

Definitive GuideTM

to

Network Visibility and Analytics in the Hybrid Cloud

Continuous Network Visibility and Control,
Both on Premises and in the Hybrid Cloud



Dave Shackelford

FOREWORD BY:

Michael Dickman

Compliments of:

Gigamon[®]

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.

Definitive GuideTM to ***Network Visibility and Analytics in the Hybrid Cloud***

Continuous Network Visibility and Control,
Both on Premises and in the Hybrid Cloud

Dave Shackelford

Foreword by Michael Dickman



CYBEREDGE
P R E S S

Definitive Guide™ to Network Visibility and Analytics in the Hybrid Cloud

Published by:

CyberEdge Group, LLC

1997 Annapolis Exchange Parkway

Suite 300

Annapolis, MD 21401

(800) 327-8711

www.cyber-edge.com

Copyright © 2022, CyberEdge Group, LLC. All rights reserved. Definitive Guide™ and the CyberEdge Press logo are trademarks of CyberEdge Group, LLC in the United States and other countries. All other trademarks and registered trademarks are the property of their respective owners.

Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of the publisher. Requests to the publisher for permission should be addressed to Permissions Department, CyberEdge Group, 1997 Annapolis Exchange Parkway, Suite 300, Annapolis, MD, 21401 or transmitted via email to info@cyber-edge.com.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on CyberEdge Group research and marketing consulting services, or to create a custom *Definitive Guide* book for your organization, contact our sales department at 800-327-8711 or info@cyber-edge.com.

ISBN: 978-1-948939-31-7 (eBook)

ISBN: 978-1-948939-32-4 (Printed book)

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgements

CyberEdge Group thanks the following individuals for their respective contributions:

Editor: Susan Shuttleworth

Graphic Design: Debbi Stocco

Gigamon Contributors: Bassam Khan, Chris Borales, Gordon Beith, and John Gudmundson

Table of Contents

Foreword	v
Introduction	vii
Chapters at a Glance.....	viii
Helpful Icons	ix
Chapter 1: Navigating the Hybrid Cloud Journey	1
Shifting from On-premises to Cloud-based Infrastructure.....	1
How Software-defined Data Centers Are Changing Infrastructure ...	3
Understanding the Shared Responsibility Model.....	4
Top Risks in the Hybrid Cloud.....	5
Chapter 2: Exploring the Roles of NetOps, CloudOps, and SecOps	9
Exploring Network Operations Today	9
Defining CloudOps and Engineering	12
Adapting Security Operations and Engineering to Cloud	14
Chapter 3: Understanding Network Visibility in the Hybrid Cloud	17
Traditional Network Visibility Strategies	17
Challenges in Adapting Network Visibility to the Public Cloud.....	20
Building a Cloud Network Visibility Program	23
Chapter 4: Reviewing the Cloud Visibility and Analytics Fabric	27
What is a VAF?	27
Traffic Brokering in the Cloud.....	35
Chapter 5: Threat Detection and Response in the Hybrid Cloud	41
Modern Cloud Attack Patterns and Workflows.....	41
Cloud Security Response.....	46
Chapter 6: Examining Common Network Visibility, Analytics, and Security Use Cases	49
Improving NetOps and CloudOps Efficiency.....	50
Protecting Resources and Sensitive User Data While Meeting Compliance Mandates.....	51
Eliminating Tool Sprawl and Enhancing Tool Efficiency.....	55

Building a Hybrid Cloud SOC56

Establishing a Zero Trust Security Model..... 57

Chapter 7: Selecting the Right Cloud VAF and NDR Vendor..... 59

 The Need for Complete Coverage..... 60

 Cloud Integration and Compatibility 61

 Core Features Every Mature VAF Should Have63

 Core Features Every Mature NDR Should Have 72

Glossary 75

Foreword

In hybrid cloud, the sum is not yet greater than its parts.

When managing risk across hybrid and multi-cloud infrastructure, you must face the challenge of too many workloads in too many places with too many tools and dashboards. One of the biggest obstacles to monitoring and managing efficiently and effectively is the lack of consistent, granular, end-to-end visibility across the hybrid cloud infrastructure. Network-based tools do not see cloud workloads. Cloud-native tools are unable to leverage the richness and ground truth of network visibility.

Whether through intentional strategy or practical reality, nearly all mid-to-large organizations are running a hybrid cloud. Somehow, in a world of simplification and streamlining, we have more infrastructure options than ever — infrastructure as a service (IaaS), colocation facilities for off-premises hosting, private cloud virtualization from classic hypervisor to hyperconverged infrastructure, containerization in private or public cloud, and existing on-premises infrastructure. IT must balance agility with cost and flexibility with legacy requirements.

IT requires a new approach built on a clear, consistent view into applications, services, and infrastructure. In today's hybrid cloud world, each of these elements offers its own view of what's happening via metrics, events, logs, and traces (MELT). *Observability* tools capture this self-generated telemetry to inform performance monitoring and application troubleshooting. *Deep Observability* completes the picture by providing immutable, complete, and correct network-level intelligence about what each element is actually producing and communicating. Because this completeness can produce large data streams, intelligent analytics are required to pull signals from the noise and direct the right information to the right tool at the right time. The sum of deep observability is greater than what either observability or visibility enable on its own.

Success in hybrid cloud requires a visibility and analytics

fabric for all data in motion. Service level objectives (SLOs) require end-to-end visibility that doesn't stop at the border of a particular hypervisor, cloud platform, or physical datacenter. Security requires deep observability, without blind spots where cloud-based tools don't see network data and network-based tools don't see cloud data.

Network visibility and analytics in the hybrid infrastructure empowers NetOps, CloudOps, and SecOps to fulfill their charters more effectively, faster, and at lower cost. NetOps can ensure reachability and performance between on-premises and public cloud, and from cloud to cloud. CloudOps can troubleshoot problems between overlay network segments, digging as deeply as needed to bridge the gap between siloed, single-platform tools. SecOps can apply consistent tools and processes to ensure no threat actor slips through a gap in the complexity of hybrid cloud infrastructure.

This guide provides an overview of hybrid cloud challenges and implications for various roles in IT, then describes how to build a cloud network visibility program to operationalize a visibility and analytics fabric. Special focus is given to security where network-based threat detection and response builds directly on top of deep observability. You will learn how to address common use cases around operational efficiency, tools optimization, security operations center (SOC) enablement, compliance, and more. Finally, the guide closes with criteria to select a network visibility and analytics solution for your hybrid cloud. These criteria include foundational capabilities like traffic forwarding, decryption, de-duplication, and cost optimization, as well as advanced capabilities like application intelligence, metadata generation, full visibility into diverse IaaS environments, and integration with the full ecosystem of both network-centric and cloud-native tools.

We at Gigamon hope you find this guide useful and informative. Our entire company is built around getting the right information and insight to the right place at the right time. Only then is the sum truly greater than the parts. Best wishes for success in your hybrid cloud journey!

Michael Dickman
Chief Product Officer, Gigamon

Introduction

Hybrid cloud adoption is expanding due to a shift to cloud-based applications and remote work scenarios, a plethora of widely differing device and access types, and more. At the same time, with the increasingly advanced attack landscape they're facing, overtaxed infrastructure and security operations (SecOps) teams need centralized, pervasive visibility into their hybrid environments.

In private and public cloud, teams often start with basic observability of the infrastructure and applications. However, the observability data does not provide context and lacks visibility into unmanaged devices, such as bring your own device (BYOD), Operational Technology (OT), and Internet of Things (IoT) devices. Hence, operations and SecOps teams have recognized that network traffic is actually the ultimate source of truth, and can be combined with observability data to provide a complete picture.

Fortunately, there are advanced hybrid cloud visibility solutions that present a unified monitoring and control fabric and security analytics. These solutions offer flexible and scalable brokering and analytics for all network traffic, selective filtering and payload removal, header and payload field obfuscation, and protocol and application metadata generation. Application performance monitoring and security tools can ingest, analyze, and process the metadata to build behavioral and early detection models of network activity, and detect and investigate threats.

Hybrid cloud visibility fabrics should be considered essential to gaining insight into infrastructure and application behavior across all environments. This guide is designed to help security and infrastructure administrators, managers, and executives leverage the power of these solutions.

Chapters at a Glance

Chapter 1, “Navigating the Hybrid Cloud Journey,” discusses how IT teams are shifting from purely on-premises infrastructure to hybrid cloud-based network environments, and what that means.

Chapter 2, “Exploring the Roles of NetOps, CloudOps, and SecOps,” reviews how operational roles and tasks are changing in today’s hybrid infrastructure, along with examples of tools and use cases.

Chapter 3, “Understanding Network Visibility in the Hybrid Cloud,” outlines what is changing in the move to cloud-based network architectures, how blind spots can be easily introduced, and how these blind spots affect visibility.

Chapter 4, “Reviewing the Visibility and Analytics Fabric,” describes the core elements of a hybrid cloud visibility fabric and its benefits.

Chapter 5, “Threat Detection and Response in the Hybrid Cloud,” explores the network threat landscape in the cloud, and what controls are needed to combat these threats, including ransomware.

Chapter 6, “Examining Common Network Visibility and Analytics Use Cases,” considers a variety of common use cases for enterprise network visibility and analytics platforms.

Chapter 7, “Selecting the Right Network Visibility and Analytics Vendor,” highlights the kinds of capabilities security and operations teams should look for when evaluating network visibility and analytics solutions and the vendors that provide them.

Helpful Icons



TIP

Tips provide practical advice that you can apply in your own organization.



DON'T FORGET

When you see this icon, take note as the related content contains key information that you won't want to forget.



CAUTION

Proceed with caution because if you don't it may prove costly to you and your organization.



TECH TALK

Content associated with this icon is more technical in nature and is intended for IT practitioners.



ON THE WEB

Want to learn more? Follow the corresponding URL to discover additional content available on the web.

Glossary

When you stumble across a word or phrase that is italicized, the odds are pretty good that you'll find a helpful definition in the Glossary.

Chapter 1

Navigating the Hybrid Cloud Journey

In this chapter

- Shifting from on-premises to cloud-based infrastructure
- Understanding the shared responsibility model
- Evaluating top risks in the hybrid cloud

When organizations deploy assets to the cloud, a number of changes in security controls and architecture design are often needed. Many controls currently managed in house will now be managed by the cloud service provider; scalable access to traffic will become more difficult; and long relied-upon third-party solutions may function differently, or not at all, in public clouds.

Shifting from On-premises to Cloud-based Infrastructure

Moving infrastructure to a cloud-based model requires an understanding of what a *hybrid cloud* represents, and what changes are involved. Figure 1-1 shows that organizations are running significant workloads and assets in the public cloud. However, some tools and security controls that have traditionally worked on premises won't always work in the cloud.

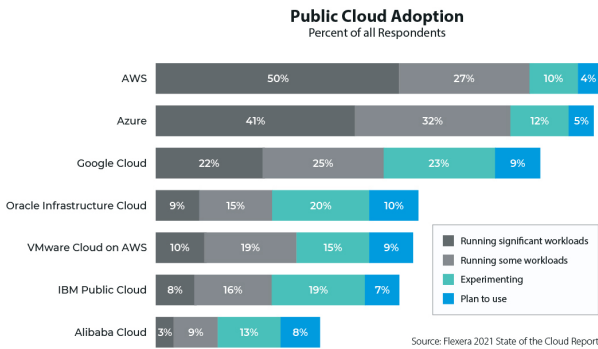


Figure 1-1: Public cloud adoption trends for enterprises.

Defining the hybrid cloud

Most security and IT teams agree that a “hybrid” cloud exists when some assets are maintained internally and others are moved to one or more public cloud environments (multi-cloud scenarios are increasingly common). Figure 1-2 highlights the growth in moving workloads and data to the public cloud. Despite this trend, many assets are still connecting back to on-premises environments.

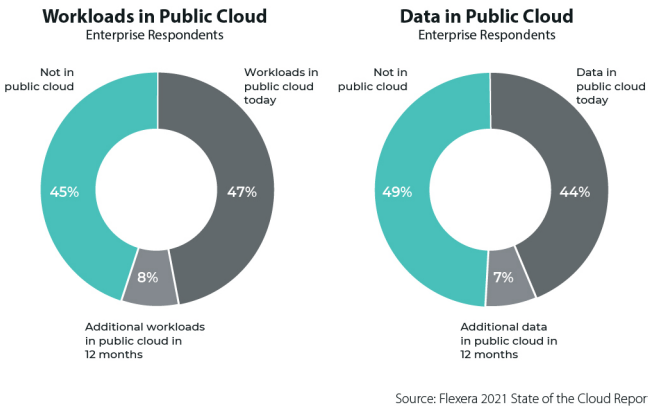


Figure 1-2: Enterprise workloads and data in the public cloud.

How Software-defined Data Centers Are Changing Infrastructure



Cloud environments represent the concept of the *software-defined data center (SDDC)*, which differs in several important ways from traditional physical data centers:

- ✓ Systems are virtual, not physical.
- ✓ The cloud provider fabric is a software platform that all cloud assets are connected to.
- ✓ A wide variety of APIs are available to enable operational activities such as instance provisioning, infrastructure monitoring and system configuration.

These changes brought about by SDDCs are forcing organizations to modify operational and security practices in many ways.

Challenges in adapting security controls

Many organizations are struggling to adapt their internal security controls to cloud environments. In some cases, existing tools and products aren't well suited to cloud provider environments. Also, many security processes and workflows built around traditional on-premises tools and controls will likely need to be adapted and reconfigured when an organization moves to the public cloud environment.

