Achieving Network Visibility
5 Key Requirements

cpacket NETWORKS

VISIBILITY YOU CAN TRUST

Who Should Read?

You can benefit from reading this eBook if your role is:

- IT Infrastructure and Operations (I&O) Leader
- IT Instrumentation
- Network Operations (NetOps)
- Application and Development Operations (AppOps/DevOps)
- Security Operations (SecOps)
- Cloud Operations (CloudOps)
- Site Reliability Engineer (SRE)
- Managed Service Provider (MSP)
- System Integrator (SI)

About the Author

Nadeem Zahid serves as Vice President Product Management & Marketing at cPacket Networks. He has spent 24 years in the IT networking industry strategizing, building and marketing hardware, software and cloud products.

Prior to cPacket, Nadeem has held several leadership positions in strategy, product management, engineering, marketing, and business development with companies like Alcatel-Lucent (now Nokia), Cisco Systems, Foundry Networks/Brocade (now Broadcom), Juniper Networks, Extreme Networks, LiveAction, and tFinery.

Nadeem is a thought leader and a published author of several articles, blogs and papers and has also published a book on Product Management. He holds a Master of Science in Technology Management from Boston University, a Bachelor of Electronics Engineering from N.E.D University of Engineering & Technology, and a Product Management certification from M.I.T. Additionally he is an ex Cisco Certified Internetwork Expert (CCIE Routing & Switching No, 15793).

About cPacket Networks

cPacket enables IT Infrastructure and Operations (I&O) through **network-aware application performance and security assurance** across the distributed hybrid environment. Our single-pane-of-glass AIOps-ready analytics provide the deep network visibility required for today's complex IT environment. With cPacket, you can efficiently manage, secure, and future-proof your network - enabling digital transformation.

Based in Silicon Valley, CA, cPacket enables organizations around the world to keep their business running. Leading enterprises, service providers, and governments rely on cPacket solutions for improved agility, higher performance, and greater efficiency. Our cutting-edge technology enables network, application, and security teams to proactively identify issues in real-time before negatively impacting the business.

cPacket inspects the network traffic to the packet level detail delivering the high-quality, high-resolution, and reliable data to the tools, enabling 20/20 visibility. cPacket solutions are **fully reliable**, **tightly integrated,** and **consistently simple** across the distributed hybrid environment: branch, data center, and cloud. Whether you need greater network visibility for security or performance monitoring, our solutions are designed to overcome scalability issues and reduce troubleshooting times. The result: increased security, reduced complexity, and increased operational efficiency.

Learn more at www.cpacket.com, the cPacket blog, or follow us on LinkedIn, Twitter, Facebook, YouTube, and BrightTalk. You can also Talk to an Expert, Contact Sales, or Request a Demo.

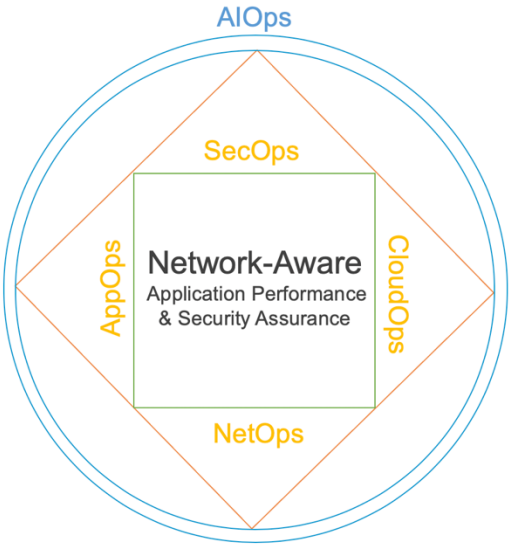

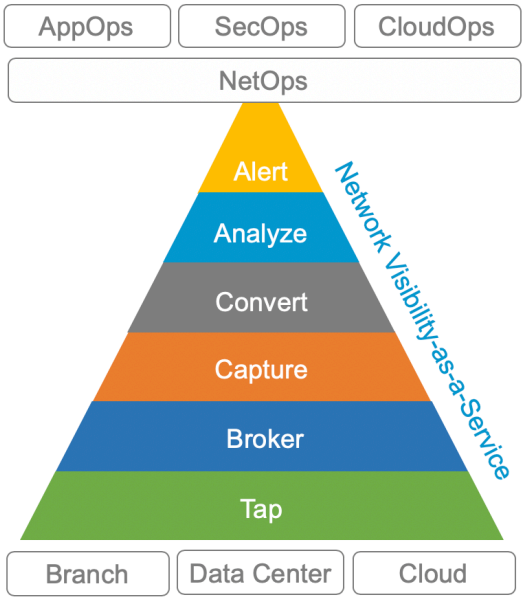

Table of Contents

You Cannot Control

What You Cannot See

INTRODUCTION: Network Visibility – More Important Than Ever

Enterprise and service provider networks have become increasingly complex: virtualization, hybrid-cloud architectures, proliferation of connected devices, higher port speeds, and densities, and an ever-expanding cyber threat landscape all make it challenging to ensure business continuity. Amidst all of this change, the network plays an increasingly strategic role, as it connects users with services. That's why **network visibility** – the ability to understand what's happening with the networks – has become so important.

Network visibility encompasses traffic patterns and trends, connected device types, end-user experience assurance, application responsiveness, performance measurements, and suspicious activities. It provides valuable details about network events – what happened, what is likely to happen, when, why, and where. Advanced network analytics are used to provide those extremely useful insights and makes them actionable in the form of *prescriptive* recommendations – or *predictive* corrective measures.

Let's face it: averting network downtime is paramount because it's expensive. According to the ITIC 2019 Global Hardware Server OS Reliability Report, 86% of surveyed businesses say that the cost of one hour of downtime is $300,000 or higher. For one-third of these companies, the cost of a single hour of downtime can reach $1 million to $5+ million. Security breaches can be even more costly. As of 2019, the global average cost of a data breach was $3.92 million, according to a report from IBM and the Ponemon Institute. In the U.S., a data breach, on average, costs a company $8.19 million - more than twice the global average. Service disruptions and security breaches adversely impact customer and partner loyalty too; resulting in customer churn.

