



Taps vs. SPAN

The Forest AND the Trees: Full Visibility into Today's Networks



Table of Contents

The First Line of Defense: Access	5
Problem #1: Dropped Packets	5
Problem #2: Switch Configuration	6
Problem #3: Switch/SPAN security vulnerability	7
Problem #4: Not Passive	7
The Tap Alternative	7
Conclusion.....	9

These days, your network is as important to your business as any other item—including your products. Whether your customers are internal or external, you need a dependable and secure network that grows with your business. Without one, you are dead in the water.

IT managers have a nearly impossible job. They must understand, manage, and secure the network all the time against all problems. Anything less than a 100 percent working network is a failure. There is a very familiar saying: Don't miss the forest for the trees. Meaning don't let the details prevent you from seeing the big picture. But what if the details ARE the big picture? Today's IT managers can't miss the forest OR the trees!

Anything less than a 100 percent working network is a failure.



To gain competitive advantage, enterprises race to implement the latest advances in virtualization, software as a service (SaaS), and software-defined networking. But as today's networks become larger and more complex, the challenge to view, manage, and protect them becomes more complex as well. Poor network visibility obscures insight into your applications, infrastructure, and security. And the hampers your ability to serve your customers.

Network visibility is a prime tool in properly monitoring your network. You need an end-to-end visibility architecture to truly see your network. This visibility architecture must reveal both the big picture and the smallest details to present a true view of what is happening in the network.

The first building-block to your visibility architecture is access to the data. That comes in one of two forms: a network tap, or a switch port analyzer (SPAN) port (also known as port mirroring).

But which is the right one? This document should help answer that question.

Why is network visibility so essential?

Simply put, the network is always vulnerable. Total insight into all traffic is critical to identifying and quickly resolving issues. With the trend to virtualization, largely unmonitored virtual machine (VM) traffic increasingly exposes networks to attacks, noncompliance, loss of availability, and impaired performance. Most network issues are still discovered by people, rather than by the technology designed to detect these issues. Often it's your customers that discover the issues!

Security threats are ingenious and opportunistic, always probing for network or application weaknesses. Everyday, a new security breach is reported—often about lost customer data.

A vulnerable network is unacceptable.

Spiraling pressure on monitoring tools led to a plague of the notorious “blind spots”— areas of the network that the tools cannot see to monitor. These fragmented areas are a perfect distillery for latent errors and pre-attack activity.

The First Line of Defense: Access

So how do you get visibility of all network traffic? There are many ways of constructing network monitoring, but the first and most important aspect is access to the data. This means getting the traffic flowing through your network to monitoring tools. There are two main ways of doing this: taps and SPAN ports.


SPAN or mirror ports are an inexpensive way to access traffic moving through a network switch. You can configure SPAN-supporting switches to mirror traffic from selected ports or virtual local area networks (VLANs) to the SPAN port. From there, you attached monitoring tools. At first glance, it seems that a SPAN port could be a good way to connect an intrusion detection system (IDS), forensic recorder, or other security monitoring device.

Unfortunately, SPAN ports can be troublesome and risky in application. This includes:

- Dropped packets
- Switch configuration
- Switch/SPAN security vulnerability
- Not passive

Problem #1: Dropped Packets

The visibility of network traffic with SPAN ports is less than perfect. A fundamental network monitoring requirement today is that you must be able to see every single packet. SPAN ports cannot meet this requirement because they drop packets. Spanning is the switch’s lowest priority task, and SPAN traffic is the first thing to go when the switch gets busy. Any port on a switch can drop packets because network protocols are specifically designed to be robust in spite of dropped packets. Dropped packets are inevitable in any network.



Different switches may be more or less prone to drop SPAN packets depending on their internal architecture. It’s unlikely, however, that SPAN port performance was an important criteria when the switching gear was selected.

There are many ways of constructing network monitoring, but the first and most important aspect is access to the data.

SPAN ports by their very nature are not transparent for Layer 1 and Layer 2 information.

Let's suppose you have switches with the best possible SPAN performance. Dropped packets are still an issue depending on how much traffic you need to send through the SPAN port. If you need to see all of the traffic on a full-duplex 1 Gigabit (Gb) link, a 1Gb SPAN port won't do the job. Full duplex link traffic exceeds the 1Gb SPAN port capacity when link utilization goes above 50 percent in both directions. To see all the traffic, you need to dedicate a 10Gb port for SPAN. Now the SPAN port doesn't seem so inexpensive any more.

SPAN port visibility issues go beyond simply dropping packets. SPAN ports by their very nature are not transparent for Layer 1 and Layer 2 information. For example, they drop undersized and oversized packets, and packets with cyclic redundancy check (CRC) errors. They usually remove VLAN tags, too.

In addition, SPAN ports do not preserve the packet timing of the original traffic, or even the packet order. This type of information can be critical for certain types of network visibility and security.