Migrating from on-premises infrastructure to cloud environments



Common challenges in migrating from on-premises to cloud infrastructure and controls include:

- ✓ Lack of feature parity in cloud versions of on-premises products and services
- ✓ Unacceptable performance for cloud security controls and services that increase costs (cloud providers charge for more utilization)
- ✓ Lack of compatibility with certain cloud providers' automation, orchestration, and deployment architecture

Understanding the Shared Responsibility Model

In a hybrid cloud model, some security responsibility is assumed by the cloud provider, while other security controls and capabilities are managed by customers. Figure 1-3 highlights the controls and security requirements that organizations retain responsibility for versus those the cloud providers assume for infrastructure as a service (IaaS)-based workloads.

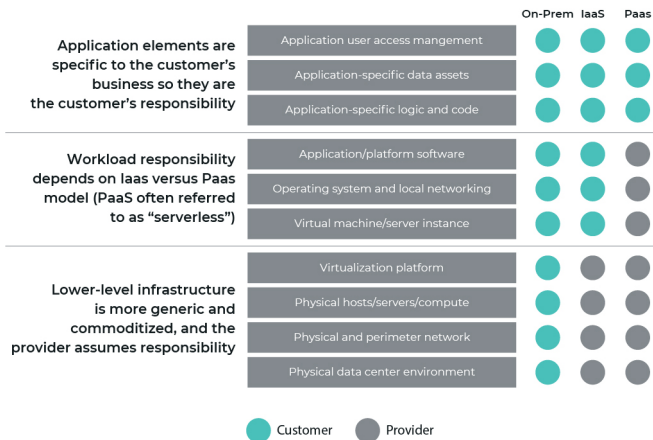


Figure 1-3: Public cloud shared responsibility model.



Shared responsibility for security is a foundational, and often underappreciated, model for public cloud. Once organizations understand this model, they can begin designing the security controls and architecture that they are responsible for.

Controls that cloud providers manage

In IaaS clouds, service providers maintain the data centers and underlying SDDC infrastructure (physical systems, virtualization hypervisors, and layers of software integration), while customers must secure all assets they create (workloads, storage, and more).

Controls that are the enterprise's responsibility

In a shared responsibility model, there are always some controls that enterprise tenants will be responsible for maintaining.



In IaaS clouds, customers are responsible for the security of most assets and workloads, as well as cloud services. Network and identity security are the responsibility of the customer, as well. While cloud providers maintain the infrastructure that facilitates a SDDC, any other infrastructure is entirely under the control of the customer. This fact should drive all security control and process requirements in the IaaS cloud.

Limitations of IaaS providers' security offerings

Cloud providers' security offerings may be easy to implement and affordable, but usually don't have extensive capabilities like mature third-party solutions do. Cloud providers' security tools and services often provide basic, foundational functionality that may be adequate for smaller organizations. However, larger organizations and those with more-extensive compliance and security requirements may find them less flexible, less compatible with existing tools and processes, and restricted in scope to a single cloud environment, thus limiting their usefulness in multi-cloud deployments. IaaS providers offer some security benefits, mainly in the areas of resilience and physical security, but leave the vast majority of security controls and management to customers.

Top Risks in the Hybrid Cloud

While cloud services offer many advantages, there are significant risks both cloud providers and customers need to address, as well.

Why attackers are targeting the hybrid cloud

The cloud opens up an entirely new attack surface. Data could potentially be exposed in transit, in storage within the cloud,

while applications and systems are operating in IaaS and platform as a service (PaaS) environments, and so on. Attackers are targeting cloud environments more than ever before, primarily because organizations are moving more-sensitive data and workloads there, and many are not skilled at properly securing their cloud deployments...a perfect storm!

Many types of attacks against cloud-based assets are common today, and some can be exploited for use in ransomware attacks, as well:

- ✓ Attacks against cloud admins and users (account hijacking)
- ✓ Hijacking of resources to run cryptomining servers, leaving cloud customers with the compute charges
- ✓ Discovery and theft of exposed cloud storage assets
- ✓ Exploitation of vulnerable cloud applications with zero-day attacks against the common OWASP Top 10 vulnerabilities that many cloud-native web application firewalls may not catch.
- ✓ Exploitation of exposed cloud provider APIs
- ✓ Compromise of misconfigured cloud control plane services

Attacks against cloud fabric configuration

The cloud provider environment itself is a complex software platform, with many different settings and options that can be misconfigured. For example, failing to enable multifactor authentication for administrative portal access could lead to brute force attacks against accounts. Exposed provider APIs may afford attackers a foothold to enumerate cloud assets or even gain remote access to cloud environments.

Network security in the hybrid cloud

In hybrid cloud networks, network engineering and security architecture teams need to look at new controls and new methods for developing and implementing network security practices in their environments. Organizations should imple-

ment a layered defense model that includes cloud-native network access controls for workloads, network traffic capture and security monitoring capabilities, and third-party platforms and services that can help to augment and centralize network security functions.

Cloudy with a Chance of Risk

In cloud environments, numerous risks can potentially manifest. While many are similar to traditional on-premises risk scenarios, a number are specific to cloud environments:

1. Poorly configured cloud fabric services and APIs
2. Lack of comprehensive monitoring for the cloud fabric, applications, and services
3. Excessive privileges and vulnerable accounts in cloud service environments

Chapter 2

Exploring the Roles of NetOps, CloudOps, and SecOps

In this chapter

- Exploring network operations today
- Defining cloud operations and engineering
- Adapting security operations and engineering to cloud

With a shift to cloud, the traditional job functions of network operations (NetOps) and security operations (SecOps), and the newer realm of cloud operations (CloudOps), are converging.

In this chapter, we'll review common roles and responsibilities of IT operations professionals today and discuss how they've evolved to meet the needs of our new, dynamic, cloud-based world.

Exploring Network Operations Today

As more organizations move to the cloud, the nature of many IT operations roles is changing, and new roles are rapidly emerging as well.

Current NetOps roles and responsibilities

In most NetOps teams, there are a variety of different roles and responsibilities:

- ✓ Network architect
- ✓ Network engineer
- ✓ Network operations specialist

Regardless of titles, NetOps teams are responsible for building network infrastructure that is highly available, responsive, and capable of supporting all necessary traffic for the organization, ideally exceeding service level agreements (SLAs) in these categories.

How network architecture is changing in the hybrid cloud

Network architecture can be significantly different in the cloud compared to on-premises environments, primarily due to different connectivity requirements between cloud services and assets, as well as more-flexible options coupled to the cloud provider backplane. For example, a mix of both cloud-native and third-party access controls may be implemented, and isolation strategies will involve concepts like creating separate security *virtual private clouds (VPCs)* and *virtual networks (VNETs)*.

Network operations and controls in the public cloud

Most major networking functions and controls, such as routing, load balancing, network traffic monitoring, and network segmentation/isolation, are readily available in the public cloud, although the implementation of these controls and technologies usually differs somewhat from traditional on-premises networks. Today, all critical network functions in the cloud are embedded in virtual infrastructure and are software based. Network operations teams will need to get used to having all network functions defined in software versus hardware

and adapt some technologies and processes to accommodate a hybrid architecture with both on-premises and cloud-based platforms and controls.

Today's top cloud network skills and functional needs

Traditional networking skills in routing, switching, load balancing, and network security still have value for cloud deployments. However, definitive changes in the realm of network operations are needed when organizations move to the cloud.

DON'T FORGET



- ✓ With a move to *infrastructure as code (IaC)*, network configuration needs to be defined in a template format such as AWS CloudFormation, Azure Resource Manager, or Terraform by HashiCorp, and network teams should be comfortable with this.
- ✓ All networking platforms in the cloud are virtual appliances, often deployed through a cloud provider marketplace. These virtual appliances may have configuration and operations requirements that are different from similar systems in physical data centers.
- ✓ Cloud-native networking functions, such as load balancing, network access control, and cloud-based routing, are all particular to a specific cloud provider's environment. Network operations teams will need to learn the specific options and API capabilities available for their chosen environment.

TIP

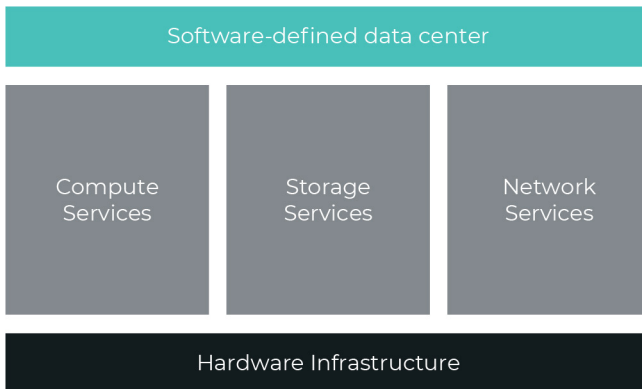


Figure 2-1: Components in a software-defined data center (SDDC).

Defining CloudOps and Engineering

Moving to the cloud will require new architecture and operational controls and platforms. Let's dive into them.

Core cloud architecture concepts

As technology services and assets are defined in software in the cloud, and cloud data centers tend to be much more dynamic than traditional data centers, a number of important architecture concepts need to be implemented:



- ✓ Build in security at every layer of the environment (storage, workloads, networking, and control plane). Often referred to as the principle of defense-in-depth, this is still critically important in the cloud, although some of the controls may be different.
- ✓ Design for failure scenarios by implementing automated detection of failure conditions that trigger high availability and failover controls for all major network functions.
- ✓ Design for elasticity, as the environment will scale up and down to meet business needs. Network platforms and controls will need to scale at the same time based on operational thresholds.
- ✓ Centralize operational management wherever possible, especially in a hybrid cloud that requires centralized controls and services.

Moving to software-defined infrastructure

When shifting to a software-defined infrastructure model, organizations will need to modify many traditional networking concepts to adapt to an environment where hardware management is the responsibility of the cloud provider. All platforms and network functions will be entirely software based and will interact with software-based assets and objects, services, and cloud provider APIs.



Cloud automation and orchestration

As more of an organization's infrastructure is based in software, with numerous new APIs and interaction methods for infrastructure components, coordination can become challenging due to the larger number of workloads involved. This is why a unified fabric for orchestration and automation of operations across network functions like monitoring and detection/response, traffic control, and workload communications becomes more important than ever.

Building a sustainable cloud fabric

Security and network operations teams will need to focus on centralized solutions that are highly scalable as organizations make use of multiple cloud services and enable various types of cloud workloads and infrastructure capabilities. A sustainable cloud fabric should integrate with leading cloud services like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), as well as workload agents and monitoring, standalone monitoring nodes, and automation and orchestration tools and capabilities. Figure 2-2 highlights a comprehensive orchestration and automation model that focuses on workload and network visibility, monitoring, and security.

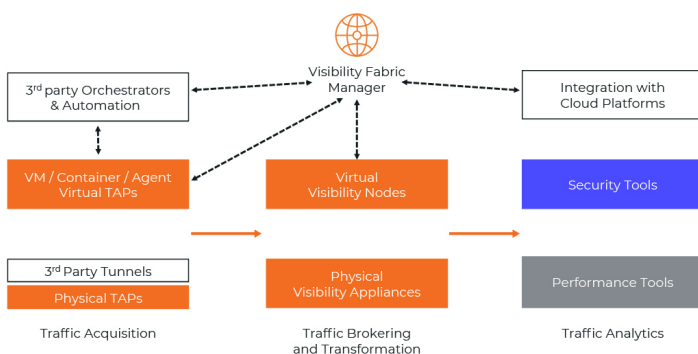


Figure 2-2: Orchestration and automation model.

Adapting Security Operations and Engineering to Cloud

With the shift to cloud, SecOps teams will need to adapt a variety of tools and processes to account for all the differences between cloud and on-premises environments.

Current roles and responsibilities of SecOps



A robust and capable SecOps team should provide continuous monitoring capability, as well as implementing and maintaining a full cycle of prevention, detection, event analysis, and incident response.

In the cloud, many SecOps functions and processes are also handled by multidisciplinary teams, and some distinct specialty skills for each major cloud provider may also be needed.

Key cloud security functional requirements

Designing a SecOps program for the cloud involves several important functional requirements for key solutions:

- ✓ Compatibility with leading cloud services through integration in the provider's marketplace
- ✓ Powerful performance tuning controls to minimize overhead in cloud environments
- ✓ Integration across other public cloud and on-premises infrastructures, particularly for securing applications and services when and as they migrate
- ✓ Integration with cloud provider APIs and automation and orchestration controls
- ✓ Centralized administration and reporting with role-based access controls

Defining and implementing cloud security controls

Security operations teams should define a core stack of cloud security controls that can be managed and adapted across cloud providers, where possible. Areas of coverage should include hypervisors and converged infrastructure (for private clouds, primarily), network security, data protection, workload protection, and application security. Many of these categories have numerous subcategories; for example, network security should include:

- ✓ Network isolation and segmentation
- ✓ Network identity and access management
- ✓ Network monitoring and traffic analysis
- ✓ Regulatory and corporate compliance
- ✓ Network detection and response (NDR)
- ✓ Distributed denial of service (DDoS) protection
- ✓ Availability and resiliency

Cloud guardrails and the “security fabric” of the cloud



Cloud architects and security engineering professionals have been discussing the idea of “guardrails” in the cloud, which can be a confusing concept. Most professionals in the cloud security space define guardrails as defensive services and controls that are automated and continuously operational, and that directly feed into detection and response processes and practices.

No More Blinking Lights!

In the cloud, only software-based networking is available. It consists of two primary implementations. First, cloud-native networking options can be implemented as services and configurations for objects and assets. Second, virtual appliances can be integrated into

the environment using cloud provider or third-party orchestration methods in combination with cloud provider APIs. Both are often used in mature enterprise environments, and neither comes with blinking lights.

