Executive Summary

IT organizations today face unprecedented challenges. Internal business customers continue to demand rapid delivery of innovative services to respond to outside threats and opportunities. At the same time, it is not unusual for organizations to issue broad mandates to cut back on spending budgets. In these tough times, the adoption of disruptive technologies holds the key to improving efficiency and truly accomplishing more with less. Cloud computing and virtualization as a technology is one such innovation that allows creation of a more dynamic and flexible infrastructure by maximizing resource utilization while increasing IT service delivery.

Virtualization stands to bring enormous cost savings by streamlining server management while also improving the efficiency in space and power usage. With virtualization technology in place, enterprises can be more agile than they have been in the past. In one series of case studies presented by VMware, for example, companies using virtualization achieved a 67% reduction in total cost of ownership for IT operations after implementation.[1] It is therefore no wonder that adoption of virtualization is proceeding at a rapid rate and is most likely being accelerated even further by tough economic times and cost cutting mandates. With virtualization as the driving technology standard, public reports estimate that in five years, the cloud market will exceed $300 Billion.[2]

Even with the undeniable cost and scalability benefits of virtualization and cloud computing, the elastic and dynamic nature of software-defined cloud computing introduces a whole new set of challenges to IT professionals. While responsiveness to immediate needs has improved, diagnosing problems and analyzing performance has become more complex. With more of the path of application data being shrouded in virtual networks, managing and monitoring network operations with traditional approaches is becoming difficult. IT leaders and stakeholders are constantly struggling with gaining back visibility, maintaining and improving application performance, and enforcing corporate regulatory policies across this new type of network while also leveraging benefits brought by virtualization. Their challenge is to adjust expectations based upon “exaggeration” and “hype” so that the true value and return on investment are understood.

Introduction

Virtualization offers compelling benefits. The financial justification comes from maximizing capacity utilization by hosting more and more virtual machines (VMs) on a single server thereby enabling cost savings, management flexibility, and business agility—to name just a few. With a literal click of a button, a new server can be deployed and be in production in minutes, often at no hard cost to the organization. However, as workloads and servers become virtualized, the tools once used to monitor, analyze, and secure data center assets and network traffic are now dark, unable to see beyond the physical links connecting the server iron. Once monolithic, large binary applications are now distributed with modern scripting languages, Java, APIs, and running in distributed function architecture with REST, JSON and Ruby. This change to the distributed application architecture has created more and more East/West traffic handling application calls, and much of that inter-application traffic is also riding encapsulated overlay networks. As virtualization occurs, visibility for tools recedes.

[1]VMware, Reducing Server Total Cost of Ownership with VMware Virtualization Software, http://www.isdsecurity.com/VMWARETCOWhitepaper.pdf (2006).
[2]Market Info Group, The Future of Virtualization, Cloud Computing & Green IT Global Technologies & Markets Outlook – 2011-2016 (October, 2010).

© 2012-2014 Gigamon. All rights reserved.

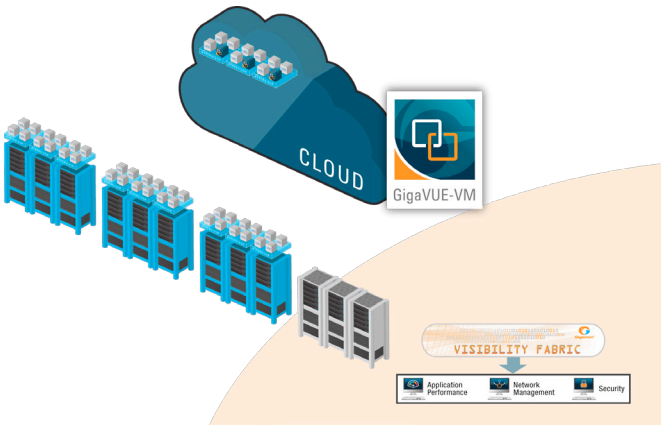

Figure 1: As virtualization occurs, visibility for tools recedes

Intra-host Traffic

Virtualization is creating blind spots, or invisible networks, within the server infrastructure. With a large share of traffic flowing across software-defined cloud infrastructure, being encapsulated across virtual tunnel endpoints, and in many cases not even hitting the physical network at all, VM and network administrators are losing the visibility and control over this communication. This lack of comprehensive visibility is causing reticence among IT professionals seeking to wrap up complex workloads in virtualized infrastructure.

