Security Reference Architectures for Public Clouds Using CloudGuard Network Security

Guide for a Successful Lift and Shift Secure Migration Model for Microsoft Azure, Amazon Web Services, and Google Cloud Platform

CHECK POINT™

ABSTRACT

This white paper outlines use cases, architecture diagrams, and a Zero Trust approach that will allow organizations to build the best strategy for a public cloud data center. CloudGuard Network Security will be used to design the strategy, according to the organization's business needs, within a variety of cloud service providers.

AUDIENCE

The desire to transition from a hardware-centric to an application-centric network construct is driving more and more organizations to embrace the cloud as part of their IT strategy. As a result, businesses are rapidly adopting cloud-based solutions to virtualize their data centers, as well as extending applications and data to public cloud environments.

This white paper aims to provide the reader with reference architectures using different technical examples taken from Microsoft Azure, Amazon Web Services, the Google Cloud Platform, and Check Point Software Technologies, as well as from a variety of technical blogs. The information presented in this paper is intended to educate and enable security and networking engineers, solution architects, and designers who would like to integrate public cloud IaaS solutions and Check Point technology for advanced security. To get the most from this paper, the reader should be well versed in cloud computing, network and security design, as well as Zero Trust methodologies.

TABLE OF CONTENTS

INTRODUCTION

According to Gartner Forecasts for Worldwide Security and Risk Management Spending[1], in 2020 investments in cloud security grew 33.3% versus 2019. As more and more organizations are convinced that cloud transformation will lead to greater business opportunities and operational agility, they should therefore become aware of the cyber security implications within the process. The transformation should not be seen as a 1-to-1 shift, which only considers the traditional approach; rather it should be aligned with the business strategy and risk appetite.

Under this approach, organizations should understand three different migration models relevant for cloud migration, taking cyber security as the main driver:

- **Rehost (lift-and-shift) –** The organization migrates their workloads as-is, with no refactoring, or rebuilding, and using a single VPC. Cyber security controls are also migrated 1-to-1 using almost the same security policies. This strategy can be high risk due to the lack of proper visibility and configuration management.
- **Refactoring and containerisation –** The organization's applications are individual components consuming different libraries and dependencies to transform data information. Application micro-segmentation has separate containers associated with frontend, backend, and shared services, where traffic flows are split among the ingress, egress, east-west, and backhaul. Specific security controls are considered for the right enforcement and visibility.
- **Rebuilding (shift-and-lift) –** The organization's business process needs a complete redesign to create cloud-native applications. At this point cloud native application protection platforms are essential to enhance cyber security policies, and to provide a more significant advantage in the cloud.

What is Lift-and-Shift?

The main objective of lift-and-shift is to preserve the same architecture organizations already have in the public cloud, without making any significant changes in the design. In other words, it is the process of migrating an identical copy of a workload (including the operating system, applications, and data), network design, and management, as-is. This makes it the fastest and least expensive path. From a security perspective, it also preserves the same management systems and even keeps the same security policies, at least in the initial stage of the cloud transformation. Lift-and-shift is the most common first stage of a general cloud transformation journey since it is relatively easy and fast to achieve.

Lift-and-shift brings several benefits to the overall security posture and operations of organizations, such as:

- Autoscaling, agility, and speed.
- Deployment of dynamic infrastructure.Zero Trust and micro-segmentation.
- Adaptive and dynamic cloud-native security.
- A transition from CAPEX to OPEX.

However, organizations should beware of confusing the lift-and-shift migration model with a copy-and-paste strategy. Such a misunderstanding may lead to disaster if design errors are migrated, especially in security systems without the right controls and policies impacting the level of service.

Lift-and-Shift Optimized Model

Check Point recommends a new migration model, enabling organizations to have greater flexibility, agility, speed, scalability, dynamic security, and posture management, for a better shared-responsibility model in their cloud data center strategies. This model, called **Lift-and-Shift Optimized**, enables the harmonization of hub-and-spoke principles and a Zero Trust extended framework[2] to deliver full visibility and control of security and compliance. Consequently, it helps to minimize the

[1] Gartner Forecasts Worldwide Security and Risk Management Spending Growth to Slow but Remain Positive in 2020, URL: https://www.gartner.com/en/newsroom/press-releases/2020-06-17-gartner-forecasts-worldwide-security-and-risk-managem

[2] What ZTX Means for Vendors and Users, URL: https://go.forrester.com/blogs/what-ztx-means-for-vendors-and-users/

attack surface and protects against vulnerabilities, identify theft, and data loss. Our suggested model can be used as the first step for migrating workloads that are candidates for migration to the cloud.

In the following sections, we will present different use cases with reference architectures where CloudGuard NS (Cloud Network Security, aka CloudGuard IaaS) can provide a robust solution to secure all communication flows in the organization's VPC for a multi-cloud strategy deploying Azure and Amazon Web Services. Additionally, we will explain the importance of deploying CloudGuard Posture Management as a single pane of glass to provide security posture management for IaaS deployments in multi-cloud architectures, thus simplifying cloud security operations.

Shared Responsibility for Public IaaS

In traditional IT environments, the organization owns the whole stack, and the dedicated security team makes the necessary infrastructure changes. In the public cloud IaaS, some responsibilities are transferred to cloud service providers, and some are transferred to application owners.

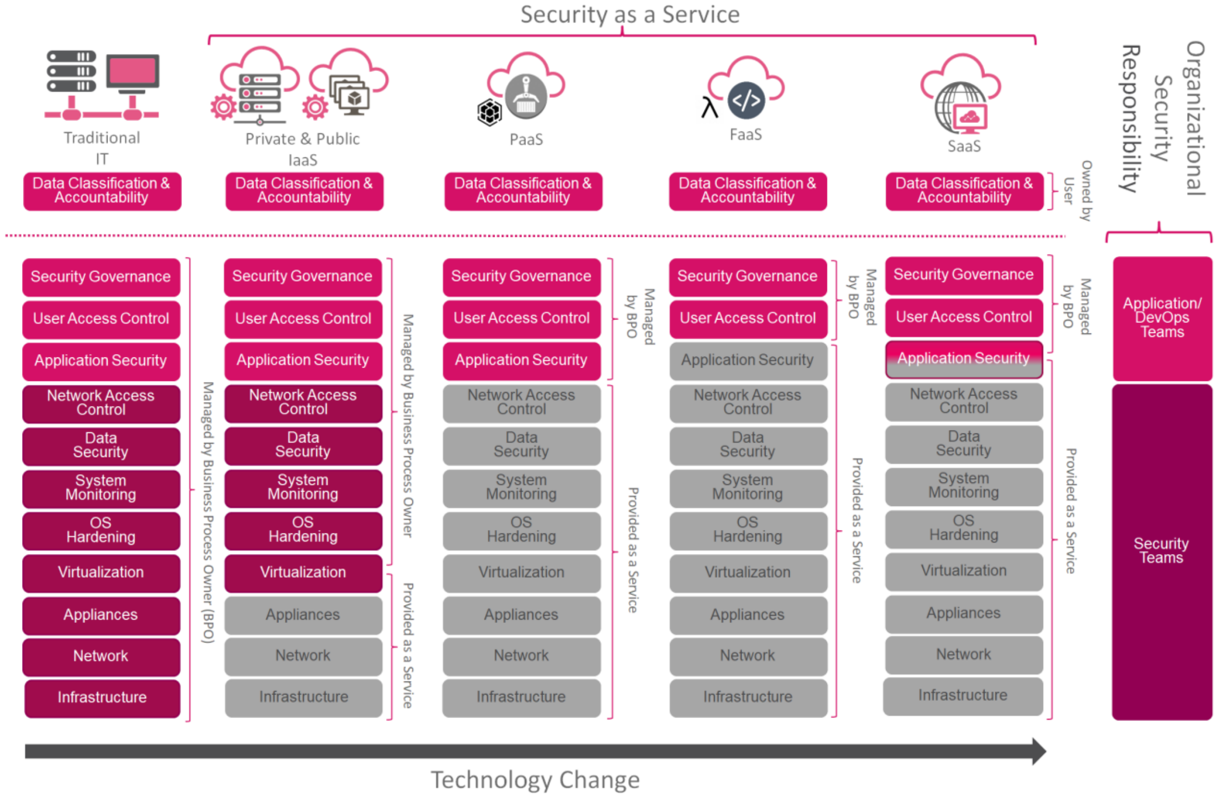

Figure 1: *Shared-Responsibility Extended Model for Public IaaS*

Cloud service providers are responsible[3] for ensuring the security of the cloud environment itself, however, IT security teams are responsible for the security controls of the infrastructure under their responsibility. Once an organization moves to the PaaS/FaaS and SaaS, some responsibilities will be transferred to the DevSecOps groups. Nevertheless, leading research and advisory company, Gartner, stated that "through 2020, 99% of cloud security failures are the customer's fault."[4] This means that the network security team is still responsible for the constant maturity of all the configurations related to the plumbing of the cloud data center. Cloud native network security and cloud security posture management tools for public IaaS therefore provide a single pane of glass for the right deployment of Zero Trust controls.

[3] Shared Responsibilities for Cloud Computing, URL: https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91
[4] Is the Cloud Secure, URL: https://www.gartner.com/smarterwithgartner/is-the-cloud-secure/

Zero Trust Model

Shared responsibility, aligned with Zero Trust principles, provides better harmonization of security controls, helping to minimize the potential risks in the migration process, especially in day-to-day operations. In this section, we will explain how the Zero Trust network and workload principles are less complicated, considering the hub-and-spoke approach and the service-oriented[5] architecture (SOA) used to protect the applications and services of the organization.

A Zero Trust framework considers the following pillars:

- **The Data:** The core of the business.
- **The Workloads:** Spokes that transform data in the information.
- **The Networks:** Hubs where data and information are transported using micro-segmentation and end-to-end encryption mechanisms.
- **The Devices:** Endpoints or IoT devices that upload or have access to data in the Hubs.
- **The People:** Who consume information using applications provided by the spokes. Also includes the administrators for the management of the cloud security operations through the security posture.

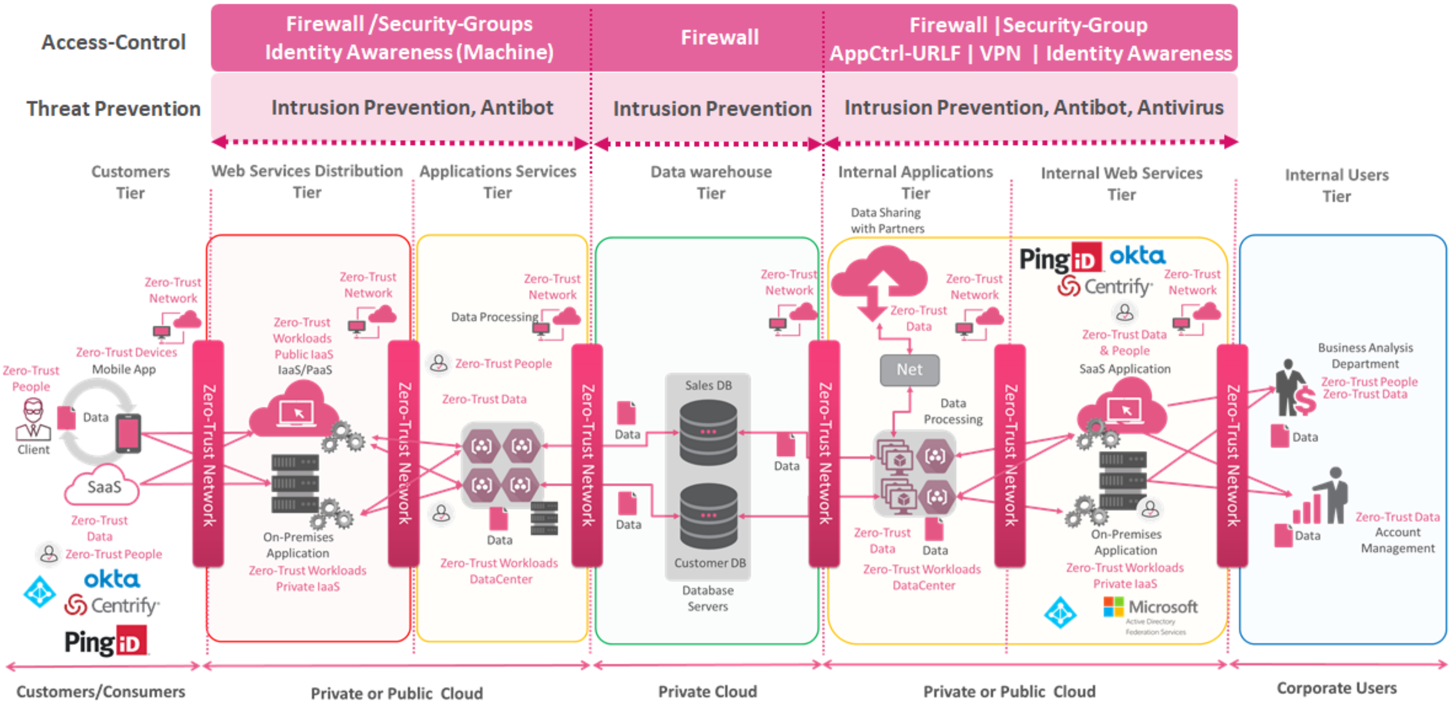

Figure 2: *Zero-Trust Architecture Reference: Architecture for Public and Private IaaS*

Proper security policies between the web tier, service-oriented application tier, and the database tier allows the organization to have a much better posture in the cloud environment to protect the different assets in the public IaaS, therefore minimizing risks. While SOA is the traditional approach that organizations follow in the migration process, once they start to migrate to the microservices, the SOA model should be transformed too.

[5] Service-Oriented Architecture (SOA), URL: https://www.ibm.com/support/knowledgecenter/en/SSMQ79_9.5.1/com.ibm.egl.pg.doc/topics/pegl_serv_overview.html

IaaS Security Segmentation

IaaS segmentation aims to reduce the blast radius and allow security teams to enforce perimeter security controls. Two powerful options facilitate these capabilities: micro-segmentation and macro-segmentation. In the following table, we will review the differences between the tools and their practical uses for IaaS segmentation.

- **Macro-segmentation:** Creates security zones within the overall network to prevent attacks between primary server segments, using multiple workloads with the same functionality and security classification.
- **Micro-segmentation:** Logically divides vNET/VPC into distinct security segments up to individual workload levels. Such a granular level at which micro-segmentation controls workload traffic, minimizes security threats and creates a Zero Trust security model. For the public cloud, the Check Point CPM provides centralized management and can be applied to individual workloads, enabling a more secure environment without the additional overhead of workload-specific configuration.

	Public IaaS Micro-Segmentation	Public IaaS Macro-Segmentation
Use Case	Used to logically divide the VPC/vNET into different security zones, up to individual workload level	Used to segregate between major groups of workloads with similar functionality and security classifications (such as web servers, application servers, and databases), preventing attackers from moving inside the perimeter and attacking the production workloads
Scope	More granular since it controls lateral movement across hosts	More on the perimeter level and across security zones
Policies	Granular host-to-host policies	Network/segment level policies
Policy Enforcement	Computing instances	Subnet/VLAN
Management and Control	Host-to-host security policies for access control or threat prevention	Functional vNET/VPC security policies for access control or threat prevention
Host-to-Host Communication Control	Between workloads in the same segment	Network or security zone level
Traffic Path Control	East-west or lateral traffic	North-south and east-west (inspect traffic between web, app, and DB zones)
Benefits	- Enforce granular tier-level segmentation within the same application group. Critical applications will remain safe even in the case of a breach - Enforce policy up to layer 7	- Enforce security at the perimeter to protect against attacks - Simpler to implement than micro-segmentation
Disadvantages	High-level skills are required, including application-level visibility, to employ micro-segmentation	Advanced network and security skillsfor deploying network-based segmentation policies

Figure 3: *Micro-Segment and Macro-Segment Attributes*[6]

[6] Matrix from networkinterview.com, URL: https://networkinterview.com/micro-segmentation-vs-network-segmentation/

Using macro-segmentation and micro-segmentation can help organizations make better choices regarding the security controls that can be used according to the applications flows. In addition, access control can be deployed in three different scenarios: network-based, agent-based, or host-based and API, using cloud-native tools.

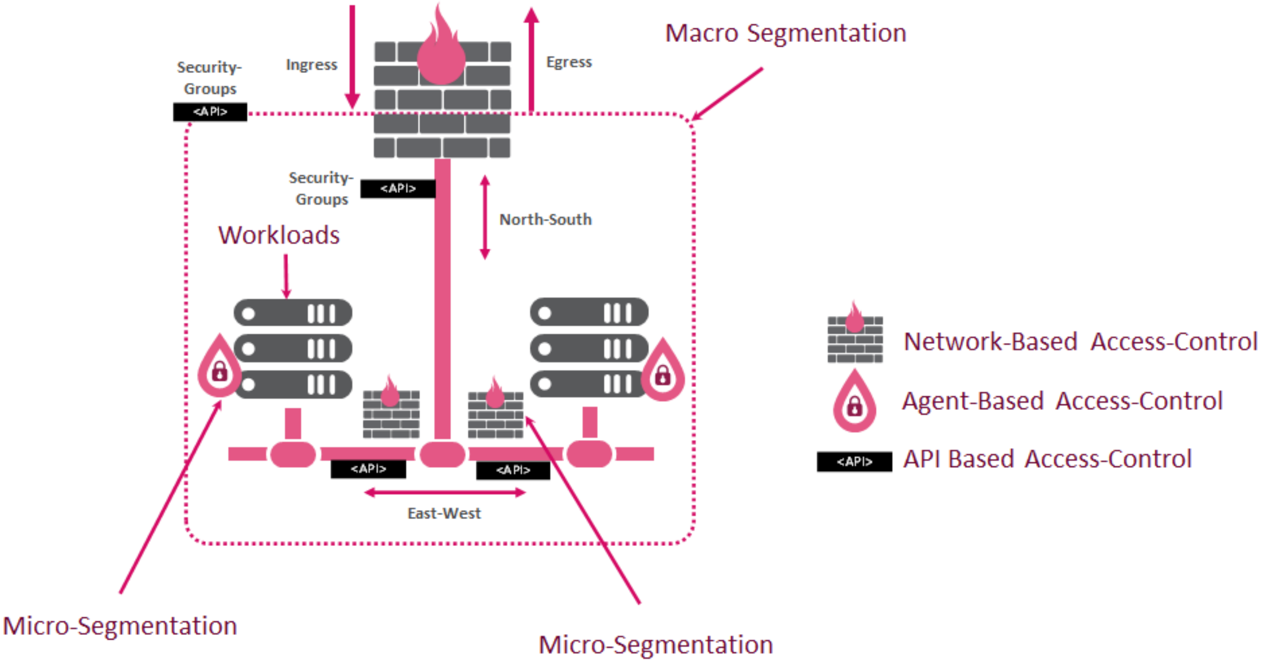

Figure 4: *Macro-Segmentation and Micro-segmentation*

The diagram above provides a visual representation of all the flows protected in the public IaaS. With this perspective, the following table proposes different security segments according to the flows, enabling you to select the controls needed more accurately.

Security Segment or Security Hub	Flow	Security Blades
Ingress Traffic from the Internet	North-south access control and traffic inspection	Firewall, rule-based IPS, SSL inspection
Egress Traffic to the Internet for Computing Instances, Azure Virtual Desktop or Amazon Web Services Workspaces	North-south access control and traffic inspection	Firewall, application control, URLF, Antibot, Antivirus, SSL inspection or HTTP categorization
Traffic Between Different vNETs/VPC and Workloads	East-west access control	Network security groups or firewall
Traffic Between Different vNET/VPC and Workloads	East-west traffic inspection	Firewall, rule-based IPS
Traffic from SD-WAN/MPLS (Backhaul)	North-south	Firewall, rule-based IPS, Identity-Awareness
Traffic Between OnPremises Data Center (Backhaul)	North-south	Firewall, rule-based IPS
Traffic Between Multi-Cloud Service Providers (Backhaul)	North-south	Firewall, rule-based IPS, VPN

Figure 5: *Aligning Security Blades With the Security Segments*

Under this approach, the shared responsibility model is more accessible for day-to-day operations. Transit Security Services vNET (Azure), Transit Gateway (Amazon Web Services), and Shared VPC (Google) provide the cloud data-centers with better scalability and enable the Zero Trust network principles related to the functional segmentation.

HUB-AND-SPOKE PRINCIPLES

This section, will explore the hub-and-spoke model recommended by Microsoft[7]. In short, the cloud environment is set up as a system of connections in which all spokes are connected to a transit hub, and all traffic to and from the spokes traverses the transit hub.

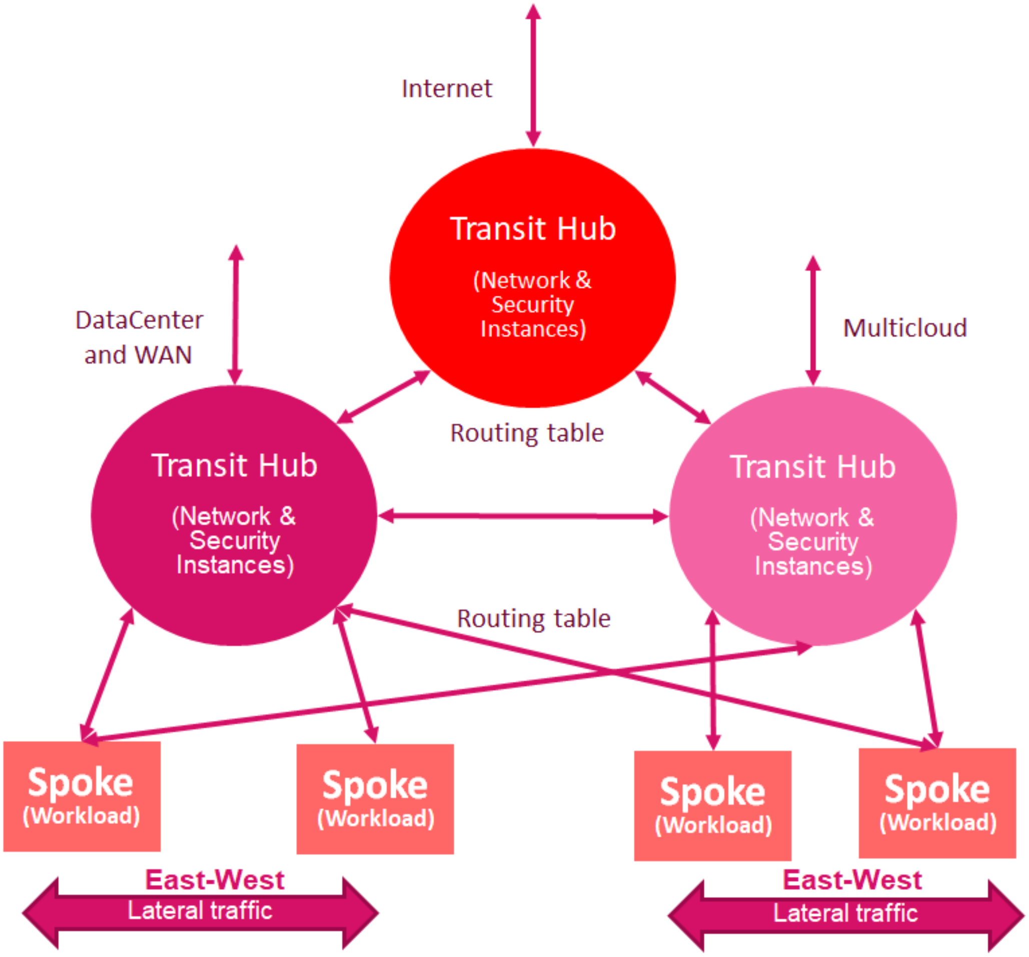

***Figure 6:** Hub and Spoke Architectural Principles*

The main focus of this principle is to provide a more practical segmentation of the lift-and-shift strategies when vNET/VPC is used to provide an easier setup for Zero Trust networking in the cloud. While we can segment inside a vNET/VPC, there is no easy way to enforce traffic inspection as cloud service providers control all routing inside the vNET/VPC perimeter.

Important definitions:

A Spoke is an isolated network environment that contains a collection of one or more network subnets from which typical workloads can be installed and run. A typical use case is a spoke that contains several virtual servers that make up either a part of, or an entire application stack (web, application, and database). Another use case is a spoke which acts as an extension of existing on-premises networks, such as a set of QA servers for testing purposes

[7] Hub-spoke network topology in Azure – URL: https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/hybrid-networking/hub-spoke

or a set of data processing servers that utilize the cloud's on-demand provisioning for lower cost and improved agility.

However, from the security perspective, we have different spokes that can be deployed in the public IaaS:

1.1. Transit Hubs
1.2. Computing Instances
1.3. Container On-Demand or as a Service
1.4. Kubernetes Clusters
1.5. Service EndPoints or VPC EndPoints
1.6. Serverless or Service Functions

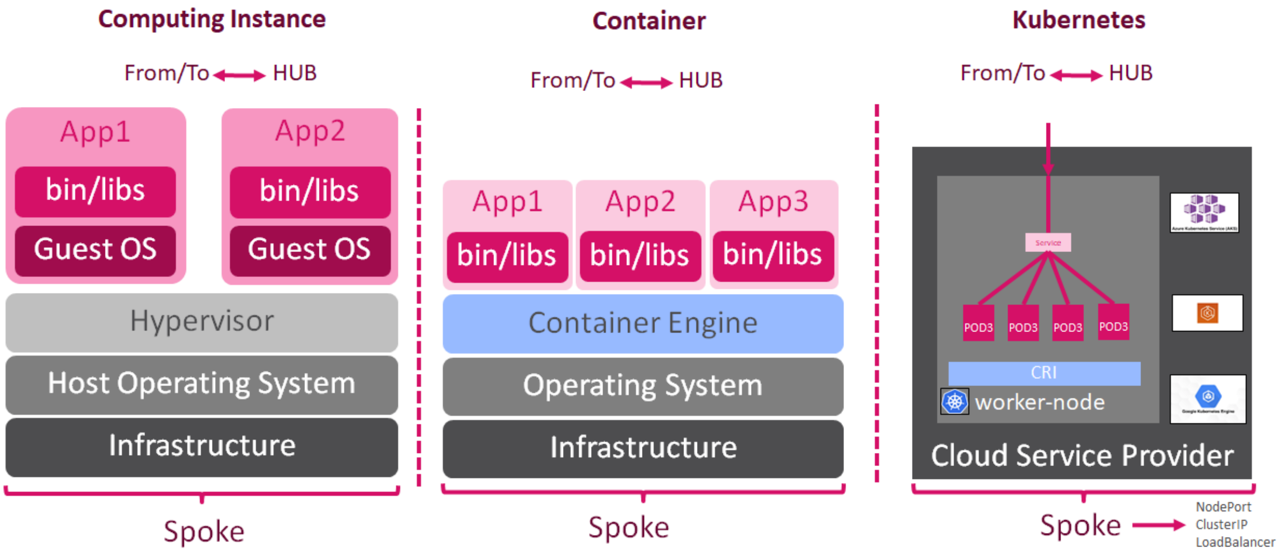

Figure 7: *Examples of Different Spokes for the Public IaaS*

A Transit Hub enables flexibility and systematic separation of communication flows through the environment. It can be designated for ingress traffic, lateral traffic between spokes, traffic in/out of the corporate network, or for outgoing traffic to the internet or other cloud environments. The routing traffic can be easily configured according to the traffic flows* in the applications.

*In the following section, we will explain the flows within the transit hubs and the interactions used, to build the correct plumbing for the infrastructure

A Transit Security Hub is an Azure Transit vNET or Amazon Web Services **Transit Gateway**, and GCP **Shared VPC** that interconnects all virtual cloud and on-premises networks. The transit security concept is defined as the security control point for cloud network interconnections and inter-spoke security. The hubs are the only way in/out of the environment as well as the only way to traverse inside and between spokes in the environment. This is due to spokes not being connected directly but only being accessible through one of the hubs. A key element is the routing and connection configuration between the hubs and spokes (UDR, static routing, or BGP for more complex environments).

High-Level Security Design

A hub-and-spoke network security design provides a central component connected to multiple networks around it, enabling different security controls. Setting up this topology in the traditional on-premises data center can be expensive, however, in the cloud, there is no extra cost. The lift-and-shift optimized model enables one to build powerful networking and security scenarios in the public IaaS that in turn enable organizations to have different scenarios providing an agnostic approach:

- Setting up separate development and production environments with different security controls enabled.
- Isolating the workloads of different customers using micro-segmentation and threat prevention capabilities.
- Segregating environments to meet compliance requirements, for example, PCI, GDPR and HIPAA.
- Segregating environments using cloud-native security controls, which use cloud security posture management tools.
- Providing shared IT services, like active directory, DNS, and file servers.
- Multicloud security as a code with automation in the provisioning of cloud security infrastructure as code, using CI/CD pipelines with Terraform, Jenkins, Puppet, and Ansible.

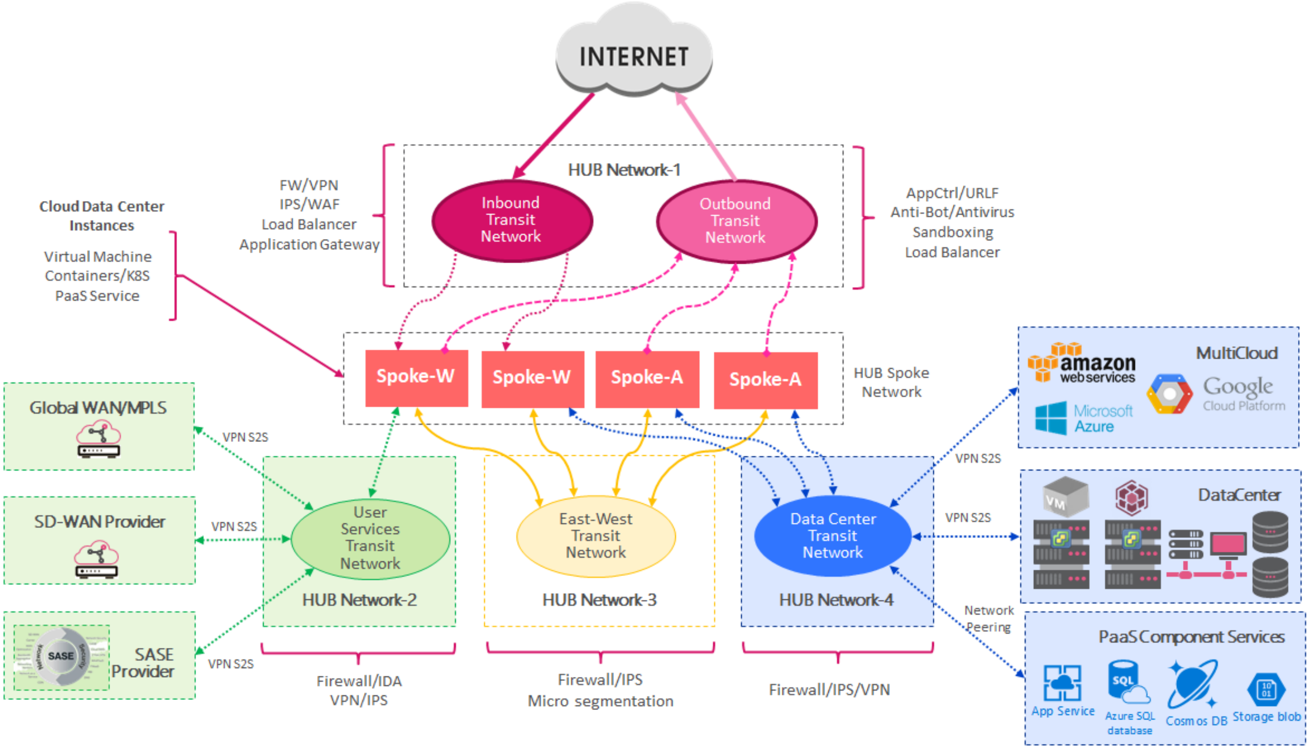

Figure 8: *Transit Hubs-and-Spokes: Lift-and-Shift Optimized Model*

In the lift-and-shift optimised model, we have four different types of hubs with different security functionalities:

A. Frontend Hub

Focused on providing communications to the public networks with two types of traffic: ingress and egress.

- The ingress perspective requires only access-control like the firewall and IPS with an SSL inspection mechanism, or web application and API protection.
- The egress perspective requires access-control for applications and websites. The common scenario is to provide secure internet access to the workloads to download fixes, patches, libraries, etc.

B. Data Center Extension Hub

Focused on providing communications from on-premises data centers and multicloud communications.

- Access-control like firewall/VPN and threat prevention, integrating the IPS and antivirus for traffic inspection from on-premises to the cloud data center.

C. East-West Hub

Focused on all lateral traffic between different spokes.

- Access-control micro-segmentation: using the cloud-native access-control tools to provide basic firewalling capabilities to allow or deny specific services.

- Threat prevention micro-segmentation: using advanced traffic inspection capabilities, providing virtual patch management.

D. User or Shared Services Hub

Focused on interconnecting users from remote branches to the resources deployed in the cloud data center.

REFERENCE ARCHITECTURE FOR PUBLIC CLOUD IAAS

In Azure, we have the Transit Security Services Hub, which is focused on providing security services as a central point of connectivity for all ingress, egress, east-west and data center extensions, multicloud peering and user traffic. This allows for a highly scalable and resilient design for large organizations. For example, ingress traffic should have dedicated auto-scale clusters in order to allow dynamic capabilities due to fluctuating throughput (similar scenarios that can be found in egress or east-west). In the case of the data center extension, we can deploy high-availability clusters for S2S VPN.

Organizations can choose their best strategy according to their analysis', and classify the flows according to the following scenarios:

- Scattered applications.
- Fluctuating throughput.

Additionally, in AWS, we have the transit gateway that provides a shared resource to allocate the routing domains to distribute the traffic between different VPCs. A route domain is a conceptual group of VPCs and/or VPNs attached to a single route table. From a security perspective, this tool provides flexibility to distribute the traffic between ingress, egress, data center, multicloud, and user traffic flows. AWS provides a highly flexible capability to enable new cloud network architectures and replace many point-to-point peering connections.

Macro-Segments and Micro-Segments With Transit Security Hubs

In this reference architecture, which follows the hub-and-spoke principles, there are five macro-segments (security hubs) where traffic flows should be protected according to their respective behavior.

Macro-segments

- **Frontend Security Hub:** Traffic related to productive systems and public services (ingress) and traffic related to providing secure internet access (egress) for different spokes, and for maintenance purposes (download patches, service packs, etc.).
- **Data Center or Multicloud Traffic Security Hub:** Traffic related to backhaul communications with the data center and multicloud.
- **Remote Branches or MPLS Security Hub:** Traffic related to users and remote branches to access internal productive systems and corporate services.
- **Cloud Security Management and Operations Security Hub:** Traffic related to day-to-day operations in the cloud.

Micro-segments

- **East-West Security Hub:** Traffic-related to the intercommunication between different spokes.

The following table shows the different terminology used by cloud service providers.

