Gigamon®

Gain Control and Observability into Your Amazon Web Services Hosted Applications

Introduction

If, as part of digital transformation, your organization is migrating existing applications and initiating new born-in-the-cloud workloads to public cloud vendors' Infrastructure as a Service (IaaS) offerings, you're not alone.

According to RightScale, 69 percent of surveyed organizations are using public or hybrid clouds as they make their digital transformation, with a third of their workloads now located in these clouds.[1]

The advantages of these moves are well known: economies of scale and access to a shared pool of resources that can be provisioned and deployed quickly and easily are two examples. Because compute resources are owned and hosted by the IaaS provider and offered to customers on demand, you can effectively outsource IT operations and minimize CapEx. But a migration to the cloud brings network and application visibility challenges along with it — and you're still responsible for your own security in this new hybrid environment.

Challenges: Obtaining, Processing, and Distributing Cloud-Based Traffic

As part of this transition, CloudOps teams scale their deployments with an ever-growing number of applications and services; the servers and virtual machines they run on are dispersed and compute nodes continually relocated throughout multiple virtual private clouds. IT routinely adds various security and monitoring tools and leverages container technologies. The result is a labyrinth of traffic sources and destinations.

There are a variety of methods for accessing network data in on-premises infrastructure, including SPAN sessions and physical and virtual TAPs. But because they need to ensure privacy in a multi-tenant environment, IaaS providers do not allow customers to deploy their own virtual TAP functions outside of their assigned VM or containers. In that scenario, you're forced to use the vendor's virtual network TAP services, if they are both offered and available with the specific type compute instances you're using, to acquire and send raw packets directly to your cloud or on-premises-located security and network monitoring tools.

Administrators need full packet and application visibility in the cloud for both North-South and East-West traffic. This requirement extends to multiple public cloud and hybrid deployments with the ability to simultaneously support any such scenario.

While on-prem topologies benefit from network packet brokers (NPB) that acquire, aggregate, process, and distribute traffic to the proper security and networking tools, cloud vendors do not offer NPB functionality. Nor can cloud solutions identify and filter content based on applications and generate advanced L4–7 metadata as their on-premises cousins can. This leads to complex network designs, excessive bandwidth usage, overwhelmed tools that lose effectiveness, and needless scaling. As a result, your IT staff will be limited in its ability to analyze network traffic and customer experience, and will have difficulty evaluating infrastructure health.

In addition to granular application visibility, organizations must deploy a solution that not only solves the aforementioned challenges but overcomes many varied demands including:

- Complete security for their apps and data, operating systems, firewall configurations, etc.
- Maximum security and monitoring tool efficiency and accuracy
- Infrastructure automation with deep integration into orchestration tools
- Generation of NetFlow and advanced metadata attributes
- Network traffic consolidation and elimination of duplicated data flows
- Backhauled traffic when security and/or monitoring tools are on-prem

Challenges: Ensure Security of Enterprise Cloud-Based Resources

When enterprises first started leveraging IaaS, they began by migrating Tier 2, test/dev, or other infrequently used applications that consumed expensive resources when run on-premises. But now, in the next stage of digital transformation, organizations want to move Tier 1 or mission-critical applications to IaaS. These applications deal with sensitive data and information that needs to be safeguarded and protected from unauthorized access and potential cybersecurity attacks.

IaaS providers emphasize mutual responsibility in the cloud: The provider is responsible for the security of the cloud infrastructure itself, but the customer is responsible for assets within the cloud. See Figure 1.

The assets that the customer's IT, cloud, and security architects must protect include data and applications; these teams are also charged with organizational and regulatory compliance. They must ensure that applications and workloads are being deployed securely by everyone within the organization. Enterprises that migrate to the cloud typically rely on techniques such as workload security, perimeter security, prevention-only solutions including access lists or security groups, and identity and access management to mitigate security risks.

Today's evolving threat landscape has rendered prevention-only security techniques insufficient. Over 80 percent of network traffic is now East-West — that is, between VMs or containers — so malware can more easily spread undetected. Any solution needs to be complemented with additional detection and response techniques to uncover early signs of security anomalies and deviations from expected behavior. For this to happen, organizations need to implement a multi-tiered security model and have accurate visibility into virtual machine network traffic. Without such visibility, moving mission-critical applications to the cloud jeopardizes their safekeeping.