With security and compliance being top of mind in virtualization and cloud deployments, organizations are struggling with how to reconcile competing priorities to virtualize their environments, while still satisfying the existing requirements for visibility. It should therefore come as no surprise to see traffic visibility, compliance, and data security consistently listed among the top inhibitors to cloud adoption.

vMotion

Developed by VMware, vSphere vMotion technology enables the live migration of running virtual machines from one physical server to another. vMotion allows the creation of a dynamic, automated, and self-optimizing data center with continuous and automatic optimization of virtual machines. This technology, which includes fault tolerance, high availability, and DRS, is the first step towards keeping downtime to a minimum. However, with this enhanced agility come changes to the server infrastructure as a completely new layer of complexity is added. To monitor these environments effectively, administrators need to ensure monitoring can seamlessly and automatically be updated to reflect these changes. Further, monitoring solutions need the capability to retain monitoring continuity and history, so administrators can better track and assess these issues over time and set better policies. Without the capability to track and monitor these vMotion events as VMs get reallocated, the resulting configurations can potentially go askew impacting the availability and performance of application and services.

vSwitch

The most common way to provide Virtual Machine (VM) switching connectivity is a Virtual Ethernet Bridge (VEB), commonly referred to as a vSwitch. A vSwitch is a software component associated with a hypervisor that functions like a Layer 2 hardware switch providing inbound/outbound and inter-VM communication. By default, every VM can communicate directly with every other VM on the same host through the simple virtual switch, without any inter-VM traffic monitoring or policy-based inspection and filtering. Intra-host VM traffic, handled internally by the vSwitch, does not transit the physical network. This communication is not visible to many network-based security and monitoring appliances residing outside a virtual server.

As a result, consolidating multiple physical servers into a single virtual server platform significantly impacts all of the network and application monitoring, firewall, intrusion detection, and other compliance tools that were in place prior to the physical to virtual migration. Simply put, traditional network monitoring and security measures may be unable to effectively manage the growing volume of inter-VM traffic, leaving VMs highly vulnerable to attack. This lack of visibility complicates fault isolation and resolution, potentially erasing cost savings associated with the physical to virtual migration.

Visibility of Inter-VM Traffic on a Single Host

Enabling Visibility with GigaVUE-VM Visibility Fabric™ Node

As mission-critical workloads migrate to virtual servers, an increasingly large share of critical network traffic is occurring between VMs residing on the same host. Visibility into this virtual switching infrastructure becomes critical to managing end-to-end service delivery. A solution is therefore required to push only interesting data streams flowing between virtual machines, on the same host, out to external monitoring tools without introducing any security concerns. The Gigamon® GigaVUE-VM Visibility Fabric node addresses these requirements by providing an intelligent filtering technology allowing specific inter-VM traffic flows of interest to be selected,

© 2012-2014 Gigamon. All rights reserved.