Feature	aws	Microsoft Azure	Google Cloud Platform
Geography	Geography	Geography	Geography
Availability Zone	Availability Zone	Availability Zone	Availability Zone
Network	VPC	VNET	VPC-Cloud Virtual Network
Subnet	Subnet	Subnet	Subnet Network
Resources Management	Across Specific Account	Across Specific Subscription	Global, Regional, and Zone Specific Resources
Virtual Machine (VM)	Instance	Virtual Machine	Virtual Machine Instance
Image Type Format	AMI	VM Images	Public / Private / Custom Image
Public IP Addresses	Public / Elastic IP	Basic / Standard IP	Ephemeral / Static external IP
Load Balancing	Application / Network / Classic Load Balancer / ELB	Azure Load Balancer, Application Gateway	External Network and HTTP Load Balancing, Internal Load Balancing
Native Security / Security Groups	Security Groups / NACL	Network Security Group (NSG)	Computer Engine Firewall Rules
Scalable Computer Instances (Servers)	Elastic Computer Cloud (EC2)	Azure VM	Computer Engine
Domain Name System (DNS)	Route 53	Azure DNS or Traffic Manager	Cloud DNS
Network Address Translation (NAT)	NAT Gateways	NAT Gateways	Cloud NAT
Network Peering	VPC Peering Connections	Virtual Network Peering	VPC Network Peering
Network Routes / Routing	Route Tables	Azure Virtual Network Routing	Routes
Region	Region	Region	Region
Virtual Private Cloud (VPC)	Virtual Private Cloud (VPC)	Virtual Network (VNET)	Virtual Private Cloud (VPC)
VPC Endpoints	VPC Endpoints	Virtual Network Service Endpoint	Private Services, Private Google Access and/or Shared VPC
VPN Gateway	Virtual Private Gateway	Azure VPN Gateway	Cloud VPN
Object Storage	S3 Buckets	Blob Storage	Cloud Storage
Identity and Access Management (IAM)	Identity Access Management (IAM)	Azure Role-Based ACL (RBAC) or Azure AD	Cloud IAM
Content Delivery Network (CDN)	Cloudfront	Azure CDN	Cloud CDN or CDN Interconnect
Autoscaling	Auto-scaling group	VM Scale Sets	Computer Engine Autoscaler
API Endpoints	API Gateway	API Management	Cloud Endpoints

Table 1: *Equivalent Terms Between Cloud Service Providers*[8]

[8] Overview - Cloud Feature Terms, URL: https://community.checkpoint.com/t5/Cloud-Network-Security-IaaS/Overview-Cloud-Feature-Terms/m-p/85610

Security Gateways - Deployment Modes

Single VPC/vNET Gateway

This is the most basic deployment used to provide advanced security **for small-to-medium** workloads. This scenario does not provide high-availability or scalability capabilities. It should only be considered for environments where resilience is not a major concern, and for testing purposes.

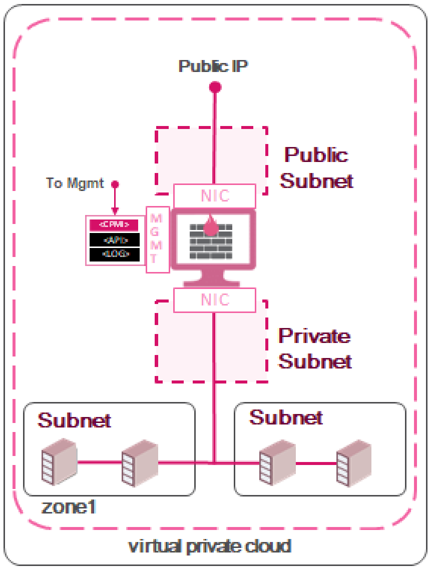

Cloud Service Provider	Supported Scenario
Azure	**Yes, sk109360**
Amazon Web Services	**Yes, sk120534**
Google Cloud Platform	**Yes, sk114577**
Oracle Cloud Infrastructure	**Yes,**
Huawei	**Yes**
Alibaba	**Yes**

Figure 9: *Single VPC with Single Security Gateway*

High-Availability Cluster Single VPC/vNET

A high-availability cluster is a group of virtual machines that work together, where one cluster member is the active, and the second cluster member is the standby. The cluster failover from active cluster member to the standby cluster member when necessary. This scenario provides advanced security for east-west traffic when regulatory requirements (NIST, ISO, PCI) demands visibility and enforcement. Additionally, this is an excellent scenario for handling VPN traffic between clouds or backhaul communications to on-premises data centers, especially for high demanding traffic like database replicas.

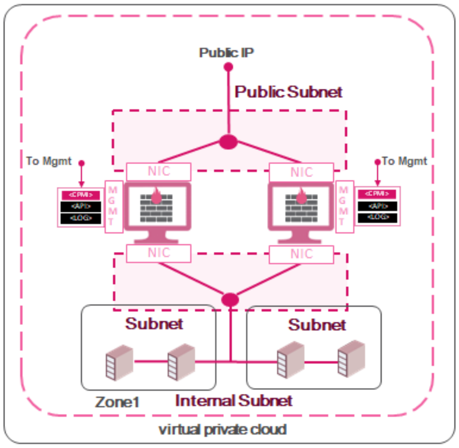

Cloud Service Provider	Supported Scenario
Azure	**Yes, sk109360**
Amazon Web Services	**Yes, sk120534**
Google Cloud Platform	**Yes, sk114577**
Oracle Cloud Infrastructure	**Yes, sk168202**
Huawei	**No**
Alibaba	**No**

Figure 10: *High-Availability Cluster Deployed in 1 Zone*

Two Gateways in Two Availability Zones, Single VPC/vNET

An availability zone is a high-availability offering that protects your applications and data from data center failures. Availability zones are unique physical locations within the cloud service provider regions, with each zone being made up of one or more data centers equipped with independent power, cooling, and networking systems. To provide an excellent resiliency, a typical design considers a minimum of three separate zones. Computing instances deployed and the security gateways in different zone-redundant services allow for the replication of applications and data across availability zones, protecting from SPOF (single-points-of-failure).

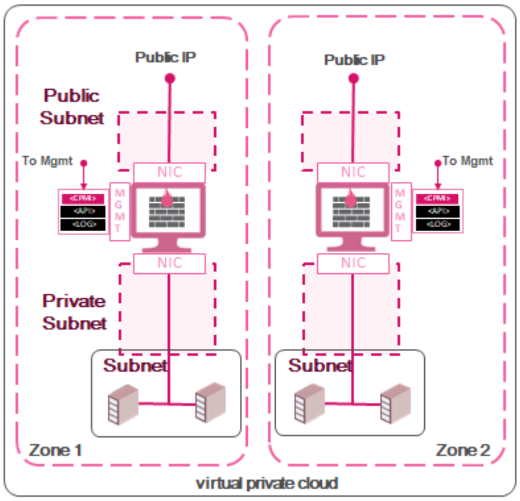

Cloud Service Provider	Supported Scenario
Azure	**Yes, sk109360**
Amazon Web Services	**Yes, sk120534**
Google Cloud Platform	**Yes, sk114577**
Oracle Cloud Infrastructure	**Yes, sk168202**[9]
Huawei	**Yes**
Alibaba	**Yes**

Figure 11: *High-Availability Architecture Considering 2 Different Zones*

Autoscale in Single VPC/vNET

Autoscale is a grouping of computing resources that you can use to deploy and manage sets of identical virtual machines (VMs). The scale sets increase or decrease the number of virtual machines based on current needs. This type of implementation is a perfect fit for workloads with fluctuating throughput, for example, public services or DMZ's where users access different services.

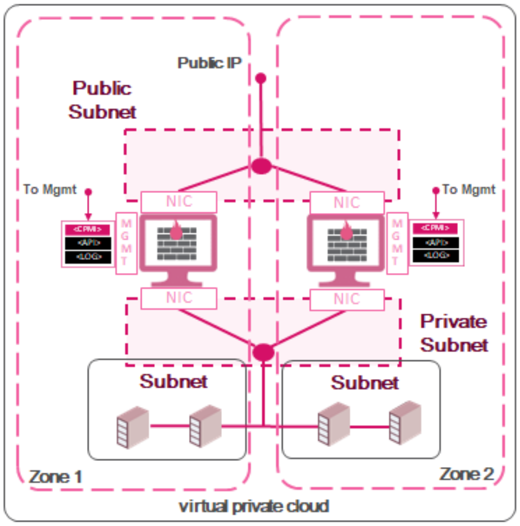

Cloud Service Provider	Supported Scenario
Azure	**Yes, sk109360**
Amazon Web Services	**Yes, sk120534**
Google Cloud Platform	**Yes, sk114577**
Oracle Cloud Infrastructure	**No**
Huawei	**No**
Alibaba	**No**

Figure 12: *Autoscale Cluster With Security Gateways Distributed in Different Availability Zones*

[9] CloudGuard for Oracle Cloud Infrastructure (OCI) – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk168202&partition=Basic&product=CloudGuard

Autoscale With Transit vNET or Shared VPC for Ingress-Traffic

This scenario is focused on inspecting ingress traffic through transit or shared service VPC/vNET with autoscale capabilities and fluctuating throughput. A typical scenario is to deploy an external load-balancer for internet connections, process traffic by the security gateway and place SNAT to keep the traffic symmetric, after which routing can forward the traffic to the relevant VPC/vNET's through the peering.

Important to note is that in Amazon Web Services, we have the Transit Gateway which is responsible for forwarding all traffic to the production VPC once the ingress VPC processes the traffic through TGW attachments.

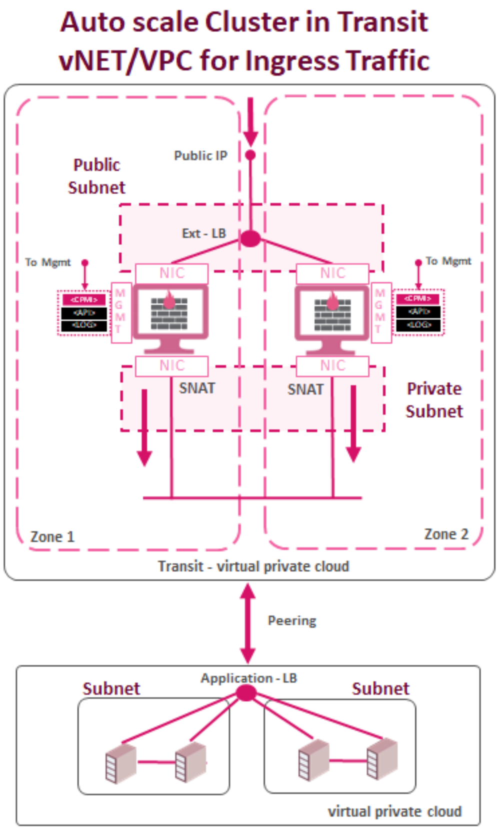

Cloud Service Provider	Supported Scenario
Azure	**Yes, sk109360[10]**
Amazon Web Services	**Yes, sk120534 (through TGW[11])**
Google Cloud Platform	**Yes, sk114577[12]**
Oracle Cloud Infrastructure	**No**
Huawei	**No**
Alibaba	**No**

Figure 13: *Autoscale Cluster for Ingress Traffic Using Transit or Shared Services VPC/vNET*

[10] Check Point Reference Architecture for Azure – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk109360&partition=Basic&product=CloudGuard

[11] TGW – Transit Security Gateway – URL: https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html

[12] Check Point CloudGuard IaaS reference architecture for Google Cloud Platform – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk114577&partition=Basic&product=CloudGuard

Autoscale With Transit vNET or Shared VPC for Egress and East-West Traffic

This scenario inspects egress traffic through transit or shared service VPC/vNET with autoscale capabilities and fluctuating throughput. A typical scenario is to deploy the internal load-balancer for internet connections, process the traffic by security gateway and place SNAT for public IP address. Additionally, it is possible to inspect east-west traffic between different vNET or VPC. This scenario should be used for small or medium environments where the egress traffic cannot share the gateway with the east-west.

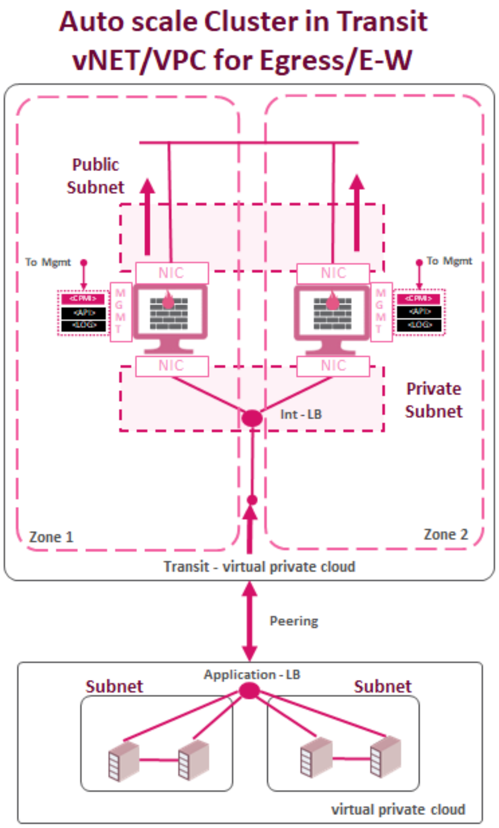

Cloud Service Provider	Supported Scenario
Azure	**Yes**
Amazon Web Services	**Yes, sk120534[13]. TGW[14] VPN with ECMP**
Google Cloud Platform	**Yes**
Oracle Cloud Infrastructure	**No**
Huawei	**No**
Alibaba	**No**

Figure 14: *Autoscale Cluster for Ingress Traffic Using Transit or Shared Services VPC/vNET*

[13] CloudGuard for AWS - Transit VPC Architecture – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk120534&partition=Basic&product=CloudGuard

[14] TGW – Transit Security Gateway – URL: https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpn-attachments.html

Autoscale With Transit vNET or Shared VPC for East-West Traffic ONLY

This scenario focuses on inspecting different vNET or VPCs (east-west traffic), while providing autoscale capabilities and processing fluctuating throughput. The internal load balancer is deployed to process the traffic between the vNETs with no direct connection to the internet (this scenario applies to Azure). However, in AWS, we can use the Geo-Cluster configuration or TGW Appliance Mode[15] to provide advanced capabilities. On the other hand, Google provides east-west to inspect traffic between VPC or internal subnets – please refer to the **sk114577**[16] for more use cases.

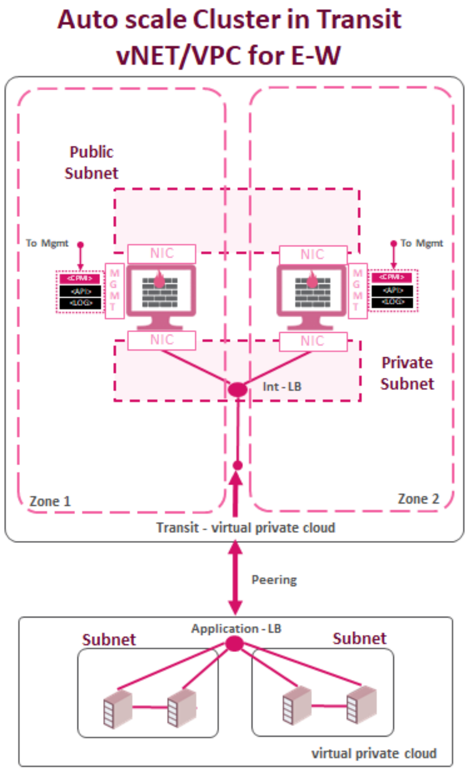

Cloud Service Provider	Supported Scenario
Azure	**Yes**[17]
Amazon Web Services	**Yes, sk120534** [18] **Also, please refer to the Appliance VPC with TGW Appliance Mode**
Google Cloud Platform	**Yes, sk114577**[19]
Oracle Cloud Infrastructure	**No**
Huawei	**No**
Alibaba	**No**

***Figure 15**: Autoscale Cluster with Security Gateways Distributed in Different Availability Zones*

[15] TGW Appliance Mode for Shared Services – URL: https://docs.aws.amazon.com/vpc/latest/tgw/transit-gateway-appliance-scenario.html

[16] Check Point CloudGuard IaaS reference architecture for Google Cloud Platform– URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk114577&partition=Basic&product=CloudGuard

[17] Virtual Machine Scale Sets (VMSS) for Azure R80.10 and Higher Administration Guide – URL: https://sc1.checkpoint.com/documents/IaaS/WebAdminGuides/EN/CP_VMSS_for_Azure/Content/Topics-VMSS-for-Azure/Overview.htm

[18] CloudGuard for AWS - Transit VPC Architecture – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk120534&partition=Basic&product=CloudGuard

[19] Check Point CloudGuard IaaS reference architecture for Google Cloud Platform – URL: https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk114577&partition=Basic&product=CloudGuard

Microsoft Azure Security Architecture

A typical security architecture operates as follows:

- The transit hub is a virtual network in Azure that acts as a central point of connectivity for your on-premises network.
- Check Point CloudGuard NS gateways act like security hubs or macro-segments, providing access control and threat prevention capabilities.
- Spokes are virtual networks that peer with the transit hub, and can be used to isolate workloads.
- Traffic flows between the on-premises data center and the transit hub through ExpressRoute or VPN gateway connection, implementing traffic inspection with the intrusion prevention system.
- Workloads deployed in different environments, such as development (dev), testing (QA or pre-prod), and production (prod), require shared services such as DNS, NTP and security (e.g. firewalls, IPS, WAF/WAAP, directory services, SIEM) that can be placed in a specific functional spoke (e.g. operations and management vNET) or the transit hub virtual network.

Each environment can be deployed as a functional spoke to maintain isolation. For example, workloads located in the frontend (production) or backend vNET (production, development, or QA) will use different security services, therefore access to shared services provides great flexibility, scalability, and better distribution of the responsibilities between the teams. As a best practice, production workloads are restricted to have connectivity to one another (micro-segment) provided by network security groups (NSGs) and require access to shared services through the transit hub (macro-segment). Under this approach, the security hubs provide central control in the transit hub, offering better isolation according to the traffic flows between workloads and traffic inspection.

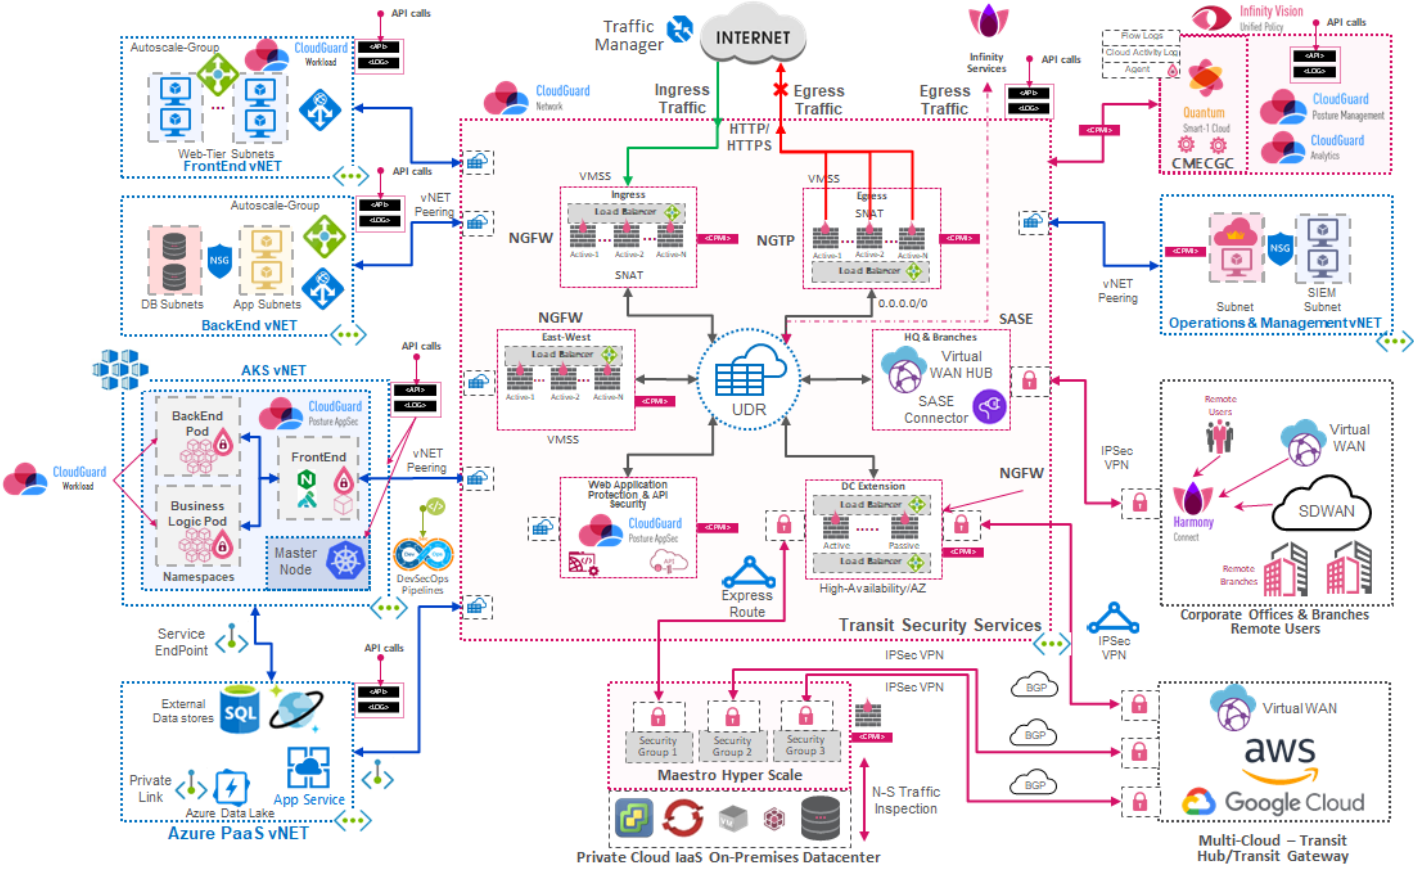

Figure 16: *Reference Architecture for Microsoft Azure*

The public cloud IaaS can be used in several cases. With this approach, it's important to understand different behaviors relating to the organization's needs, including:

- Hosting complex websites.
- High computing performance.
- Program testing and development.
- Disaster recovery or backup solutions.
- Big data analysis.

Google Cloud Platform (GCP) Security Architecture

In a similar approach to Azure, Check Point CloudGuard for GCP easily extends comprehensive threat prevention security to protect assets in the cloud from attacks, while at the same time enabling secure connectivity.

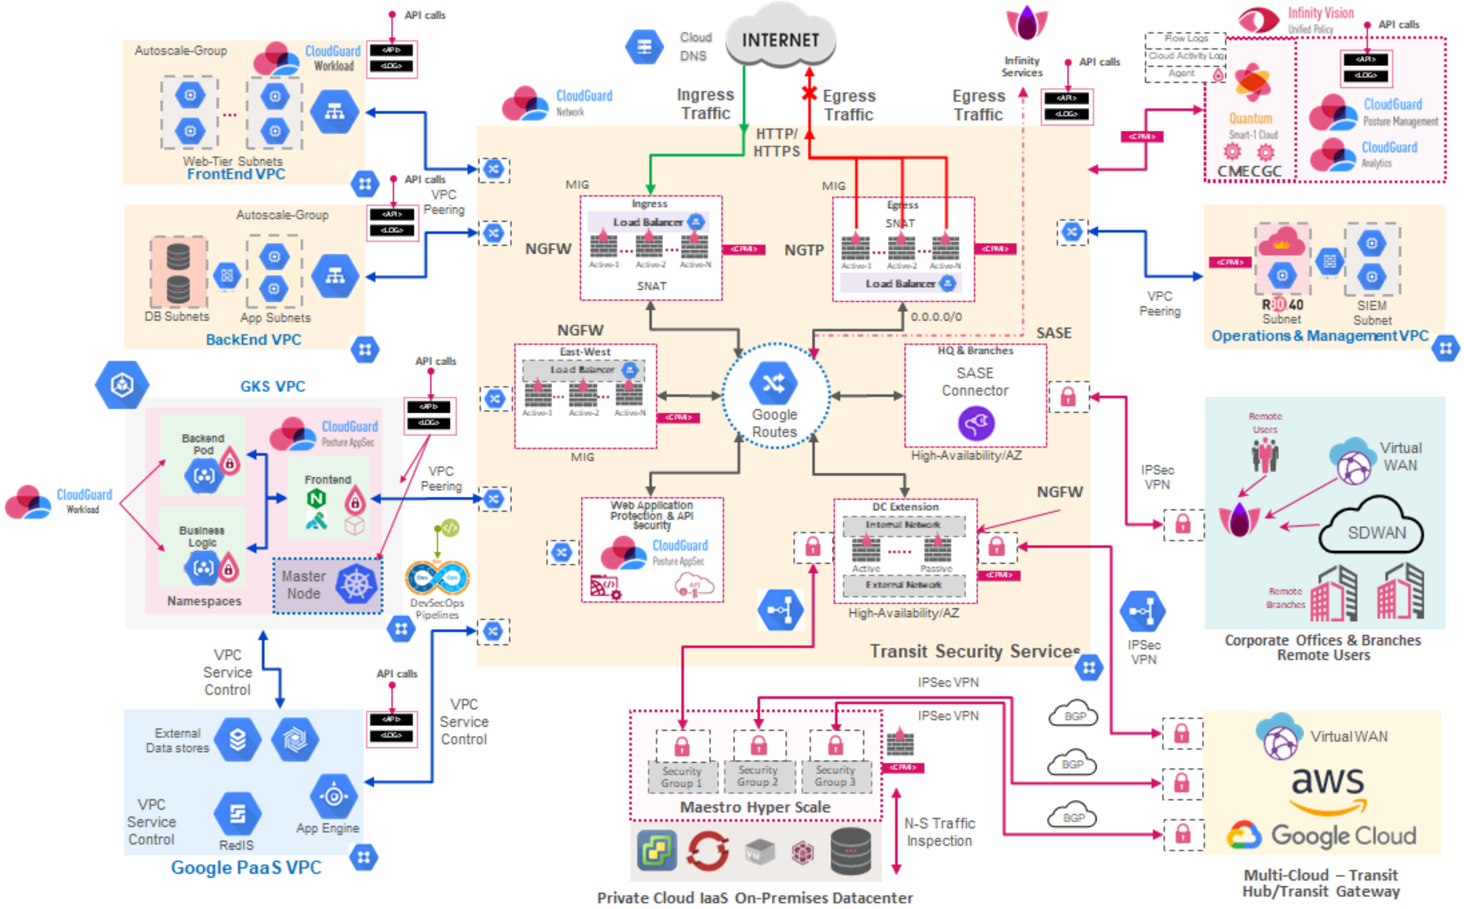

Figure 17: Reference Architecture for Google Cloud Platform

Use Cases for CloudGuard Network Security for Microsoft Azure and Google Cloud Platform

1. Ingress and Egress Hub

1.1. I2V: Internet to vNET/VPC

Ingress traffic for public services or corporate users is the most common scenario for organizations where the access-control (firewall) and threat prevention (IPS, web application and API protection, and inbound SSL inspection) are considered to protect the workloads. Traditional monolithic traffic inspection should be replaced for a rule-based approach, providing only protection for relevant ingress traffic.

1.2. V2I: vNET/VPC to the Internet

1.2.1. The main focus of this traffic flow is for maintenance purposes: egress traffic to allow workloads to securely access the internet, and to allow the download of patches, service-packs, etc. Mixing with the ingress traffic is only recommended in cases where the number of workloads requiring secure internet access is considerable. Additionally user traffic should not be placed here.

1.2.2. Internal segmentation is done by NSG.

2. Data Center Extension Hub

2.1. D2V: Data Center to vNET/VPC

This traffic is related to communications between the on-premises data center and the specific functional vNETs. Common scenarios include the data base replicas where access control and simple traffic inspection (for normalization purposes) may be required.

2.2. C2C: Cloud to Cloud

This traffic is related to multicloud communications between different cloud service providers. For example, peering communications between Amazon Web Services, Google, Oracle, or Huawei. This case is commonly used for database replicas sharing information between APIs.

3. East-West Hub

3.1. V2V: vNET/VPC to vNET/VPC – East-West Access Control

This traffic is related to the communication of several workloads. Access control only provides the isolation to prevent communication between them with basic protection against lateral attacks. Using network security groups is the best option here. However, it is important to use cloud security posture management tools to validate that all configurations and segmentation have been done properly.

3.2. V2V: vNET/VPC to vNET/VPC – East-West Threat Prevention

This traffic is related to the communication of several workloads where traffic inspection is required to provide more advanced capabilities like intrusion prevention or Antibot. This approach should be used only for specific requirements, not for all the internal traffic.

4. Remote Branches and Users Hub

4.1. B2V: Branch to vNET/VPC

This traffic is related to the communication from remote branches to specific services located in the functional vNETs (frontend or backend), for corporate applications or virtual desktops

4.2. VW2V: Virtual WAN to vNET (applicable only to Azure)

This traffic is related to global peering between different regions where organizations needs to deploy several cloud data centers to access corporate applications.

4.3. B2I: Branch to Internet

This traffic is related to secure internet access for the remote branches; it can also be categorized as egress traffic. However, the main difference from the frontend hub is that this traffic is more related to users and used only for specific scenarios.

For all scenarios, it is important to note that there is a direct correlation between concurrent connections (CC), connections-per-second (CPS), and packets-per-second (PPS), as each cloud service provider has "rate-limits" defined per computing instance. Traffic flow analysis should therefore be done before the migration process to ensure the use cases and the right configurations are applied properly.

Amazon Web Services (AWS) Security Architecture

Check Point CloudGuard NS for AWS easily extends comprehensive threat prevention security to the AWS Cloud by protecting assets in the cloud from attacks. The service also enables secure connectivity and allows for the enforcement of consistent security policies across your entire organization. This is done by protecting data traveling between the corporate network and the Amazon Web Services VPC (Virtual Private Cloud). It also inspects data that enters and leaves the private subnet in the VPC to prevent attacks and mitigate data loss or leakage. CloudGuard NS can therefore protect services in the public cloud from the most sophisticated threats, unauthorized access, and also prevents application layer Denial of Service (DoS) attacks.

Check Point CloudGuard NS for AWS meets organizational cloud security needs in the following ways:

- Automatically deployed tag-based IPsec VPN between AWS Transit Gateway and the security VPC.
- Automatic configuration of AWS VPN Gateways on spoke VPCs. This includes planning of IP addresses to prevent subnet IP address conflicts.
- Next Generation firewall with application control, data awareness, HTTPS inspection, NAT, and logging.
- IPS and virtual patching of cloud resources, Anti-bot and Anti-virus, and zero-day threat emulation and threat extraction.
- Application control and URL filtering for internet-bound traffic.
- Remote access VPN to connect remote clients.
- IPsec VPN for VPC-to-VPC, and VPC-to-on-premises connections with optional direct connect support.

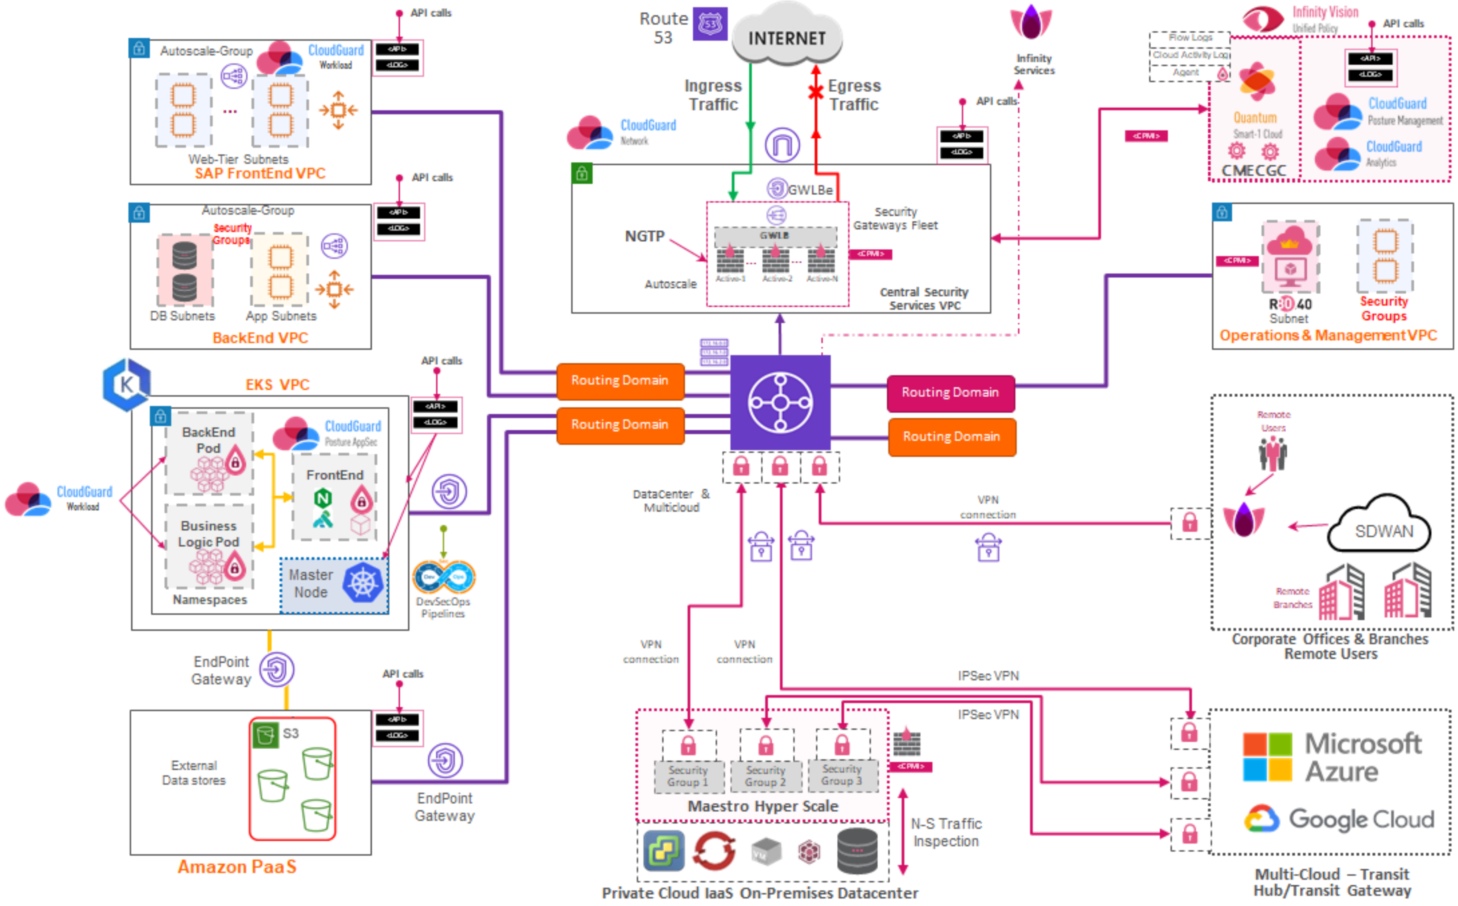

Figure 18: *Reference Architecture for Amazon Web Services*

The diagram above shows the Transit Gateway architecture of Check Point CloudGuard Network Security for AWS, an end-to-end solution that includes:

- Amazon Web Services Transit Gateway (TGW) object.
- Spoke (consumer) VPCs attached to the AWS Transit Gateway with their routing domains.
- Ingress security VPC with CloudGuard security gateways auto scaling group, attached to the AWS Transit Gateway.
- Egress security VPC with the CloudGuard Transit Gateways auto scaling group.
- Automatic provisioning of VPN tunnels.
- BGP routing configuration between the AWS Transit Gateway and the CloudGuard NS security gateways.
- Corporate VPN between on-premises perimeter and the AWS.