For example, network consultant Betty DuBois observed, "... losing the VLAN tag information when Spanning, if there is an issue with ISL or 802.1q, how will I ever know with a SPAN port?"¹

Problem #2: Switch Configuration

SPAN ports require switch configuration to send specific traffic to the SPAN port. This leads to a host of complications:

- **Incorrect configuration of the SPAN port.** For example, if the switch manager configures the SPAN port to show specific information, you lose insight into the traffic that isn't forwarded. You must check the configuration to know why.
- **Sharing the SPAN port.** A switch typically supports only one or two SPAN ports. The switch may need to reassign the SPAN port to other purposes for one reason or another—without notifying anyone. IT Manager Bob Huber recalled, "SPAN was a huge issue ... on the IDS team where I used to work. We had constant issues with the SPAN going up and down. ...the network engineers have priority to the limited number of SPAN ports available. Hoping they remember to reconfigure your SPAN port was a waste of time."²
- **Switch configuration is not always available.** If you need to change the SPAN traffic profile, or change it back after someone else used the port, it's not always easy to get the switch owner's to do it. In larger organizations, you may also need change authorization through a Change Control Board, and then wait for a maintenance window to get it implemented.
- **Changes to the network switches impacts the SPAN traffic.** Networks are constantly being reconfigured to optimize applications or support new requirements. If the monitoring solution depends on SPAN ports, it is vulnerable to changes (planned or unexpected) any time the network is reconfigured for any reason.
- **Switch configuration itself is a security vulnerability.** The network's security is paramount. Switches are a highly vulnerable network point, and the ability to reconfigure them must be tightly controlled. Does it make sense to require switch reconfiguration to configure SPAN, when reconfiguring a switch can accidentally or deliberately expose or bring down the network?

¹ <http://www.lovelytool.com/blog/2007/08/span-ports-or-t.html>

² <http://taosecurity.blogspot.com/2007/12/expert-commentary-on-span-and-rspan.html>

Even Cisco recognizes that SPAN port misconfiguration can be an issue, as evidenced by this note in the Cisco Catalyst 6500 Series documentation: “Connectivity issues because of the misconfiguration of SPAN ports occur frequently in CatOS... Be very careful of the port that you choose as a SPAN destination.”³

Problem #3: Switch/SPAN security vulnerability

SPAN ports are usually configured for unidirectional traffic, restricted to transmitting traffic to the monitoring device. However, in some cases they can receive traffic as well (a feature Cisco calls ingress traffic forwarding), in order to enable management of the monitoring device over the same switch port and monitoring device network interface card (NIC) as the mirror traffic. When this configuration is used, the SPAN port becomes an open ingress port to the switch, creating a serious security vulnerability. Therefore, this configuration should be avoided as a best practice.

Problem #4: Not Passive

A final important consideration when using SPAN ports for monitoring access is that SPAN ports are not passive: They affect the performance of the switch’s other ports. For example, Gerald Combs, the father of Wireshark, warns, “Some switch families (e.g., the Cisco 3500 Series) don’t set a lower priority on SPAN traffic, and will slow down the backplane in order to deliver packets to a SPAN port.”⁴ This effect violates a primary principal of monitoring: monitoring should not affect the traffic being monitored. It may have legal as well as practical implications.

The Tap Alternative

To avoid the problems that SPAN ports bring, experts recommend using network taps for access to the network traffic. Taps are specifically designed to provide 100 percent traffic visibility without any impact on monitored traffic. Optical taps for fiber links use optical splitters to divert part of the light from the link to a monitor port, creating a true copy of the link traffic all the way down to Layer 1 and Layer 2 errors. Taps for copper links perform a similar function electronically. Optical taps do not use any power at all, while copper taps include relays that ensure that link traffic continues to flow even when the tap loses power.

Taps avoid all of the pitfalls of SPAN ports:

- Taps send the monitoring tool an exact copy of the link traffic, including Layer 1 and Layer 2 errors and malformed packets, no matter how busy the link is. They never drop packets.
- Taps require little or no configuration. Once a tap is installed in a link, monitoring access to the link traffic is always available, consistently and persistently.
- Taps do not use a switch port, leaving it to do the job it was intended to do (switching).
- Taps are secure. They do not have an IP address so attackers cannot see them, and they cannot inject traffic into the network under any circumstances. In fact, a tap actually hides the monitoring tool from the network as well, providing true “stealth” monitoring.
- Taps are completely passive. They cannot affect the link traffic—even if they lose power.

³ http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.


⁴ <http://www.lovemytool.com/blog/2007/08/span-ports-or-t.html>

To avoid the problems that SPAN ports bring, experts recommend using network taps for access to the network traffic.

The wide range of tap devices available today enable appropriate monitoring access to be built into all parts of the network architecture, at the edges, distribution, LAN, and core.