Network visibility enables you to see what is happening with the network and how it impacts application performance, user experience and security.

Given the growing numbers and types of high-availability and mission-critical applications, systems, and networks – and our increasing reliance on them – it is clear that the process by which network performance and security issues are detected and resolved must evolve. Implementing a comprehensive strategy to assure **network-aware application performance and security assurance** is crucial. Maximizing business continuity and security necessitates the strategy to include proactive "predict-prevent" measures to replace the break-fix approach as much as possible.

Doing so will significantly reduce the resolution time of incidents. Timely problem isolation is imperative for business continuity.

> *Network-aware application performance and security assurance means recognizing the dependencies that applications and security posture have on the network – and managing those network dependencies for the best outcome.*

When driving a vehicle, visibility is essential and blind spots are risks and may become the cause of accidents. The same applies to IT operations. Network blind spots are a risk that could result into unanticipated failures or downtime costing businesses real losses. The only way to consistently achieve network visibility without blind spots is by reliably collecting the **network data**. The foundation of all of those network visibility insights is data. In order for analytics to be useful for **IT infrastructure and operations (I&O)** teams (such as CloudOps, AppOps, DevOps, SecOps, and NetOps), the data must be trustworthy – complete, clean, and consistent. Starting with data you can trust; you get the *visibility you can trust.*

While data is essential, too much of it or not having the right data is also an issue. According to a recent research report by Enterprise Management Associates, nearly 25 percent of large enterprises have eight or more **network performance monitoring (NPM)** tools currently installed, with some supporting as many as twenty-five. It's a big challenge for IT I&O teams to: familiarize themselves with numerous tools and user interfaces, consume and interpret mounds of data, and monitor countless alerts. Making sense of it all to take appropriate, timely action is daunting, to say the least. To save valuable time and resources, enterprises are demanding new ways to leverage infrastructure data that can provide complete visibility, while eliminating existing tools (or the need to invest in new ones) to reduce tool sprawl.

> *Starting with data you can trust; you get the visibility you can trust.*

With digital connectivity and services at the heart of most businesses, maintaining them is mission critical. This eBook focuses on building the right approach towards network visibility which results in an architecture that is modular, efficient, and produces results rather than causing equipment sprawl, overwhelming data, and alert flooding. Here are the five major elements of architecting an effective network visibility solution:

- Reliable Collection, Processing and Distribution of High-Resolution Network Data
- Recording Network Data for History and Playback
- Converting Network Data to Consumable Formats
- Correlating and Analyzing the Network Data
- Using Data for Predictive Analytics and Prescriptive Recommendations

Organizations everywhere are adopting a **cloud-first** or **cloud-smart** architecture, distributing their business applications across private and public cloud infrastructures. At the same time, private data centers continue to be consolidated, while more and more branch offices are connecting to data centers and the public cloud simultaneously. Hybrid is the new normal. According to analyst firm IDC, by 2021 over more than 90% of enterprises worldwide will rely on a mix of on-premises private clouds, public clouds, and legacy platforms to meet their infrastructure needs. This poses both opportunities and challenges.

In a Cloud-First model, the public cloud is preferred and the first priority to run the business. In a Cloud-Smart model, business applications and services are balanced across the private and public cloud in the most optimal way.

The opportunity: Under the hybrid model, users enjoy a consistent, seamless experience, and can maintain high uptime expectations. The challenge: The underlying infrastructure, capabilities, and policies of hybrid-cloud architectures can be challenging for IT I&O teams to manage. A network visibility solution must be able to reliably monitor traffic across an organization's current and future hybrid network architecture – with physical, virtual, and cloud-native elements, which can be deployed across the data centers, branch offices, and multi-cloud environments.

Read the **Business Impact Brief**: Distributed Network Monitoring in a Hybrid World for a survey outcomes and insights offered by 451 Research (S&P Global) and cPacket Networks.

Effective monitoring of physical network traffic requires a high-performance solution built for a range of network topologies, line speeds, locations, and traffic flows. Ideally, network visibility components can support all throughput rates, natively, at line speed, and without dropping packets. When making a purchase decision, IT teams must also take cost, rack-space, scalability aspects into account. A flexible port speed assignment, for example, can enable smooth migration and scaling as network links are upgraded.

Enterprise networks today are not only more distributed, they're also increasingly virtualized, making them

difficult to monitor. Traffic moving between virtual resources does not pass through a physical switch, making it invisible to traditional network traffic collection and analysis tools. Without access to those packets, these areas become blind spots to your network monitoring solutions. To address this challenge, consider leveraging a purpose-built **virtual appliance** to collect, consolidate, monitor, and forward selective network traffic to the appropriate virtual tools. For maximum flexibility, select a solution that can be deployed in various network locations, including the **software-define branch office (SDBO)** and **software-defined data center (SDDC)**.

Monitoring workloads in public clouds is even tougher. To the AppOps team, the cloud can be a "black box" and a major blind spot. How can they measure, much less assure, application performance and dependencies, for traffic they can't see? Cloud-native monitoring tools can help observe infrastructure and application layers but come with significant limitations. They are vendor-specific, often lack features and visibility, and typically do not integrate well with on-premises tools.

How can IT teams collect, consolidate, and analyze traffic in the cloud? One successful approach comprises a software-only solution, natively integrated with leading **Virtual Private Cloud (VPC) traffic-mirroring** services. Advanced functions such as filtering, load balancing, slicing, etc. can be realized and applied to the cloud application workloads. This not only enables seamless access to the VPC's network data, it also reduces complexity and cost. By natively replicating and monitoring network traffic to tools within their VPC, IT teams can avoid using forwarding agents or container-based sensors. Those techniques will be discussed in the following sections.

Let's get started with the **5 key requirements** to build a **network visibility architecture**.

