

Understanding Network TAPs

The First Step to Visibility



Introduction

A network TAP is a simple device that connects directly to the cabling infrastructure to split or copy packets for use in analysis, security, or general network management. Although the term “tap” predates the networking industry by decades, the IT industry has generally adopted the term to mean test access point. Thus, TAP is considered an acronym.

This paper covers

- Background
- TAP vs. SPAN
- Network TAPs Overview
- Types of TAPs and How They Work
 - Passive TAPs
 - › Optical Splitter Types
 - › Specialized 40Gb BiDi TAP
 - › Split Ratios
 - › Optical Speeds and Types
 - › Power Budgets and Light Loss
 - Active TAPs
 - Other TAPs
 - › Bypass Technology
 - › Aggregation TAP vs. Aggregation Node
 - › Standalone vs. Embedded TAPs
- 10GBASE-T
- TAP Best Practices
- Summary

Background

The heart of networking relies on a common set of communication protocols. Simple interconnections work regardless of the information passed in the payload. The brilliant simplicity of the system has evolved into a worldwide internet, enabling a huge array of applications, including everything from online banking to international telephone calls.

However, this same simplicity also creates complications.

Since every package looks the same from the outside, how do you know if a packet (or frame) contains the correct information? Did a specific banking transaction or an online sale close for the correct amount? Was a healthcare record properly filed according to compliance and tracking regulations? Was a client properly authorized to access a database?

Proper analysis of any of the above situations requires visibility into the actual packets running on the wire. All information passed resides in the packets. Of the hundreds of analysis tools available, most rely on packet information. As such, network monitoring and IT security have become key components of every industry.

TAP vs. SPAN

There are two common methods to extract traffic directly from the system: TAPs and SPANs. A network TAP is a hardware component that connects into the cabling infrastructure to copy packets for monitoring purposes. A SPAN (switch port analyzer) is a software function of a switch or router that duplicates traffic from incoming or outgoing ports and forwards the copied traffic to a special SPAN port (sometimes called a mirror port). In general, network TAPs are preferred over SPAN ports for the following reasons:

- SPAN ports are easily oversubscribed and have the lowest priority when it comes to forwarding, which results in dropped packets
- The SPAN application is processor-intensive and can have a negative performance impact on the switch itself, possibly affecting network traffic
- Because SPAN traffic is easily reconfigured, SPAN output can change from day to day, resulting in inconsistent reporting

However, there are some situations where inserting a TAP is not practical. For example, traffic could be running on a physical infrastructure outside your direct control, or maintenance windows may not allow for timely TAP deployments. Perhaps a remote location may not be able to justify a permanent TAP, but has SPAN access for occasional troubleshooting needs since a SPAN can be added without bringing down a link. There are also speed or interface types for which a SPAN may be the only option. So a combination of TAP and SPAN usage is fairly common. To quote a networking axiom: TAP where you can, SPAN where you can't.

Network TAPs Overview

Since a network TAP provides the most effective means to copy actual traffic running across a system, the remainder of this paper is dedicated to TAP types, usage, and functionality. It should be noted that TAPs are available for a wide variety of network speeds and cable types. Instead of two switches or routers connecting directly to each other, the network TAP sits between the two endpoint devices, connected directly to each of them. Then traffic is seen and copied, providing visibility into the networked traffic. See Figure 1.

TAPs are straightforward devices that run for years and are generally placed in secured locations. Once the traffic is tapped, the copy can be used for any sort of monitoring, security, or analytical use. Thus, TAPs are a key component of any visibility system.




Figure 1. Direct cabling vs. TAP cabling.

Types of Network TAPs and How They Work

There are many different types of TAPs. The two primary types of network TAPs are:

- Passive TAPs
- Active TAPs

Passive TAPs

A passive TAP requires no power of its own and does not actively interact with other components of the network. It uses an optical splitter to create a copy of the signal and is sometimes referred to as a photonic TAP. Most passive TAPs have no moving parts, are highly reliable, and do not require configuration.

