Whitepaper

Gigamon Intelligent Flow Mapping

In today's competitive world where more and more business-critical applications are moving from the physical confines of the corporate organization to the Internet, these applications are expected to be available and accessible from anywhere, and at any time. From home, work, and all points between, we are leading "always-connected" digital lifestyles. The responsiveness and availability of business-critical applications and essential IT services are of paramount concern to IT organizations everywhere. As the market moves to an always-connected existence, the insatiable demands from customers for ever-higher bandwidth, lower latency, and superior user experience require the communications industry to transform itself and IT organizations to rethink how they design and manage their networks.

Obtaining high-quality, fine-grained relevant data is more important than ever for real-time insights into end-to-end application interdependencies. To this end, service providers and IT organizations are turning to the best-in-class, end-to-end visibility and performance analysis tools to effectively manage and monitor the security and performance of their infrastructure.

Introduction

The rise of multimedia, social media, and mobility is fueling an exponential growth in data (aka Big Data). Service providers and IT organizations are investing in monitoring solutions that interpret consumer behavior, detect fraud, monitor performance, and even predict the future with trending analysis. However all of these systems are only as effective as the information and traffic that they can see. Limit visibility to the traffic, and the value of these systems is equally limited. With the increasing volume and detail of information being moved across the infrastructure, these tools find themselves drowning under the volume of traffic—traffic that might not even be relevant to the tool!

It is not just the volume that defines Big Data—it is the velocity, variety, and complexity of that data. Between now and 2020, the volume of digital information is predicted to increase to 35 trillion gigabytes—much of it coming from new sources including social media, internet search, and sensor networks, as well as from existing video traffic and new types of mobile video services. It is all about finding a needle of value in a haystack of unstructured information. With millions of traffic flows and thousands of changes occurring within the infrastructure on a daily basis, visibility needs to be pervasive, dynamic, and scalable. A key factor in ensuring application availability and network performance is having a traffic visibility solution that can efficiently handle huge volumes of data in real time and thus deliver relevant traffic to the relevant tool.

Visibility solutions can differ greatly, employing a variety of filtering mechanisms with varying degrees of efficiency and performance to deliver the desired set of packets to one or more monitoring tools. However, with the magnitude and complexity of current network infrastructures, the challenge is to develop visibility solutions that can scale to allow thousands of diverse traffic streams originating from dozens or hundreds of network traffic sources to be granularly filtered and forwarded to a variety of monitoring tools and analyzers with zero packet loss. In this age of Big Data, efficient and scalable distribution of traffic is key for the monitoring tools and analyzers to focus on the relevant traffic they were optimized for.

Connection-based Traffic Filtering

Traditional approaches to visibility typically employ traffic forwarding based on statically defined connections. These connections are simple one-to-one flows between network and tool ports where traffic can be filtered with allow or deny operations at both the ingress and egress sides of the connection, thus achieving simplistic packet distribution.

On the surface, it may appear that this method provides sufficient flexibility to achieve the desired packet distribution. Closer examination reveals significant limitations that cripple the device's capability to work with large volumes of data and large numbers of distinct traffic streams. In some cases, connection-based filtering can be inadequate even with a single, moderately loaded ingress network feed when sending traffic from high-speed network segments to low bandwidth tools.

Ingress Filters

Ingress filters, also known as pre-filters, are used to allow or deny traffic on network or ingress ports. Any traffic allowed by the filter is sent to all the tool ports at the other end of the connection. This is adequate when all the tools need to view the exact same packet streams, but allows no flexibility or scalability. Connections with ingress filters send all tools the same packets whether they need them or not. For example, referring to Figure 1, if the tool on port A is monitoring Web traffic and the tool on port B is monitoring VoIP traffic, both Web and VoIP traffic must be sent

to both tool ports. Not only does this waste the tools' processing and storage resources on unwanted packets, but also it can oversubscribe the tool port. If the combined traffic exceeds the tool port capacity, the device will indiscriminately discard excess traffic and the tools will not get all the packets they should.

Another way to utilize ingress filters is to drop specific types of traffic and forwarding the rest to the monitoring tools. This prevents the dropped traffic from being used by other tools, like tool C in Figure 1, and other user groups who may need to monitor it. If the drop filters allow too much traffic through, the broadcast and oversubscription problems described above occur.

