NIAGARA NETWORKS

# Network visibility

A clearer view of your network

# About this paper

Managing network security means to be on top of things. That is – to be aware of everything that is taking place in the network, and having every skill and resource at your disposal to control and react instantly to each potential problem. In the information technology (IT) world of network connectivity and security, we call this Network Visibility – and all that it comprises. This is all that we need, to get a handle on things, on what is going on over the enterprise network. The mechanism is a combination of autonomous system modules (hardware and software) and human intervention (when needed) to monitor, gather data and analyze it, and make decisions (control), as required. Today, it's mostly automatic, but the human factor is often the definitive one.

**The IT personnel have to take care of two ongoing concerns in parallel:**

1. Keep the network up-and-running, and solve all system and human-computer related issues (a great load in and of itself).
2. and: Fend off all attempts at cyber breaches, take care of viruses and malware, prevent data leaks, etc. Not an easy task, despite all the hi-tech software and hardware that IT implements to assist them in their work.

Network visibility – basically, being on top of things – provides the enterprise with the precise view, correct leverage, and potential to getting a handle on everything and making sure it all runs the way it should.

This paper will primarily focus on part 2 - the security issues. We present here an overview of the inherent problem scenarios, and the solutions that empower network visibility to enable a company's users to receive the best network access services while the organization maintains its integrity, and protects its data, intellectual property, and other business information.

In short: to offer the enterprise and its clients the best, stable, and secure network possible.

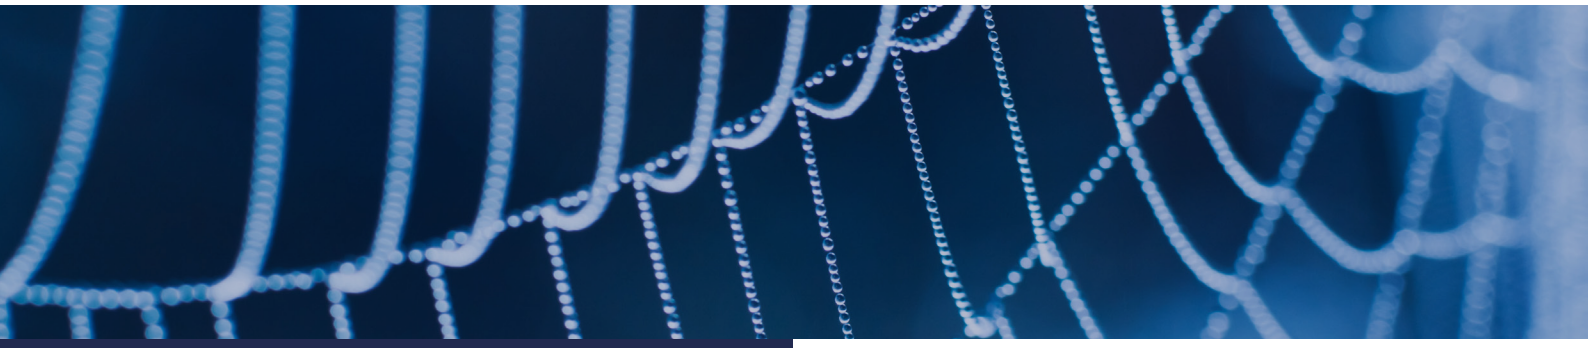

# An insight into network visibility

## Why it's enterprise-critical

Networks are getting ever more multifaceted with a greater number of physical and virtual devices and services connected and linked up to them, and therefore, maximum visibility is essential. If the enterprise is not in touch with, and aware of what's on the network, then it will likely lose control of it—and that opens up a Pandora's box of potential problems that makes the enterprise vulnerable.

Today, there is tremendous increase in BYOD (Bring Your Own Device) and proliferation of the IoT (Internet of Things), with numerous devices connecting to and communicating with and through networks growing by the minute. Some are internal – set up by IT, with all the inherent and correct anti-malware, and security devices and protocols. Some are external – such as guests trying to connect up, or casual passersby and visitors roaming with their tablets and smartphones – and sometimes (unfortunately, only too often) attempts at breaching the network by malicious agents (that could be either physical or virtual).

> It was the year nothing seemed safe. Bombshell hacks were revealed one after another in 2017, from an Equifax breach that compromised almost half the country to global ransom campaigns that cost companies millions of dollars.
>
> CNN Tech, Dec 2017

These occasional rogue and unknown devices often show up on the networks, but even more disconcerting – they very often do not!

How do we handle the ones we do detect: Do the suspicious fingerprints belong to employees going about their everyday business, or is it possibly an indication that something more nefarious is taking place? Typical suspicious users can range from dishonest employees and industrial spies and criminals, to state-sponsored cyber-hackers — each hunting for unprotected points of access with vulnerabilities to exploit.