A Typical TAP Installation Involves:

1. Placing the TAP on a shelf or in a rack
2. Connecting the cables
3. Verifying everything is working

It is really that simple. If the TAP fails to work, there is probably a cabling issue or a bad connection. Do be aware that installing or replacing a TAP in an existing environment does bring down the link while the cables are reconnected. So TAP installations are typically scheduled during pre-defined maintenance windows or during the network architecture design phase, prior to running live traffic.

Optical fiber sends light from a transceiver through a thin glass cable to a receiver on the other end. Instead of connecting directly to each other, each of the two endpoint nodes (switches, routers, databases, and so on) are connected to network ports on the TAP. These special ports are physically wired in pairs

such that traffic continually passes through them. In addition to the network ports are monitoring ports. The monitoring ports send out complete copies of the traffic seen, as shown in Figure 2.




Figure 2. TAP diagram showing logical flow.

Unlike network ports with both TX (transmit) and RX (receive) traffic, monitoring ports are unidirectional and only send traffic. They have no ability to receive traffic and never pass traffic back into the system. You will notice there are two monitoring ports in the diagram. Since each network port both sends and receives traffic, a 10Gb link could have 20Gb running across it. If all this traffic were put into one monitor cable, the link could quickly be oversubscribed. By running two separate monitor links, oversubscription is eliminated. The monitored traffic is thus separated into two transmit (TX-only) signals, one copy from endpoint A (Switch X) and one copy from endpoint B (Switch Y).

As depicted in Figure 2, a passive network optical TAP leverages a simple internal design. The external connectors lead to sets of glass fibers, splitters, and more glass fibers leading back to the external connectors. Each splitter has one fiber coming in and two going out.

Optical Splitter Types

Internal to the TAP, between the network port pairs, lies a small piece of hardware called an optical splitter. The splitter does exactly as the name implies; it splits an optical stream into two paths. A portion of the light continues onto its original destination; the second path is directed to a monitor port.

A traditional method to split the light is to fuse (or melt) two cables together such that a portion of the light is funneled off to the secondary stream. This technology is called fused biconical taper (FBT) and is shown in Figure 3. The concept is similar to when a river hits a fork. A portion of the water continues in the original direction while the rest takes an alternative path. Both forks of the river continue to flow downstream. Like water, light is also directional. As a result, the FBT tends to pass the traffic one way. FBTs tend to be low cost and work well for lower-speed cable plants.




Figure 3. Fused biconical taper (FBT).

A second splitter type uses thin film technology. The concept here is similar to shining a flashlight through a clear glass window. Although the majority of the light continues through the window, a portion of the light is reflected back as it hits the glass. If angled properly, a semipermeable membrane cutting across the fiber will copy a portion of the optical signal to the monitor port, as shown in Figure 4. Thin film's reflective technology tends to have a lower loss rate when working with high-speed links, such as 100Gb, where hot spots tend to occur due to uneven light distribution across the fiber. The FBT slice sees only the portion of the light where it is fused. Thin film is more evenly distributed because it sees the reflected light across the entire diameter of the cable.




Figure 4. Thin film splitter technology.

Specialized 40Gb Bidi Tap

Thin film is also preferred for TAP bidirectional links, such as 40Gb Cisco BiDi, because multiple wavelengths can be reflected simultaneously to break out each lambda (or wavelength) of light. Cisco BiDi leverages 40Gb technology using standard LC-based cabling to minimize the overall costs of deploying 40Gb links. This is a growing trend, especially with regard to Cisco leaf/spine configurations. See Figure 5 for an example of how reflective technology is used within this highly specialized passive TAP.




Figure 5. Thin film used in a bidirectional implementation.

Split Ratios

Regardless of the method used, the passive splitter physically diverts a portion of the light from its original source. The proportional share of light for each path is known as the split ratio. The split ratio is written as a combination of two percentages. The first number is designated as the network percentage, the second number is the monitor percentage. They always add up to 100 percent. For example, a common split ratio for traditional 1Gb short-range links is 70/30, where 70 percent of the light continues to the network and 30 percent is allocated to the monitor port.