For these reasons, ingress filtering is a very poor method of managing visibility traffic and it is not recommended for accurate monitoring solutions.

Egress Filters

Egress filters, also known as post-filters, allow only specific traffic to be sent out to the monitoring tools for analysis. When compared to ingress filters, egress filters provide more granular control over the traffic flows sent out to the monitoring tools. However, if a tool is sent traffic from multiple network ports, the egress ports can become oversubscribed and drop packets before they can be filtered. Relying on egress filters to handle all the filtering needs requires multicasting all the traffic to all the egress ports, which can overload the backplane capacity of the visibility node, resulting in dropped packets. Lastly, the number of egress filters that can be configured on the system is often limited, further diminishing any value that egress filters have to offer.

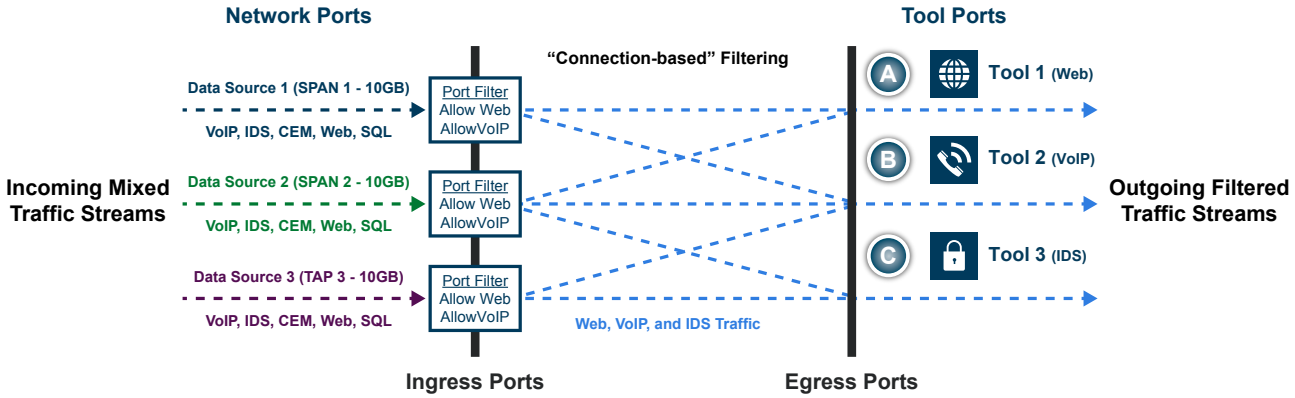

Figure 1: Ingress filters

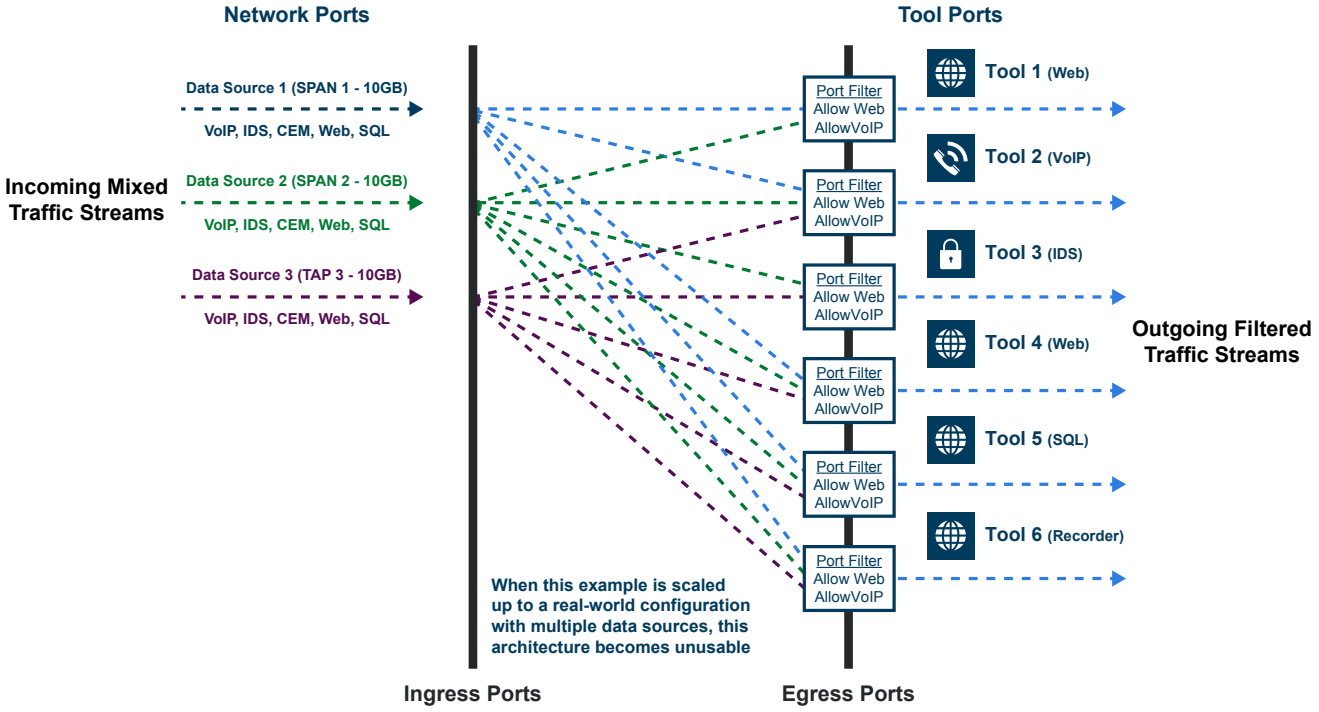

Figure 2: Egress filters

Egress filters can be a useful and easy means of narrowing down traffic sent to monitoring tools, but they have inherent scalability limitations and are not a satisfactory traffic control mechanism for full visibility. Combining ingress and egress filters can mitigate some of the limitations of each approach. Such solutions quickly encounter issues of scalability and flexibility, resulting in a rigid, overly structured and underengineered solution that does not meet the visibility demands of today's complex, diverse, and constantly evolving networking infrastructures.

Flow Mapping Technology

