WHITE PAPER

To TAP or SPAN?

Introduction

TAP and SPAN technologies provide direct access to the actual packets navigating across networks, but which is the preferred methodology given today's infrastructures? If both options work, when should one technology be used over the other? To TAP or SPAN? That is the question.

Today's networks are getting larger and more complex, carrying unprecedented volumes of data at increasing speeds. For example, 400Gb and 1Tb Ethernet are in development, new initiatives such as Internet of Things (IoT) and cloud computing add new layers of complexity, and all make pervasive packet-level access problematic. At the same time, cyber threats are getting more and more sophisticated. As a result, network visibility is essential to monitor, manage and protect your network.

Accessing the data in motion down to the packet level is the first step to acquiring that visibility, as nothing else provides a similar level of depth and granularity. The two most common methods to extract this information are SPAN and TAP technologies. For a given situation, how do you decide which to use?

This paper dives into the details of the two technologies to clearly categorize where each should best be utilized within today's modern infrastructures.

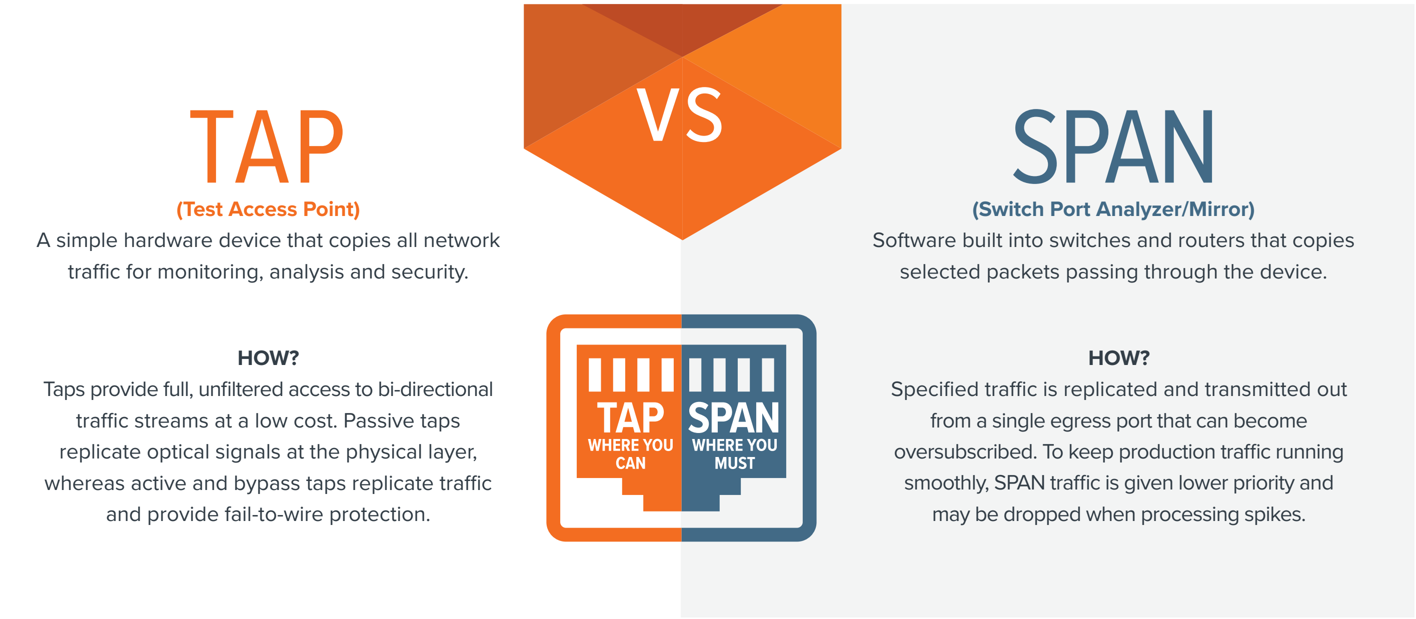

TAP

(Test Access Point)

A simple hardware device that copies all network traffic for monitoring, analysis and security.

HOW?

Taps provide full, unfiltered access to bi-directional traffic streams at a low cost. Passive taps replicate optical signals at the physical layer, whereas active and bypass taps replicate traffic and provide fail-to-wire protection.

SPAN

(Switch Port Analyzer/Mirror)

Software built into switches and routers that copies selected packets passing through the device.

HOW?

Specified traffic is replicated and transmitted out from a single egress port that can become oversubscribed. To keep production traffic running smoothly, SPAN traffic is given lower priority and may be dropped when processing spikes.

Figure 1: TAP vs SPAN

Basic TAP and SPAN Technologies

A network TAP (Test Access Point) is a simple device that connects directly to the cabling infrastructure. Instead of two switches or routers connecting directly to each other, the network TAP sits between the two devices and all data flows through the TAP. Using an internal splitter, the TAP creates a copy of the data for monitoring while the original data continues unimpeded through the network.

Data from a network is transmitted (Tx) from device A to be received (Rx) by device B. At the same time data can travel in the reverse direction where device B transmits data to device A. Most TAPs separately copy the transmit signals from A and B and send them to separate monitoring ports (TxA and TxB).

