WHITE PAPER

Best Practices for Visibility Architecture Tap Planning

Data filtering, and the conversion of data into actionable information, is dependent on the initial data being correct and relevant.

Getting the Right Data Is Key

When it comes to data monitoring, ensuring proper access to network data is the most critical thing you can do. Everything else, data filtering and the conversion of data into actionable information, are all dependent on that initial data being correct and relevant. If this practice is ignored, you can, and will, end up realizing the well-known adage – “Garbage In equals Garbage Out”. This white paper is designed to be a generic guide to help assist you in optimizing the access layer of your Visibility Architecture so you can capture that correct information.

As you can see in Figure 1, there are three different frameworks that can be incorporated into a visibility architecture: the access framework, the out-of-band visibility framework, and the inline security framework. These three frameworks take the raw data from the network and manipulate it as necessary to get the appropriate data to the correct monitoring and analysis tool(s).

The access layer framework is naturally focused on creating access to network data. This is the base framework that then feeds data to packet brokers in either the out-of-band visibility or inline security frameworks where these data points can then be filtered and manipulated before being be sent on to the appropriate monitoring tools.

KEYSIGHT TECHNOLOGIES

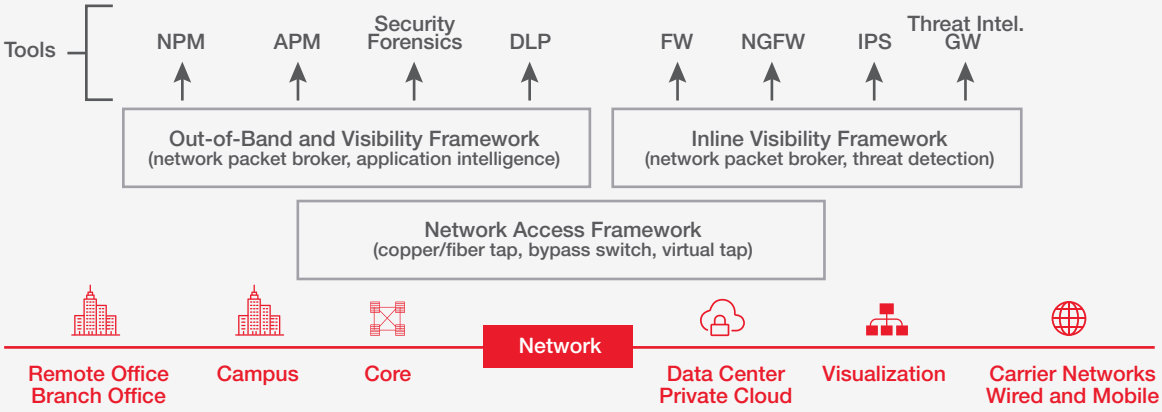

Figure 1. A generic visibility architecture.

Taps and SPANs

Proper visibility starts with proper data access. However, this activity also happens to be one of the least thought about activities by IT engineers. When it comes to data monitoring, many IT professionals simply use SPAN (switched port analyzer) ports from their network routing switch because “it was there” and “it was free”. While this may not be the worst decision someone can make in their lifetime, it is often a very unwise decision. This is because they did not give any preference to data integrity. The SPAN port may, or may not, provide the right data or be the right place to collect the necessary data.

Proper visibility starts with proper data access. However, this activity also happens to be one of the least thought about activities by IT engineers.

So what are your options? There are two main options – SPAN ports (as just mentioned) and taps. While both can be used to provide monitoring data streams, taps often have the most advantages. A free resource (What is a Tap, and Why Are Taps Critical to Network Visibility and Security?[1]) can explain the differences. A short synopsis of the capabilities of each solution are summarized in the following chart:

Functionality	Tap	SPAN
Provides access to monitoring packets	X	X
Delivers a complete copy (100%) of data (including bad data vital for diagnosis)	X	
Has full systems resource priority during a crisis (i.e. does not drop frames)	X	
Less vulnerable to security attacks	X	
Does not create unnecessary, duplicate packets	X	
Does not create time stamp issues	X	
Recommended for lawful intercept	X	
Relieves SPAN port contention	X	
Plug & play: no configuration needed	X	

Another consideration is that SPAN port data is isolated to the data capabilities of your network switches. If a problem occurs towards the edge of the network or if you have more tools than available SPAN ports and sessions, then you will miss critical data. Figure 2 gives you an example of this.

