WHITE PAPER

The Importance of Lossless Network Visibility

Why Lossless Visibility Matters

IT security and analytics tools are only as good as the data they are seeing. IT's fundamental challenge is to ensure that the infrastructure behind these tools delivers applications that are reliable, fast, and secure. This means that IT needs total visibility of the network. With the current level of network security threats and a complete dependence by the business on the data network, you cannot afford partial network visibility. You need lossless visibility.

According to a Cost of Data Center Outages study conducted by the Ponemon Institute, the average cost of a network outage is $7,793 per minute.[1] When the network is down or impaired, minutes matter. This is especially important as global IP traffic levels will triple from the 2017 level of 122 EB per month to 396 EB per month in 2022,[2] according to the Cisco Visual Network Index, Forecast and Methodology: 2017-2022. Better network visibility can improve mean time to repair by helping pinpoint problems.

A visibility architecture helps eliminate these concerns by organizing and integrating your monitoring strategy with your security architecture and problem resolution processes. However, there are three fundamental items to consider when creating a visibility architecture:

- Performance matters when you are trying to control costs
- Missing data is an unnecessary security risk
- Missing data leads to longer and more costly troubleshooting efforts

So, how do you eliminate these issues?

Network visibility matters

When the network is down or impaired, minutes matter. This is especially important as global IP traffic levels will triple from the 2017 level of 122 EB per month to 396 EB per month in 2022, … Better network visibility can improve mean time to repair by helping pinpoint problems.

[1] Ponemon Institute – Cost of Data Center Outages, January 2016
[2] Visual Network Index, Forecast and Methodology: 2017-2022. Cisco Systems, 2017

KEYSIGHT TECHNOLOGIES

Monitoring Component Performance Effects Network Visibility

Simply put—performance matters. Poor execution of a brilliant strategy will not get you too far in a hyper-competitive marketplace. However, the positive performance of enterprise data networks can have a profound effect on deals closed and the fulfillment of sales orders.

An enterprise's monitoring solution needs to operate at high performance levels. This includes being able to process data at line rate so that no data packets get dropped. This level of performance is becoming increasingly important as core network speeds move from 10 GE to 40 GE and 100 GE. The monitoring solution needs to support these speeds natively at full speed. This means the data access, network packet brokers (NPBs), and the monitoring tools all need to be able to operate at peak performance.

Your monitoring equipment needs to be able to handle weekly, daily, and hourly fluctuations in traffic load so that you can capture the proper data. Otherwise, you can miss critical data. As illustrated in the image below, it is one thing for the monitoring equipment to operate under a steady state condition, it is another when it is running under loaded conditions that change during the course of time.

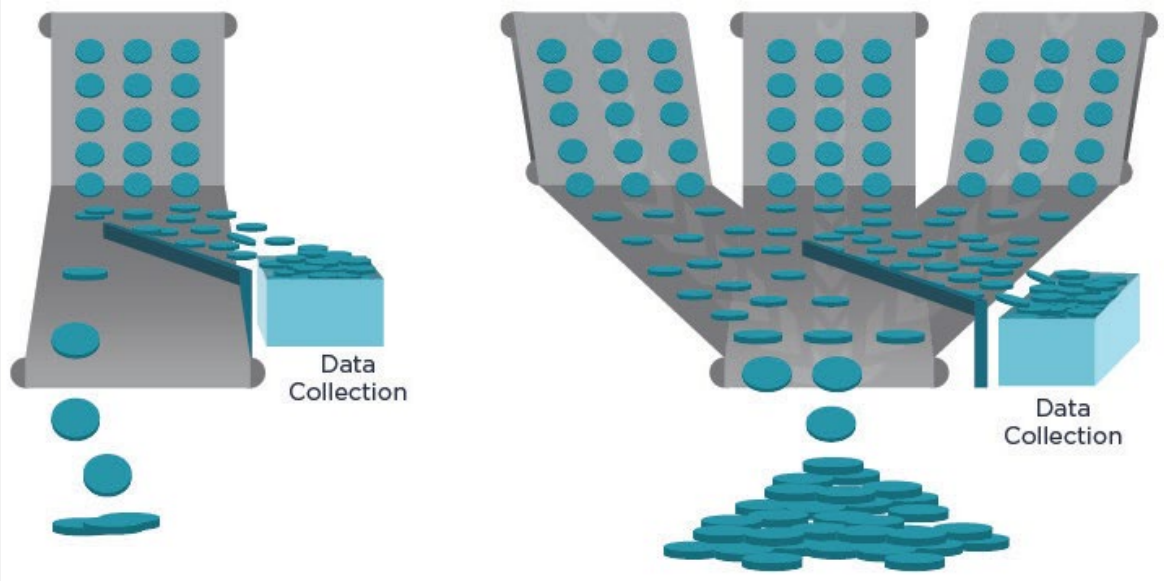

Monitoring data collection at low speeds — some solutions drop packets.

Monitoring data collection at high speeds — small data losses can become serious.

Figure 1. The effect of packet loss at different data speeds.