It is of the utmost importance that the enterprise knows who and what is on the network - and at all times. If the network is unprepared (until now) – it is critical that the enterprise does something about it – and soon!

# Network visibility

## Facing the new network challenges

In the past, networks were hard-wired and ran through the building's floors and walls. Nowadays, network security has become a lot more critical, and its deployment is complex. Implementing network security infrastructure based on a Network Visibility Layer is the effective and robust way to keep the security of your network well managed.

Wireless networks and mobile devices have become pervasive. This includes network devices such as wireless gateways and switches, as well as endpoint devices such as mobile or Internet of Things (IoT) devices. For each device, the enterprise needs clear visibility to know to whom the device belongs, and what segment of the network it is trying to access - and also what it is permitted to access.

With this information, IT can allow authorized users to go deeper into the network, while keeping unidentified devices at bay by limiting access. In a scenario where a potential breach or threat may occur, the enterprise needs immediate threat detection, and a proactive response that can only take place

if the enterprise is in control – if the network is persistently visible to IT, and it is truly on top of things. The 'Network Visibility Layer' addresses three main needs : Visibility, Availability, Efficiency.

## Visibility

Visibility deals with the ability to know what is happening in the network. To achieve this, network inspection and monitoring devices are deployed at key nodes in the network to carry out surveillance of the traffic, gather and pass on the data to the security monitoring tools for further processing and – where required – take measures to deal with any anomalies or suspected transmissions. Full visibility, however, is achieved by deploying the tools through a layer of visibility devices.

Network taps – as described further – are one of the key tools in fulfilling visibility.

## Availability

Another key factor in today's 'time-is-money' business activities, is ensuring continuous service to all users. This is known as business continuity or high availability (HA) and includes making sure that there is uninterrupted control and unobstructed data traffic (with zero loss) across the entire network. Here is where network visibility flexes its muscles. Designed around continuous monitoring, and when added to that, the capacity to circumvent issues by bypassing and rerouting – the visibility solutions built into your network architecture secure your traffic's continuous and stable service.

When your IT personnel can be network-aware via a clear, up-to-date, on-screen network status, they then become empowered, and are able to rapidly identify, isolate, and resolve problems through root cause analytics, and optimize support for critical applications. This offers them complete visibility of the

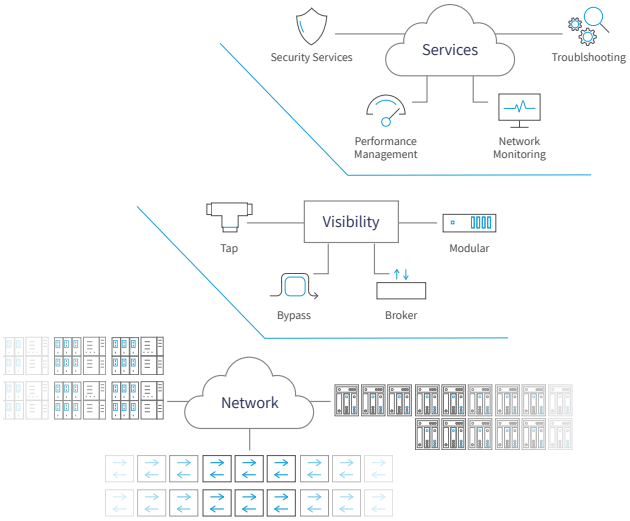


This figure depicts the visibility layer between the network data and the network tools, ensuring that the network tools will operate efficiently and with high availability without impacting the underlying network.

exact network paths taken by all data, and they are able to monitor both the state and the impact of the running system applications when network failovers occur, during recovery, or even during standard maintenance upgrades.

Unimpeded network system performance is essential for business continuity, and network visibility is a key factor for guaranteeing this Network Bypass – as described further – are one of the key tools in fulfilling availability.

## Efficiency

By offering the means to handle the vast amounts of data and inherent traffic loads, network visibility management features enhance the system efficiency. Such mechanisms as filtering, aggregation and load balancing create a more in-tune, leaner system allowing more data to cross the to network tools effectively.

Network Packet Brokers (NPBs) – as described further – are one of the key tools in fulfilling efficiency.

# Visibility adaptation layer

The visibility adaptation layer, is the layer that includes devices such as taps, bypasses, and packet brokers, as well a software defined network (SDN) architecture overlay to turn all the individual building blocks into a single virtual switching fabric and to assist in keeping everything running smoothly and seamlessly. There are two basic deployment architectures Inline (in band) and Monitoring (out of band).