[1] https://www.ixiacom.com/sites/default/files/resources/solution-brief/915-6855-01-inlinetapsvsspanports.pdf

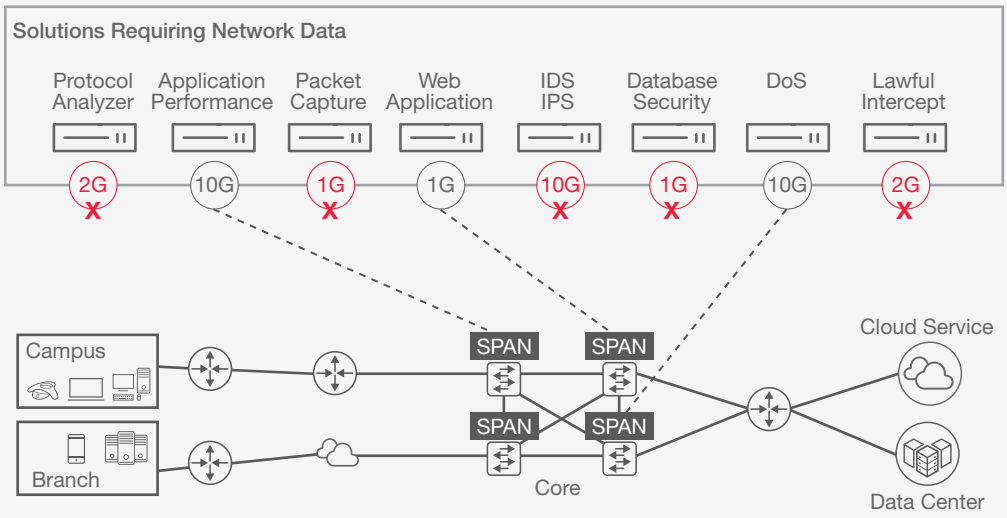

Figure 2. What is the core access layer problem?

As you can see, data is isolated to just one part of the network, the core network switches. In addition, a lack of necessary SPAN ports is going to cause several of the tools to sit idle and perform no usable function.

At the same time, the data that does come out of the SPAN port is probably going to be summarized data, i.e. you cannot export all of the data from a switch across two SPAN ports. For instance, say you have a 48 port switch with 10 gigabit Ethernet (GE) transceivers. If you run the switch at capacity, that is 480 GE of data. There is no way you can send 480 GE of data out on two 10 GE (20 GE total bandwidth) links. This is why you need to filter the data that goes to the SPAN ports. Even if you only want half the data, you are still extremely far outside of the monitoring bandwidth. With two dedicated SPAN ports, you can only get access to 1/12th worth of the data.

An alternative is shown in Figure 3. If SPAN ports are absolutely required, the insertion of taps at key points in the network can be extremely useful to gather the additional data required. However, depending upon your monitoring solution configuration, taps may not solve all of your problems. There may still be a mismatch between the supported data rate of your tool and the speed of the network data. This is where a packet broker can add value. These devices are inserted to help load balance and filter the data from the tap to prevent security and monitoring tool overload and improve tool performance.

A lack of necessary SPAN ports is going to cause several of your tools to sit idle and perform no usable function.

In addition, when you combine SPAN and tap data, you are mixing two very different sets of data. The tap sends a complete and full copy of the original data. SPAN ports change the data packets (time stamps, check sums, etc.). For troubleshooting purposes, you will need to keep track of the packet source (if you mix the two access types) so that you know what parts of the data may, or may not, be original.