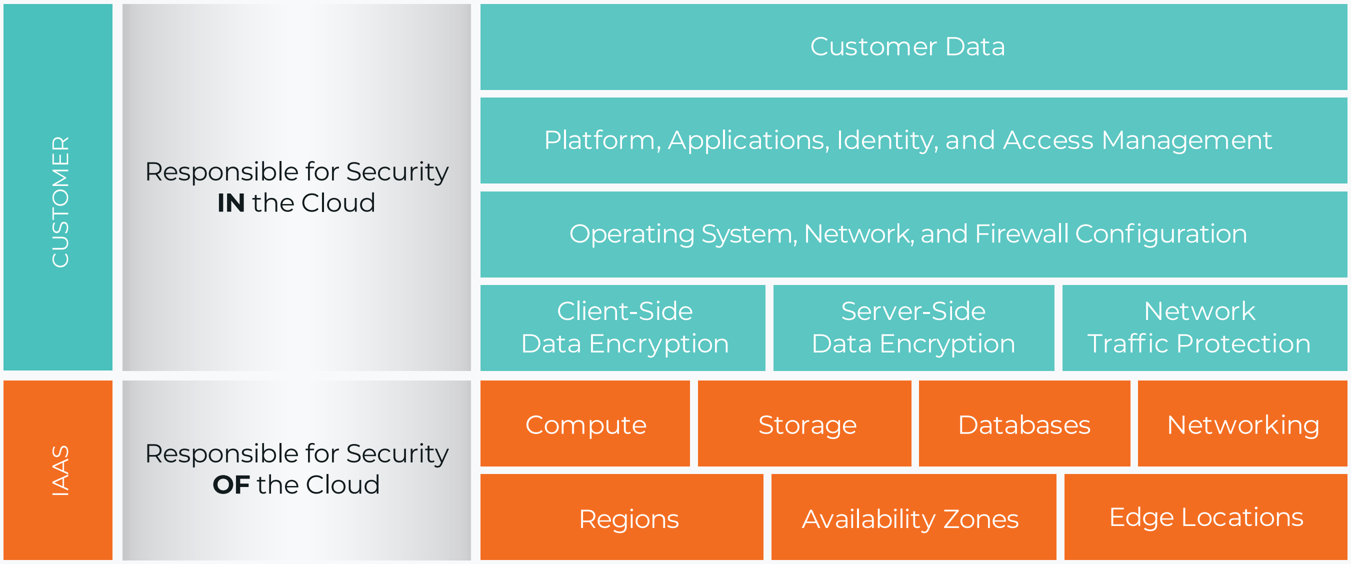
CUSTOMER
Responsible for Security IN the Cloud
Customer Data
Platform, Applications, Identity, and Access Management
Operating System, Network, and Firewall Configuration
Client-Side Data Encryption
Server-Side Data Encryption
Network Traffic Protection
IAAS
Responsible for Security OF the Cloud
Compute
Storage
Databases
Networking
Regions
Availability Zones
Edge Locations

Figure 1. Shared responsibility model for public cloud.

Legacy Approach to Visibility and Security

Historically, cloud customers were simply unable to obtain copies of traffic and direct it as desired. Recently, several cloud vendors have introduced virtual network TAPs, such as AWS's traffic mirroring service. These tools provide raw, unprocessed packets with minimal filtering, which can be distributed only to the end security and monitoring tools. However, these services do not allow packet replication and thus can only send a copy of one packet to one tool, not to multiple tools. In mirroring traffic from multiple locations, packets can be duplicated with no ability to remove this superfluous content that needlessly floods the network. These vendor offerings also don't support all compute node types and don't enable secure tunneling such as with IPsec.

Other traffic acquisition methods include approaches such as AWS CloudWatch, which triggers packet capture on specific events; but this is a reactive technique, primarily intended for troubleshooting. CloudWatch also provides VPC Flow Logs. Users of such tools must accept a lack of advanced NetFlow/IPFIX metrics and suffer from reduced network efficiency and traffic insights. When faced with these limitations, some users may decide that the cloud isn't ready for mission-critical applications and will choose to run those on-premises; this leads to expensive backhauling of all cloud traffic to on-prem tools.

Another approach is to deploy custom agents on every compute node for each and every tool (see Figure 2). But workloads that deploy numerous agents may suffer from agent overload as a result. As each agent copies traffic, more compute cycles and network bandwidth are used. Bandwidth ceilings limit the number of compute instances allowed in a given virtual cloud, which means you'll incur additional expense as you need ever more virtual cloud-based resources. This technique also results in a significant performance impact and requires manual intervention when new tools are added.