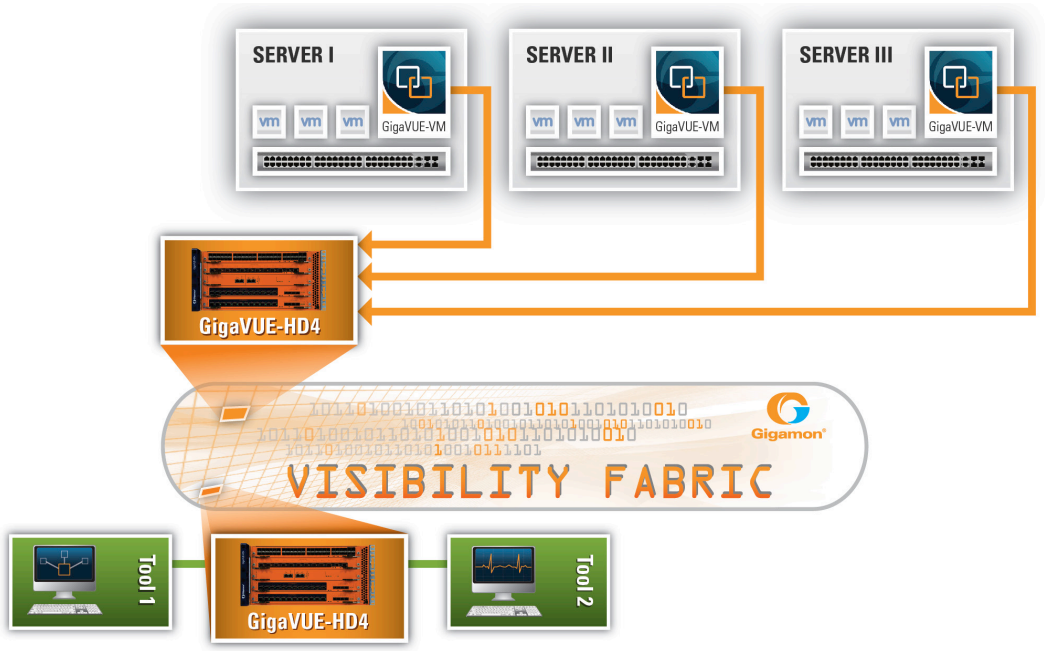

Figure 2: GigaVUE-VM provides traffic visibility into virtualized environments

forwarded, and delivered to the appropriate monitoring, analysis, or security devices (see Figure 2). Options provided by virtualization vendors such as placing the guest's network adapter in promiscuous mode have their own set of concerns and limitations.

Promiscuous mode allows a monitoring tool connected to the virtual switch to receive all frames passed on the virtual switch including traffic destined to other guests or host operating systems. This mode can raise legitimate security concerns (even more so with multi-tenant environments) since any adapter in promiscuous mode has access to the packets regardless of whether the traffic is destined to that particular network adapter. The alternative option of using the port-mirroring capability available in VSphere 5.0 is identical to this functionality in physical switches. This would result in all the virtual machine-to-virtual machine traffic being sent out to the physical network, thereby potentially overloading the physical NIC as well as the network with traffic that might not even be interesting from the end-user's monitoring perspective.

A native VMware vSphere 5 Virtual Machine, the GigaVUE-VM fabric node is installed without the need for invasive agents, or changes to the hypervisor, allowing system managers to achieve the same packet-level traffic visibility between virtualized applications as is normally available between discrete physical applications and servers. End-users can selectively filter traffic flows between virtual machines on the same ESXi host, based on specific application criteria and forward these to GigaVUE® fabric nodes resident in the physical network to be additionally aggregated, replicated and made available to network performance, application performance and security monitoring systems.

Enabling Visibility in Cisco Nexus 1000V Deployments with GigaVUE-VM Virtual Visibility Fabric Nodes

Cisco Nexus 1000V Series represents the first example of third-party distributed virtual switches that are fully integrated with VMware virtual infrastructure, including VMware vCenter for the virtualization administrator. When deployed, the Cisco Nexus 1000V Series not only maintains the virtualization administrator's regular workflow; it also offloads the vSwitch and port group configuration to the network administrator, reducing network configuration mistakes and helping ensure that consistent network policy is enforced throughout the data center. In the Cisco Nexus 1000V Series, traffic between virtual machines on the same host is switched locally without ever hitting the physical switch, thus potentially creating blind spots for monitoring and management tools. With GigaVUE-VM fabric

© 2012-2014 Gigamon. All rights reserved.

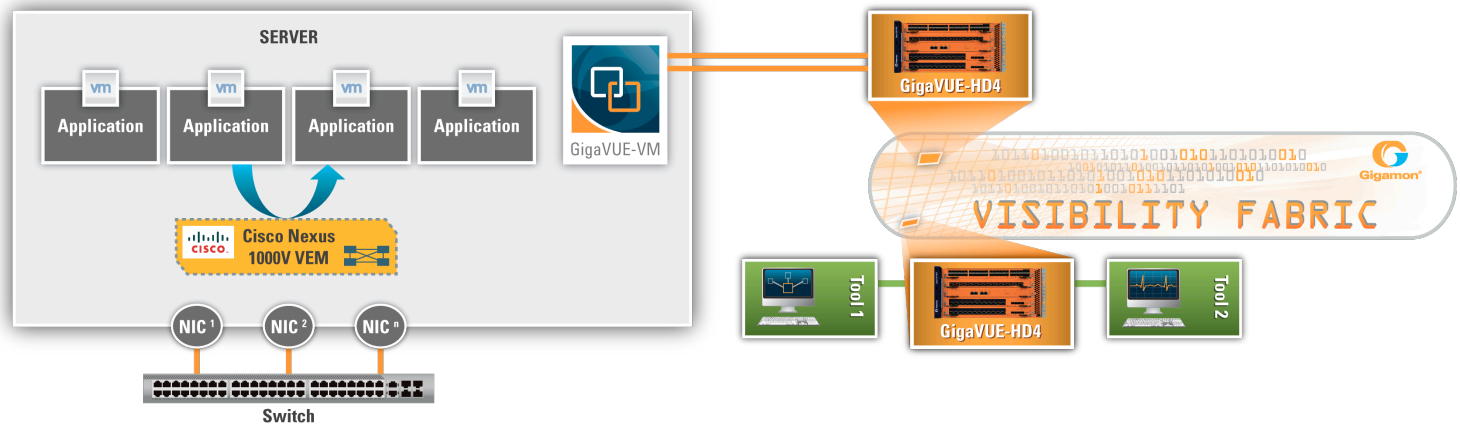

Figure 3: Visibility in Cisco Nexus 1000V deployments with GigaVUE-VM Visibility Fabric nodes