The concept is to allocate more light to the network to reduce the risk of dropping network traffic. Speeds such as 10Gb, 40Gb, and 100Gb have different technical requirements and tend to use more of an even split ratio such as 50/50 or 60/40. The most common split ratio deployed in networks today tends to be 50/50, provided the proper light levels are available. When light levels are marginal, the safe option is to move to better optics offering higher safety margins.

Gigamon tests every TAP manufactured and provides the actual tested loss values with each Gigamon-branded TAP shipped. In addition, Gigamon data sheets for TAPs describe the maximum acceptable network and monitor loss values (including connections) for each split ratio as listed in Figure 6.

Multimode Passive TAPs

Split ratio	50/50	60/40	70/30
Max network loss	3.9dB	3.15dB	2.2dB
Max monitor loss	3.9dB	5.15dB	6.2dB

Singlemode Passive TAPs

Split ratio	50/50	60/40	70/30
Max network loss	3.7dB	3.05dB	2.0dB
Max monitor loss	3.7dB	4.9dB	6.1dB

Figure 6. Gigamon published charts for maximum loss (which includes connections) for various split ratios.

Optical Speeds And Types

Fiber TAPs are available for a wide variety of speeds and cable types. Most networks rely on IEEE 802.x standard-based optical cables. Speed is shown in gigabits per second, or Gbps. However, it is commonly shortened to Gb or G. The most common speeds in use today are 1Gb, 10Gb, and 40Gb, but the trend is quickly moving toward higher-speed networks of 100Gb. Speeds of 400Gb are on the horizon and expected to be available in the next few years. Since different transceiver technologies are leveraged for each speed, passive fiber TAPs do not change speeds midstream. If traffic is coming in at 10Gb with a wavelength of 1550nm, the traffic after the split has the same speed and wavelength.



For best results, cable types should be consistent across the flow. Match the cables to the need. In general, there are two major categories of fiber cable:

- Multimode
- Singlemode

Shorter-distance links often run over multimode cable, while longer-distance connections tend to use singlemode cable. The main difference between the two is that multimode has a larger core diameter (up to 62.5µm), which allows for a broader dispersion of light. This permits lower-cost, LED-based optical transmitters to be used, keeping the overall cost down.

Since the light is dispersed across multiple modes on a larger core, the light has a tendency to bounce around a lot while it is traveling through the cable. Since different modes travel different lengths, the signals arrive at different times, making it difficult to distinguish one pulse from another. This leads to higher attenuation, or loss of signal, as the light travels down its path. Because of this, multimode is only rated for shorter runs of up to a couple hundred meters,

depending on the cable type. It should be noted that the larger core multimode cabling (62.5µm) should only be used for 1Gb and below.

Singlemode fiber runs at higher bandwidths over smaller cores. This requires higher-precision instrumentation and higher-priced laser diodes to transmit the signal. Vertical-cavity surface-emitting lasers (VCSELs) are small flat emitters commonly used for short and medium distances. Longer distances such as 40km and beyond require more precise (and higher-temperature) Fabry-Perot lasers. Figure 8 (next page) depicts common cable types.




Figure 7. Multimode and singlemode cables.



Power Budgets and Light Loss

When you look to the night sky, you see stars. However, there are many stars the eye cannot see. Light is obstructed by clouds or pollution. It is also dependent on the brightness (or power) generated by the star itself and the sensitivity of your eyes. Unlike starlight that has traveled billions of miles in a relatively straight line through the vacuum of space to your eye, the low-powered light in an optical cable is bouncing off the walls of the cable shielding and must transfer through multiple connections, which all inflict light loss. As a result, cable light levels degrade fairly quickly.

Optical power loss of fiber optic cable is measured in decibels (dB). A quality light source is required for the receiving end to properly understand the signal. If the signal is too weak, the message will not be properly interpreted and packets will be dropped. The optical power budget is the difference between the Transmitter Power output and the receiver sensitivity as shown in Figure 8. All passive TAPs divert a portion of the light without boosting the signal. So it is important to understand how much loss is incurred to preserve proper light margins. It is also critical to recognize that the devices transmitting and receiving the light are completely external to the TAP, and all have their own degrees of variance. Most optic vendors provide specific power and receiver sensitivity information about their products, which may vary significantly from the industry specifications. In many cases, the actual numbers are much better than those shown in the Institute of Electrical and Electronic Engineers (IEEE) Standards 802.3 documents.

