Scaling Network Traffic Visibility to 100 Gbps

Gordon Beith October 20, 2014 No Comments »

As enterprise networks scale to 40 Gbps and 100 Gbps, network managers encounter a serious problem: many of today's traffic monitoring, management and security tools are incapable of operating well, or even at all, at these high data rates. Complicating this challenge is the need to have a sufficient number of interface ports, supporting data rates ranging from 100 Mbps to 100 Gbps, to access all of network links needed to get full visibility into traffic and sessions.

To accommodate these needs, network TAPs and network packet brokers (NPBs) must be able to scale to higher port densities. Also, NPBs need to work both smarter and faster to be able to capture, consolidate, manipulate and otherwise optimize traffic in ways that enable existing 1 Gbps and 10 Gbps, and even some of the newer 40/100 Gbps, traffic monitoring, management and security tools to operate with full effectiveness in the 40/100 Gbps networks that most data centers employ.

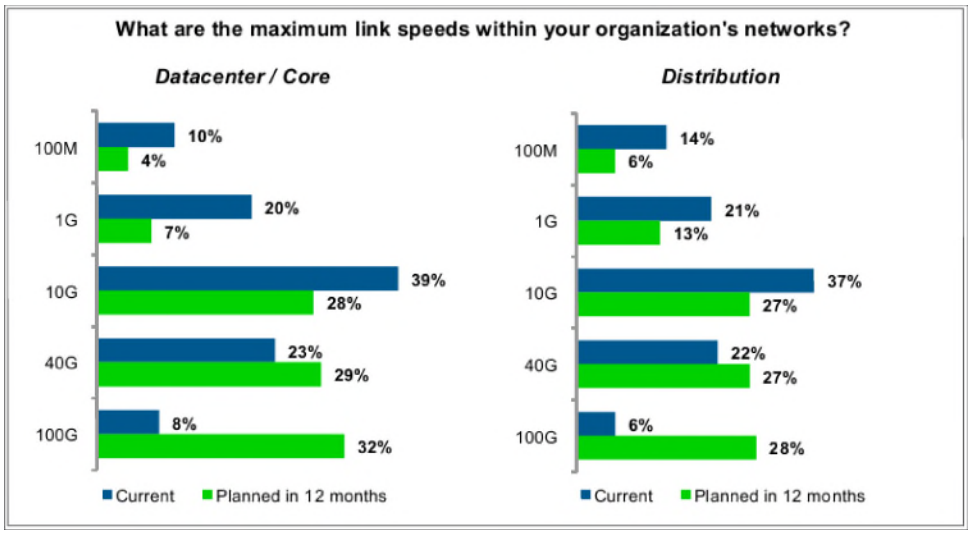

Enterprise Management Associates discovered that planned use of 1 Gbps and 10 Gbps is diminishing as networks move to 40 Gbps and 100 Gbps (October 2013 survey).

The potential problems caused by the lack of full traffic visibility, when using 1/10 Gbps tools to monitor, manage and secure 40/100 Gbps networks, goes well beyond protecting the investment in those tools. The old adage that "you can't manage what you can't measure" (or monitor in this case) applies fully to the network performance management (NPM) and application performance management (APM) systems that data centers use today. Here, the lack of full visibility inevitably limits the effectiveness of these systems.

According to Enterprise Management Associates, "Solutions for packet-based network monitoring and security technologies have only recently reached maturity for fully loaded 10G networks, and as of mid-2014 were only beginning to take the very first incremental steps toward making those products work in 40G settings. Even then, maximum throughputs of monitoring and security technologies will not support a fully loaded 40G link. As of mid-2014, none of those technologies were capable of operating directly in a 100G environment—the best, in fact, were just reaching 25 Gbps of total processing capability."

Further exacerbating the scalability challenge is the widespread virtualization of servers, storage and networks that is adding enormous complexity to data center infrastructure. Packet data is increasing exponentially in some environments, but only packet-level data can provide the full visibility and granularity needed by some of the more sophisticated forensic and troubleshooting applications that both physical and virtual environments employ. These tools must also contend for a limited number of the available switched port analyzer (SPAN) or mirroring ports available on switches.

As with many challenges, success in scaling network visibility both holistically and end to end requires a solid foundation.

The Visibility Foundation: TAPs and NPBs

Although the data rates are higher and the traffic flows are more complex today, the problems associated with providing full network visibility at the packet level were solved long ago with TAPs and NPBs. And it is the latest generation of these systems that are now making it possible to scale visibility into 40 Gbps and 100 Gbps networks.

A basic TAP is a passive device that copies or replicates traffic flows in their entirety from network links to monitoring output ports on a 1:1 basis. Some TAPs can aggregate traffic from multiple network ports to a single monitoring port (N:1) and/or regenerate traffic from a single network port to multiple monitoring ports (1:N). But because TAPs can only replicate, aggregate and regenerate traffic, they lack the intelligent traffic-optimization capabilities to enable 1 Gbps and 10 Gbps tools to monitor 40/100 Gbps traffic flows.

Intelligent traffic optimization is available in so-called smart TAPs, which are more commonly called network packet brokers. As the name implies, NPBs serve in a third-party mediator or "broker" role capable of mapping packet flows from any network TAP, switch or router port to any monitoring and management tool's port.

Although NPBs can replicate, aggregate and regenerate traffic like TAPs do (but in a far more selective manner), their real value derives from their ability to preprocess or transform packet flows and even individual packets from any network port to any monitoring port in ways that enable the attached tools to function more efficiently and effectively. For instance, they can balance traffic for multiple instances of the same tool. In effect, NPBs create a layer of abstraction and intelligence between the network infrastructure and the tools that monitor, manage and secure the various network functions and networked applications.