Chapter 3

Understanding Network Visibility in the Hybrid Cloud

In this chapter

- Reviewing traditional network visibility strategies
- Understanding the challenges of adapting network visibility to the public cloud
- Building a cloud network visibility program

When organizations move to the cloud, they often realize that network visibility is significantly reduced. Fortunately, a variety of tools and technologies are available now to help improve cloud network visibility for operational and security monitoring.

Traditional Network Visibility Strategies

In most mature enterprise network environments, security and network operations teams work diligently to provide adequate network visibility for performance monitoring, troubleshooting, and threat prevention and detection.

Key network monitoring, security, and control platforms

These enterprise teams use a variety of common network monitoring and control platforms within their data centers:

- ✓ Network *test access points* (TAPs) and packet brokering infrastructure
- ✓ Network detection and response (NDR) platforms
- ✓ Network intrusion detection/prevention systems
- ✓ NetFlow collectors and monitoring tools
- ✓ Full packet capture systems for forensics and traffic analysis
- ✓ Application performance monitoring (APM) tools
- ✓ Network configuration management systems

Network traffic monitoring and visibility on premises

Security analysts today need to be aware of everything going on within their environment. Focusing on network traffic is a major element in intrusion detection and prevention efforts.

Incident responders and security analysts have a critical need for packet analysis and inspection capabilities. Network traffic, and individual packets or streams of packets, can be captured and analyzed in many ways. Many types of network devices may be able to inspect traffic. Common examples include application delivery controllers, proxies, enterprise firewalls, and network intrusion sensors.

In addition to common methods like packet filtering and stateful filtering, deeper analysis of network traffic has become feasible. As hardware has become more powerful, we've seen a shift toward application proxies that can perform this type of in-depth analysis. The latest "next-generation" firewalls are offering far more advanced capabilities than ever before, including behavioral and anomaly detection, as well as event generation and filtering using end-user IDs.



Traffic access and packet capture

Security analysts who need to see traffic at wire speed have a few options for getting the traffic to their monitoring platforms. These include inline devices like firewalls and intrusion prevention systems (IPS), which normally have “port pairs” for inbound and outbound traffic that enable the data to flow through the devices. In addition to inline options, the two most common solutions seen in enterprises today are switched port analyzer (SPAN) ports and network TAPs (physical or virtual devices that split or copy traffic control and replicate traffic). Passive traffic mirroring is much more common than inline options in cloud, as well.

A *SPAN port* function, in essence, copies traffic from one or more network ports on a switch to a designated SPAN port on the same switch. Any network detection device can thus receive traffic associated with a switch and alert on it. While SPAN ports are a simple and relatively inexpensive option, they can lead to switch backplane congestion, SPAN port oversubscription, variable latency, and dropped frames due to buffer overflow. They also, by definition, require dedicated ports on the switch.

TAPs, on the other hand, tend to be more efficient and provide a 100% copy of both directions of a link, while freeing up switch ports. Most fiber optic TAPs are also unpowered.

In cloud environments, virtual equivalents of SPANs and TAPs are also available via software-based mirroring services and tapping agents – some requiring close integration between the mirroring/tapping function and the receiving element.

Types of network monitoring: packets and flow data

Any seasoned security analyst should have a good understanding of how packets can reveal details about intrusion attempts or successful hacking scenarios. All security analysts should have the ability, at a minimum, to acquire and assess packet capture files in a standard format. Gathering packets can be done in a number of ways, as discussed above. However, ensuring that the files are compatible with a variety of tools, and can be interpreted properly, is a skill unto itself.

Further, analysts should know the basics of what to look for in packets and flow data when searching for the major indicators of malicious activity.

Network flows are records of communications between systems in a network environment. NetFlow records a common set of basic information, including:



- ✓ Source and destination MAC addresses
- ✓ Source and destination IP addresses
- ✓ Source and destination ports (type and code for ICMP)
- ✓ Layer 3 protocol in use
- ✓ Type of service
- ✓ Ingress interface for traffic
- ✓ Ingress and egress byte counts

Collecting flow data may provide some insight into what is really going on in your network. However, enabling NetFlow on switches can overload these devices or lead to missed flows, and out-of-band NetFlow generators are more common in large enterprises today.

Network security teams should sniff traffic in specific locations to gather flow content and then focus on anomaly detection in traffic patterns. Anomalous behavior might include unusually short or long lengths of session connections, connections to systems that are unexpected, non-standard port usage or addresses seen in the network, and more.

Challenges in Adapting Network Visibility to the Public Cloud

The last several years have been interesting for most security teams as they sought to adapt traditional on-premises controls to cloud-based infrastructure. These professionals have tried to install tool-based agents, one agent for each tool, on workloads, set up monitoring sensors to collect cloud network

traffic, and observe all communication flows across a broad array of cloud systems and services.

Agents can add enormous overhead to workloads, which is not ideal. The move to cloud security operations has been cumbersome to achieve, largely due to the operational headaches involved in installing agents, enabling network traffic inspection and monitoring, and configuring centralized tools that can help the organization to keep track of it all.

There's still a need to provide in-depth coverage for network visibility at the workload, network, and cloud fabric levels. Figure 3-1 illustrates how all of these monitoring elements need to be considered in cloud environments today.

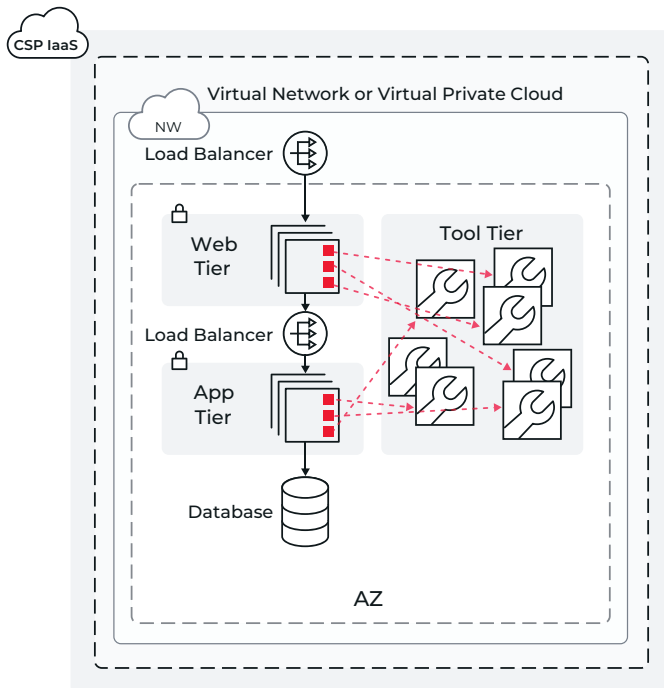


Figure 3-1: Cloud visibility challenges.

Gaps in cloud provider services and capabilities

Unfortunately, cloud providers rarely have mature services available to access and collect network traffic. In contrast, most on-premises environments have some combination of network tapping and traffic mirroring on switches to enable packet capture for security and operational purposes. Also, network flow data has been readily available to feed to analysis tools. In the cloud, a variety of packet capture methods have only very recently become available, and while flow data can be collected, no native tools or services can adequately parse and analyze this network data.

Lack of tooling and third-party provider support



As switching is no longer available for configuration by public cloud tenants, all access control and segmentation is applied via policies that usually only focus on layers 3 and 4 (IP addresses and ports) as well as tags or object references. Additional networking tools can be integrated into the SDN controller or other infrastructure management systems to provide more-granular micro-segmentation, more-capable routing, and intrusion detection. Some types of network controls and tools traditionally used in physical data centers have not adapted well to cloud provider environments either, complicating efforts to build a comprehensive network security stack.

Performance and telemetry challenges

In addition to security visibility challenges, many network teams have struggled with network traffic in the cloud that is: a) lagging in performance and quality of service (QoS); and b) difficult to monitor and track for telemetry. For operations teams that need to carefully track and control network performance and continuity, these issues can be significant. Cloud-native tooling for network traffic flows is rarely designed with optimal performance and tracking in mind.

Building a Cloud Network Visibility Program

Achieving network visibility in the cloud today requires a different set of tools than is natively available in the cloud.

Assessing current network monitoring tools and capabilities

The first step toward network visibility in the cloud is assessing your current network security and operations controls. What do you have? And perhaps most importantly, what do you need?

The need for a “single pane of glass”



In building a network security strategy for the hybrid cloud, it's important to enable networking controls and capabilities everywhere to prevent a disjointed and ineffectual control model. This means looking for solutions that can span both on-premises and cloud environments, while emphasizing centralized management for teams that are already managing numerous technologies and processes. With a “single pane of glass” to implement and manage controls irrespective of location (cloud or data center), operations teams will be more successful overall.

Cloud-native services and capabilities

Virtual private clouds (VPCs) and network subnets should be carefully created and linked together, including using transit gateways as needed to create the network scale and architecture desired.

Then, the first type of network monitoring control that organizations should enable within the hybrid cloud is the collection of network flow data for monitoring communication to, from, and between workloads within network segments. Flow logs can be used to monitor and track network events and behaviors on a large scale. Flow logs can help security teams in troubleshooting and analyzing security group rules, monitoring traffic communicating with workloads, and determining

the direction and patterns of traffic to and from cloud network interfaces.

Next, network access controls like security groups should be used to their fullest extent, with logging of all events related to enabled rules for workloads and subnets. After this, packet capture should be enabled and the packets sent to a centralized network visibility and analytics fabric before being selectively forwarded to monitoring and security tools.

Capturing and processing network traffic and information



TIP

A critical cloud capability sought by many network security teams is full control of network packet capture, which most providers have not facilitated until very recently. Traffic mirroring services are now available in several leading IaaS clouds. Traffic mirroring permits network traffic to be copied from the data plane in a network zone and sent to a suitable destination such as a monitoring or security tool, network load balancer, network visibility and analytics fabric, etc. For IaaS clouds that do not have native traffic mirroring services, a network visibility and analytics (NVA) fabric can provide a single mirroring or tapping agent per workload instance.

Advanced NVA fabrics can now leverage one of the available mirroring capabilities to pull traffic to/from workloads in the IaaS cloud and selectively optimize or transform the traffic flows. This capability allows security operations teams to effectively and efficiently perform deep packet inspection and network forensics.

It is possible for many organizations to adopt a centralized network visibility and analytics fabric that forwards critical workloads' traffic to a virtual (or physical) network packet broker (aka: a visibility node) (Figure 3-2).

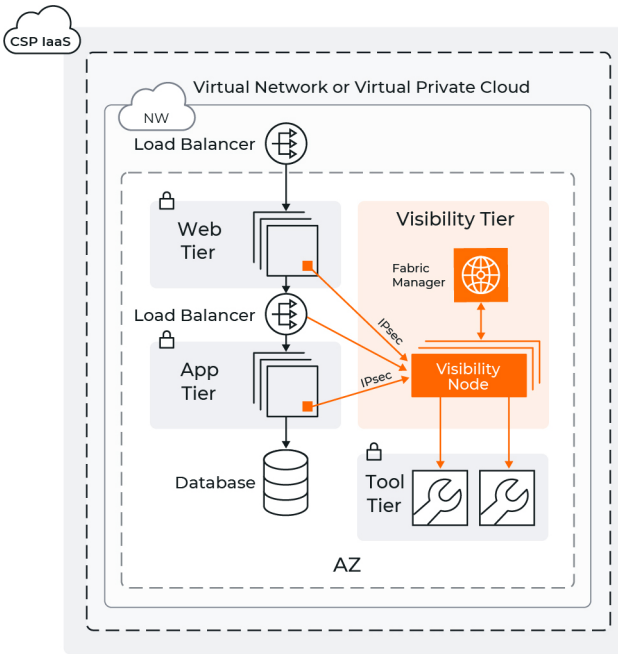


Figure 3-2: Addressing cloud visibility challenges.

Chapter 4

Reviewing the Cloud Visibility and Analytics Fabric

In this chapter

- Defining the visibility and analytics fabric (VAF)
- Understanding traffic acquisition
- Examining traffic analysis in the cloud

To achieve full network and application visibility into cloud in addition to on-premises environments, it's important to enable a new set of tools and services. A purpose-built *visibility and analytics fabric (VAF)* can help to provide IT and security operations teams with a more comprehensive view of heterogeneous network environments.

What is a VAF?

The concept of a VAF is growing in importance in hybrid cloud infrastructure environments. The VAF is the visibility layer between network infrastructure and network monitoring and security tools. It helps ensure each tool receives all the relevant traffic it was designed for (acting as a comprehensive next-generation network packet brokering platform, in essence).

Full network visibility for the cloud

As network operations moves from monolithic physical appliance deployments to self-managed virtual and cloud-based workloads, and application deployments shift to cloud-based

scenarios, managing and securing networks and applications become more challenging. Many organizations have relied on log files and other data from network functions and applications themselves to provide complete visibility into network behavior.

However, this approach is unreliable due to changes in log formats and generation of logs from specific services. Also, it requires constant development as functions and applications change over time. Logs, combined with traces and metrics, are insufficient as they lack the actual packet header and payload details that must be analyzed for the presence of potential malware and for traffic levels to help evaluate network performance. In addition, logs exclude advanced metadata necessary to address key use cases.



Whether you're deploying workloads in the hybrid cloud or physical data centers, highly available operation, strong security, and regulatory compliance of network data and applications require a sound strategy and a unified fabric for monitoring and analytics. To automatically and proactively identify and remediate security and performance limitations, you should have granular and comprehensive packet-level visibility into all network environments.