Within a cable infrastructure, it is important to understand what components can negatively affect the light as it travels from one point to another. Some degradation, such as attenuation, is simple math and cannot be avoided. These types of degradation tend to be linear and relatively small compared to other loss factors, such as dirty connections, poor splices, or mixed cable types, that can have a serious loss impact and should be avoided. Handheld optical time-domain reflectometer (OTDR) devices are available to test cable plants by injecting a series of optical pulses and then measuring light as it is reflected back from points along the fiber. Although these are often used for troubleshooting, they can quickly validate some of your calculations.

	Cable Type	Diameter (µm)	Color	Connector	Typical Usage
OM1	Multimode	62.5/125	Slate	LC	FE/1Gb/10Gb
OM2	Multimode	50/125	Orange	LC	FE/1Gb/10Gb
OM3/OM4	Multimode	50/125	Aqua	LC/MPO	FE/1Gb/10Gb/40Gb/100Gb
OS1/OS2	Singlemode	9	Yellow	LC	FE/1Gb/10Gb
Cat 5e/Cat 6A	Copper/twisted pair	N/A	Various	RJ45	FE/1Gb

Figure 8. Common cable types.

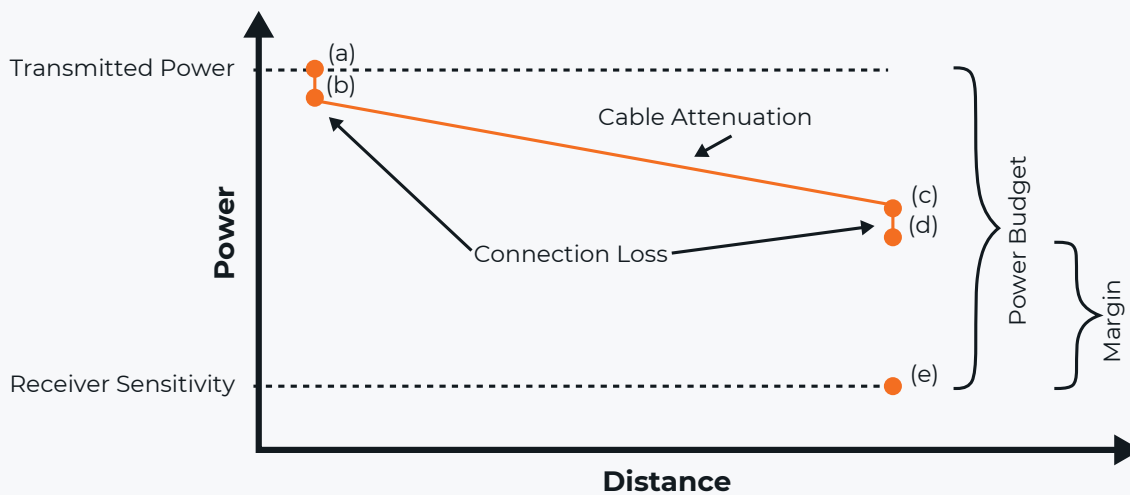


Figure 9. Theoretical loss chart.