Ensuring security can be equally daunting. Cloud-native security services commonly deployed in the public cloud include identity and access management (IAM), security groups, logs, and web application firewalls (WAF). But these have limitations:

- IAM: Once an attacker has successfully hacked credentials, they won't need to undertake noticeable activity that gets alerted by cloud dashboards; they could sit there silently and do just enough to not trigger any alerts. The time to detection in this case is many weeks or months.
- Security groups: Despite opening access to only necessary ports, security group configurations have no application context and no visibility into higher layers (beyond L4). Attacks could happen on those ports in the application layer and could result in malware being deployed or data exfiltrated.
- Logs: Logs only convey high-level metrics about conversations and application access points: You'll know who communicated with whom, but won't have a record of what the communication was about. No packets or payload are included. In case of silent attacks where attackers use the infrastructure and try to operate within limits of threshold violations, logs are of no help.
- WAF: Cloud-native WAFs are very limited in their functionality when compared to industry-leading WAFs, and only protect apps from the OWASP top ten attacks.

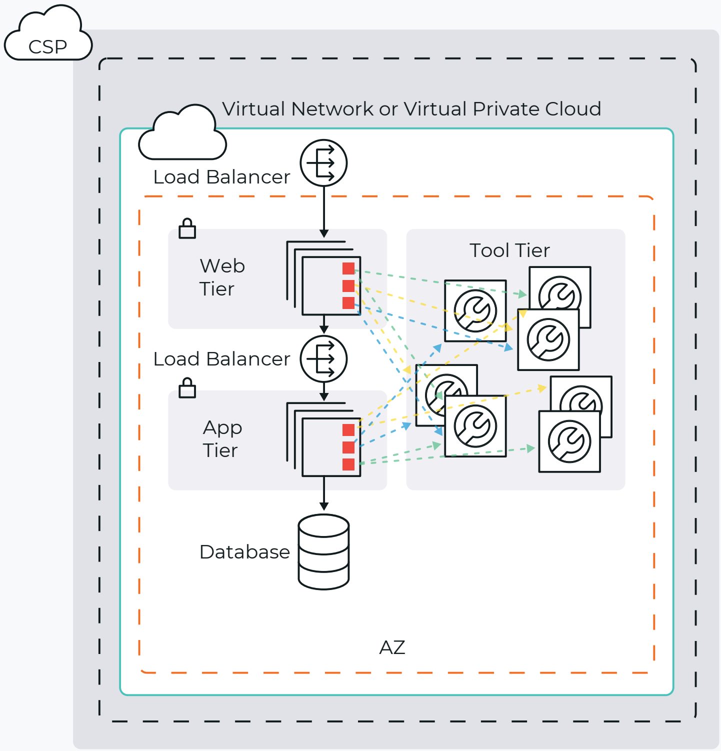

Figure 2. Legacy traffic acquisition in public clouds by installing an agent on every VM for every security and monitoring tool.

Drawbacks include:

- Inability to access all traffic
- Discrete vendor monitoring agents per instance
- Excessive loads placed on compute instances
- Excessive duplication of identical network traffic flows
- Cannot process traffic prior to sending to tools
- No ability to replicate and send packet to multiple tools at once

Network and Application Traffic Visibility in the Cloud: The Missing Link

Organizations with on-prem operations have successfully deployed next-generation network packet brokers (NGNPB), in both physical and virtual form factors, for many years. These platforms help NetOps and SecOps teams obtain the necessary visibility into traffic throughout their data center. IT can only ensure superior security and network performance while minimizing costs if they have full access to all data-in-motion, including from VMs and containers, along with an ability to properly identify applications, filter and distribute traffic to the right tools and shield tools from needless processing. Public cloud customers can achieve the same results by leveraging cloud-native versions of the on-prem NGNPBs they rely on.

Vendor-certified and Marketplace-listed solutions enable CloudOps and DevOps teams to have the same NGNPB capabilities for the cloud that they depend on in their own environments. With these tools, IT won't have the baggage of running this infrastructure, but can enjoy complete North-South and East-West traffic visibility — including at the application layer. Cloud-hosted NGNPBs can be automatically scaled to any level required and provide the packet processing critical to removing superfluous content and easing the burden on security and monitoring solutions. Automation of the infrastructure is ensured with deep API integration into the cloud vendor's orchestration tools, minimizing manual efforts and errors. NetFlow, IPFIX and advanced metadata can now be generated and used to feed SIEM solutions and other tools. With a critical mass of cloud-based security tools, traffic need not be expensively backhauled to on-prem infrastructure, but instead will continue to reside in the cloud.

Organizations often use multi-cloud or regional deployments, so NGNPB vendors' orchestration tools must support simultaneous multiple public or hybrid clouds. Centralized management, monitoring and control can be simplified through a single-pane-of-glass GUI. This is important as cloud vendors may have dozens of regions and availability zones spread over dozens of countries and geographic regions. Typically, enterprises distribute their cloud infrastructure across these multiple regions and accounts. Having a security policy for such a distributed infrastructure — let alone enforcing that policy — is challenging. In such a scenario, an inconsistency in security configurations anywhere could lead to a weak spot that can be attacked and compromised.