Step 1: Reliable Collection, Processing and Distribution of High-Resolution Data

Starting with the lossless data mining is the very first step toward building an overlay visibility architecture. This starts with **tapping** the right, or important, or strategic places in the network for data acquisition across physical, virtual and cloud infrastructures. You want to make sure that you are collecting the copy of the data from all critical locations without missing any part of it or the picture you will see will be distorted or incomplete. The speeds and feeds, scale, and cost matter at this stage.

A physical network TAP (Test Access Point) or vTAP (Virtual TAP) mirrors or creates a carbon-copy of the network data.

But what kind of data is needed for proper network visibility? Analyzing log, flow, and telemetry data to monitor application performance is useful, but not always sufficient. Data from different network components do not provide a uniform level of detail. Log, flow, and telemetry data are in varied formats, making data capture and analysis more complex. Event logs are not always updated in a timely fashion, and sometimes turned off. They provide a snapshot in time; not designed to capture all information. Flow data is high-level and good for conversational or transaction information but does not provide deep enough information for advanced troubleshooting. In contrast, **packet data** is always consistent, complete, and never lies, making it indispensable for network visibility. It is the highest possible resolution available in the data world. With full packet data, IT I&O teams can avert long and costly troubleshooting times, investigative complexity, and poor end-user-experiences.

In a **distributed hybrid environment** that many enterprises maintain these days, this means you must use the right tools and techniques to fetch the packet data you need. It's not a piece of cake. In the branch offices or remote sites, you may need to tap the WAN or *north-south traffic* going in and out of the office through the firewalls and edge routers. You may also need to tap the local wired and wireless network connecting worker machines, point-of-sale systems, WiFi access points, printers, and other devices. The devices commonly used to access the packet data or *wire-data* in the network and creating a copy of it, or *mirroring* it, is called a **test access point (TAP)**. Additionally, many branch offices are going virtual, meaning hardware devices like firewalls, wireless controllers, and edge routers are being replaced with software-only versions as part of **network function virtualization (NFV)** – resulting in **software-defined branch offices**. This requires **virtual-tap (vTap)** rather than physical Tap.

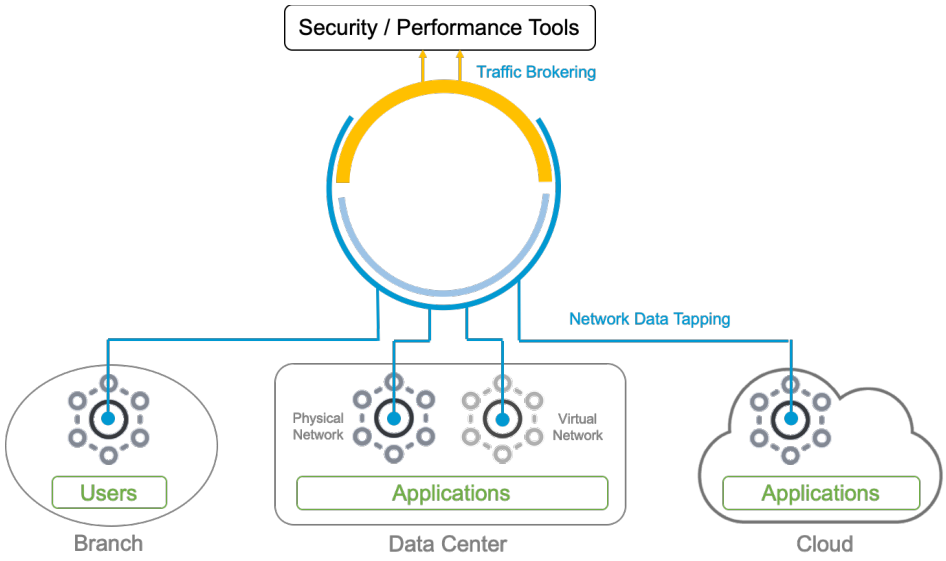

The data centers are larger, denser, and complex and encompass both physical and virtual networks. Just like the branch offices, there is north-south traffic in and out of the data center traversing the spine-leaf, core, firewalls, and edge/border routers which needs to be tapped for having the network visibility. Additionally, there is increasingly more *east-west* traffic in the **software-defined data centers (SDDC)** between application and database components deployed as **virtual machines (VM)** that requires virtual-tapping.

Public **cloud** is a highly virtualized infrastructure and there is no access to the physical network of the cloud data centers referred to as the **infrastructure-as-a-service (IaaS)**. Tapping network data inside the **virtual private cloud (VPC)** space of a public cloud therefore requires virtual-tapping and mirroring but in a specialized way and usually specific to the cloud provider.

cPacket cTap® series TAPs provide a scalable and cost-effective way to tap the physical networks in the branch offices and data centers. For tapping the virtualized networks.

cPacket cVu-V® series can be deployed as VM on top of mainstream hypervisors such as VMware ESXi, Redhat KVM, Microsoft Hyper-V and Cisco NFVIS. For tapping the cloud infrastructure,

cPacket cCloud® cVu-V® series integrates with mainstream cloud VPC traffic mirroring services such as Amazon Web Services (AWS) and Google Cloud Platform (GCP).

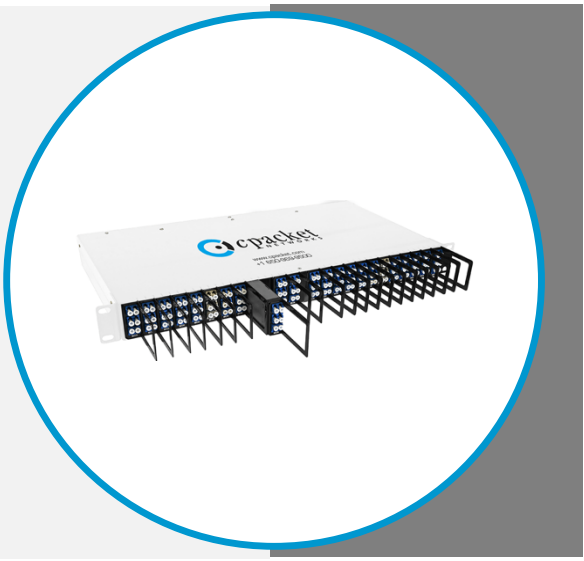

Once you have mirrored the network data or network traffic – you must make sure to consolidate it at a central place so that no *data islands* exist. That is the second common cause for having a distorted or incomplete visibility picture. Since network TAPs do not have the intelligence to distinguish the important data from unimportant data and mirror every packet, too much data can accumulate in a short period of time. You must process the collected data to remove unnecessary and duplicate parts and only forward the right data, in the right format, to the right destinations – such as performance and security tools for analysis.