Out of Band (Network Monitoring)

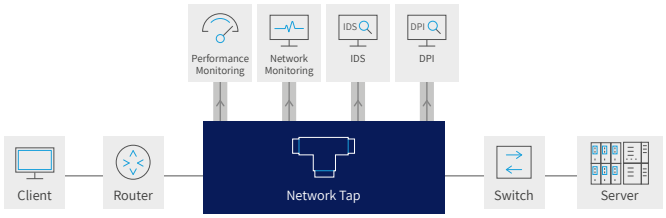


Inline (Security)

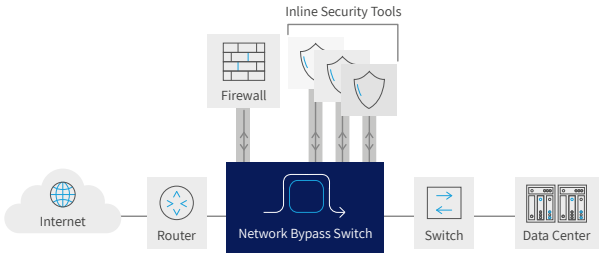


The deployment of inspection and analysis devices, and the nature of the monitoring, tracking, and security activities affect the nature of their deployment, whether out-of-band or inline. Both Inline and Out-of-Band deployments may include NPB for improved tool efficiency and reduced ToC.

A network security device such as a firewall that needs to inspect and possibly block traffic will be deployed inline, while a network performance tool that only needs to inspect a copy of the traffic will be deployed out-of-band, as its actions are not intended to impact the traffic flowing through the network.

While there is a typically an association between the type of network tool and its deployment, in many real world use cases the Network and Security Manager will work together on a network design that may well combine both deployment types. For example, the need to monitor the traffic going through an inline device is a use case combining both deployments scenarios.

The following sections describe in greater detail some of the network visibility devices, that can be used and combined as building blocks to satisfy the connection of any network tool and any business need.

When designing the architecture of the network visibility layer - consider that an inline visibility device such as a bypass (see below) can prevent failure of an inline tool from taking part or all of the network down.

## Port mirroring / SPAN

Also known as SPAN (Switched Port Analyzer) or roving analysis, port mirroring is a method for catching and observing network traffic in a non-intrusive manner. It 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 mirrored (SPAN) port. Using software, we can easily configure what data is to be monitored.

While useful and flexible, the number of SPAN ports per device is limited, and they typically cannot be used to provide visibility into all of the traffic. It should also be noted, that SPAN ports are not lossless, and that they can reduce the performance of switches in high traffic environments.

## Test access points (taps)

In order to enable full exposure and control, 100% of the traffic flowing through the network must be accessible, with no data loss. This is where the taps come into play, because without that full exposure, the network security systems will not ‘see’ all the traffic and will not be able to handle or block suspicious traffic. In addition, any performance inspection devices will be hampered in detecting and flagging system degradations and pinpointing system performance issues if they do not have visibility of the network traffic.

Well-designed architectures with multiple built-in access points throughout the network will also empower system administrators to significantly reduce their mean time to getting a handle on every issue they investigate. Thanks to the deployed network taps, the administrator will have the ability to quickly connect to the network at its various endpoints for essential troubleshooting activities.

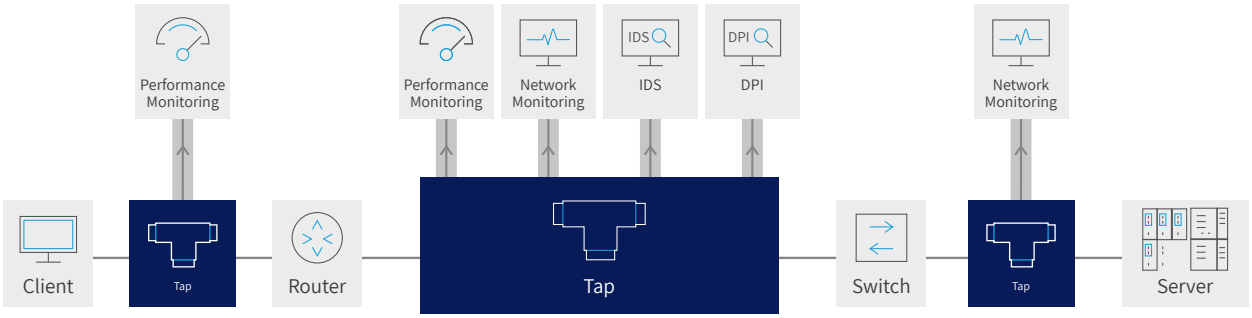


Best practice for network design recommends using a network tap to provide connected tools and devices with 100% of the data (full stack) passing through the network. The monitoring devices will be able to compare packets on a granular level before and after they reach key infrastructure elements (such as web servers, load balancers, VoIP systems, switches, etc.).

