WHITEPAPER

TAP vs SPAN

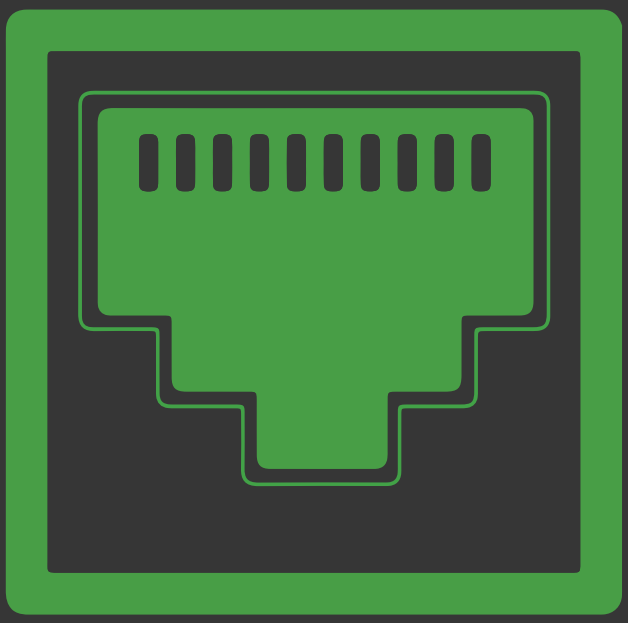

Best Practice Guide to Improving Network Visibility

GARLAND TECHNOLOGY

See every bit, byte, and packet®

TABLE OF CONTENTS

TAP VS SPAN: BEST PRACTICE GUIDE TO IMPROVING NETWORK VISIBILITY

3 | **Introduction: How to Navigate Network Visualization**

4 | **Brief History on Network Access**

5 | **2 Methods for Network Access**

5 | Network TAP [Test Access Point]

6 | SPAN [Switched Port Analyzer]

7 | **Today's Data Access Requirements**

8 | **4 Fundamental Considerations for Network Visualization Requirements**

9 | **SPAN Port Considerations**

10 | **SPAN Reality Check**

12 | **When Is SPAN Port Methodology Ok to Use?**

13 | **Putting TAP vs SPAN to the Test**

14 | **TAP to SPAN Comparison**

15 | **Access Best Practices Conclusion**

INTRODUCTION

HOW TO NAVIGATE NETWORK VISUALIZATION

This whitepaper is an in-depth look into network visualization access.

Today's requirements for security, monitoring, management, compliance, deep or historic captures and auditing of our networks require full and real time access to the packets that flow through our network. Today network, security and management personnel must have full access and using test access point (TAP) technology is the only viable and reliable technology for that job.

We will cover the value TAP access will provide, some fundamental considerations, requirements, best practices and highlight some disadvantages about SPAN or monitor access through switches.

The goal of network, security, compliance and application managers require full visualization of the network and the packets therein. Real visualization is everything. If you cannot see an issue, like an attack, misusage, inefficiency, etc., how are you going to understand it and resolve it?