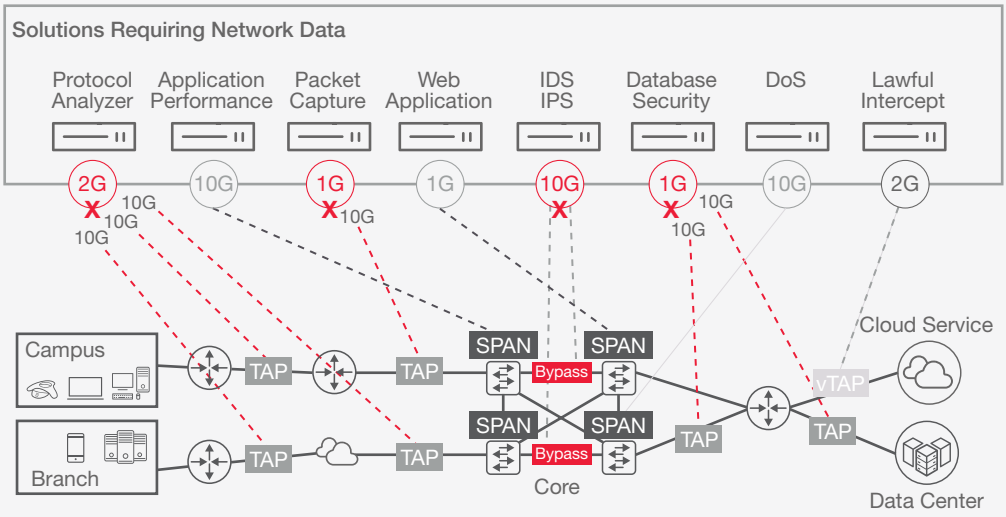

Figure 3. Added monitoring points help but they are only a partial solution.

Virtual Taps

Another consideration is that data access is not just about the physical environment but should address your virtual environment as well. Do you have virtualized servers or a virtualized data center? Do you have a cloud network? If so, do you have access to the inter- and intra-virtualized machine data? If not, you could, and probably are, missing key performance and security information. For this type of data access, you need a virtual tap, especially one that lets you filter the data you export so you do not blow out your LAN. A SPAN port definitely will not give you the data you need in this situation.

Data access is not just about the physical environment but should address your virtual environment as well.

Figure 4 shows three possible monitoring points for a virtualized server. Monitoring point 1 is an easily solved issue by placing a physical top at the top of your rack. However, this monitoring point misses 80% or more of the virtual traffic leaving the two common pain points, inter- and intra-VM data, notated as “2” and “3”. These pain points can be solved with a virtual tap, which is simply a software version of a physical tap, installed on the virtualized server. Now you have full access to that east-west data.

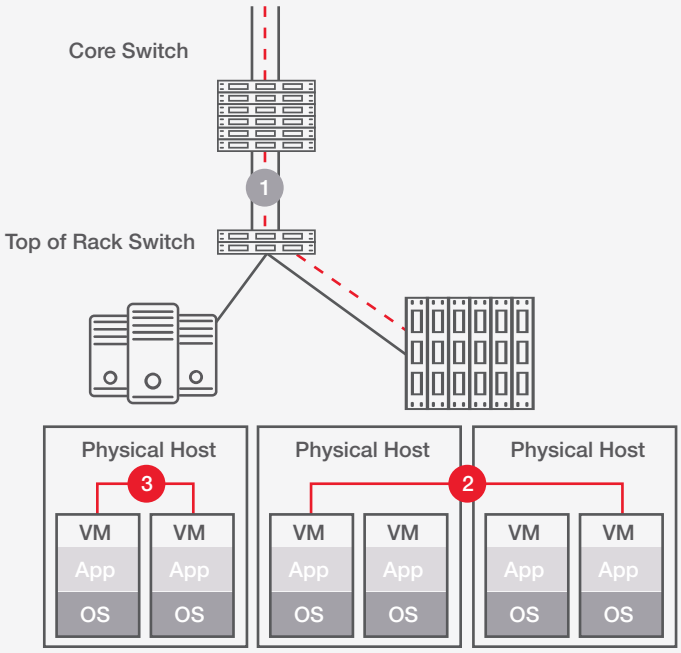

Figure 4. Exposing virtual data center data capture issues.

Bypass Switch (Inline Tap)

A third common access consideration is whether tools need to be deployed in the direct path of network traffic, referred to as inline. Typical inline tool examples are firewalls, intrusion prevention system (IPS), next generation firewalls (NGFW), web caches, and SSL encryption / decryption devices. If you have this type of architecture, a special type of fail-safe tap, called a bypass switch, is needed to preserve network reliability and uptime. Rather than making a copy of the traffic and funneling that to a monitoring port, the bypass switch forwards the main stream of data to the monitoring port for transmission to the tools. The bypass switch then receives the analyzed data back from the tool(s) and lets it pass downstream at that point.

