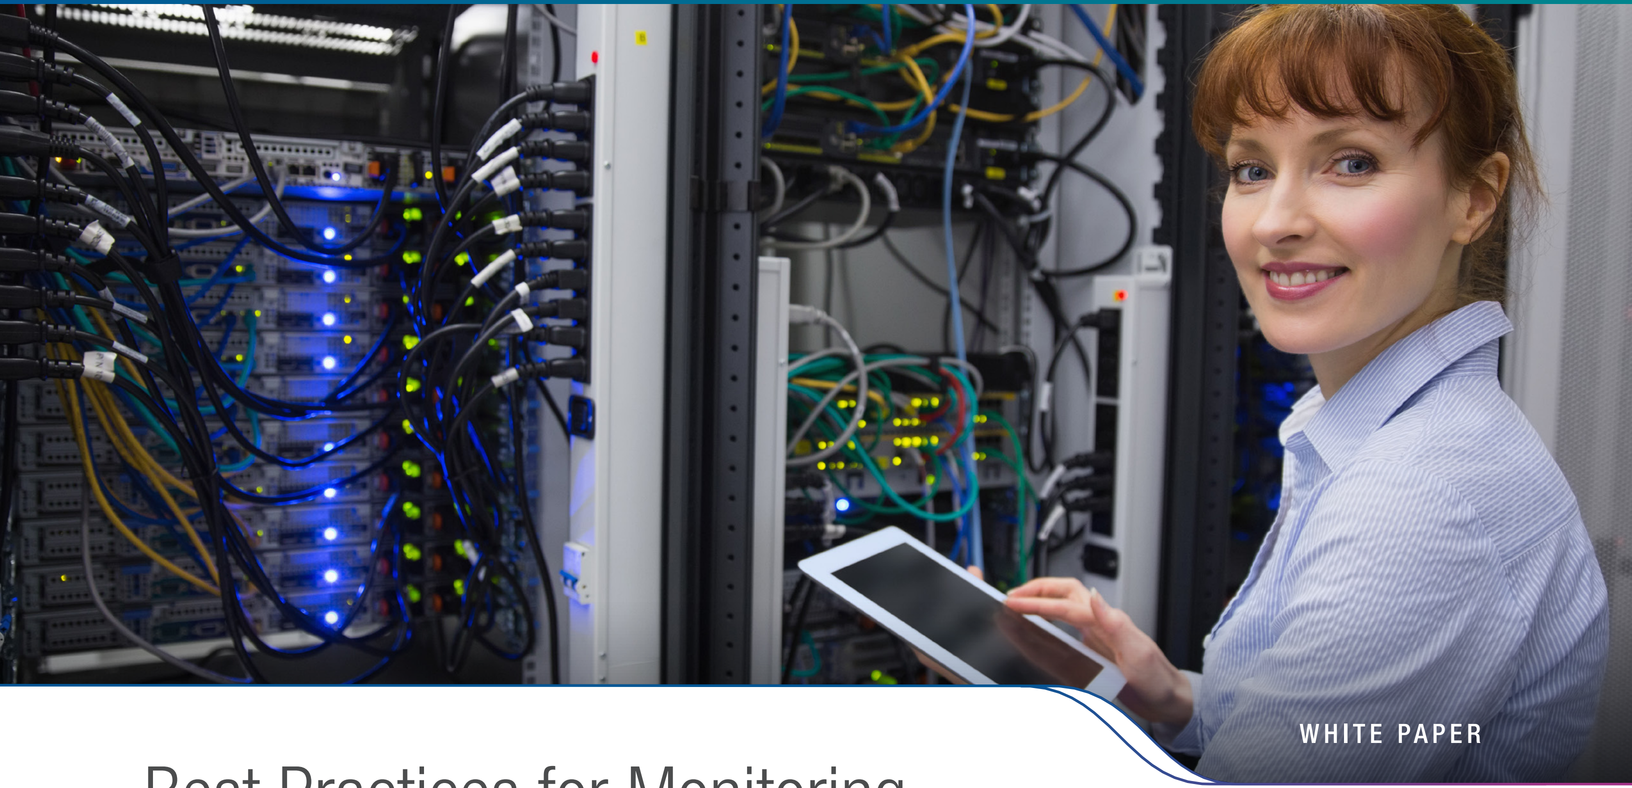

WHITE PAPER

Best Practices for Monitoring Distributed Service Provider Networks

Network modernization is needed to create a coherent strategy to acquire, aggregate, and process the large volumes of monitoring data that exist.

