default. This setting discards packets deemed to be duplicates of packets already received and processed if they arrive within a few packets of each other.

During initial port mirroring configuration, you can temporarily disable the global setting for packet deduplication. Disabling the setting lets you see duplicate packets, which can help you eliminate duplication from mirrored sessions.

More information:

[Avoid Packet Duplication](#) (see page 11)

[Limit the Traffic Sent to Collection Devices](#) (see page 12)

Chapter 6: Best Practices for UC Monitor

In a Cisco environment, you can connect the collector to a [SPAN](#) (see page 25) port on the switches that carry VoIP traffic on your network. The collector must be able to inspect call setup-related packets that pass between the call servers and endpoints in your system. Connect the collector at an appropriate network location, using a properly configured SPAN switch port.

If possible, SPAN only VoIP traffic and call server traffic

- Add VLANs dedicated to voice traffic to the list of spanned ports.
- Configure an [ACL](#) (see page 39) to discard non-VoIP traffic before it is sent out of the SPAN port. An ACL has significant overhead on some switches, so use your best judgment in selecting a method to separate traffic.
- Use a [VACL](#) (see page 39) to filter for all traffic going to and from the IP addresses of the Cisco Unified Communication Managers.

Install UC Monitor collectors near monitored call servers or voice gateways

Calculations for server response times assume that the collector and the monitored server receive an inbound frame at approximately the same time. Similarly, traceroute investigations provide the most accurate path results when the collector is located close to the call server, and the collector and call server share a router. This configuration helps ensure that traceroute results reflect the path taken by packets sent to and from the call server.

Consider NIC usage

A distributed collector and a standalone server use one NIC for the collector service and one NIC for management.

When a network tap is used, one NIC is used for management and two NICs are used for the collector service. A tap splits the transmitting and receiving flows of a full-duplex link across two physical ports. The collector receives the separated traffic from both NICs.

The collector works better with SPAN ports than with network taps

Plug the monitor NIC on the collector into the SPAN port. Configure a network tap only when the SPAN port is in use. You can configure remote spanning of switch traffic to another switch and connect the collector to the SPAN port on the remote switch. Having a dedicated collector for each switch is the recommended configuration.

Forward the following VoIP protocols to the collector:

- **SIP.** Session Initiation Protocol performs call setup and teardown, manages sessions, and determines user location, availability, and capabilities. Forward TCP and UDP traffic on port 5060 to the collector to get SIP flows. UC Monitor supports authenticated SIP flows, but it cannot interpret or report on calls from encrypted SIP flows.

Note: Collectors can monitor only one port per protocol. To monitor multiple SIP ports on a single SPAN, contact [CA Technical Support](#).

- **H.323.** This family of protocols supports real-time transfer of data over packet networks by enabling communications between H.323-enabled devices, such as VoIP gateways. Forward the following traffic to the collector:
 - for H.225, TCP traffic on port 1720
 - for H.245, TCP traffic on ports above 11000 (dynamically selected)
- **SCCP.** The Skinny Call Control Protocol is a proprietary Cisco protocol used for messaging between skinny clients and a Cisco Unified Communications Manager. Forward TCP traffic on port 2000 to the collector to get SCCP flows. If SCCP signaling flows are encrypted, UC Monitor cannot interpret or report on calls.
- **MGCP.** The Media Gateway Control Protocol enables call-control devices, such as media gateway controllers, to control VoIP calls. Performs similar functions to H.323. Forward UDP traffic on port 2427 to the collector to get MGCP flows.
- **Q.931.** Call setup traffic is backhauled over a TCP connection to Cisco Unified Communications Manager to control the ISDN PRI for the voice gateway. The collector must see the Q.931 Setup and Alerting messages that are sent between the Unified Communications Manager and the gateway. Forward TCP traffic on port 2428 to the collector to get the PRI backhaul flows.

Verify that the SPAN port is not congested or misconfigured

Review the dropped packets statistics on the SPAN port where you want to connect the collector.