nodes deployed, virtual traffic across these environments can be intelligently detected, selected, filtered, and forwarded locally or remotely, without any changes to the operational procedure or adding any further complexity to the underlying infrastructure. Currently deployed monitoring and management tools can thus be utilized to analyze traffic flowing across the virtual infrastructure using best-of-breed virtual switching including vSphere Distributed Switch (VDS) and Cisco Nexus 1000V (see Figure 3).

Maintaining Awareness with vMotion

The benefits of virtualization are unassailable: increased agility, scale, and cost savings to name a few. However, so too are the monitoring challenges posed by these environments—including complexity, lack of visibility and control, and potential inefficiency. To monitor these environments effectively, administrators have to leverage visibility solutions that have an integrated, synchronized awareness of these automation technologies. Tightly integrated with the VMware vCenter infrastructure, while leveraging VMware open APIs, the GigaVUE-VM fabric node can track agility across VMware High Availability (HA) and Distributed Resource Scheduler (DRS) cluster environments. As part of this support, the visibility policies are tied to the monitored VMs and migrate with the VMs as they move across physical hosts in the virtual clusters. Closed loop feedback through standards-based APIs to vMotion events and an automation framework that enables sync-up of visibility policies facilitates seamless, real-time adjustment of monitoring and security posture in an agile virtual infrastructure.

Enabling Visibility in Cisco Deployments with VN-Link

In addition to the distributed virtual switch, Cisco also provides the option of forwarding virtual traffic streams to an external network switch, after tagging the original packets with a unique VN-Tag ID. The most important components of the tag are the source and destination virtual interfaces (VIF) IDs which identify multiple individual virtual interfaces on a single physical port. However, now that the packets have been modified with tags, the primary challenge is to access this encapsulated traffic without making hardware or software changes or wasting processing cycles.

Encapsulation awareness enabled by Adaptive Packet Filtering, allows operators to filter and forward incoming traffic streams based on VN-Tag source or destination VIF_IDs and/or the inner (encapsulated) packet contents. For monitoring and analytic tools that do not understand VN-Tag headers, Adaptive Packet Filtering can also be used in combination with header stripping to remove VN-Tag headers before forwarding the packets. The advanced processing available with Gigamon GigaSMART® technology provides the flexibility of conditionally filtering and forwarding traffic and alternatively stripping out VN-Tag headers based on specific contents found across the packet. With the preprocessing provided by GigaSMART technology, the network and security monitoring devices can now inspect the traffic source from the virtual network without expending precious resources (see Figure 7).

© 2012-2014 Gigamon. All rights reserved.