Monitoring Challenges for Service Providers

Information technology (IT) has become a strategic part of the world's most successful companies. Service providers are capitalizing on this valuable technology as well. This includes modernizing their existing network infrastructure by using Internet Protocol (IP) and Ethernet to reduce costs and improve security. The re-architecting of the central office (CORD) is another initiative driving service providers as they try to unify software-defined networking (SDN), network function virtualization (NFV), and cloud technologies to increase business agility, improve the user quality of experience (QoE), and offer new services and capabilities.

Part of the service provider strategy needs to include the modernization, or creation if it doesn't exist, of the service provider's network visibility architecture. This modernization is needed to create a coherent strategy to acquire, aggregate, and process the large volumes of monitoring data that exist. This is especially important with all of the different public and private initiatives happening: like SDN and cloud-based networks, Internet of Things (IoT), Openstack forum recommendations, next generation mobile networks, CORD, and network security modernization.

KEYSIGHT TECHNOLOGIES

A visibility architecture with network packet brokers (NPBs) can be used effectively to:

- Aggregate data from multiple sources (SPAN ports, taps, virtual machines)
- Lower the cost of data aggregation and backhaul
- Optimize security and monitoring tool performance
- Improve security and compliance initiatives

Service Provider Architecture Overview

The service provider network, whether it is a wireline or mobile network, typically has two distinct regions – the network core and distributed/edge locations. As depicted in the following diagram, network monitoring is required for both areas to create optimum visibility:

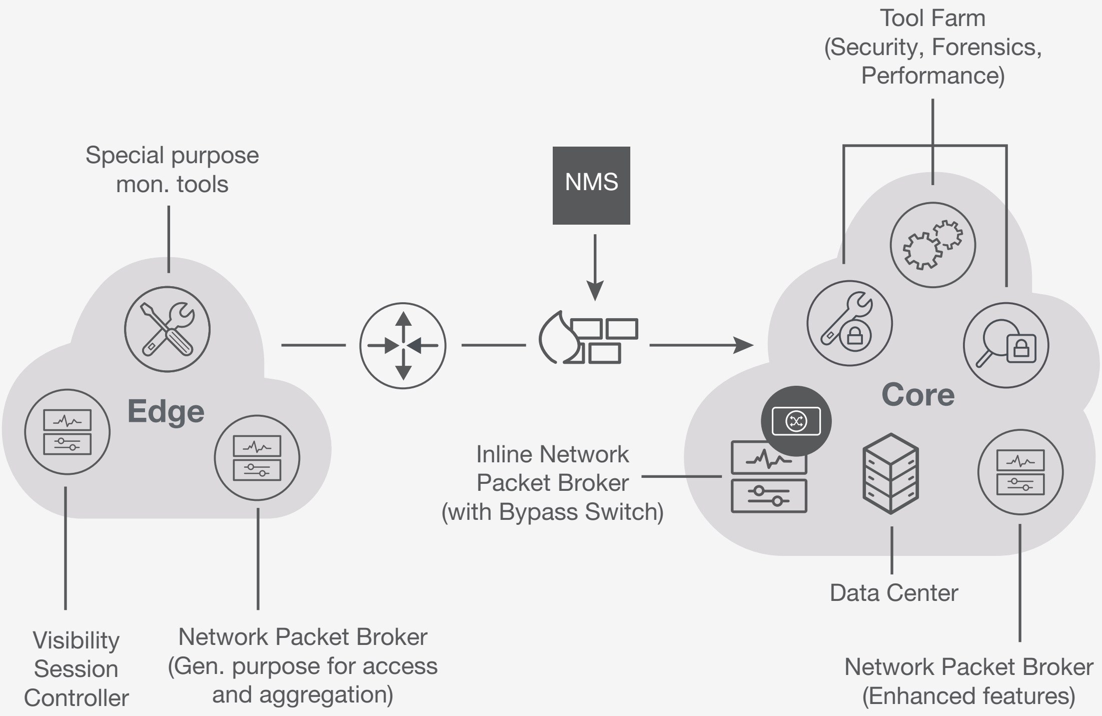
Tool Farm
(Security, Forensics, Performance)
Special purpose mon. tools
NMS
Edge
Core
Inline Network Packet Broker (with Bypass Switch)
Data Center
Visibility Session Controller
Network Packet Broker (Gen. purpose for access and aggregation)
Network Packet Broker (Enhanced features)

Figure 1. General Visibility for Service Providers

The exact type of data monitoring functionality needed depends upon many factors including: whether the monitoring functions take place at the edge or in the core, the type of service provider network (wireline or mobile), and any particular objectives or needs that the service provider has. This paper focuses on visibility solutions for the Distributed (edge) portion of a service provider's network.

