savvius

WHITE PAPER

Best Practices for 10G and 40G Network Forensics

On highly utilized 10G and 40G networks, capturing network traffic from individual SPAN ports on switches and routers typically results in spotty visibility, compromising an IT team's ability to respond to network outages, performance degradations, and security threats. How should IT organizations re-think their network forensics strategies, now that 10G and faster networks are becoming the norm? Read this paper to learn about best practices for recording, storing, and analyzing network data.

Contents

Introduction

Network forensics is the capture, recording, storage, and analysis of network events. A network forensics solution records network traffic, stores it in a searchable repository, and provides IT engineers with filters for mining stored data to discover and analyze network anomalies. Using network forensics, IT engineers can discover both the cause of an anomaly and its effects on IT services and assets.

Some IT organizations mistakenly believe that the primary purpose of network forensics is to find proof of security attacks. It's certainly true that network forensics can enable IT engineers and security teams to find proof of attacks that other IT tools, such as Intrusion Detection Systems (IDS), can only hint at. But it's also true that network forensics can be used daily to analyze far more common issues on networks, such as spikes in utilization, drops in VoIP call quality, and increased latency in specific applications or in network traffic overall.

In fact, for organizations that have deployed 10G and 40G networks, network forensics provides the only practical way to analyze network traffic systematically in detail. Traffic is flying by far too quickly on 10G and 40G networks for IT engineers to monitor and analyze in detail through real-time dashboards. Only by analyzing captured traffic can IT engineers really understand what has taken place on a high-speed network, which problems, if any, are occurring, and how they might be solved.

This paper considers the special challenges that 10G and 40G networks create for network operations centers (NOCs) and other IT organizations responsible for monitoring, managing, and securing enterprise networks. It also presents best practices for applying network forensics to these high-speed networks. The paper concludes with an introduction to the network forensics solutions offered by Savvius, Inc. — solutions that have been designed expressly to meet the needs of IT engineers working with 10G and 40G networks.

The Challenge of Analyzing Highly-Utilized 10G and 40G Networks

Increasing network speeds from 1G to 10G or 40G does not affect all IT monitoring infrastructure the same way. Device monitoring is hardly affected at all. In flow-based monitoring systems, IT engineers will see incremental increases in flow metrics as applications such as video and VoIP start to take advantage of the increased bandwidth available from the faster networks. With deep packet inspection, however, IT engineers will notice a big difference almost immediately — and the difference is significant enough to affect the IT organization's workflow.

To understand this difference, it's helpful to review the process for monitoring networks running at 1Gbps or slower. On these slower networks, IT engineers have considerably more flexibility when it comes to network monitoring and troubleshooting. Traffic streams across the network, monitored in real time and characterized by stats and charts on dashboards. If a problem arises, an engineer can connect a network analyzer to the network, start a trace, and if the issue is an ongoing problem, reproduce it for analysis. Data volumes are low enough that almost any network interface card (NIC) and almost any computer, even a laptop, can be used to capture and analyze traffic. Data capture and analysis have little to no impact on the network traffic being analyzed.

> "Packet monitoring really is the most definitive, most complete source of performance data you can get for managing networks and for troubleshooting in particular."
>
> **JIM FREY,** MANAGING RESEARCH DIRECTOR
> ENTERPRISE MANAGEMENT ASSOCIATES

The flowchart below represents a typical problem-solving workflow for troubleshooting 1G networks.