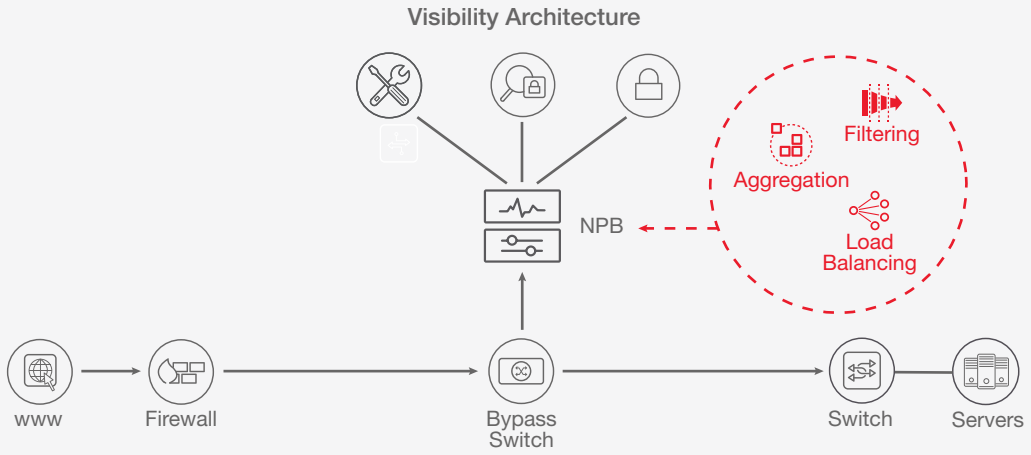

Figure 5. Bypass switch example.

In the event of a tool failure or if the tool needs maintenance/upgrades, the bypass switch enables the bypass of the inline tools so that traffic can pass downstream. This allows for limited, if any, no disruption to network performance during normal operation but also supports multi-tool and multi-network deployments, maximizing security and network up time. Adding a network packet broker after the bypass switch adds even more flexibility and capabilities.

The bypass switch enables the bypass of the inline tools so that traffic can pass downstream.

Generic Guidelines For Tap Deployments

A key point to understand is that every network and every architecture is unique. The data you need and the location of that data vary from company to company. A full analysis of what and where to tap may require the services of a consultant. For instance, are you trying to integrate real-time security tools or is a time delay for security data acceptable? Certain tools will need specific types of data to be effective. Other tools are less stringent on requirements. The answer depends upon what problem you are trying solve.

The data from these questions allows you to quickly get an overview of your needs and better formulate a direction. The first part of this activity is to give you some ideas so you really think through your architecture. There can be significant differences in the monitoring capabilities you deploy based upon on your needs. Some examples were given earlier like the need for security data, compliance data or troubleshooting data. Another example is whether you are concerned about the core network, or edge data, or both. Application level data can also require a different monitoring infrastructure. Once you get some basic ideas for all of your needs, other questions can start to help you narrow down the details of what you need.

That being said, a best practice in this area is often a general four step process based upon evaluation of the following criteria:

- What are your needs?
- What do you already have in place?
- What is the delta between the two?
- Where do you need to place the new equipment or where do you need to move the existing equipment to?

Start by asking yourself the following five basic questions:

- Q1 – What do you want to accomplish? e.g.- better troubleshooting data, better compliance data (PCI, HIPAA, financial data, etc.), obtain security-related data, collect performance monitoring information for trending, gather application-related data, or is it something else?
- Q2 – Where is your data located? e.g. – server farm, entrance to the network, spread across the network, contained within the data center, etc.
- Q3 – Does the tap need to be inline, i.e. do you need real-time data or is an information delay acceptable?
- Q4 – Do you need data from your virtual data center?
- Q5 – What is your budget?

The second step should be much easier, but maybe not. What equipment do you already have in place and where is it currently located? This includes the number of physical taps and virtual taps, the number of SPAN ports, the number of SPAN sessions that have been set up along with the data they are providing, the types of monitoring tools you have and their location, and the network packet brokers you have deployed and their corresponding locations.

There can be significant differences in the monitoring capabilities you deploy based upon on your needs.