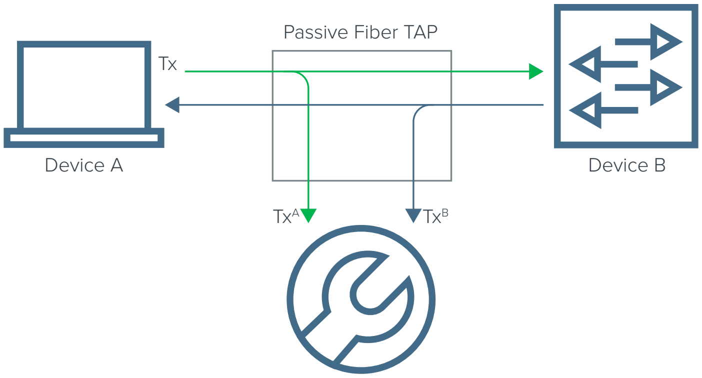

Figure 2: TAP flow diagram

This ensures every packet of any size will be copied. This technique also eliminates any chance of oversubscription. Once the data is TAPed, the duplicate copy can be used for any sort of monitoring, security or analytical use. Thus TAPs are a key component of any visibility system.

It should be noted that inserting a TAP into an existing network link requires a brief cable disconnect, so TAPs are typically installed during a maintenance window.

A SPAN port (sometimes called a mirror port) is a software feature built into a switch or router that creates a copy of selected packets passing through the device and sends them to a designated SPAN port. Using software, the administrator can easily configure or change what data is to be monitored.

Since the primary purpose of a switch or router is to forward production packets, SPAN data is given a lower priority on the device. The SPAN also uses a single egress port to aggregate multiple links, so it is easily oversubscribed.

Both of these situations can lead to dropped packets. SPANs were never intended for long-term monitoring. Rather, they work best for ad hoc monitoring of low volumes of data in locations where TAPs have not been installed. SPANs still represent the only means for accessing some types of data, such as data crossing port-to-port on the same switch.

A Historical Perspective

In the early days of networking, networks were slow and it was easy to capture packets. Networks ran 10Mb Ethernet using coaxial cable and shared media hubs to move data. Since every port on a hub contained all the data within the domain, a protocol analyzer could be placed on any data port on the hub to see everything coming or going from any other local device.

As LANs and WANs expanded and began to experience greater use, there were scaling challenges and serious security concerns. In 1995, the IEEE standardized Fast Ethernet, and the era of switches began. Switching technology forwards data to ports where specific addresses reside instead of having stations receive all data while looking for their own address. Switch scalability is much more robust, and use of the internet exploded around the same time they appeared.

However, since switches forward packets only to the appropriate port, easy access to all data was lost and network TAPs became the de facto standard to see all data crossing specific links. This led to a public outcry from network administrators. The major switch and router manufacturers (such as Cisco) ultimately responded by adding software within their devices to mirror data to SPAN monitoring ports. Since this feature was only intended for occasional troubleshooting, and its implementation impacted switch performance, vendors had to prioritize production data over mirrored packets.

For years TAP and SPAN technologies competed head-to-head with each other. Network professionals took sides based on a variety of networking and security perspectives that morphed into religious-like zeal. Some shops would exclusively deploy TAPs; while others would ban them outright within their infrastructure.

As networks evolved from 10Mb to 10Gb speeds, another shift took place. SPAN ports frequently became overwhelmed due to the sheer volume of data. Vendors clearly documented that SPAN ports were for low-volume data only and could negatively impact production traffic if improperly configured. Some high-speed switches were designed without SPAN capabilities while others openly discouraged its use. At 10Gb some vendors even implemented rate limiting on the SPAN output port as a default setting to reduce adverse effects. For this reason, TAPs re-emerged as the preferred access technology for contemporary networks.

Why are Network TAPs Preferred Over SPAN ports?

In modern networks a TAP tends to be preferred over a SPAN port for a variety of reasons. Passive TAPs are unpowered and run for years without failure. Their hardware design eliminates oversubscription while capturing every packet, in perfect order. Although SPAN ports work fine at low utilization levels, both transmitted (Tx) and received (Rx) data streams are forwarded to a single SPAN output port.

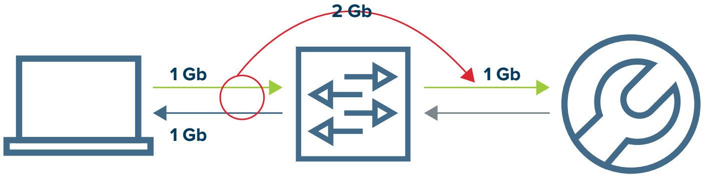

Figure 3: Why SPAN ports are easily oversubscribed

That means a single switch port running bidirectional traffic at 60 percent utilization would send data to the SPAN port at 120 percent. Since an Ethernet port can never go beyond 100 percent, at least 20 percent of the data would be immediately dropped. If a dozen or so similar ports are aggregated to a single SPAN port, only a fraction of the data would ever get through to the monitoring tools.