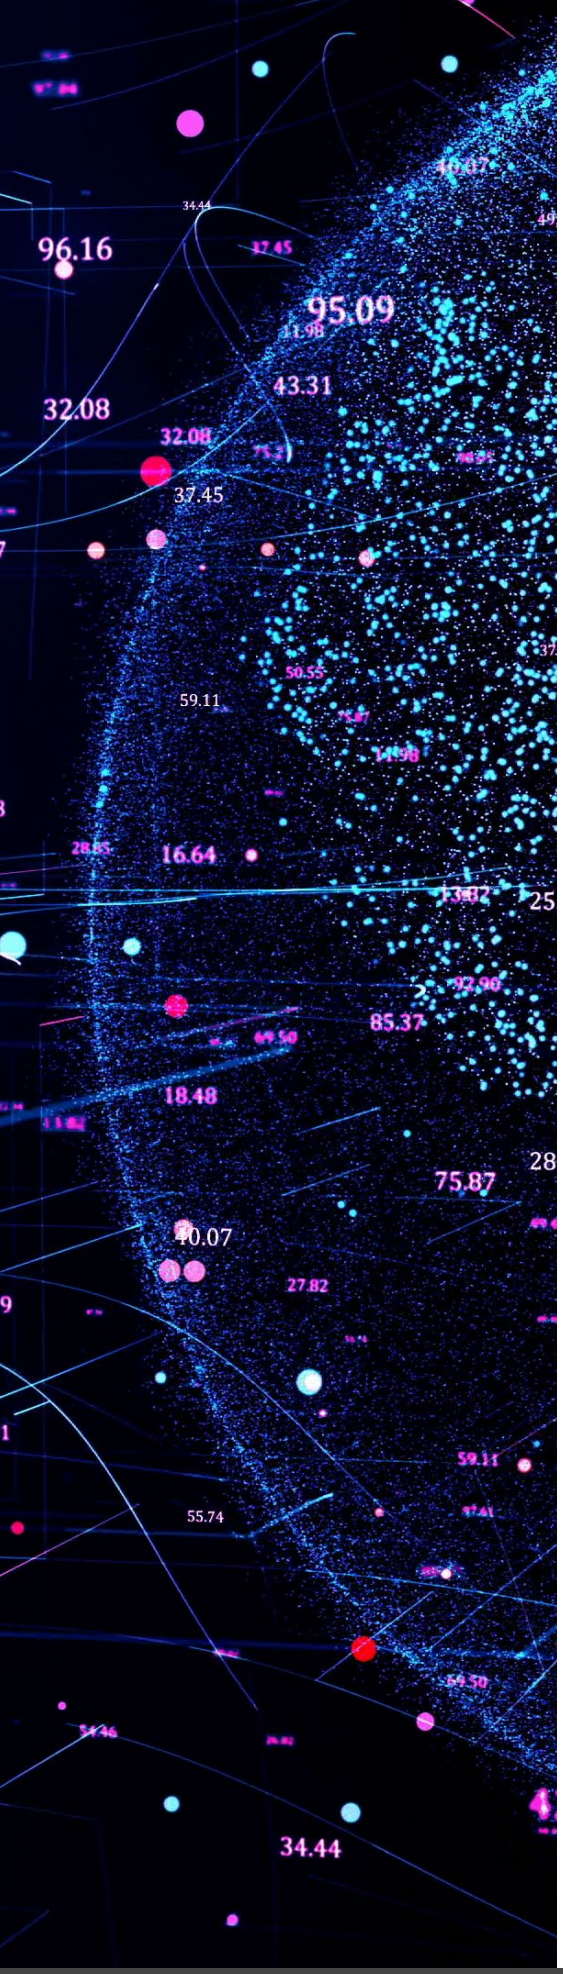

BRIEF HISTORY ON NETWORK ACCESS

Until the early 1990's, using a network TAP (test access point) from a switch patch panel was the only way to monitor a communications link. Most links were WAN so an adaptor like the V.35 adaptor from Network General or an access balun for a LAN was the only way to access a network. Most LAN analyzers had to join the network to really monitor it.

As switches and routers developed, there came a technology we call SPAN/Mirror ports — and monitoring was off and running. Analyzers and monitors no longer had to be connected to the network directly; engineers would use the SPAN (mirror) port and direct packets from their switch or router to the test device for analysis.

SPAN was a great way to effortlessly and non-intrusively acquire data for analysis. By definition, a SPAN Port usually indicates the ability to copy traffic from any or all data ports to a single unused port. The SPAN or Monitor port was originally a quality assurance test point for their manufacturers and became a visualization access point as an afterthought.

SPAN usage proliferated as an easy method for quick packet access. Though as networks have evolved and speeds have increased exponentially, architects now have to consider the limitations of SPAN, and use a mixture of network TAPs and SPAN depending on the scenario.

> **"Research found that 83 percent of current network visibility fabrics use TAPs for at least half of the fabric access layer."**
>
> **-EMA [Enterprise Management Associates]**

2 METHODS FOR NETWORK ACCESS

Network TAP [Test Access Point]

A network TAP is a purpose-built hardware device that allows you to access and monitor your network traffic by copying packets without impacting or compromising network integrity. Network TAPs sit in a network segment, between two appliances (router, switch or firewall), and allows you to access and monitor the network traffic. The TAP allows network traffic to flow between its network ports without interruption, creating an exact copy of both sides of the traffic flow, continuously, 24/7, 365. The raw packet copies are then used for monitoring and security analysis.