The third step in the process becomes important to help you understand and decide the time criticality of your data. For instance, many large enterprises want to connect their tools inline in their network. This is usually just after the firewall but before the core network switches. Examples of security tools located here are often Next Gen firewalls, intrusion prevention systems, threat analysis tools and malware scanning tools. Other security tools (like sniffers, data recorders, SIEMs, and data loss prevention) are often inserted further down the line after the routing switches.

For inline tool deployments, a bypass switch is the preferred choice. A SPAN does not work here as you want to be able to connect the tool inline – so a bypass switch is the obvious choice. However, you need a bypass switch that will be able to allow failover to the network downline, should the bypass switch or any of the tools connected to it

fail. This is often called “failing open.” The last thing you want is a self-imposed network outage that costs the company thousands, or millions, of dollars. At the same time, this needs to be balanced with the probability that data passed downstream in a failover situation might contain malicious content and contribute to a security breach. If that probability dominates, then you will want the bypass switch to “fail closed” and stop all data flow in or out of the main company ingress/egress point.

The last thing you want is a self-imposed network outage that costs the company thousands, or millions, of dollars.

The next inline consideration is to make sure that the fail-over capability is as fast as possible (on the order of milliseconds or so) to ensure that there is minimal data disruption downline. Closely associated with this is the concern that you will probably want a bypass switch that has a fast heartbeat signal (on the order of 1 microsecond or so) so that the bypass switch knows as soon as a tool (that is connected to it) has failed. This failure notice then allows the bypass switch to initiate the fail-over process.

Another inline security consideration is whether you need High Availability. You may want to place your IPS or other security tools inline to maximize network protection. At the same time, you cannot afford any device failures. High Availability technology using multiple equipment solves this problem for you. This technology also adds extra cost and some complexity. However, there are bypass taps and packet brokers that support high availability which makes this easy and much more cost effective. Once you know what you need, the solution gets a lot easier.

The fourth major consideration is what information do you need from your virtual data center to achieve your monitoring goals? For instance, if you want to get better compliance data and you have virtualized servers (as most companies do nowadays), are you already using virtual taps to gather compliance data from your virtual environment, as well as using a physical tap on your physical data network. Most virtual traffic (up to 80%) stays within the hypervisor (as inter- or intra- VM traffic) and never reaches a physical tap or SPAN monitoring point. This means that there is a lot of data you may know nothing about. So, is your virtual data center following your company and industry standards for compliance? Without the right monitoring data, who knows?

Your budget also plays an important role. Maybe you cannot buy taps for all of your equipment. Maybe you need more expensive types of taps, like bypass switches. At the same time, you will need to compare costs correctly. For instance, by deploying the right type of bypass switch, you might actually decease other types of equipment you need. And then there is the Holy Grail possibility. If you deploy the right types of access equipment at the right places in your network, then maybe your costs will be lowered so that you can actually buy more equipment (more taps or more monitoring tools) for the same amount of budget that you already have.

Another item to consider when performing your cost analysis is whether you have properly accounted for any hidden costs. For instance, some engineers are under the impression that SPANs are “free”. This almost always is not the case. Each SPAN session needs to be provisioned so that the correct monitoring data is delivered to the appropriate tool. This is necessary every time you need different data. If you refer back to Figure 2, each of those three SPAN connections had to be configured. If you move the SPAN port connection over to a different tool (there were five tools not getting any information), then you will probably need to reconfigure the SPAN session so that it provides the right data to that tool. This applies each time you need to make changes, which can be often, if you have different data needs (compliance, security, troubleshooting, application performance, mobility tools, etc.). This SPAN programming cost can end up being significant. Ixia did an analysis and discovered that it could amount to over $6,000 per year for an enterprise. That $6,000 could buy quite a few taps. Taps also have the advantage that they are dedicated and do not typically require any programming. They are as easy as “set and forget.”

In regards to network visibility access, there are two common types of projects: a data center refresh or new construction.

Three Tap Deployment Examples

In regards to network visibility access, there are two common types of projects. The first is a remediation effort for some sort of problem that is either at the data center or at the edge of the network. This effort focuses on either remediating the problem directly or solving the problem as part of a data center refresh, but the refresh is not a complete re-architecture of the network. The second type of project is a Greenfield project that is new construction or a complete re-architecture of the network.