The above data consolidation, processing and delivery role is handled by the **network packet broker (NPB)**. Packet brokers, or traffic brokers, or data brokers process packet data and broker them to deliver the right packets to the right tools. They can filter out unneeded or duplicate packets and slice, decapsulate, mask, timestamp, load-balance, and intelligently distribute network traffic in the required format. Some packet brokers can also extract meaningful network insights while processing the data which complements the overall visibility. Learn about the key challenges and solutions of traffic brokering.

A physical Network Packet Broker (NPB) or Virtual Packet Broker (vPB) aggregates, consolidates, processes and forwards the network data.

A packet broker built on a flexible, intelligent, and distributed architecture helps ensure reliable collection, consolidation, and distribution of network data. This is the stage that requires a solution that is scalable as the traffic grows, or as you turn on different features requiring additional processing power. The solution must not get stressed or it will undo all the work you did so far by dropping the data i.e. **packet loss** – resulting in incomplete picture again. Pick a solution that has a distributed architecture and is simple and cost-effective to scale.

To understand why, let's think about a busy airport. Services such as curbside check-in, self-check-in kiosks with baggage drop, and trusted traveler programs (e.g., TSA Precheck in the U.S.) distribute – and therefore, *expedite* – the processing of travelers. Imagine what would happen if these services didn't exist. What if all travelers – whether flying domestically or to an international destination – had to use the same centralized resources for check-in, security scanning, etc., before departing at their respective gates? The process would take a lot more time, significant mistakes would be made, and many people would end up missing their flights. Some of those people were enroute to make important decisions; missing their flights could have sizable negative consequences. You get the picture.

With the objective of comprehensive visibility, high-resolution data, and real-time processing, it is important to choose a packet brokering solution carefully. Yet, most visibility solutions based on packet brokers are built using a centralized architecture that is much like a network switch. More specifically, most packet broker devices use a single processing unit and shared memory for all processing, which is easily overloaded when trying to simultaneously process data from multiple ports at high speeds. The problem worsens at higher network speeds and even more so when advanced features requiring more processing – such as deduplication, filtering, slicing, flow exporting, and load-balancing – are enabled. When such packet brokers get overloaded, they act just like a network switch and start dropping packets. This is a huge problem, because once a packet broker fails to capture a packet, one piece of the puzzle is missed. More packet drops mean more pieces lost. The result is a very incomplete picture that the performance or security tools see, causing them to make incorrect assessments, raising the risks that could hurt the business.

In contrast, a packet broker built on an intelligent, distributed architecture will capture all data, even when processing it at the same time. As such, a packet broker with distributed, smart-port architecture, each with its own set of dedicated processing resources, enables both pre-ingress and post-egress processing at the port level. This allows every packet to be fully inspected, filtered, timestamped, decapsulated, deduplicated, load-balanced, aggregated, and prioritized *before* leveraging shared resources. It also means that each packet can be customized for optimal use by NetOps or SecOps tools before being forwarded to its appropriate destination.

cPacket cVu® series physical packet broker+ provides physical network TAP aggregation, consolidation and processing of traffic.

cPacket cVu-V® series virtual packet broker provides the same for the virtualized branch offices and data centers and can be deployed as VM on top of mainstream hypervisors such as VMware ESXi, Redhat KVM, Microsoft Hyper-V and Cisco NFVIS.

cPacket cCloud® cVu-V® series provides the same inside the public cloud and integrates with VPC traffic mirroring services from Amazon Web Services (AWS) and Google Cloud Platform (GCP).

Step 2: Recording Network Data for History and Playback

While tapping and brokering are fundamentally important, they only provide real-time or live data for the analysis. There is no record of data saved for any historical analysis or evidence – which is usually required under distress situations – both for troubleshooting and security incident forensics. Also, the data brokered by the packet brokers provides you a snapshot in time, in a single location and not end-to-end session or transaction level information, which is very important.

A physical or virtual traffic or packet capture solution captures the live network data from the network and saves it to the storage for later retrieval and analysis.

For this purpose, you need to capture and save the network data to some sort of storage – referred to as the **capture-to-disk (CTD)** What you are doing is essentially converting the *data-in-motion* to *data-at-rest* or static data. Learn more about the key challenges and solutions of traffic capture and analysis.

For a session level analysis, you need at least two reference points. It is recommended that you capture the data at multiple key locations to see the proper visibility picture. Just like the tapping and brokering, to have the full ground covered in a hybrid environment, you need the physical, virtual, and cloud-native capture solutions. The physical packet capture devices get deployed in the data centers or service provider locations to capture the data from north-south direction, edge and spine-leaf network. The virtual packet capture appliances, usually as VMs or containers, get deployed in the virtualized branch offices and data centers to capture north-south and east-west traffic respectively. Cloud-native capture solutions do the same inside a cloud VPC.

To perform its job properly, the packet or data capture solution must be high performing enough to ingest the data at sustained rates without dropping the packets, and efficient in terms of structuring, searching, and retrieving the saved data for playback and analysis. Search and query capabilities are important as well as the usability aspect. Capturing the network data losslessly at higher network speeds becomes quite tricky as you need to grab every packet fast enough and write it to the disk while organizing it in some sort of file system. Data compression techniques may also be used to maximize the storage capacity of the disks or cloud storage which has a cost associated with it. It is also important that the capture solution can handle multi-tasking.