FIGURE 1: TYPICAL 1G NETWORK ANALYSIS WORKFLOW

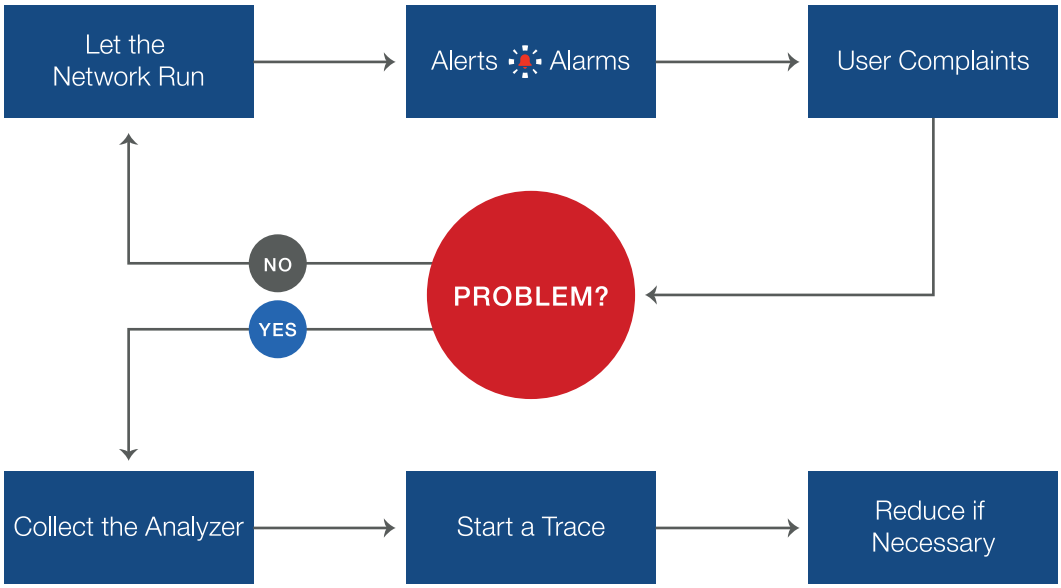

Now consider the flow chart below, which depicts the workflow required for monitoring and troubleshooting a high-speed network. Because the traffic is flowing so quickly, the only way to analyze it in detail is to capture it first with a network forensics solution and replay it for inspection. To do this, though, IT organizations must already have identified key points on the network and deployed 24/7 monitoring solutions (specialized hardware, not generic NICs in PCs) that continuously capture traffic for analysis.

FIGURE 2: 10G/40G NETWORK ANALYSIS WORKFLOW

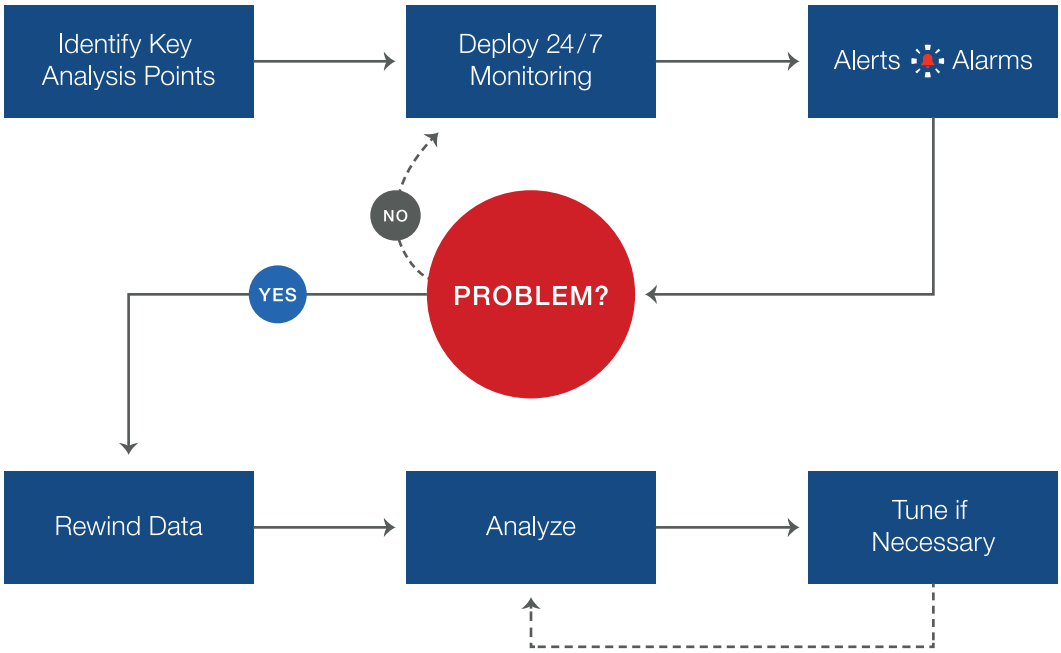

There are many significant operational and technical differences between network monitoring at 1G and network monitoring at 10G and 40G. These differences include:

- **Continuous Data Capture**
  Instead of initiating traffic capture only in response to a user complaint or alert, IT engineers managing 10G and 40G networks must identity key locations on the network and capture traffic continuously at those locations 24/7. Otherwise, IT engineers risk not having a capture set up fast enough to keep up with high- speed network events.
- **Reviewing Data as the New Norm**
  Alerts and user complaints prompt IT engineers to review data that has already been captured, rather than to initiate a capture in time to record data for analysis.
- **High-Performance I/O Interfaces**
  IT engineers need network interface technology capable of reliably capturing high-speed traffic. Traditional NICs, such as 10/100 or 10/1000, are no longer useful for capturing traffic. These cards were designed to route and move packets on slower networks. They were not designed to capture packets reliably on higher speed networks.
- **High Performance I/O Buses and Disks**
  The high throughput of 10G and 40G networks can tax all parts of a system being used to capture and analyze network traffic. I/O busses and disk arrays must be able to keep up with high data rates reliably, hour after hour, day after day.

- **Increased Processing Power**
  IT engineers need specialized processing technology to be able to process the number of packets. IT engineers should not expect to be able to connect a laptop with network analysis software to a 10G or 40G link and expect it to process the number of packets that will flow through in a given second or millisecond.
- **Increased Storage**
  Capturing data from a 10G, let alone a 40G, network can easily consume tens or even hundreds of terabytes of storage. Network forensics systems need to provide fast, scalable storage solutions that can accommodate these high volumes of data.
- **The Increased Importance of Search Tools and Filter**
  Given the high volumes of data to be searched and analyzed, precise filtering and high-performance searches become even more important for network analysis. A few hours of traffic can easily generate tens of terabytes of data. Searching through that data quickly is paramount for resolving issues, such as security threats and performance degradations, in a timely manner.

10G networks are becoming the norm for new network investments, and 40G networks are gaining popularity as well.[1] How should IT organizations meet the technical and operational challenges of monitoring, analyzing, and optimizing these high-speed networks?

Best Practices for Network Forensics on 10G and 40G Networks

Fortunately, with the right network forensics tools and a methodology designed to accommodate the challenges of high-speed networks, IT organizations can monitor, manage, and optimize networks, even as network speeds increase tenfold or more.

This section sets forth some best practices for designing and implementing network forensics solutions on 10G and 40G networks.

Best Practices for Capturing Network Traffic

Best Practice #1: Capture traffic continuously and comprehensively.

To be able to investigate any type of network activity or anomaly on 10G or 40G networks, record all network traffic — everything from email to VoIP — to a single repository that can be examined after the fact.

Comprehensive data captures enable IT engineers to examine every packet and every unencrypted communication payload on the network. Whether troubleshooting an application performance problem or investing an HR policy violation, this comprehensive record-keeping is invaluable.

Best Practice #2: Deploy a traffic continuously and comprehensively.

Select a solution that captures traffic reliably, even at speeds exceeding 10Gbps.

In a recent survey of enterprise IT organizations about their network analysis solutions, TRAC Research found that more than half of IT leaders (59 percent) were concerned about the number of packets dropped by network recorders, and a little over half (51 percent) were concerned about the reliability of data being captured.[2] Before purchasing a network forensics solutions, an IT organization should test the solution with live data to ensure it captures and stores data as reliably as it promises to.

[1] Infonetics expects sales of 10G network ports to grow ten-fold between 2013 and 2017. See source.

[2] See our blog post on 10g network performance monitoring tools.

Organizations should be aware that in many cases, "line rate" doesn't guarantee loss-less packet capture at 10G or faster. Unfortunately, some network analysis products drop packets at high speeds. IT organizations are advised to test before purchasing.

Best Practice #3: Set up filters to catch anomalies.

To facilitate the detection of anomalies, especially those associated with security events, set up filters to capture unexpected traffic automatically.