Acquiring and accessing network traffic

To analyze network information across a hybrid cloud environment, organizations need to access and acquire network traffic regardless of where it flows. For many, there's a definite need to implement sound network monitoring via traffic acquisition. Network monitoring may take many forms: some are operationally focused, and others are aimed at security and compliance. What types of systems and controls are typically included in a network operations and security monitoring architecture? Here are some of the most common implementations found in physical on-premises networks:

- ✓ **Software-based monitoring:** In some cases, software-based network monitoring relies on protocols such as the Simple Network Management Protocol (SNMP) to coordinate periodic polling of network devices and other systems to gather configuration and statistical data about their status and performance. Many software-based solutions, often in virtualized formats, are available today.
- ✓ **Network traffic capture hardware:** NetOps teams may implement hardware-based traffic analysis tools such as TAPs in network links and mirroring ports on switches to replicate network traffic to monitoring and security tools for capture and analysis.
- ✓ **Performance monitoring functions:** Monitoring traffic flows and throughput, as well as uptime and availability for workloads and applications, is a critical element of network telemetry today. Operations teams rely on network performance monitoring to track SLAs.
- ✓ **Security monitoring functions:** Security teams often deploy or leverage existing network traffic capture hardware, usually in the form of an intrusion detection system (IDS) or network forensics capture device. In addition, many organizations send traffic flow data and events to aggregation and correlation systems such as log management platforms and security information and event management (SIEM) products. Recently, some have begun sending network traffic to network detection and response (NDR) solutions.

DON'T FORGET

Without full access to packet data, organizations are unable to build a complete picture of network activity for operational and security monitoring. In cloud environments that don't currently offer native traffic mirroring, it's vital that a VAF platform include agents and other methods to access traffic and replicate the desired traffic types to specified systems and services. Any enterprise-class VAF must be able to access and acquire traffic in a diverse set of environments and provide intake for a variety of workloads.

Using cloud network TAPs and mirroring capabilities

There are many pitfalls associated with the way organizations are gathering traffic for network monitoring today. Most organizations that are monitoring traffic in physical networks are enabling SPAN/mirror ports on switches or deploying inline network TAPs. The use of SPAN ports creates a number of potential operational challenges. First, most switches cannot feasibly support more than one or two SPAN ports, which are typically inadequate for the total network traffic being copied and result in arbitrarily dropped packets. Also, SPAN ports can create performance issues by overloading the switch backplane and reducing overall throughput.



TAPs, on the other hand, capture at full line rate, don't impact the network, and are more scalable than SPAN ports. But TAPs may be limited in their features and centralized control options – which are not necessary with passive network TAPs.

In virtualized private or public clouds, depending on the infrastructure in use, there may be built-in mirroring or tapping services that can be used to acquire both north-south and east-west traffic (that is, traffic between virtual machines or containers, which constitutes over 80% of the total traffic in most scenarios) See Figure 4-1.

A network visibility vendor can also provide a tapping agent that is instantiated alongside the workloads being monitored. In containerized infrastructures, it may become more complex, requiring a level of data translation integration between the mirroring/tapping functions and the receivers of the copied traffic.

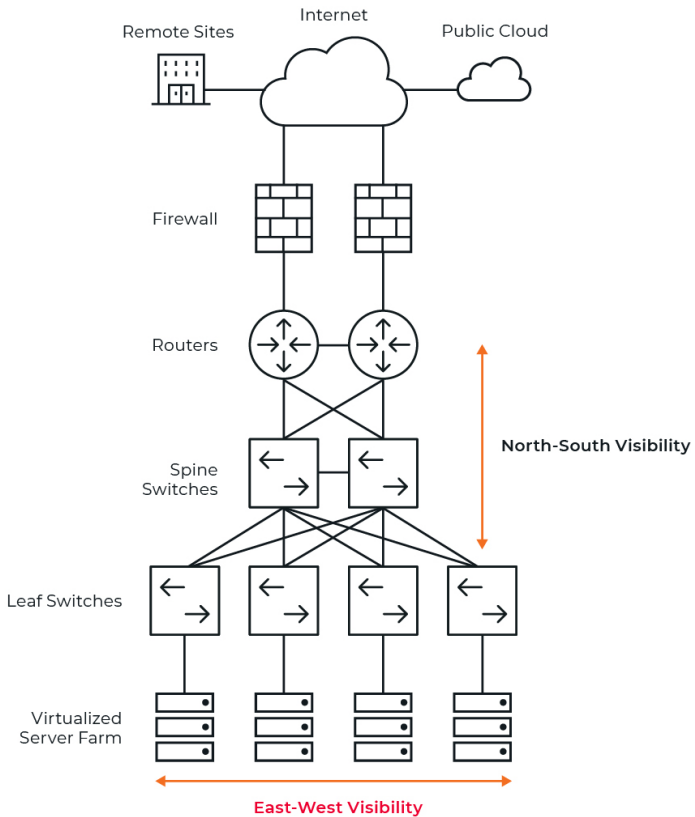


Figure 4-1: East-west and north-south traffic monitoring. Traffic analysis

The landscape for effective traffic access and analysis involves a shift from standalone TAPs or SPAN ports to a unified “monitoring layer” that incorporates a dedicated TAP/access fabric model capable of extremely rapid and detailed traffic analysis. In the past, traffic acquisition tools and methods largely copied traffic from one place to another without performing any analysis of the traffic being processed. That model is no longer sufficient, as the volume and diversity of traffic types in many environments necessitate a processing/filtering layer that can control which traffic types are forwarded and which are left alone.

Virtual machines

Virtual machines are the most common and consistent implementation of traditional server platforms today, and operate similarly in both on-premises and cloud-based environments. All network analytics solutions need to accommodate virtual machine workloads.



An important consideration in VAF effectiveness is operating system support. A VAF should support Windows and Linux platforms equally for maximum workload coverage. Additionally, minimal resource utilization is a key benefit for any VAF agents installed with a shared resource model like virtualization or cloud virtual machines.

Containers

Containers are increasingly common in both on-premises and cloud workload deployments. Capturing network traffic from containers is tricky, though – all containers share a host platform, and that’s the only place an agent can be installed. Additionally, many containers running in cloud-based service environments don’t give tenants control over the operating system, so the only way to capture traffic is to integrate with cloud provider APIs, deploy a unique monitoring container, or integrate with each container image.



Any enterprise-class VAF should provide ample capabilities to monitor container traffic, and also integrate with industry-leading orchestration and automation platforms like Kubernetes. Large-scale deployments of containers can include thousands of microservices, and the pace of change in these environments can be daunting. The more fully VAF solutions are integrated with orchestration tools, the better aligned and automated traffic acquisition, analysis, and processing will be.

Built-in traffic mirroring services from cloud providers

In recent years, leading cloud providers have begun offering native traffic mirroring capabilities to organizations that need to access and monitor network traffic in their cloud environments. These solutions target longstanding challenges:

- ✓ The software-defined network fabric of cloud providers lacked native functionality to easily monitor network behavior.
- ✓ Leading network security vendors didn't have compatible solutions within the cloud for monitoring network traffic.
- ✓ Full packet capture network TAPs hadn't yet materialized.

Today, these challenges have diminished significantly. More network monitoring platforms are compatible with leading cloud services, and (more importantly) some cloud providers have created native tools and services that can mirror traffic to their platforms simply and efficiently. AWS VPC Traffic Mirroring permits network traffic to be copied from any compatible system in a virtual private cloud to suitable endpoints such as elastic network interfaces, or ENIs, network and gateway load balancers, etc.

GCP Packet Mirroring captures all traffic and packet data from virtual machines, including payloads and headers. The capture can be configured for both egress and ingress traffic, or one or the other. It's important to note that GCP Packet Mirroring executes on the virtual machine instances, not on the network. Azure's virtual TAP (vTAP) is still in development.

A drawback in using traffic mirroring arises when cloud processing charges are added in. The mirroring service can have a limit to the number of sources (e.g., only 10) per target, which can become prohibitively expensive as compute instances mount with workload levels. The VAF needs to support essentially an unlimited number of traffic sources that can target the visibility node function, which may need to independently scale out to handle the volume of traffic. A further concern with traffic mirroring services is that they can often only send

packets to one tool. The VAF should automatically scale its elements based on traffic volume, not the number of traffic sources or destinations, thereby improving cost efficiency.

IT has typically used cloud-native orchestration tools as much as possible, but if, for instance, the organization has deployments in multiple clouds, they need to bounce from dashboard to dashboard to undertake investigations. Such complexity is the enemy of security. This method doesn't scale, and organizations can miss breaches.

A centralized fabric manager is the overarching VAF management solution that works with various clouds simultaneously. Whether using a combination of multiple, built-in mirroring services or VAF-based virtual TAPs or third-party devices and tunnels, the fabric manager console provides a unified dashboard from a central site with details of all acquired traffic.

Large organizations can potentially establish thousands of virtual private clouds in the IaaS environment, as well as on premises, by making use of transit gateways for the interconnects. This simplifies the network and puts an end to complex peering relationships.

Yet these gateways are complex infrastructure elements that involve transit gateway maximum transmission units, routing tables, and route propagation aspects. The VAF must fully support transit gateways to ensure visibility across interconnected virtual private clouds and on-premises data centers, and forward traffic to tools in any of these environments.



A true VAF can leverage these mirroring capabilities to pull traffic from instances in cloud virtual networks, allowing SecOps teams to perform deep packet inspection, network forensics, and selective packet filtering.

The Need for Unified Traffic Acquisition across the Cloud

As organizations transition to multi-cloud deployments, a reliance on cloud-native technologies and services can become operationally unwieldy due to different types of implementation and controls for tuning and management, and a lack of central monitoring and reporting.

For organizations with critical workloads and data running in numerous cloud platforms, acquiring and analyzing network traffic for operational and security use cases need to be centralized in a single model. This consolidation will help to streamline operational processes and practices and allow a single technology to facilitate traffic acquisition, troubleshooting and performance monitoring, packet capture, threat hunting, network forensics, and reporting.

Trying to manage different solutions for on-premises versus cloud network traffic capture and analysis can lead to significant increases in operational overhead because NetOps and SecOps teams will need to track two or more different tools and dashboards to observe network behavior and events.

By implementing a centralized network visibility and analytics fabric, these teams can leverage the same console and system to track network flows, longer-term behavioral patterns, specific events, and performance data. Consolidating effort into a single product or service that aggregates all the data needed for multiple teams, while providing a unified reporting engine, can greatly enhance a hybrid cloud monitoring strategy.

Traffic Brokering in the Cloud

Today's complex architectures require a new approach to network traffic capture and analysis that encompasses cloud-focused technologies and methods. A critical ability for a unified network visibility and analytics fabric is operating as a "layer," which simplifies remote administration and correlation/aggregation across the entire environment.

Once traffic is acquired, it needs to be brokered to meet various users' requirements for operational and security assurance. Ideally, such a centralized traffic brokering plane will also align with core configuration and orchestration tools to help streamline deployment and maintenance. This centralized plane should incorporate several capabilities, as well as

increasing operational efficiency when managing visibility platforms and applying unified policies where necessary.

Aggregating traffic

The unified VAF process includes aggregation of traffic acquired from various sources in the cloud. Aggregation brings all the acquired flows and packets together so possible issues from asymmetrical routing can be addressed and further processing can be consolidated.

Replicating traffic



Replication is key to forwarding the desired traffic to any number of monitoring and analysis tools in the cloud or on premises.

Tagging traffic in the cloud

Hybrid cloud network visibility solutions should be capable of tagging traffic to identify where it was captured to aid in identification of traffic source(s). Tagging can help in identifying specific network flows, as well as differentiating traffic in a busy network environment. It can also help to reduce costs and operational overhead.

Filtering cloud traffic

Organizations that begin monitoring cloud network traffic are often surprised by the volume of traffic they're seeing. The reason for this high volume is that, in addition to expected system and application traffic, a significant quantity is automatically generated for metadata service lookups and internal cloud communications.

This is completely normal but may present challenges when organizations are seeking to minimize the traffic collected for analysis. Fortunately, leading VAF platforms can identify specific traffic based on a variety of attributes, and then selectively apply filters to differentiate which should be monitored, which should be captured and/or mirrored, and which should be ignored entirely. In order to apply filters efficiently, a VAF must have built-in analytics intelligence and support packet analysis at extremely high speeds. Despite filtering, a massive amount of traffic will still be generated. Use of Data Plane

Development Kit (DPDK) processing methods can substantially increase the processing of traffic, often by more than an order of magnitude.

Delivering cloud traffic

The traffic now needs to be delivered to the monitoring and security tools. Due to the high volumes involved, the cloud VAF needs to be able to balance the load of traffic across multiple instances of each tool based on each instance's capability. Load balancing must also ensure that all packets of a given flow are forwarded to the same tool instance.

Packet forwarding must be done in a way that matches the location and capabilities of each tool, as well as the requirements of the SecOps team. Delivery methods include tunneling to cloud-based, remote, or externally networked tools and directly connecting physical SPAN ports to the tools.



TIP

A VAF's built-in analytics intelligence should include capabilities like packet deduplication, slicing or extraction, and generation of metadata records for forensics analysis tools such as SIEM. Organizations that use cloud-based VoIP services could leverage the VAF to separate this traffic from other data traffic and send it to a voice/video quality analysis tool. More on this topic of advanced, application-based traffic flow later.

Discovering cloud workloads and assets

Leading VAFs offer a critical capability – automated workload and asset discovery. Without a discovery capability, network monitoring solutions will be limited in their ability to locate, acquire, filter, and control traffic flows, and ultimately, to centrally monitor overall performance and security.



TIP

In the physical data center, VAF platforms rely on the direct physical connections from TAPs and SPAN ports, but in the cloud an API-driven solution will likely be necessary (once again underscoring the importance of a platform that is well integrated into the cloud environment) to determine the traffic sources or appropriately deploy virtual traffic agents.