Network taps can be inserted in the network at different points, to provide full exposure and visibility. All of the enterprise's network security and monitoring tools and devices will then be continuously provided with a complete copy of the network traffic data. By using network taps, this takes place seamlessly and without interruption – and without disrupting any of the network's operations or performance.

Most network engineers agree, that the network taps (that are a purpose-built piece of hardware whose only function is to copy network data and send it to any and all devices connected to it) provide the optimum foundation for network visibility. Not only do they replicate all of the network data that flows through them, they can be placed pervasively throughout the network to enable a far and wide reaching range of visibility into the traffic passing through essential system and network elements and OSI layers.

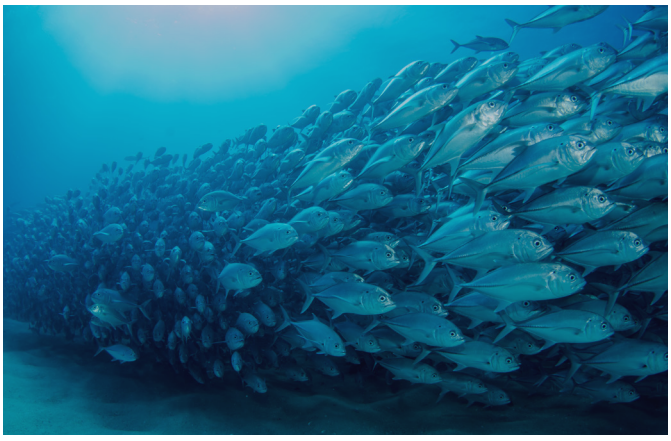

### Taps come in two categories: active and passive

- An active tap is a network element that monitors the traffic that passes through it and then duplicates and passes that data along to an endpoint such as an inspection or analysis device. Active taps can be used in fiber-cabling environments, but they are required when the transmission media is copper. Active taps, need to be powered at all times, regenerate the traffic to a connected "monitoring" tool while keeping traffic on the tapped network ports flowing through at all times (see above). In the event of a power outage, the network traffic flow is not affected, but the connected "monitoring" tool no longer receives a regenerated copy of the traffic. these taps may have battery backup to keep them up-and-running. Alternatively, a bypass tap is a special type of visibility layer device that can be configured as a bypass or as an active tap.
- A passive tap is an optical tool that 'splits' the light (the data traffic) passing through a fiber cable. It enables the data to flow through it, while at the same time, splitting that light, thus passing some of the light along to an endpoint such as an inspection or analysis device. Depending on the loss budget you can select the split ratio, i.e how much light is going to the monitoring appliance and how much goes through to the network. It cannot work in a copper cabling environment (this is also due to speed negotiation), and will only function with fiber-optic cabling. Passive taps have no need for any power source.

When deploying the taps, it is essential to match them to the network's physical characteristics and behavior, taking into account network speed capabilities and cabling types. For example, If the network cabling is made up of copper in a 1Gb network, it will need to use an active network tap (as previously explained, passive network taps will not work in environments made up of copper cabling – only fiber).

Network architects can use both passive and active network taps in fiber network infrastructure, but there are different requirements for singlemode and multimode implementations. Network taps are often built for specific 1Gb, 10Gb, 40Gb, and 100Gb environments. It is essential, therefore, to match the taps with the devices' capabilities. Taps further need to be properly configured and matched with the appropriate network connectors (LC, MPO, etc).

In addition, a company's network manager (or network architect) should take into consideration the judicious placing of the network taps at such points within the network, where the most important core business operations of the enterprise take place. A typical (and critical) tapping point is on traffic going "out" from the enterprise, or coming "in" from the outside (what is sometimes termed as north-south traffic). For example, the network taps should be positioned just before and just after the web servers, or in case of global presence, then it should include an access point within the communications network.

In addition, a company's network manager (or network architect) should take into consideration the judicious placing of the network taps at such points within the network, where the most important core business operations of the enterprise take place. A typical (and critical) tapping point is on traffic going "out" from the enterprise, or coming "in" from the outside (what is sometimes termed as north-south traffic). For example, the network taps should be positioned just before and just after the web servers, or in case of global presence, then it should include an access point within the communications network.

### Tap added value: Regeneration

Often, network traffic needs to be simultaneously tapped by multiple security devices and monitoring and inspection systems. Network taps that have regeneration capabilities to copy traffic and send it to multiple connected systems, can do this efficiently. Network engineers may install network taps with extra ports to ensure that new security or monitoring system that will be added in the future has access to the network traffic it needs. Additionally, this enables IT troubleshooters to quickly access data from different network links, to speed root cause analysis work.