For applications and workloads in cloud IaaS, security tools need to be able to access the right data. But as organizations deploy multiple security tools across their infrastructure to ensure an effective security and performance monitoring strategy, the NGNPB needs to support the tools no matter where they reside. Scenarios include:

- Tools are in on-premises infrastructure and traffic is backhauled from the cloud to these tools
- Tools are in a cloud IaaS tool tier and traffic needs to be moved across compute instances and/or tiers
- Different users in an enterprise may have multiple virtual private cloud (VPC) instances and a common set of tools may be required to inspect traffic across these VPCs

An Observability Platform for the Public Cloud

GigaVUE® Cloud Suite is a a part of the Gigamon Deep Observability Pipeline and provides cloud-native network and application traffic observability solutions dedicated to specific cloud environments, including Amazon AWS, Microsoft Azure, and Google Cloud Platform. Gigamon offers the industry's only true deep observability pipeline (sometimes referred to as next generation network packet broker) on those vendors' marketplaces. The software suite elements reside fully in the cloud; they acquire traffic from every compute site, either through the aforementioned native traffic mirroring services, from AWS external load balancers, or via Gigamon G-vTAP Modules (agent-like instances provisioned on each VM), UCT (container pod instances provisioned on each node), or infrastructure mirroring services. UCT supports any Kubernetes environment including Amazon Elastic Kubernetes Service (EKS).

Network packets are copied and directed to GigaVUE V Series virtual appliances, where they are aggregated and processed: Duplicate packets are eliminated, irrelevant application content deleted, sensitive payload material masked, and headers transformed. Advanced L2–7 NetFlow and application metadata attributes are generated. Optimized traffic is replicated and, along with metadata, is then load balanced and steered to the proper tools. GigaVUE-FM fabric manager is integrated into the cloud tool suite to provide full automation.

With GigaVUE Cloud Suite, you can now extend your security posture to the public cloud, ensuring compliance and detecting threats to crucial applications more quickly. This suite makes it possible to:

- **Improve tool capacity.** Virtual security and monitoring tasks are offloaded from tools and irrelevant application content dropped to improve their effectiveness, reduce scaling, and minimize costs.
- **Choose the proper traffic acquisition method.** Flexibly select between your cloud vendor's traffic mirroring offerings for more simplified operations or Gigamon lightweight agents for added security and filtering.
- **Fully automate the infrastructure.** Automatically identify new and relocated workloads, instantiate and scale visibility nodes and configure new traffic policies as needed.
- **Simplify operations.** Centralize orchestration and management with a single-pane-of-glass visualization portal across any hybrid network.
- **Reduce risk** by leveraging a common Deep Observability Pipeline across your entire IT environment.
- **Empower your SIEM** and monitoring tools with advanced metadata to identify and resolve security and performance issues.