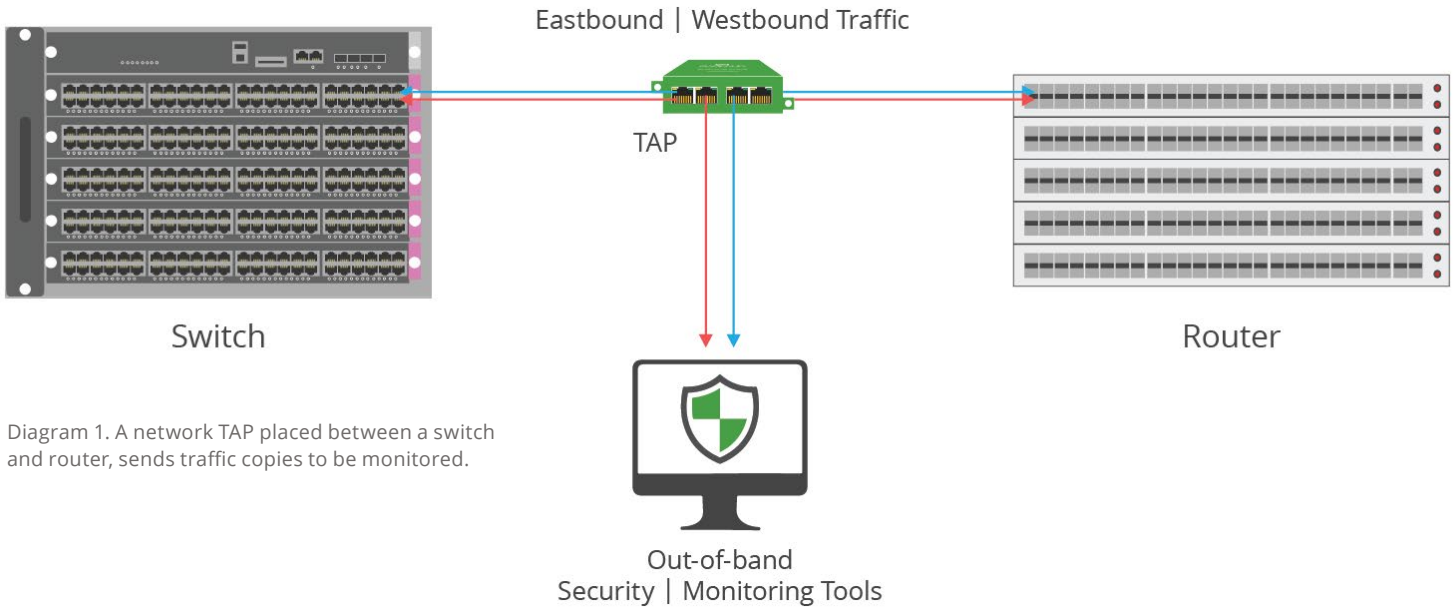

Diagram 1. A network TAP placed between a switch and router, sends traffic copies to be monitored.

A network TAP provides strategic, persistent monitoring capabilities. Installing a TAP during deployment means you have a permanent method of access to network traffic.

- Ensure 100% full duplex copies of network traffic without altering the data.
- Support 10M, 100M, 1G, 10G, 40G, 100G, and 400G.
- Are scalable and can either provide a single copy, multiple copies (regeneration), or consolidate traffic (aggregation) to maximize the production of your monitoring tools.

> **"EMA recommends that enterprises use TAPs as much as possible in the access layer to avoid network performance impacts and assure packet fidelity."**
>
> **-EMA [Enterprise Management Associates]**

SPAN [Switched Port Analyzer]

Port Mirroring also known as SPAN (Switch Port Analyzer), are designated ports on a network appliance (switch), that are programmed to send a copy of network packets seen on one port (or an entire VLAN) to another port, where the packets can be analyzed.

Many switches have a limit on the number of SPAN monitoring ports that you can configure. This limit is often a maximum of two monitoring ports per switch.

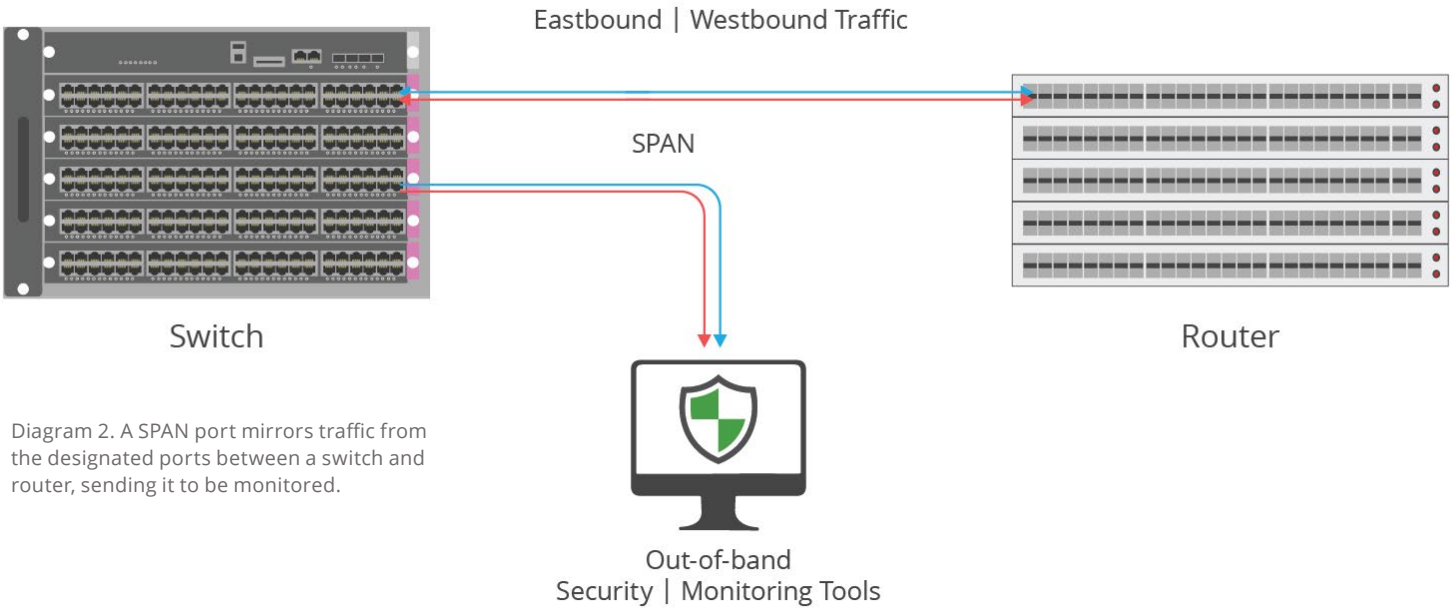

Diagram 2. A SPAN port mirrors traffic from the designated ports between a switch and router, sending it to be monitored.

Port mirroring best practices vary by switch vendor, as many architectures use non blocking methods that drop overages if you overrun a port mirror, depending on the switch you use, there can be an adverse effect on traffic or switch performance.[1]

- SPAN are programmed ports on a switch, that provide access to packets for monitoring.
- SPAN sessions do not interfere with the normal operation of the switch.
- Low priority processing — the switch will drop SPAN packets if heavily utilized or oversubscribed.