Unfortunately, IT is often forced to make potentially unwise trade-offs. According to a ZK Research survey, 45% of respondents admitted to turning off features in security devices in order to improve performance.[3] One reason was that while one security tool might be able to handle multiple tasks (intrusion prevention, firewall, malware detection, DDoS, decryption, etc.) like a "Swiss army knife," the processor in the single appliance could not handle the tasks at full load. In other cases, the security tools were being pushed to do non-core tasks (like deduplication, data masking, etc.) which overloaded the tool's CPU and made it run slow.

While performance is definitely a high priority task, is it really more important than network security? Why should IT even be asked to make this trade-off? The answer is you should not, and you do not have to be. You simply need to understand the question, and then you can optimize your monitoring solution.

Another performance need involves creating a non-blocking architecture. Some packet brokers cannot support line rate when multiple features are turned on (e.g. deduplication plus NetFlow or SSL Decrypt in a single module). This is especially true of software based NPB systems. The use of CPU and software-based processing results in a solution capability limited by the CPU's capability. This shortcoming is even more pronounced with oversubscription, where multiple standard ports share a single port's resources for advanced packet processing.

There are superior monitoring solutions available that satisfy the requirements of no dropped packets, no need to turn features off, and a non-blocking architecture. For instance, there are field-programmable gate array (FPGA)-based NPB systems that are able to support line rates with minimal latency and no dropped packets. This is because FPGA-based packet brokers can be purpose-built to process monitoring functions (like deduplication, packet slicing, protocol header stripping, etc.) and still run at line rate. This eliminates the need for performance compromises and trade-offs.

Missing Data Can Lead To Missed Security Threats

Disregarding the performance question for a minute, is your architecture even capturing all the necessary information you need, or is it missing critical information required to capture security-related incidents? According to a research survey[4] by EMA, 78% of respondents say it is very important that their monitoring tools receive all the packets they need. At the same time, 29% said that they are not completely confident that their tools are receiving all of the data. This skepticism is for good reason. The Tolly Group ran a comparison between two network packet brokers and found that one packet broker was indeed dropping packet data and not reporting it. According to the report, the vendor in question "demonstrated packet loss at every data size. At 256-bytes and below, the loss ranged from 20% to nearly 75%."[5] The only thing worse than missing data is not knowing that you are missing the data in the first place.

This missing data is an extremely important concern because data loss means that you can experience false positives and/or miss real positive indicators of a breach. According to the 2021 Verizon Data Breach Investigation Report[6], most victimized companies do not discover security breaches themselves. Approximately 80% have to be informed by law enforcement and 3rd parties (customers, supplier, business partners, etc.) that they have been breached—they had no idea the breach had happened. It's

[3] Zeus Kerravala. "Simplified Programming of a Visibility Layer Can Have a Big Impact on Application Performance," ZK Research 2016.
[4] Enterprise Management Associates conducted research, October 2016
[5] Tolly Group, "Tolly Test Report - Ixia Network Tool Optimizer (NTO) 5288", 2016
[6] 2021 Data Beach Investigations Report. Verizon. 2021.

hard enough to defeat modern network security threats, you do not want to start off with limited network visibility and “fumble around in the dark” to figure out what is missing.

As an example, many security tools need “session stickiness.” This allows them to correlate all components of a session to evaluate the risk and analyze the data for security threats. What happens when data is missing? There are usually two outcomes—both of which are very bad. Outcome 1 is that the device, an intrusion prevention system (IPS) or some other inline security analysis tool, does not detect that the session has closed. If too many sessions remain open, the tool’s memory cannot track anymore sessions. In some cases, the security tool will shift from an “inline blocking mode” state to an “out-of-band detection mode” state. It then sends a trouble alert but ignores additional sessions, allowing them to pass downstream without inspection. This means the device is not actively analyzing those potential security threats. It can also be a manual process for an IT engineer to issue a command to move the tool back to an inline state. Then the engineer needs to perform some sort of analysis to see what triggered the incident, which costs more time and effort.

Outcome 2 is for the security tool, like a web application firewall (WAF), to simply ignore the data when the session does not end. This means it lets the data, which could be malware or some other security threat, just pass on through without warning. In addition, the evidence (packet data and missing data information) is inadvertently thrown away by the device, so now the engineer will not even know about the problem or how to debug the situation.

In a second type of example, the loss of data can actually help hackers hide themselves. For instance, the hacker would start off with a DDoS attack. As the targeted network equipment gets loaded down, security tools and monitoring equipment would get loaded down as well. If the NPB starts dropping packets that it is supposed to send to the security tools, then the loss of packets can provide a type of smoke screen cover for the hacker as he begins probing for weaknesses.

A third example could be an out-of-band intrusion detection system (IDS) that is monitoring data from a switched port analyzer (SPAN) port. SPAN ports are a well-known entity that forwards summarized data, not a complete copy of all data. Bad and missing data is dropped by the SPAN. This can allow an attacker to hide threats, like embedded malware in the packet’s payload, within unmonitored gaps from the SPAN port. In contrast, a tap would have forwarded a complete of all the packets, which could have indicated that there were data gaps that should be analyzed.