Monitoring cloud traffic with a VAF

To enable comprehensive cloud traffic monitoring, a mature VAF should be able to perform the following basic functions (see Figure 4-2):

1. Leverage cloud provider APIs to discover assets, network security configurations, and cloud security-related events
2. Register the existing environment with the VAF
3. Deploy or leverage workload traffic acquisition methods (containers or virtual machines)
4. Align desired security policies with cloud-native network security options, like AWS Security Groups
5. Aggregate and replicate the acquired traffic
6. Filter traffic based on defined policies for each monitoring and security tool's needs
7. Deliver traffic to each tool type in a coherent and load-appropriate manner
8. Provide continuous monitoring and reporting through a single console

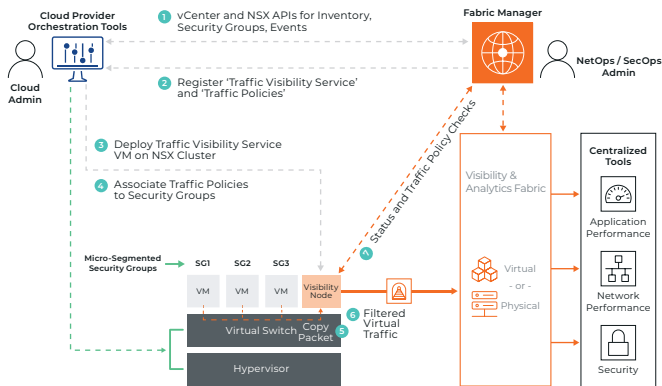


Figure 4-2: Cloud traffic monitoring functions, private cloud example.

Benefits of a VAF

Organizations with a diverse network footprint stand to gain the following distinct benefits from implementing a VAF, particularly in cloud environments:



- ✓ *Unified monitoring and control:* For organizations operating in diverse environments (using multiple cloud services and on-premises data centers), centralizing and coordinating monitoring efforts through a single platform can reduce operational overhead and investment significantly.
- ✓ *Integration with cloud provider APIs:* Good API integration is important for enterprises looking to implement a VAF, and a strong alignment with cloud infrastructure providers means smooth automation that facilitates easy deployment and dynamic operation.
- ✓ *Orchestration and DevOps integration:* With the shift to more-agile and -dynamic DevOps deployment models, and the use of orchestration platforms like Kubernetes, cloud environments are changing more rapidly than ever. VAF solutions can help organizations continuously monitor and keep up with changes in the operating environment.
- ✓ *Reduced complexity and cost:* A unified solution for all traffic monitoring can help to reduce both capital and operational costs over time, especially as IT environments grow more diverse and complex.
- ✓ *Reduced application downtime:* By implementing a robust VAF, SecOps teams gain in-depth insight into traffic patterns and behaviors that should improve planning, troubleshooting, and availability/uptime monitoring processes.



Chapter 5

Threat Detection and Response in the Hybrid Cloud

In this chapter

- Understanding modern cloud attack patterns and workflows
- Exploring cloud visibility and detection
- Assessing cloud security response

The cloud expands the footprint of technology and assets that security teams must discover, monitor, and defend against both insider and external threats. While none of these technologies is inherently higher in risk than those traditionally used on premises, new cloud resources, assets, and workloads expand the attack surface and can become new targets for bad actors.

Modern Cloud Attack Patterns and Workflows



Many attackers today focus on access to and exfiltration of sensitive data in the cloud, and the most sophisticated adversaries often use advanced, malware-based espionage that operates stealthily in pursuit of targets. Attackers also attempt to maintain a persistent presence within cloud networks, escalating privileges as needed and moving laterally within the environment to extract sensitive information to locations under their control.

MITRE ATT&CK Framework for Cloud

The MITRE ATT&CK® framework for cloud is helping us better understand attack patterns and methods used by attackers today. This framework's Cloud Matrix provides valuable insights into cloud-based techniques.

- **Initial Access:** Threat actors find an initial means of gaining access to an organization's assets and/or environment.
- **Execution:** Once initial access has been accomplished, threat actors will entice users to execute malicious code or other commands if they have not been able to execute malicious code themselves.
- **Persistence:** This stage involves setting up backdoors and methods to maintain access on the system or in the environment over time.
- **Privilege Escalation:** In the cloud, privilege escalation is usually tied to unauthorized access to and use of cloud accounts and privileges.
- **Defense Evasion:** This phase includes disabling or modifying cloud firewalls, manipulating cloud workloads, employing unused cloud regions, and manipulating cloud accounts or authentication types.
- **Credential Access:** Accessing cloud accounts without authorization, discovering credentials in files, and abusing cloud metadata APIs for privilege acquisition are common tactics seen in cloud threat scenarios.
- **Discovery:** In a nutshell, the discovery phase is when threat actors enumerate and identify where valuable corporate information or controls exist.
- **Lateral Movement:** Many lateral movement tactics in the cloud may leverage cloud APIs, access tokens, service accounts and privileges, and cloud metadata services.
- **Collection:** Focusing on cloud storage and secrets, such as API keys, is common, and top focal areas for attackers are likely to include cloud storage objects.
- **Exfiltration:** Whether for ransomware, cyber espionage, or a data breach, actors look to exfiltrate collected data from the environment.
- **Impact:** Whether to cause brand damage, hold a company at ransom, or sell off secrets, this is the end goal for the attacker.



Despite the fact that threat actors are executing cloud attacks, much of their activity appears as insider behavior (as it may originate from a compromised end-user workstation or account), and can look legitimate and often fly under the security team's radar. Many enterprise security teams have struggled to keep pace with these newer tactics and techniques. As a result, they are facing pressure to detect and respond to threats more rapidly. This can be difficult when trying to find evidence of privilege escalation, discovery, lateral movement, collection, exfiltration, and other stealthy behavior.



New and different threats facing cloud environments today, such as ransomware as a service (RaaS), necessitate changes in the controls and processes SecOps teams must use to monitor and respond to attacks in software-defined cloud data centers.

- ✓ Increased attacker emphasis on cloud administrators, engineers, and DevOps teams that have privileged access to deployment pipelines and internal services for many cloud environments could be missed without network and event visibility.
- ✓ Lack of network monitoring could cause teams to overlook the use of cloud services as a delivery mechanism for malware and other malicious content. Similarly, without cloud traffic visibility, SecOps teams can't detect, hunt, or investigate attacker behavior effectively.
- ✓ Abuse of overly permissive identity policies in cloud environments can easily be missed when cloud network and security visibility is poor. Most cloud service providers offer minimal tools and services to help pinpoint creation of overly permissive policies.
- ✓ Misconfiguration of cloud environments is an ever-growing attack vector that security teams are struggling to keep up with. To address this challenge, security teams need effective policy to manage cloud configurations across the enterprise.



Additionally, advanced malicious actors are readily using cloud services for major elements of their targeted campaigns. These could include denial-of-service efforts, command and control (C2) systems set up in the cloud, cryptocurrency mining, password cracking systems, and many more.

Cloud Visibility and Detection

To gain full visibility into cloud network environments, it's important to tap into network traffic traversing the cloud fabric.



As an organization's cloud presence and footprint grow, its NetOps and SecOps teams need to provide complete parity with on-premises packet inspection methods and monitoring scope. This means providing deep packet inspection and analysis at all layers based on rich metadata extracted from network flows – both inbound-outbound (North-South) and lateral (East-West) traffic. In the cloud, metadata extraction, inspection, and analysis should be implemented as workloads in traditional virtual machines as well as container pods.

A particularly painful, as in costly, shadow activity is cryptocurrency mining. Fast becoming the top activity for threat actors, it is the fastest and “safest” route to monetization. Crypto mining servers are stealthy, do not generate log files, and do not have command-and-control communications with known bad IP addresses. That is why monitoring network metadata is effective for detecting crypto mining activities.

Cloud detection controls



For most enterprises, prevention and detection tooling consists of a combination of network and endpoint controls, coupled with event management platforms and solutions. In recent years, a newer model of detection and response called extended detection and response (XDR) has emerged. It combines network detection and response (NDR), endpoint detection and response (EDR), and SIEM, as depicted in Figure 5-1. The goal of XDR platforms is to extend NDR and EDR to incorporate deep correlation and analytics and create a more holistic analysis capability than any one of these technologies alone. EDR and SIEM provide additional correlation to what the NDR solution detects, investigates, and responds with.

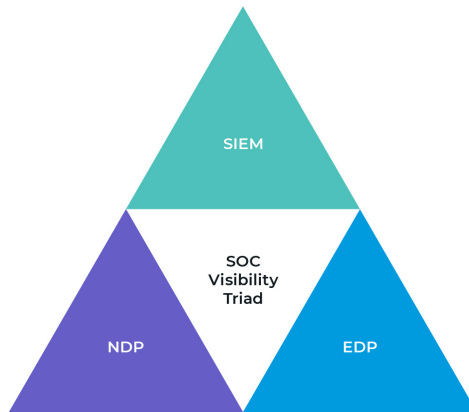


Figure 5-1: The security operations center triad of tools.

Key network cloud detection techniques

For any enterprise-class cloud-centric detection platform, machine learning and artificial intelligence (AI) processing from solution providers can greatly augment detection of new and emerging threats. Some of the common use cases today for cloud detection with machine learning and AI include:

- ✓ **Threat intelligence analysis:** Threat intelligence feeds can be aggregated, analyzed at scale using machine learning engines in the cloud, and processed for likelihood/predictability models.
- ✓ **Security event management:** Machine learning and AI can augment event data processing technology to provide additional intelligence detection and alerting tactics.
- ✓ **Network behavior modeling:** Network flow modeling is a great use case for in-cloud AI processing. There are massive quantities of traffic between systems and the cloud provider backplane (control plane) that could and should be developed into “normal” baselines for monitoring.
- ✓ **Malware detection:** Cloud-native event processing of data and file attributes could likely help with detection of ransomware and other malware variants, particularly those without known signatures.

- ✓ **Data classification and monitoring:** Leveraging known content types and patterns, AI-based cloud analysis engines can classify and tag all data uploaded to and created in the cloud environment against predefined policies, and then monitor it for unauthorized access.

Today, highly advanced network VAF and monitoring solutions can greatly enhance overall cloud visibility.



When building a cloud-focused detection and response program, it's a good idea to track metrics over time because they can help improve your security posture overall. Start with some easy-to-capture metrics:

- ✓ Number of incidents/events logged
- ✓ Number of incidents/events per priority, impact, and urgency
- ✓ Number of incidents/events per type and category
- ✓ Average time to achieve incident resolution
- ✓ Average cost per incident

Cloud Security Response

A mature cloud security program should also include triage, hunting, and investigation elements of a security operations function. By collecting historical network metadata, organizations can rapidly analyze events in context within the cloud environment, and also look for indicators of compromise for follow-up during investigations.



Today, cloud engineering and security teams are actively looking to automate tedious and repetitive processes that consume too much of highly skilled analysts' time and may provide little value in investigations. Activities that many teams consider for automation include:

- ✓ Identifying and correlating alerts
- ✓ Identifying and suppressing false positives

- ✓ Performing initial investigations and threat hunting
- ✓ Opening and updating incident tickets/cases
- ✓ Producing reports and metrics

Building a hybrid detection and response framework

To build a hybrid cloud security model, operations teams need versatile tools and processes that work well in on-premises networks and cloud environments.

In addition, enabling SSL/TLS inspection for both cloud and on-premises traffic is critical for modern network security. To adequately monitor encrypted traffic, network monitoring solutions must be able to identify the encryption in place, then forward this encrypted traffic to network monitoring devices specifically slated for rapid decryption and inspection of the traffic.

By developing on-premises and cloud network security intelligence through monitoring and deep packet analytics, organizations can build a more comprehensive and sustainable network security strategy.

Chapter 6

Examining Common Network Visibility, Analytics, and Security Use Cases

In this chapter

- Improving network and cloud operations efficiency
- Protecting resources and sensitive user data while meeting compliance mandates
- Eliminating tool sprawl
- Building a hybrid cloud SOC using a VAF
- Establishing a zero trust security model

As the threat landscape evolves in today's networks, information security teams are scrambling to keep up. As attackers use new and stealthier methods to infiltrate organizations and steal data, the complexity of most networks makes it easier than ever for them to send malicious traffic in and out without being noticed.

Attackers are also going after a wider variety of attack surfaces, making monitoring more difficult. For example, attackers are taking advantage of social media to infiltrate devices used by today's workforce to access email and other work applications. They're also exploiting enterprise website vulnerabilities by means of SQL injection, cross-site scripting, and other common attack methods. Many networks are running a variety of traffic beyond just web traffic, including VoIP,

video, and audio traffic, for example, which expands the attack surface and adds complexity.

Network and security teams can't even keep pace with the growing size of their networks, let alone with all these new attack surface areas to monitor and protect. As a result, current deployment and operation of security monitoring infrastructures today leave significant gaps in coverage. Teams are wondering: Are we seeing all the traffic? What threats are passing us by? Do we have remote segments that we're completely missing? How do we extract more information from our network traffic monitoring to use for correlation? How do we see into encrypted traffic to make sure it's not carrying a malicious payload, exposing sensitive information, or following remote commands?



On the other hand, operations teams need to ensure that advanced monitoring functions don't impede the flow of business by adding even small amounts of latency. In fact, network visibility tools should help improve the performance and management of network operations.

Improving NetOps and CloudOps Efficiency



By implementing advanced network visibility and analytics, organizations can improve visibility into network traffic as well as reduce operational overhead and time-consuming tasks such as troubleshooting. By observing traffic events and patterns across the entire network environment (both on premises and in the cloud), they can determine and report key metrics such as mean time to respond/resolve/recovery (MTTR), a flexible and multifaceted metric that is useful for tracking and improving network events and activities. This operational metric can represent any of the following:

- ✓ Mean time to respond: Average time to fully engage and initiate a response process based on specific events
- ✓ Mean time to resolve: Average time to fully resolve an event

- ✓ Mean time to recovery: Average time to respond to and resolve an event or incident (from time of detection through full resolution)

By tracking network and application flows and collecting metadata on a large scale, organizations can track metrics more effectively and benefit from these metrics over time, which can also help reduce operational time spent on troubleshooting by understanding common behaviors and patterns in the network environment.