> **"SPANs can add overhead on a network device, and that SPAN port will often drop mirrored packets if the device gets too busy. Therefore, TAPs are a better option."**
>
> **-EMA [Enterprise Management Associates]**

TODAY'S DATA ACCESS REQUIREMENTS

To add more complexity and challenges to SPAN port as a data access technology:

1. We have entered a much higher utilization environment with many times more frames in the network.

2. We have moved from 10 Mbps to 40 and 100 Gbps Full Duplex.

3. We have entered into the era of data security, deep packet capture, legal and policy compliance, network auditing and Lawful Intercept approved by Commission on Accreditation for Law Enforcement Agencies (CALEA) which requires that we must monitor all of the data and not just "sample" the data, with the exception of very focused monitoring technologies (i.e. application performance monitoring).

These demands will continue to grow since we have become a very digitally focused society and are all connected via the Internet of Things (IoT). With the heavy use of VoIP and digital video we have revenue generating data that is connection oriented and sensitive to bandwidth, loss and delay.

The older methods need reviewing and the added complexity requires that we change some old habits to allow for real 100% Full Duplex real time access to the critical data. Being able to provide real access is not only important for Data Compliance Audits and Lawful Intercept events, it is the law.

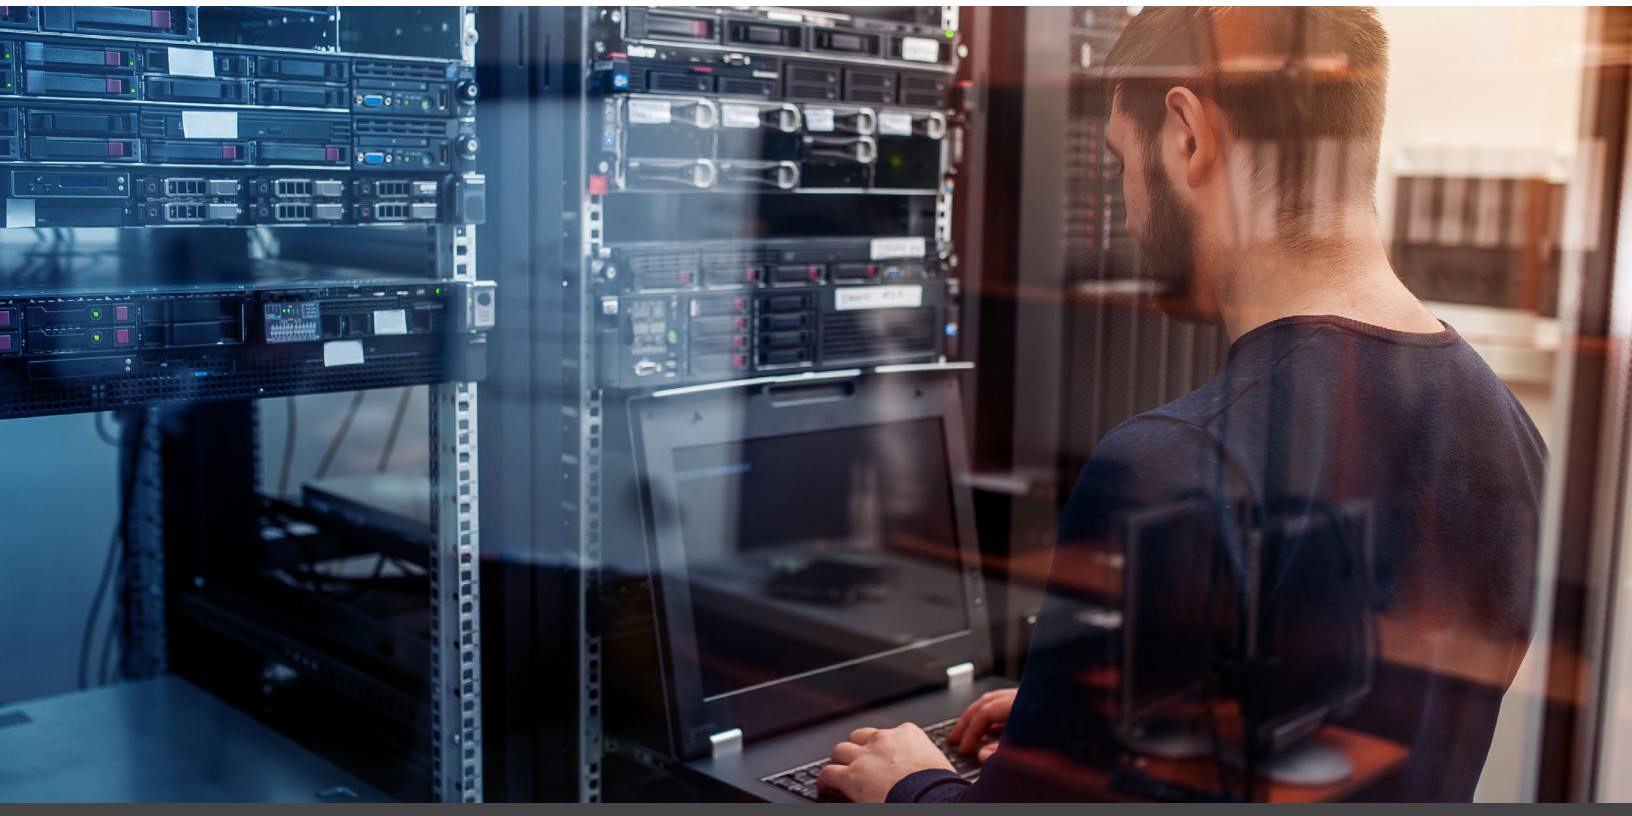

4 FUNDAMENTAL CONSIDERATIONS FOR NETWORK VISUALIZATION REQUIREMENTS

During any deployment there are many different factors that go into your connectivity strategy and design. Here are some considerations that may seem like common sense but should always be a part of your mental checklist.