If you are running Application Delivery Analysis on the network, connect Application Delivery Analysis to a SPAN port on a different switch

In some situations, however, both systems must monitor the same switch. Therefore, you can use a network tap to separate the traffic intended for Application Delivery Analysis from the traffic intended for UC Monitor. The collector can be connected to a standard tap (copper or fiber) or an aggregating tap rather than a SPAN port. Application Delivery Analysis supports a dual-NIC collector, which a standard tap requires. UC Monitor does not support this type of collector. Purchase a tap that sends the request and the response traffic over the same connection on the tap.

Consider traffic from Publishers and Subscribers

Cisco Unified Communications Manager servers act in multiple roles to provide failover and load balancing capabilities. Therefore, configure the SPAN port so that VoIP-related traffic is associated with both Publishers and Subscribers.

Voice gateways can register with either the Publisher or Subscriber and are often configured to perform failover duties. Signaling between a voice gateway and a call server must reach the collector so that all call traffic on the network is monitored.

In a large Cisco deployment, each cluster typically includes as many as four servers: one Publisher, two Subscribers to handle call processing, and one TFTP server. The Publisher and TFTP server act as backups for failover situations. In such a configuration, SPAN the traffic that goes to two of the servers in the cluster and send it to one collector. Then, send the traffic for the other two servers to a second collector. If the phones are load-balanced among the servers in the cluster, each collector sees data from only half the phones, which is a manageable load.

Chapter 7: What is Port Mirroring or SPAN?

A port mirroring session sends a copy of packets from one switch to a port on the destination (or monitor) switch. Mirror individual ports whenever possible. Ideally, mirror ports that are directly connected to the servers that host the applications of interest. The ideal location is a core switch in a network operations or data center. However, any switch with maximum visibility into the traffic of interest is acceptable.

In a Cisco environment, port mirroring is accomplished with the Switch Port Analyzer (SPAN) feature. SPAN lets you copy traffic from physical ports on a switch to another port on that switch.

You configure port mirrors by creating a monitor session consisting of a source and destination.

A session source consists of the following attributes:

- **Session number:** Differentiates a monitor session from others on the switch
- **Session source:** The physical ports or VLANs from which the SPAN copies data
 - Source ports can be L2 or L3 LAN ports.
 - Trunk and non-trunk ports can be used at the same time.
 - Do not configure WAN interfaces to be source ports (such as ATM interfaces).
 - Do not configure EtherChannel ports as source ports. IOS versions 12.1(13)E and later do not permit such configuration.
 - Do not mix physical ports and VLANs as sources within the same monitor session. Configure either physical ports *or* VLANs.
 - When you specify the source information using a VLAN or VLAN list, the SPAN function is known as VLAN SPAN or VSPAN. Sourcing from a VLAN adds every interface in the VLAN to the monitor session.
- **Session direction:** The direction of the traffic you want to copy: receiving (RX), transmitting (TX), or both (the default)

A *session destination* specifies the physical port to which the mirror port copies data. A destination port can be any physical port.

- With release 12.1(13)E and later of Cisco IOS, you can configure the destination port to be a trunk port. This configuration lets you forward VLAN tags to the collection device. You can use the **switchport trunk allowed vlan** command to filter the data that leaves the destination port.
- A destination port can service only one SPAN session and cannot be an EtherChannel port.
- A monitor session can have up to 64 destination interfaces.

Mirroring with Trunk Ports

When the source port of a SPAN or port mirror session is a trunk port, you can monitor specific VLANs on that port. Use the following command:

```
monitor session session-number filter vlan-list
```


You can configure an interface as a trunk port to filter the VLANs that leave a destination port. Use the following command:

```
switchport trunk allowed vlan vlan-list
```

This command lets you specify which VLANs the interface transmits. An interface configured as a destination interface filters out traffic. By filtering VLANs, you reduce contention for bandwidth and lower the probability of packet discards due to buffers filling.

Port Mirroring at Access-Layer Switches

Port mirroring and SPAN are ideally suited for use at access-layer switches because you can select individual interfaces as sources. For the source of the monitor session, use interfaces that are connected to production servers that host business-critical applications. This configuration helps ensure that data destined for other servers is not sent to the collection device and does not contend for bandwidth at the destination port.