Distributed Service Provider Visibility Architectures

The need for greater packet manipulation and data reduction capabilities at the edge of the network is driving the need to push network visibility solutions to the edge as well. This is especially important as data bandwidth needs explode. The monitoring data at the network edge needs to be "processed" first to eliminate unnecessary content before it is forwarded to the tools. Another problem arises as the speed of the network backbone increases faster than the speed of security and monitoring tools. Common examples of out-of-band monitoring activities that can be performed at the edge include: data aggregation, data distribution, packet filtering, mobile user session correlation, deduplication, header stripping, packet slicing, load balancing, and specialized monitoring functions (troubleshooting, suspicious data analysis, specific data storage, QoE investigation and validation, etc.)

Visibility activities that are performed at the edge include: data aggregation, data distribution, packet filtering, packet correlation, deduplication, header stripping, packet trimming, and load balancing.

The general benefits from implementing a visibility architecture at the edge of the service provider network are numerous and include the following:

- Ability to monitor services specific to the edge, for example eNodeB handoffs, voice quality inspection, etc.
- Significant cost reductions for data monitoring
- Additional monitoring data storage options (e.g. monitoring data can be captured and sent straight to a storage area network (SAN) for a delayed analysis in the future by the tool farm in the core)
- Distributed failure and better uptime statistics, since capabilities are distributed across the network with fail-over capabilities possible
- Faster ability to react to business needs since capabilities already exist at the edge (e.g. troubleshooting, quality of service investigations, etc.)

Distributed Visibility for Wireless Service Providers

Wireless service providers have specific needs for distributed network visibility. This is especially true as new technologies, like 4G and 5G, increase the reliance on low latency networks and place higher bandwidth requirements on the network. Some specific wireless service provider use cases for network visibility include:

- Using an NPB to control data monitoring costs with filtering, aggregation, and packet manipulation

- Capturing and correlating customer data with a visibility session controller for GTP session data and then distributing that data to special purpose tools for analysis of customer data (dropped calls, quality of service (QoS) issues, etc.)
- Capturing and filtering data for special purpose remote tools (e.g. data loss prevention (DLP), application performance monitoring (APM) and network performance monitoring (NPM), voice over Internet protocol (VoIP) QoE analysis, and troubleshooting)
- Capturing and routing monitoring data to a SAN for future data analysis
- Support for automation response capabilities to commands coming from network management systems (NMS) for configuration control and accuracy
- Bandwidth throttling to control network loading

The NPB sits between network SPANs/taps and the monitoring tools where it can capture the data packets and manipulate them as needed.

Figure 2 shows where visibility equipment (including NPBs, visibility session controllers, and monitoring tools) fit within a typical wireless network both in the evolved packet core (EPC), which is really at the edge of the network and not the core, and then the IP multimedia subsystem (IMS), which is located within the network core.