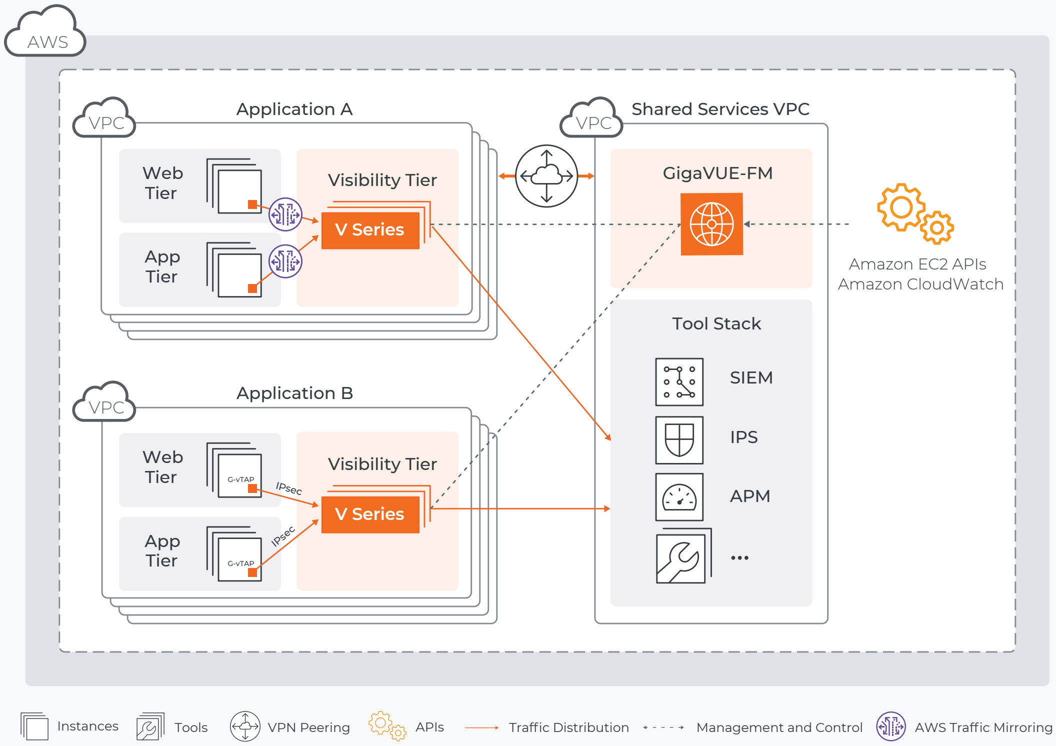

Figure 3. GigaVUE Cloud Suite for AWS supports multiple VPCs including transit gateways and has tight integration with AWS cloud management tools to enable automation. Either AWS's agentless native VPC traffic mirroring or Gigamon GigaVUE G-vTAPs can collect all traffic streams.