Flow Mapping® is a technology that takes line-rate traffic at 1Gb, 10Gb, 40Gb, or 100Gb from a network TAP or a SPAN/mirror port and sends it through a set of user-defined rules to the tools and applications that secure, monitor and analyze IT infrastructure. Flow Mapping technology provides superior granularity and scalability above and beyond the capabilities of connection- and filter-based technologies by addressing the problems inherent when you go beyond small numbers of connections or when where more than one traffic distribution rule is required.

In the example illustrated in Figure 1, the user would define a map with rules that takes traffic from the three data sources and directs all Web traffic to Tool 1 and all VoIP traffic to Tool 2. Since all of the forwarding decisions are made after the ingress port has delivered the packet to the Flow Mapping engine, ingress filters are not needed and no extraneous traffic is ever forwarded to any tool port. Egress filters can be added, but they are not necessary.

Granular Control over Distinct Traffic Streams

With Flow Mapping technology, users can combine thousands of map rules (each with multiple packet criteria) to achieve the exact packet distribution desired. Each rule provides the ability to configure up to 13 unique criteria based on over 30 predefined Layer 2, Layer 3, and Layer 4 parameters including IPv4/IPv6 addresses, application port numbers, VLAN IDs, MAC addresses, and more. Additionally, users can define custom rules that match specific bit sequences in the traffic streams, letting users apply Flow Mapping to tunneled traffic, specialized applications, and even higherlayer protocols.