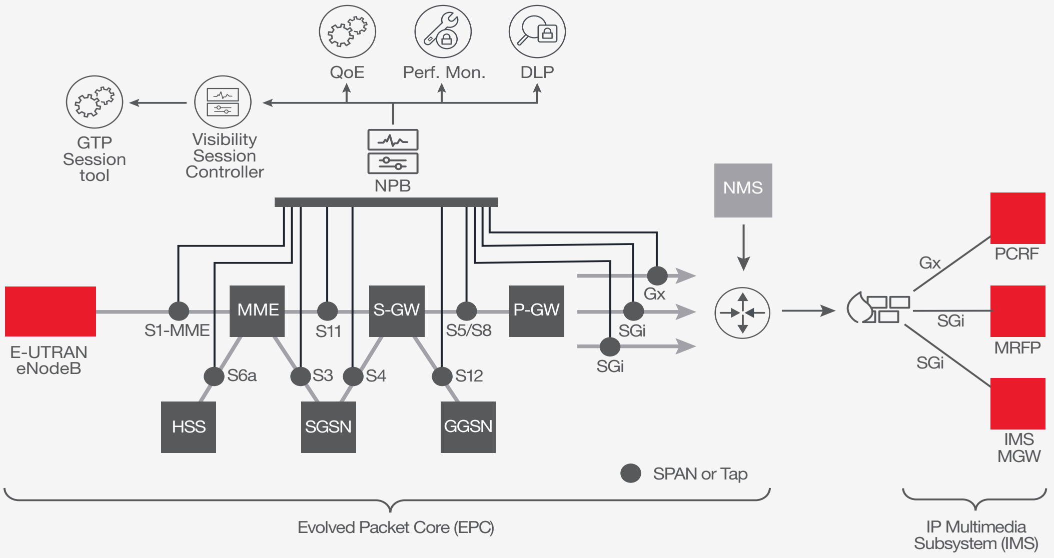

Figure 2. Visibility Solution for Wireless Service Provider Edge Nodes.

Mirror ports (also known as SPAN ports) on network equipment and network taps create data access points that provide network data to the NPB. While network access points could be connected directly to the tools, the number of data access ports on security and monitoring tools is limited. In addition, the data needs to be streamlined so that only the key monitoring data is sent to the core. The NPB solves both of these types of data aggregation problems. The NPB sits between network SPANs/taps and the monitoring tools where it can filter and forward data packets as needed.

Distributed Visibility for Wireline Service Providers

Just like wireless service providers, wireline service providers have specific needs for distributed network visibility as well. There are actually similar, but different, scenarios in existence. One scenario is the legacy central office connected to the telecom network core. The second scenario is the relatively new Central Office re-architected as a datacenter (CORD) architecture. Some of the use cases for these two architectures include:

- Using an NPB to control data backhaul costs from the central office (CO) Ethernet aggregator to the core with filtering, aggregation, and packet manipulation
- Capturing and filtering data from PODs for special purpose tools (e.g. DLP, APM and NPM, VoIP QoE analysis, and troubleshooting) located at the CO
- Support for automation response capabilities to commands coming from an NMS, including business support systems (BSS) and operations support systems (OSS), for configuration and policy control
- Data aggregation of CO servers by NPBs that connects to the leaf/spine switch architecture and then transmits the monitoring data to the core

An NPB can be used to control data backhaul costs from the central office Ethernet aggregator to the network core by using filtering, aggregation, and packet manipulation.

Figure 3 shows where visibility equipment (including NPBs and monitoring tools) fit in a conventional CO-based architecture.

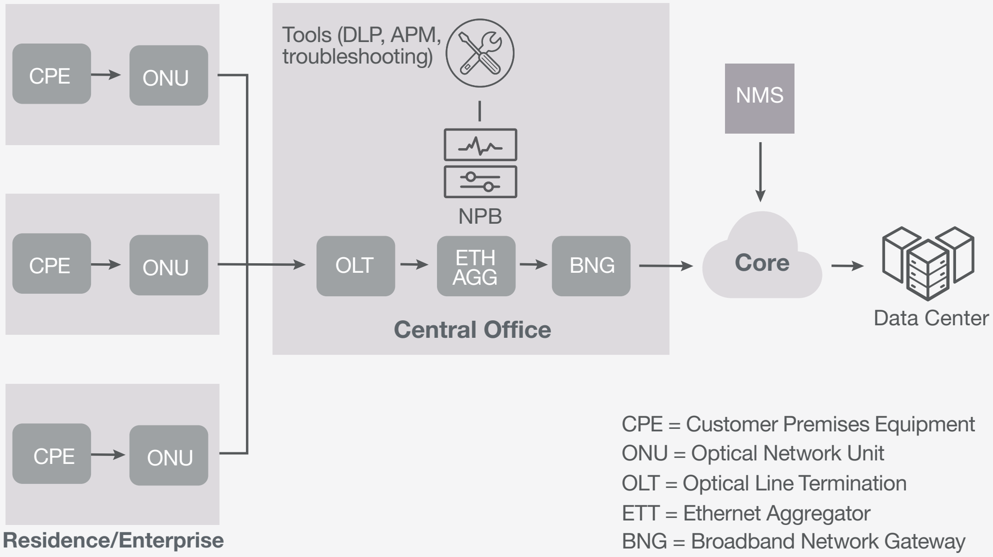

CPE = Customer Premises Equipment
ONU = Optical Network Unit
OLT = Optical Line Termination
ETT = Ethernet Aggregator
BNG = Broadband Network Gateway

Figure 3. Visibility for Wireline Service Provider Traditional Central Office.