That is, it can capture and write the data to the disk while you are searching the already saved data and analyzing it simultaneously. Not many solutions can do that, just like they cannot handle a sustained rate of capture over a longer period.

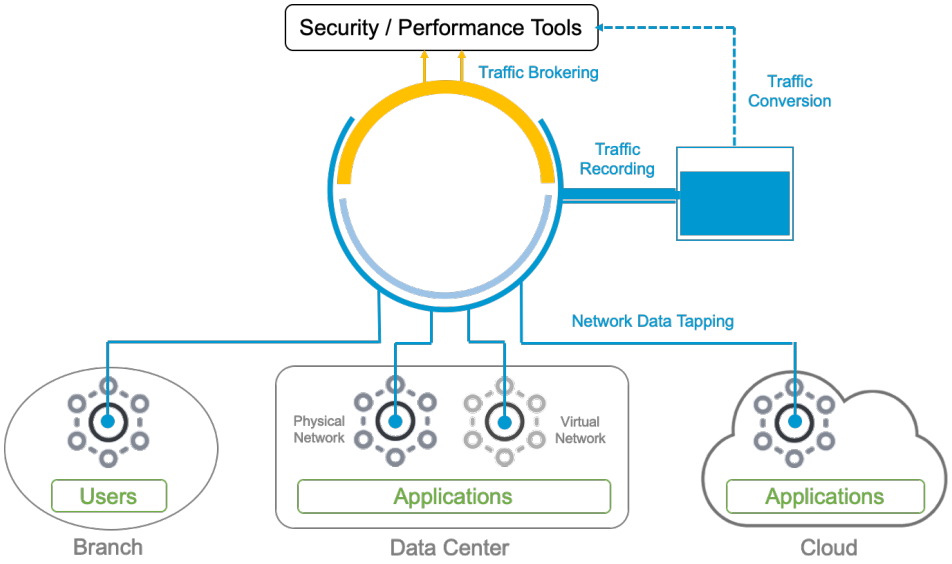

One of the ways to detect malicious activity is by capturing and analyzing the suspicious east-west traffic. **Network Traffic Analysis (NTA)** is the process of intercepting, recording, and analyzing network traffic communication patterns in order to detect and respond to security threats. NTA may or may not leverage the **Deep Packet Inspection (DPI)** technology in the network visibility devices which can identify specific attacks that an organization's firewall and intrusion prevention/intrusion detection systems (IPS/IDS) cannot adequately detect. DPI is also an important way to prevent worms, spyware, and viruses from getting into your corporate network through end devices.

The packet data captured and stored is useful in many ways – sometime even more useful than the live data. It lets you travel back in time at an exact moment and playback the network data just like a DVR would for a movie. You can pause, rewind, and playback as many times as you like without the pressure of analyzing data in real time. This provides the capability to troubleshoot deep application and network issues and reduce the **mean-time-to-resolution (MTTR)**. Most importantly, it is used for **Network Forensic Analysis (NFA)** for gathering, monitoring, and analyzing network activities to uncover the source of attacks, viruses, intrusions, or security breaches that occur on a network or in network traffic. While real-time analysis can cut storage costs and troubleshooting time, it can't replace forensics altogether. Analyzing network packet and flow data helps IT teams quickly and accurately reconstruct events, to pinpoint precisely what took place. Lossless packet capture and analysis enables IT SecOps teams to successfully perform security audits, ensure regulatory compliance, and enforce policies.

User and Entity Behavior Analytics (UEBA) uses machine learning, algorithms, and statistical analyses to: detect deviations from established patterns, identify which of these anomalies could result in a real threat and, alert IT SecOps to the occurrence of a successful attack. As such, UEBA can reduce an organization's vulnerability to common cyberattacks, including phishing, whaling, social engineering, malware, and ransomware

It can take weeks, even months, for organizations to discover that a security breach has occurred, and the extent of damage inflicted. In fact, according to the 2019 Cost of a Data Breach report by IBM and the Ponemon Institute, it takes 279 days, on average, to identify and contain a data breach. Why does it take so long? Security vendor solutions deployed within the same infrastructure typically do not correlate suspicious activity with one another nor do they have access to detailed threat information. It's up to the IT SecOps team to undertake the time-consuming and costly endeavor of collecting and correlating information from these disparate sources. Real-time analysis of network packet data can help to mitigate this problem. Immediately correlating a rich, relevant set of packet-level information with existing security device metrics provides a more complete and detailed view of a potential security issue. This insight enables IT SecOps teams to deliver a targeted response to threats, saving valuable time and reducing the organization's overall risk.

cPacket cStor® series physical packet capture appliances capture the live data to local disks on-board and provide session level analysis using Wireshark.

cPacket cStor-V® series virtual packet capture appliances provide the same services for the virtualized branch offices and data centers and can be deployed as VM on top of mainstream hypervisors such as VMware ESXi, Redhat KVM, Microsoft Hyper-V and Cisco NFVIS.

cPacket cCloud® cStor-V® series provides the same for capturing the traffic inside the public cloud and integrates with VPC traffic mirroring services from Amazon Web Services (AWS) and Google Cloud Platform (GCP).

3. Converting Network Data to Consumable Formats

The network data captured in the previous step is in the packet format. Packet data is the richest form of data one can collect but it is also intense. Unless, you need to retain the packet data for longer durations for reasons such as compliance or forensics, most use cases require either capturing the packets over a window of time when an issue arises or extracting the **metadata** out of the packet data and discarding the full packets. It's the metadata that is stored and takes much less space. Metadata contains most of the useful information extracted from the packet headers and sometime from the payload as well if **deep packet inspection (DPI)** capabilities are supported. The metadata can be shared among the tools and exported to analytics platforms to correlate and construct important network **key performance indicators (KPI)** or metrics which are displayed across different types of dashboards – usually in a **network operations center (NOC)** or **security operations center (SOC)**.