The configuration in the illustration captures all traffic to and from the production server, while ignoring traffic to and from the backup server. The following commands represent the configuration in the illustration:


```
(config)# monitor session 1 source interface f2/17
(config)# monitor session 1 destination interface f2/13
```

Avoiding Duplication with Multi-Tier Applications

A packet that crosses multiple source interfaces is copied twice to the monitoring application, causing additional load on the destination interface and the collection device. The duplication can skew metric calculations. To avoid duplication in this scenario, use only the front tier as the source of a session.



The preceding figure identifies the port to use as the source when a monitoring application talks between two servers. When the front-end server is the source, the monitoring application sees transmitted and received data at the client and at the back-end server. When more than two tiers are talking, use the interface of every alternate tier as a source.



The preceding figure shows how to avoid packet duplication by sourcing the port mirror from every other tier of a multi-tiered application. Notice that each flow of data is only circled (sourced) once, and each flow crosses only one source port.

Impact of SPAN on Switch Performance

On the Cisco 6500 Series switch, each packet received by a port is automatically transmitted on the internal switching bus to every line card. Each line card buffers the packet while the Encoded Address Recognition Logic (EARL) determines the destination of the packet. After the EARL process is complete, the packet is flushed from the buffers for which it is not destined. If the EARL process cannot determine the destination, the packet is sent out of all ports. Because each packet is already copied to each line card, there is no performance impact when configuring a SPAN destination.

The non-blocking architecture of the Cisco 4500 switch also allows for the use of SPAN without performance impact.

Verify Session Status

Use the following command to verify the status of monitor sessions:

```
show monitor sessions
```

```
QARouter-6509-8#show monitor
Session 1
-----
Type                : Local Session
Source Ports        :
   Both             : Fa2/14
Destination Ports   : analysis-module 9 data-port 2

Session 2
-----
Type                : Local Session
Source Ports        :
   Both             : Fa2/14, Fa2/20
Destination Ports   : Fa2/17-18
```

IOS Configuration for Port Mirroring

A monitor session lets you specify the direction in which the port mirror is sourced. You can configure source ports to copy ingress traffic or egress traffic. If you do not specify a direction, the port mirror session copies both directions.

```
(config)# monitor session 1 source interface fa2/17
(config)# monitor session 1 destination interface fa2/13
```

Because this example does not specify a direction, interface f2/17 is sourced from the ingress and egress directions of the switch port.

IOS Configuration for Ingress Only

The number of port mirror sessions that a switch can support depends on the switch model. For example, the Cisco 6500 Series switch supports four sending (TX) sessions and two receiving (RX) sessions. Therefore, a Cisco 6500 Series switch can support up to six separate SPAN sessions, when each session monitors one direction. However, the switch supports only two bidirectional sessions because each direction consumes a TX session and an RX session.


```
(config)# monitor session 1 source interface fa2/17 rx  
(config)# monitor session 1 destination interface fa2/13
```

The Cisco 6500 Series switch also supports multiple destination interfaces on a SPAN session, allowing you to copy production traffic to more than one collection device. A common use for this feature is to send one copy of each packet to a performance monitor and another copy to an IDS or other security system.

What to Mirror in a Multi-tier Application

Use port mirroring to gain visibility into each tier of a multi-tier application. Place collection devices on individual switches to monitor server communications. You do not need to mirror all servers to obtain the appropriate metrics for monitoring. Mirroring all servers yields more duplicate packets and reduces monitor capacity.

For example, consider the following network:



Because the mirroring configuration includes both transmitted and received packets, the number of ports required to be mirrored is reduced. When you mirror a tier, you see the traffic to that tier and to the next tier after it.

In this scenario, the ports to mirror are the Citrix Load Balancer, the NetReport, NetWeb, and NetApp servers. The data for the Oracle back-end server is obtained from the NetReport and NetApp server mirrors. Traffic mirrored from Citrix Load Balancer and NetReport provides visibility into the NetCitrix tier.

Chapter 8: What is VSPAN?