New architectures, like CORD, are transforming service provider networks to give them the economic efficiencies of the data center combined with the agility of a cloud provider that allows them to replace aging hardware with software. Figure 4 shows where visibility equipment (including NPBs and monitoring tools) fits in a CORD-based architecture.

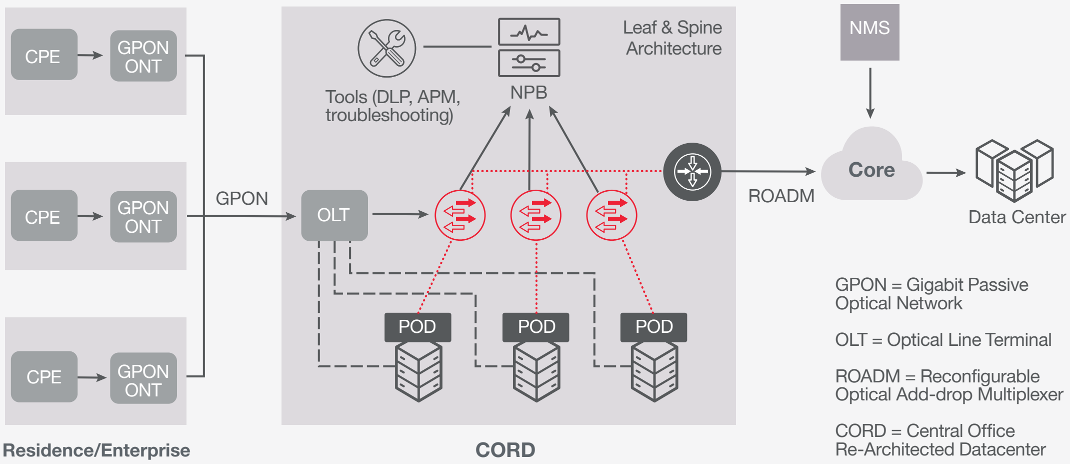

GPON = Gigabit Passive Optical Network

OLT = Optical Line Terminal

ROADM = Reconfigurable Optical Add-drop Multiplexer

CORD = Central Office Re-Architected Datacenter

Figure 4. Visibility for Wireline Service Provider CORD.

Optimizing the Monitoring Architecture with an NPB

Service providers make huge investments in edge network infrastructure, core network infrastructure, monitoring tools, and policy management systems. As a result, it is essential that service providers get the most out of their investments.

Network packet brokers are a key piece of functionality to help service providers optimize their visibility architecture and maximize the return on their investment. The following list shows the most common out-of-band NPB features used by service providers:

- Data Filtering
- Aggregation
- Regeneration
- Deduplication
- Load balancing
- Filter libraries
- Header striping
- Packet slicing
- Media speed conversion
- Automation
- NetFlow generation
- Visibility session correlation

The following sections detail the specific features.

A benefit of removing duplicate packets is that the bandwidth of the network can be conserved, allowing service providers to postpone costly network upgrades.

Data filtering

Once the data is captured and fed to the NPB, filtering can be used to remove non-relevant monitoring data. These filters can be created based upon various criteria like IP address, VLAN, port number, or application signature.

When considering an NPB, it is important to understand its filtering capabilities. Filtering is usually performed in three stages. The first stage is performed at the port where the network data enters the NPB. This is called the network port. Filtering at this location permanently removes the data for all operations downline. Once this traffic is removed, it is no longer available for analysis. So filtering at this location should be used judiciously.

The second stage of filtering is performed by a highly capable, port-independent filter that is located between the network port and the ports to which the monitoring tools, or backhaul connections, are attached. The port-independent filter, also called a dynamic filter engine, is the ideal place to perform the bulk of the filtering, as it is possible to understand exactly what is happening by looking at this location. Multiple filters can be created to effectively transfer the same data to multiple tools without clipping the data, assuming the filters are created correctly. This is what makes this location the perfect point to conduct the majority of data filtering.

The third stage of filtering is performed at the exit point of the NPB, which is called a tool port. This location does an effective job of removing unnecessary traffic on a per tool basis. However, filtering at this location can cause two problems. First, the tool port can be overrun by the volume of traffic coming from the network ports. Second, the interaction between the network filters and the tool filter is complex and not obvious unless you are well versed in set theory.