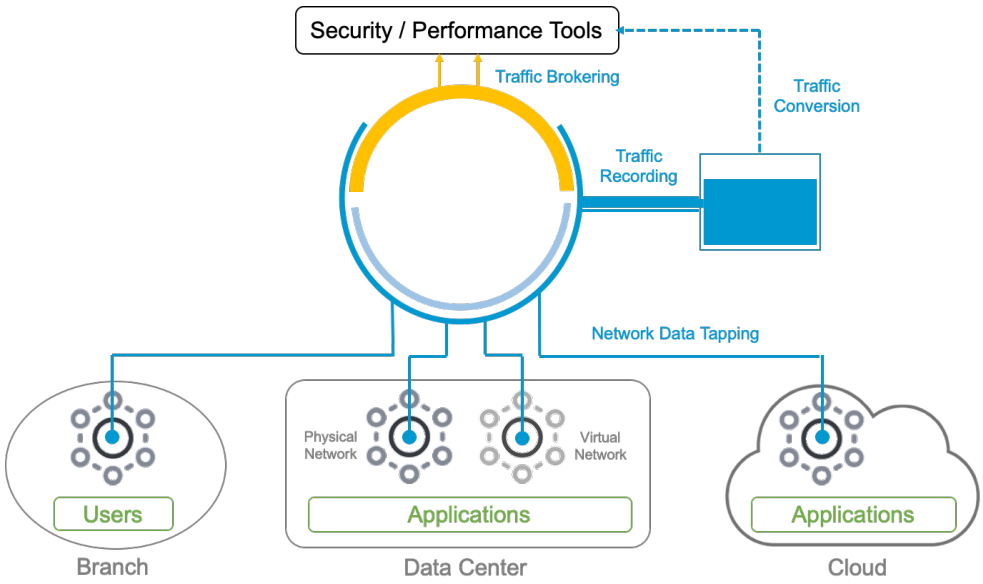

The packet data can also be converted into what is called the **flow data**. Flows carry end-to-end conversational or transactional information between a source and a destination. Several hundred or thousands of packets may be exchanged between a source and a destination, a single flow, in one direction, and in the reverse direction. Therefore, there are much smaller number of flows than there are packets and hence, you have to deal with consuming less information. Flow data gives you a different view of things. In this rather top-down approach you can decide which transactions or conversations over the network you are interested in from investigation or measurement perspective, and then look into the packet data associated with those flows only. This narrows down the set of packets to be analyzed – greatly simplifying the analysis.

A flow generator device takes the packet data and converts it into the flow data in standard format to be exported to other devices and tools.

You can convert the packet data to other forms of data such as flow data but not in reverse. Flow and packet data are not mutually exclusive but rather complementing when it comes to building a network visibility architecture. They provide different types of and different levels of information. Some tools or dashboards like to consume the data in flow format rather than packet format. Your next step is to have the capability to generate the flow data out of packet data in the standard formats most tools want – such as security and performance tools. Most common formats include *Netflow*, *IPFIX*, and *sFlow.*

When it comes to generating and exporting the flow data from the packet data, you may have different strategies. One common way is to use the packet brokers for this additional role and add one more function to what they are tasked to do. Most packet broker vendors charge an additional license fee for that. Still, it is not the best and a scalable way since as discussed earlier, most packet brokers in the industry are constrained by the central processing which is already limited. Adding the flow generation and export functionality adds additional burden on processing and prevents the packet broker from doing its basic job. The same argument applies to other network devices that generate and export flow information such as routers and switches. Their primary job is forwarding packets to the right destinations and they are not designed to scale with such additional responsibilities.

The preferred way to do this is to use dedicated appliances or probes in the network that do it on the fly as a dedicated function at high speeds. The flow generator device can take a data feed straight from a network TAP or a packet broker. The performance parameters that generally matter include the number of new flows and the number of sustained flows a device can generate per minute – referred to as **flows-per-minute (FPM)**. Using a dedicated flow generator and exporter device has clear advantages over other approaches as you don't have to worry about them being chocked up and feeding misinformation constructed out of missed packets.

cPacket cProbe® series flow generator and exporter solution are exactly that. cProbe physical appliance can capture the live packet data and generate flow data out of it in Netflow and IPFIX formats to export to third-party devices and tools.

4. Correlating and Analyzing the Network Data

You are now done with collecting, processing, distributing, recording, and converting the network data – but that is only half the job. The other half is what you do with the data. One approach IT teams can take is to ensure that the network visibility solution they select delivers required KPIs and analytics in a single-pane-of-glass. Your benefits come from network data that is collected end-to-end from a hybrid environment, then it is correlated, analyzed and presented, ideally in a single-pane-of-glass to increase timely decisions.

Single-Pane-of-Glass analytics collects the network data from end-to-end hybrid environment, correlates it and presents it in a single place.

Correlating data collected from heterogeneous environments, making sense out of it, and presenting it in a format that is easy to understand and useful to make timely decisions, is super critical. At this stage, you must be able to make sense out of the packet, flow and metadata and able to perform baselines, set thresholds for normal behavior, map dependencies, and generate alerts for the service level monitoring. This is where the value-add part resides. Moreover, you benefit most when you have visibility and actionable insights that span your entire hybrid environment.

Establishing a proactive network monitoring strategy starts with defining network health KPIs – such as latency, packet errors, and connection errors – using expected baselines. Once these KPIs have been defined, automated comparisons to their baselines can proactively alert IT teams. These details and actionable insights empower IT teams to address potential issues before end users and the business are negatively impacted. To be successful, it's important to identify the correct network KPIs for given circumstances.

As described earlier, there are foundational reasons to collect the richest possible data – packet data – so that you don't miss the fine-grain information you need to extract out of it. But what you use is the information and not the packets. It is much like drawing the blood from the vein to test for a problem with the body since the rich and most reliable information a blood sample can provide, cannot be obtained through alternative tests. But once a sample has been analyzed under the microscope and the information has been extracted, there is little reason to keep the original blood sample in the storage bank.