The above chart shows the assumed loss associated between two endpoints with a transmitter at one end and a receiver at the other, with two connectors (at each end). The following formulas may be used:

Power Budget = Transmitter Power – Receiver Sensitivity = a – e

Cable Attenuation = Decrease in signal strength due to absorption and scattering per kilometer of a given cable type = b – c

Connection Loss = Signal degradation due to connectors in the system = (a – b) + (c – d)

Total Cable Plant Loss = Cable Attenuation + Connection Loss = (a – b) + (b – c) + (c – d)

Power Margin = Additional power that could be consumed while still providing a valuable signal = Power Budget – Total Cable Plant Loss

Whenever possible, it is best to run the calculations using the actual numbers from the transceivers and cables in use. An alternative method is to take the worst-case scenario and plug in the minimum numbers as established in the IEEE specifications. If we were to pull numbers for a 10-meter run of OM2 multimode fiber running 1Gb (according to IEEE 802.3-2012 section 3 specifications) we would find:

1000BASE-SX Transceiver Average Launch Power (Min) = -9.5dBm

1000BASE-SX Receiver Sensitivity = -17dBm

Attenuation rates of multimode cable (for 10 meters) = 3.5dB/Km = .035dB/10m

Connection loss of multimode connectors = .5dB

Plugging in the worst-case numbers into the original equations, we would come to the following conclusions:

Power Budget = (-9.5) – (-17) = 7.5dBm

Cable Attenuation (10 meters) = 3.5/100 = .035dB

Connection Loss = .5 x 2 connectors = 1dB

Total Cable Plant Loss = Cable Attenuation +
Connection Loss = .035 + 1
= 1.035

Power Margin = 7.5 – 1.035 = 6.465

Thus, with a power margin of 6.465dB, a TAP will fit nicely into this network. The TAP with the highest maximum loss in Figure 6 is 6.2dB (including connections to the TAP). So there is ample margin to insert a 50/50, 60/40, or 70/30 split-ratio TAP into this environment.

However, the user should be aware that all environments are different. The 1Gb example shown above provides for a much larger margin than higher-speed optics such as 10Gb, 40Gb, and 100Gb. As an example, the entire power budget allocated for some short-range 40Gb transceivers is less than 2dBm. Best practices dictate running the numbers for each installation. As a general rule, Gigamon does not recommend using a 70/30 split ratio for 10Gb multimode infrastructures as the light margins are too low for the monitored traffic.

To quickly summarize light calculations determining passive TAP placements, there are four primary considerations that come into play:

1. Transmit power (the starting light signal)
2. Receiver sensitivity (residual light seen at the other end)
3. Light loss within the cable plant (prior to TAP insertion)
4. Impact of the TAP (the actual TAP signal loss)

Active TAPs

Active TAPs are not passive. They require their own power source to regenerate the signals. There is no split ratio consideration because the TAP receives the message and then retransmits it to both the network and monitoring destinations. From a high-level perspective this would appear to be a positive feature. Even so, passive TAPs are preferred. During a power outage, an active TAP cannot regenerate the signal, so it becomes a point of failure. Since a passive TAP is not powered, it would be unaffected during a power outage, and the packets (originating from a source that still has power) would continue to flow. Some active TAPs do incorporate bypass or failover technologies to mitigate this issue, but more on that toward the end of this section.

When are active TAPs preferred? Active TAPs are commonly used for the following applications where passive TAPs are not a good alternative:

1. Locations where the light levels are too low to use a splitter → regeneration provides a viable solution
2. Copper infrastructures → where electricity is used to move electrons (instead of photons)
3. Signal conversions → since an active TAP regenerates the signal anyway, it can also be designed to create a signal of a different type (such as 10Gb SR converted to 10Gb LR)
4. SFP-based links that cannot otherwise be broken (such as TwinAX cabling) → regeneration works here as well

As long as the drawbacks of power failure are fully understood, active TAPs provide excellent value and extend visibility to sections of the network that would otherwise go unmonitored. Sophisticated active TAPs offer battery backup to extend usage during power failures. When the battery begins to die, some TAPs will offer additional failover capabilities. For example, when certain active copper TAPs lose power, electromagnetic relays fall into place to physically close a link to allow traffic to continue flowing through the network. The monitoring traffic stops, but at least the network

traffic is protected. Do be aware that when the relay closes, a renegotiation takes place, so a few packets would be affected. TCP transmissions would normally accommodate for the loss, but be aware higher-speed networks are more susceptible to routing table changes and other effects.




Figure 10. An example of a G-TAP® ATX always-on active copper TAP providing multiple power options, battery backup, and failover capabilities.