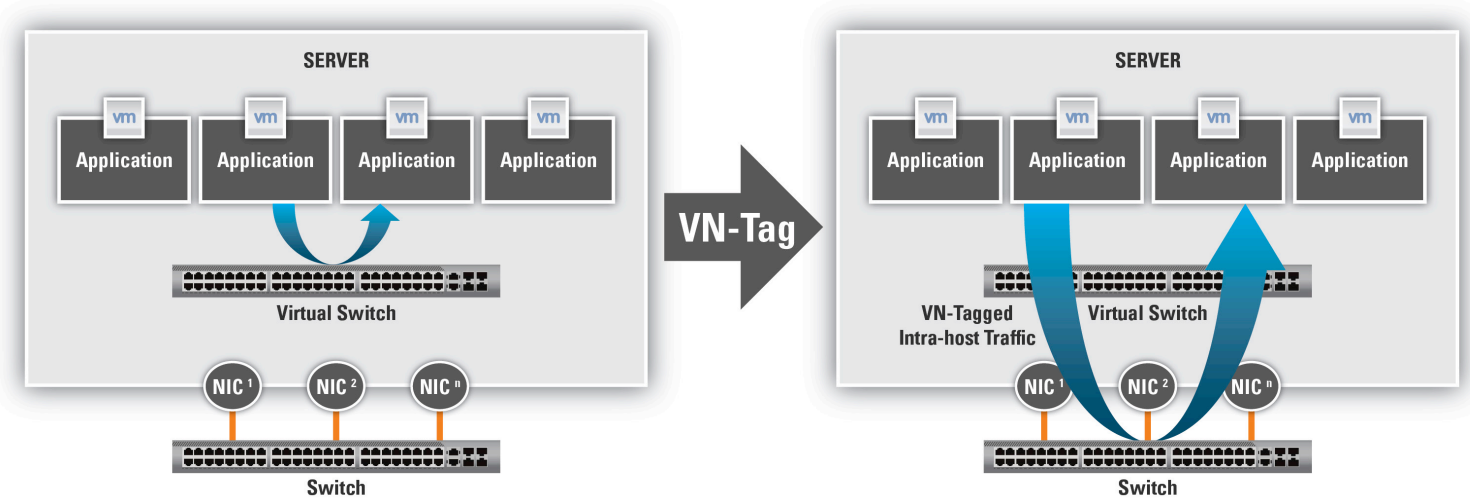

Figure 4: VN-Tagged intra-host traffic where switching is managed by the physical switch

VN-Tag

The VN-Tag standard was proposed as an alternative solution to provide access layer extension without extending management and STP domains, end-to-end policy enforcement, and therefore better network visibility in virtualized environments—specifically VM to VM traffic on the same host.

Using VN-Tag, an additional header is added into the Ethernet frame which will be used by a VN-Tag-aware external (switching) device to uniquely identify VIF and forward data after enforcing the necessary security policies (see Figure 5). VN-Tags thus provide virtual networking awareness, allowing for individual configuration of each virtual interface as if it were a physical port.

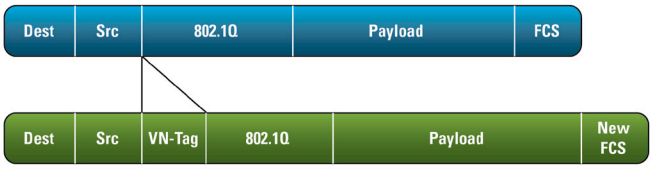

Figure 5: VN-Tag added as an additional header

This approach completely removes any switching function from the hypervisor and locates it in an external hardware network switch physically independent of the server—packet switching is completely decoupled from the hypervisor (see Figure 4). The versatility of VN-Tags allows this technology to be applied in existing physical network infrastructures. For example, VN-Tags are used in bridge extensions, where a unique identifying tag is inserted into each frame exchanged between the Cisco fabric extenders and the Nexus parent switch to uniquely identify the originating port.

One of the disadvantages of VN-Tags is that they utilize additions to the Ethernet frame. Standard monitoring tools which do not understand VN-Tag would be completely blind to this traffic and therefore rendered useless. VN-Tags also increase traffic on the host server's physical network links.

© 2012-2014 Gigamon. All rights reserved.

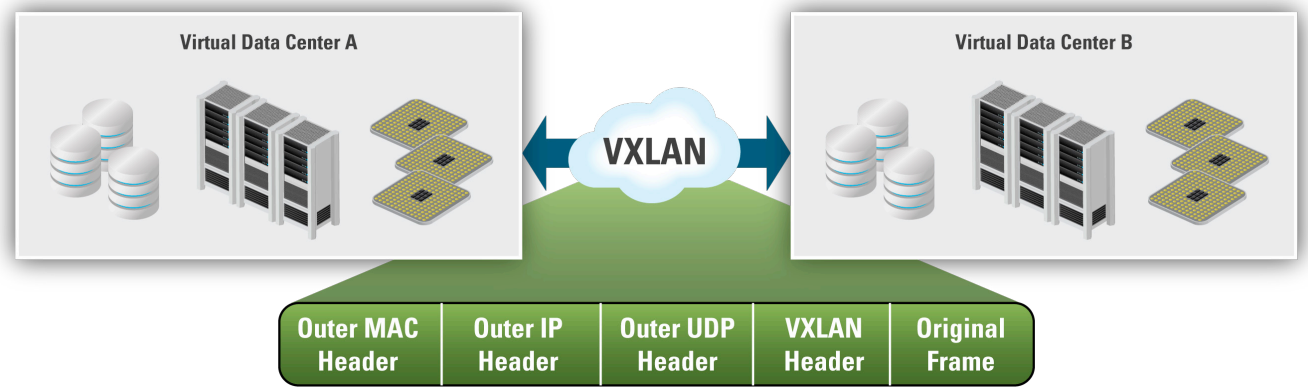

Figure 6: VXLAN allows Layer 2 virtual networks to span across physical boundaries