In the networking world, metadata provides what is equivalent to the plasma. Metadata is extracted from the packet and flow data and is sorted and saved in a database, usually a *time-series* database,

The metadata needs to be exported or collected from multiple devices or probes spread throughout the hybrid network such as physical packet brokers, packet capture and flow generator devices, virtual appliance running as VMs and cloud-native instances running inside multiple cloud environments. The metadata may travers public internet or WAN and may have security and cost issues to be addressed. Therefore sometimes, collect-and-store maximum locally, export minimal externally makes sense and sometimes collect-and-store maximum locally and analyze locally makes sense.

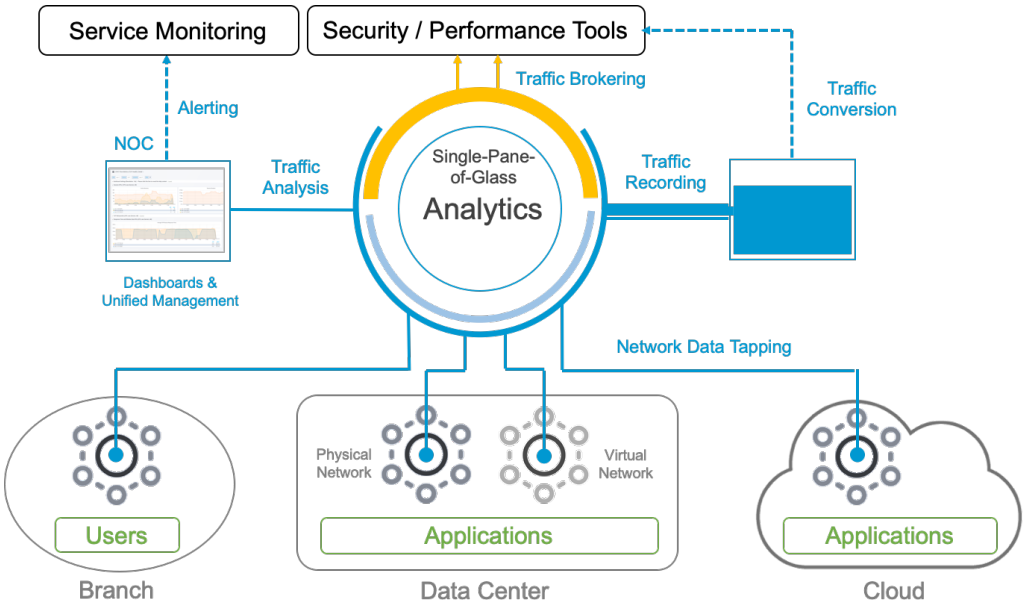

Once the central correlation and analysis engine collects the metadata from multiple sources, it starts to look for the information it needs and calculate what it must per the required KPI. Once it has generated the required network metrics, they must be visualized in an easy to understand way using an intuitive **user interface (UI)** for a richer user experience. But this is not the only part expected out of the analysis engine. The ability to *filter* the data to only display the meaningful information is a basic required function. Setting up the *policies* to collect only certain data under certain circumstances may be an additional requirement.

Other features may include setting up the thresholds based on normal behavior or parameters and sending alerts when those thresholds are crossed. Per our earlier example, this is like having minimum and maximum ranges for key health indicators on a blood test report such as for blood cholesterol, white and red blood cells etc. When the results fall below or above those ranges, it is indicative of something going wrong in the body. Those thresholds are set based on normal state of the human body. In a network, certain parameters are easier to be defined by default as what is considered normal. Others may need to be defined by the network administrator. This is called **baselining**. Once you have baselined your network behavior, it becomes rather easier to track when a deviation or change occurs and then track what changed, when, and why.

Another important part is determining and mapping the relationships between application or security components and the network. In other words, determining how an application depends on the network and how it can be impacted by it if certain parameters in the network change. This is called **dependency mapping** and is a key step in building modern network or application monitoring systems. Both applications and security do get impacted in significant ways when network performance or network security is compromised. Instead of finger-pointing at each other, AppOps, SecOps, and NetOps teams need to define those dependencies and monitor them. This is why the concept of **network-aware application and security assurance** is important.

Finally, the network KPI based metrics that are measured and generated in terms of **analytics** need to be defined. Those metrics may include things like top talkers, bandwidth utilization, one-way latency, jitter, micro-burst, gap analysis as well as session level measurements such as TCP response time, errors, voice and video analysis, and other parameters. The breadth and depth of those measurements, sampling resolution and ability to customize the dashboards, as well as ability to share those with other visibility platforms through open **application program interface (API),** are some of the important considerations at this stage.

cPacket cClear® series single-pane-of-glass analytics solution is exactly that. cClear series physical appliance collects, stores and analyzes the meta data from branch offices and data centers.

cPacket cClear-V® series virtual appliance provides the same services for the virtualized branch offices and data centers and can be deployed as VM on top of mainstream hypervisors such as VMware ESXi, Redhat KVM, Microsoft Hyper-V and Cisco NFVIS.

cPacket cCloud® cClear-V® series provides the same inside the public cloud and integrates with mainstream cloud VPC traffic mirroring services such as Amazon Web Services (AWS) and Google Cloud Platform (GCP).

5. Using Data for Predictive Analytics and Prescriptive Recommendations

By now, you have architected a functional visibility infrastructure that is great for day-to-day operations. It provides everything that is needed to monitor a mission-critical network to run the business. If you like to have a visibility model that you can setup, benchmark, and use for reactive monitoring based on troubleshooting, analysis, and resolution, then you are done. This visibility model still depends on human intervention, analysis, and intelligence to figure out what is wrong, where and why and what corrective actions need to be taken. This means, IT I&O teams' significant resources and time required to be allocated towards network monitoring.

Artificial Intelligence for IT Operations (AIOps) leverages machine learning and artificial intelligence along with the data to enhance the ITOps with proactive, prescriptive and predictive insights.