The specific visibility access needs for these two types of projects will vary depending upon various factors such as size of the business, business focus, and the monitoring needs of IT personnel. While these different factors customize the visibility access needs for each customer, we can still look at three different deployment examples to get a better feel for how a generic network diagram for tap deployments might look. Here are the three examples we will look at:

- Large financial trading company
- Large software development company
- Medium-sized enterprise

Example 1 – Large Financial Trading Company

In this example, the goals of the monitoring solution are primarily to validate the service level agreement (SLA) for transmission time. The requirement stems from both optimizing the speed of transmission times for financial trades but also to document trading times for liability purposes. Security is obviously another top concern.

This solution requires taps to be deployed at the company demarcation point to capture ingress and egress times for the packet data. At the same time, some core network data may be required to understand packet travel times for the traders. While SPAN ports may be used by some IT departments in the core, SPAN ports are not adequate for either location because data needs to be captured at the very edge of the network and any data captured within the core of the network needs to be 100% accurate. There cannot be any timing issues as the time stamp data is critical for this application. Duplicate packets from SPAN ports can also be an unnecessary distraction.

For financial trading companies, the perimeter is any ingress or egress connection to the data center. This includes cross-connects from other financial institutions and links to internet service providers. For egress traffic, this includes any un-quarantined out going traffic. Examples include data exiting from the DMZ out to another data center, traffic sent over fail-over pairs, and basically anything that is not east-west data center traffic.

Virtual taps are typically not an option as virtual machines are not used in financial trading environments. Inline taps are not commonly required either. They may be used for some security tool purchases but the most common solution here is a physical fiber tap.

Again, while the business specifics will dictate the exact placement of the taps, there are some general recommendations that can be employed:

1. Tap the ingress/egress points to give you all of the client to server traffic.
2. Then tap at choke points.
3. If you have branch locations, you will probably want to tap at the ingress/egress to the core feed handler to get performance metrics out and back from the branches. By combining this with the data center transmission information, you will have the latency calculations for the service provider, especially regarding fill and order issues.
4. After that, things get specific to the customer.

Figure 6 gives an example of this type of deployment.

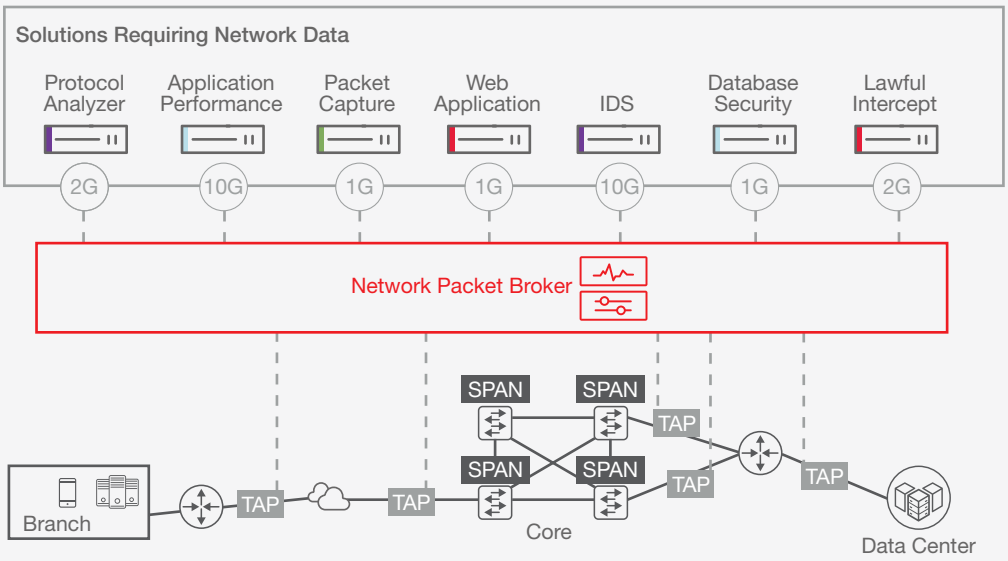

Figure 6. Large financial trading company example.

Example 2 – Large Software Development Company

In this example, the business is transmitting a high volume of data across the network and is typically looking at performance information. This may involve a combination of physical taps and virtual taps. It will also involve gathering a lot data within the core that may come from either physical or virtual machines. Virtual machines will probably be more of the norm nowadays.