1 - Frame Changes

Any active device that touches a frame, has changed the frame timing, even if nothing more than changing its absolute timing reference to the network needs to be noted.

It is essential to keep all changes by a device, linear. If the frame offset was 10ms then all frames should have the same offset, if not, the device is interfering with the real time analysis capability of that access point.

SPAN access is a great example of variable offset and the impossibility of doing authentic time based analysis from a SPAN/Monitor port. A network TAP with a tested algorithm handles the send and receive integration with consistent timing for the best visualization.

2 - Passing All Packets

A network TAP is the only device that will pass every bit, byte, and packet, including the interframe gap, bad, large, small and other error packets. Even if one uses a higher technology filtering device, network TAPs as your media access guarantees you pass all packets.

There is debate about the viability of passing bad packets for capture and post capture analysis. A leading point of view is counting the bad packets/types are acceptable and many times a requirement for baselining analysis.

3 - Access Analysis

Before deploying access technology, a good best practice to always consider:

- Test more than one device to make sure you are getting what you really need for your tools and that you and your company can really use the device and the data it provides.
- Be sure to test the network before and after the access device to compare and get a real baseline of the access device effects on the frames.
- Always buy quality. TAPs are relatively inexpensive compared to the tools and infrastructure. Look for a trusted vendor with customers to back it up, not online overseas knockoffs. Before making a decision on which TAP to go with, look into the testing, where they are manufactured, hardware warranties, optical transceivers, Mean Time Between Failure (MTBF) rates and first time pass rate (FTPR) and durability.

4 - Lawful Capture

One major and important consideration about access technology is legal and lawful capture. Please do not forget that any access device can be called into question in civil and criminal cases. When using the data captured as the evidence in employee misuse or for CALEA/Lawful capture situations a network TAP is your best ally. It presents the evidence with no chance of changing anything, providing a solid time reference. We call this forensically sound data and is a must for court evidence. Another plus to consider in our security conscious world - a network TAP cannot be hacked, so any evidence gathered is usually foolproof.

SPAN PORT CONSIDERATIONS

Cisco warns that "the switch treats SPAN data with a lower priority than regular port-to-port data." In other words, if any resource under load must choose between passing normal traffic and SPAN data, the SPAN loses and the mirrored frames are arbitrarily discarded. This consideration applies to preserving network traffic in any situation. For instance, when transporting remote SPAN (RSPAN) traffic through an Inter Switch Link (ISL), which shares the ISL bandwidth with regular network traffic, the network traffic takes priority. If there is not enough capacity for the remote SPAN traffic, the switch drops it.

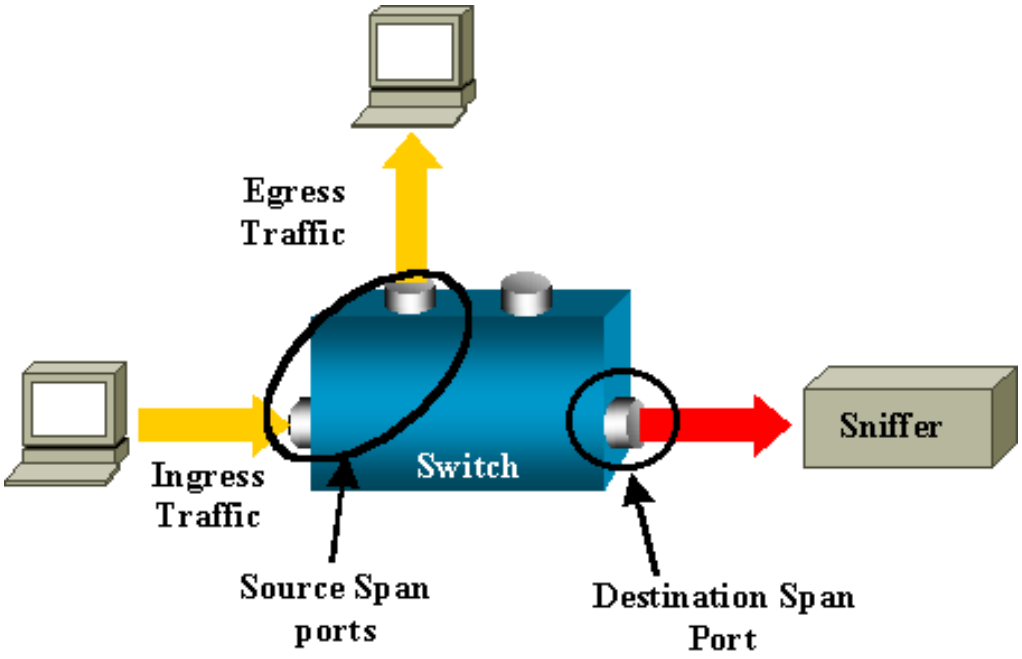

Diagram 3: Cisco reference diagram for their Catalyst switch SPAN port.