If however, you would like to take it to the next step where the visibility model has intelligence of its own to learn and adapt from day to day data and operations and train itself, then you need to move to the next step which is incorporating **machine learning (ML)** capabilities. Machine learning can learn and train itself from the situations it encounters and gets better and better at it. The more you expose it to different scenarios; better and more accurate it becomes. Machine learning based visibility models can do operations such as data pattern matching, correlation, and analysis in a much more efficient way than humans can – saving IT resources for high-value activities. Machine learning based visibility models are still diagnostic and prescriptive in nature but are more efficient and accurate. They cannot predict what kind of issues could come up if the network situation heads in a certain direction or if certain actions are not taken.

Yet there is a next level of sophistication which is based on using **artificial intelligence (AI)** technology. Artificial intelligence-based algorithms can analyze the data and metrics collected under a context much like human analysts use logic to infer what could be going on and suggest or even take the corrective actions such as tuning up certain network parameters or shutting down the connectivity to suspicious hosts. So, in this sense they are also prescriptive. In fact, mostly machine learning and artificial intelligence work together to provide prescriptive resolutions as artificial intelligence builds upon the work done by the machine learning. Artificial intelligence however has the ability to be *predictive* and can predict, within certain margin of accuracy, what could happen in near future based on the data it sees and the analysis it does. It's kind of how an experienced

network engineer can predict an expected outcome based on what they have learned over the years.

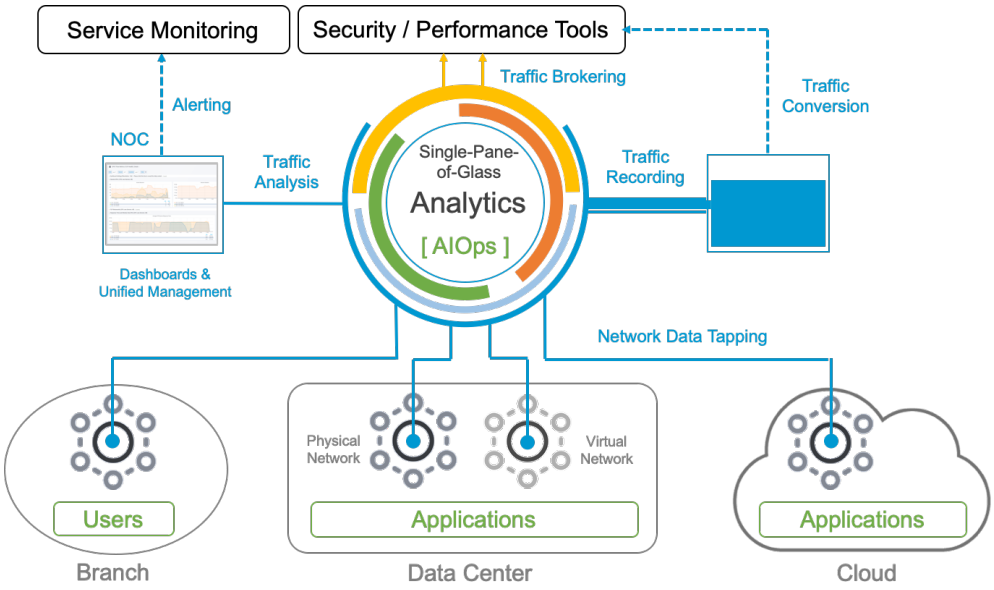

Predictive analytics is required for comprehensive network visibility. Predictive analytics empowers IT I&O teams to efficiently resolve issues *before they happen*. By ingesting and correlating data from diverse hybrid-cloud environments, predictive analytics provides actionable insights to prevent unplanned performance degradation and service disruption. This is where **artificial intelligence for IT operations (AIOps)** is creating the hype in the IT industry. AIOps is getting lots of attention by the IT I&O leaders that is CIOs and CISOs since it greatly helps them to improve IT efficiency by reducing pressure on IT human resources, improving the accuracy and productivity, and enhancing the agility. Therefore, AI/ML capabilities in building a network visibility architecture to support overall AIOps capabilities are becoming highly desirable.

cPacket cClear® series single-pane-of-glass analytics solution is exactly that. cClear series physical appliance collects, stores and analyzes the meta data from branch offices and data centers.

cPacket cClear-V® series virtual appliance provides the same services for the virtualized branch offices and data centers and can be deployed as VM on top of mainstream hypervisors such as VMware ESXi, Redhat KVM, Microsoft Hyper-V and Cisco NFVIS.

cPacket cCloud® cClear-V® series provides the same inside the public cloud and integrates with mainstream cloud VPC traffic mirroring services such as Amazon Web Services (AWS) and Google Cloud Platform (GCP).

Conclusion

Our increased reliance on always-on digital connectivity has made managing, optimizing, and securing enterprise and service provider networks a mission-critical exercise. Investing in a solid network visibility solution that meets the criteria discussed in this eBook will empower IT teams to ensure business continuity and better client experience, even under the most demanding of circumstances. More importantly, it will provide them with significant efficiency uplift.

Building a network visibility architecture that scales with growth, helps reduce the complexity rather than introducing its own. It is also cost-effective for pay-as-you-grow but are not readily available from every visibility solution provider. The strategies discussed in this eBook and associated solution pointers will help you make the right choices. An effective network visibility solution is an indispensable part of the IT infrastructure and operations. To make this strategic investment, IT managers should work with an experienced vendor – one with a demonstrated commitment to developing innovative, scalable, and cost-effective network visibility solutions.

cPacket Networks delivers visibility you can trust through network-aware application and performance assurance per the best practices discussed in this eBook. Its cutting-edge technology enables IT I&O teams to proactively identify issues in real-time before negatively impacting end-users. Leading enterprises, service providers, healthcare organizations, and governments rely on cPacket solutions for improved agility, higher performance, and greater efficiency.

Learn more at www.cpacket.com, the cPacket blog, or follow us on LinkedIn, Twitter, Facebook, YouTube, and BrightTalk. You can also Talk to an Expert, Contact Sales, or Request a Demo.