A VLAN SPAN (VSPAN) configuration is a monitor session whose source is a VLAN rather than specific interfaces. When a VLAN is the source of a monitor session, all physical interfaces that are members of the VLAN are also sources.

In a VSPAN configuration, the monitor session is still sourced from physical interfaces.

Advantages


- VSPAN is easy to configure and does not require you to know the interfaces to which every server is connected.
- You use one command to add a group of physical interfaces as sources.
- If servers move ports within the switch or VLAN, the change is transparent, requiring no change to the SPAN session.

Disadvantages

- VSPAN does not let you control which interfaces are the source of a monitor session. Instead, you collect traffic on *every* interface in the VLAN.
- The destination SPAN port may not be able to handle the data it receives, which can result in congestion and discards at the outbound interface.
- VSPAN can cause duplicates in many scenarios. VSPAN is best suited for use on Layer 2 devices where one VLAN is sourced, either ingress-only or egress-only. This configuration allows each packet to be captured only as it enters the switch, or only as it leaves the switch.

Packet Duplication

When a VSPAN is the source of a monitor session, every physical interface of the VLAN is also a source. Packets that travel between two servers on the VLAN are duplicated.




Duplication also occurs when more than one VLAN is sourced and servers in both VLANs communicate with each other. The following commands represent the configuration in the illustration:

```
(config)# monitor session 1 source vlan 10
(config)# monitor session 1 destination interface f2/13
```

VSPAN with Trunk Ports

Packet duplication can occur when a Layer 2 device uses a trunk port to communicate with other switches. By default, a trunk port is a member of all VLANs. Therefore, traffic crossing the trunk is mirrored when a VLAN is the source for the monitor session. When a mirrored VLAN is on an Access-Layer switch attached to the server, traffic is seen as it crosses the trunk port and as it crosses the port connected to the server.



The following commands represent the configuration in the illustration:


```
(config)# monitor session 1 source vlan 10
(config)# monitor session 1 destination interface f2/13
```

In the preceding illustration, traffic leaving Fa2/17 is destined for the host in VLAN 20 or destined for an end user. This traffic also crosses Trunk Port Gig3/4, which is considered a member of VLAN 10. Additional replication occurs if multiple VLANs are sourced.

If VLAN 20 is also a source for the monitor session, traffic that travels between VLAN 10 and VLAN 20 is copied:

- As it is received by Fa2/17 on VLAN 10
- As it is transmitted by Trunk Port Gig3/4 on VLAN 10
- As it is received by Trunk Port Gig3/4 on VLAN 20
- As it is transmitted by Fa2/15 on VLAN 20

This configuration leads to quadruplicate packets for certain conversations




The preceding illustration shows a Layer 2 device that uses a trunk port to send traffic to a Layer 3 device. This configuration leads to duplication of all traffic between client and server and intra-VLAN. Also, traffic is quadruplicated when hosts in VLAN 10 and VLAN 20 communicate with each other.

In this scenario, VLANs are sourced in only in the RX direction, because there are more RX sessions available. Data is mirrored only as it enters the switch. The following commands represent the configuration in the illustration:

```
(config)# monitor session 1 source vlan 10
(config)# monitor session 1 source vlan 20
(config)# monitor session 1 destination interface f2/13
```

VSPAN with Layer 3 Switches

On devices with Layer 3 routing, do not source a monitor session from only one direction of a VLAN. SPAN only monitors data transmitted or received by a physical port. When a packet enters a device from Layer 3 and is routed into the VLAN, the packet is not seen until it is transmitted. As a result, when a VSPAN is sourced from only one single direction, one side of a TCP conversation is missing.



The preceding illustration shows that the SPAN configuration does not see traffic that enters VLAN 10 from Layer 3, because no physical interface is involved. Packets are seen only as they are received by an interface in VLAN 10. Packets destined for hosts in VLAN 10 (TX) are missed. Traffic coming from VLAN 20 or from end users is missed.