These three sets of examples illustrate how missing data can help hide/exacerbate security threats.

Complete Visibility Means Faster Troubleshooting

A complete copy of all your monitoring data also allows you to get to root cause analysis faster. Some solutions, like SPAN ports and command line interface (CLI)-based packet brokers, can clip data or just provide a summarized version. Clipped data is obviously bad because you do not know what is missing, unless you perform an extensive evaluation. Summarized data is just as bad. This means that data considered irrelevant (mal-formed packets, packets with corrupted checksums, etc.) which could actually be very useful for troubleshooting purposes, never gets passed along. The direct effect is an unknown loss of data that often results in delays and misdiagnosis in solving network problems because of missing/misleading information along with missed security threats.

This missing data can lead to longer and more costly troubleshooting efforts. Specifically, it can result in false conclusions, longer resolution times, a poorer quality of experience for users, and lower customer satisfaction ratings. For instance, missing data will look identical to packets dropped over the network. This can quite often generate a “false positive” kind of result where an IT admin or a Network Operations Center (NOC) may end up going through a troubleshooting workflow to solve some underlying problem and immediately conclude that the network is dropping packets as the cause. This is problematic for two reasons. First, they may declare “success”, without actually having identified anything going wrong over the actual network.

Second, the organization may spend considerable time, effort, and money trying to “fix” the packet loss issue because they have assumed it is a network issue. As an example, a first remedial action might be to increase the bandwidth over the observed link. After a certain amount of time and effort, they will discover that this action did not solve anything. The false conclusion actually turned a single problem into two problems now, and their mean time to repair (and possible customer satisfaction) scores are suffering because they still did not fix the initial problem. In addition, the IT admin and team will start to experience a lot of frustration and doubt the monitoring solution because of the false positive.

In a second example, the missing data makes it harder on monitoring tools to perform their job. For instance, using a performance analysis tool as an example, at a certain amount of packet loss the tool will start to fail in its ability to monitor the data. This is because of the heavy drain on memory required to watch for conversations that never complete, since the “end session” data never came through. So, the tool becomes less effective as the memory buffer hits its maximum threshold and the performance data becomes either worthless or of minimal utility.

Another example is that maybe some of the data does end up making it to the tool. However, the data is out of order because buffers reordered the data packets. This can lead to loss as well. Even if the probe can make up for the lost/incorrectly ordered data, it puts a higher load on the tool and detracts from the tool’s core purpose of analyzing data. Tool CPU and memory resources are wasted trying to make corrections for errors that a poorly designed packet broker, or poorly designed visibility architecture, has introduced.

How To Deliver Lossless Visibility

If you cannot use your monitoring equipment to its fullest potential, then why use it? There are several ways to prevent the loss of monitoring data on your network. First, validate existing and future NPB solutions with a traffic generator at load, i.e. at speeds of 20 to 40 Gbps. This is where rubber meets the road.

You cannot go by the statistics listed in a graphical user interface (GUI). You need impartial verification. The truth is that some network packet broker manufacturers cannot support running all of their features together, much less at full line rate. Test your solution at 60% or more load to see the truth. Why buy a monitoring solution only to run it at half speed? Your internal and external customers will not settle for excuses.

A second consideration is to look for solutions that use FPGA components. You want to run all of the NPB features at line rate, correct? FPGAs are purpose-built microprocessors that can be programmed to focus on specialized activities. This gives them a performance advantage over CPUs, especially when it comes to advanced feature processing capability like packet deduplication, protocol header stripping, packet trimming, data masking, and timestamping. A CPU and software approach to performing these functions has inherent issues since every line of code steals cycles, making the CPU plus software approach slower. Faster is better.

A third consideration is to make sure that your monitoring solution has a GUI that is intuitive to use. Data from ZK Research shows that 20% of CLI filters created have errors in them.[7] Many self-inflicted performance, security, and troubleshooting errors can be avoided. A GUI that uses point and click, drag and drop technology eliminates this error source. Why add more complexity and more effort to your workload? A GUI interface makes life easier and less error prone.

Conclusion

Not all network packet brokers are created equal. NPB technology continues to rapidly evolve with increasing network and security requirements. To maximize the performance of any solution, consider evaluating the actual performance (with 60% or higher load) for current and prospective NPBs. You should also consider investigating if the packet broker is a zero-loss solution and whether it uses a GUI interface for all management functions.

The Keysight architecture uses FPGAs to process data (instead of a CPU running software) and has a fully integrated GUI. This allows our customers to operate without any restrictions at line rate. They can monitor their network at full throttle, giving the business a competitive edge.

[7] Zeus Kerravala. "Simplified Programming of a Visibility Layer Can Have a Big Impact on Application Performance," ZK Research 2016.

Learn more at: www.keysight.com

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

KEYSIGHT TECHNOLOGIES

This information is subject to change without notice. © Keysight Technologies, 2016-2021, Published in USA, July 20, 2021, 3121-1283.EN