Sometimes security attacks use common protocols such as CIFS or SMTP to spread malware and issue commands. IT organizations can accelerate the detection of this anomalous behavior by defining network captures like the following:

- Non-DNS traffic to or from DNS servers
- SMTP traffic among non-SMTP servers
- Spikes in CIFS traffic

One advantage to defining special captures like this is that the resulting files are usually small and can be searched quickly.

Best Practices for Storing Traffic

Best Practice #4: Allocate sufficient storage for the volume of data being collected.

An effective network forensics solution has to do more than capture traffic; it also has to store it so that it is always available for analysis. IT organizations should ensure they have provisioned adequate storage for the networks being monitored.

10G and 40G networks significantly increase the amount of data storage required for network forensics. For the sake of comparison, here's a formula for calculating the storage required for monitoring a fully utilized 1G network or a 10G network running at only 1Gbps:

1 Gbps x 1 bits / 8 bytes x 60 s/min x 60 min/hr x 24 hr/day = 11 TB/day

With a 32TB appliance, an IT organization would be able to go back in time 2.9 days to analyze traffic. If problem occurred over the weekend, they would probably be able to find out what happened and recreate the problem on Monday. They would have every packet that was transmitted through the link being monitored, so data reconstruction at any level would be possible. After 2.9 days, the storage system will start overwriting the data it initially recorded almost 3 days earlier. On a Monday, traffic from the previous Thursday is no longer available.

Now let's look at a 10x increase in traffic, 10 Gbps on a fully utilized 10G network.

10 Gbps x 1 bits / 8 bytes x 60 s/min x 60 min/hr x 24 hr/day = 110 TB/day

At this higher speed, a 32 TB appliance is able to go back in time only 7.0 hours. If a problem occurred overnight, engineers might be able to find out what happened, but if the problem occurred over the weekend, they would not have the packet data they need for their investigation. And on a 40G network? Even if the network is only half utilized, a 32 TB appliance can store only about 3.5 hours of traffic. To store a day's worth of traffic would require over 220 TB — nearly a quarter of a petabyte of storage per day.

Of course, not all networks are fully utilized, so storage systems are not likely to fill up quite as quickly as these numbers suggest. Futhermore, IT engineers can reduce storage requirements by filtering out types of data that are unlikely to be useful in analysis.

If, though, after monitoring network utilization and selecting appropriate filters, data volumes remain high, an IT organization might want to consider deploying a network tap that could split recording over multiple network forensics appliances. Or they might want to consider connecting network analysis and recorder appliances to a SAN for additional data storage.

Best Practice #5: Adjust file sizes for the desired performance optimization.

The two most common formats used for storing recorded traffic are standard packet files and databases. In either case, two metrics to manage closely are file size and frequency of disk writes.

Increasing the file size can reduce disk writes but increase memory requirements when the file is eventually opened for analysis. Files that are too large simply become unworkable on whatever system is being used for analysis.

Smaller files, however, typically lead to more frequent disk writes, and this can rob the system of precious resources for performing the actual packet capture. Optimum performance is achieved by balancing these two demands.

Also, try to use the lowest number of simultaneous captures as possible. Several network forensics solutions allow engineers to create as many captures as they want. Engineers should remember, though, that each capture reserves its own memory for buffering, reducing the system's overall memory for data processing.

Best Practices for Analyzing Traffic

Best Practice #6: Select a network forensics solution that supports filters and searches that are fast, flexible, and precise.

The faster the network, the more important it is for a network forensics solution to be scalable and reliable — capable of supporting additional tens of TBs of storage without dropping packets or storing them in such a way that searches become overly complex or time-consuming.

Best Practice #7: Record baseline measurements of network performance.

To get a sense of "normal" conditions before trouble arises, IT engineers should take baseline measurements across specific network traffic such as HTTP, VoIP, and key business applications over typical cycles, such as an hour, a day, and a week, for the network as a whole. They should also note other metrics, such as packet size distribution as well as protocol and node usage over time, uncovering cycles in these metrics, which provide a "fingerprint" of the network's typical utilization. By archiving these measurements, IT engineers ensure that they always have an accurate model of the network's "normal" functioning for comparison when trouble arises.

Best Practice #8: Use filters to zoom in on the problem at hand.