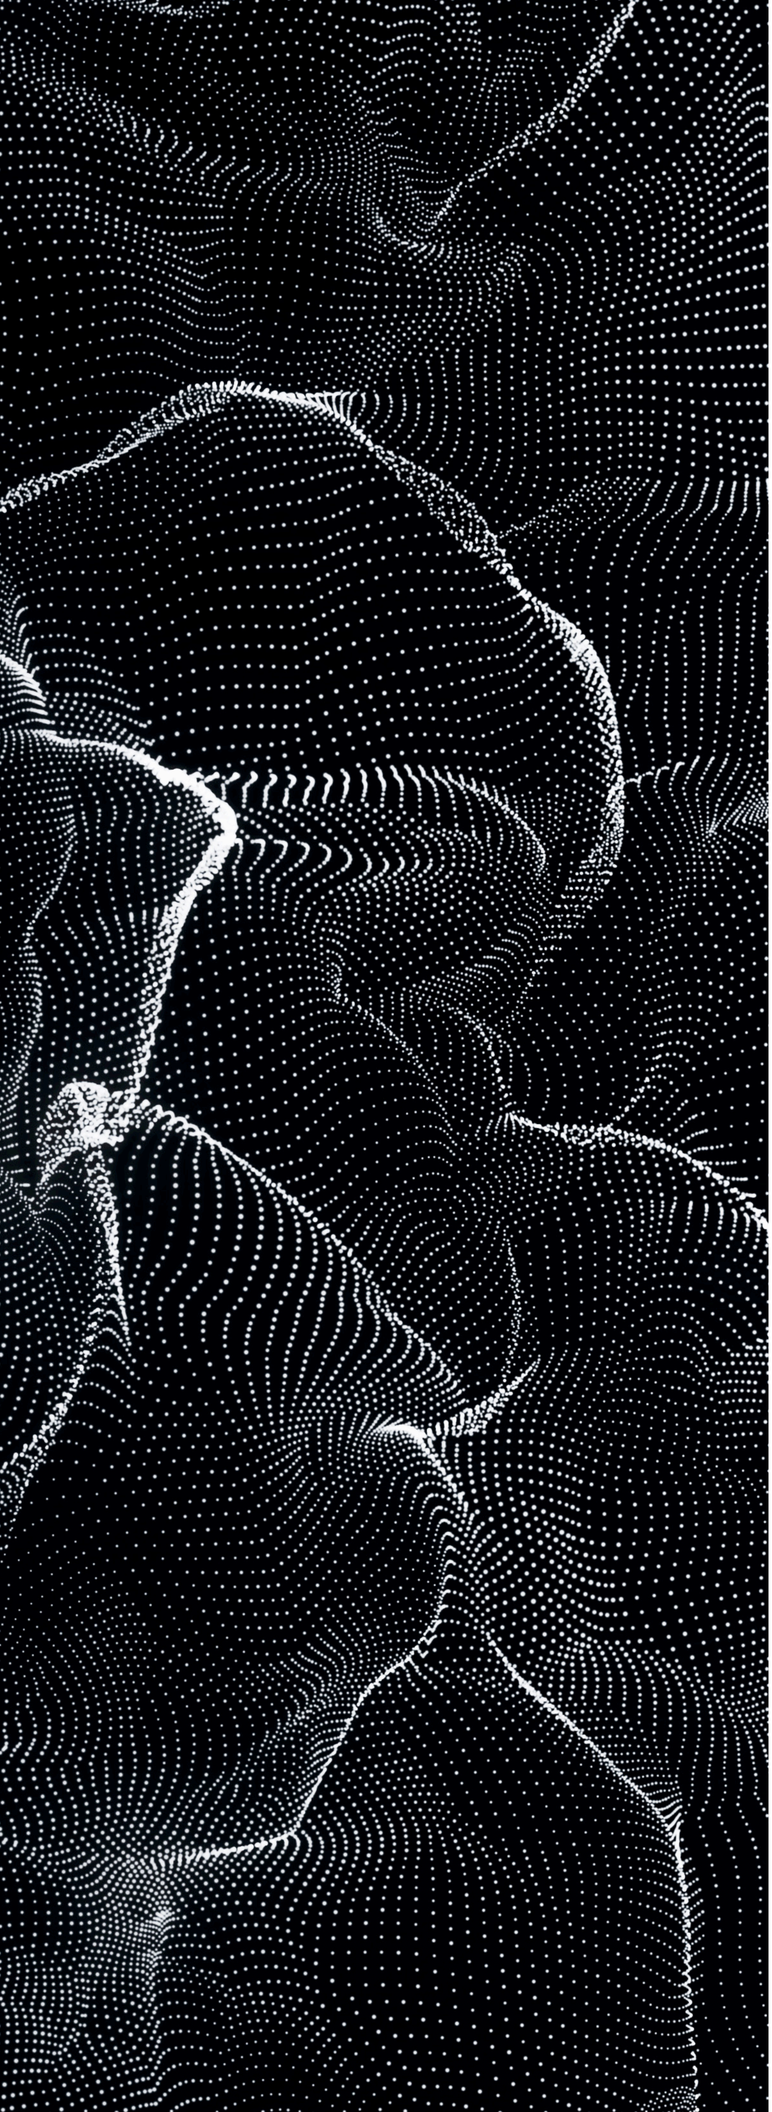

GigaVUE Cloud Suite Illuminates Public Clouds and Enhance Security

This suite comprises multiple elements that enable traffic acquisition, aggregation, intelligence, and distribution, along with centralized, single-pane-of-glass orchestration and management. The solution consists of these components:

G-vTAP Module

This lightweight agent is deployed in various compute instances to mirror production traffic and send to GigaVUE V Series nodes for further processing and distribution to monitoring and observability tools.

Key features and benefits:

- Minimize VM overload. Only one module is necessary per instance, lowering the impact CPU and throughput.
- Automatic Module scaling. As new workloads are spun up, GigaVUE-FM interoperates with the compute instance APIs and the cloud vendor's management tools to instantiate new modules.

Universal Container TAP (UCT)

This light-weight container pod is deployed in container workload nodes to mirror production traffic and send to GigaVUE V Series intelligent visibility nodes for further processing and distribution to monitoring and observability tools.

Key features and benefits:

- Minimize node overload. Only one UCT is necessary per worker node and traffic does not pass through UCT, minimizing the impact on CPU and throughput.
- Automatic UCT scaling. As new worker nodes are spun up, GigaVUE-FM interoperates with the compute instance APIs and the cloud vendor's management tools to instantiate new UCT instances.

GigaVUE V Series Nodes

These are visibility nodes that aggregate, select, optimize, and distribute traffic to the tool tier, which may be located in the public cloud IaaS or in an on-prem data center. These nodes, usually deployed as a cluster, reside within the public cloud VPCs.

Key features and benefits:

- Traffic acquisition: Acquire traffic from multiple VM and container pod instances, using G-vTAP Module and UCT or AWS infrastructure sources such as VPC Mirroring and Network Load Balancers. The acquired traffic is forwarded to V Series visibility nodes to conduct core intelligence and additional GigaSMART processing.
- Core intelligence: Aggregate, replicate, select or exclude traffic of interest based on Layer 2 to 4 policies (e.g. MAC address, IP address, VLAN ID), optionally send to GigaSMART service functions, balance the traffic load across multiple destinations, and then forward to monitoring and security tools anywhere, as raw packets or tunneled (e.g. L2GRE, VXLAN).
- Traffic intelligence (GigaSMART): Remove duplicate packets from aggregated traffic sources, strip unwanted protocol headers (e.g. GENEVE, MPLS, VLAN), slice off unnecessary payload data, modify header information to obfuscate identifying network information (e.g. MAC address, IP address), mask specific payload data to obfuscate sensitive or private information.
- Application intelligence (GigaSMART): Select or exclude specific applications based on DPI detection of over 3,000 applications, transform flows into application rich metadata selected from over 7,000 attributes.
- GigaSMART service chaining: Apply multiple traffic and application intelligence operations to the same traffic, dynamically, based on tool needs.
- Elastic scale and performance:
  - Use automatic target selection to extract traffic of interest in the infrastructure being monitored.
  - Automatically scale based on varying number of compute instances, without impacting performance.