Knowing that the SPAN port arbitrarily drops traffic under specific load conditions, what strategy should users adopt so as not to miss frames? According to Cisco,

> **"the best strategy is to make decisions based on the traffic levels of the configuration and when in doubt to use the SPAN port only for relatively low-throughput situations."[2]**

Also, consider that a switch's SPAN access is not fault tolerant and can be a major fault or failure point for your monitoring and management vision. A network TAP is not a failure point.

SPAN REALITY CHECK

Is a SPAN port a passive technology?

No. Some may call a SPAN port a passive data access solution, but passive means "having no effect" on the packet. Spanning (mirroring) does have measurable effect on the data packets themselves as well as all the packet timing is affected.

Isn't SPAN always available?

No. SPAN retention is a real challenge. Many switches have minimal ports capable of spanning. If a span port is committed for permanent monitoring duties, sometimes it needs to be reassigned for troubleshooting on a VLAN or other aspect of the traffic, denying traffic to your monitoring tools.

Do SPAN ports drop packets?

Yes. Dropped packets or packet loss can occur when packets of data traveling across a network fail to reach their destination. Packet loss can be caused by network congestion, hardware capacity and bottlenecks, errors in data transmission, or overloaded devices.

Packet loss is measured as a percentage of packets lost compared to packets sent. The Transmission Control Protocol (TCP) is the standard that defines how to establish and maintain a network conversation through which application programs can exchange data. TCP detects packet loss and performs retransmissions to ensure reliable messaging. Monitoring packet loss is a fundamental best practice for network performance.

Network congestion, hardware capacity, bottlenecks and overloaded devices can all be symptoms of an over reliance on SPAN ports for network monitoring. SPAN is commonly known to drop packets if a port is oversubscribed. The issue arises when the total amount of traffic aggregated from the source ports exceeds the physical limitations of the destination port, it results in some dropped packets on the destination port.

This does not cause any switch performance degradation, disruption or traffic flow on the source ports. The only affected port is the destination port, and it drops packets on a first in first out (FIFO) basis once the egress buffer limit is exceeded.

- Monitoring tools may miss packets due to SPAN port oversubscription.
- SPAN will not pass corrupt packets or errors (bad packets), these are dropped.

Do SPAN ports alter or duplicate packets?

Yes. Duplicate packets are commonly seen when packets are traversing multiple switches or routers and are copying traffic to the SPAN/mirror port.

You can SPAN packets in or out of a switch port, but typically most applications require a copy of both sides. SPANs are known to result in trace files with duplicated packets when the SPAN port is set up to capture both ingress and egress traffic flows. This common problem when both the ingress and egress ports are spanned, end up sending duplicate packets to the monitoring tool, which becomes a whack-a-mole type headache.

- SPAN can change the timing of the frame interactions, altering response times
- The timestamps are can read different but the packet contents are the same
- Can duplicate packets if multiple VLANs are used

Is a SPAN port scalable technology?

No. When we had only 10Mbps links and a robust switch, one could almost guarantee they could see every packet going through the switch, except for bad frames. With 10Mbps fully loaded at around 50% to 60% of the maximum bandwidth, the switch backplane could easily replicate every good frame. Even with 100Mbps one could be somewhat successful at acquiring all the good frames for analysis and monitoring and if a frame or two here and there were lost, it was no big problem.

This has all changed with 1, 10, 40 and 100 Gigabit technologies starting with the fact that maximum bandwidth is now twice the base bandwidth – so a Full Duplex (FDX) Gigabit link is now 2 Gigabits of data and a 10 Gigabit FDX link is now 20 Gigabits of potential data flows.

No switch or router can handle replicating/mirroring all of that data, plus handling its primary job of switching and routing. It is difficult if not impossible to pass all frames (good and bad ones), including FDX traffic at full time rate with the interframe gap, in real time at non-blocking speeds. All of this, on say 16 ports, is a whole lot of data to go through one port. Furthermore, to this FDX need we must also consider the VLAN complexity and finding the origin of a problem once the frames have been analyzed and a problem detected.

Is RSPAN (remote SPAN) a viable access technology?

No. Especially if the packets are passed over the WAN as it will gobble up all your bandwidth passing frames back to your local switch that have already passed through the network. RSPAN is not considered an acceptable nor viable visualization access method.

WHEN IS SPAN PORT METHODOLOGY OK TO USE?

Many monitoring products can and do successfully use SPAN as an access technology. These monitoring products are looking for low bandwidth application layer events like “conversation or connection analysis,” “application flows,” and applications where real time and knowing real delta times are not important.

These monitoring requirements utilize a small amount of bandwidth and grooming does not affect the quality of the reports and statistics. The reason for their success is that they keep within the parameters and capability of the SPAN ports, and they do not need every frame for their successful reporting and analysis. In other words, SPAN port is a usable technology if used correctly and the companies that use mirroring or SPAN are using it in a well-managed and tested methodology.

SPAN could also be used in a remote location that doesn’t justify a permanent deployment, offering temporary access for troubleshooting. After all, SPAN ports were not intended for long-term use.