A typical security architecture operates as follows:

- AWS Transit Gateway (TGW) connects VPCs and on-premises networks through a central hub, simplifying the network designs and eliminating complex peering relationships.
- TGW acts as a cloud router where each new connection is only made once.
- Integrating the security gateways with TGW, Check Point can deliver comprehensive security for cloud workloads and assets with VPC perimeter security services.

Other benefits of this type of architecture include seamless security segmentation between VPCs, and automatically established IPsec VPN connectivity between cloud environments.

For higher demands in traffic and to provide a more centralized approach for shared services, similarly to the transit hub in Azure, the TGW has the capability to configure an appliance (such as a security appliance) in a shared services VPC. This scenario provides the scalability and flexibility required when traffic demands grow. To facilitate such growth, traffic is routed between Transit Gateway attachments after first being inspected by the appliance in the shared services VPC. This creates excellent flexibility and scalability options for the growing demands of huge cloud data centers.

Once the security management server and shared security hub are deployed, every new or existing VPC that is specifically tagged is automatically configured to route all traffic, providing a more simplified way of deployment. For more complex environments where BGP should be considered, traffic is routed via an AWS managed VPN gateway into the security hub. The VPN gateways are also added as IPsec interoperable devices. The gateways are then added to the automatically established route-based VPN process, (powered by Border Gateway Protocol (BGP)), to propagate network route changes to on-premises and in-cloud tenants.

Transit Gateway Appliance Mode and GWLB – Gateway Load Balancer

AWS has launched the GWLB (Gateway Load Balancer), which combines the functionality of a L3 gateway and a L4 load balancer, allowing organizations to transparently insert virtual network appliances into security gateways for deep packet inspection. The following diagram explains the flows and functionality of the GWLB and the Security Gateways fleet.

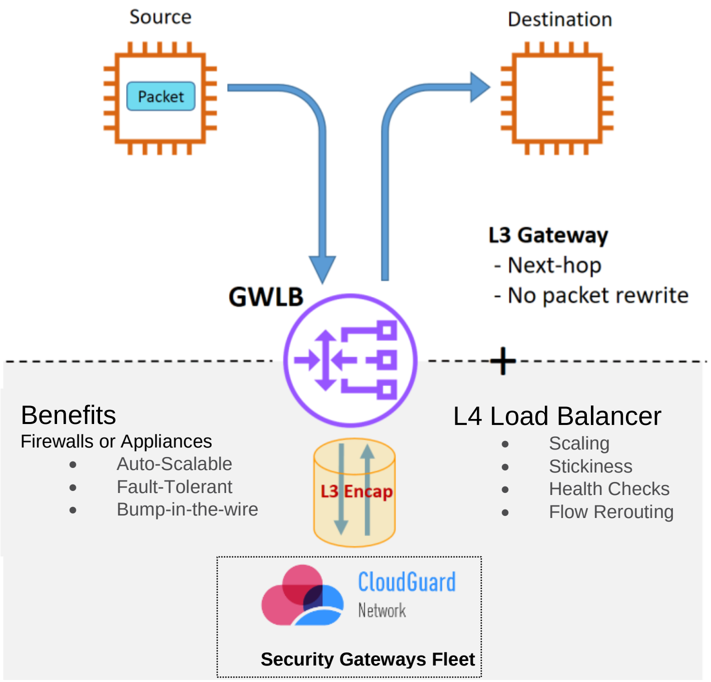

***Figure 19**: AWS Gateway Load Balancer - Flow and Functionality (source: Amazon Web Services)*[20]

A typical security architecture operates as follows:

- Traffic that would otherwise flow from the source directly to the destination is routed to the GWLB using route tables.
- GWLB has a frontend, which faces the traffic source and destination, and operates in bump-in-the-wire mode, acting as the next-hop gateway, with no packet rewrite.
- GWLB's backend, which faces the fleet of virtual appliances, operates as a load-balancer for routing traffic flows to and from one of multiple equivalent target appliances.
- GWLB ensures stickiness of flows in both directions to target appliances, performs regular health checks of the appliances, and reroutes flows if the selected appliance becomes unhealthy.
- When GWLB receives a traffic packet, processes it at layer 3, it encapsulate with a Geneve protocol (Generic Network Virtualization) encapsulation and sends it to the appliance without any change to the original traffic packet (also known as "packet-in, packet-out") to achieve transparent forwarding.

The GWLB Endpoint (GWLBE) is another important component of the GWLB. The GWLBE (one GWLBE per VPC) behaves as a mirror image of a GWLB's frontend, thus acting as the next-hop gateway for each subnet that needs traffic inspection, which is different from the VPC that hosts the corresponding GWLB. Thus, the GWLBE enables the deployment of a multi-tenanted traffic filtering service in a Service Provider VPC (different from the VPC/s of the GWLBE/s).

[20] Amazon Web Services Gateway Load Balancer – URL: https://aws.amazon.com/elasticloadbalancing/gateway-load-balancer/

The diagram below shows a simple architecture using GWLBE.

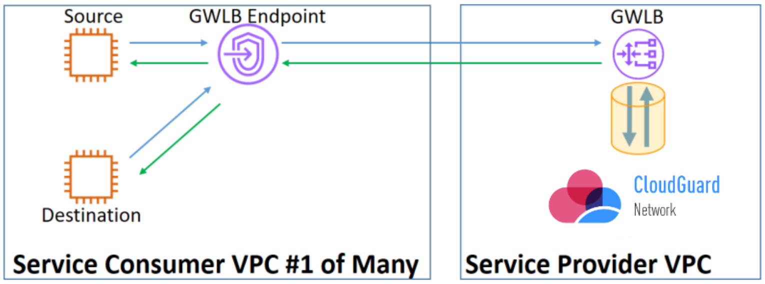

***Figure 20**: Access via GWLBE (source: Amazon Web Services)*

With c onsideration of the Zero Trust principles, GWLB allows organizations to design fault-tolerant architectures in an easier and more intuitive way; specifically the addition of multiple virtual appliances and functional Service Provider VPCs. This approach allows the organization, especially the IT and Networking teams, to maintain consistent security practices between in-cloud and on-premises deployments. This provides a huge benefit of being able to leverage the existing skill-set of security engineers who understand and trust Cloud Guard Network Security virtual appliances with NGFW or NGTP.

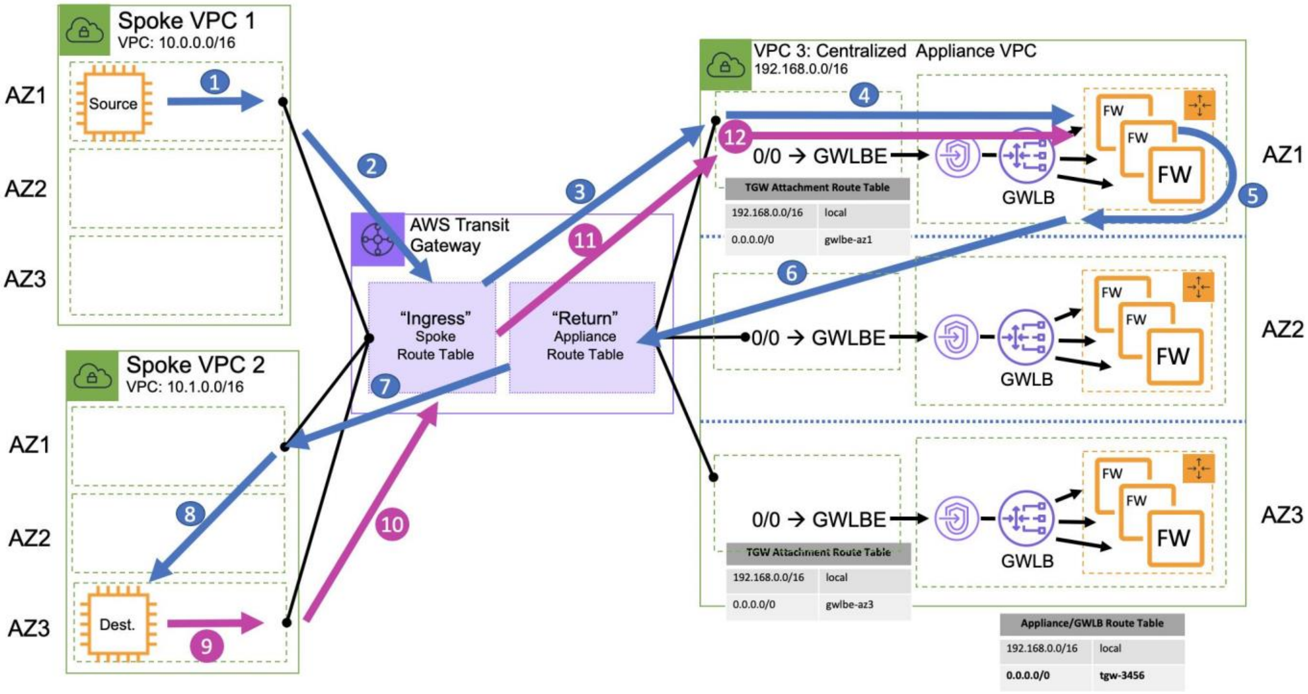

***Figure 21**: Symmetric Traffic for Security Gateways (source: Amazon Web Services)*

Below, are the top use cases for CloudGuard NS for AWS.

1. Transit Gateway (TGW)

1.1. Acts as a regional virtual router for traffic flowing between all virtual private clouds (VPC) and VPN connections. The TGW scales elastically based on the volume of network traffic. Routing through a TGW operates at layer 3, where the packets are sent to a specific next-hop attachment, based on their destination IP addresses.

1.2. Functionalities

1.2.1. AWS TGW - Is a cross-account, per-region service available to VPCs and VPNs in those regions.

1.2.2. AWS TGW Route Tables - The dynamic and static routes that decide the next hop for VPCs/VPNs based on the destination IP address. AWS TGW comes with one default table. There is an option to add many route tables.

1.2.3. AWS TGW Attachment - AWS TGWs can have VPCs or VPNs as attachments. A VPC or VPN can be attached to one or more AWS TGWs.

1.2.4. AWS TGW Route Domain - Similar to virtual routing and forwarding (VRFs) in traditional networks, a route domain is a conceptual group of VPCs and/or VPNs attached to a single route table.

1.2.5. AWS TGW Appliance Mode - All traffic routed between TGW attachments is first inspected by the security stateful appliance in the shared services VPC.

1.2.6. AWS Gateway Load Balancer (GWLB) – The GWLB makes it easy to deploy, scale, and manage your third-party virtual appliances. It gives you one gateway for distributing traffic across multiple virtual appliances, while scaling them up, or down, based on demand. This eliminates potential points of failure in your network and increases availability.

1.3. Scenarios for Deployment:

1.3.1. Inter-Account Communication

1.3.2. Inter-Region Communication

1.3.3. On-Premises Data Center to AWS Cloud Communication

1.3.4. Multicast Communication

2. Frontend Hub

2.1. I2V: Internet to VPC

Ingress traffic related to public services for customers or corporate users, is the most common scenario for organizations where the access-control (firewall) and threat prevention (IPS, web application and API protection and inbound SSL inspection) are deployed. Traditional traffic inspection should be replaced by a rule-based approach, providing protection only for relevant traffic, for better performance.

2.2. V2I: VPC to the Internet

Egress traffic is used to allow workloads to securely access the internet, download patches, service-packs, and more. The main focus of this traffic flow is for maintenance purposes. It is not recommended to mix it with the ingress traffic, except if the number of workloads that requires secure internet access is considerable. Additional user traffic should not be placed here.

3. East-West Hub

3.1. V2V: VPC to VPC – East-West Access Control

This traffic is related to the communication of several workloads. Access control only provides the network isolation to prevent communication between workloads, with basic protection against lateral attacks. Using network security groups is the best option here, however, it is important to use cloud security posture management tools to validate that all configurations and segmentation have been done properly.

3.2. V2V: VPC to VPC – East-West Threat Prevention

This traffic is related to communications of several workloads where traffic inspection is required to provide more advanced capabilities like intrusion prevention or Anti-Bot. This approach should be used only for specific requirements and not placed for all internal traffic.

4. Remote Branches and User Hubs

4.1. B2V: Branch to VPC

This traffic is related to communications from remote branches to access specific services located in the functional VPCs (frontend or backend) for corporate applications or virtual desktops.

4.2. B2I: Branch to Internet

This traffic is related to secure internet access for remote branches. It can also be categorized as egress traffic, however, the main difference to the frontend hub is that the traffic is more related to the users, and used only for specific scenarios.

5. Data Center Extension Hub

5.1. D2V: Data Center to VPC

This traffic is related to communications between the on-premises data-center and specific functional VPCs. A common scenario is the database replica, where access-control and simple traffic inspection (for normalization purposes) may be required.

5.2. C2C: Cloud to Cloud

This traffic is related to multicloud communications between different cloud service providers. For example, peering communications between Amazon Web Services, Google, Oracle, or Huawei. This case is commonly used for database replicas, sharing information between APIs.

USE CASES

In this section, we will discuss the scenarios in more detail using the reference architecture blueprints for Azure and Amazon Web Services, as well as the following use cases:

- Ingress and egress.
- East-west and data center extension (backhaul).
- Remote access and virtual desktop infrastructure with SASE.

The following diagram shows the vNETs in which the "lift-and-shift" infrastructure is located.

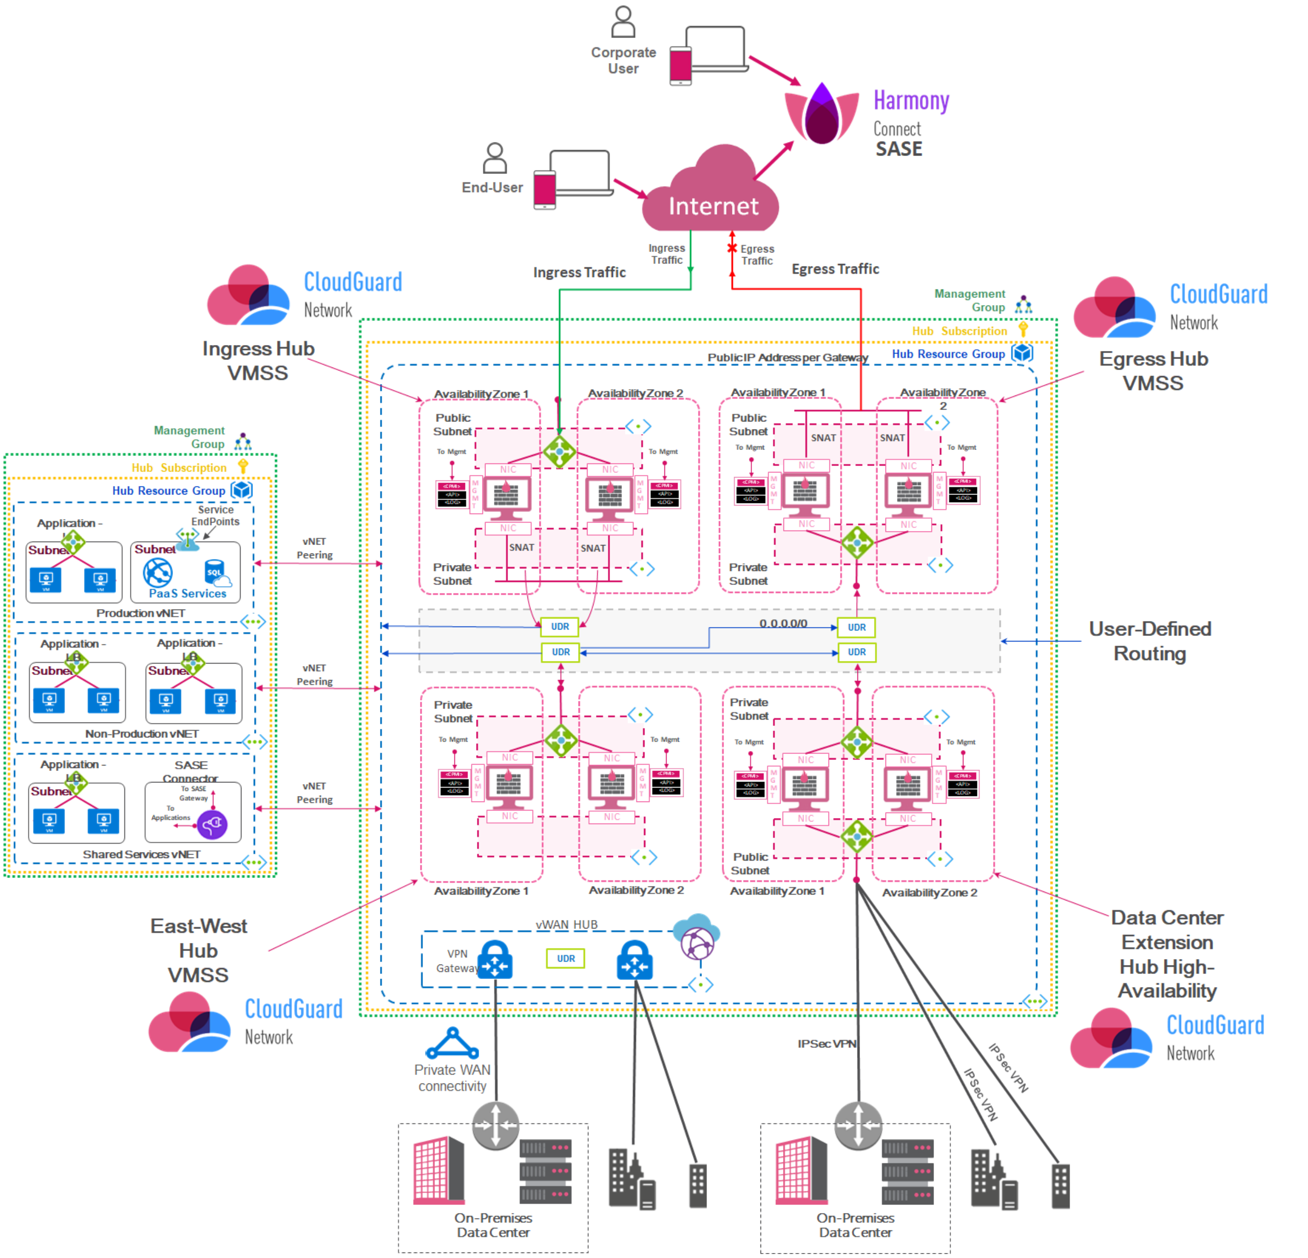

Figure 22: *Transit Security Services Hub - Reference Architecture for Microsoft Azure*

Next, we have the Google Cloud Platform reference architecture that illustrates how CloudGuard NS can protect resources deployed in the Google Cloud Platform. A common use case involves a web application environment deployed in a virtual network in the Google Compute Engine. Such web applications can consist of multiple tiers, including a web tier, application tier, database tiers, and containers.

Traffic that flows in this environment typically has the following attributes:

- Traffic arriving from the internet to the web tier, known as ingress.
- Traffic between the web tier and the application tier, known as east-west.
- Traffic between the environment and an on-premises network for administration and backend services, known as data center extension, or backhaul.
- Outgoing connections from the environment to the internet for software updates and access to external web services, known as egress

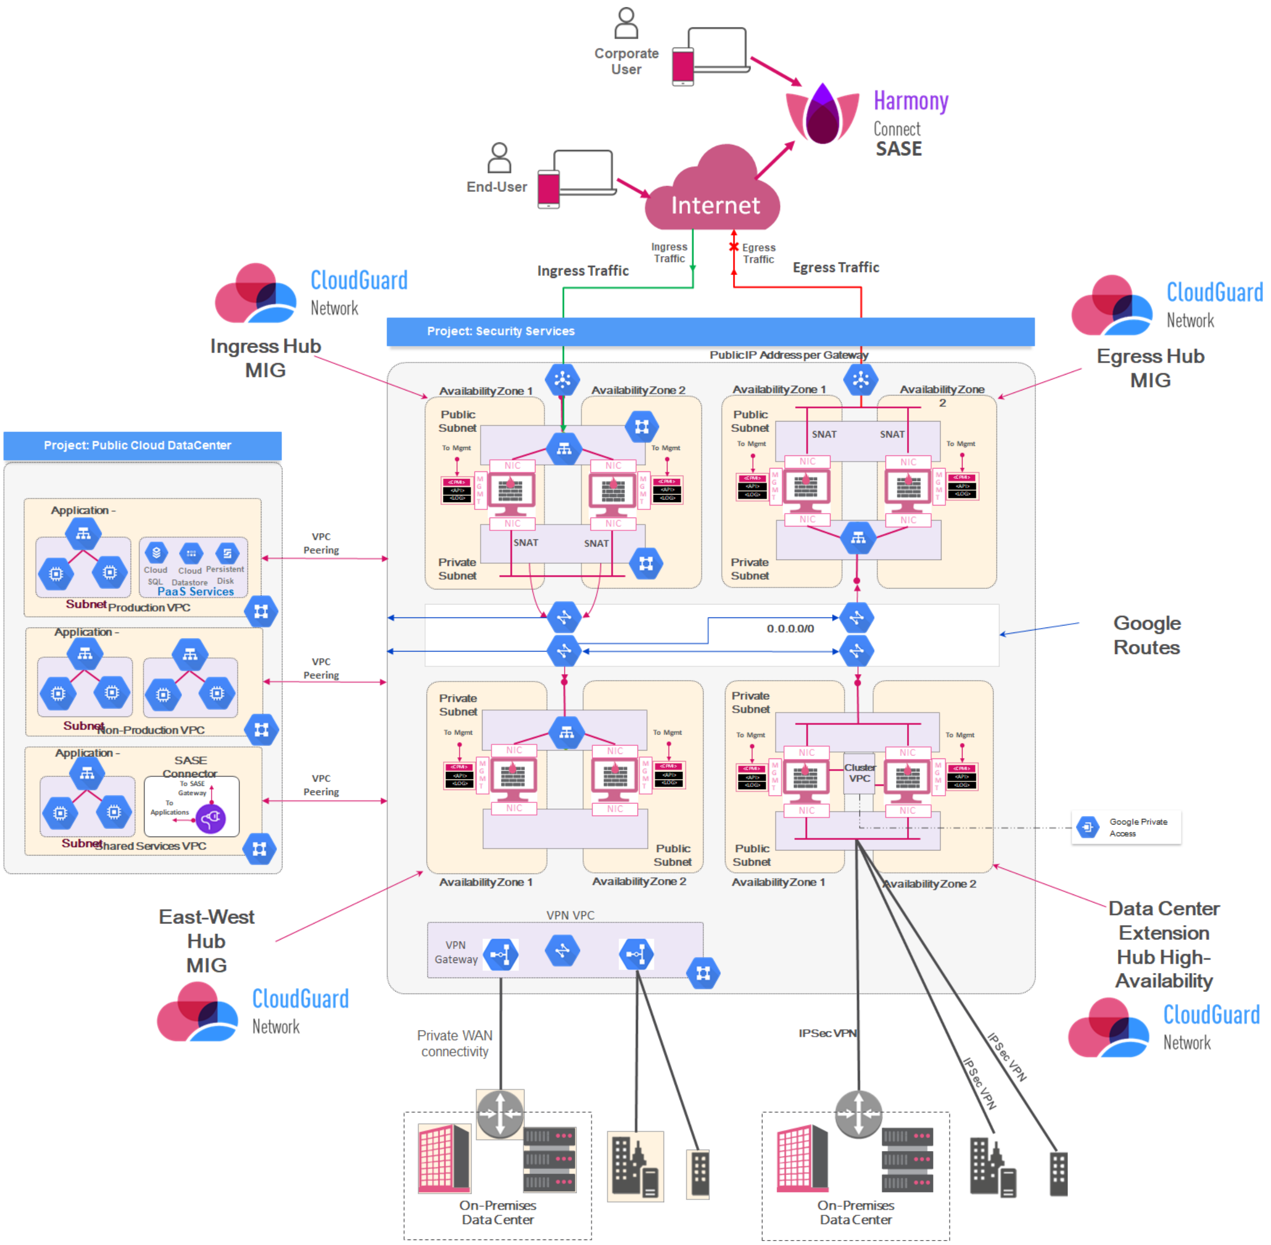

Figure 23: Reference Architecture for Google Cloud Platform

The reference architecture for AWS also has the ingress and egress hubs for public services. This includes the relevant security gateways and their security policies, as well as the east-west, or the data center and backhaul hubs. The TGW is located in the middle of the diagram, displaying the routing domains.

Using this approach, in AWS, we have a different architecture that allows for a highly scalable, flexible, and distributed architecture integrating the gateway load balancers, as can be seen in the following diagram.

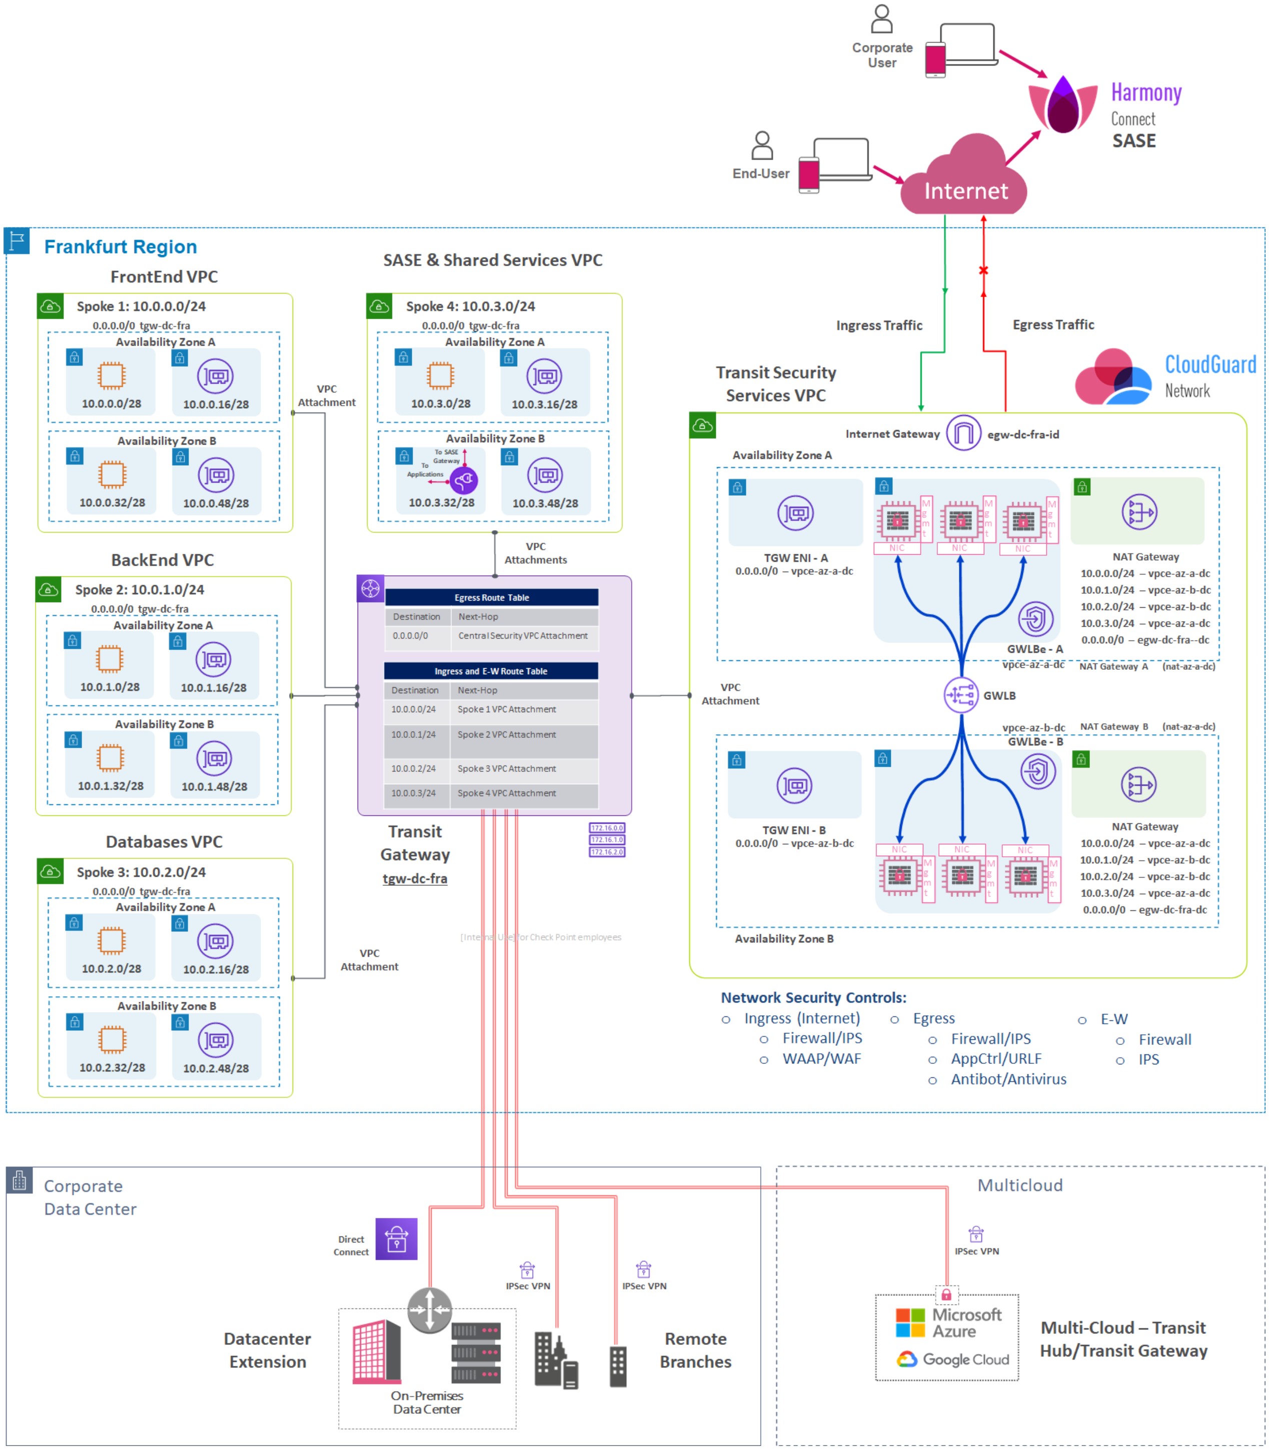

Egress Route Table

Destination	Next-Hop
0.0.0.0/0	Central Security VPC Attachment

Ingress and E-W Route Table

Destination	Next-Hop
10.0.0.0/24	Spoke 1 VPC Attachment
10.0.0.1/24	Spoke 2 VPC Attachment
10.0.0.2/24	Spoke 3 VPC Attachment
10.0.0.3/24	Spoke 4 VPC Attachment

NAT Gateway (Availability Zone A)
10.0.0.0/24 – vpce-az-a-dc
10.0.1.0/24 – vpce-az-b-dc
10.0.2.0/24 – vpce-az-b-dc
10.0.3.0/24 – vpce-az-a-dc
0.0.0.0/0 – egw-dc-fra--dc

NAT Gateway (Availability Zone B)
10.0.0.0/24 – vpce-az-a-dc
10.0.1.0/24 – vpce-az-b-dc
10.0.2.0/24 – vpce-az-b-dc
10.0.3.0/24 – vpce-az-a-dc
0.0.0.0/0 – egw-dc-fra-dc

Network Security Controls:

- Ingress (Internet)
  - Firewall/IPS
  - WAAP/WAF
- Egress
  - Firewall/IPS
  - AppCtrl/URLF
  - Antibot/Antivirus
- E-W
  - Firewall
  - IPS

Figure 24: *Reference Architecture for Amazon Web Services (source: Amazon Web Services)*[21]

[21] Centralized inspection architecture with AWS Gateway Load Balancer and AWS Transit Gateway – URL: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-inspection-architecture-with-aws-gateway-load-balancer-and-aws-transit-gateway/

Ingress Traffic to the Public Cloud

This traffic is related to the production systems and services, with the common scenario being the publishing of web servers. There are three different paths for this purpose:

- Ingress traffic for workloads using the traditional 3-layer approach.
- Ingress traffic for web applications provided by PaaS.
- Ingress traffic for containers acting like web servers, to provide highly dynamic services.

The following diagram is an example for Azure:

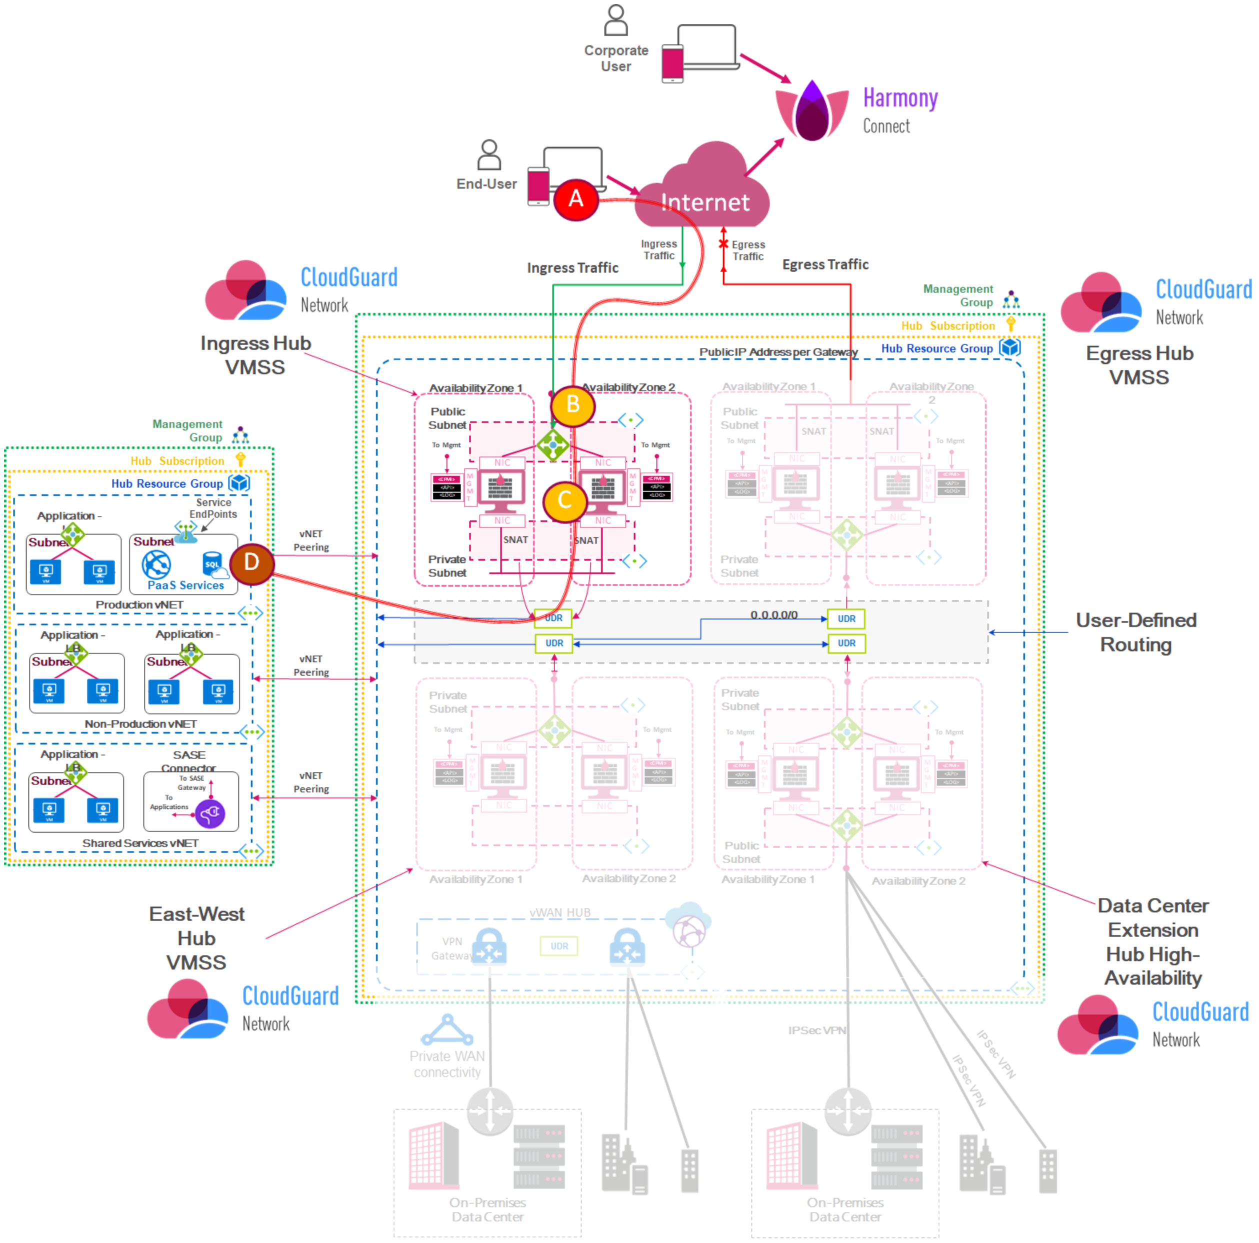

Figure 25: *Ingress Traffic Flow Between Transit Security Services Hub and Production Hubs – Microsoft Azure*

As previously discussed, a Transit Security Services hub is a virtual network (vNET) in Azure acting as a central point of connectivity for all cloud data centers. The following table explains how the ingress traffic works.