Libraries of filter definitions can be saved, allowing IT teams to create common filter definitions and disseminate these libraries for use among the teams to improve productivity.

Aggregation

Network packet brokers allow the aggregation of multiple ingress links to one egress port. This allows IT to connect multiple network segments to an NPB and then send the relevant data out as a single stream to security and monitoring tools.

Regeneration

Regeneration is the feature that allows the NPB to take data that is coming in and make a copy of that specific data so that it can be sent to multiple outputs, i.e. tools. One use case is that the same data may be needed by multiple tools (like a troubleshooting tool and an NPM tool). Another use case would be to send the data to a local tool for analysis but also send a copy to a designated SAN for additional reviews at a later time.

Packet de-duplication

Duplicate packets are most commonly caused by the use of SPAN/mirror ports. While some monitoring tools are capable of removing duplicate packets, others do not have this capability. However, even if a monitoring tool has the ability to remove duplicate packets, doing so is an extremely resource intensive task. Off-loading duplicate packet removal to the NPB can cut the CPU load of a monitoring tool in half.

Another key benefit of removing duplicate packets is that the bandwidth of the network can be conserved. This reduces the cost of data backhaul, allowing service providers to postpone costly network upgrades.

Load balancing

As the volume of data grows within a service provider's network, IT teams often find that the network data flow is increasing faster than the capabilities of their monitoring tools. A single monitoring tool that previously performed well may now be out of capacity. Through load balancing, some NPBs have the ability to send data across multiple, similar tools and do it in a way that consistently sends all the data from a particular session to the same monitoring tool. The load balancing feature keeps session data together for better analysis, yet balances the total network load across multiple monitoring tools. This function is used extensively with network data recorders. Since session data is kept together, only one data recorder needs to be accessed to analyze any given session at a later time.

A second use case utilizes the load balancing feature to send data across multiple, redundant links within the network. Should one path fail, the NPB can send all data across one link, instead of splitting the data across two links, as during normal operations.

Monitoring tools, network management systems, and IT automation systems can dynamically change a NPB. Using a RESTful interface, software-based systems can control various aspects of the NPB.

Filter libraries

Another time-saving feature of a NPB is the import and export of configuration information with granular control over what gets saved or loaded. Libraries of filter definitions can also be saved, allowing IT teams to create common filter definitions and disseminate these libraries for use among the teams to improve productivity.

Configurations can also be pushed down from an NMS. This includes running scripts to program filters and creating filter libraries on the NPB. The automation of the filters reduces errors and makes filter creation faster. Scripts can be tested first in a lab environment using the Keysight NPB simulator before rollout to a live NPB.

Header stripping (E.G. MPLS stripping)

Removing MPLS labels increases the capability of many monitoring tools. This is because most monitoring tools are not capable of understanding MPLS-tagged packets, which means that they are unable to monitor MPLS networks. An NPB is used to remove the MPLS headers and forward the original packet contained within the MPLS tagged packet. Standard network monitoring tools can then be used to monitor activities with the MPLS network.

Packet trimming/slicing

Packet trimming removes payload data from the packet, leaving the header information, prior to sending the packet to monitoring tools. Some monitoring tools do not require packet payload information, in which case, removing payload data allows more data to be sent across the link from the NPB to the monitoring tool. As a result, the monitoring tool can receive a far greater amount of network data. Additionally, for compliance reasons, it may be desirable to "trim" or remove sensitive payload data from packets before they are sent to a monitoring tool.

If the data is being backhauled, removing the payload also significantly decreases the size of the bandwidth required for the data.

A GTP session controller can be used to effectively identify and track mobile subscribers. At the same time, it can correlate data from network probes which can be used to load balance bandwidth to enforce capacity and rate limits for each customer, even as mobile traffic rates fluctuate.

Media speed conversion

With network backbone speeds increasing to 40 and 100 Gbps, it is increasingly common that some network equipment cannot handle this high of a data rate. Upgrading a network is expensive, and becomes even more costly when network monitoring tools must be upgraded at the same time. Traditionally, network engineers did not have a choice. When they upgraded their network they also had to upgrade monitoring tools. With the advent of the NPB, network engineers now have a choice. The NPB can be used to "downshift" the speed of network data to match the speed of the available monitoring tool. Network engineers now have the ability to monitor higher-speed backbone networks with their lower-speed tools, protecting their original monitoring tool investment.