USE CASE

Network TAPs can enhance SPAN deployment

Yes, network TAPs provide 100% visibility to your network monitoring and security tools, but as SPAN is still being used, there are many TAP use cases that can enhance your current SPAN deployment. There are many situations when there are not enough SPAN/mirrored ports available on a router or switch to allow access to all of the monitoring tools that need to see the traffic of the link. Introducing a Regeneration/SPAN Mode TAP provides a way to distribute a link’s traffic to up to multiple network tools.

SPAN Regeneration: Garland’s network TAPs have SPAN or regeneration mode, which allows you to take one SPAN link and copy the same traffic to multiple tools (1:3, 1:5).

SPAN Aggregation: Another good best practice to follow if SPAN port usage is required, Aggregator TAPs allow you to take those SPAN and consolidate them into just one or two links (2:1, 8:1, 22:1) . This optimizes and reduces network complexity.

SPAN Data Diodes: Garland’s Data Diode TAPs are designed to secure SPAN links, ensuring no bidirectional traffic is sent to monitoring tools.

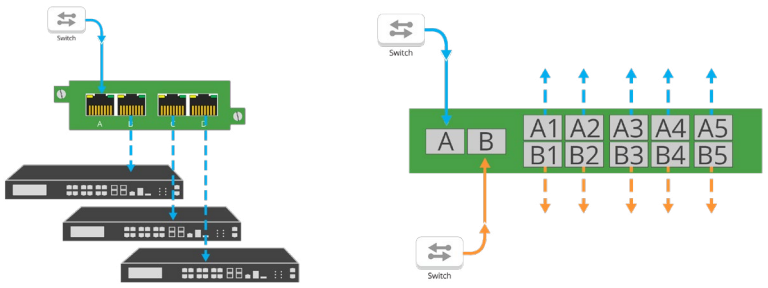

Diagram 4: SPAN regeneration mode for network TAPs

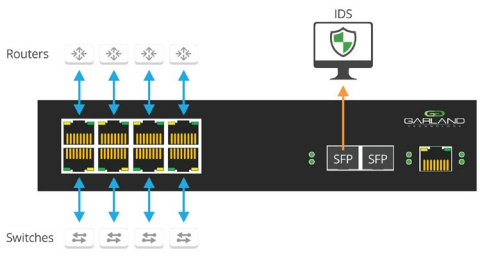

Diagram 5: SPAN aggregation with network TAPs

PUTTING TAP VS SPAN TO THE TEST

In a test conducted by Packet Pioneer, Chris Greer set out to see the difference between a data stream captured on a network TAP versus a SPAN port.

The test connected two PCs to a basic Cisco Catalyst Switch at 100Mbps. A throughput test using iPerf was configured and run between the two machines. On one of the PCs, he placed a 100Mbps TAP, and a hardware analyzer to capture. Lastly, he configured a SPAN on the switch to forward all traffic to and from this port to another hardware analyzer.

The throughput test finished with a result of 93.1Mbps sustained for 10 seconds between the two PCs.

TAP vs SPAN	Packets captured	Delta Time at TCP Setup
TAP Capture Results	133,126	243uSec
SPAN Capture Results	125,221	221 uSec

The SPAN data capture showed almost 8,000 packets missing from the trace. This represents almost 8% of the total packets that were captured by the analyzer from the network TAP. We should also point out that this was on a 100Mbps interface, not a Gigabit interface, and there were no errored frames. The switch bus was not in a near overloaded state.

Also, the difference in the timing between the TCP SYN and SYN ACK in the two traces shows us that the switch is not treating both the SPAN and Destination ports the same. In fact, it was forwarding traffic to the SPAN port faster than the true destination. While the difference is only 21 uSec, it shows that the switch is affected when SPAN is enabled. It is not as seamless as it would appear, and this delay was under no load test. With the switch loaded with traffic, the losses and timing will show greater differential and also dropped packets.

Considering the results of their test, Chris Greer, a network analyst at Packet Pioneer, said, “I am now a full believer in using a real [network] TAP whenever possible, especially when timing and total view of the data is important!”