Overlay Networks for Virtualized Data Centers

Virtual eXtensible Local Area Network (VXLAN)

Current limitations of physical networks tie an increasingly pooled dynamic virtual world back to rigid, complex network architectures creating artificial barriers to realizing the full agility organizations expect from private clouds. While a virtual machine can be provisioned in a matter of minutes, "surrounding" that VM with all the necessary network and security services still takes days. Highly available virtualization technologies such as VMware Fault Tolerance work best with "flat" Layer 2 networks, but creating and managing this architecture can be operationally difficult, especially at scale. True software-defined networking (SDN), where a data center or cloud operator would not care about the hardware and device OS, really starts with "plumbing" projects like OpenFlow and VXLAN.

VXLAN helps solve the data center networking challenge. It provides the capability to create isolated, multi-tenant broadcast domains across data center fabrics and enables customers to create elastic, logical networks that span physical network boundaries. Thus virtualizing the network and creating networks that meet the agility, performance, and scale requirements of virtualized applications and data.

VXLAN works by creating Layer 2 logical networks that are encapsulated in standard Layer 3 IP packets thus allowing the extension of Layer 2 virtual networks across physical boundaries (see Figure 6). A "Segment ID" in every frame differentiates the VXLAN logical networks from each other without any need for VLAN-Tags.

This method not only allows very large numbers of isolated Layer 2 VXLAN networks to coexist on a common Layer 3 infrastructure, it also allows virtual machines to reside on the same Layer 2 virtual network but be on two different Layer 3 networks. VXLAN runs over standard switching hardware, with no need for software upgrades, enabling multi-tenancy while extending virtual data centers across different physical locations of the cloud/data center networks.

However, there is potentially a huge disadvantage of this approach: the traffic is hidden in a tunnel, making network monitoring difficult. Individual applications flowing within the tunnel therefore cannot be monitored. Even with hardware or software modifications to make monitoring tools compatible with VN-Tag labels and VXLAN encapsulated tunnels, these tools would still be spending precious resource cycles stripping out encapsulations while they can be more efficiently used for what they were originally designed to do, i.e. analyze and monitor traffic.

Pervasive Visibility into the Virtualized Data Center and Cloud

Today, as organizations drive toward the adoption of further virtualization and cloud solutions, there is no longer any need for them to let performance or security concerns hinder that advance.

© 2012-2014 Gigamon. All rights reserved.

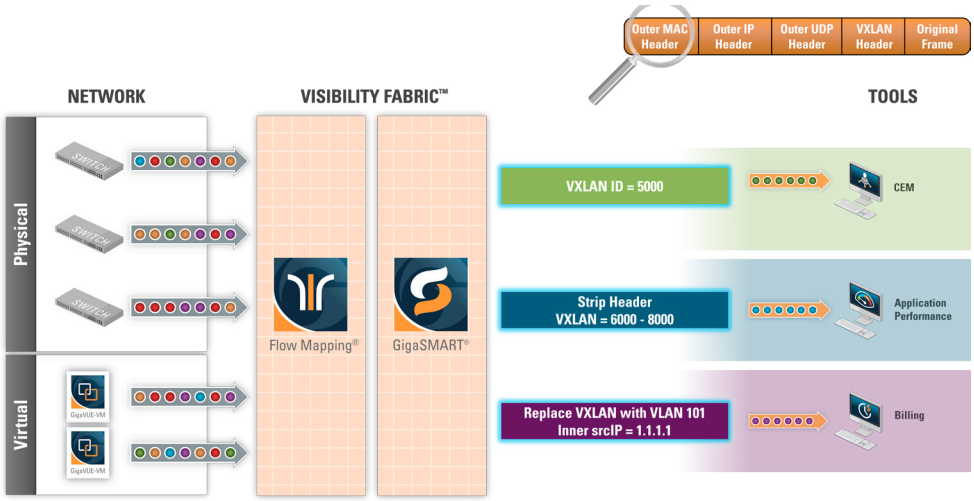

Figure 7: VXLAN awareness

Visibility in Overlay Networks

Enabling Visibility of VXLAN Encapsulated Traffic

With a 24 bit segment ID to uniquely identify broadcast domains, VXLAN enables multi-tenant environments at cloud scale and extends the Layer 2 network across physical boundaries by encapsulating the original frames in a MAC-in-UDP encapsulation. Monitoring performance of VXLAN SDNs and virtual tunnel endpoints is the key to enabling network operations teams to control and comprehend the "virtual" domains floated on top of the common networking and virtualization infrastructure. The VXLAN encapsulated traffic can be sent to the Visibility Fabric architecture, which can utilize the processing capabilities provided by GigaSMART to filter on the segment-ID to forward and/or de-capsulate specific traffic flows before forwarding to the monitoring tools that need access to this information (see Figure 7). Thus the network and security analyzers that need critical visibility into the UDP-encapsulated traffic can now monitor the security and performance of the VXLAN virtual overlay network without any hardware or software modifications.