Automated scalability is key. It has been stated that cloud is not about building software — it’s about operating it at scale. “Infinite” scalability is at the heart of public cloud’s value, but it must be viable. With potentially tens of thousands of workloads constantly adding new or deleting existing instances and moving them within and across virtual private networks, automation is paramount. On top of this, dozens of tools and associated traffic forwarding must be managed in real time and in different environments.

A VAF can seamlessly interoperate with automation and orchestration management suites, such as AWS CloudFormation and Azure Resource Manager, or leverage open-source solutions such as Terraform or Ansible. Virtual traffic mirroring services, virtual TAPs, and visibility nodes can be automatically instantiated, configured, and monitored. This is done through a single-pane-of-glass fabric manager with visualization across any hybrid network.

Protecting Resources and Sensitive User Data While Meeting Compliance Mandates



One prevalent use case that most organizations will want to explore is the use of network flow and metadata to identify data exfiltration activity. Intelligent network monitoring platforms can flag suspected data exfiltration based on type and quantity of data flows, DNS records, and lookups associated with suspicious activity. These can indicate data “tunneling” or covert channel use, and reveal assets involved in communication patterns (certain network subnets and assets can be

categorized as sensitive or as having access to sensitive data). Figure 6-1 shows an example from a SIEM platform of suspicious DNS activity that may indicate DNS is being used as a tunnel for data exfiltration.

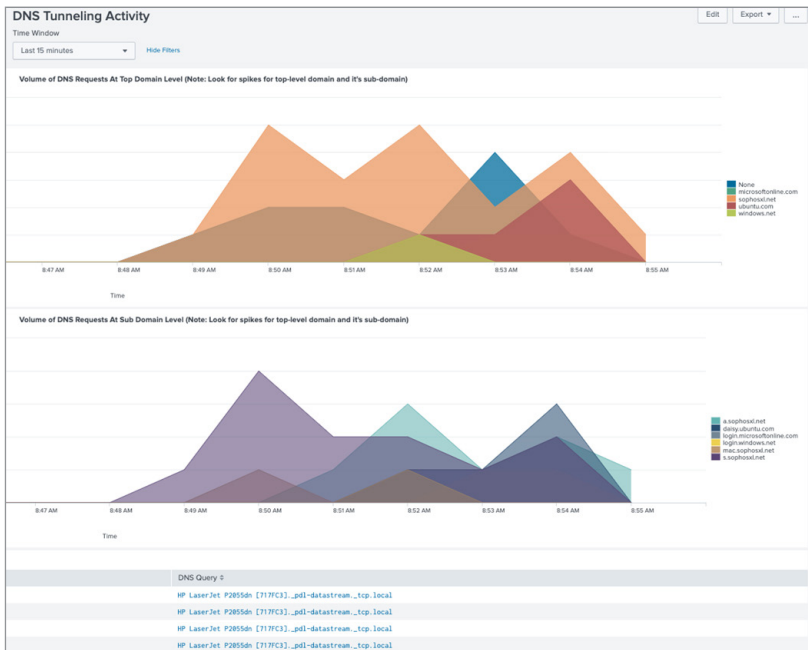


Figure 6-1: SIEM display of DNS activity driven by application metadata.



Another common use case for network metadata and application flow monitoring is checking SSL/TLS certificate use, along with the cryptographic ciphers associated with the handshakes invoked. This is an excellent way for security analysts to periodically review the types of certificates deployed and used in the environment, and to report on weak ciphers, expired certificates, and self-signed certificates that haven't been updated or replaced. Figure 6-2 shows application metadata in a SIEM dashboard that indicates expired TLS certificates detected in the environment.

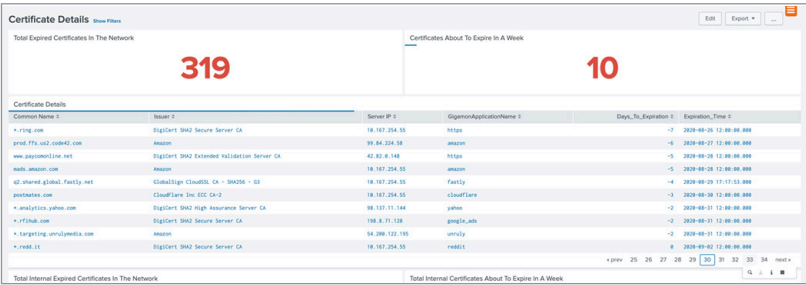


Figure 6-2: SIEM display of expired certificates driven by application metadata.

A third way that application metadata can improve monitoring in the security operations center (SOC) is flagging suspicious and unauthorized external remote connections, as shown in a SIEM console in Figure 6-3.

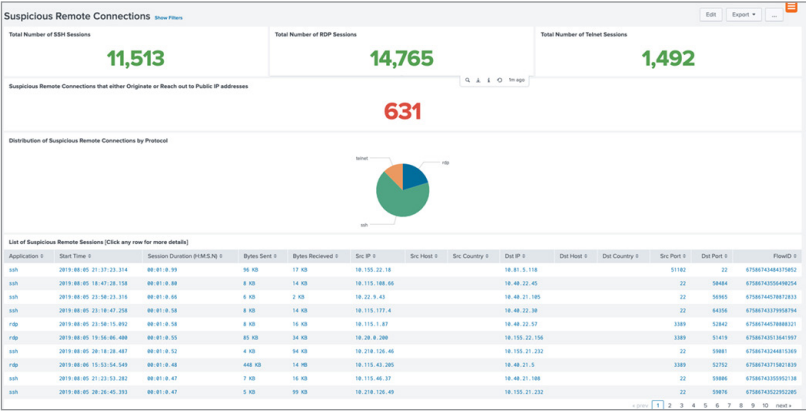


Figure 6-3: SIEM display of suspicious connections driven by application metadata.



Identifying and remediating security breaches and vulnerabilities are primary goals for all SOC teams. Network analytics platforms can play a significant role in this effort by helping teams find and eliminate rogue apps and shadow IT operations, which are often vulnerable targets for attackers.

There are numerous additional benefits and use cases for intelligent network security monitoring platforms that extract and analyze metadata:

- ✓ Helping heavily regulated industries like financial services and healthcare meet compliance requirements
- ✓ Obfuscating internal compute source IDs, account numbers, birth dates, credit card numbers, personally identifiable information (PII), etc.
- ✓ Detecting rogue DNS and DHCP servers
- ✓ Exposing nefarious communications hidden within encrypted flows

Deeper network visibility and improved network threat intelligence can provide a multitude of benefits to SOC teams. Then there is the need to comply with privacy regulations or internal confidentiality policies regarding private and sensitive data, which can even occur at network layers of the traffic (versus the content).

Servers use countless naming structures that may need to be opaque to external environments or monitoring tools; otherwise, malicious actors can obtain keen insights on where resources reside and how they are structured, even identifying potential applications and their known vulnerabilities. Tools do not necessarily need to know the actual ID of the traffic source.



Of course, the data carried in certain traffic types may contain personal, private data or even sensitive or proprietary information. To ensure obfuscation of such data or information, a VAF can employ header transformation to modify fields within the network layer, such as MAC address, IP address, and VLAN ID.

Eliminating Tool Sprawl and Enhancing Tool Efficiency



The first key element of an advanced monitoring solution is the ability to identify and eliminate superfluous traffic that can overload traditional network monitoring and security tools. An intelligent network monitoring fabric should be capable of parsing all traffic and filtering it according to specified policies (any identifiers, attributes, and/or patterns) to extract only the traffic types and application flows that an organization wants to analyze. Figure 6-4 depicts advanced filtering and extraction capacity, which feeds only specified traffic and data types to the appropriate monitoring and analysis tools – and ignores the rest.

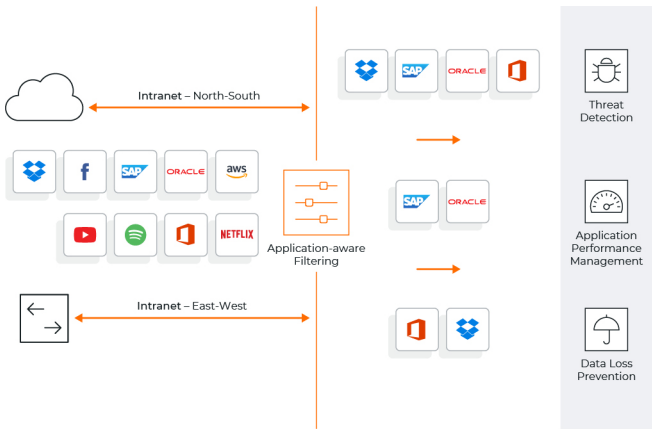


Figure 6-4: Extracting application flows of interest.

Monitoring and traffic control capabilities can help to optimize traffic flows to tools to ensure their availability and prevent saturation within the network environment. This can often mean filtering out high-volume, low-risk traffic that doesn't need in-depth analysis or treatment. It's important that network monitoring tools provide accurate packet and metadata inspection and analysis at high speeds to properly filter the appropriate flows.



By focusing on only the specified, relevant flows detected in acquired and analyzed traffic, organizations can optimize what is sent to security tools while preventing excessive load on the bandwidth and processing (CPU and memory) capacity of tools in use. For cloud-based environments, controlling flows can lower costs due to processing capacity requirements and/or traffic export restrictions, whereas in on-premises environments, the benefits will be more evident in preservation of processing capacity and optimization of security tool productivity.

Building a Hybrid Cloud SOC

Many types of security processes and practices will need to be updated as an organization progresses further into cloud deployments, whether entirely native or hybrid. Several primary factors are driving this need to update and change technologies and processes for a cloud-focused SOC:

- ✓ Additional services and assets: First, the cloud introduces a broader technology footprint, incorporating new assets for security teams to discover, monitor, and defend.
- ✓ Lack of visibility: Enterprises have been notoriously challenged with a lack of visibility into cloud infrastructure and modern application stacks being built and run in the cloud, as these new environments require new technologies and skillsets. Increasingly, they are deployed across multiple and hybrid clouds.
- ✓ More data: As sources of data increase, so does the burden of processing data into practical intelligence that can be applied to identifying and remediating threats.



To ensure that organizations are prepared for security detection and response in the cloud, SOC teams will need to adapt or replace some of their network monitoring tools and their processes for detecting and responding to events and incidents. With the availability of cloud APIs and cloud-compatible network visibility fabrics and monitoring platforms, organizations can now send network traffic to monitoring tools in the cloud or on premises from anywhere. Leveraging uniform

tools for traffic interception, filtering, and monitoring can greatly facilitate detection and response across the entire network environment, while streamlining network and security operations.



By collecting and analyzing network metadata from traffic flows, SOC teams can also uncover network blind spots and potentially malicious traffic patterns by enriching the data with threat intelligence feeds and research from leading providers. This information can highlight DNS and WHOIS records of interest, patterns of command and control (C2) behavior, and more. This enriched network data can also aid in streamlining investigation workflows by pinpointing unusual traffic in busy networks and highlighting events and flows of interest quickly and efficiently. This assistance can eliminate the need for extended (and often fruitless) digging by security analysts who may not know exactly what they're looking for.

Establishing a Zero Trust Security Model



Organizations have begun thinking about network design and topology from a data-centric viewpoint. Rather than model network and security architecture around the classic ideas of an external perimeter, they are shifting toward a model based on workloads, storage, and data in a hybrid cloud, with an emphasis on the concept of “zero trust.” Zero trust makes data the central focus of all isolation and protection tactics and considers all assets in an IT operating environment to be potentially untrusted by default until network traffic and behavior are validated and approved.

With the move to software-defined infrastructure, enterprise teams can now perform more-effective (and continuous) discovery of assets within the environment, which greatly facilitates building a zero trust model that requires automatic discovery of network topology.



Another aspect of zero trust is control over access and permissions, so strong role-based access control (RBAC) and authentication management capabilities are also important.

The zero trust approach does not involve eliminating the perimeter; instead, it leverages network micro-segmentation to move the perimeter as close as possible to privileged apps and protected surface areas. It includes continuous assessment of identity relationships and privileges to help detect and respond to active threats.

A zero trust monitoring architecture should include a policy engine that defines the types of interactions that are permitted, a data plane that includes workload agents and gateway platforms to inspect network traffic, and a monitoring plane that comprises a network visibility and analytics fabric and network monitoring, security, and tracking tools, as shown in Figure 6-5.

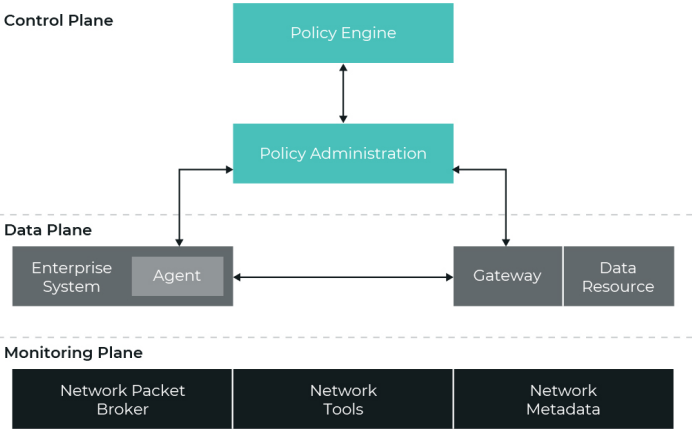


Figure 6-5: Zero trust network monitoring architecture.