While a selective SPAN or tap may work in the interior of the network. There will probably be a heavy requirement for data that necessitates the need for top of rack switch information on application performance. This often involves integrating data from multiple applications spread across different cabinets within the data center that need to talk and communicate with each other.

Here are some general recommendations for this configuration:

1. Tap the ingress points to give you all of the client to server traffic.
2. Then use taps or SPANs in the physical data center to capture application performance information. A key question is where are the applications located within the data center and how many locations are there?
3. If application performance data details are absolutely required, consider adding a packet broker that can perform application filtering and metadata collection to get more performance information.

4. For virtual data center deployments, you will need a virtual tap. Chose one that has minimal impact on the hypervisor but still gives you all of the east-west data from inter- and intra-VM communications. After you collect that data, make sure the virtual tap has built-in filtering capabilities so you can segment out only the relevant data and then forward that on to a packet broker for consolidation and transmission to the appropriate physical tool.

Figure 7 gives an example of this type of deployment.

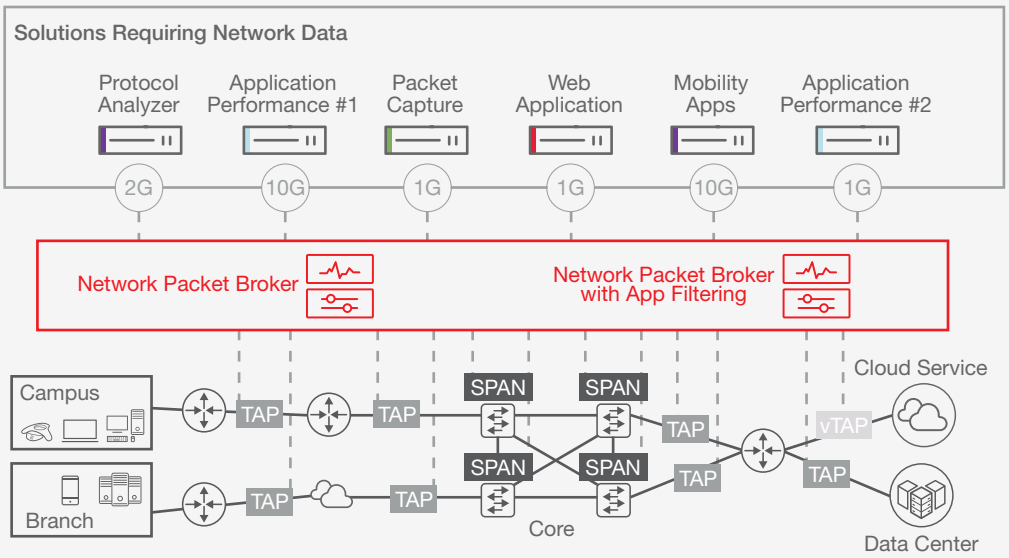

Figure 7. Large software development company example.

Example 3 – Medium-Sized Enterprise

In this example, more emphasis is placed upon collecting troubleshooting data. Due to cost pressures, these organizations may see a mixture of SPAN and taps. In addition to the configuration costs for SPAN sessions (mentioned earlier), the network engineer will need to be cognizant of the fact that they may get duplicate packets, summarized data, and altogether missing bad packet data from their SPAN ports. These issues can be eliminated/overcome with the proper placement of taps throughout the network and the use of a packet broker to help aggregate, filter and de-duplicate unnecessary data.

One counterpoint that supports the use of SPANs is that most SPAN connections will usually provide enough troubleshooting data to allow you to uncover P1 issues in the data center. This is especially true where VLANs are used, as it is often not cost effective to deploy a tap for every VLAN.

Due to CPU resource priority on the data switches, SPAN ports can drop potentially useful troubleshooting information, withou any notification to yo of what data was lost.