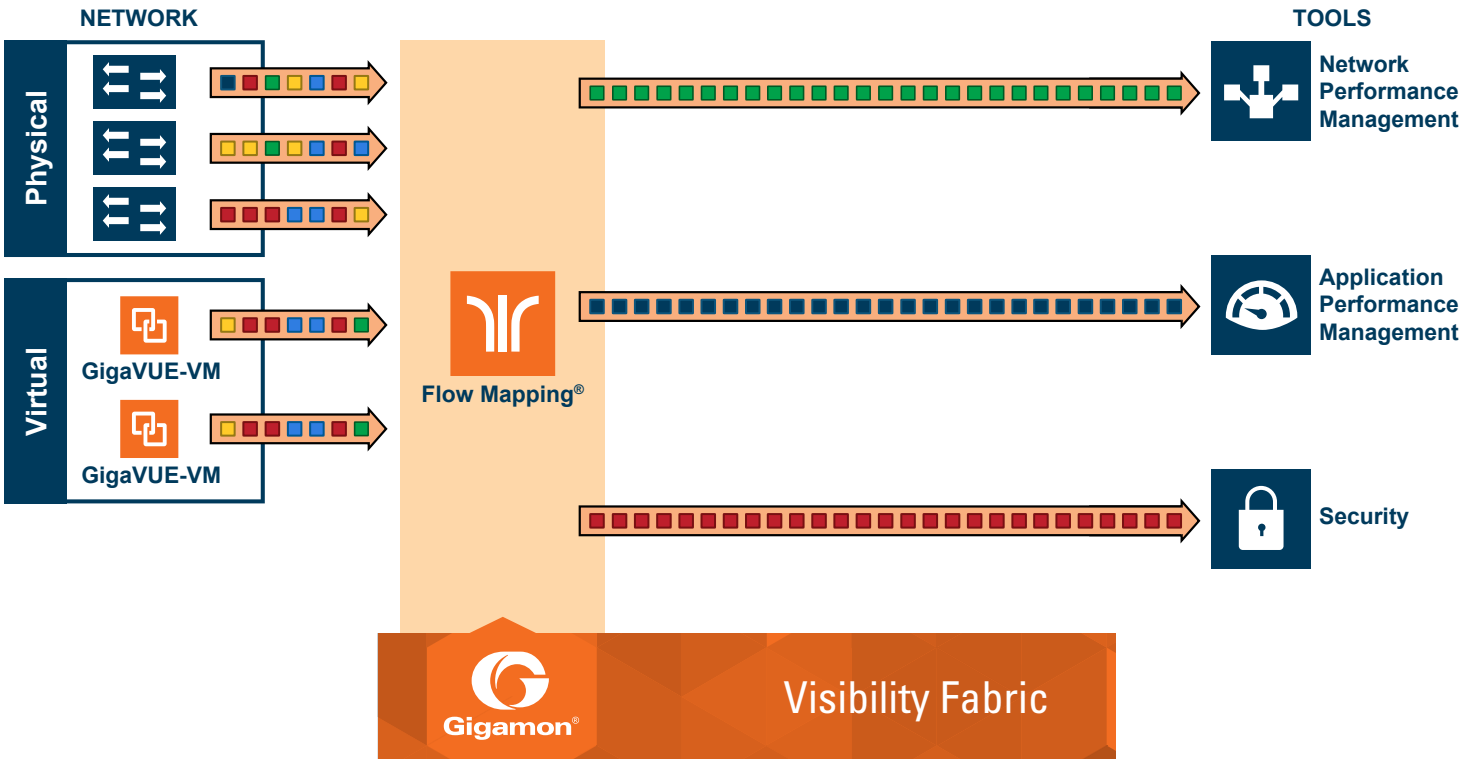

Figure 3: The Gigamon Flow Mapping technology

Optimize Tool Performance

Applying maps to data ensures that each tool sees only the traffic that best suits its individual strengths and nothing else. Tools are made more efficient since they are presented with only the traffic they need to see—therefore maximizing their effective throughput and being better able to process more of the Big Data load per connected tool.

Visibility into All Traffic

Users can also define a collector rule, an Everything Else bucket. Packets that don't match the criteria specified by any of the other map rules are sent to the designated collector port, or they can be simply dropped. Collectors are a simple way to ensure all traffic is available for inspection or archiving, even though there is not a dedicated tool for that traffic.

Replicate or Distribute Traffic across Multiple Tools

Map rules can replicate traffic to multiple monitoring tools. For example, the same web traffic can be captured by both an application performance analyzer and an intrusion detection system. Additionally, traffic can be distributed across multiple tools rather than overburden a single tool. Each tool process a subset of the traffic, based on IP address, application port, or other user-configurable distribution options. As network speeds increase, older, slower tools can be grouped to extend their life and reduce tool expenditures.