FROM	TO
Traffic arriving from the internet ("A")	- Traffic for WebApp1 ("D") is sent to the public IP address ("B") allocated for that web application. - The Azure external load balancer ("C") is set up with an inbound NAT rule that forwards all HTTP (port 80) traffic arriving at that public address to the Check Point gateway's external private address (e.g. 10.0.1.10) on port 8081. - Traffic for WebApp2 ("B") is sent to the public IP address allocated for that web application. - The Azure external load balancer ("C") is set up with an inbound NAT rule that forwards all HTTP (port 80) traffic arriving at that public address to the Check Point gateway's external private address (10.0.1.10) on port 8082. The WebApp2 is located similarly to PaaS service where the service endpoint should be configured. The Check Point security gateway uses SNAT for the following flows (A to D): - From "A" to "D-WebApp1" → forward traffic arriving on TCP port 8081 to Web1 on port 80. - From "A" to "D-WebApp2" → forward traffic arriving on TCP port 8082 to Web2 on port 80.

A. How a scale-in event works

A scale-in event occurs as a result of a decrease in the current load. When a scale-in event triggers, Azure autoscale designates one or more gateways as candidates for termination. The external load balancer then stops forwarding new connections to these gateways and autoscale ends them. The Check Point security management server is then able to detect that these CloudGuard NS security gateways are stopped and automatically deletes the gateways from its database.

Note: We recommend you to have at least two security gateways for redundancy and availability purposes.

B. How a scale-out works

A scale-out event occurs if the current load increases. When a scale-out event is triggered, Azure autoscale launches one or more new instances of the Check Point CloudGuard NS security gateways. The new instances of CloudGuard NS security gateways automatically run the Check Point First Time Configuration Wizard and then reboots.

Throughout the scale-out, the Check Point Security management server detects that new instances of CloudGuard NS security gateways have been launched. The security management server then waits until the CloudGuard NS security gateways finish deploying. Afterwards, the security management server automatically:

- Initializes a secure internal communication (SIC) channel with the CloudGuard NS security gateways.
- Installs a security policy on the CloudGuard NS security gateways.

After a security policy has been installed, the CloudGuard NS security gateways start to respond to health probes, and the load balancer starts to forward new connections to them. Finally, the newly created CloudGuard NS security gateways report their status and send logs to the Check Point security management server.

We can observe a similar approach in the Google Cloud Platform using the autoscaling managed instance group (MIG). The MIG is a GCP compute engine resource that is a collection of VM instances managed as a single entity. Here an external load balancer sends incoming traffic to a Check Point autoscaling MIG residing on the external VPC. The gateways in the group can then inspect the traffic and, if allowed by policy, forward the traffic to an internal load balancer. The internal load balancer then sends incoming traffic to a group of servers residing on an even further internal network. Finally, GCP autoscale is configured to increase or decrease the number of Check Point CloudGuard security gateways in the MIG.

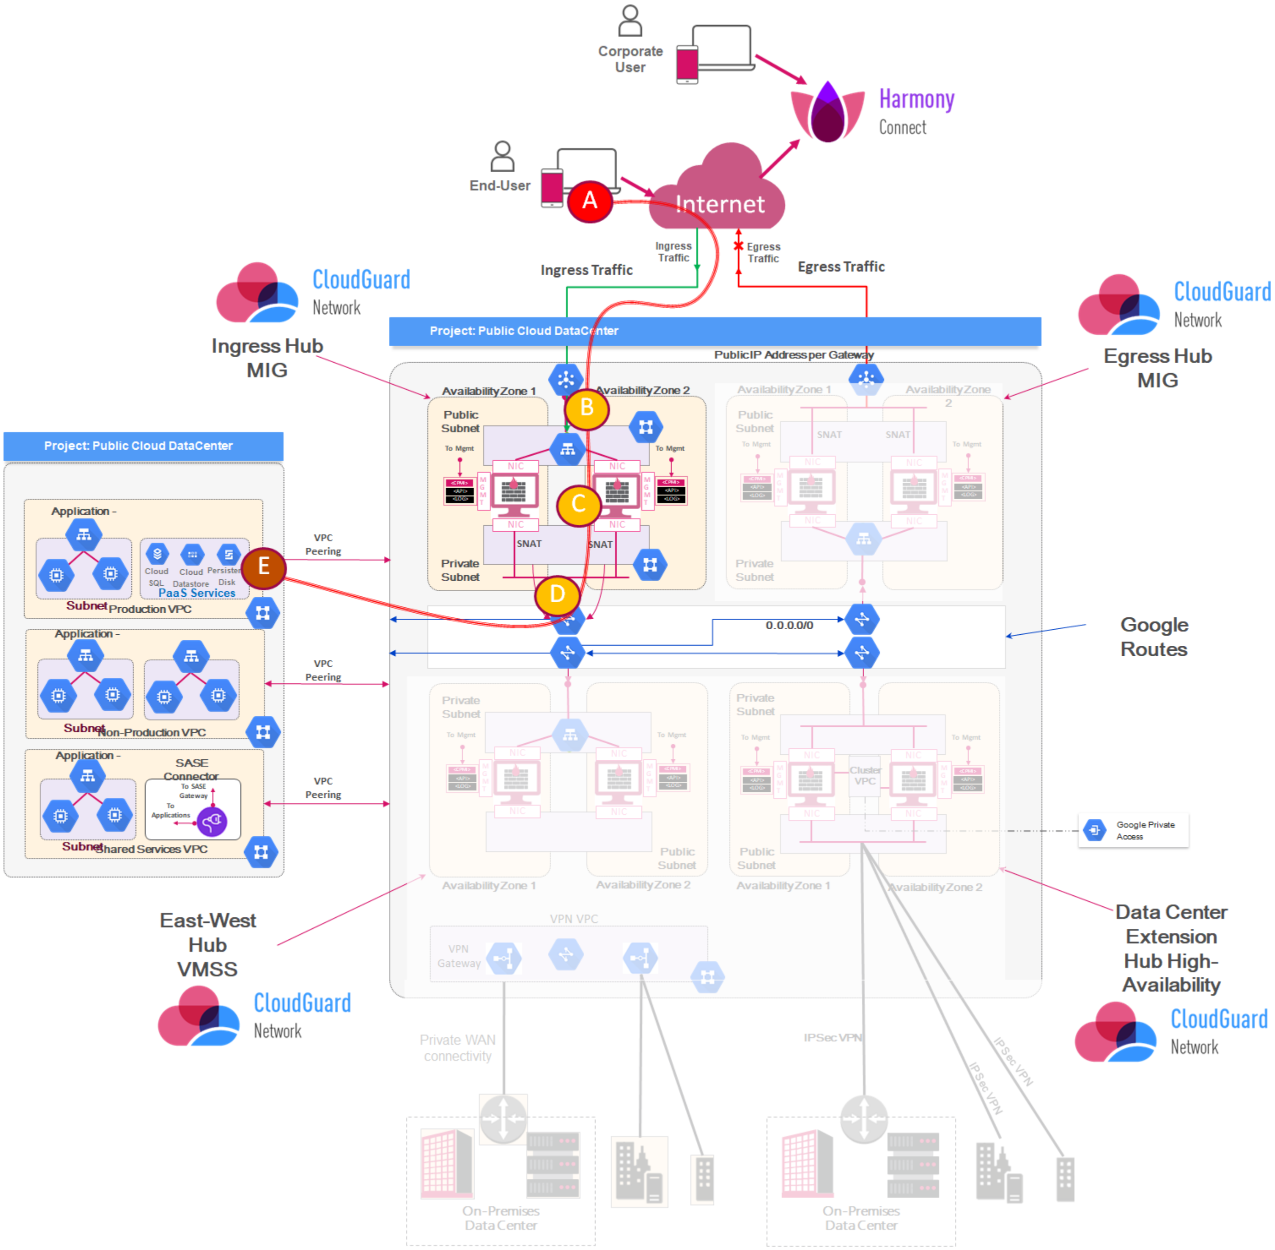

Figure 26: Ingress Traffic Flow Between Transit Security Services Hub and Production Hubs – Google Cloud Platform

Amazon Web Services has a different approach to deploying the GWLB:

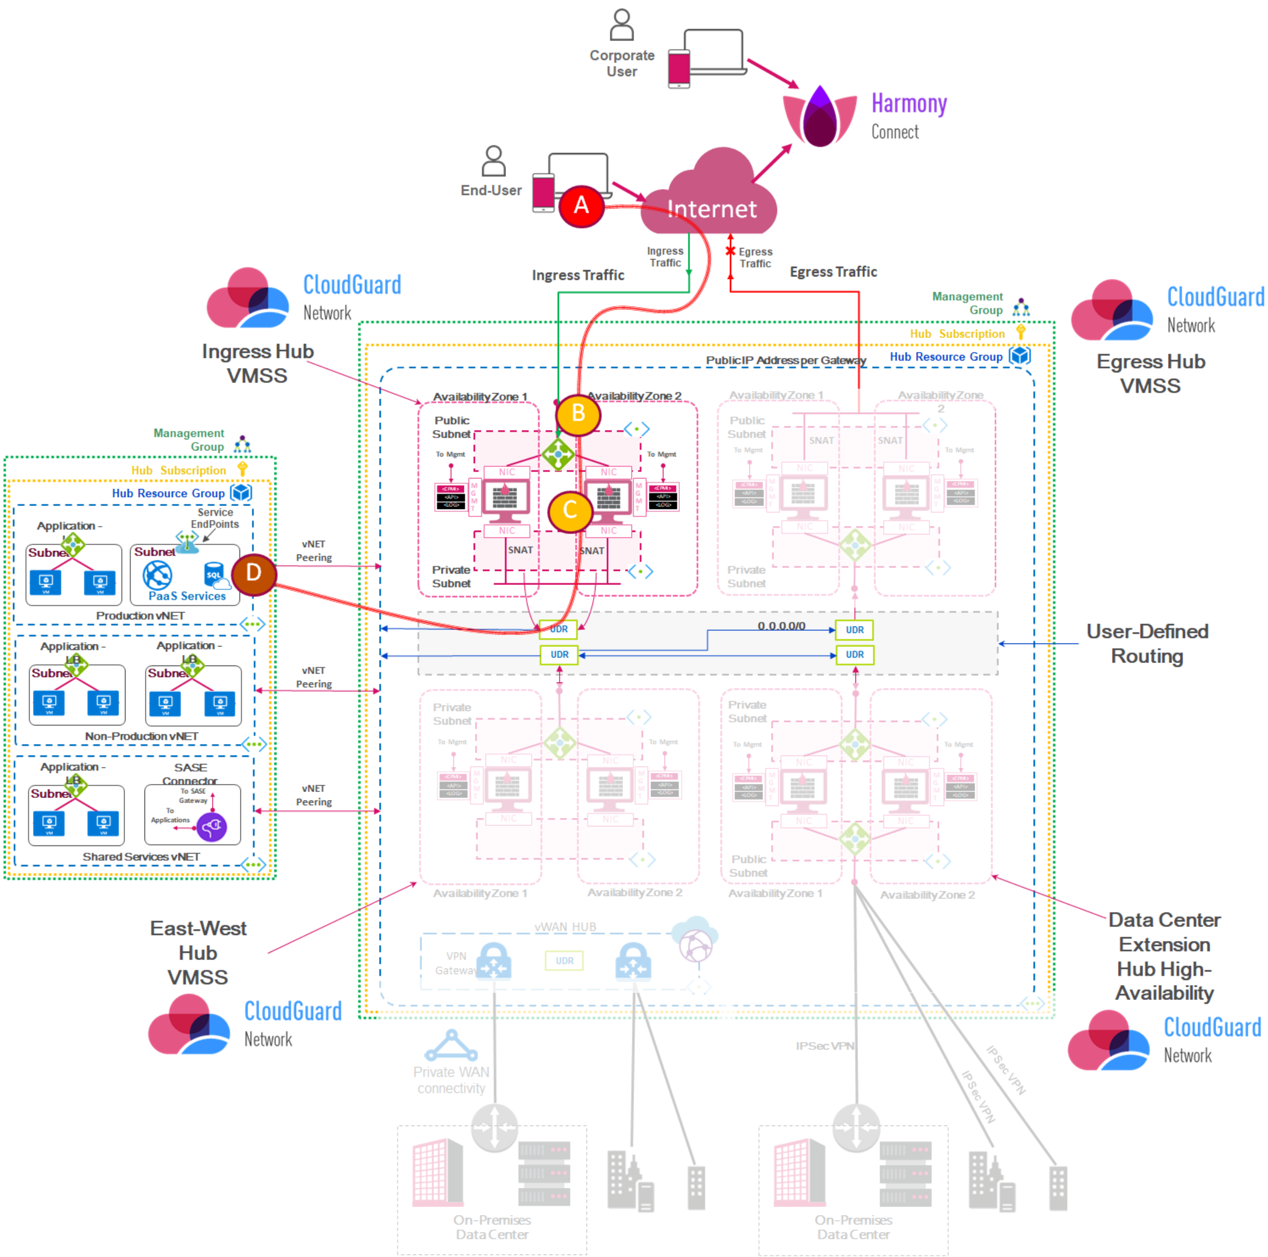

Figure 27: *North-South Traffic Inspection With Centralized Appliance VPC: Use Case for Ingress Traffic*

The above traffic flow can be described using information found in a blog published by AWS, in relation to the centralized inspection architecture with R80.40 or R81 security gateways:

- **Flow A:** End-user traffic arrives at the internet gateway. The NAT Gateway is translated by the source IP address. The internet gateway routes the traffic back to NAT gateway A.
- **Flow B:** NAT gateway A uses frontend VPC's network address route in NAT gateway route table A and sends traffic to GWLBE (gateway load balancer endpoint) A.
- **Flow C:** GWLBE A, using AWS PrivateLink, routes traffic to GWLB (gateway load balancer) securely over the Amazon network.
- **Flow D:** Since this return packet is associated with an existing flow, GWLB encapsulates the original IP traffic with a GENEVE header and forwards it to the R80.40 or R81 security gateway chosen for this flow.
- **Flow E:** The R80.40 or R81 security gateway behind the GWLB decapsulates the GENEVE header, inspects the traffic, and decides how to handle the traffic depending on the security policy configured.

  - In R80.x, we can enable rule-based traffic inspection for IPS, Antibot, and Antivirus. Under this approach, we can create policies according to the flows and select the inspection needs for different applications more accurately.
  - The addition of GENEVE headers doesn't count towards the overall MTU limit of GWLB.
- **Flow F:** Assuming the traffic is allowed and inspected by the security policy, the R80.40 or R81 security gateway re-encapsulates with GENEVE headers and forwards the traffic to the GWLB.
- **Flow G:** GWLB, based on GENEVE TLV, selects GWLBE A, removes GENEVE header, and forwards traffic to GWLBE A.
- **Flow H:** GWLBE A uses a frontend VPC's network address route in the appliance route table A and routes the transit gateway traffic.
- **Flow I:** Since R80.40 or R81 security gateway VPC is associated with the transit route table, transit gateway uses the frontend VPC's network address route in the transit route table to send traffic to the frontend VPC.
- **Flow J:** Finally, once the traffic is at the frontend VPC, the packet's destination is within the VPC CIDR range, where the local route is used to deliver traffic to the application instance that sourced the traffic.

The flows described above allow the security gateways to maintain the connections without SNAT (Source NAT). GWLB's ability to use 5-tuples or 3-tuples of an IP packet to select a specific appliance behind it for the life of that flow combined with transit gateway appliance mode, also provides session stickiness irrespective of the source and destination AZ. This includes the AZs that the transit gateway attachments and GWLB are deployed in – while still providing autoscaling and automatic health checks.

Security policies should be focused per-flow, which means it's important to define the protections according to the referred workloads and the policies with the source, destination, services, and relevant signatures, for the IPS. For example, a workload using a Windows/Linux server and a Microsoft IIS/Apache server, should be protected only by specific access-control, allowing the HTTP/HTTPS traffic and enabling the traffic inspection relevant signatures for the workload.

Using the traffic flow logic, the following diagram displays how a simple definition of the security policy can protect the web servers.

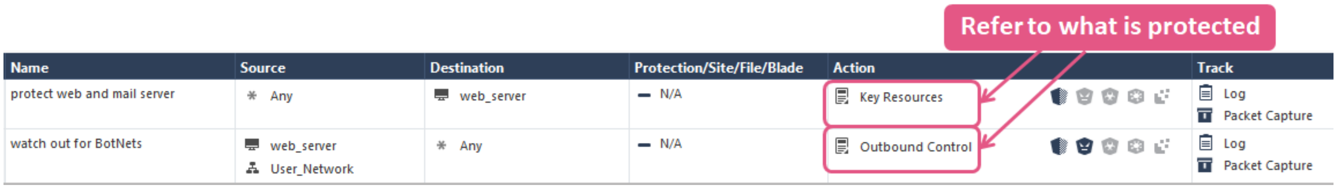

Refer to what is protected

Name	Source	Destination	Protection/Site/File/Blade	Action	Track
protect web and mail server	Any	web_server	N/A	Key Resources	Log Packet Capture
watch out for BotNets	web_server User_Network	Any	N/A	Outbound Control	Log Packet Capture

Figure 28: Sample Policy for Ingress/Egress Controls Related to the Traffic Flows.

Egress Traffic

Egress traffic is commonly used to allow access to the computing instances in a secure way that requires the download of patches, service packs, or upgrades. This traffic shouldn’t be focused on end-users, and only used for maintenance purposes. It is also important to clarify that this traffic shouldn’t be confused with the "traffic return" in the ingress perimeters, as both flows are entirely different. Here, egress filtering controls the traffic attempting to leave the vNET or VPC through the transit security hub or transit gateway. Therefore, the filtering is essential and focused on providing secure internet access to the computing instances located in the vNET or VPC and prevents outbound connections to dangerous and unwanted hosts by:

Disrupting malware

- Suppose one of the computing instances (virtual machines) located in the subnets on the vNET/VPC is infected with malware. In that case, egress threat prevention policies can prevent it from connecting to the malware's command server. If the malware tries to export the machine's data, the egress filter can prevent it from connecting to the destination.

Blocking unwanted services

- In cloud data centers, egress traffic is not focused on corporate users. As a consequence, users aren't allowed to browse the internet or chat using social network sites. An egress security policy can also block the ports, protocols, applications, and URLs used to allow specific services for maintenance purposes or any other sites they cannot access.

Stopping contribution to attacks

- Egress threat prevention is also good at protecting the organization's reputation. Blocking certain traffic types prevents the computing instances in the vNET/VPC from being used for DDoS attacks, malware hosting, spamming, and botnets.

The egress threat prevention solution offers a multi-layered, pre- and post-infection defense approach and a consolidated platform that enables enterprise security to detect and block modern malware:

1. Anti-Bot - Post-infection detection of bots on hosts. Prevents bot damages by blocking bot C&C (command and control) communications. The Anti-Bot Software Blade is continuously updated by ThreatCloud, a collaborative network to fight cybercrime. The Anti-Bot is able to discover infections by using multiple detection methods.
2. Anti-Virus - Pre-infection detection and blocking of malware at the gateway. The Anti-Virus Software Blade is continuously updated by ThreatCloud. It detects and blocks malware by using multiple detection engines before users are affected.
3. SandBlast - Protection against infections from undiscovered exploits, zero-day and targeted attacks:
    a. Threat Emulation - This innovative solution quickly inspects files and runs them in a virtual sandbox to discover malicious behavior. The Malware is then prevented from entering the network. The ThreatCloud Emulation service reports to the ThreatCloud and automatically shares the newly identified threat information with the Check Point network.
    b. Threat Extraction - Protection against incoming malicious content. The Threat Extraction Blade capability removes exploitable content, including dynamic content and embedded objects, reconstructs files to eliminate potential threats, and promptly delivers sanitized content to users to allow for business to continue uninterrupted. The Threat Extraction Blade also creates a safe copy of the file to remove possible threats, while the Threat Emulation Software Blade inspects the original file for potential threats.

In practical scenarios, we can deploy computing instances or containers to allow reverse-proxy connectivity to enable SASE solutions, which allow clientless remote access to access corporate applications. On the other hand, VDI desktops can also utilize this use case to allow users access to the internet (if the acceptable usage policy allows it). It is highly advisable to use rule-based threat prevention (IPS, Anti-Bot, Anti-Virus) and rule-based access control (application control and URLF) for specific sources.

The following is an example of Egress traffic flow for Microsoft Azure:

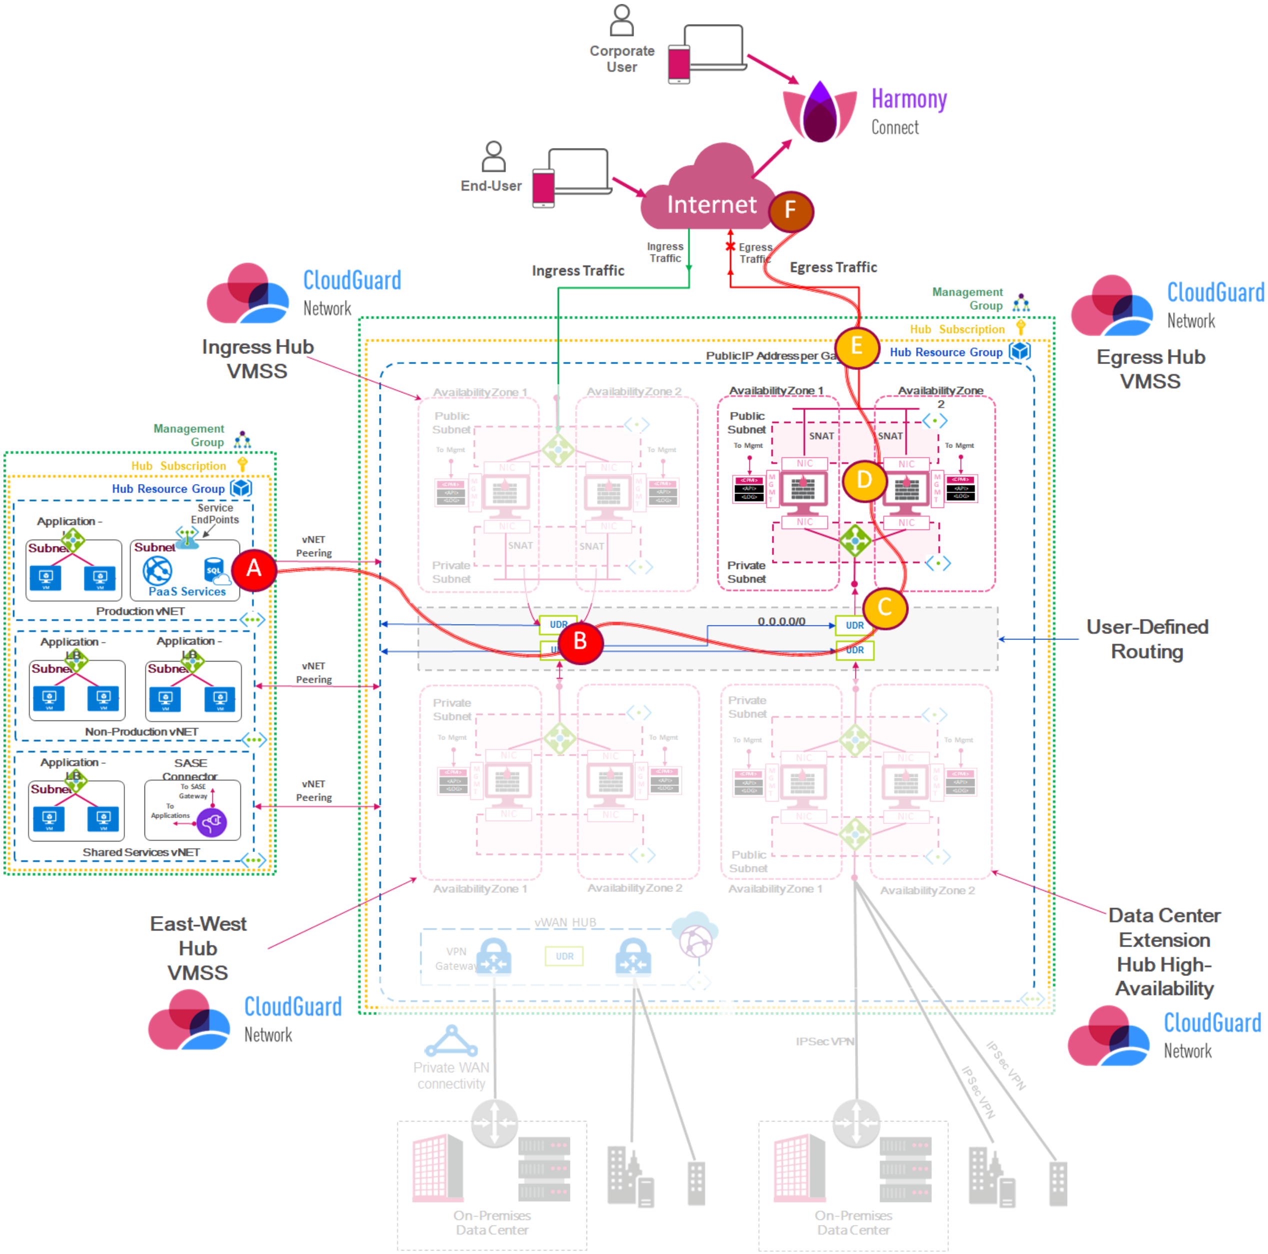

Figure 29: *Egress Traffic for Computing Instances – Microsoft Azure*

FROM	TO
Traffic from the backend subnets to the internet (from "A" and "D")	- Internal traffic ("A") is routed through the Check Point gateway using user defined routes (UDR) pointing to the internal load balancer ("B"). The traffic is then forwarded to the security gateway to process and inspect it using rule-based threat prevention policies ("C"). - The gateway uses SNAT ("D") to hide this type of traffic behind its external private address (10.0.1.10). As the traffic leaves the virtual network, Azure replaces the private address with the gateway's public address to access the internet ("E").

In a similar approach, we have this use case for egress on Google Cloud platform,

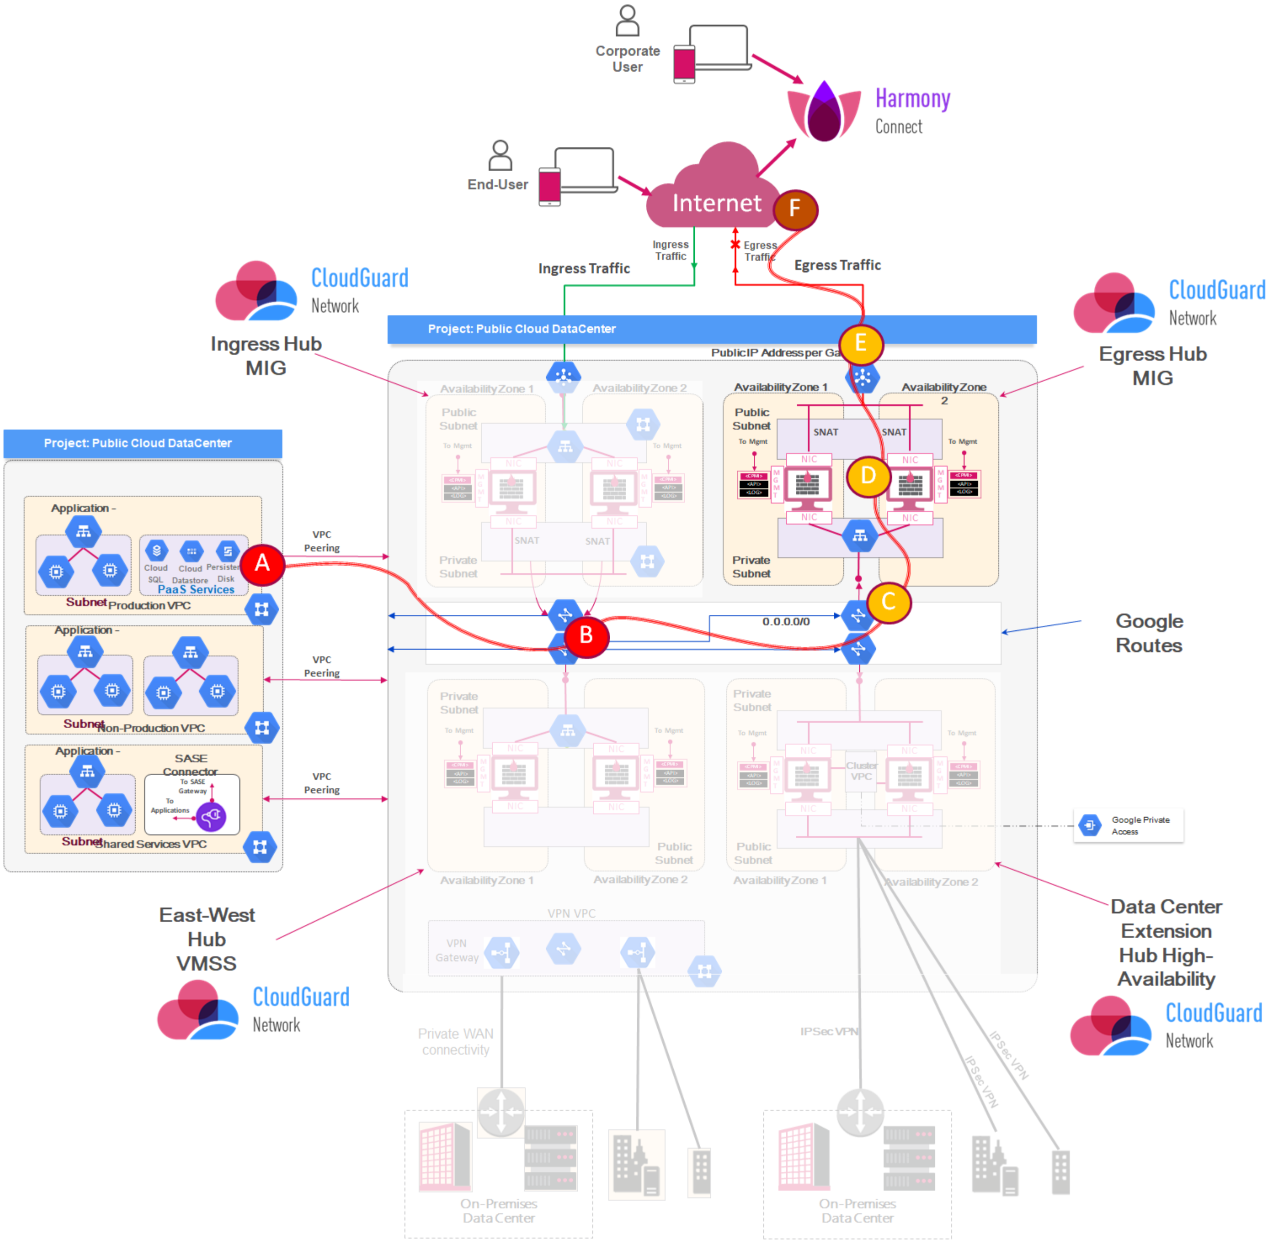

Figure 30: Egress Traffic for Computing Instances – Google Cloud Platform

Here we can see how the route-based load balancing distributes outbound connection requests between the Check Point CloudGuard security gateways instances. The CloudGuard security gateways instance then receives the request, inspects it, and, if allowed, it forward out to the internet.

A computing instance (flow “A”) tries to connect to the internet, using (flow “B”) multiple routes to 0.0.0.0/0 with next-hop (flow “C”) to the security gateway in the managed instance group (flow “D”). However, the Google Cloud Platform doesn’t allow the setting of a load balancer as the next hop.

As a result a separate route is required for each security gateway in the MIG. However, when specifying multiple routes to the same 0.0.0.0/0 with a different next-hop (different security gateway instances in the MIG), Google Cloud Platform will automatically perform load balancing on this traffic, and provide redundancy load distribution capabilities for outbound traffic inspection.

As Google Cloud Platform does not support next-hop to a load balancer, you must specify a route per security gateway instance. We therefore recommend that you turn off autoscaling on the outbound MIG. This is to avoid scenarios of outbound routes becoming invalid (flow “E”) when the MIG autoscaler automatically terminates the security gateway instance defined as the next hop. CloudGuard security gateways can also be set up to perform deep packet inspection of encrypted HTTPS traffic using the HTTPS Inspection feature. With this feature enabled, the web clients should be set up to trust a CA certificate issued by the Check Point Security Management Server during the HTTPS inspection configuration.