G-vTAP & UCT Controllers and V Series Proxy

For hybrid and multi-VPC deployments, GigaVUE uses a controller-based design to proxy the command-and-control APIs while preserving existing IP addressing schemes or Network Address Translation (NAT). G-vTAP and UCT Controllers proxy commands from GigaVUE-FM to the G-vTAP Module and UCT instances. GigaVUE V Series Proxy is optionally used to proxy commands from GigaVUE-FM to the GigaVUE V Series nodes. (see Figure 4).

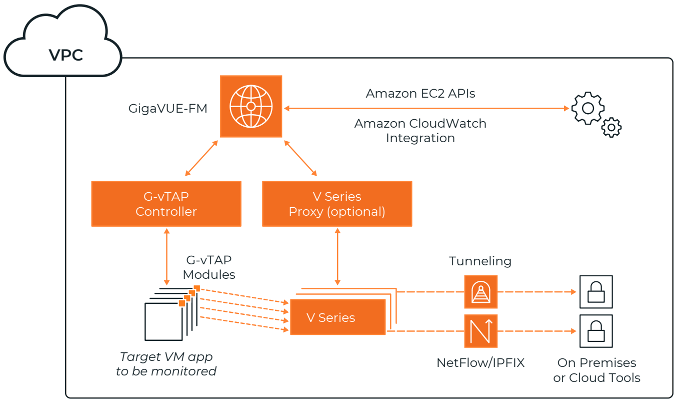
VPC
GigaVUE-FM
Amazon EC2 APIs
Amazon CloudWatch Integration
G-vTAP Controller
V Series Proxy (optional)
G-vTAP Modules
Tunneling
V Series
Target VM app to be monitored
NetFlow/IPFIX
On Premises or Cloud Tools

Figure 4. GigaVUE Cloud Suite for AWS VMs is composed of five components: G-vTAP, G-vTAP Controller, V Series, V Series Proxy (optional), and GigaVUE-FM.

GigaVUE-FM Fabric Manager

GigaVUE-FM provides centralized orchestration and management across the entire organization, including on-prem, public, private, and multi-clouds. FM eliminates manual processes by utilizing auto-discovery methods to identify new workloads in real time and configuring the G-vTAP Module policies to copy and direct traffic to the appropriate GigaVUE V Series node. FM uses AWS APIs for detecting VM changes to dynamically scale these nodes. Further integration with third-party systems automatically adjusts received traffic and configures new traffic policies as needed.

GigaVUE-FM generates an end-to-end topology view via a single-pane-of-glass GUI, which gives you insights into which cloud instances are or are not part of the deep observability pipeline. A single instance of GigaVUE-FM can manage hundreds of visibility nodes across on-premises and multi-cloud environments. Traffic policies are configured using a simple drag-and-drop user interface.

Key features and benefits:

- Centralized orchestration and management:
  - Leverages a single-pane-of-glass GUI for end-to-end topology visualization. Traffic policies are defined using a simple drag-and-drop user interface.
  - Software-defined networking constructs are used to configure traffic policies.
  - Steers packet flows from the G-vTAP Module and UCT or cloud traffic mirroring sources to the V Series, as well as subsequently from V Series to either the monitoring and security tools or to an on-premises physical visibility node.
- Automation:
  - Tight integration with cloud APIs provides auto-discovery of instances, detects changes in the VPC, and automatically adjusts the visibility tier.
  - Under guidance via APIs from the cloud management suite, FM automatically instantiates, configures, scales, and monitors the V Series nodes as needed based on the varying number of compute instances deployed.
  - Open REST APIs published by GigaVUE-FM can be consumed by tools to dynamically adjust traffic received or to orchestrate new traffic policies.