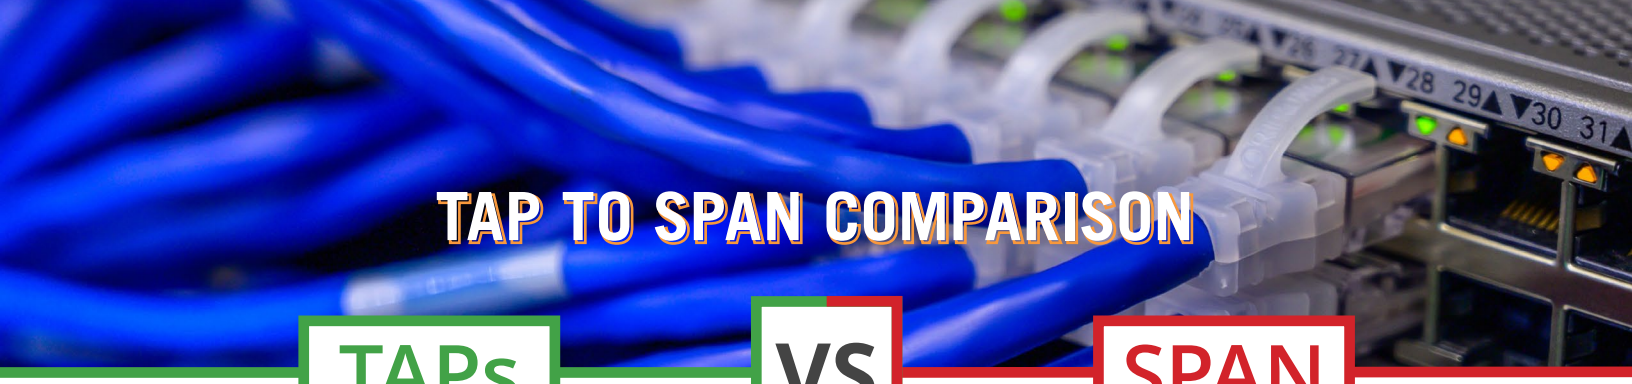

TAP TO SPAN COMPARISON

TAPs VS SPAN

- Provides 100% full duplex copies of network traffic, ensuring no dropped packets for monitoring
- TAPs are passive or failsafe, ensuring no single point of failure (SPOF)
- TAPs do not alter the time relationships of frames, spacing and response times, which is especially important with VoIP and Triple Play analysis including FDX analysis
- TAPs do not introduce any additional jitter or distortion which is important in VoIP / Video analysis
- VLAN tags are not passed through the SPAN port so this can lead to false issues detected and difficulty in finding VLAN issues
- TAPs pass all traffic: IPv4 or IPv6, error packets, short or large frames, bad CRC frames, interframe gap is not dropped nor altered, packets are not dropped regardless of the bandwidth
- TAPs are fault tolerant
- TAPs are secure, do not have an IP address or MAC address, and cannot be hacked
- CALEA (Commission on Accreditation for Law Enforcement Agencies) approved for lawful intercept, providing forensically sound data, ensuring 100% accurate data captured with time reference
- Simple, plug and play. TAPs typically have little setup or command line issues, data is assured and saves users setup time
- TAPs are timeless - They never need to download or be upgraded, they do not have access to anything except the LAN they are monitoring
- Scaleable for traffic optimization, can regenerate one link to multiple or aggregate multiple links down to one

- Provides access to packets for monitoring
- SPAN traffic is the lowest priority on the switch
- Switch will drop SPAN packets if heavily utilized or oversubscribed
- Corrupt packets and low layer errors can be dropped out by a switch/SPAN
- SPAN can duplicate packets if multiple VLANs are used
- Using SPAN/Mirror ports can change the timing of the frame interactions, altering response times
- SPAN is not legally compliant for lawful intercept cases
- SPAN ports can easily be incorrectly configured impacting network performance, and even cause outages
- Bidirectional traffic opens back flow of traffic into the network, making switch susceptible to hacking
- SPAN retention. SPAN ports are limited in number compared to how many are needed for monitoring and can be costly as available ports come at a premium
- Admin/programming costs for SPAN can get progressively more time intensive and costly

ACCESS BEST PRACTICES CONCLUSION

Spanning (mirroring) technology is still viable for some limited situations. But as IT teams are now migrating from 10Mb to Gigabit, to 40 and 100 Gigabit networks, along with the demands of seeing all frames for data security and policy compliance, deep packet capture and Lawful Intercept, they must consider the use of real access TAP technology to fulfill the demands of today's complex analysis and monitoring technologies.

Creating a foundation of visibility is key for network management. Once deployed, a network TAP allows you to access that point in your network at any time. Many organizations have adopted the best practice of tapping all critical links for easy access during troubleshooting or inevitable security breaches.

Setting Yourself Up for Visualization Success

Looking to add TAP visibility to your deployment, but not sure where to start? Join us for a brief network Design-IT consultation or demo. No obligation - it's what we love to do.

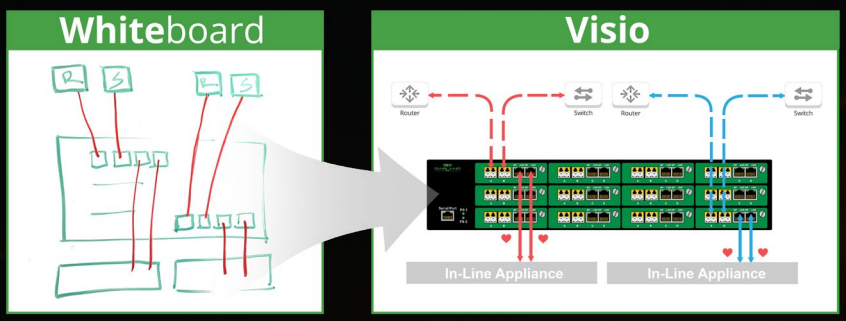

For more info, please visit: https://www.garlandtechnology.com/design-it

Contributions from Tim O'Neill OLDCOMMGUY™ and Chris Greer

1 - Port Mirroring and SPAN (Riverbed) | https://support.riverbed.com/bin/support/static/tku8mot0iaoa67qben06tukj4h/html/afhi95mcqoa01gcejeafeknach/sc_may2016_dg_html/index.html#page/NPM%2520Deployment%2520Guide%2Fpacket_collection.06.3.html%23

2 - Cisco | https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/10570-41.html

3 - EMA [Enterprise Management Associates] | https://www.garlandtechnology.com/wp-ema-security-visibility