## Bypass

A bypass switch (sometimes also referred to as a bypass tap), is a more evolved hardware than a network tap. It is used to connect a monitored network segment to an active, inline device (for example a security tool) and monitor that device's health. It eliminates points of failure by automatically bypassing traffic around an inline network security device, should that device become incapable of processing or passing the traffic, for example during a tool failure.

The bypass device contains a relay switch. If there is a power outage in a security device it is linked to, this will cause the relay to automatically close, and the switch will go to 'bypass mode'. This mode ensures continuity and will prevent an inline tool from bringing down part or all of the network, by having the bypass switch the data packets, circumventing the no longer responding inline device. At the same time that your network connection continues to remain 'alive and operative', you will be able to resolve the issue (replace, upgrade, carry out maintenance, or just disconnect the non-functioning device).

**Fail-open** - is the terminology used when configuring the network to continue to pass network traffic incase of inline appliance failure or in case of power failure in the bypass switch itself

**Fail-close** - is the terminology used when configuring the network to 'disconnect' the network traffic in case of inline appliance failure or power failure in the bypass switch itself. In certain network deployments the user may prefer to cut network traffic in case a critical security element fails so as to protect the network from vulnerabilities until the inline security device is running again.

## Bypass switches come in two categories:

- Passive Bypass - a passive bypass includes optical switch or copper switch and only supports inline deployment. In this use case traffic is coming from the network, through the inline appliance to the other connected network port. In case of power failure the passive bypass can be configured fail-open or fail-close. In some of the passive bypass devices there is an external port (often a USB port) through which the device receives a 'signal' on whether the inline appliance is active or inactive. This signal will also trigger the fail-close or fail-open configuration. Passive bypass are not entirely power-free in the strict sense of the word, as they need some power to enable configuration and setting of the optical or copper switch. As with network taps, the network manager needs to select and match the interfaces and speeds between the passive bypass, the network link and the connected inline appliance.
- Bypass Switch - is an active bypass switch that includes the same capabilities (and components) as a passive bypass, but in addition has a 'switching' fabric. The switching fabric endows the bypass switch additional functionality, and thus makes it useful for additional use case scenarios.

Some of the most important capabilities provided in a bypass switch over a passive bypass are: accurate heartbeat, interface speed and connectivity matching and last but not least bypass-tap configurations.

## Bypass Switch Heartbeat

In order to know what the state of a device connected to the bypass is – that is, to decide if it indeed needs to be bypassed – the network bypass switch generates a continuous series of bi-directional heartbeat (HB) packets, to monitor the health of the inline device.

As long as the HBs are returned, the traffic will continue to flow through the inline device. If the HB is not returned by the bypass switch, or in the event that the inline device loses power, is disconnected, or otherwise fails, the bypass switch will reroute the traffic directly between its other network devices, bypassing the out-of-service one. This will ensure that traffic continues to flow on the network link and data packets are not dropped.

## Interface and speed connectivity matching

In a bypass switch the traffic from the network ports is input to the switching fabric, and the traffic to the appliance ports is coming from the switching fabric. The appliance ports are typically connected via user selectable transceivers, so that the Network Manager can select the transceivers that he needs, for example connecting a network multimode interface to a copper/RJ45 connected appliance.

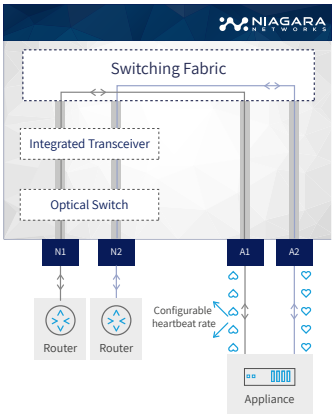


(a)

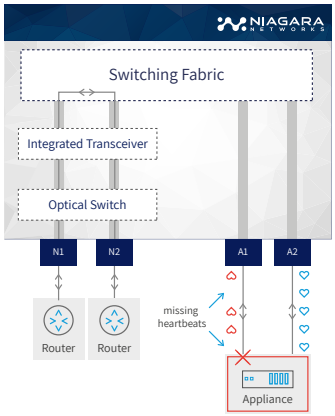


(b)

The figure (a) depicts user-configurable heartbeat packets transmitted on the appliance ports. In the event of an appliance malfunction (such as a software crash, system failure or loss of power depicted in (b), the failure is detected, and the traffic intended for the inline appliance is bypassed to the network ports, allowing it to continue to flow through the network link. Once the appliance is back up, or the power is restored to the appliance, it is detected by the heartbeat mechanism, and network traffic is seamlessly diverted back to the inline appliance, allowing it to resume its critical functions.