On other hand, in Amazon Web Services we use a different approach with the GWLB for egress traffic:

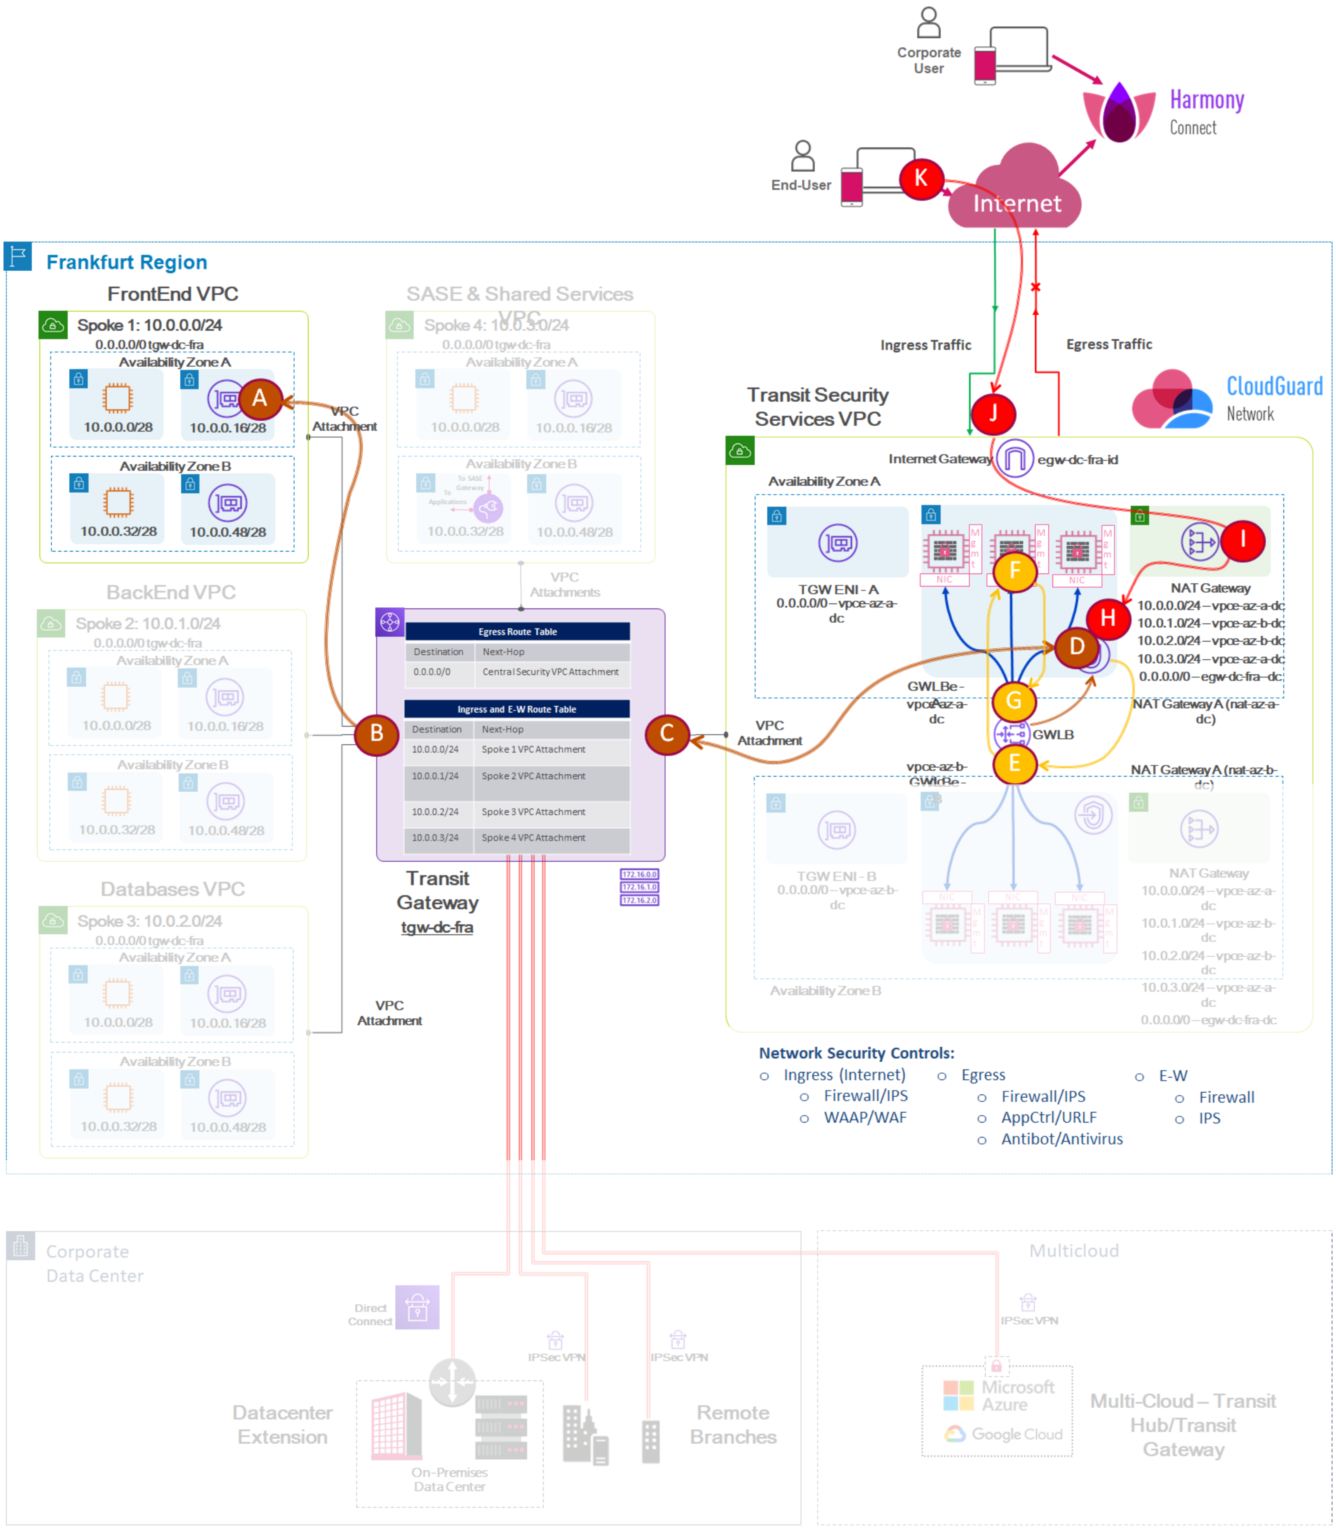

Egress Route Table	
Destination	Next-Hop
0.0.0.0/0	Central Security VPC Attachment

Ingress and E-W Route Table	
Destination	Next-Hop
10.0.0.0/24	Spoke 1 VPC Attachment
10.0.0.1/24	Spoke 2 VPC Attachment
10.0.0.2/24	Spoke 3 VPC Attachment
10.0.0.3/24	Spoke 4 VPC Attachment

Network Security Controls:

- Ingress (Internet)
  - Firewall/IPS
  - WAAP/WAF
- Egress
  - Firewall/IPS
  - AppCtrl/URLF
  - Antibot/Antivirus
- E-W
  - Firewall
  - IPS

***Figure 31**: North-South Traffic Inspection With Centralized Appliance VPC: Use Case for Egress Traffic*

The above traffic flow can be described using information found in a blog published by AWS, in relation to the centralized inspection architecture with R80.40 or R81 security gateways:

- **Flow A**: An application in the frontend VPC wants to access the internet, typically downloading patches or fixes. The application uses the default route (0.0.0.0/0) in the frontend VPC route table A to send traffic to the transit gateway.
- **Flow B**: Since the fronend VPC is associated with the egress route table, the transit gateway uses a default route in the egress route table to send traffic to the transit security services VPC.
- **Flow C**: In the transit security services VPC, the transit gateway subnet A uses the default route in the transit gateway route table A to send traffic to GWLBE A, which is in the same availability zone (AZ).
- **Flow D**: GWLBE A, using AWS PrivateLink, routes traffic to GWLB. Traffic is routed securely over the Amazon network without any additional configuration.
- **Flow E**: GWLB uses 5-tuples or 3-tuples of an IP packet to pick an appliance for the life of that flow. This creates session stickiness to an appliance for the life of a flow required for R80.40/R81 security gateways.
  - GWLB encapsulates the original IP traffic with a GENEVE[22] header and forwards it to the appliance over UDP port 6081.
  - This encapsulation allows all IP traffic to be delivered to the R80.40/R81 security gateways for inspection, without specifying listeners for every port and protocol.
- **Flow F**: R80.40/R81 security gateway behind the GWLB decapsulates the GENEVE header and decides to allow the traffic based on the security policy configured. For egress traffic the relevant blades that should be enabled are: Anti-virus, Anti-bot, and Sandblast.
- **Flow G**: The R80.40/R81 security gateway then re-encapsulates the traffic and forwards it to the GWLB.
- **Flow H**: GWLB, based on the GENEVE type-length-value (TLV)[23], selects GWLBE A, removes the GENEVE header, and forwards the traffic to GWLBE A.
- **Flow I**: GWLBE A uses the default route in appliance route table A and routes traffic to NAT gateway A.
- **Flow J**: NAT gateway A uses the default route in NAT gateway route table A, performs source IP address translation, and routes traffic to the internet gateway (igw-id). From there, traffic egresses out to the internet.

East-West Traffic in the Public Cloud

East-west traffic is related to the communication between the production systems under the 3-tier architecture of applications (web servers, application servers, and databases). Traffic can originate from a virtual machine, PaaS service, or container, and is destined to another virtual machine, PaaS service or container.

For east-west traffic, there are two different scenarios:

- Access-control.
- Threat prevention.

In some scenarios, the workloads can communicate to each other without NAT policies due to all traffic being generated by private subnets between them. For that approach, the security groups provide the access-control capabilities for basic micro-segmentation to limit the communications between workloads.

CloudGuard enhances the native micro-segmentation and elastic networking of cloud environments to dynamically deliver advanced security and consistent policy enforcement that automatically grows and scales with your cloud environments. CloudGuard can easily secure workloads and applications running in hybrid and public cloud environments, thus mitigating risks from breaches, data leakage, and zero-day threats.

Security groups provide strong protection and allow for micro-segmentation, as they can be associated with a cloud object's network interface. However, they are limited to traditional layer 3-4 access control and are not suitable for "threat prevention". Security groups also cannot recognize if a forbidden protocol is being tunneled inside of an allowed one, or if traffic is truly clean SQL or something malicious (such as an SQL injection attack) riding the known SQL port. Moreover, security groups can’t allow or prevent access based on assigned metadata tags with values like "database servers" or "web servers."

Check Point, therefore, typically recommends internal segmentation and the use of "threat detection/prevention" solutions at the inside perimeter edge that can monitor traffic moving between the internal network segments. This scenario should be used if the organization requires specific compliance requirements to have visibility and prevention in the network. Such

[22] Geneve: Generic Network Virtualization Encapsulation – URL: https://tools.ietf.org/html/rfc8926

[23] TLV vs Bit Fields – URL: https://tools.ietf.org/html/draft-ietf-nvo3-encap-05#section-6.6

a security solution can reside inside each host/instance and/or it can be inserted at the network level, considering the premises of the transit hub, transit gateway, or shared VPC.

Important definitions to understand before deploying east-west access control or traffic inspections:

- Macro-segmentation: Generally refers to traffic leaving your network or zone.
- Micro-segmentation: Generally refers to traffic flowing between resources within your network or zone.

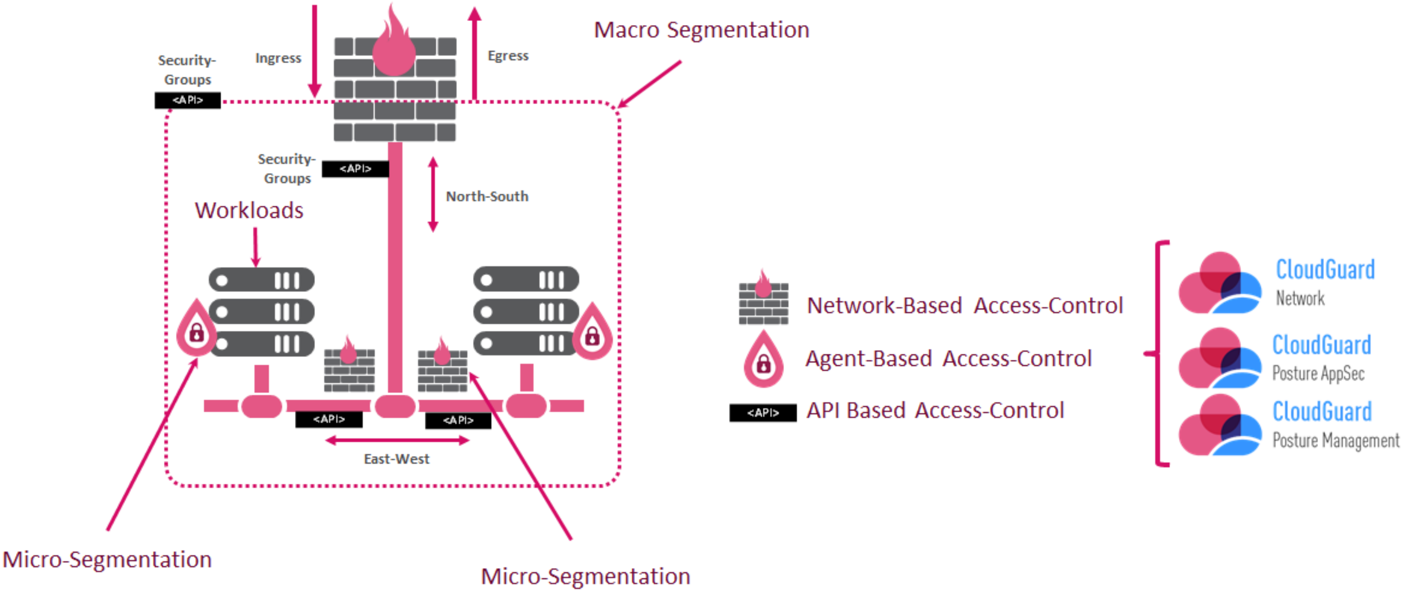

***Figure 32**: Macro-Segmentation and Micro-Segmentation Principles for East-West Access-Control and Traffic Inspection*

East-west, or lateral movement security, is a pillar of the Zero Trust model. Zero Trust begins with the premise that a breach has already occurred (or that an internal threat is already present). Under such circumstances, internal perimeters are a must, because there is no possibility of an internal "trusted" zone. An east-west security policy that lacks deep packet threat inspection (versus access control only) can leave internal networks wide open to lateral propagation of malware after an initial breach.

The following diagram is an example of east-west traffic inspection for Microsoft Azure:

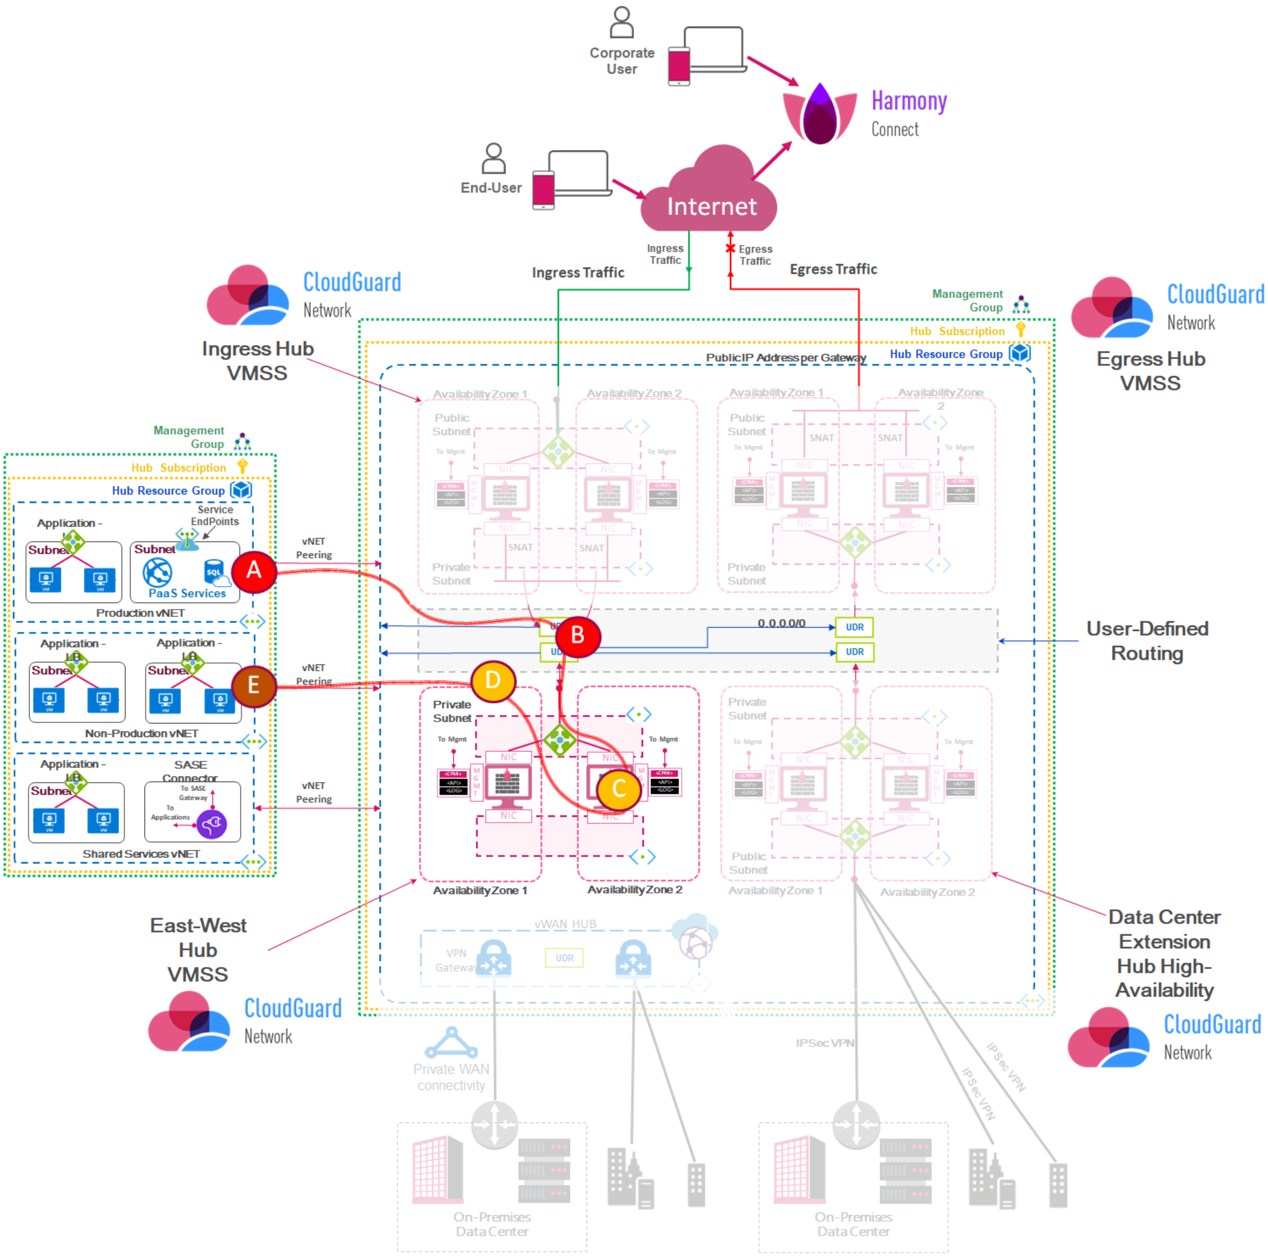

Figure 33: Threat Prevention for East-West Traffic Inspection – Azure

FROM	TO
Traffic between the web and application tiers	- This traffic originates from computing instance ("A") and is routed through the Check Point security gateway (VMSS) by the internal load balancer ("B"). Then, the traffic is inspected according to the flows through ("C") routed through of the User Defined Routes (UDR) to access the destination computing instance (Point "D"). - Note: this scenario is focused on providing east-west traffic inspection when the cloud data center has hundreds or thousands of virtual machines. - For his particular scenario, we do not recommended to mix the east-west traffic with egress traffic, as traffic inspection should be done according to security needs. Such needs should be aligned with the process of vulnerability and patch management, as well as with regulatory requirements. - For trusted sources, such as database replicas, cloud-native networking security policies should be used instead of the traffic inspection unless a regulatory requirement needs to be met.

In Google Cloud Platform, we have a similar approach. The following is the flow analysis:

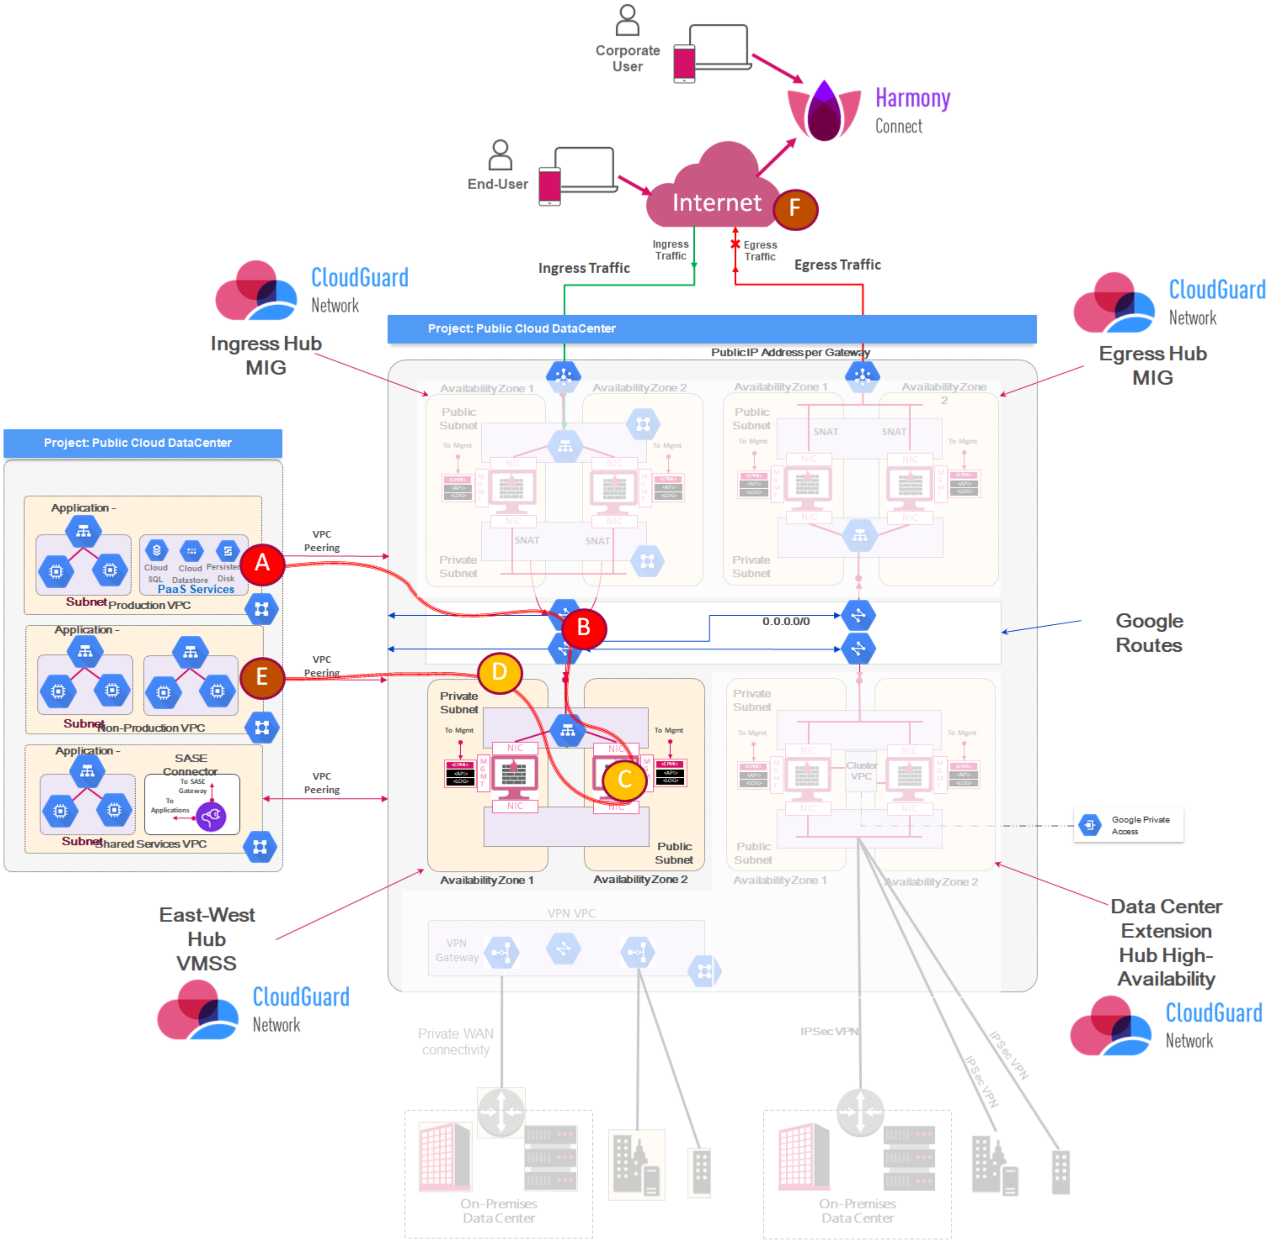

***Figure 34**: Threat Prevention for East-West Traffic Inspection – Google Cloud Platform*

Similarly to the use case described in Azure, in Google Cloud Platform the traffic that originates from computing instance ("A") placed in production VPC, is routed through the Check Point security gateway (MIG) by internal load balancer ("B"). Afterwards the traffic is inspected according to the flows through ("C") the use of Google Routes (UDR) to access the destination computing instance ("D"). It is important to note that this scenario does not need SNAT for the traffic inspection between VPC's.

In Amazon Web Services we have a different scenario as illustrated in the following diagram:

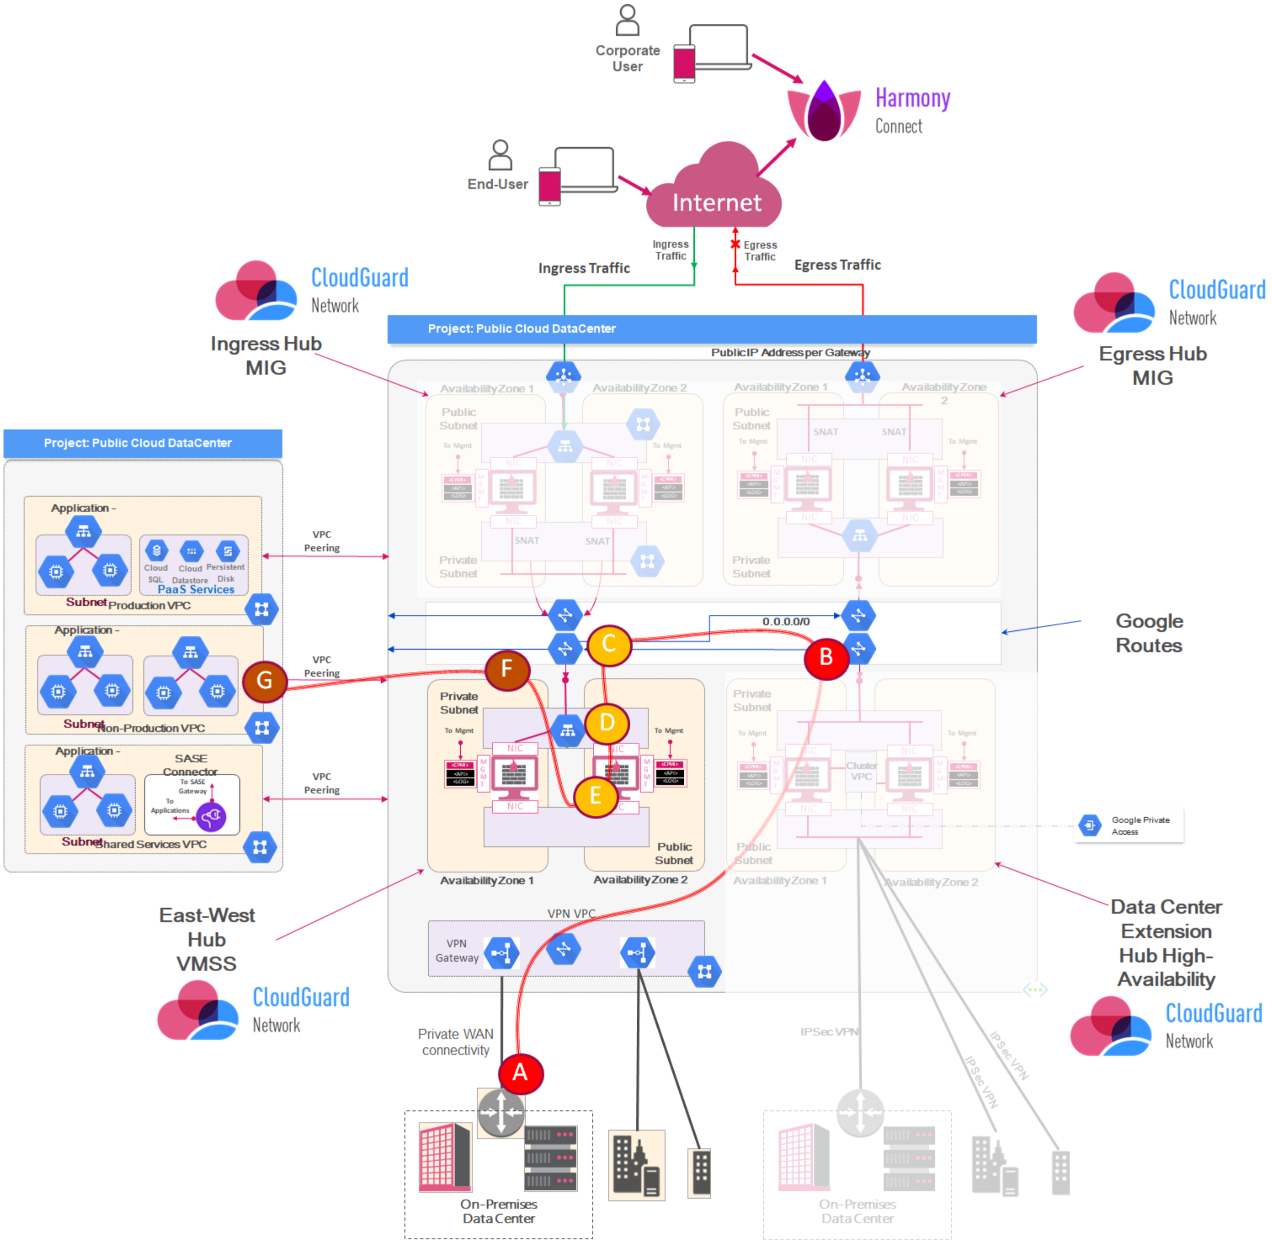

Figure 35: Threat Prevention for East-West Traffic Inspection – Google Cloud Platform

In this scenario, for Amazon Web Services to provide east-west traffic inspection between the production VPCs, we can use the transit gateway in appliance mode to forward the traffic to the GWLB and be inspected by the security gateways fleet.

This new capability will enhance the requirements to inspect the east-west traffic instead to do Access Control only, simplifying the design of cloud architectures and allowing organizations to design fault-tolerant architectures in an easier and more intuitive way, specifically when adding multiple virtual appliances. In the east-west VPC (the transit security services VPC), where all security traffic inspection can be placed, traffic is load balanced in security gateway fleet allowing for a very flexible and scalable model for hyperscale data centers. It is important to note that the security gateway fleet can be placed in different availability zones and can be attached to different transit security gateways for global deployments. An example of this is data centers interchanging traffic between regions such the Americas, Europe, and Asia.

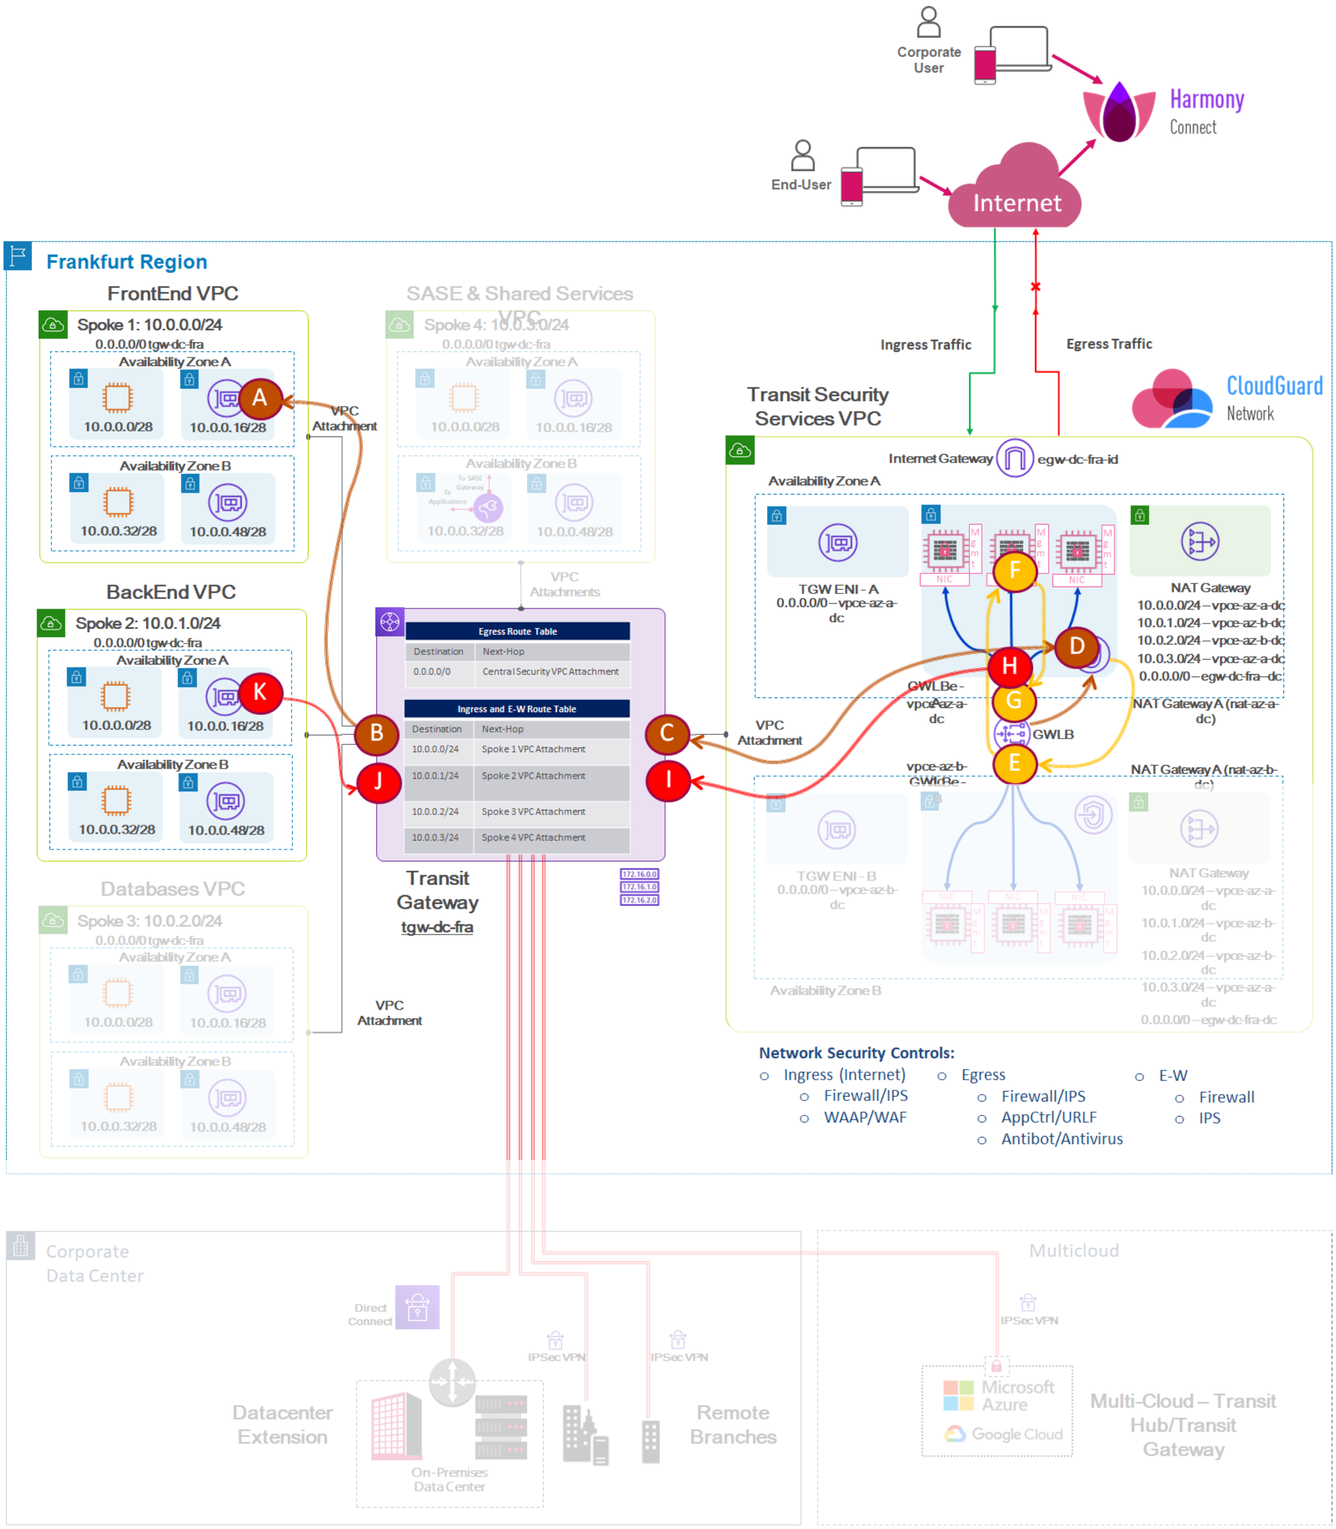

Egress Route Table

Destination	Next-Hop
0.0.0.0/0	Central Security VPC Attachment

Ingress and E-W Route Table

Destination	Next-Hop
10.0.0.0/24	Spoke 1 VPC Attachment
10.0.0.1/24	Spoke 2 VPC Attachment
10.0.0.2/24	Spoke 3 VPC Attachment
10.0.0.3/24	Spoke 4 VPC Attachment

***Figure 36**: Threat Prevention for East-West Traffic Inspection – Amazon Web Services*

In the above flow, the traffic originates from an instance in spoke1 VPC and is placed in availability zone A. The destination (computing instance, PaaS or CaaS) can be located in availability zone B of spoke 2 VPC. The scalable fleet of security gateways can also be located in availability zones A and B, which are found in the transit security services VPC. A key element in the architecture is the configuration of the routing of the transit gateway using the VPC attachment, spoke 1, spoke 2, spoke 3, and spoke 4 VPCs, which are attached through availability zone A and availability zone B.