Chapter 7

Selecting the Right Cloud VAF and NDR Vendor

In this chapter

- Determining what to look for when selecting a VAF and/or NDR solution that can meet your company's needs
- Examining core features that every mature VAF and/or NDR solution should have

For enterprises seeking to implement a comprehensive cloud visibility and analytics fabric (VAF) and network detection and response (NDR) solution, choosing the right provider is of paramount importance. Because the capabilities of the VAF and NDR are critical elements of a network security and operations strategy, choosing the right solution requires some due diligence. This process involves evaluating providers and ensuring their solutions include important features.

Any scalable and robust VAF needs to be capable of acquiring network traffic from a wide variety of sources and supporting a range of acquisition types, both physical and virtual. The primary type of traffic acquisition for on-premises environments is physical, and the most common physical option for traffic mirroring and processing seen in enterprises today is network TAPs.

Many network platforms and core data center networking technologies are still comprised of physical appliances and cabling, so network analytics should be capable of accommodating physical networks that range in speed and traffic capacity. As an alternative to network TAPs, network switch mirror or SPAN ports may be a necessary and acceptable option in some circumstances.

Cloud and virtual networks also have a number of traffic acquisition options available. First, virtual TAP agents can be deployed on each workload’s virtual machine or Pod to acquire and forward traffic to processing nodes within the VAF and on to tools. Second, agentless virtual machines can be added to each hypervisor to mirror traffic running on the underlying virtual switch. Third, some cloud infrastructure providers now offer traffic mirroring capabilities to copy packets to and from targeted workloads. And finally, “TAP control” functions can be deployed on each host to control the available infrastructure’s mirroring services to copy and forward the traffic.

With the ability to access all traffic comes the opportunity to detect threats operating in that traffic. That is why an integrated NDR is the fastest route to adding detection and response capability to any VAF deployment. NDRs that are linked to VAFs allow for instant and complete visibility into all traffic, encrypted or otherwise, whether running on-prem or in the cloud.

The Need for Complete Coverage



Any enterprise VAF solution should offer compatibility and integration with numerous environments. On-premises network monitoring of workloads should be possible via virtual and physical TAPs that feed virtual and physical visibility nodes, which in turn feed virtual and physical monitoring and security tools.

Flexible cloud implementation is also important today as more organizations build extensive IaaS cloud infrastructure. Some leading cloud service providers may offer packet mirroring services as a way to direct network traffic to a VAF, but all IaaS cloud environments should facilitate the deployment of virtual tapping agents for acquiring traffic from cloud workloads, or for directly managing the mirroring services.

Any leading VAF platform should accommodate all of these approaches. Figure 7-1 demonstrates a flexible cloud traffic acquisition strategy that includes workload agents, traffic mirroring, and centralized management and control of all tools and policies.

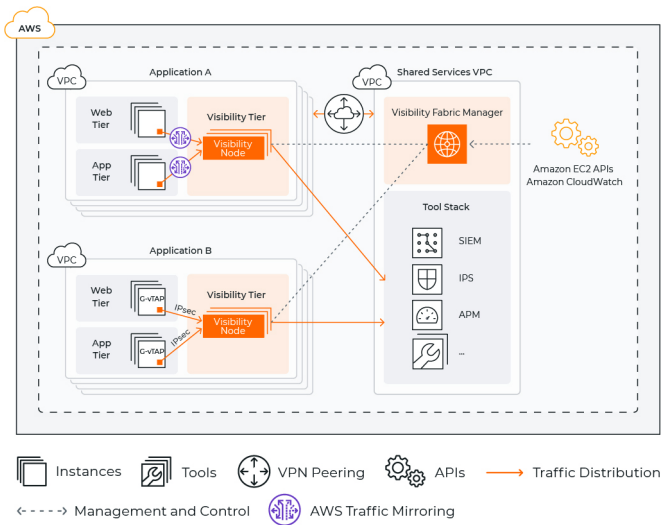


Figure 7-1: Cloud traffic acquisition methods.



In addition to deployment flexibility, leading VAF and NDR solutions should offer partnerships and compatibility with other leading solutions and services, ranging from network firewalls and software-defined networking infrastructure to network intrusion detection and SIEM solutions, security orchestration, automation and response (SOAR) systems, observability solutions, network performance management tools and products, and many others.

Cloud Integration and Compatibility

To cover both on-premises and cloud environments in a hybrid model, it's vital that VAF and NDR solutions be compatible with all major IaaS/PaaS cloud platforms, namely Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. As more organizations shift to a multi-cloud architecture, compatibility across clouds grows in importance. Likewise, VAF platforms ideally should extend coverage to private cloud environments built on all major platforms, including Kubernetes (which can be private or public), AWS Outpost, Azure Stack, Nutanix, OpenStack, and VMware.

Private cloud infrastructure that makes extensive use of virtualization technology presents unique challenges that VAF solutions should address. In scalable virtual infrastructure, workloads can be added or migrate among different hypervisors to scale applications as needed, improve performance, optimize overall traffic routing, and load balance across hypervisor clusters.



Network monitoring for workloads should be dynamic and integrate with migration technologies such as VMware vMotion or Nutanix Live Migration to provide a seamless transition in network monitoring by automatically identifying new or relocated workloads and dynamically ensuring their visibility. Figure 7-2 illustrates how each hypervisor includes a dedicated visibility node that monitors virtual machines hosted on them. Integration with solutions like vMotion can trigger automation of changes to the VAF within the private cloud environment.

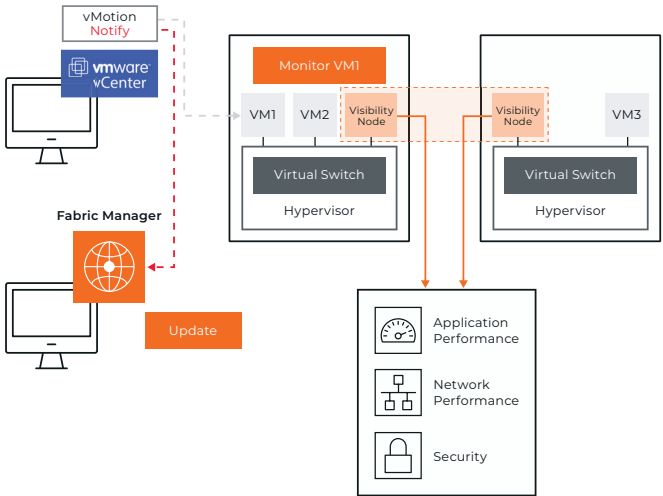


Figure 7-2: A VAF with dynamic workload visibility

Not All Cloud Network Monitoring Is Created Equal

When evaluating VAF and NDR platforms, it's important to keep in mind that this market is highly evolved, with a wide variety of solutions and products available. Some are focused entirely on cloud-based implementations, others are only available in on-premises models, and many are limited to very particular types of implementations and integrated with specific cloud solution providers.

Enterprises with rapidly changing infrastructure, applications, and operational practices should look for mature network monitoring platforms that provide maximum flexibility and capabilities to parse and control traffic flows. These platforms should place some emphasis on filtering capabilities and metadata generation and correlation, as well as value-added threat intelligence and security intelligence.

Core Features Every Mature VAF Should Have

DON'T FORGET



One of the foundational capabilities of any VAF solution is SSL/TLS interception. In the past, network TAPs usually just forwarded encrypted traffic to specialized decryption platforms before handing the packets off to analysis tools, but a modern VAF solution should be capable of SSL/TLS decryption directly onboard, as shown in Figure 7-3. With the advent of perfect forward secrecy encryption and TLS 1.3, typical network taps are not sufficient. Instead, an inline bypass capability is necessary for the SSL/TLS decryption function to intercept the communication.




Figure 7-3: Decrypting SSL/TLS traffic to send to numerous applications and security tools.

Lastly, do ensure the VAF tool is able to decrypt TLS 1.3, which is rapidly gaining adoption.



In most enterprise networks, TAPs collect packets from multiple points along a network path. As a result, duplicate copies are sent to your tools for analysis. Duplicates can also be caused by inter-VLAN communication, incorrect switch configuration, or unavoidable SPAN/mirror port configurations. A mature VAF solution should be able to detect and remove duplicate packets before sending the remainder to analysis tools. This process is illustrated in Figure 7-4, where a deduplication function is embedded in a VAF node.

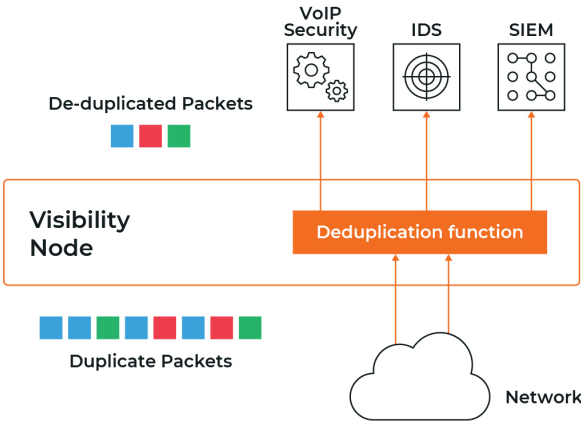


Figure 7-4: Packet deduplication in a VAF.

Since VAF solutions are adept at parsing and analyzing network packets at high speeds, they should be fully capable of “carving” or “slicing” specific packet flows and/or payloads from the traffic that are intended for monitoring, performance tracking, and security analysis tools. Flow slicing can greatly improve monitoring efficiency by analyzing a set number of packets and then slicing the rest of a flow’s packets or payloads that aren’t needed. As shown in Figure 7-5, network VAF nodes apply highly tuned content filters to traffic, forwarding only a subset to select monitoring tools.

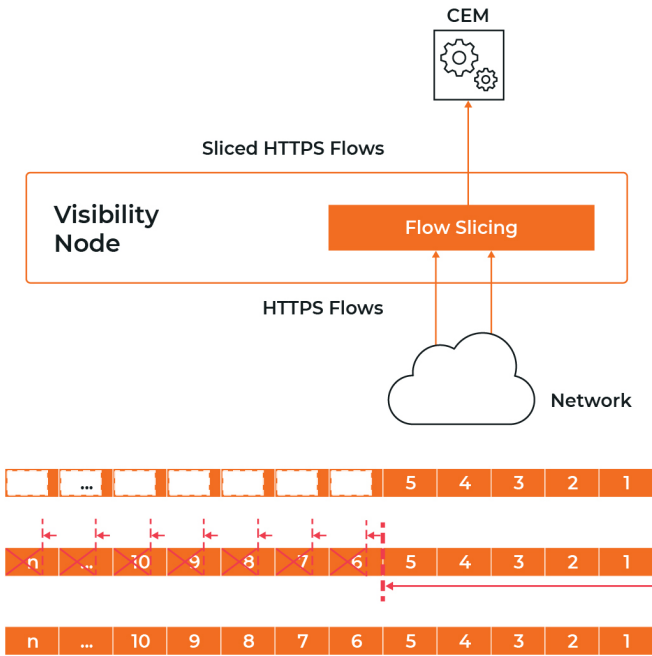


Figure 7-5: Removing unnecessary packets or payloads from flows.

Along the same lines, VAF solutions should support flexible traffic filtering and forwarding based on configured policies. In essence, the VAF should be capable of acquiring, aggregating, replicating, filtering, and forwarding network traffic to tools of any sort. The granularity of traffic parsing and selective packet handling is a differentiator of a robust VAF platform. Figure 7-6 shows how VAF nodes can capably control traffic flow to different security and network operations tools.

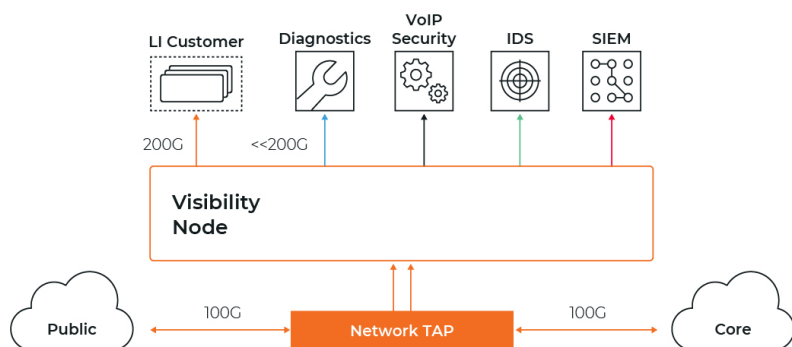


Figure 7-6: Selective traffic parsing and forwarding to tools.

A mature network visibility fabric should have deep and comprehensive layer 7 intelligence built in, as relying on superficial protocol and application indicators like network ports doesn't provide granular operational and security filtering and control. Packet data should also be matched against analysis from third-party researchers and kept up to date.

For example, the platform should send only relevant application traffic to security, performance monitoring, or data loss prevention (DLP) tools. Threat detection tools would receive all application traffic, while DLP tools would receive only email, cloud communication, and file transfer data. This traffic delineation helps increase the ROI of many tools by limiting the traffic they need to process. It also allows existing tools to better monitor and secure the network by analyzing traffic in more depth without wasting processing cycles on irrelevant traffic.

For all traffic that is tunneled in some way, particularly TLS and VPN traffic, monitoring solutions should be capable of tracking tunnel initiation and termination details to look for unusual destinations or sources, as well as patterns of connectivity that may indicate unusual behavior.



Layer 7 application and protocol awareness at wire speed should also enable VAF solutions to perform application-specific visualization and filtering, as shown in Figure 7-7.

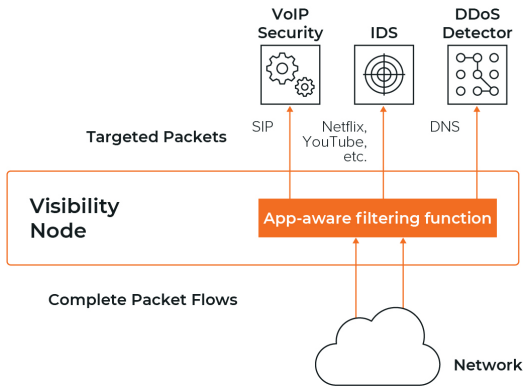


Figure 7-7: Dropping or forwarding specific applications.