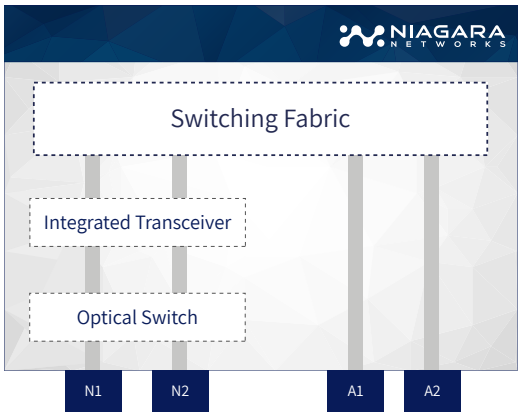


A full bypass switch segment comprises two network and two appliance ports. The network ports offer direct single mode (SM) or multimode (MM) connectivity. The appliance ports utilize customer pluggable transceivers. The network ports connect to the non-blocking switching fabric via integrated transceivers.

## Bypass-tap

In a bypass switch, because all the traffic from the network ports and appliance ports are connected to the switching fabric, and because the switching fabric allows the user to configure different connectivity paths between inputs and outputs; it is possible to configure a bypass switch segment to act as an active tap.

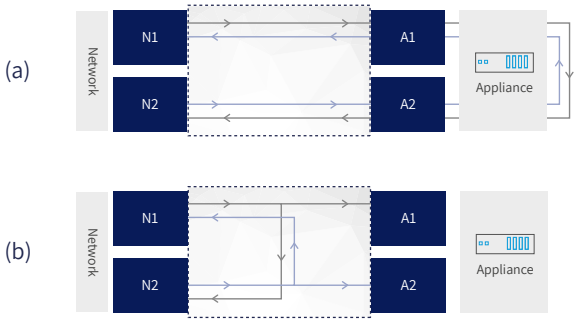


This figure depicts two common configurations of a bypass switch segment (a) inline configuration (b) active tap configuration

# Network packet brokers (NPBs)

(NPBs) provide a collection of functions and features that enable monitoring and security tools access to traffic, from single to multiple links across the network. Network Packet Brokers role has been evolving over the past several years, but for a good foundation on Network Visibility, we will dedicate this section to the "plain vanilla" Network Packet Broker and build on that to more complex deployments and features.

The NPB is a hardware device or appliance that receives network traffic sources on input ports and has various network tools and appliances such as monitoring tools, performance management, and security tools connected on output ports. The NPB has a switching fabric which is is used to configure and map paths between the 'input' traffic sources and the 'output' tools. The following three basic NPB features will better help in understanding its role. These are:

### Filtering

Denotes the capability to apply a filter on the traffic path, so that the performance management, monitoring or security tool only receive the traffic that they need. Filtering is defined by a rule, where the user is able to configure any field in the header in (OSI) Layers 2-4.

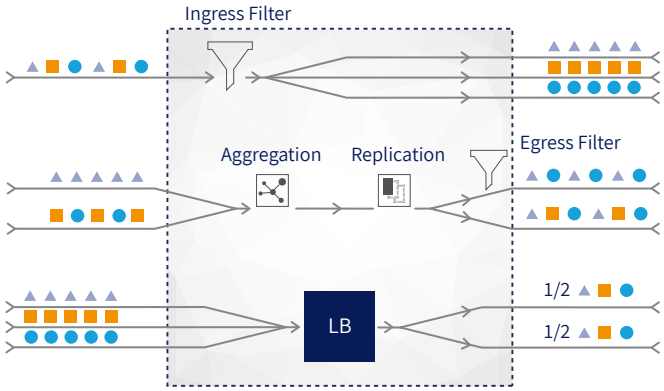


### Aggregation / Replication

Denotes the capability to forward traffic from any port to any port in a one to many and many to one configuration. Some of the use cases of these features would be to aggregate multiple low traffic capacity input ports to a single appliance port for an efficient utilization of a network tool or to replicate input traffic to multiple tools so that the traffic could be inspected and handled by different tools and departments.

### Load Balancing

Denotes the ability to load balance input traffic between multiple network tools. A use case for load balancing would be the ability to connect multiple network devices, each with lower processing capacity, and enable them to process a higher speed link by load balancing the high speed link between the connected network devices.