To troubleshoot problems quickly, IT engineers should use their knowledge of the network and its typical behavior to filter on specific data types or IP subnets for analysis. The more specific the network forensics search, the more quickly the data can be analyzed. Often a variety of conditions can be ruled out immediately. By filtering out non-applicable protocols and analytics, IT engineers can accelerate root cause analysis.

For example, if engineers are analyzing a 10G network link, they are not capturing wireless traffic. They can turn off the wireless analysis for this particular link (that is, turn off analysis of Wi-Fi protocols and transmissions, while continuing to analyze all traffic traversing the 10G link, even if it originated on a WLAN).

If the issue being investigated does not directly pertain to VoIP or video, engineers may be able to turn off VoIP or video analysis as well, saving CPU cycles, memory, and disk space for analytics that are more pertinent.

Even after analysis has been streamlined to only essential areas of the network, data capture for network analysis on 10G networks quickly generates a great deal of data, and managing that volume of data can be a challenge. For best results, IT engineers should thoroughly understand the performance and storage capabilities and limitations of whatever network forensics solution they are using to record and analyze network traffic.

Something else to consider, is whether you'll be performing real-time analysis or post-incident or forensics analysis. Real-time analysis is tricky at 10G, but you can still use real-time data to pinpoint developing problems, highlighting a period of time you want to look at further, drilling down later using forensic analysis.

Finally if you're just doing network performance analysis or network performance tuning, you may not need the packet payloads and can slice the payload from your data, significantly increasing the data that you can store.

Conclusion

The network analysis tools that organizations have invested in over the past decade or so are simply not able to keep up with today's high-speed networks. New tools and IT practices are necessary if IT organizations are going to keep new networks running as well and as securely as old ones.

Network forensics enables organizations to realize the full benefits of 10G and 40G networks: high performance with the control and security IT organizations take for granted on 1G networks. By investing in network forensics solutions and following the best practices listed in this paper, IT organizations can ensure that speed does not come at the expense of visibility, control, or security.

Savvius Solutions

Savvius solutions combine real-time analysis, seamless packet capture, and sophisticated, easy-to-use software for monitoring networks, uncovering the root cause of performance issues, and enabling deep understanding of security incidents. Savvius solutions match requirements ranging from the largest central datacenters to far-flung offices in distributed organizations.

Network Monitoring

savvius Spotlight™ Appliance

Savvius Spotlight is a new technology that provides actionable network visibility for performance monitoring and troubleshooting in real time at over 20 Gbps. Learn more.

Capture and Analysis Appliances

savvius Omnipliance Ultra™

Savvius Omnipliance and Omnipliance Ultra (Includes Spotlight technology) deliver powerful, precise, and affordable packet capture and analytic solutions for 1G, 10G, and 40G networks. Learn more.

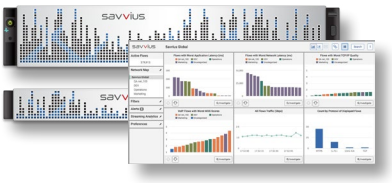

Network Analysis

savvius Omnipeek®

Savvius Omnipeek delivers visual packet intelligence based on sophisticated deep-packet analysis. Customizable work flows and visualization across multiple network segments drive faster mean time to resolution of network and security issues. Learn more.

Long-term Packet Storage

savvius Vigil™

Savvius Vigil automates the selective packet capture of network traffic needed for network forensics and security investigations. Learn more.

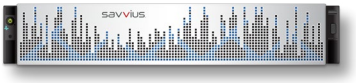

Remote Locations

savvius Insight™

Designed for deployment in a variety of distributed office and retail environments, Savvius Insight and Insight Plus are compact, fanless, ELK-compatible mini-appliances that provide all the benefits of Savvius' enterprise-grade datacenter solutions. Learn more.

WiFi

savvius Omnipliance WiFi

Savvius Omnipliance WiFi is the only WLAN analysis solution that enables network engineers to monitor, analyze, store, and troubleshoot multi-Gigabit speed 802.11ac traffic. Learn more.