Given the current focus on data privacy and control of access to personal information, regulated information like financial records, and intellectual property, a modern VAF solution should be capable of masking specific content fields within packets. In essence, a VAF should offer a “search and replace” function for all packets of particular types, obfuscating sensitive fields and details before forwarding to analysis tools (Figure 7-8).

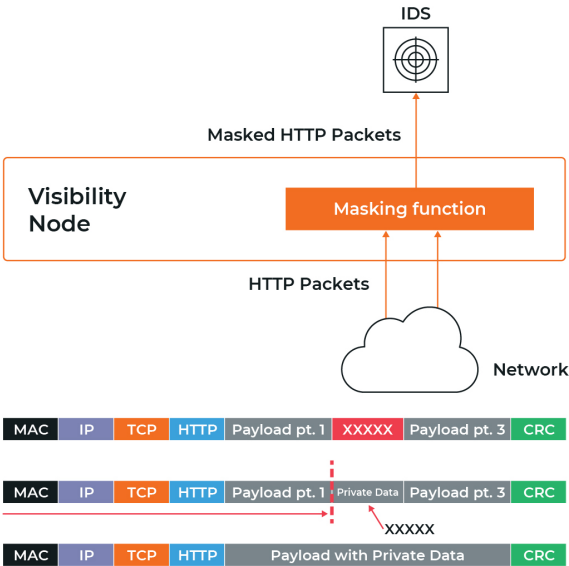


Figure 7-8: Obfuscating sensitive and private information.

Network metadata based on network flows and application traffic patterns can be useful in many types of monitoring and analysis workflows, including performance monitoring and troubleshooting, as well as network behavioral monitoring for application tuning and suspicious traffic pattern detection. As shown in Figure 7-9, generation of application-aware metadata for all traffic traversing the VAF in standardized formats like NetFlow, Internet Protocol Flow Information Export (IPFIX), and Common Event Format (CEF) should be possible before forwarding to monitoring tools.

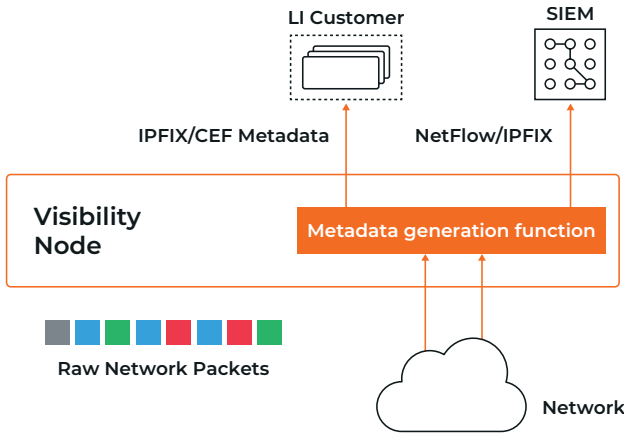


Figure 7-9: Generating and forwarding flow and application metadata.



Today's exponential volume increase in voice, video, and data traffic is significant when it comes to monitoring and managing networks. As more organizations transition to 5G wireless networking and control/user plane separation (CUPS), the need to process and control specific traffic types in large environments grows and further complicates this ongoing challenge.

In most subscriber services, like VoIP and 4G and 5G, there are distinct traffic types for actual voice (user) traffic as well as control plane traffic related to the services themselves. Detecting these distinct traffic types and patterns is important so that network operations teams can prioritize and optimize performance configuration within the network. In addition, security teams may seek to combine and correlate both traffic types to better analyze end-user behavior and activity when needed for investigations. This correlation has been notoriously challenging in high-speed environments, so a network analytics fabric should be capable of mapping user traffic/voice data flows to specific teams, aggregating traffic to develop and report on behavioral patterns, and tracking sessions to help in developing performance metrics.

Subscriber intelligence can also be derived from the generation of metadata. Metadata attributes may include device or

equipment ID, subscriber ID, network identifiers, and more. They can be aggregated and correlated at scale to provider filtering, whitelisting, session sampling, and load balancing.

In addition to operational and security monitoring, VAF platforms are invaluable in helping to improve traffic efficiency and performance. They do this by using selective VoIP- and subscriber-aware traffic handling. In the case of VoIP traffic, selectively handling the signaling traffic for management (commonly SIP) versus the call data (RTP) can help to prioritize and filter traffic based on user ID, user phone identifiers, and specific traffic flows. For mobile carrier traffic leveraging the GPRS Tunneling Protocol (GTP), a mature VAF platform can help to unify and correlate control plane protocol (GTP-c or HTTP/2) and user plane protocol (GTP-u) traffic that carries subscriber application traffic from the subscriber device to the Internet. GTP correlates user and data plane traffic, after which both can be coherently forwarded to various tools for analysis (maintaining session stickiness towards each instance of a tool), as shown in Figure 7-10.

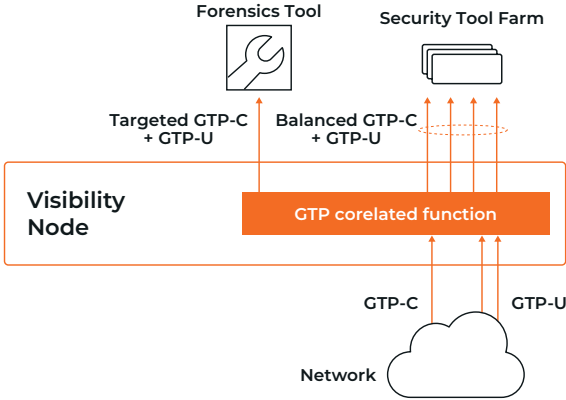


Figure 7-10: Subscriber-aware forwarding of mobile core user traffic.



Note: GTP-c is the control plane protocol for 3G and 4G mobile core networks, whereas HTTP/2 is the control plane protocol for 5G core networks – but the same principles apply.

environment), organizations will continue to face limitations in obtaining centralized and comprehensive visibility into network traffic and behavior. Figure 7-11 shows a hybrid network that includes a deployment in a public cloud and an on-premises data center with both physical and private cloud infrastructures all covered with a unified VAF that brings relevant data and traffic back to a central monitoring plane.

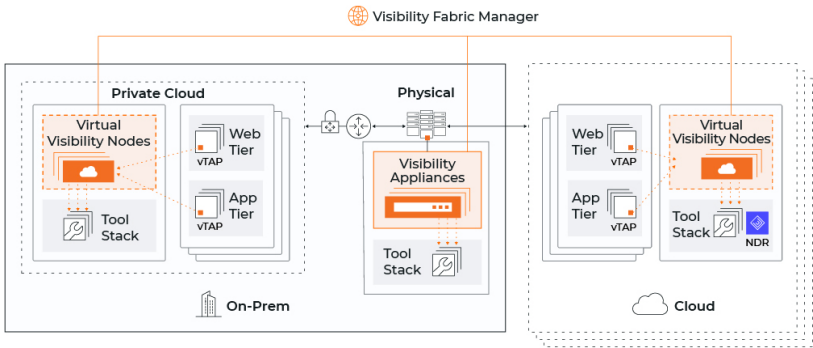


Figure 7-11: Enabling visibility across hybrid cloud infrastructures.

Deep integration with all types of network environments enables a VAF to facilitate automation and orchestration with a high degree of policy-oriented specificity, which likely accelerates creation of both detection and response workflows and playbooks.

VAFs are synergistic to cloud vendors' native services and provide a host of complementary enhanced network packet brokering capabilities. Table 7-1 shows the key features that organizations can take advantage of as they embrace the cloud.

Function	VAF	AWS	Azure
Traffic Acquisition (basic)	Y	Y	Y
Filter / Eliminate Traffic	Y	Y	Y
Traffic Forwarding	Y	Y	Y
NetFlow v5 Generation	Y	Y	Y
Traffic Acquisition (Advanced); Prefiltering, IPsec	Y	N	N
Container Traffic Visibility	Y	N	N
Automatic Target Selection with Traffic Forwarding and Filtering	Y	N	N
Single Management for Hybrid / Multi-Cloud	Y	N	N
Packet Deduplication	Y	N	N
Packet Slicing	Y	N	N
Header Transformation and Data Masking	Y	N	N
Application-layer Filtering	Y	N	N
NetFlow v9 and IPFIX	Y	N	N
Application Metadata	Y	N	N

Table 7-1: Key traffic handling capabilities of a VAF vs. public cloud services.

Core Features Every Mature NDR Should Have

Leading enterprise platforms ideally employ machine learning for deeper/better threat analysis and detection across a range of disparate event types at scale. This capability can help reduce false positives, discover new types of assets and services in both on-premises and cloud environments, and detect subtle behavioral patterns of network traffic that likely indicate malicious activities and tactics, techniques, and procedures (TTPs). Building a mature artificial intelligence function, including machine learning, is not trivial; it requires a huge computing stack coupled with multiple input sources.

Any leading provider of network security intelligence should definitely apply AI and machine learning capabilities to large quantities of network and security event and behavioral data.

By investing in a centralized VAF solution, security and operations teams should be able to take advantage of a dedicated vendor research team that feeds observed attacks and threat data into sophisticated machine learning algorithms, distilling the outputs for dissemination to customers. In view of the increasingly sophisticated and stealthy threats facing organizations today, current detection models based on real-world patterns and behaviors are more critical than ever before.

An NDR should facilitate incident response capabilities by providing SOC teams with tools and visibility into historical network data that allow them to identify adversary activity across the MITRE ATT&CK framework.

With attacker dwell times lasting more than 280 days, mature NDR solutions should provide 365 days of extended historical visibility into network activity.

In today's threat landscape, it's critical that organizations arm their security teams with an NDR solution that matches - and beats - adversaries' sophisticated techniques. It should give SOC teams the tools necessary to not only detect a potential threat but also to respond to that threat when time is of the essence. This gives SOC teams a level playing field to defend their organization with more time, data, and insight into an attacker's behavior.

Glossary

application metadata: Application data extracted from network packet data that summarizes application traffic at a higher, abstracted level. This information provides much greater visibility into how applications are performing, behaving, and being used across the network.

endpoint detection and response (EDR): Endpoint security tools and monitoring, usually delivered through an agent.

hybrid cloud: A combined cloud environment consisting of on-premises private cloud infrastructure and one or more public cloud services.

infrastructure as a service (IaaS): Full virtual machines and deeper networking and identity controls in a cloud-based delivery model.

infrastructure as code (IaC): Managing and provisioning of infrastructure through code instead of through manual processes.

mean time to respond/resolve/recovery (MTTR): An operational metric tracking mean time to respond to, recover from, or resolve incidents.

network detection and response (NDR): An enterprise platform that offers network traffic visualization, monitoring, and threat detection and response capabilities.

network flows: A set of network traffic attributes and characteristics (source and destination ports and addresses, as well as locations) that define network communications.

software as a service (SaaS): Software applications delivered in a cloud-based format.

software-defined data center (SDDC): Core data center technologies like virtual machines, networking, and storage in a software-based format.

SPAN port: A designated port on an enterprise switch that is set to copy traffic from network ports on the switch.

test access point (TAP): A physical or virtual device for copying traffic from a network for the purpose of monitoring, analytics, and/or troubleshooting, and an important element of a VAF.

virtual network (VNet): An isolated virtual network in Microsoft Azure.

virtual private cloud (VPC): An isolated virtual network in public cloud provider environments like AWS, GCP, and Azure (see VNet).

visibility and analytics fabric (VAF): An arrangement of one or more virtual or physical devices for the purpose of acquiring, aggregating, replicating, filtering, optimizing, transforming, and/or forwarding traffic to network performance monitoring and security analytics tools, ensuring complete and optimal visibility.

zero trust: A concept that defines how assets can communicate in a network environment without trusting each other by default. Zero trust can include network perimeter controls, identity controls, and traffic flows to define a comprehensive network access control strategy.



Deep Observability from Core to the Cloud

We harness actionable network-level intelligence to amplify the power of your observability tools so you can assure security and compliance governance, deliver complete performance management, lower operational overhead, and realize the full transformational promise of the cloud.

READY TO LEARN MORE, VISIT
gigamon.com



Discover how network visibility and analytics enhance security operations by enabling in-depth monitoring and security response for physical and virtual data centers.

Security and network architecture and controls are rapidly changing today, with expansion into cloud data centers and containers. Simultaneously, there's more need than ever before for deep network packet metadata analysis in high-speed, high-volume network environments. Advanced network visibility and analytics offers security and operations teams many capabilities.

- **Updating core operations functions and roles** — explore how NetOps, SecOps, and CloudOps are changing and converging today
- **Exploring network visibility in the hybrid cloud** — learn about key network and security visibility concepts and controls
- **Understanding the network visibility and analytics fabric** — examine what kinds of controls and capabilities an enterprise network visibility and analytics platform should have
- **Evolving threat detection and response** — understand how network visibility and analytics fits into the rapidly changing threat landscape facing modern network environments
- **Exploring network visibility and analytics use cases and vendor criteria** — know what to look for when evaluating network visibility and analytics solutions

About the Author

Dave Shackleford is a faculty member at IANS, CEO of Voodoo Security, and a SANS senior instructor and course author. He has consulted with hundreds of organizations on security, compliance, and network architecture. Previously, he served as CTO for IANS, CSO for Configuresoft, and CTO for the Center for Internet Security, and as a security architect, analyst, and manager for several Fortune 500 companies.



CYBEREDGE
PRESS

Not for resale

ISBN 978-1-948939-31-7



9 781948 939317 >