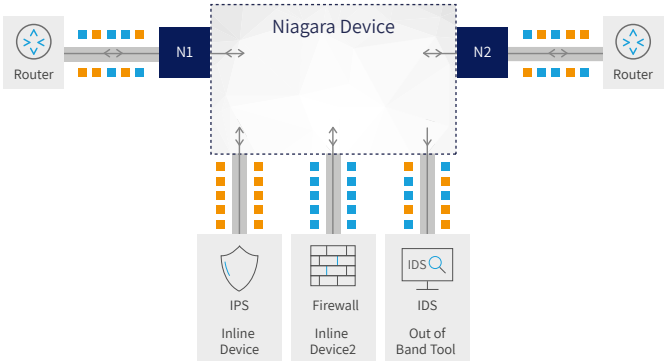

More evolved NPB feature include different level of packet header manipulation. These may include, for example, VLAN stripping and tagging, or MPLS stripping, or stripping in general of headers, to increase the efficiency of the connected tools. Often network tools such as monitoring, performance management and security tools are designed to receive packets for processing only in a specific format. Another use case for VLAN tagging is to use the NPB to "color" the traffic so that the connected network tool can apply different traffic processing to different traffic profiles.

We can see from these examples and feature functionality that the main purpose of the NPB is efficient processing, deployment and utilization of Network devices and tools.

The NPB sees the network traffic via links coming from SPAN ports or Network Taps. These are connected to the packet broker's port. This leads us to the next section in understanding combined deployments of Network Packet Brokers and other Network Visibility Layer building blocks

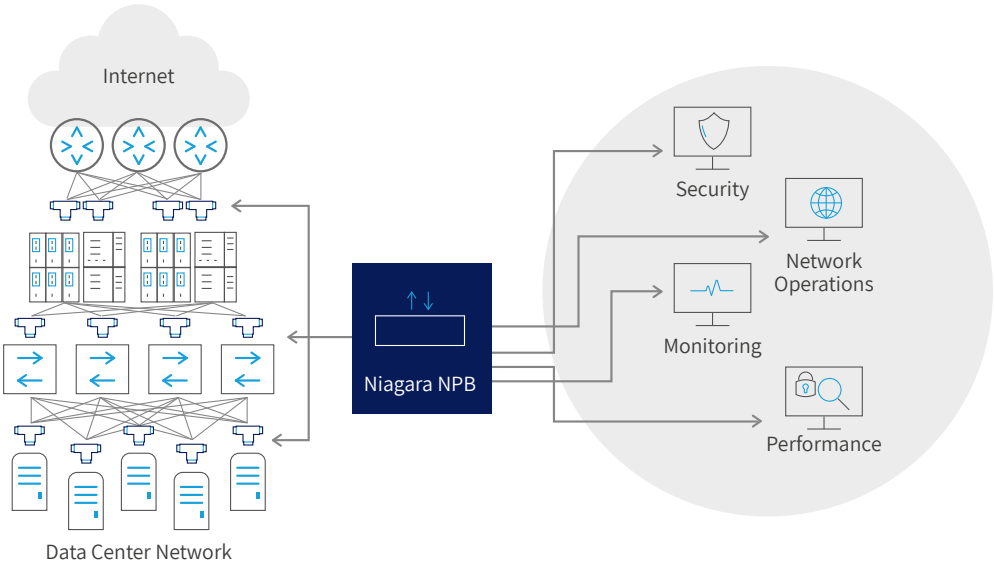


## Combining network packet brokers and network taps

Packet brokers, like the taps, are hardware devices. The coupling of the two (NPBs and taps), can provide an additional, more sophisticated, level of functionality that will ensure enhanced visibility across all the devices that are connected throughout the network. When the packet broker is placed between the network tap and the intended security or monitoring solutions, it can route the traffic in an intelligent manner, by using different port mapping schemes.

The combination of network taps and packet brokers, enables network architects to design a complete connectivity solution that includes the ability to regenerate, aggregate, filter, and load balance data. This will ensure that each and every element in a network monitoring and security architecture has maximum and optimum visibility to carry out its multiple functions and tasks in the most efficient manner.

## Combining network packet brokers and network bypass switch

Inline bypass protects network traffic from failure of inline tools such as Intrusion Protection Systems (IPS) – i.e. if traffic begins to show signs of a failure, then physical or logical bypass can properly divert the traffic to keep it flowing smoothly.

There is a new breed of network packet brokers that provide an all-in-one connectivity solution incorporating that bypass functionality. It combines built-in packet broker capabilities that (as previously described) are often an asset in enhancing various visibility functions. However, instead of deploying multiple solutions with many different functions, the new combined network bypass-broker already incorporate functionality that supports filtering, aggregating, regeneration, load balancing and bypass all in a single device.

Hybrid options enable network designers and administrators the flexibility that is required to provide 100% network visibility to the wide range of security

and monitoring solutions that are necessary to run and maintain the network and its many and varied components on a day-to-day basis.