Here are some general recommendations for this configuration:

1. Tap the ingress and egress points to give you all of the client to server traffic. It is not typical to troubleshoot the core network, although there are times when it is necessary to isolate something to a core switch. When you do need to troubleshoot the core, there is typically no need to tap at the top of every rack switch.
2. SPANs may be useful in the physical data center to capture application performance information but the use of taps is highly recommended instead as useful troubleshooting data that you need (bad packets, bad frames, timestamp data, etc.) may be missing or altered by the SPAN port. This is especially true if your data is heavily loaded. Due to CPU resource priority on the data switches, SPAN ports can drop potentially useful troubleshooting information, without any notification to you of what data was lost. This will be especially irritating in a troubleshooting situation. In addition, the use of SPAN ports often requires configuration for troubleshooting purposes which will require Change Board approvals and time delays. Once taps are installed, you can access the data whenever you want without delays.
3. For virtual data center deployments, you will need a virtual tap. The same note from recommendation 4 under the large software development company example applies here as well. This information is often used for performance trending, compliance or security analysis rather than troubleshooting information.
4. If you are troubleshooting security issues, a good recommendation is to skip the use of SPAN ports for the reasons mentioned previously. The last thing you want is incorrect data, or any SPAN issues, when it comes to your security equipment. Most of the requisite security data can be collected by taping at the perimeter.
5. For inline security monitoring deployments that also require data for out of band monitoring tools, there are bypass taps that can provide both functions in a single unit and simplify things for you. This is especially useful if you need to calculate the latency time of the security tools on the network as the out of band copy of that data can be routed to monitoring tools for comparison of the latency into and back from the security tools.
6. Troubleshooting connections to internal users can be accomplished with the placement (usually temporary) of a tap in the wiring closet near the suspected problem area.

Proper planning for the access layer of a visibility architecture is one of the most important considerations you can make for network monitoring.

Figure 8 gives an example of this type of deployment.

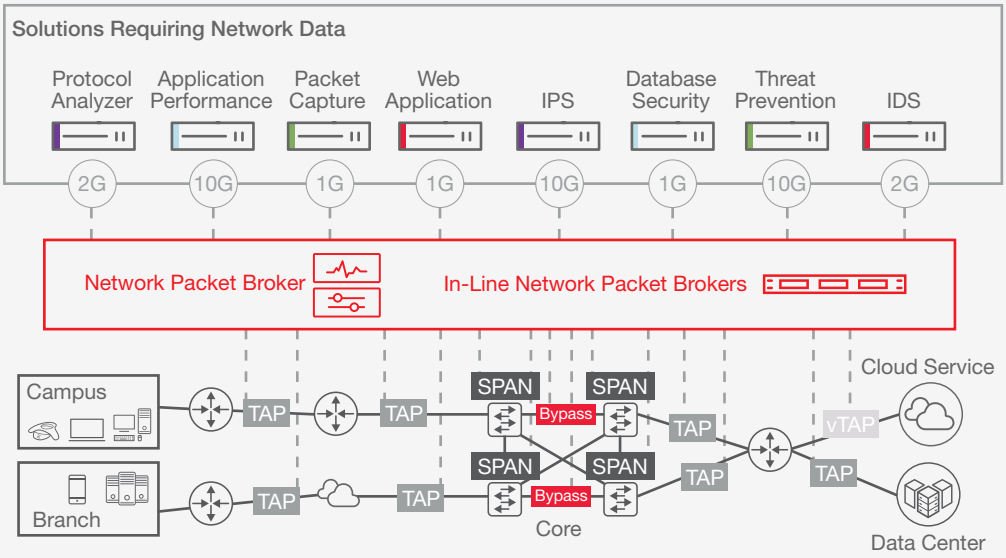

Figure 8. Medium-sized enterprise example.

Conclusion

Proper planning for the access layer of a visibility architecture is one of the most important considerations you can make for network monitoring. What mixture of taps and SPANs should you use? Where should you tap? How many taps will you need? While specific answers depend upon your network and your needs, almost every monitoring solution will benefit from the following five recommendations:

- Use taps where you can to ensure that you get the best data possible as fast as possible
- Tap your network ingress and egress points
- Tap any known choke points
- Deploy virtual taps to gather information from your virtual data center
- Make sure you understand your security tool requirements – do you need a high availability solution, have you set up the bypass tap correctly for your needs, and have you considered deploying an inline packet broker to help your bypass tap support the proper number and types of monitoring tools that it needs to?

Find out more information on visibility architectures by visiting www.ixiacom.com/solutions/out-band-monitoring.

Learn more at: www.keysight.com

For more information on Keysight Technologies' products, applications or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus

KEYSIGHT TECHNOLOGIES