Savvius Distributed Network Diagram

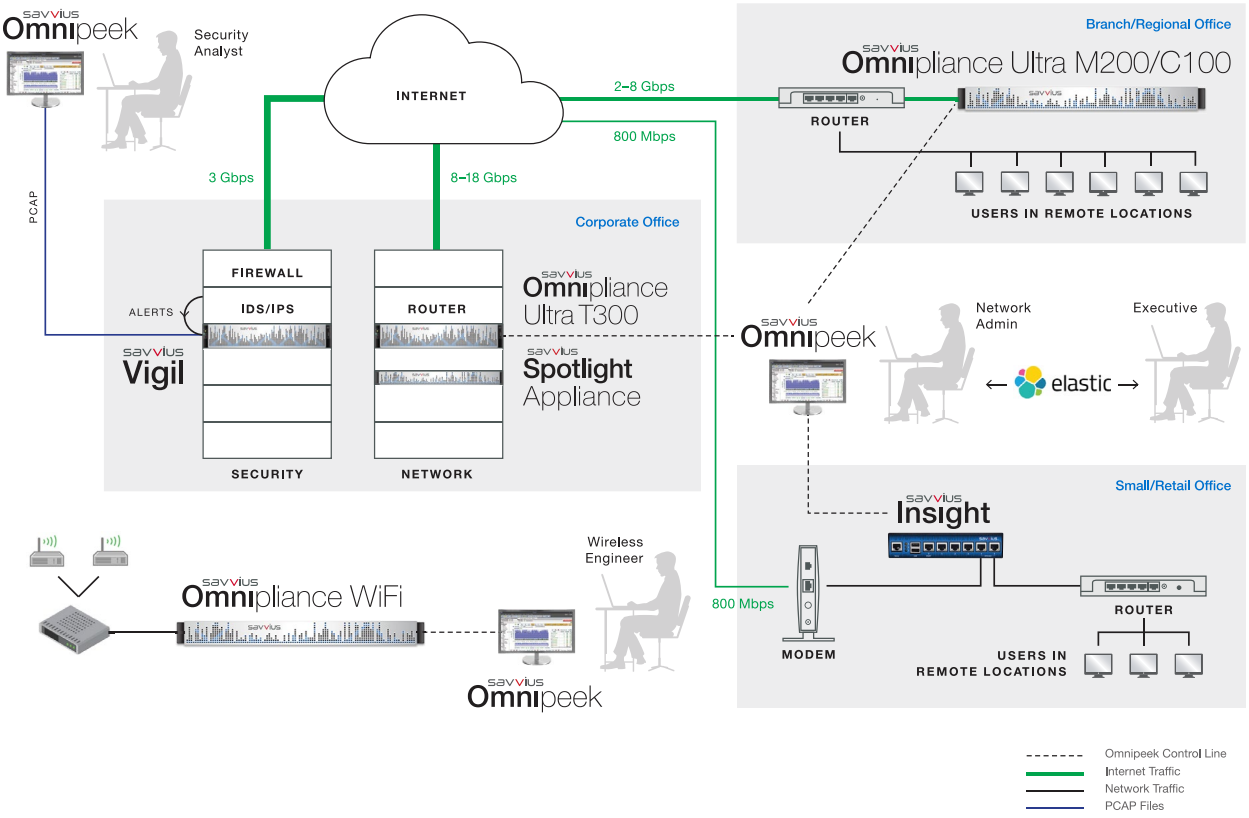

More Resources

You'll find the latest information on industry trends, best practices, and Savvius products here: https://www.savvius.com/resources

About Savvius, Inc.

Savvius sets the standard for actionable network visibility with software and appliance offerings relied on by leading enterprises around the globe. Trusted by network professionals at over 6,000 companies in 60 countries, Savvius solutions provide unparalleled insight into network performance with real-time analysis and seamless packet capture. Visit https://www.savvius.com to learn more about Savvius Omnipliance®, Savvius Omnipliance Ultra™, Savvius Spotlight™ Appliance, Savvius Omnipeek®, Savvius Vigil™, and Savvius Insight™, and to leverage Savvius technology and channel partners. Follow us on Twitter, Facebook, and LinkedIn.