Prior to the transit gateway appliance mode, when traffic is routed between VPC attachments, the transit gateway will keep the traffic in the same availability zone as it originated, until it reaches its destination. Traffic then crosses availability zones between attachments only if there is an availability zone failure or if there are no subnets associated with a VPC attachment in that availability zone.

This architecture also considers the appliance mode enabled in the transit gateway, which allows it to use 4-tuples of an IP packet, and then select a single transit gateway ENI in the transit security services VPC for the life of a flow to send traffic to. Once at the transit gateway ENI, traffic is routed to the GWLBE (endpoint) and then on to the GWLB, in the same availability zone, providing stickiness to all flows.

For return traffic, transit gateway ensures symmetry by using the same selected transit gateway ENI. This ensures bi-directional flow being processed by the same security gateway behind the GWLB irrespective of the availability zones of the three elements: the source, the destination, and the security gateways. Such a characteristic enables you to avoid source/destination NAT in configurations by removing complexity during deployments.

Data Center On-Premises Traffic to the Cloud

Another important use case is to allow communications to the data center and the cloud. This is a critical component due to organizations' need to enable replicas in the databases and application servers. In the following diagram, we can see different scenarios in the backhaul communications, and the different VPCs. For example, if we take a look at the flow of A to B, we will see that this is the most common scenario where replicas are originated from the on-premises data center. Using ExpressRoute or VPN site-to-site communicated in the data center "hub", this cluster of security gateways provides the capability for access-control, however the value here comes from the rule-based traffic inspection. Considering the traffic flow, the IPS provides the capability of virtual patch-management and protects the communication between the cloud data center and the on-premises data center.

For the data center we have two scenarios:

1) Cloud-native VPN communication without security gateways, east-west hub or shared services VPC, is used.

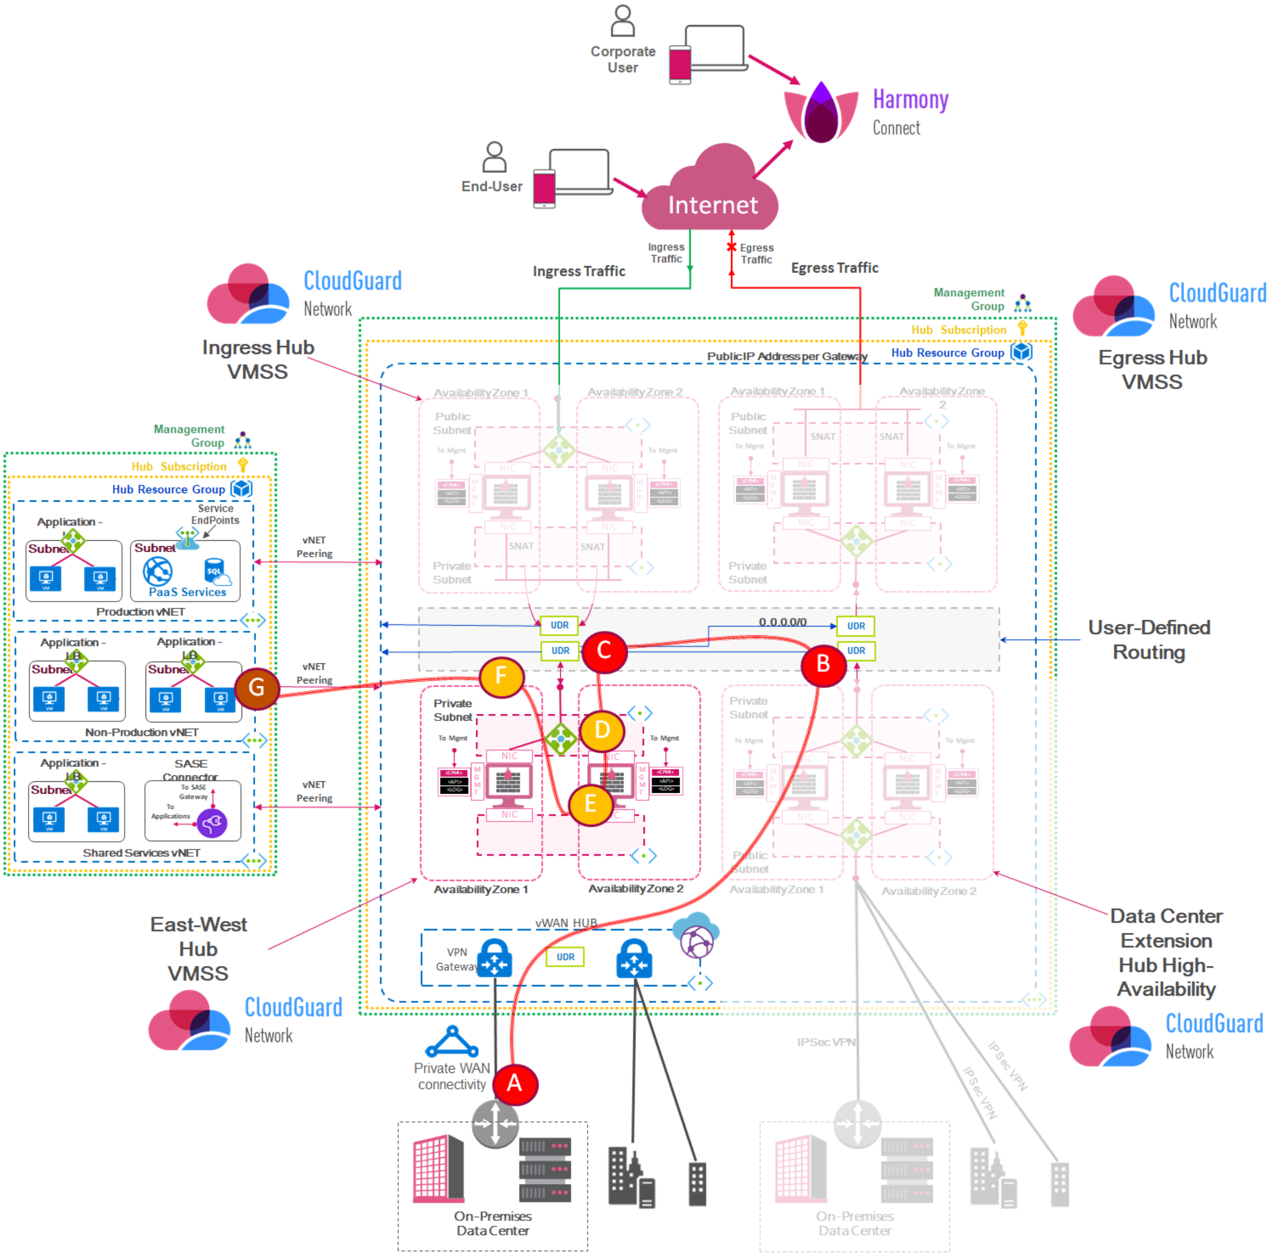

Figure 37: *Data Center Communications to Public Cloud IaaS Using Cloud-Native VPN Gateways - Microsoft Azure*

Azure service endpoints allow Azure customers to keep data private, even if it is hosted in PaaS like Cosmos, SQL, etc. They allow VMs in the customer vNET (private address space) to access supported Azure PaaS services without traversing the internet. Service endpoints are usually enabled on the spoke vNET, and subnets close to the computer VMs. However, this means the traffic does not pass through Check Point.

CloudGuard NS which might be used to accomplish requirements for compliance, regulations, visibility, inspection of traffic, or controlling access to specific PaaS services from specific VMs. This can be achieved by deploying a Check Point CloudGuard NS solution in the environment, be it a single vNET, or a more complex hub-and-spoke topology. If service endpoints are already in use by the VMs, they will need to be disabled in their current location in the event of a cutover. The procedure is almost entirely based around Azure, with the only steps for Check Point being to add a rule accepting the traffic, and to verify that Hide NAT is configured.

FROM	TO
On-premises data center (backhaul)	- Encrypted IPsec traffic is sent to the gateway's public IP address ("A"). - The gateway decrypts the traffic and sends it into the virtual network ("B"). Another scenario is the database replicas ("C") or other outgoing traffic for B2B or data center that needs to be encrypted, is routed to the Check Point gateway through the use of User Defined Routes (UDR) ("D"). A similar approach can be used for AKS services, where the containers can establish their egress traffic to the data center using the same path. - The gateway encrypts this traffic and sends it over a site-to-site VPN tunnel to a Check Point gateway on the perimeter of the on-premises network. - Note: In normal conditions this traffic is encrypted end-to-end where we use inter-pod communications, however, the VPN tunnel allows the transport for this traffic, avoiding any connection to the internet.

2) Check Point security gateways for site-to-site VPN communication between the data center on-premises and the public cloud data center[24].

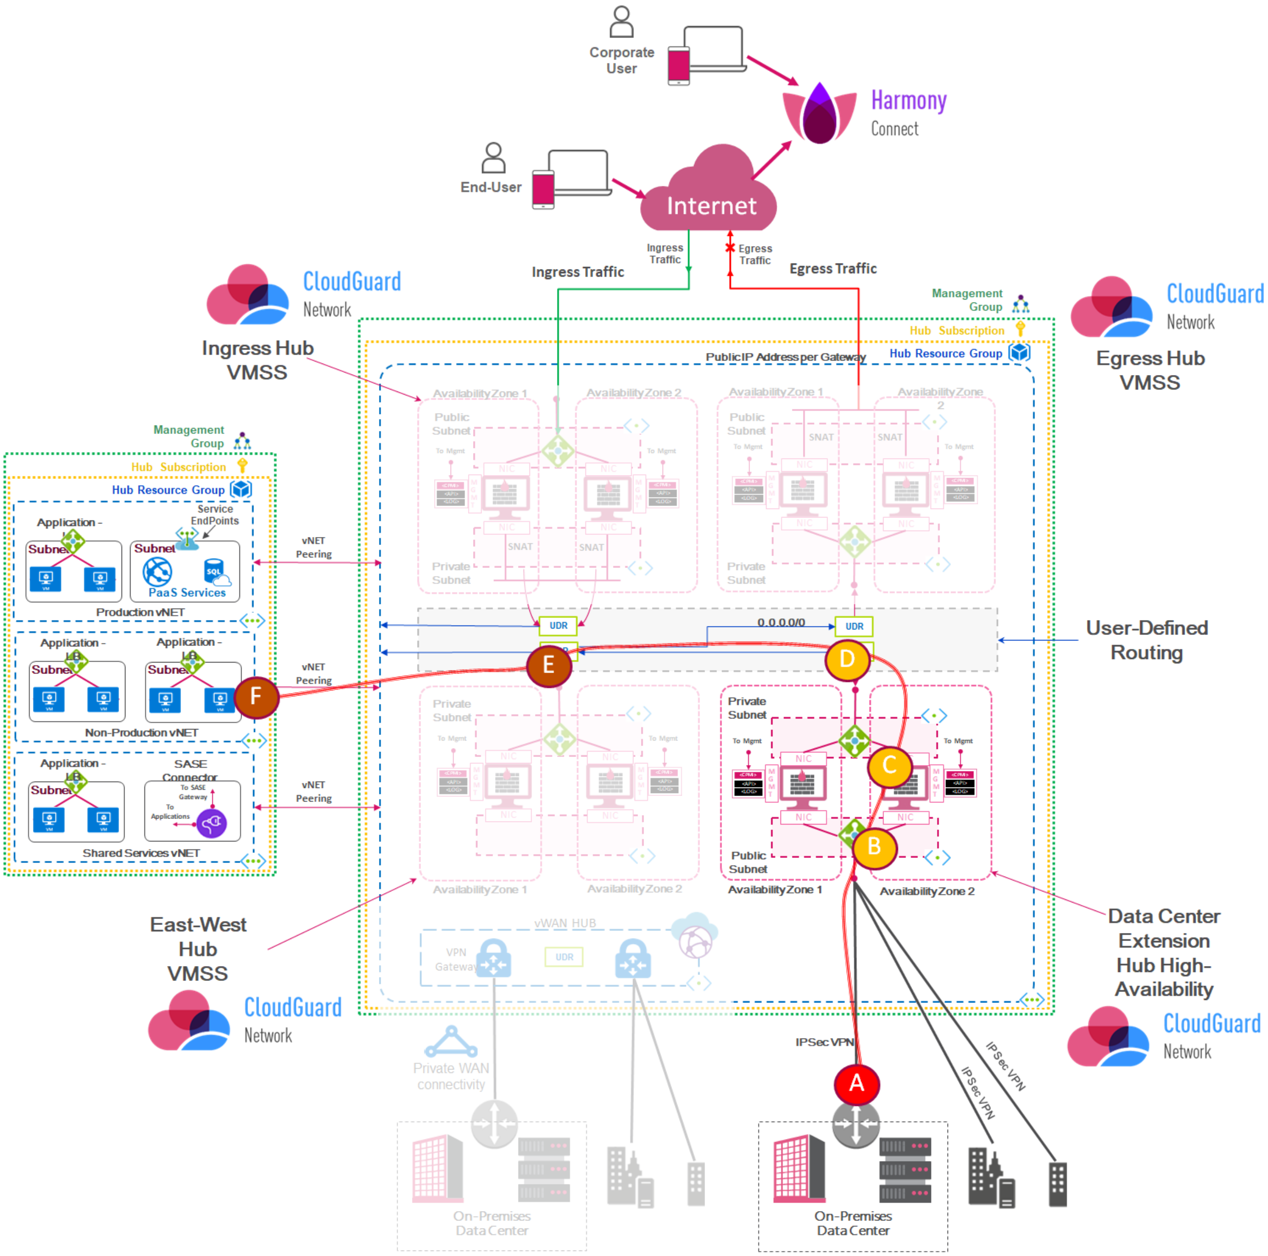

Figure 38: Data Center Connectivity to the Public IaaS Cloud Data Center Using Check Point Security Gateway – Microsoft Azure

In this case, a cluster is deployed in high availability mode: one cluster member is active, and the second cluster member is the standby. The cluster fails over from the active cluster member to the standby cluster member when necessary, and the cluster members communicate to each other with unicast IP addresses. Load balancers only forward traffic to the active cluster member. For VPN traffic between the data center on-premises and the public cloud data center, the load balancers (internal and external) use API calls to Azure to communicate the failover from the active cluster member. The standby cluster member then promotes itself to 'active'.

[24] Virtual Machine Scale Sets (VMSS) for Azure R80.10 and Higher Administration Guide – URL: https://sc1.checkpoint.com/documents/IaaS/WebAdminGuides/EN/CP_VMSS_for_Azure/Content/Topics-VMSS-for-Azure/Overview.htm

During cluster failover, the standby cluster member associates the private and public cluster IP addresses of the active cluster member with its external interface.

Encrypted IPsec traffic is then sent to the gateway's public IP address, where the gateway decrypts the traffic and sends it into the virtual network. Outgoing traffic that needs to be encrypted is routed to the Check Point gateway through the use of User Defined Routes (UDR). The gateway encrypts this traffic and sends it over a site-to-site VPN tunnel, to a Check Point gateway on the perimeter of the on-premises network.

Note: If the organization has express route and site-to-site VPN, we can use a similar approach following the recommendations located in the sk110993.

The Google Cloud Platform also uses a similar approach, as can be seen in the following diagrams:

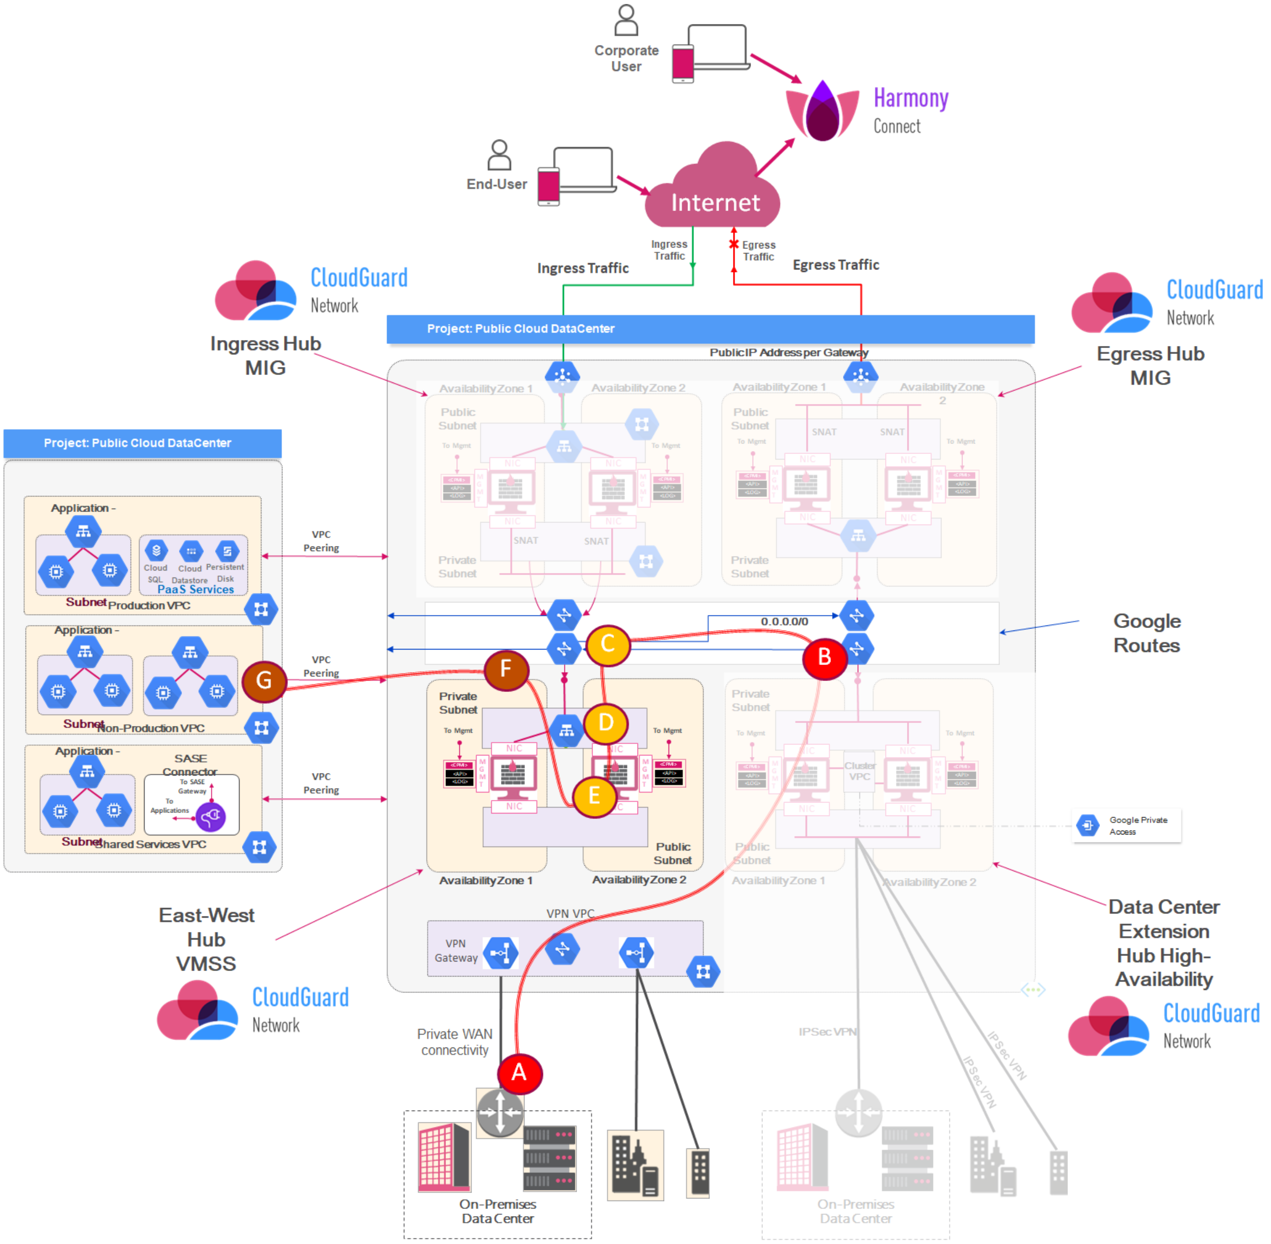

Figure 39: *Data Center Communications to Public Cloud IaaS Using Cloud-Native VPN Gateways – Google Cloud Platform*

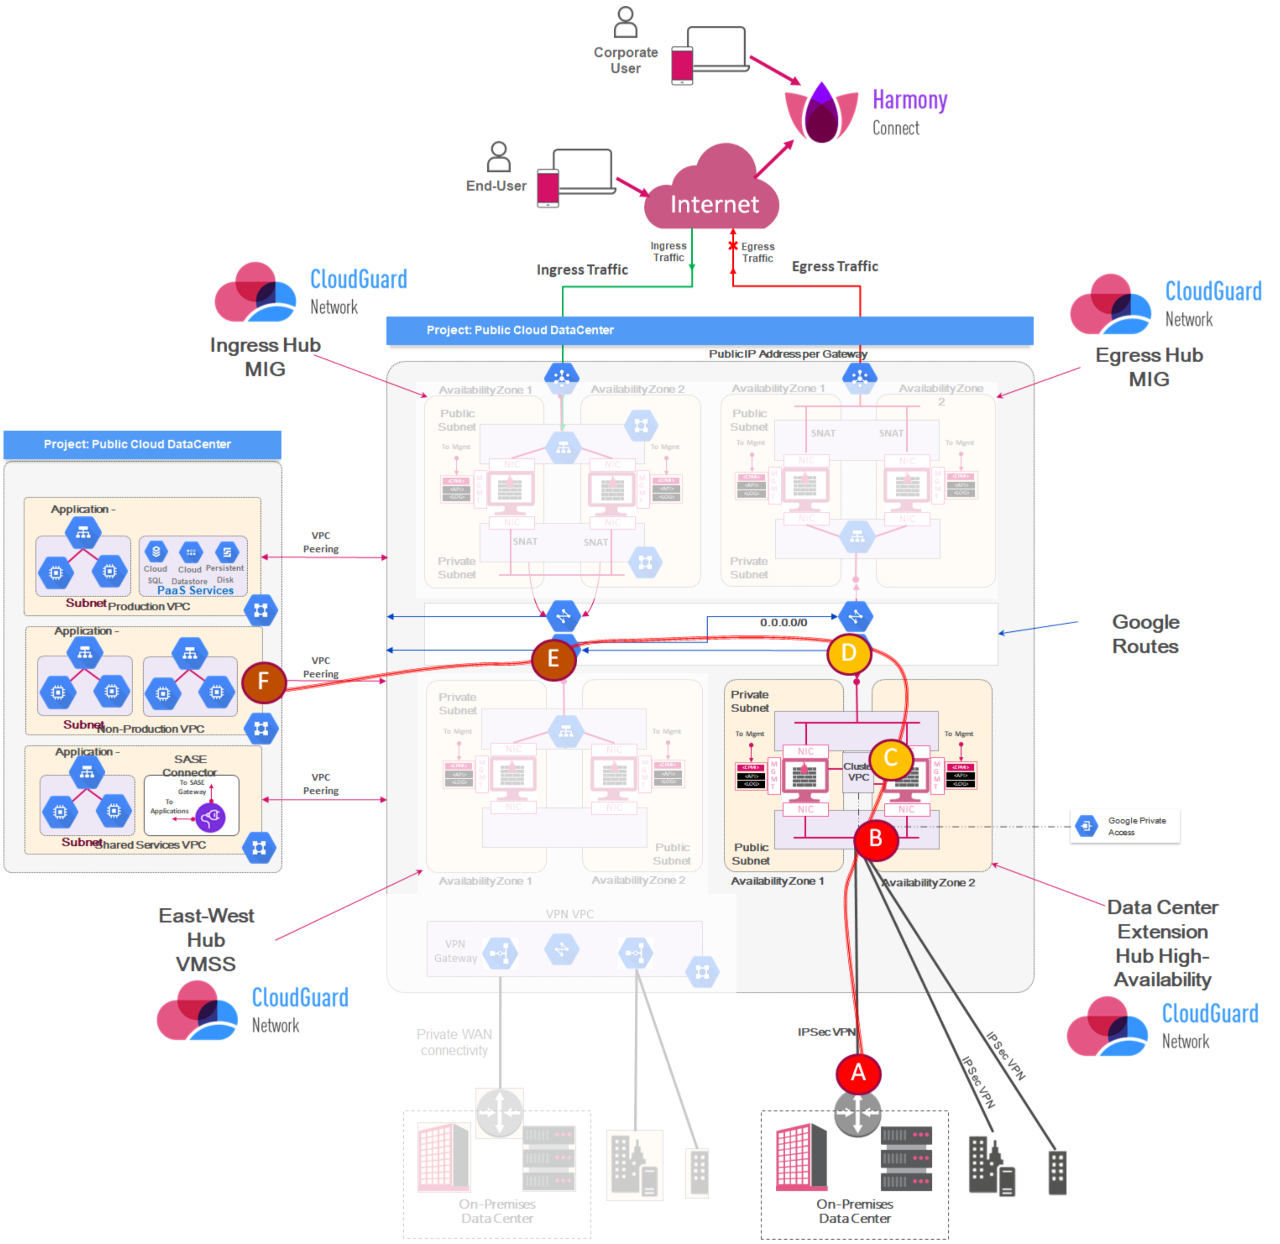

Figure 40*: Data Center Connectivity to the Public IaaS Cloud Data Center Using Check Point Security Gateway – Google Cloud Platform*

In the Amazon architecture, VPN traffic is terminated through the transit gateway to forward the traffic into the proper routing domain. For example, in a traditional infrastructure where we need to communicate the data center to the VPC with database computing instances; routing domain 2 forwards the traffic to routing domain 1 and delivers it. However, if traffic inspection is required, we can use the east-west use case to inspect the traffic, in order to protect the cloud environment. In this hybrid cloud setup, customers can connect the on-premises data center and cloud environments, where all cloud assets can have secured access to on-premises assets.

The connection is established via a secured VPN connection between your Check Point security appliance and a CloudGuard NS for AWS.

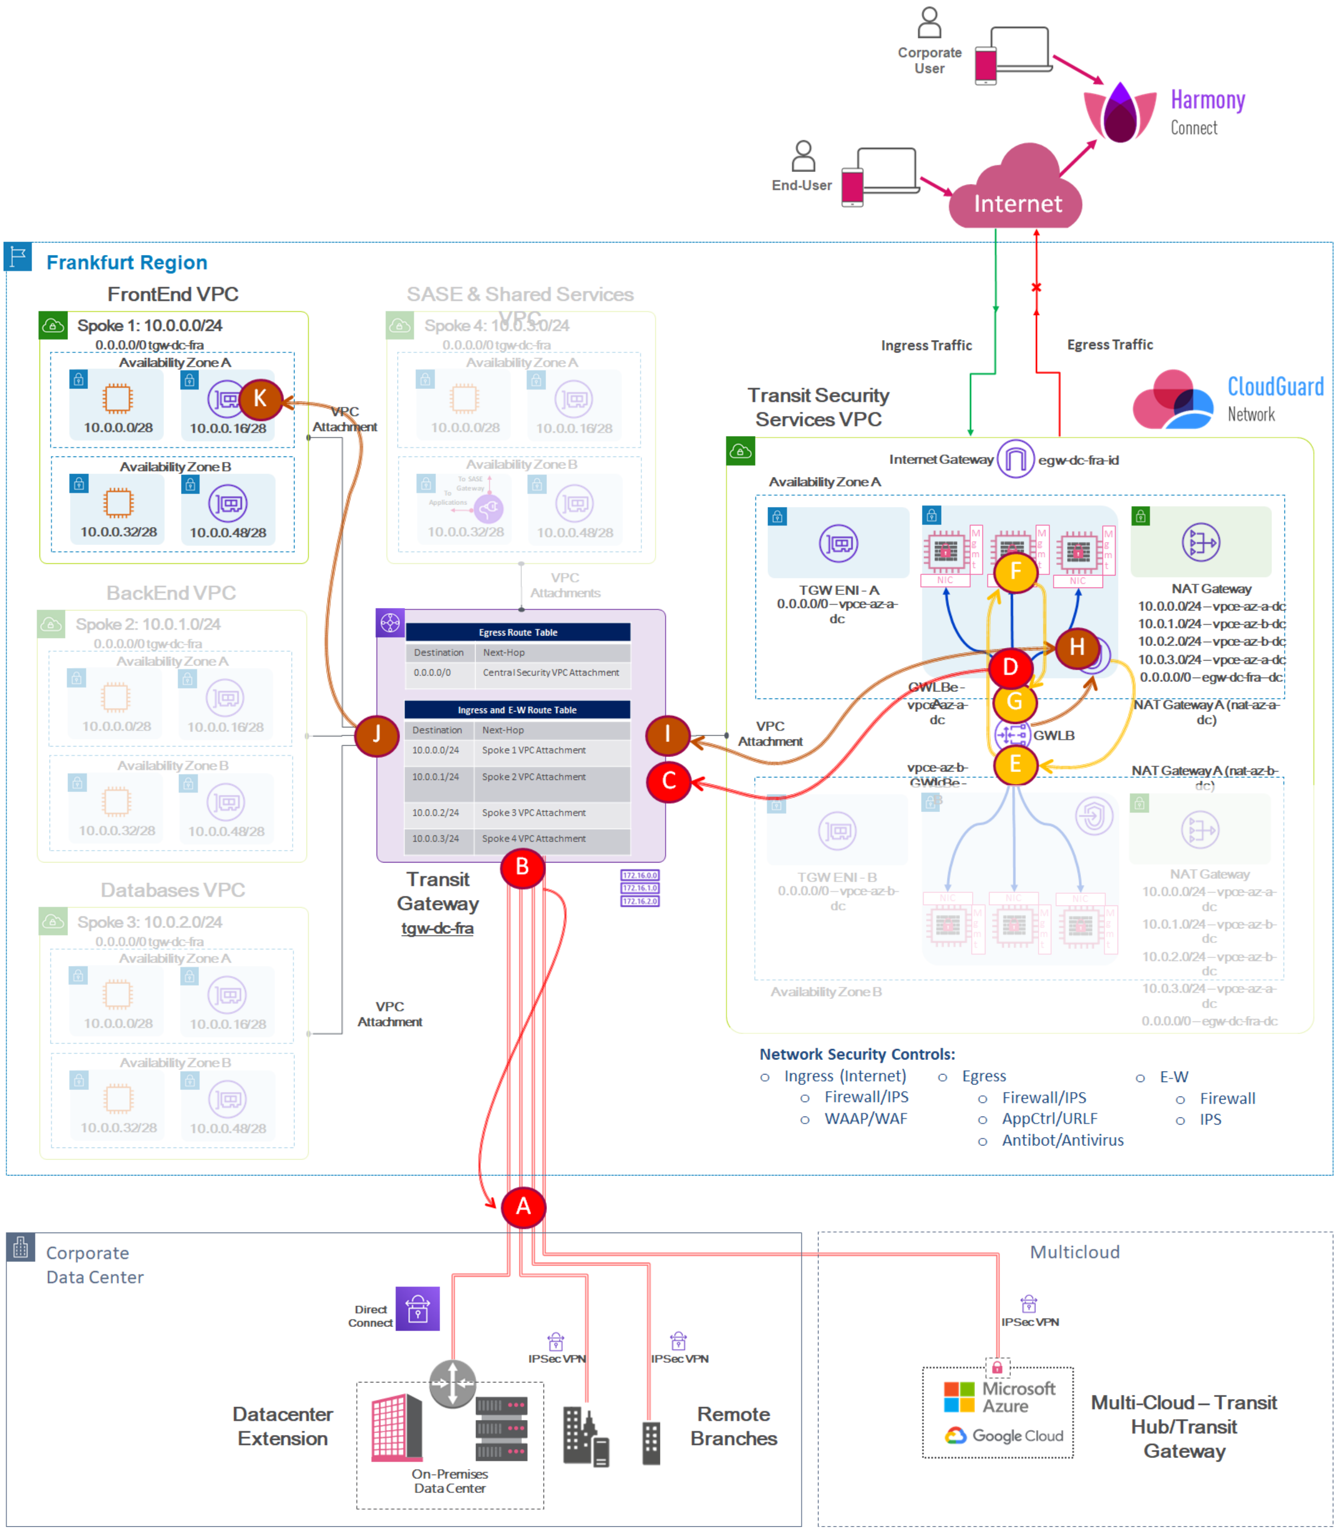

Egress Route Table	
Destination	Next-Hop
0.0.0.0/0	Central Security VPC Attachment

Ingress and E-W Route Table	
Destination	Next-Hop
10.0.0.0/24	Spoke 1 VPC Attachment
10.0.0.1/24	Spoke 2 VPC Attachment
10.0.0.2/24	Spoke 3 VPC Attachment
10.0.0.3/24	Spoke 4 VPC Attachment

Figure 41: Backhaul Communication From Data Center On-Premises to Different Services – Amazon Web Services

In the diagram above, we have illustrated the flows to connect the data center to the different production VPCs, and integrated the east-west VPC, or shared VPC. In this example the flow originated from the data center ("A") to the databases located in the cloud ("K"). Under this approach, traffic is transported over direct connect VPN (between "A" and "B") and forwarded to the security gateway fleet called east-west hub (between "C" and "I"). Here, traffic is allowed to the destination located in the production VPCs ("K"). Another common scenario is database replicas (flow "A" to "K") that can be configured in the routing domain to bypass the traffic inspection.