## The Modular Multi Purpose NPB

Network packet brokers are taking an evolving role and by combining them with other devices such as taps and bypasses, they assist in greatly enhancing overall network visibility.

Modular NPBs may offer support for both passive and active tap modules and bypass modules of different speeds and feeds. As a result, rather than having distinct taps and NPBs in the network, network architects can deploy a modular NPB and add tap modules as needed. This simplifies wiring, total cost of ownership and simplifies network design issues.

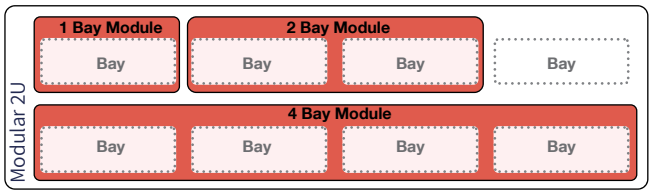


Each row in the chassis has 4 single width bays. Users can dynamically hotswap different bay-width modules to mix and match module functionality and and module bay width size according to their needs. The double width bay module can dynamically fit into the space of any two adjacent single width bay locations as depicted above.

## New directions in NPBs

As networks continue to evolve and the processing needs and functionalities of Network devices increase, Network Architects are finding that Network Visibility Layer in general and the NPB in particular are good locations to off-load sophisticated packet processing.

Some of the more advanced and up-and-coming capabilities incorporated into next generation NPBs include identifying known, suspicious, and unknown traffic passing through the network, providing SSL decryption capabilities and more.

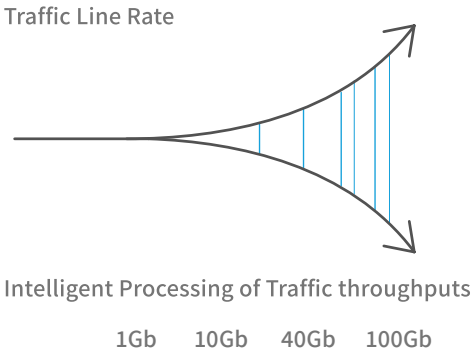


The above diagram represents how next generation network packet brokers are well suited to address the growing gap in providing effective application layer processing, as network traffic throughput is increasing.

# Software Defined Network (SDN) role in Network Visibility

SDN, operating with the OpenFlow protocol for remote communication with network plane elements (for determining the path of network packets across network switches) is a more recent mechanism associated with network management. This trend is also impacting the network visibility layer. Making the Visibility Layer devices such as the network taps, network bypass switches and network packet brokers, SDN-aware has multiple advantages. It turns an enterprise deployment into a virtual 'single pane of glass' switching fabric where all the building blocks can be connected to forward the right traffic to the right tool across your entire network. It empowers network management and facilitates efficient network configuration in order to improve the performance and reliability of network monitoring, performance and security tools.

SDN aware Visibility Layer is not only crucial for large scale networks, but is also the harbinger of dynamic visibility layer. In the dynamic visibility layer paradigm the traditional semi-static mapping configurations between 'sources' and 'destinations' are responsive and can automatically adapt the visibility layer to evolving traffic and network conditions.

# Virtual taps

The need for traffic-based visibility, along with increased correlation of traffic and increased controller states will become an integral part of SDN deployments. These functions are naturally implemented within the Visibility Adaptation Network. Virtual Taps are evolving from implementations of network virtualization (NFV) where a visibility device cannot always be readily deployed in a virtualized environment.

## Summary

A pervasive visibility layer not only increases the security of your network, but also holds key advantages in reducing downtime in maintenance periods, improving network service recovery time and increasing your overall ROI.

End-to-end network visibility is attained by companies understanding the need of their network tools, their network topology and how to connect the two together - how to bring the right network traffic to the right tool in the most efficient manner. This requires strategically placing Visibility Layer devices throughout the network design, eliminating blind spots and the number of visibility gaps to create a connected visibility layer that covers all possible network access points – both physical and virtual.

The right combination of network taps, NPBs, and bypasses, produces a network visibility solution that offers a truly versatile (smart, useable, and manageable) means to secure your enterprise network.

# About Niagara

Niagara Networks provides high performance network visibility solutions to allow seamless administration of security solutions, performance management and network monitoring. Niagara Networks products provide advantages in terms of network operation expenses, downtime, and total cost of ownership.

A former division of Interface Masters, Niagara Networks provides all the building blocks for an advanced Visibility Adaptation Layer at all data rates up to 100Gb, including taps, bypass elements, packet brokers and a unified management layer. Thanks to its integrated in-house capabilities and tailor-made development cycle, Niagara Networks are agile in responding to market trends and in meeting the customized needs of service providers, enterprise, data centers, and government agencies.