The following customer-generated graph depicts a real network at the point where ExtraHop's source was converted from SPAN to TAP. Upon making the change there is an immediate spike in data. Bear in mind the same data was being collected both before and after the conversion, but the original SPAN configuration was clearly oversubscribed and dropping packets.

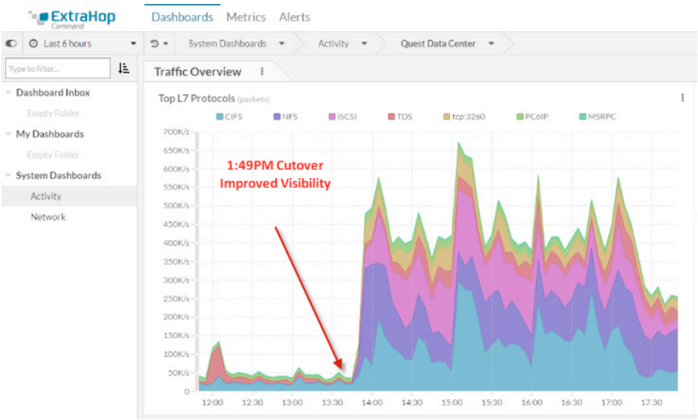

Figure 4: Conversion from SPAN to TAP

Other reasons explain why TAPs are preferred over SPAN ports. Because switches queue SPAN data as a low priority, in addition to dropping data, it is possible for packets to arrive at the monitoring tool out of order or to have latency variations so that SPAN data is delivered sooner or later than production data. This never happens with a TAP as every packet is forwarded in exact order at line rate.

To summarize, here are the top 10 reasons why network TAPs are preferred over SPAN ports:

1. TAPs create an exact copy of the bi-directional network traffic at full line rate, providing full fidelity for network monitoring, analytics and security.
2. Passive TAPs provide continuous access to traffic and require no user intervention or configuration once installed — a true set and forget solution.
3. SPAN ports are easily oversubscribed — resulting in dropped packets leading to unsatisfactory or inconsistent results for monitoring and security purposes.
4. SPAN traffic has the lowest priority when it comes to forwarding and may not achieve full line rate. In some situations, low priority can cause packet drop even on a SPAN port operating at single digit utilization.
5. The SPAN application can have a negative performance impact on the switch itself, sometimes affecting network traffic.
6. Because SPAN traffic is easily reconfigured, SPAN output can change from day to day or hour to hour — resulting in inconsistent reporting.
7. Legal regulations or corporate compliance sometime mandate that all traffic for a particular segment be monitored. This can only be guaranteed with a TAP.
8. Incorrectly configured SPAN ports have been known to impact network performance, or even cause network outages.
9. SPAN ports are limited in number compared to the number of ports that may require monitoring, and they consume ports that could otherwise be carrying production traffic.
10. TAPs do not care what protocol is carried in the traffic or if it is IPv4 or IPv6. All traffic is passed through a passive TAP, including packets with errors. Active TAPs typically block errors but forward everything else.

The bottom line is TAPs should be used wherever 100 percent visibility and traffic fidelity is required. Anytime traffic volumes are moderate to high, deploy network TAPs. As a best practice solution, install TAPs during the early design phase and pass the traffic directly to a Gigamon visibility node. Even if the traffic is not yet required for daily inspection, it will be available for ad hock troubleshooting or security inspection within seconds and without needing to involve change management.

When Should SPAN Ports Be Used?

As described above, network TAPs are preferred over SPAN ports in today's networking environments. However, there are still locations where a TAP is not practical. Consider using SPAN ports for the following exceptions:

- Limited ad hoc monitoring in locations with SPAN capabilities where a network TAP does not currently exist.
- Locations with limited light budgets where the split ratio of a TAP may consume too much light. (Another possibility here would be to use an Active TAP or more powerful optics capable of longer distances.)
- Production emergencies where there is no maintenance window in which to install a TAP.
- Remote locations with modest traffic that cannot justify a fulltime TAP on the link.
- Access to traffic that either stays within a switch or never reaches a physical link where the traffic can be TAPed.
- Low-cost troubleshooting alternative where links have low utilization.

In summary, both network TAPs and SPAN ports can provide valid access to data if properly positioned.

So, TAP where you can, and SPAN where you must.

Keep in mind that TAP and/or SPAN are just the first step in the process of getting pervasive visibility across your entire network infrastructure. Once traffic has been captured over either TAPs or SPAN ports, you can send it to a Gigamon Visibility Platform to be monitored, managed and secured. A range of GigaSMART® applications is available to optimize the traffic and deliver the relevant data to the specific tools you rely upon to improve the performance and efficiency of these tools.

About Gigamon

Only Gigamon enables you to see, control and secure what's happening across your entire physical, virtual and cloud network. We offer the market-leading network visibility and control solution to accelerate threat detection and incident response while optimizing infrastructure performance. Gigamon solutions enable our customers to embrace network architectures, application frameworks, and cloud deployments of today and tomorrow. For more information visit www.gigamon.com.

Reference: For more detail on how a TAP works see:
https://www.gigamon.com/sites/default/files/resources/whitepaper/wp-understanding-network-taps-the-first-step-to-visibility-3164.pdf.