A third example is when the data center communicates to the PaaS or Kubernetes cluster: the flows are similar, traffic is allowed, inspected, and forwarded to the routing domain, and delivered to the endpoint gateway. The endpoint gateway is the component that allows us to communicate to the PaaS.

A VPC endpoint enables the private connection of the VPC to supported AWS services and VPC endpoint services powered by AWS PrivateLink. Such an endpoint doesn't require an internet gateway, NAT device, VPN connection, or an AWS Direct Connect connection. Instances in your VPC do not require public IP addresses to communicate with resources in the service, as traffic between the VPC and the other services does not leave the AWS network.

As can be seen, endpoints are virtual devices that can be horizontally scaled, redundant, and highly available VPC components, allowing communications between instances in your VPC and services without imposing availability risks or bandwidth constraints on your network traffic.

Remote Access to Corporate Resources Using the Public Cloud (ZTNA)

Following the worldwide outbreak of COVID-19 in March 2020, remote access solutions have become a critical requirement for organizations. Under today's circumstances many workers have their own broadband connections at home with WiFi networks capable of supporting network speeds like those found in the office. This makes it possible for employers to leverage their workers' home networks for connectivity to the edges of corporate resources. However, home networks are often shared with family members, fully rely on untrusted public backbones, and can be entirely unprotected. Connecting remote employees to corporate systems therefore requires additional layers of security to ensure the confidentiality, integrity, and availability of corporate data and systems.

Under a new approach, Check Point's ZTNA strategy is to allow remote access communications without deploying traditional gateways in the public IaaS. This new method facilitates global coverage to be reached through proximity algorithms.

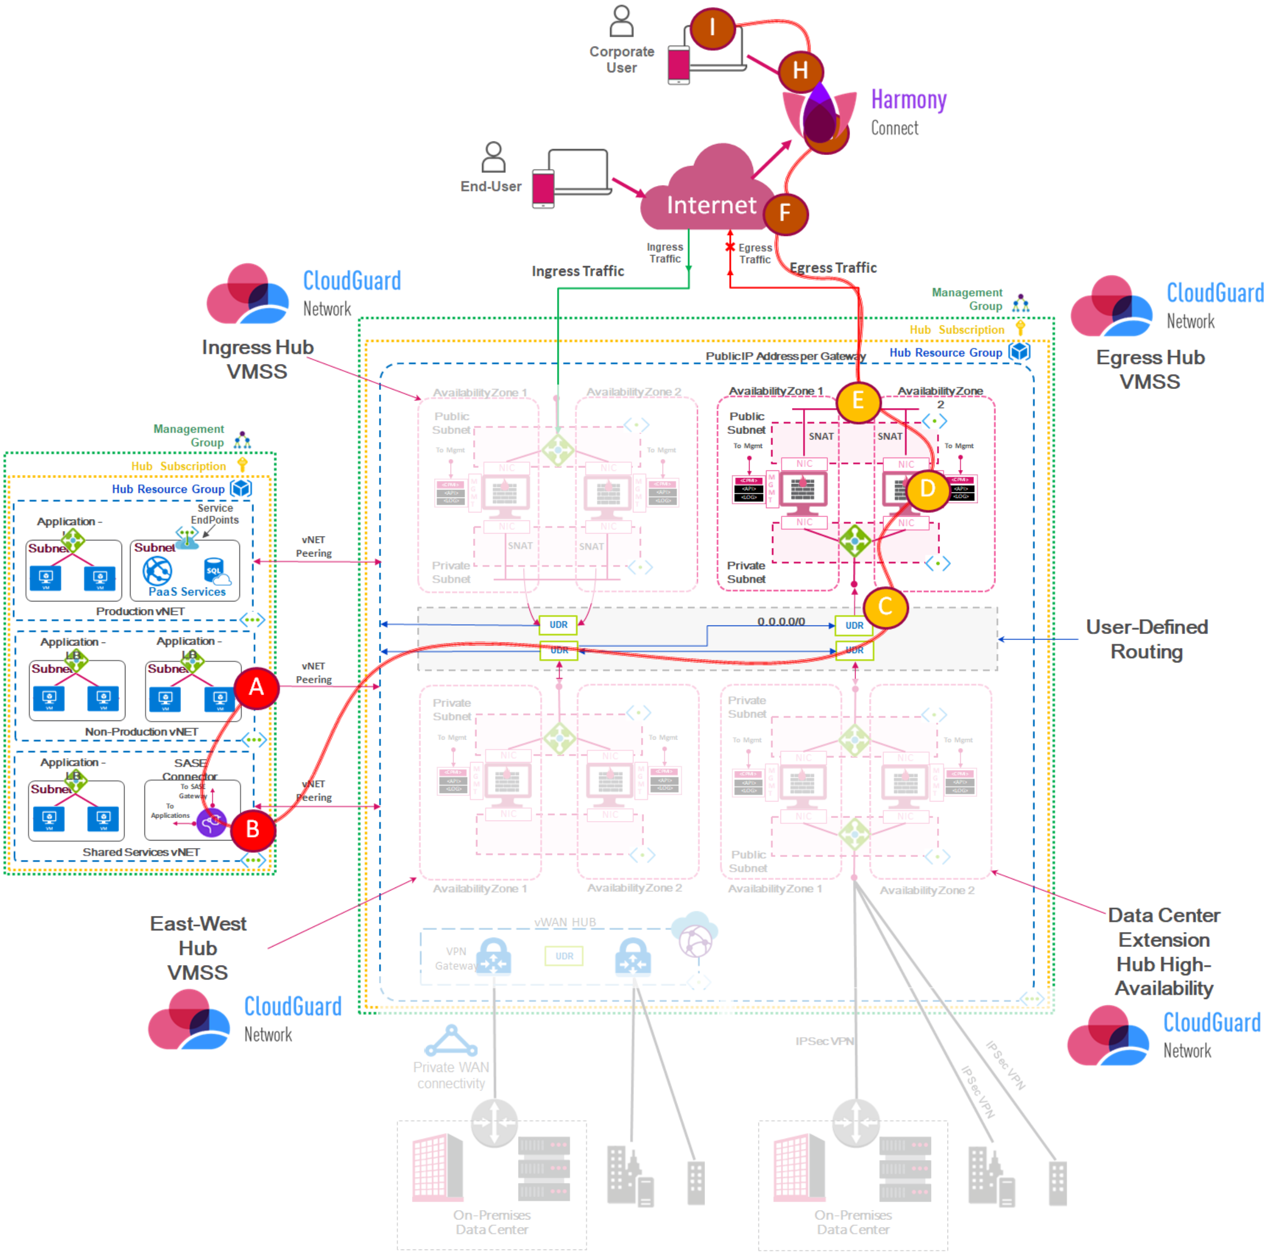

Figure 42: *Remote Access through Check Point SASE – Microsoft Azure*

In the flow description visualized above we can see that traffic, which originated from point A, is located by external users with a client-based, or clientless, VPN. This traffic is terminated in the Harmony Connect for users where all traffic will be inspected and the posture security device will confirm if the remote user can be considered as trusted. If the answer is 'yes', traffic is forwarded to the relevant computing instances (point B) through the Check Point Harmony Connect (ZTNA), peering to the public IaaS data center to access specific services in PaaS (such as app services in point C) or services in the Kubernetes cluster.

FROM	TO
Remote access traffic to cloud data center applications	- Remote access users ("A") connect to the Harmony Connect ("B") using its public IP address to access all the remote access services in the cloud access gateway. - The Check Point SASE connector ("C“, acting as a reverse proxy) is located in the cloud data center and has the capability to connect to different applications located in the production vNETs. The reverse proxy also has connection t through the egress traffic to the internet ("E") and is connected to the Harmony Connect services ("F").

Public cloud IaaS integrated with a ZTNA framework has many advantages for organizations, the most important of which is how it helps adopters prioritize security, no matter what other tools they're using.

For example, if a business requires a transition from one cloud service provider to another cloud service provider, the SASE platform continues undisrupted due to its platform agnostic. This flexibility also makes it easy for businesses to scale up their security infrastructure as they grow, without having to reconfigure the central architecture or deploy dedicated VPN clusters in the public cloud data center. Such an ability to customize security setting operational needs, enables organizations to create an architecture that meets their current and evolving business needs. It is important to keep in mind the following security and network components of a ZTNA architecture:

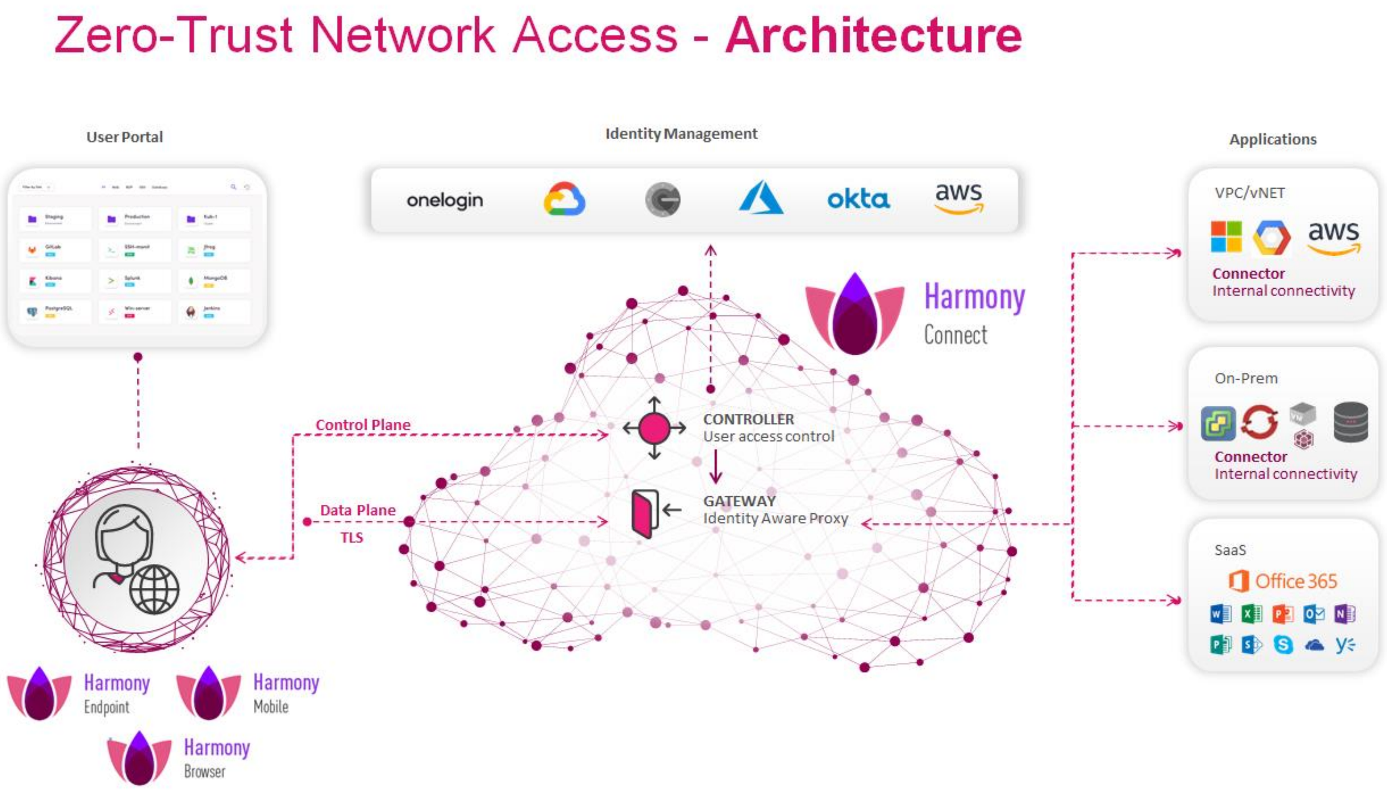

***Figure 43**: Check Point Harmony Connect - ZTNA Architecture*[25]

[25] Check Point Remote Access – URL: https://docs.odo.io/docs/odoaccess-architecture

A similar approach can be seen on the Google Cloud Platform. Following the egress use case described in Google Cloud, the new element here is the ZTNA connector (flow B) that connects to different applications (web, RDP or SSH) located in different VPC's (flow A). Once the reverse proxy is connected to the relevant applications and services are authorized, we can follow the same flow for the egress traffic starting with the Google routes (flow C) and it being processed by Check Point MIG Clusters.

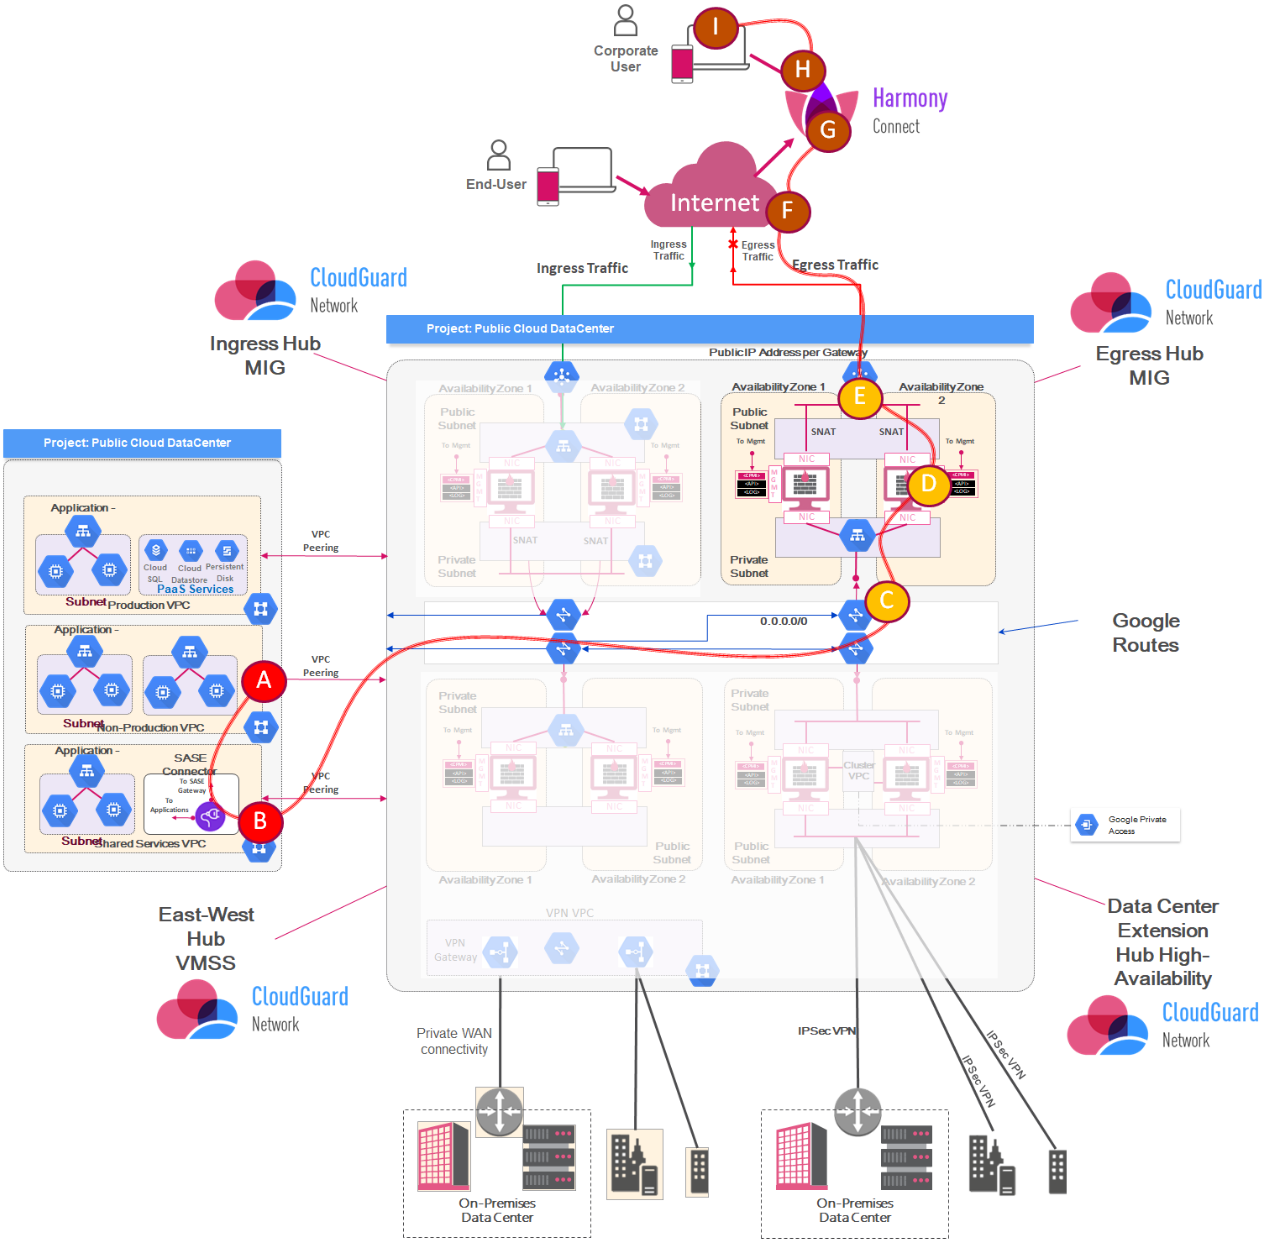

***Figure 44**: Remote Access Through Check Point SASE – Google Cloud Platform*

In AWS, we can integrate the egress and east-west use cases to allow for the connectivity between the SASE connector and Harmony Connect cloud service. The ZTNA connector is located in the SASE and shared services VPC, where it connects to the specific applications located in different VPC's (web applications, remote desktop or SSH), as a reverse-proxy (flow A). The transit gateway then forwards the traffic (flow B) to connect to the applications located in frontend VPC (flow C).

From the user perspective, the clientless VPN first originate the traffic from internet (flow 1), and then connects to the Harmony Connect cloud service (flow 2). Once the user is authenticated, according with their role previously defined in the

RBAC policies (flow 3), the reverse proxy uses the egress flow traffic (from A to D, then D to forward the connection from the connector to the Harmony Cloud service.

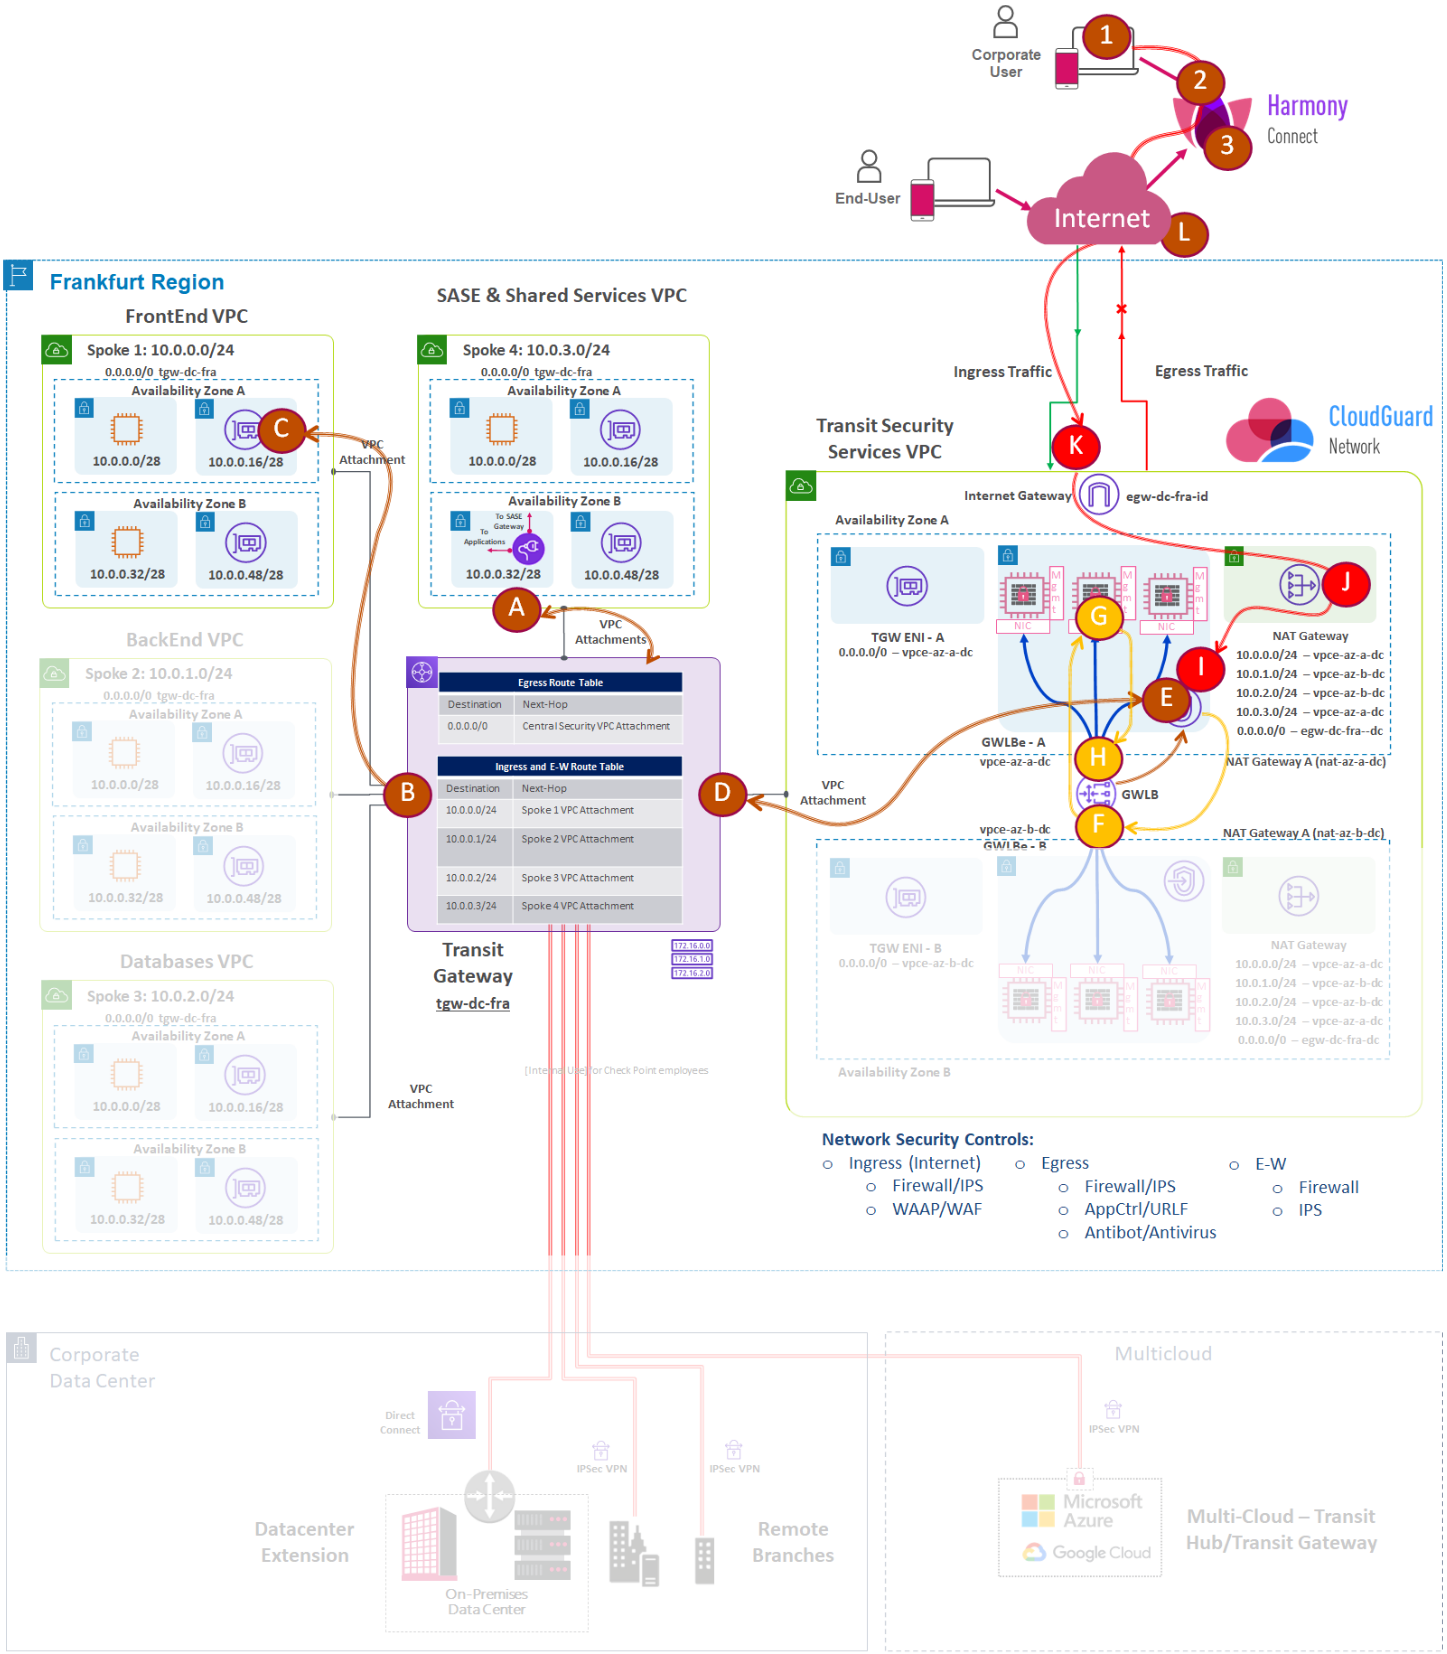

Egress Route Table	
Destination	Next-Hop
0.0.0.0/0	Central Security VPC Attachment

Ingress and E-W Route Table	
Destination	Next-Hop
10.0.0.0/24	Spoke 1 VPC Attachment
10.0.0.1/24	Spoke 2 VPC Attachment
10.0.0.2/24	Spoke 3 VPC Attachment
10.0.0.3/24	Spoke 4 VPC Attachment

Network Security Controls:

- Ingress (Internet)
  - Firewall/IPS
  - WAAP/WAF
- Egress
  - Firewall/IPS
  - AppCtrl/URLF
  - Antibot/Antivirus
- E-W
  - Firewall
  - IPS

Figure 45: *Remote Access Through Check Point SASE – Amazon Web Services*

Remote Access and Windows Virtual Desktop (RDP)

Check Point SASE can also enable the Remote working environments with Windows Virtual Desktop or Virtual Desktop Infrastructure as a Service. Windows Virtual Desktop is a comprehensive desktop and app virtualization service running in the cloud. It is the only virtual desktop infrastructure (VDI) that delivers simplified management, multi-sessions on Windows 10, optimizations for Microsoft 365 apps for enterprise, and support for Remote Desktop Services (RDS) environments. Most significantly, you can deploy and scale your Windows desktops and apps on Azure in minutes, and get built-in security and compliance features.

However, one of the major concerns is the latency to deploy ZTNA services located in different data centers worldwide. To provide a more optimized model, Check Point ZTNA can be complemented with virtual desktop infrastructures. In the following diagram, we can gain an overview of how the customer can access their cloud data center services using the WDI (the same approach could be done with VDI).

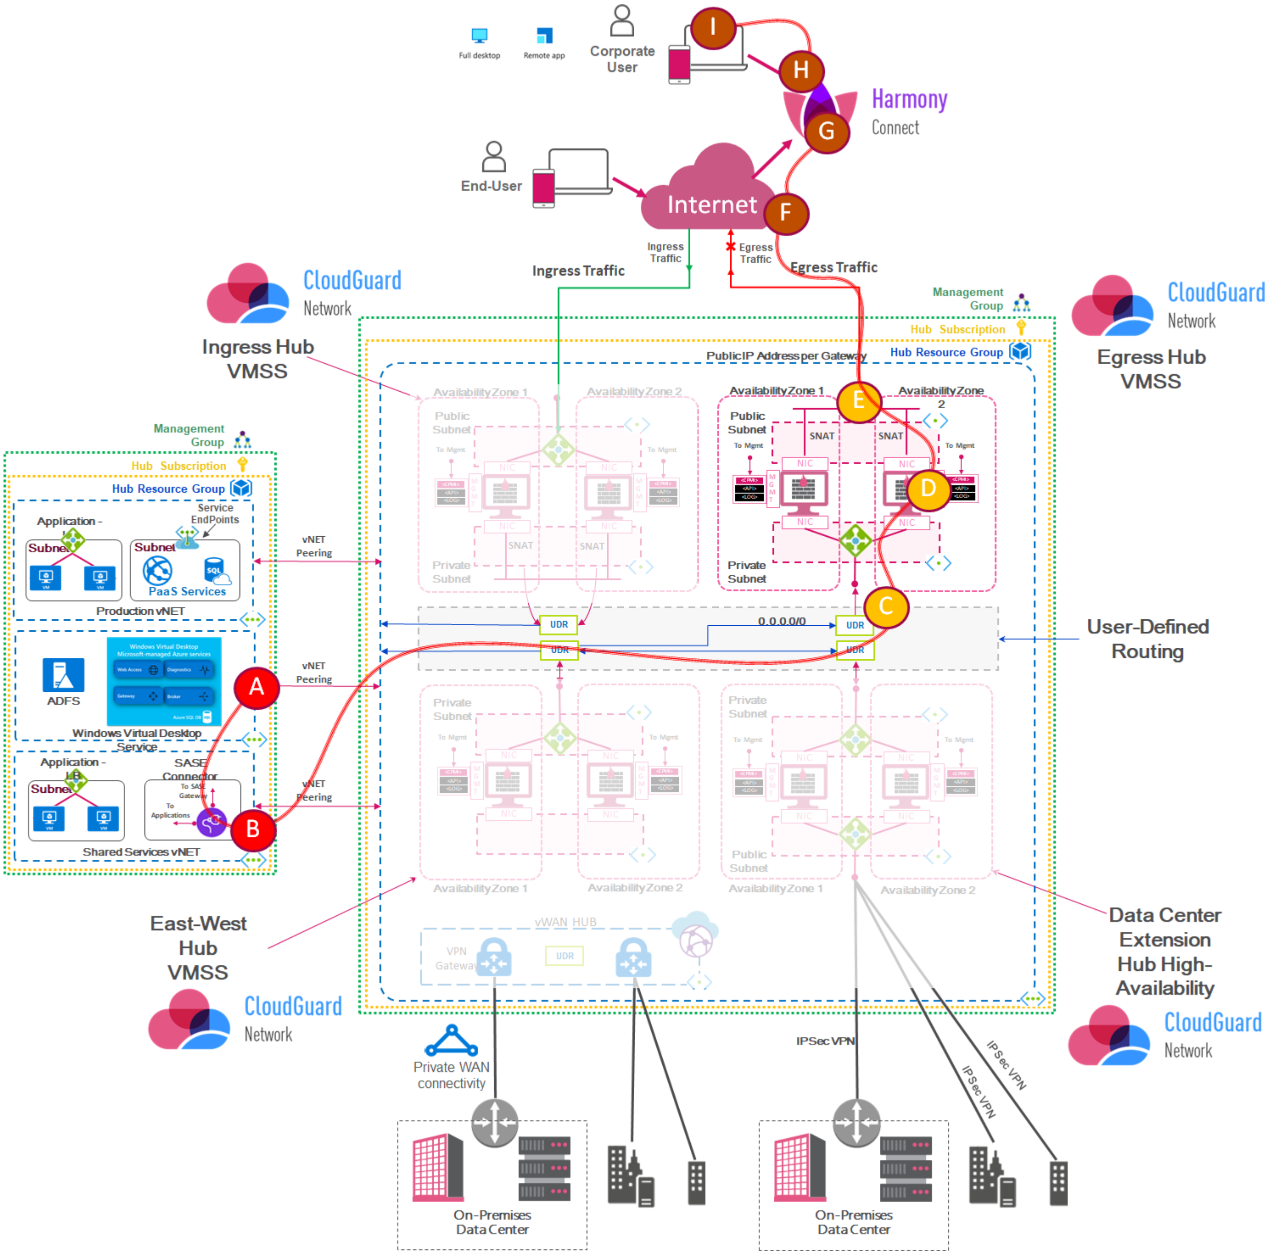

Figure 46: *Secure Remote Access With SASE and WDI to the Cloud Data Center*

FROM	TO
Remote access traffic to cloud data center applications through Windows Desktop Infrastructure	- Remote access users ("A") connect to Harmony Connect for users ("B"), using its public IP address. Check Point ZTNA service inspects the traffic and sends it to the internet or to the peering communications ("C"). This provides communication into the virtual network to access Windows Desktop Infrastructure (or another similar solution in "D"). Then, according to policies, remote users ("E") can access one specific application using the service portal or full desktop. Under this approach computing execution is done in the user virtual environment. This saves on bandwidth and latency to access applications installed in the computing instances, PaaS web applications, or container applications. - In this scenario, the routing is more simple and flexible. This is due to returning packets being managed internally between the WDI managed services and the peering to the functional vNETs. The only returning packets to warranty are the routes back to the gateway through the use of User Defined Routes (UDR) for the thin client, and the WDI managed service. - The reverse-proxy also has connection to the internet through the egress traffic to internet ("F") and is connected to the Harmony Connect services ("G").

CONSOLIDATED CLOUD SECURITY MANAGEMENT

The entire security network is unified and efficiently managed through a single pane of glass, based on modular policy management and rich integrations with 3rd party solutions through flexible APIs. This consolidation brings all security protections and functions under one umbrella. A single console manages the entire IT infrastructure and a unified policy for users, data, applications, and networks that introduces unparalleled granular control and consistent security. The single console provides rich policy management, enabling delegation of policies within the enterprise, and flexible administration rights. In addition, the consolidated management increases operational efficiencies with the automation of routine tasks, freeing up security teams to focus on strategic security rather than repetitive tasks.