A Unified Visibility Fabric Architecture

The Unified Visibility Fabric architecture is an innovative solution that delivers pervasive and dynamic visibility of network traffic traversing communication networks. A unified monitoring fabric with centralized access is required to assure independence and segregation of the monitored data delivered to multiple organizations and tools. It enables the unification of data visibility across different network architectures including physical and virtual networks, resulting in a safe deployment within a multi-tenant setting (see Figure 8).

Services Tier

Aggregation, Filtering, Replication, and Intelligent Packet Modification

The Services layer consists of distributed network appliances that provide an advanced level of filtering intelligence, including traffic forwarding, manipulation, and modification. As part of the services layer, GigaVUE-VM further extends visibility across virtualized server infrastructure by leveraging standards-based APIs to filter and forward traffic of interest into higher-functioning Gigamon service nodes. Pervasive visibility across the virtual and physical network provides a singular view across the entire infrastructure, thus helping to ensure productivity, performance, and meeting service level agreements. In addition to providing access to critical information, the services layer can be used to modify packets in flight in order to hide confidential information, add timing information, remove duplicate data, and strip out extraneous headers thereby enhancing the efficiency of monitoring and security tools. Today more and more organizations are driving toward the adoption of virtualization, cloud, and programmable networks. New technologies such as VXLAN, VN-Tag, etc., are rendering tools used to monitor, analyze, and secure the IT infrastructure essentially blind to the traffic flowing in and out of the data center. The enhanced header-stripping functionality

© 2012-2014 Gigamon. All rights reserved.

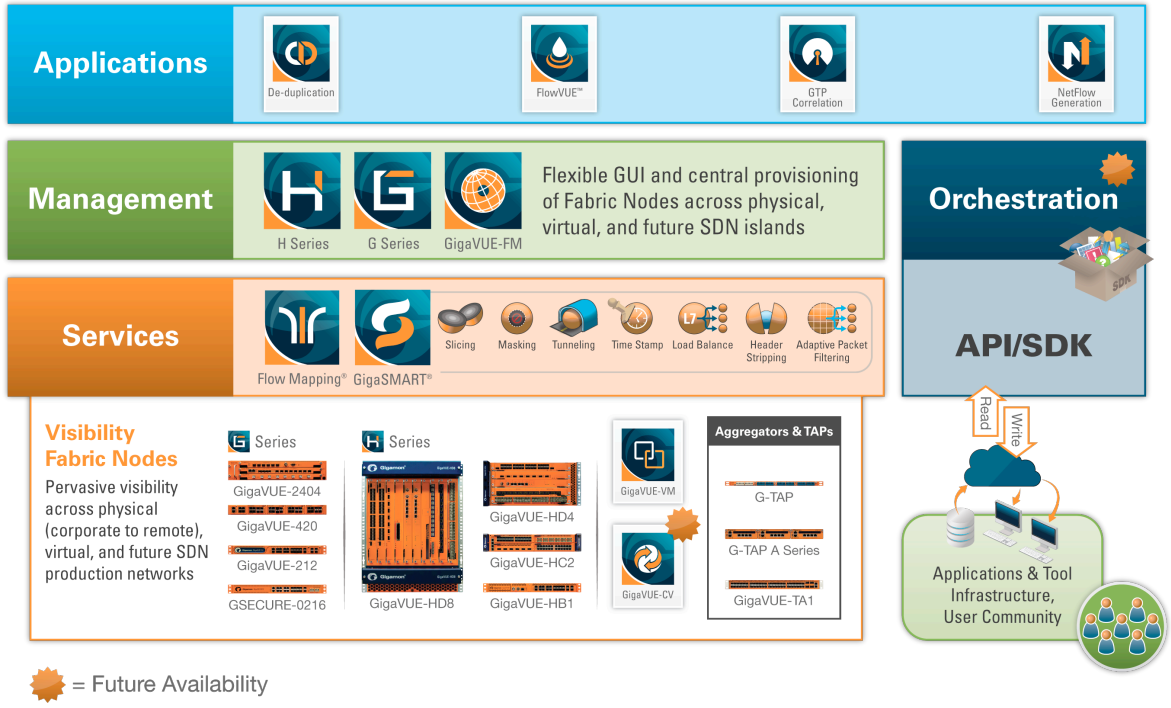

Figure 8: Unified Visibility Fabric Architecture

offered by the services layer enables organizations to overcome concerns about performance and security as they look to adopt virtualization and cloud solutions, and take advantage of the tremendous value propositions they offer.