Other TAPs

Bypass Technology

Bypass technology is the ability to take an active flow of traffic and quickly reroute the flow to “bypass” a particular process. From a hardware perspective, bypass technology is a derivative of a TAP. It uses similar port pair interconnections as a TAP to provide link protection capabilities, typically for security-based inline tools.

An inline tool passes live traffic directly through a tool to process the live traffic before it is forwarded on to its final destination. Unlike most out-of-band analysis tools that never affect live traffic, an inline tool, such as intrusion prevention systems (IPS), can drop or even add packets into the production network. Since it is running as an inline application, a tool failure could be devastating and bring down the entire system. Bypass technology was designed to protect this from happening. These tools can bypass a tool that enters into a degraded or down state. This could be due to power failure or even oversubscription. The bypass device will continually monitor the tool by issuing heartbeats. If a heartbeat does not properly pass through the tool, the device will automatically close a link, forming a bypass connection

Aggregation Tap Vs. Aggregation Node

As described earlier in this paper, one of the benefits of a TAP is having separate monitoring ports for ingress and egress. There are, however, some TAPs that combine both feeds into a single monitoring port. This is sometimes called an aggregation TAP. The benefit is it reduces the number of monitoring ports, but the drawback is the risk of oversubscription and dropped packets. As such, these devices are not commonly recommended. If links are running at low utilization (< 5–10 percent), a better option is to combine edge traffic from multiple TAPs by use of an aggregation node. These devices, like the Gigamon GigaVUE-TA Series, offer filtering rules to aggregate traffic prior to sending it to more intelligent components of the visibility platform.

Standalone vs. Embedded Taps

TAP technology has traditionally been deployed as standalone devices, remotely distributed across the network wherever traffic needs to be seen. The monitoring ports can either connect directly to the analysis tool or to the visibility nodes to efficiently filter and distribute traffic among multiple tools.

A growing trend is to deploy the TAP hardware as an embedded module within a visibility node. One major physical difference between standalone TAPs and embedded TAPs is the exposed ports. Standalone TAPs have ports for both network and monitoring connections, while embedded TAPs only expose the network ports. The monitoring ports are connected directly to the backplane of the system and routed accordingly. This simplifies the cabling infrastructure and enhances operational efficiencies. Since there are no open monitoring ports, it is impossible to arbitrarily connect a sniffer or other capture device without proper configuration. Thus it provides for a more secure environment.

10GBASE-T

10GBASE-T is an Ethernet technology designed to deliver 10Gbps rates over short-reach copper pairs. It is mainly used in top-of-rack switches, servers, or appliances in a data center to drive down the cost of infrastructure and connectivity. Current 10GBASE-T SFP+ transceivers exceed the maximum power dissipation requirement in the SFP+ specification (SFF-8431). Additionally, because of the high speed of 10Gbps and associated noise on a copper link, such links are not conducive for tapping. Recommended options in such situations are:

- Best option: Do not use 10GBASE-T if tapping is required; use a 10G short-reach (SR) interface instead, which can be readily tapped.
- Alternative option 1: Use a SPAN/port mirror on the switch hosting the 10GBASE-T link. Be aware of the limitations of SPAN sessions when using this option, as explained in the “TAP vs. SPAN” section above.
- Alternative option 2: Use back-to-back media converters with a TAP or port pair, noting this has no fail-to-wire. Note that as the 10GBASE-T PHY performs significant signal processing, this adds considerably more latency than the corresponding 10G optical PHY would. For latency-sensitive applications, the impact of the two media converters needs to be taken into account.



TAP Best Practices

A TAP is a basic building block of any visibility system. For complete coverage, many companies have adopted a TAP-all strategy as a best practice. This means that all critical links are set up with TAPs (and/or SPANs), even if the traffic is not under continuous monitoring. By having the TAP already in place, in the event of a security breach or troubleshooting requirement, the data is readily accessible.