Tap technology has evolved to offer a range of additional features as well, most of which are not available with SPAN ports:

- Regeneration taps produce multiple copies of the link traffic so multiple tools and multiple users can view the same traffic simultaneously. Your monitoring device does not need to give up access when the network administrator needs to put an additional protocol analyzer onto the link.
- Aggregator taps combine the traffic from both directions of full-duplex links and from multiple links and sends it to a single NIC on the monitoring tool. No packets are dropped as long as the aggregated traffic does not exceed the monitor port bandwidth.
- Active response taps permit monitoring tools to send response packets such as transmission control protocol (TCP) resets, Internet control message protocol (ICMP) messages, and access control list (ACL) changes into the tapped link. This feature can be used by an IDS to take action when certain types of intrusions are detected.
- Virtual taps (vTaps) provide insight and basic monitoring data about virtualized architecture traffic.
- Media conversion refers to taps that support different media types on their network and monitor ports. Many taps have pluggable SFP or XFP ports, enabling different media types to be accommodated simply by plugging in different transceiver types. Some taps even perform data rate conversion as well.
- Filter taps enable mirrored traffic to be restricted to particular protocols, source and destination IP addresses, VLANs, ports, and other criteria, making it easier to isolate or troubleshoot issues, and relieving monitoring tools from spending valuable processing cycles on pre-filtering traffic.
- Bypass switches create fail-safe access ports for in-line devices such as intrusion prevention systems and firewalls.



The wide range of tap devices available today enable appropriate monitoring access to be built into all parts of the network architecture, at the edges, distribution, LAN, and core. This monitoring access does not depend on SPAN ports for strategic information access, but in fact frees up the SPAN ports for tactical monitoring access when special needs arise. Permanent and ongoing monitoring can rely on a tap-based architecture for consistent, persistent, and secure monitoring access, immune to the vagaries of day-to-day network administration and management.

Conclusion

When it comes to running networks in the modern world, you can't miss the forest or the trees: both a long view and a close view are necessary to keep networks optimally running. Network operators need access to ALL traffic data in order to achieve the insight required to run today's stressed and complex networks. Although there are specific use-cases for SPAN ports, Taps are clearly superior for day-to-day monitoring activities.

Taps are placed between any two network devices and allow full-duplex traffic to pass through without affecting the data stream. Taps have two ports that connect directly to each monitoring device, enabling the manager to passively view all traffic without affecting the traffic.

Taps are an improvement over SPAN ports because taps send all traffic in the physical Layers 1 and 2, including errors, to the monitoring device. By using taps, network engineers gain visibility into the network to look for packet errors and bandwidth anomalies that were not visible using SPAN ports.

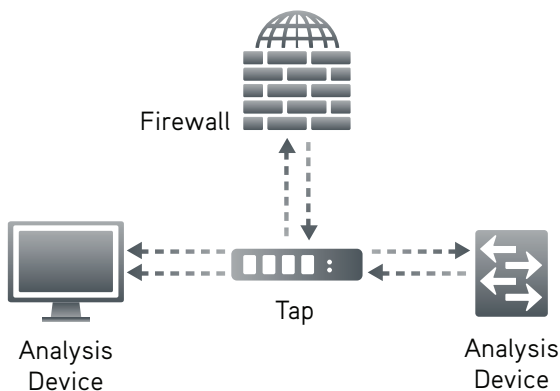
Although taps are placed inline just as monitoring devices can be placed inline, the passive design of the tap means that the device does not affect the physical stream if the tap fails. Industry-leading features include dual power supplies that increase reliability.

Ixia has many tap and visibility architecture solutions available to help support your network and business goals.

Taps are an improvement over SPAN ports because taps send all traffic in the physical Layers 1 and 2, including errors, to the monitoring device.

Tap vs. SPAN

Taps are passive splitting mechanisms located between a network device and the network. They transmit both the send and receive data streams simultaneously, so that all traffic data gets to a monitoring or security device in real time.




Pros

- No dropped packets
- Monitoring device receives all packets, including errors
- Provides full visibility
- Does not create unnecessary, duplicate packets
 - Does not create time stamp issues
 - Relieves SPAN port contention
- Plug & play: no configuration needed

Bottom Line

A tap is ideal when you need to see all the traffic, including physical-layer errors. A tap is required if network use is moderate to heavy.

Many enterprise switches can copy the activity of one or more ports (with limitations) through a SPAN port ("port mirroring"). A monitoring device can then be connected to the SPAN port to analyze network traffic.



Pros

- Low cost as it is provided with switch

Cons

- Burden is on switch's CPU to copy all data passing through ports
- Can change the timing of frame interaction, altering response times
- Switch prioritizes SPAN port data lower than regular port-to-port data
 - Filters out physical-layer errors
- Requires that you dedicate a switch port to monitoring (which should be used for switching)

Bottom Line

A SPAN port works on low-use networks, or if deep analysis is not needed (and isn't affected by dropped packets).

**Ixia Worldwide Headquarters**

26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)

1.877.367.4942

(Outside North America)

+1.818.871.1800
(Fax) 818.871.1805

www.ixiacom.com

Ixia European Headquarters

Ixia Technologies Europe Ltd
Clarion House, Norreys Drive
Maidenhead SL6 4FL
United Kingdom

Sales +44 1628 408750

(Fax) +44 1628 639916

Ixia Asia Pacific Headquarters

101 Thomson Road,
#29-04/05 United Square,
Singapore 307591

Sales +65.6332.0125

Fax +65.6332.0127