The following commands represent the configuration in the illustration:

```
(config)# monitor session 1 source vlan 10 rx
(config)# monitor session 1 destination interface f2/13
```

Example: VLAN as a Source

In the following example of the **show vlan brief** command, VLAN 10 has two ports. If you use VLAN 10 as the source of a monitor session, data from both ports are included in the session.

```
QARouter-6509-8#show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa2/2, Fa2/3, Fa2/12, Fa2/14, Fa2/21, Fa2/22, Fa2/27, Fa2/28, Fa2/48, Gi9/1
10	VLAN0010	active	Fa2/9, Fa2/17
20	VLAN0020	active	Fa2/11, Fa2/13
100	VLAN0100	active	
666	VLAN0666	active	
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup	

Chapter 9: What is a VACL?

A VACL is an Access Control List (ACL) that is applied to a VLAN rather than an interface. This type of ACL can be used to match traffic, which is forwarded to its destination for capture, and then sent to a monitoring device. VACLs are processed in hardware and applied to packets that are bridged on a VLAN or that enter the VLAN from a Layer 3 process.

VACLs are applied as a VLAN processes a packet, which prevents duplication of intra-VLAN traffic. You can create custom filters for VACLs to limit the traffic that is captured. VACLs are supported on Cisco 6500 and 4500 (IOS only) series switches. Capturing data from a VACL is not supported on the 4500 series switch.

A VACL does not consume a SPAN session

Use VACLs only when everyone involved is aware of the risk of misconfiguration and is committed to avoiding this possibility through peer-review of proposed configurations. If VACLs are not an acceptable risk for your organization, consider implementing filtered port SPAN to VACL a copy of production traffic. Or consider using Multi-Port Monitor or a SPAN aggregation tool available from such vendors as Anue Systems and Gigamon.

Important: Use VACLs to capture traffic only from Layer 3 and Layer 4. VACL capture is not recommended for GigaStor deployments, for which the intent is to capture all traffic, including Layer 2 issues.

More information:

[Use Multi-Port Monitor and Data Aggregation Tools](#) (see page 13)

What is an ACL?

An Access Control List (ACL) is an ordered list that matches traffic based on specific characteristics. Each line in the list is an Access Control Entry (ACE). Each ACE contains a condition to match and an action to take when traffic that matches that condition. A packet that does not match an ACE in an ACL is dropped. This process is referred to as the *implicit deny all* on an ACL.

An ACE can filter on:

- Source IP address
- Destination IP address
- Protocol
- Source protocol port
- Destination protocol port

Note: ACLs provide additional filtering options that are beyond the scope of this document.

An ACE follows the format:

```
Action protocol source [port] destination [port]
```

Where the source [port] is the source IP address or subnet and Layer 4 port.

Each ACE is directional and uses an inverse network mask notation. To permit traffic in both directions for a host or network, you need two ACEs. The following example shows the ACEs necessary to permit TCP port 80 traffic traveling from and to a 24-bit subnet.

```
permit tcp 192.168.0.0 0.0.0.255 eq 80 any
permit tcp any 192.168.0.0 0.0.0.255 eq 80
```


What is an Access Map?

ACLs are applied to VLANs using an access map. An access map can contain multiple ACLs, each with an assigned priority number to specify the order in which it is processed. Each ACL is also assigned an action to take when the ACL is matched.

When a packet matches a *permit* ACL entry, it is forwarded along its intended path. The packet is not inspected against ACL lines or ACL maps. When a packet matches a *deny* ACL entry, it is verified against the next entries. If a permit exists for the next entries, the packet is forwarded into the network. A packet is dropped *only* when it does not match a permit ACL entry.



The preceding illustration depicts an access map that contains two ACLs. Each ACL is associated with a separate set of actions. The priority number determines which ACL is processed first (the lowest).

We recommend that you name ACLs numerically and that you name access maps in plain text. This strategy helps to prevent confusing the two. Add a description to each ACL to denote its purpose.

Configure an Access Map

Use this procedure to configure an access map.

Follow these steps:

1. Define the access map using the following format:

```
VLAN access-map map_name [0-65535]
```

For example:

```
vlan access-map SA-Capture 10
```

2. Configure the match clause using the following format:

```
match ip address acl_name
```

For example:

```
match ip address 101
```

3. Configure the action clause using the following format:


```
Action {forward | forward [capture] | drop | redirect }
```

For example:

```
action forward
```

4. Apply the access map to the VLAN using the **vlan filter** command.

Example:



The VACL in the example is configured to deny ICMP traffic on VLAN 10. The ACL permits TCP and UDP traffic only, which implicitly denies all other Layer 3 traffic. The access map applies the ACL with a priority of 30 and specifies that traffic meeting the ACL is forwarded to its destination. The **vlan filter** command applies the access map to VLAN 10.

Using VACLs to Capture Traffic


Use VACLs to capture traffic by specifying an action of **forward capture** in the VLAN access map. This command specifies the following:

- Traffic that meets the ACL is forwarded to its destination.
- A copy of the traffic is sent out of the interfaces that are configured as capture ports.

Capturing All Traffic on a VLAN

In the following example, all traffic in VLAN 20 is forwarded to its destination and a copy sent to configured capture ports. This configuration is ideal for capturing all traffic on a VLAN because it does not create duplicates for intra-VLAN traffic, as a VSPAN would. Use this configuration to send data to Application Delivery Analysis without filtering out traffic.

This example also represents the simplest VACL to implement. This implementation poses no risk of the ACL blocking the production VLANs to which it is applied.



The following commands represent the configuration in the illustration:

```
(config)#access-list 101 permit ip any any
!
(config)#vlan access-map sa_cap 30
(config)#match ip address 101
(config)#action forward capture
!
(config)#vlan filter sa_cap vlan-list 20
!
(config)#interface fa2/13
(config)#switchport capture
```


Filtering Captured Traffic

The following example shows how to use a VACL to filter the traffic that is captured in a VLAN.

- The first ACL in the access map (ACL 101) specifies that TCP traffic is forwarded to its destination and captured.
- The second ACL (ACL 102) specifies that traffic that does not meet ACL 102 is retained.
- Interface fa2/13 is configured as a capture port.

The example filters out non-TCP traffic, which is useful when capturing data for Application Delivery Analysis. Using ACLs, you can limit traffic to specific ports or IP addresses.

Note: If the second ACL is configured improperly, production traffic that does not meet the first ACL is dropped. We recommend that multiple people review VACL configurations. The review helps to prevent a misconfiguration that can result in dropped traffic. We also recommend that you test a VACL in a lab environment before applying it to a production environment.



Time	10.10.20.104	10.10.10.103	Comment
0.000	2382 > telnet [SYN]		TCP: 2382 > telnet [SYN] Seq=0 Len=0 MSS=1460 WS=2 TSV=0 TSER=0
0.000	telnet > 2382 [RST]		TCP: telnet > 2382 [RST, ACK] Seq=0 Ack=1 Win=0 Len=0
0.449	2382 > telnet [SYN]		TCP: 2382 > telnet [SYN] Seq=0 Len=0 MSS=1460 WS=2 TSV=0 TSER=0
0.449	telnet > 2382 [RST]		TCP: telnet > 2382 [RST, ACK] Seq=0 Ack=1 Win=0 Len=0
0.950	2382 > telnet [SYN]		TCP: 2382 > telnet [SYN] Seq=0 Len=0 MSS=1460 WS=2 TSV=0 TSER=0
0.950	telnet > 2382 [RST]		TCP: telnet > 2382 [RST, ACK] Seq=0 Ack=1 Win=0 Len=0

The following commands represent the configuration in the illustration:

```
(config)# access-list 101 permit tcp any any
!
(config)# access-list 102 permit ip any any
!
(config)# vlan access-map sa_cap 30
(config)# match ip address 101
(config)# action forward capture
!
(config)# vlan access-map sa_cap 40
(config)# match ip address 102
(config)# action forward
!
(config)#vlan filter sa_cap vlan-list 10
!
(config)#interface gig2/13
(config)#switchport capture
```

Blocking Traffic for Application Delivery Analysis

A common usage for VACLs is to block backup traffic from being sent to Application Delivery Analysis. To deny specific traffic, create the access map and ACLs differently than in [Filtering Captured Traffic](#) (see page 45). When you deny traffic to and from a backup server and then use the **permit ip any any** command in an ACL, the access map both denies and permits the same traffic. In an access map, each packet is verified against each entry in the ACL before being discarded.

Instead, create an ACL that matches traffic going to or from the backup servers and a pass-through filter, as shown in the following example:

```
(config)# access-list 101 permit ip host 192.168.1.1 any
(config)# access-list 101 permit ip any host 192.168.1.1
(config)# access-list 101 permit ip host 192.168.1.2 any
(config)# access-list 101 permit ip any host 192.168.1.2
!
(config)# access-list 102 permit ip any any
!
(config)# vlan access-map sa_cap 30
(config)# match ip address 101
(config)# action forward
```

```
!  
(config)# vlan access-map sa_cap 40  
(config)# match ip address 102  
(config)# action forward capture  
!  
(config)# vlan filter sa_cap vlan-list 10  
!  
(config)# interface gig2/13  
(config)# switchport capture
```

In this example, the backup traffic is matched first and forwarded, but not captured. All other traffic is forwarded and captured. With this concept, you can create a filter that performs the following tasks:

- Forwards backup traffic.
- Captures the traffic of interest, such as TCP for Application Delivery Analysis.
- Forwards all IP traffic that does not meet the filtering criteria of the first two ACLs.

This configuration requires three ACLs.

Tip: Use this configuration to capture traffic from your production application server IP addresses without sending server backups to Application Delivery Analysis. Sending the backups to Application Delivery Analysis can overwork its collection devices.

Configuring Multiple Ports to Capture Data

VACL capture allows for multiple ports to be configured as capture ports. Use this configuration to send captured data to multiple monitoring devices. For example, you can send a copy of all traffic to GigaStor and another copy to an IDS system. Configure a monitor session when you use VACLs to filter traffic and you need an unfiltered copy of the traffic for an IDS or probe.

When configuring capture ports, use the following command, which provides several options:


```
Router(config-if)# switchport capture allowed vlan {add | all | except | remove}  
vlan_list
```

switchport capture

Instructs the switch to send captured traffic from all VLANs out of the capture port.

switchport capture allowed vlan

Specifies which VLANs to send out of a capture port. Use this command when the volume of the captured traffic is too large to buffer and send on one interface. To determine whether this situation is occurring, monitor the capture port for discards over time using an SNMP poller.



The preceding diagram shows how traffic is captured and filtered from VLANs 50 and 60. Three capture ports are configured, one to accept only VLAN 50, one to accept only VLAN 60, and a third to accept both VLANs.


Tip: Connect the first two capture ports to Application Delivery Analysis collection devices to reduce the risk of discards at the switch port and to prevent overloading of the Application Delivery Analysis server. Connect the third capture port to a device that does not need every packet to compute its metrics. For example, devices that use sampling techniques, such as EMC Application Discovery Monitor or certain IDS solutions.

Chapter 10: What is Remote SPAN?

A remote SPAN (RSPAN) is sourced in the same manner as a traditional SPAN, either from individual ports or all ports in a VLAN. An RSPAN consumes one SPAN session in the same way that a local SPAN does. However, the RSPAN uses a VLAN for a destination instead of an interface.

We do not recommend the use of RSPAN for collecting data across multiple switches. Instead, use Multi-Port Monitor or a SPAN aggregation tool, both of which let you gather data from multiple switches with less risk and more flexibility than RSPAN.


You can an RSPAN with a [VACL](#) (see page 39). With this combination, you can filter traffic before it leaves the switch, without the risk of applying a VACL to a production VLAN.



The preceding illustration shows how data from the source VLANs is copied to a new RSPAN VLAN. The following commands represent the configuration in the illustration:

```
monitor session 1 source interface gigabitethernet1/28
monitor session 1 destination remote vlan 999
```

The following illustration shows how a VACL is applied to the RSPAN VLAN.



On the RSPAN VLAN, a VACL captures only the data of interest for Application Delivery Analysis or UC Monitor. Add an ACL to forward only the traffic that is captured. Traffic that is dropped has no impact to the production VLANs. In the diagram, the VACL shown in red is a restrictive ACL that does not capture all traffic. Notice that there is no second ACL to forward all other traffic. Only one monitor session is used. Do not configure a second monitor session from the RSPAN VLAN. The VACL captures all of the data that meets the ACL when the **switchport capture** command is applied to the interface connected to Application Delivery Analysis or UC Monitor.

IOS Commands for RSPAN with VACL

- Description - RSPAN VLAN 777

```
vlan 777
remote-span
```
- Interface gigabit ethernet 9/1

```
description SA Monitor Interface
no ip address
no shut
```
- Interface VLAN 777

```
description RSPAN VLAN - Must exist for VACL RSPAN
no ip address
shutdown
```

- Interesting TCP traffic ACL

```
access-list extended Monitored-TCP
permit tcp ip any host x.x.x.x
permit tcp ip host x.x.x.x any
[...]
```

- Define VLAN Access-map

```
vlan access-map RSPAN-VACL 10
match ip address Monitored-TCP
action forward CAPTURE
```

- Map the VACL to the RSPAN VLAN

```
vlan filter RSPAN-VACL VLAN 777
```

- Monitor session 1 captures bidirectional traffic from uplink ports to RSPAN VLAN 777

```
monitor session 1 source g1/1,g2/1
monitor session 1 destination remote vlan 777
```

- Set int g9/1 for capture

```
Int g9/1
switchport capture
end
```


Index

A

- access maps
 - configuring • 42
 - defined • 41
- access-layer switches • 9, 27
- ACL
 - and access maps • 41, 42
 - and ACE • 39
 - defined • 39
 - limiting traffic • 10
- Application Delivery Analysis
 - best practices • 9, 15
 - blocking traffic • 46
 - configuring capture ports • 47
 - data aggregation • 13
 - increasing buffer depth • 13
 - limiting traffic • 12, 15

D

- deduplication • 11

F

- forward capture • 43

G

- GigaStor
 - best practices • 9, 17
 - limiting traffic • 12

I

- implicit deny all • 39

L

- limiting traffic • 10, 12, 19

M

- Multi-Port Monitor
 - best practices • 9, 19
 - data aggregation • 13
 - increasing buffer depth • 13
 - limiting traffic • 19
 - packet duplication • 19

- multi-tier applications • 10, 28, 30

N

- network taps • 10, 21

P

- packet duplication
 - and Multi-Port Monitor • 19
 - and VSPAN • 34, 35
 - avoiding • 11, 28
- permit ip any any • 46
- port mirroring
 - and trunk ports • 26
 - at access-layer switches • 9, 27
 - defined • 25
 - identifying servers • 9
 - instead of network taps • 10
 - IOS configuration • 29
 - limiting traffic • 10, 12
 - packet duplication • 11
 - port mirroring, with multi-tier applications • 28, 30

R

- RSPAN
 - defined • 49
 - IOS commands • 50
 - limiting traffic • 12

S

- show • 11
- show module • 13
- show monitor sessions • 29
- show vlan brief • 38
- SPAN
 - and trunk ports • 26
 - at access-layer switches • 9, 27
 - defined • 25
 - identifying servers • 9
 - instead of network taps • 10, 21
 - IOS configuration • 29, 30
 - limiting traffic • 10, 12
 - packet duplication • 11
 - performance impact • 29

- with multi-tier applications • 28, 30
- switchport capture allowed vlan • 47
- switchport trunk allowed vlan • 17, 26

U

UC Monitor

- best practices • 9, 21
- forwarding VoIP protocols • 21
- forwarding VoIP traffic • 21
- increasing buffer depth • 13
- limiting traffic • 21
- NIC usage • 21

V

VACL

- and access maps • 41, 42
- capturing traffic • 43
- configuring capture ports • 47
- defined • 39
- filtering traffic • 45
- limiting traffic • 12
- with RSPAN • 49, 50

vlan allowed • 17

VoIP protocols • 21

VSPAN

- and Layer 3 switches • 37
- and trunk ports • 35
- defined • 33
- limiting traffic • 12
- packet duplication • 34, 35