The best time to deploy a TAP is when the infrastructure is being built, as it is always more costly to introduce equipment after the fact. A TAP installation requires bringing down a network link, so should be done during a scheduled maintenance window. TAPs are generally preferred over SPAN ports, yet both provide value. Best practices dictate deploying physical TAPs for critical links with medium to high utilization. SPANs are best used in locations that are not conducive to TAPs. Examples include links with power budget limitations and remote sites with low-utilization links.

Traditionally, if both options were available, passive TAPs were generally preferred over active TAPs. The primary reason for this was to minimize loss during power outages. This trend, however, is changing. Not only do active TAPs boost the signal to provide longer distances, but many now include battery backups to minimize power loss and provide fail-safe operation.

Fully understand the light limitations of your environment prior to making any infrastructure change. Power budgets dictate proper TAP deployment. They are also used to determine the appropriate split ratios to deploy. Major obstacles regarding loss include distance (attenuation), connections, split ratios, splices, and dirty environments. It is always best to use the sensitivity and power ratings of the specific optics in use. When actual vendor-based light numbers are unavailable, worst-case numbers may be calculated as per IEEE specifications. If the power budget is too narrow, consider using an active TAP or SPAN port. Another option is to upgrade the optics and cabling to a standard rated for longer distances. Longer-distance optics are more expensive but tend to use higher-end lasers leading to strong signals. Having extra power margin on a critical link removes risk and is often worth the extra expense.

Most TAP failures are due to improper cabling. When connecting TAPs, always use new cabling and properly clean all connections. Never mix and match cable types within a single, end-to-end link. Verify the wiring diagrams to ensure the proper cables are plugged into each port. Match each TAP to the cable type in use, and never bend cabling beyond specifications. For newer technologies, such as Cisco BiDi deployments, only use TAPs that are rated for the exact wavelengths in use.

Not all TAPs are created equal. Check with peers within your industry for recommendations of quality vendors and ask about hardware warranties. Like optical transceivers, mileage can vary from one vendor to another. If a vendor only warrants the product for a short period of time, question the quality. There is no harm in asking for mean time between failure (MTBF) rates either.

Although a TAP can be connected directly to a monitoring tool, it is far better to connect directly to the Gigamon Visibility Platform. The visibility platform is a matrix of nodes deployed throughout the infrastructure to act as a common platform to move packets from any source to the proper monitoring, analysis, or security tools. This allows for tool consolidation to optimize your monitoring solutions while extending overall visibility across the network for pervasive security and analysis.

Traffic can be replicated, aggregated, or filtered as required. Higher-level intelligence, such as packet de-duplication, SSL/TLS decryption, or header stripping, can also be accomplished within the platform, reducing the load on each of the tools while accelerating mean time to resolution when problems emerge. NetFlow generation is also available, making more efficient use of critical switches and routers.

Summary

A TAP represents the connection point where real traffic is copied directly from the network. As such, it is the first step toward any visibility solution. TAPs can be standalone devices or integrated directly as a module inside a visibility node. In both cases, traffic is copied for monitoring, security and analysis as the traffic continues to pass through the network unimpeded.

About Gigamon

Gigamon offers a deep observability pipeline that harnesses actionable network-level intelligence to amplify the power of observability tools. This powerful combination enables IT organizations to assure security and compliance governance, speed root-cause analysis of performance bottlenecks, and lower operational overhead associated with managing hybrid and multi-cloud IT infrastructures. The result: modern enterprises realize the full transformational promise of the cloud. Gigamon serves more than 4,000 customers worldwide, including over 80 percent of Fortune 100 enterprises, nine of the 10 largest mobile network providers, and hundreds of governments and educational organizations worldwide. To learn more, please visit gigamon.com.



Worldwide Headquarters

3300 Olcott Street, Santa Clara, CA 95054 USA
+1 (408) 831-4000 | gigamon.com

© 2023 Gigamon. All rights reserved. Gigamon and Gigamon logos are trademarks of Gigamon in the United States and/or other countries. Gigamon trademarks can be found at gigamon.com/legal-trademarks. All other trademarks are the trademarks of their respective owners. Gigamon reserves the right to change, modify, transfer, or otherwise revise this publication without notice.