Visibility as a Service

Flow Mapping allows the capability for different user groups to access the same data and send it to their specific tools. IT operations is often organized into groups that support specific business functions, such as finance, supply-chain management, marketing, and sales, or various IT teams such as the network, security, desktop, server, or application performance teams. Each of these sub-teams delivers analytics and reporting to their specific business stakeholders. Flow Mapping with overlapping map rules allows the IT operations team to offer Visibility as a Service (VaaS) to different groups within the enterprise. Role-based access controls (RBAC) allow IT teams supporting specific user groups to define their own maps, independent of other departments' maps, and ensure that each business unit only receives data pertinent to their business needs. RBAC allows security units to define white groups have visibility to which traffic flows.

Industry-Leading Scalability enabled by Flow Mapping

Flow Mapping allows every network port to receive 100% line-rate traffic while each tool port can output relevant traffic up to 100% of the port's capacity. Maps can be associated with any number of network ports and are not limited by the number of available tool port filters. The hardware-based design avoids the bottlenecks and throughput limitations of processor-based solutions. With Flow Mapping, more network ports can send desired traffic to each tool port and every tool can see more traffic than otherwise would be possible, overcoming the core problems associated with Big Data.

Gigamon's Unified Visibility Fabric allows multiple discrete Visibility Fabric nodes to be combined into a single manage-as-one fabric. Users are not limited to the port capacity of a single chassis. Maps can combine network and tool ports from multiple chassis, extending traffic visibility beyond a single rack, row, or data center. Nodes can be connected in star, daisy chain, or hybrid configurations, moving as much as 320Gb of bidirectional traffic between each node. No other visibility solution provides the intelligence and scalability to manage Big Data traffic and optimize tool processing.

Packet Manipulation and Tool Optimization enabled by GigaSMART and Flow Mapping

In addition to providing access to critical information, the packet distribution capabilities of the Unified Visibility Fabric can be combined with GigaSMART® technology to process and optimize the filtered traffic streams before they are sent out to the monitoring tools. Features such as stripping extraneous headers, removal of duplicate packets in the incoming streams, and slicing packets to remove superfluous or private information can be used to optimize tool performance and improve monitoring accuracy as well as allow for greater integration between the tool layer and the data access layer. Incoming traffic streams can also be time stamped closer to the source, allowing performance monitoring tools to calculate end-to-end latency and jitter, while preserving link-layer visibility.

With many visibility solutions, these advanced packet manipulation features are typically applied to all the incoming traffic and are often limited to a subset of ports on the chassis. Using Flow Mapping, incoming traffic on any ingress port within the Unified Visibility Fabric can be directed to a GigaSMART operation. Since the GigaSMART operations are tied to the map rules, the end user has the flexibility to granularly control the traffic flows over which the GigaSMART operations are applied. This improves the throughput of tools by allowing the tool to see only the traffic of interest to it and by eliminating the manual steps needed to format the data so tool processor parsing cycles can be reduced. Thus, each tool is better able to address more of the Big Data load it is presented with (see Figure 4).