Automation

Automation is a productivity enhancing option available with some NPBs. Monitoring tools, network management systems, and IT automation systems can now issue instructions to an NPB using a RESTful interface. This allows IT teams to create extremely powerful systems using multiple network devices that work cooperatively and automatically.

A third use case involves an intrusion detection system (IDS) system that detects an intrusion as it occurs. Using the automation capabilities of the NPB, the IDS sends a command that sets up a connection between the network port being monitored by the IDS and a network data recorder, immediately capturing the intrusion event for later analysis.

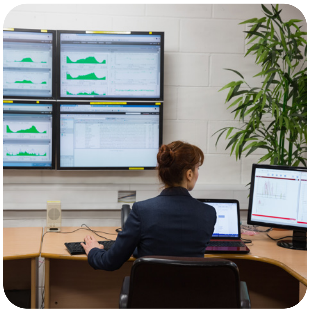

VLAN tagging

An NPB can be used to both generate and remove VLAN headers. In the case of adding the header, this can be especially useful for identifying where a packet came from, even after it has been aggregated with other traffic. This aids in troubleshooting and forensic analysis as the NPB can filter data based on a VLAN tag. The feature is also useful if you are using SPAN ports, as VLAN tags are not normally passed through a SPAN port. On the other hand, certain monitoring tools do not understand VLAN headers, so an NPB can be used to remove them in those situations.

SDN, cloud-based networks, and the Openstack forum are just some of the drivers for creating new capabilities to cost-effectively and competitively scale data centers.

NetFlow generation

Certain NPBs can generate NetFlow data and send that data to collectors on the network. This flow data can be used for multiple purposes including: network segment bandwidth overload identification, application bandwidth overload identification, application identification and filtering, problem geolocation, and performance trending. Once this NetFlow data is created, it can be sent to a dashboard or to special purpose tools that analyze NetFlow data.

Visibility session correlation

Service providers (especially wireless service providers) need good customer call data (e.g. service holes, malfunctioning radios, poor coverage, and even customer dissatisfaction) to properly plan their networks and deliver a better quality of experience. This effort leads to higher levels of service assurance and increased revenue for mobile carriers. An important element in this process is the use of sophisticated and costly network monitoring probes that allow mobile carriers to immediately detect and resolve issues that impact Quality of Experience (QoE).

While network probes can provide visibility into wireless networks, these devices have limited capacity and may not withstand fluctuating mobile subscriber traffic. At the same time, under-loading network probes can create additional costs for the carrier. A GTP session controller can be used to effectively identify and track mobile subscribers. If the controller detects faulty or overloaded monitoring probes, it automatically redistributes the load to other probes in the cluster. As a result, monitoring probes are able to focus on QoE analysis, rather than spend cycles trying to reassemble GTP session traffic.

Preparing for the Hyperscale and Microscale Evolution

SDN, cloud-based networks, and the Openstack forum are just some of the drivers for creating new capabilities to cost-effectively and competitively scale data centers. One of these new techniques is to focus parts of the network towards hyperscale and microscale clusters. Hyperscale refers to the ability to use virtualized compute resources to scale massively as needed while Microscale refers to 4 or less servers grouped together to form a purpose-built cluster.

Hyperscale datacenters are designed to bring fast scalability to three common datacenter areas: compute, storage, and networking. As service provider initiatives like CORD take hold, coupled with general issues like the meteoric expansion of IoT and Internet-based applications, hyperscale datacenter designs will be needed by service providers to offset these technology drivers.

Microscale clusters are becoming popular at the edge of mobile networks because they can be used to focus on wireless connection latency issues, processing QoE data for rapid cause determination, and increased uptime (due to distributed capabilities).

An easy way to prepare for these new types of deployments is to create the Visibility Architecture mentioned previously. This architecture can help you set up a coherent visibility solution that can take advantage of the new trends while leveraging existing capabilities to leverage economies of scale in existing networks.

Summary

Service providers are under ever-increasing pressure to modernize their capabilities to be competitive in the 21st century. This means reducing costs, improving network performance, maintaining a high level of security, and being agile enough to offer new services to customers. The right visibility solution is a key ingredient as to how well the service provider can execute on their plans.

The NPB is a necessary tool and solves many of the data monitoring issues that service providers encounter within their network including the need for:

- Cost effective data aggregation from multiple remote data points
- Removal of unnecessary data and data size reduction before transmission across the network
- Improving overall monitoring tool performance
- Providing segmented data access for special purpose tools

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