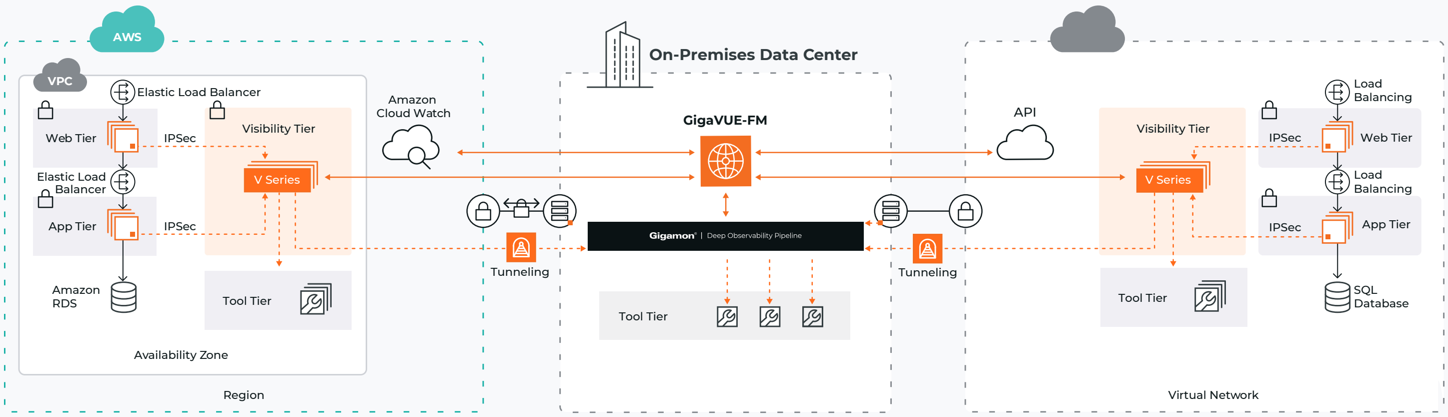

Figure 5. Multi/hybrid cloud deployment with GigaVUE-FM.

Public Cloud Plus GigaVUE Cloud Suite: A Winning Combination

Leading IaaS vendors have developed robust public cloud environments with extensive worldwide availability and solid foundational networking, storage, and security. Some even offer limited traffic mirroring capabilities to acquire network packets, NetFlow (version 5) generation, and some rudimentary filtering. However, these platforms lack the ability to significantly process data; they cannot identify applications and appropriately filter their content, eliminate duplicated packets, drop irrelevant packet header and payload content, or mask source IP addresses for security. GigaVUE works together with cloud platforms, building on their basic functionality to add expanded visibility powers.

As you leverage immense, well-architected and scalable IaaS platforms and expand your use of cloud computing, your progress through your digital transformation may unfortunately also reduce traffic visibility, diminish network efficiencies, and reduce security and monitoring tool effectiveness. These issues will make it harder for you to proactively detect threats, identify deviations from organizational policies, or ensure application performance and exceed SLAs for mission-critical applications, all while minimizing total cost of ownership. And the lack of a well-defined cloud networking and security architecture may end up delaying a move to the cloud.

Gigamon is the leader in pervasive network traffic visibility, and that expertise extends to the cloud as well. With the help of the Gigamon Deep Observability Pipeline, you can use one consistent method across on-prem or multi-cloud deployments to acquire network traffic and apply traffic intelligence — then distribute that optimized traffic to multiple tools. GigaVUE Cloud Suite is a cloud-native solution that enables you to extend your security posture to cloud IaaS, assuring compliance and helping you detect threats to mission-critical applications faster. Now is the time to ensure granular visibility to your workloads and promote an effective security posture, no matter where your data resides.

Function	Gigamon	AWS
Basic VM traffic acquisition	YES	YES
Basic Container traffic acquisition	YES	NO
Traffic direction	YES	YES
Traffic replication	YES	NO
NetFlow v5, v9, IPFIX	YES	v5 only
Basic traffic filtering/elimination	YES	YES
Automatic target selection with L4 flow mapping	YES	NO
Unified management for hybrid/multi-clouds	YES	NO
Packet deduplication	YES	NO
Packet slicing	YES	NO
Packet sampling	YES	NO
Header transformation	YES	NO
Data masking	YES	NO
Application (L3-L7) metadata	YES	NO
Application filtering/elimination	YES	NO

Table 1. Gigamon builds on cloud vendors' platforms to extend visibility and ensure infrastructure control.

About Gigamon

Gigamon offers a deep observability pipeline that harnesses actionable network-level intelligence to amplify the power of observability tools. This powerful combination enables IT organizations to assure security and compliance governance, speed root-cause analysis of performance bottlenecks, and lower operational overhead associated with managing hybrid and multi-cloud IT infrastructures. The result: modern enterprises realize the full transformational promise of the cloud. Gigamon serves more than 4,000 customers worldwide, including over 80 percent of Fortune 100 enterprises, nine of the 10 largest mobile network providers, and hundreds of governments and educational organizations worldwide. To learn more, please visit gigamon.com.

[1] Source: "RightScale 2019 State of the Cloud Report from Flexera." 2019. Flexera. https://resources.flexera.com/web/media/documents/rightscale-2019-state-of-the-cloud-report-from-flexera.pdf.

Gigamon®

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