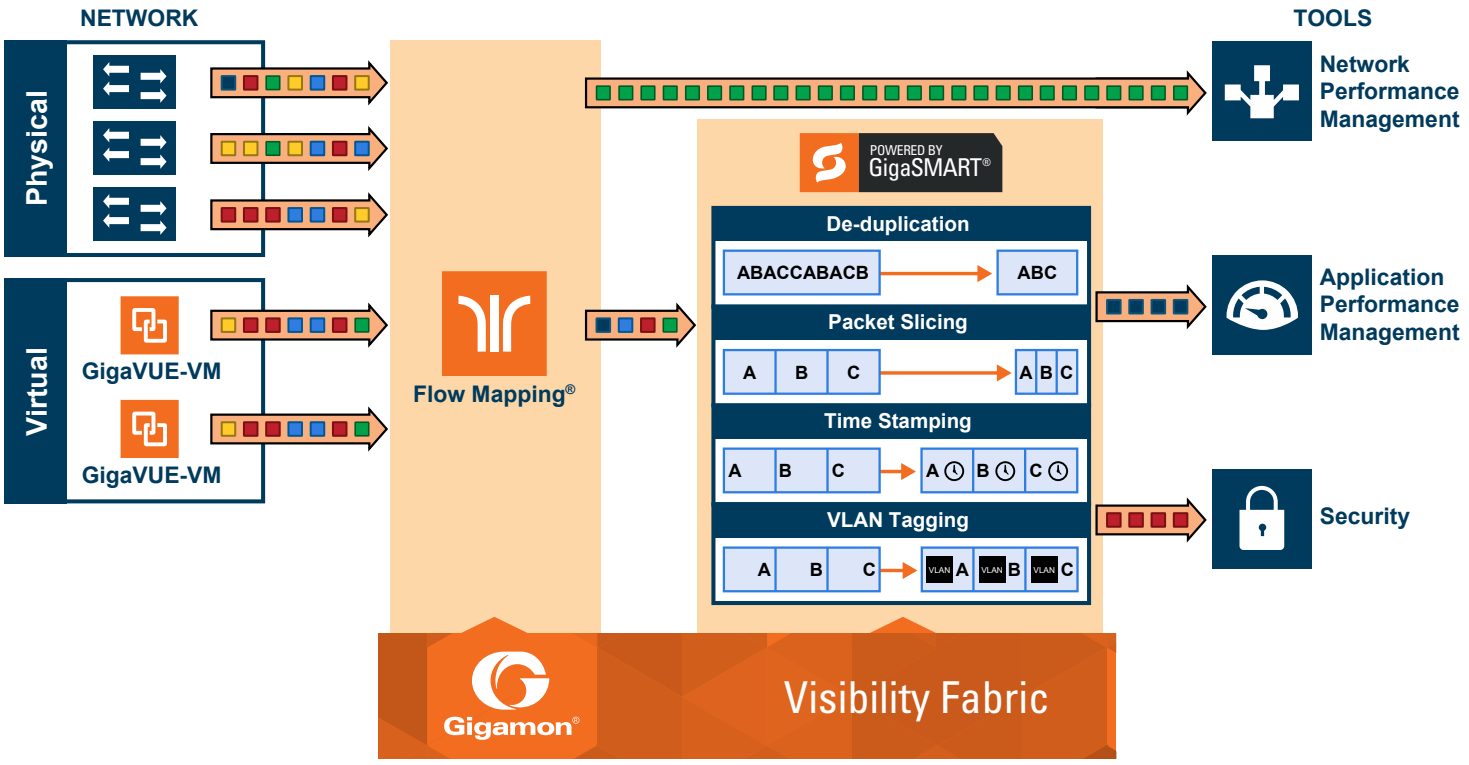

Figure 4: Packet manipulation and tool optimization enabled by Gigamon's GigaSMART and Flow Mapping technologies

Conclusion

While new web-based applications and mobile devices continue to help businesses improve productivity and empower employees, the challenges related to data mobility, complexity, and volume continue to vex organizations. The responsiveness and availability of these business applications become even more critical in the face of ever-evolving IT infrastructures and usage models. With the exponential growth in applications and end users leveraging these applications, data traffic will continue to grow more than ever along with the need to monitor that traffic.

With Gigamon's unique patented Flow Mapping at the heart of the Unified Visibility Fabric, traffic streams ranging from 1Gb up to 100Gb and flowing across both physical and virtual networks can be granularly filtered and aggregated before being replicated into management tools including application and network performance (APM/NPM), security tools like intrusion detection and prevention systems (IDS/IPS), customer experience management (CEM), data loss prevention (DLP), network forensics, and analyzers. Operators can now direct data from any number of access points to any number of monitoring tools at line rate without data loss—taking on the issues of Big Data head on. With end-to-end, access-to-core visibility and detailed analysis of performance impacting events, operators are empowered to proactively maintain a subscriber's quality of experience (QoE) while securing the integrity of the network and satisfying the issues brought on by continuing and increasing amounts of subscriber data.

About Gigamon

Gigamon provides an intelligent Unified Visibility Fabric™ to enable the management of increasingly complex networks. Gigamon technology empowers infrastructure architects, managers and operators with pervasive visibility and control of traffic across both physical and virtual environments without affecting the performance or stability of the production network. Through patented technologies, centralized management and a portfolio of high availability and high density fabric nodes, network traffic is intelligently delivered to management, monitoring and security systems. Gigamon solutions have been deployed globally across enterprise, data centers and service providers, including over half of the Fortune 100 and many government and federal agencies.

For more information about the Gigamon Unified Visibility Fabric visit: **www.gigamon.com**

3039-04 04/15

Gigamon® 3300 Olcott Street, Santa Clara, CA 95054 USA | +1 (408) 831-4000 | www.gigamon.com