The visibility fabric that NPBs typically create results in a separate, out-of-band monitoring and management network, but some NPBs can also be used in line or in band to support active network security applications and other in-line applications, such as WAN optimization. Although an NPB has multiple ports and performs basic switching functions, like packet inspection and forwarding, an NPB must also be able to optimize traffic flows by manipulating packets in ways that switches do not. These additional capabilities include full Layer 2–7 inspection, filtering, slicing, reordering, deduplication, fragment reassembly, stripping off protocol headers, and accurate time and port stamping. Some of these techniques for optimizing traffic are what make the NPB essential to scaling network visibility.

Optimizing Traffic to Scale Network Visibility

Network packet brokers have long been able to help organizations move to higher network speeds—from 10 Mbps to 100 Mbps to 1Gbps to 10 Gbps, and now to 40 Gbps and 100 Gbps with the latest generation. Described here are some of the traffic-optimization techniques that together make it possible to use 1/10 Gbps tools in 40/100 Gbps networks.

Speed conversion: As noted by Enterprise Management Associates, the performance of network monitoring and management tools lags the data rates of the latest-generation switches and routers. Speed conversion is required to enable tools with only a 1 Gbps or 10 Gbps interface to work with traffic traversing 40 Gbps or 100 Gbps links. It is important to note that speed conversion involves more than simply throttling back a traffic flow, which would inevitably result in buffer overflow and packet loss. So one or more of the other techniques described here is also needed to enable full visibility at full line rates.

Packet filtering: Full visibility does not require that every tool see every packet. Indeed, getting "slow" tools to monitor "fast" networks requires that the tools somehow divide and share the workload. Packets are normally filtered on the basis of Layer 2–4 header information (or a little beyond Layer 4) to separate the individual flows in a data stream by source/destination address, tunnel, protocol type and so on. Some NPBs can filter at higher layers to isolate individual sessions or specific applications. Filtering can also be used to discard any packets that are unnecessary for or uninteresting to a particular monitoring or management application, such as Address Resolution Protocol (ARP) requests.

Content matching: Content matching, also known as deep packet content-based filtering, is a more in-depth filtering capability available in some of the more advanced NPBs. Content matching scans the packet's entire payload for data that is pertinent to an application. It then forwards that traffic, and only that traffic, to the tool; conversely, it matches traffic that is unneeded and forwards all other traffic only. The content-matching rules can be customized by the user to be as broad or narrow as desired. This advanced capability is especially powerful in troubleshooting, analytics and security applications.

Protocol stripping: Protocol stripping or de-encapsulation removes the header or headers associated with tunneling protocols, such as Cisco's FabricPath, GRE, GTP, MAC-in-MAC, MPLS, TRILL, VLANs, VNTags and VXLANs. This capability allows tools that do not need or support these protocols or encapsulations to still be able to analyze the traffic.

Packet slicing: Packet slicing discards any unnecessary portion of packets from a defined position, such as a number of bytes from the beginning or end of the TCP header. Some NPBs can conditionally slice packets by identifying the specific packet, on the basis of Layer 2–4 rules, and only slice those packets as directed. A common use for packet slicing is the elimination of the packet's main payload content, which significantly reduces the amount of traffic that is forwarded to tools for certain applications. And in those applications where payloads cannot be stored or even seen for security or privacy reasons, eliminating the protected content before it reaches the tools can help organizations comply with pertinent security regulations and/or privacy policies.

Packet deduplication: Networks deduplicate traffic for a variety of reasons, including planned redundancies in the network's design, use of switch SPAN (mirror), load-balancing switches and traffic capture from multiple network segments. The stripping out of all duplicate packets by the NPB can make any tool more efficient, and it can also help make some security applications more effective.

Fragment reassembly: Traffic becomes fragmented in networks for two main reasons: the networks' maximum transmission unit (MTU) settings, and tunneling of jumbo TCP packets. Traffic fragmentation can prevent all fragments from being forwarded on the basis of Layer 4 or higher criteria, and it places a significant load on monitoring and security tools to reassemble for analysis. Having the NPB reassemble the fragmented traffic eliminates these and other potential problems.

Load balancing: Session-based, flow-aware load balancing is perhaps the most important role network packet brokers play in scaling network visibility. Load-balancing logically binds multiple monitoring ports, allowing traffic to be

mapped to a group rather than being manually split across individual ports. Once the traffic is mapped to a group, it is automatically spread across the ports in accordance with user-defined flow criteria. Load-balancing traffic across all of the group's monitor ports helps prevent any overflow on any single port, ensures each port (tool) receives every packet in each flow and, when configured for redundancy, can accommodate a failure in one or more of the tools attached.

By minimizing the amount of traffic and number of packets being forwarded along to monitoring and management tools, and by load-balancing what is forwarded among multiple tools, the network packet broker makes it possible to scale network visibility in even the largest of data centers.

Conclusion

As organizations move to 40 Gbps and 100 Gbps networks in their data centers, they will be confronted by challenges, both familiar and new. Network packet brokers are the solution to one of those challenges, being able to scale network visibility by optimizing traffic flows in ways that enable existing 1 Gbps and 10 Gbps monitoring, management and security tools to function more efficiently and effectively. Of course, the network visibility fabric or plane that the NPBs create must also be able to scale by adding more ports and more units as needed to maintain full visibility top to bottom and end to end.

Leading article image courtesy of error_in_excellence

About the Author

Gordon Beith is director of product management at VSS Monitoring, overseeing management of the company's entire product line. His extensive career includes over 25 years of work in telecommunications product management, marketing, hardware and software development, and related next-generation services. In addition to his work with VSS Monitoring, his background includes product, project, marketing and account-management positions for IneoQuest, Empirix, Spirent, Cisco and Ericsson in global locations including North America, Europe, Australia and China.