Management Tier

Unified End-to-End Provisioning

A unified configuration interface made available by GigaVUE-FM (Fabric Manager) will provide end-to-end policy management of monitored traffic from physical, virtual, and SDN networks, while enabling a full-lifecycle eco-system interface with the tools. At the heart of the Visibility Fabric architecture is the patented Flow Mapping® technology that identifies and directs incoming traffic to single or multiple tools based on user-defined rules implemented from a centralized management console. New enhancements to Flow Mapping help address multi-tenant access and segregation of monitored traffic and policies by providing advanced role-based management and simplified GUI workflows. This enables a more dynamic management of monitored traffic, which eliminates silos of operation—reducing both CAPEX and OPEX.

*Orchestration Tier**

Programmability, Automation, and Tool Integration Through an Open Framework

The orchestration layer offers the ability to provide 'just-in-time' responsiveness to real-time events that occur within the network. The ability to tune monitoring services dynamically without manual intervention helps minimize reactive management of the infrastructure to a more proactive approach. It enables an independent software developer community to provide applications that enable visibility as a service. The standards-based set of APIs will be designed to have full interoperability with a broad range of tools, allowing IT organizations maximum architectural flexibility in the design and augmentation of their infrastructure.

**Future availability*

© 2012-2014 Gigamon. All rights reserved.

Application Tier

Dynamic Power to Customize Traffic Selection and Enable Tool Optimization

The Unified Visibility Fabric architecture makes it simple to create value-added visibility applications, and develop new application-specific capabilities, so users can efficiently and securely address their business needs. Monitoring tools can now perform more efficiently by eliminating duplicate content with currently available applications like de-duplication, and in the future turn big data in to manageable data using intelligent flow-based sampling enabled by FlowVUE. It establishes a foundation to enable enterprises and service providers alike to develop specialized and complementary solutions, just in time for the dynamic data onslaught brought about by cloud computing, mobility, and social networking.

Conclusion

Tomorrow's data center is here now. There's no question that virtualization owns the current phase of data center transformation. The efforts to centralize, optimize and simplify the delivery of IT services in an on-demand method are driven by the focus of businesses on reducing their cost of IT as a percentage of business revenue. But virtualized servers and virtual networking capabilities are creating bigger and bigger pockets of IT that are hidden from the tools that are relied on to measure uptime, secure data and assets, and analyze performance of the network and applications. Versions of those tools can be virtualized, but because of their heavy dependency on server resources, IT solution architects won't mix their tools with their applications on the same hypervisor. Thus the tool dashboards remain blank when physical to virtual migration occurs. Using a virtual port mirror introduces great risk of over-consuming network bandwidth. End to end tunneling protocols also hide relevant packet information from the tools, or tax the tools beyond their limits with the task of header stripping and decapsulation. The solutions offered by Gigamon shed the necessary light on the virtualized and cloud environment, and extend the reach of monitoring, analytic and security tools into every corner of the data center and every silo of IT that has been hidden by the many facets of cloud computing and virtualization technologies that are in production today.

About Gigamon

Gigamon provides an intelligent Visibility Fabric™ architecture 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 Visibility Fabric architecture visit: ***www.gigamon.com***

Summary of the key benefits of the Unified Visibility Fabric architecture:

- Bridge islands of physical, virtual, and SDN* infrastructure with the tools required for end-to-end visibility across campus, cloud, and carrier
- Normalize and optimize traffic to the tools across islands of users, virtual machines, devices, and applications to enable tool optimization
- Enable parallel monitoring policies to serve multiple departments simultaneously with a flexible policy engine
- "Just-In-Time" responsiveness to real-time events that occur within the network through automation and orchestration

***Future availability**

© 2012-2014 Gigamon. All rights reserved. Gigamon and the Gigamon logo are trademarks of Gigamon in the United States and/or other countries. Gigamon trademarks can be found at www.gigamon.com/legal-trademarks. All other trademarks are the trademarks of their respective owners. Gigamon reserves the right to change, modify, transfer, or otherwise revise this publication without notice.