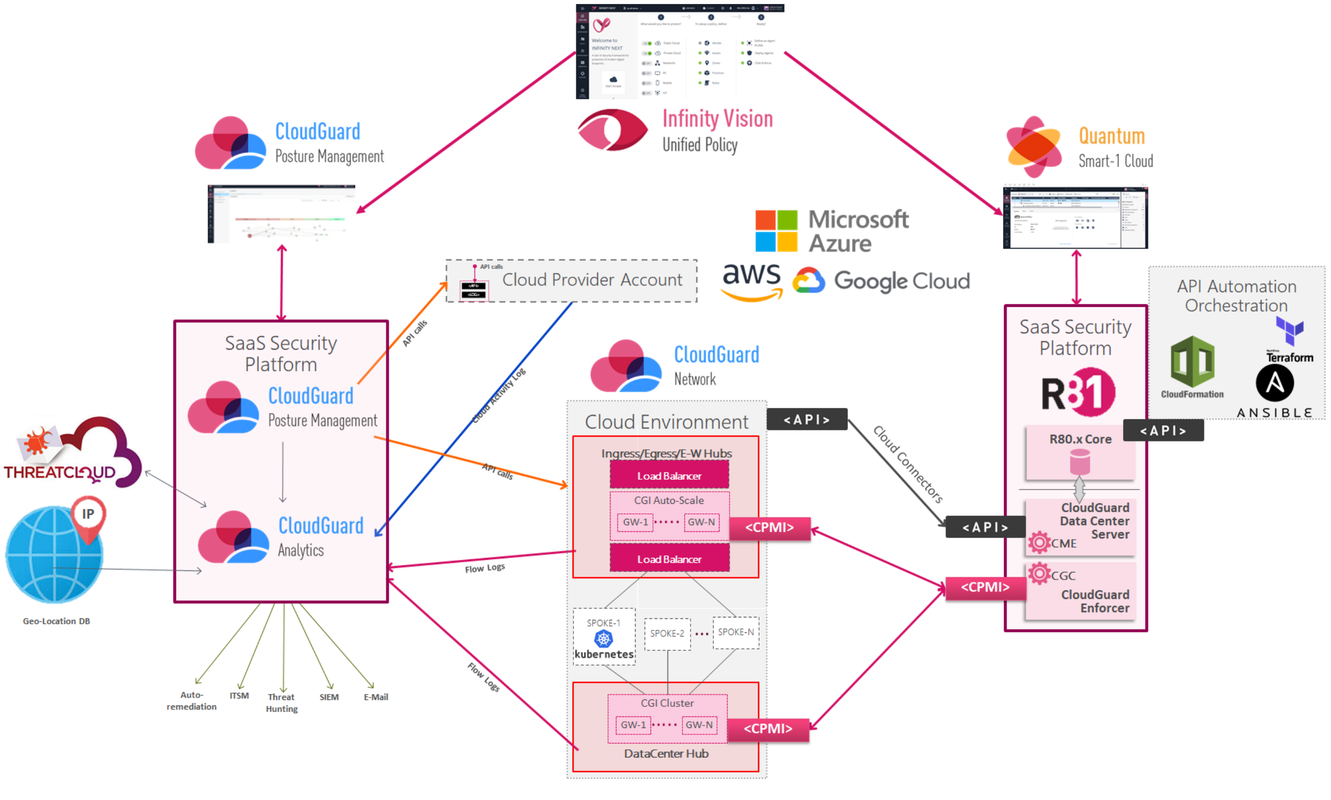

Figure 47: Consolidated Cloud Security Management Architecture

Cloud Network Security Management as a Service (Smart-1 Cloud)

Smart-1 Cloud is a new Software as a Service solution, and cloud-based management tool for organizations looking to unify all their security policies in the cloud. Instead of deploying and maintaining your own management server, it is now possible to connect to Check Point's new cloud service to manage all premise gateways or your CloudGuard NS gateways.

Key Benefits:

- **Elastic growth –** As your company grows, your security management requirements grow as well. Our backend platform scales up according to your needs and enables you to add more capacity on demand.
- **Infinite logging –** There are no physical limits for scaling. Keep your security logs in the cloud for as long as you need with the required log rate, without compromising on the search and query experience.
- **Always up-to-date –** Our service will always be up-to-date with the latest security features and software versions.
- **Zero maintenance –** No need for monitoring or backup operations of your security management system.

Unified Management

A key component of the Check Point Infinity Architecture is the ability to present a single pane of glass for all required management. This means that all security management is consolidated into one place with one console, with no more back and forth between platforms to have full-spectrum visibility. A single view also means that organizations can present a consistent security policy over multiple platforms. The example below shows an R80.x/R81 policy, in which security is being enforced across the public and private cloud, as well as in the on-premises gateways.

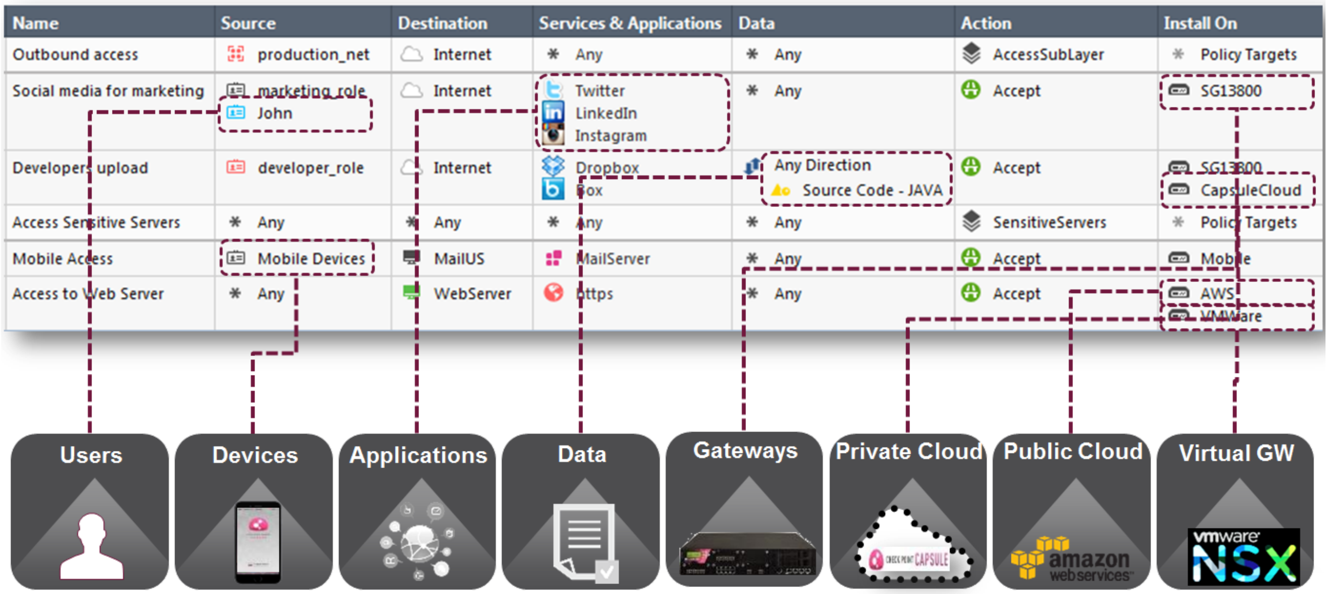

Name	Source	Destination	Services & Applications	Data	Action	Install On
Outbound access	production_net	Internet	Any	Any	AccessSubLayer	Policy Targets
Social media for marketing	marketing_role John	Internet	Twitter LinkedIn Instagram	Any	Accept	SG13800
Developers upload	developer_role	Internet	Dropbox Box	Any Direction Source Code - JAVA	Accept	SG13800 CapsuleCloud
Access Sensitive Servers	Any	Any	Any	Any	SensitiveServers	Policy Targets
Mobile Access	Mobile Devices	MailUS	MailServer	Any	Accept	Mobile
Access to Web Server	Any	WebServer	https	Any	Accept	AWS VMWare

Figure 48: *Full Security Policy for Different Realms: Users, Devices, Applications, Data, Gateways, Private/Public Cloud and Virtual /Environments*

R80.x/R81 Unified Security Policy

The next-generation security management allows for the combining of firewall, application control, and data awareness in one "access control" rulebase. All you need to do is edit your policy, edit the access control layer, and check all the relevant blades.

With that, your rulebase will have new columns available: "Services & Applications" and "Data". In the "Services & Applications" column, where you would normally select the services in your firewall rulebase, you will also be able to select web applications, as well as mobile access applications. In the "Data" column you will be able to select individual data types, and the upload/download a direction. Enforcement will then only apply to the granular selection of this rule. The "Track" column will also have options to include more information about the connection, with regard to the selected applications and data.

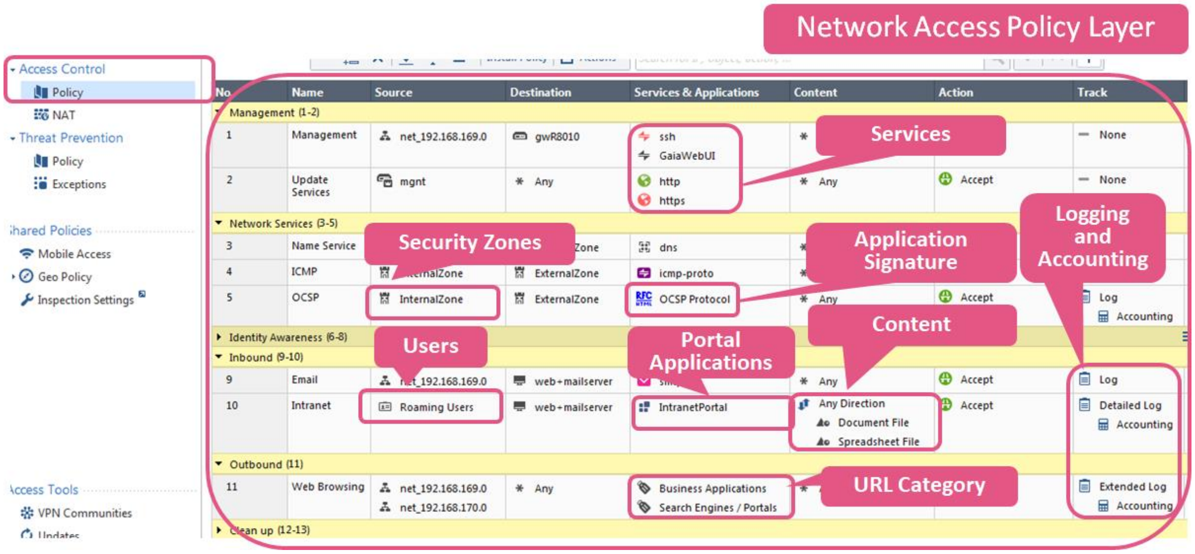

No	Name	Source	Destination	Services & Applications	Content	Action	Track
Management (1-2)							
1	Management	net_192.168.169.0	gwR8010	ssh GaiaWebUI	*		None
2	Update Services	mgnt	Any	http https	Any	Accept	None
Network Services (3-5)							
3	Name Service		Zone	dns	*		
4	ICMP	ernalZone	ExternalZone	icmp-proto	*		
5	OCSP	InternalZone	ExternalZone	OCSP Protocol	Any	Accept	Log Accounting
Identity Awareness (6-8)							
Inbound (9-10)							
9	Email	net_192.168.169.0	web+mailserver	sm	Any	Accept	Log
10	Intranet	Roaming Users	web+mailserver	IntranetPortal	Any Direction Document File Spreadsheet File	Accept	Detailed Log Accounting
Outbound (11)							
11	Web Browsing	net_192.168.169.0 net_192.168.170.0	Any	Business Applications Search Engines / Portals	*		Extended Log Accounting
Clean up (12-13)							

Figure 49: *Unified Security Policy*

Layers

R80.x/R81 introduced the concept of layers with a policy as a new way to organize security access control and threat prevention rules. Layers offer several benefits, among which 'delegation' (the ability to assign different layers to different administrators) is key. This feature can be used to provide separation between tenants in the event that they share a policy.

R80.x/R81 also organizes the policy with ordered layers. For example, gateways that have the firewall and application control blades enabled will have their policies split into two ordered layers: network and applications. Another example is gateways that have the IPS and threat emulation blades enabled, will have their policies split into two ordered layers: IPS and threat prevention. Ordered layers are enforced in the following way: when the gateway matches a rule in a layer it starts to evaluate the rules in the next layer.

The layers concept opens more options for policy management:

- Setting different view and edit permissions per layer for different administrator roles.
- Re-using a layer in different places: the same application control layer in different policy packages (sharing a layer across different policies), or the same inline layer for different scopes.
- Explaining global and local policies in multi-domains with the same feature set of layers: a domain layer will be the set of rules added in each domain by the domain administrator.

R80.x/R81 gateways and above will have the ability to utilize layers in the following new ways:

- Unifying all blades into a single policy.
- Segregating a policy into more ordered layers, not necessarily by blades.

- Allowing sub-policies inside a rulebase, with the use of inline layers.

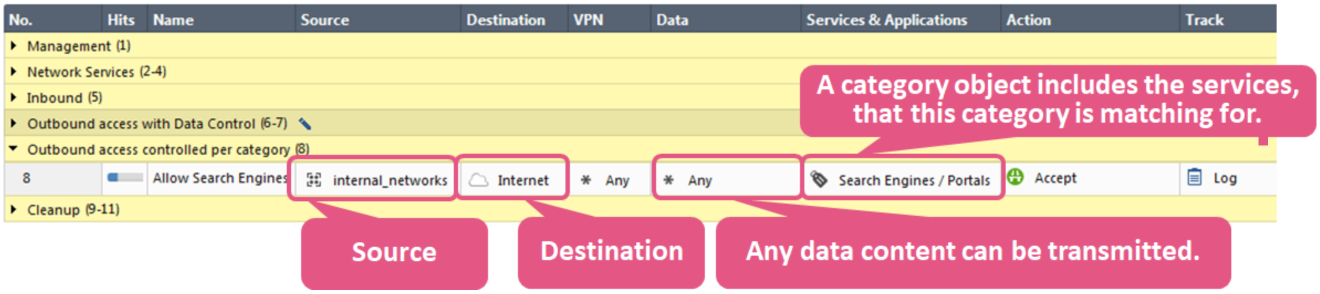

No.	Hits	Name	Source	Destination	VPN	Data	Services & Applications	Action	Track
▸ Management (1)									
▸ Network Services (2-4)									
▸ Inbound (5)									
▸ Outbound access with Data Control (6-7)									
▾ Outbound access controlled per category (8)									
8		Allow Search Engines	internal_networks	Internet	Any	Any	Search Engines / Portals	Accept	Log
▸ Cleanup (9-11)									

A category object includes the services, that this category is matching for.

Source

Destination

Any data content can be transmitted.

__Figure 50:__ Rule-Based Content Awareness Security

For the access control policy, two types of layers exist, for maximum flexibility: the inline layer and the ordered layer, where layers allow for the separating of the security policy into multiple components. In this way, better security and manageability is created. To be more efficient, support of concurrent-admin and segregation of duties allow organizations to reuse layers either inline or ordered in multiple policies.

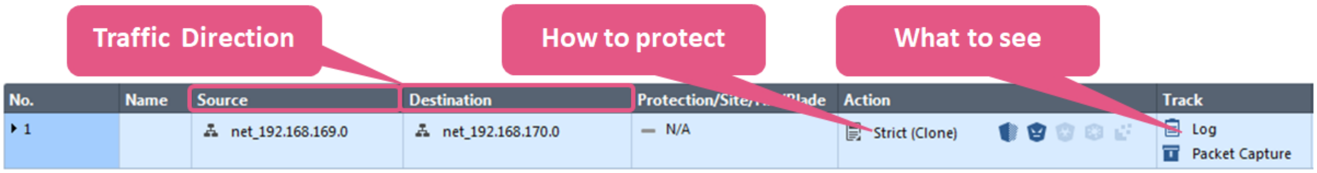

Traffic Direction

How to protect

What to see

No.	Name	Source	Destination	Protection/Site/...Blade	Action	Track
▸ 1		net_192.168.169.0	net_192.168.170.0	N/A	Strict (Clone)	Log Packet Capture

__Figure 51:__ Rule-Based Threat Prevention

1. In the case of inline layers, only traffic matched/accepted on the parent rule will reach and be inspected by the inside layer rules.
2. In the case of ordered layers, when an accept rule from the first layer is matched, the gateway goes over the rules in the next layer.
   1.1. For backward compatibility with pre-R80 gateway, you will need to use ordered layers to manage the firewall rulebase and application control rulebase. This is where the first layer needs to be a firewall layer and the second layer needs to be an application control and URL filtering layer.
   1.2. During an upgrade from pre-R80 to R80 gateways, using policy packages that are implementing firewall and application control policies, the existing policy will be separated by an ordered layer. The separation is made by the network layer, the firewall policy rules as the first layer and application layer, and the application control policy rules as the second layer.

The following is an example of traffic matching using policy layers

Policy with inline Layers	Policy with ordered layers	Policy mixed with ordered and inline layers
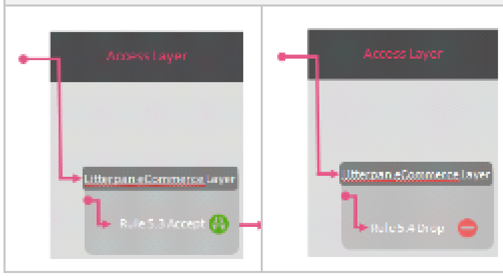	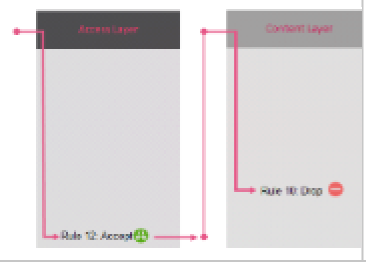	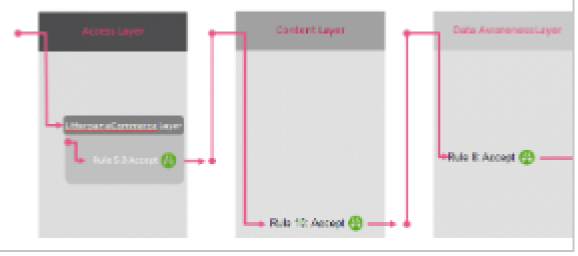

__Figure 52__: R80.x/R81 Policy Layers

Smart-1 Cloud answers the evolving needs of today's enterprise security management. By keeping up with the latest security best practices, organizations can manage threats across devices and workloads via a single management console in the cloud. They have the ability to scale and manage the increasing number of gateways as our environments grow, without having to worry about limited physical storage space or log storage capacity. Most importantly, they can effectively manage maintenance and save valuable time on onboarding new gateways, monitoring devices, and any new updates or hotfixes, as well as on facility power. This is all done automatically when you use Smart-1 Cloud.

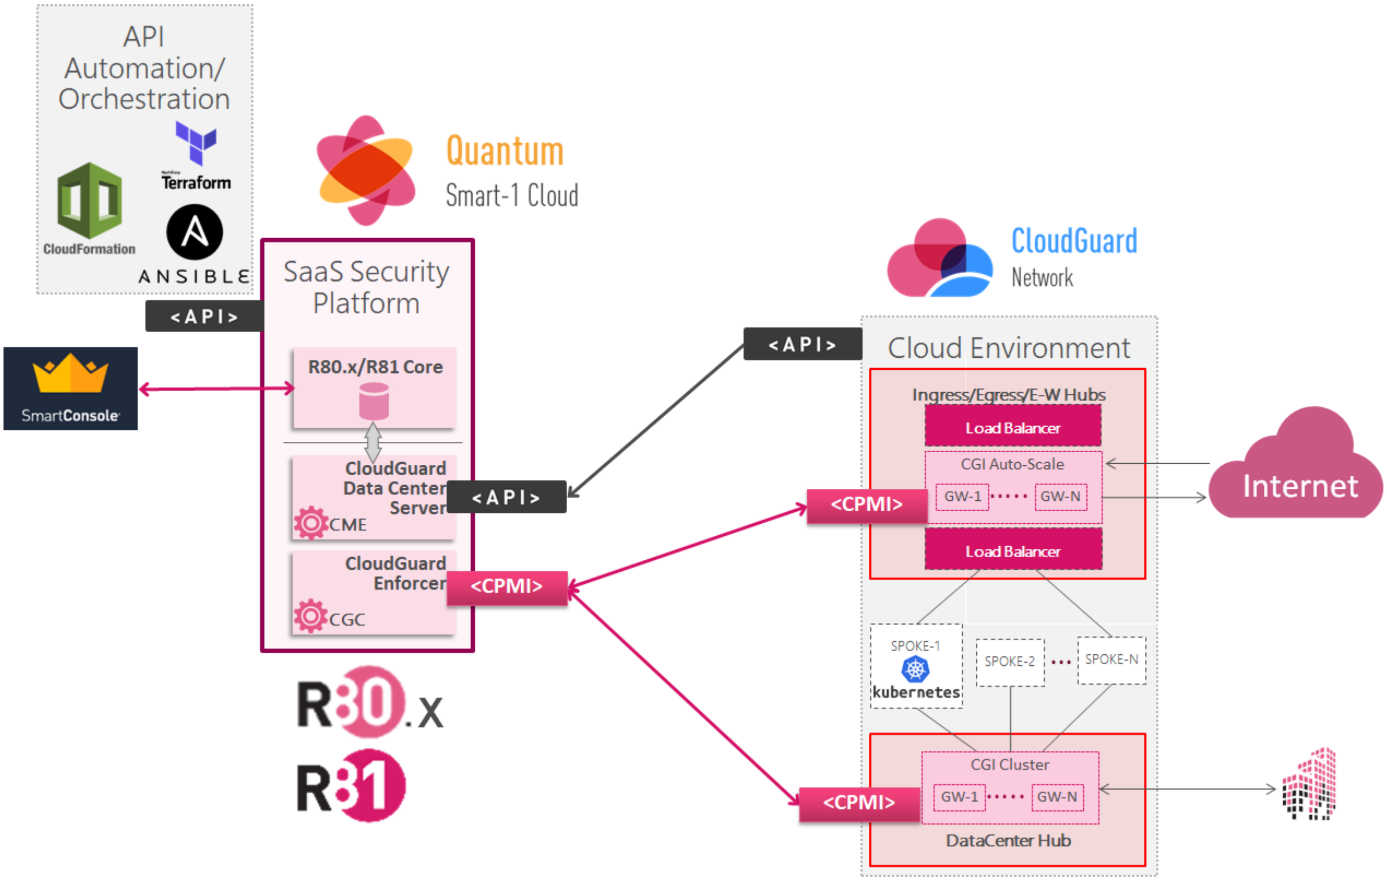

Figure 53: *Smart-1 Cloud Architecture*

CloudGuard Controller

Check Point CloudGuard integrates with all leading public cloud management solutions, to absorb and leverage contextual information about the infrastructure. Cloud-defined elements such as asset tags, objects, security groups, and more are updated in real time, allowing CloudGuard to automatically adjust security policies to any changes in your dynamic cloud environment. Organizations can create a unified and creative policy. The controller also allows for the simultaneous use of cloud objects like regions, availability zones, subnets and virtual machines in our policy, from different cloud platforms. By doing so, the cloud network security policy becomes more adaptive to the dynamic changes that occur in the cloud.

Overview of Cloud Management Extension (CME)

CME is a utility that runs on Check Point's security management server and multi-domain servers running Gaia OS. CME allows cloud-native integration between Check Point CloudGuard NS solutions and cloud platforms.

As a service that runs on Check Point management servers, CME continuously monitors CloudGuard NS solutions deployed in various cloud vendors, and synchronizes them with the security management server.

Supported solutions and features for:

- CloudGuard for Azure VMSS.
- CloudGuard for AWS ASG.
- CloudGuard for GCP MIG.
- Automatic hotfix deployment for autoscaling solutions.
- CloudGuard for AWS Transit VPC.
- CloudGuard for AWS Transit Gateway.

Attributes:

- Management module.
- Based on Check Point's identity awareness technology.
- Allows the use of cloud-management defined objects (e.g. tags, security groups) within Check Point security policies.
- Can be used independently of any CloudGuard gateways.

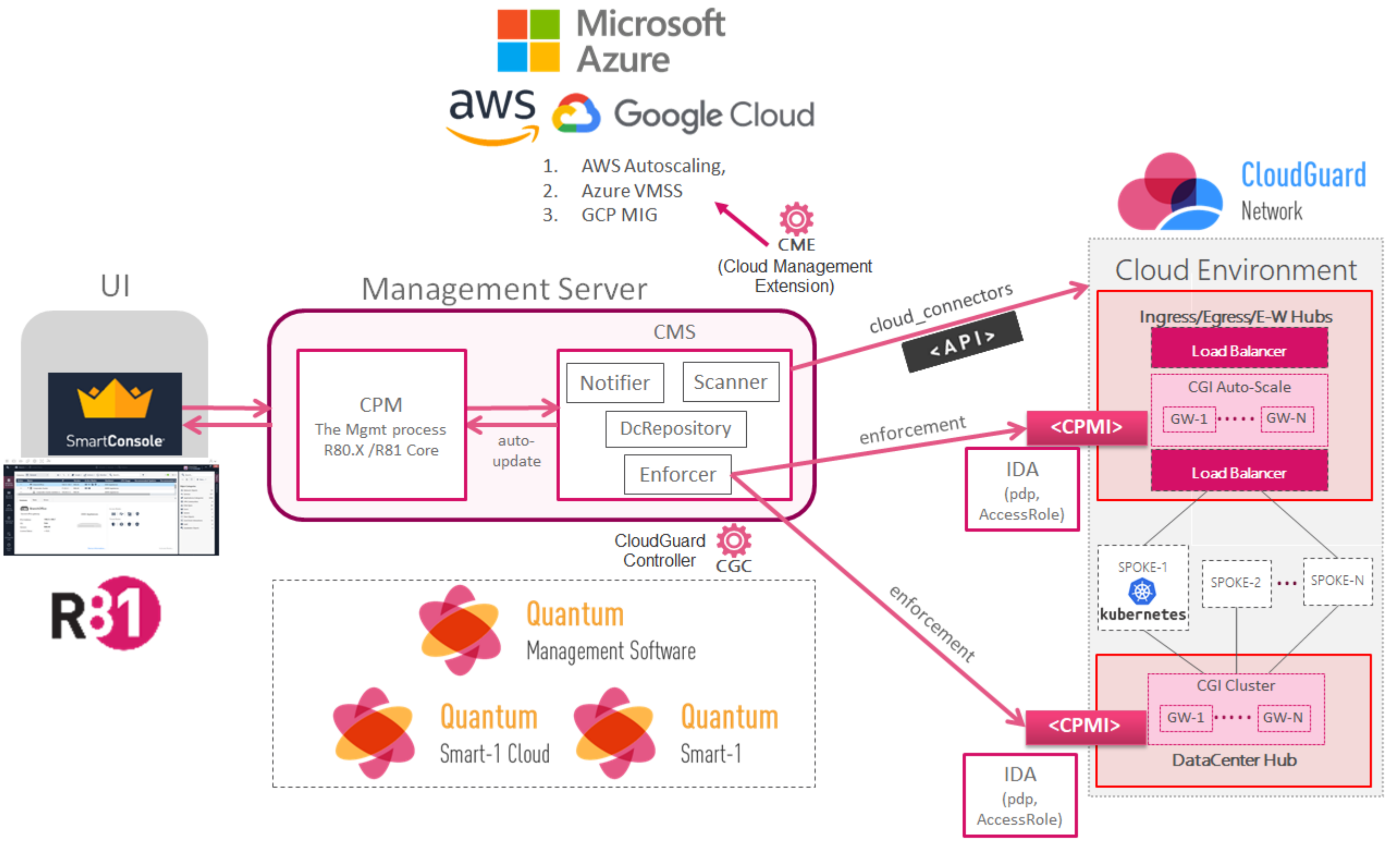

Figure 54: *CloudGuard Controller Architecture*

Quantum Smart-1 Cloud is Check Point's security management server delivered as a cloud service (Software as a Service),it provides security policy management of on-premises and IaaS security gateways, security event analytics, and reporting from a single user-friendly, web-based SmartConsole.

Additionally, it ramp-up time changes from hours to minutes when compared with on-premises deployments. Check Point provides the infrastructure and the software, ensuring installation, deployment, and maintenance times are eliminated and the latest software.

	Standalone	Open Server	Management Appliance	Smart-1 Cloud
Number of Managed Gateways	1	Static	Static	Scale on-demand
Deployment Time	+1 hour	+2 hour	+1 hour	< 2 mins
Troubleshooting	Complex, TAC + if appliance failure needs RMA	TAC + hardware vendor support	TAC + if appliance failure needs RMA	Part of service
Management Performance	Based on hardware utilization	Based on hardware size	Based on hardware size	Auto-scaled
Management High Availability	Additional hardware needed	Additional hardware needed	Additional hardware needed	Delivered from cloud

Figure 55: *Quantum Management Software, Quantum Smart-1, and Quantum Smart-1 Cloud Comparison Table*

Smart-1 Cloud and Shared Responsibility

Smart-1 Cloud is available via the Infinity Portal, Check Point's one stop shop for all SaaS services. With zero installation, deployment, or maintenance, IT teams have more control and oversight over their entire infrastructure. This makes their cloud-native, hybrid, and multi-cloud environments even more secure, manageable, and compliant.

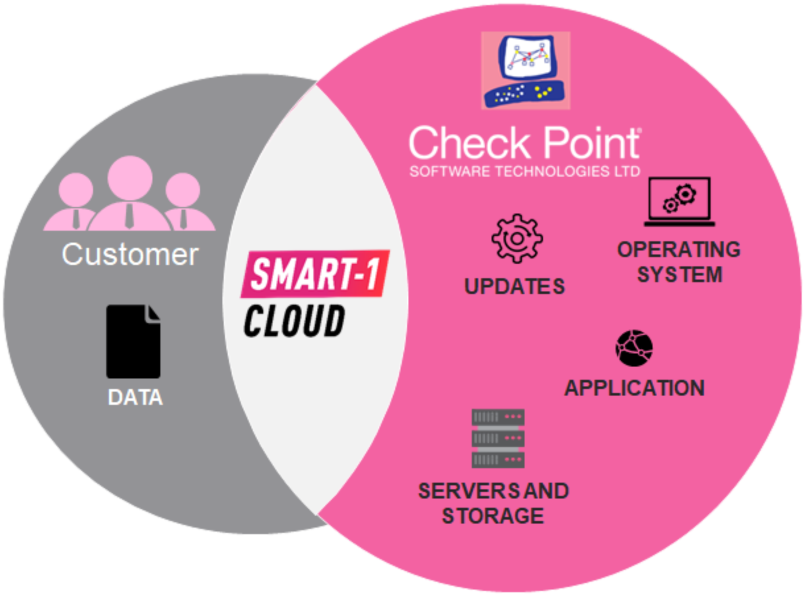

Figure 56: *Smart-1 Cloud Shared Responsibility*

Operational efficiency increases using cloud delivered management, as Smart-1 Cloud environments are always up-to-date with the latest security. Smart-1 Cloud also leverages automatic tools to make sure the system is always up and running, with routine backups every 12 hours.

INFRASTRUCTURE AS CODE

Infrastructure as Code (IaC) is the “process of managing and provisioning computer data centers through machine-readable definition files, rather than physical hardware configuration or interactive configuration tools”. Private and public cloud infrastructures are managed by such a process using scripts or declarative definitions in the code, rather than manual processes. IaC is therefore an essential part for the public cloud, used to build complex data centers through the proper modeling of the business process.

Modular development

- It allows to treat the infrastructure as a piece of software which can be written once and used multiple times, making the process of code writing much more efficient.

Software development methodologies

- It allows to advantage of proven practices like version control, modular development, testing etc. in the infrastructure world.

Problem resolution

- Debug and ascertain the root cause of the problem easily. As we maintain IaC, we can take advantage of versioning systems, so anyone trying to debug an issue or a problem with infrastructure can look through the history of changes made and trace the problem.

Agility

- It allows the entire application development and deployment process more agile by ensuring less dependence on manual work, which thereby reduces errors.

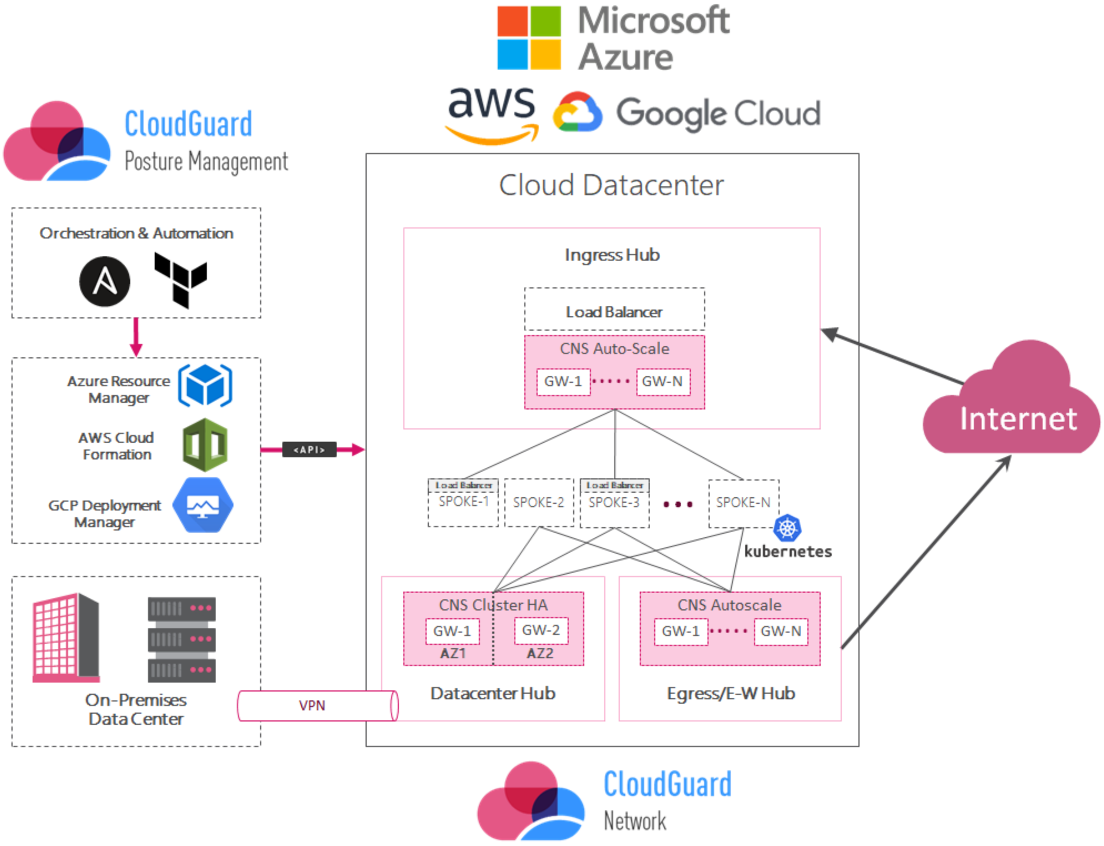

***Figure 57**: IaC – Infrastructure as a Code Secured by the CloudGuard Portfolio*

Infrastructure as Code Posture Management

CloudGuard Posture Management, part of the CloudGuard Cloud Native Security platform, automates governance across multi-cloud assets and services. This includes visualization and assessment of the security posture, misconfiguration detection, and enforcement of security best practices and compliance frameworks. CloudGuard Posture Management is a purpose-built security and compliance solution for cloud environments, providing network security policy management and automation for your cloud environment. Organizations trust CloudGuard Posture Management to ensure their network security is well defined and understood, and then to enforce the security policy on a continuous basis.

For example, in the case of an unauthorized change, (such as when a human or machine attempts to open network ports) CloudGuard Posture Management prevents the event, instead enforcing the configuration you have previously defined.

Here are the following benefits to implement and deploy Posture Management in the Infrastructure as Code:

- **Security operations:** Visualize assets, assess security posture, fix misconfigurations and threats, manage the cloud firewall, and enforce security from a single source of network authority.
- **Privileged identity protection:** Protect against compromised credentials and identity theft using a cloud's native IAM capabilities to safeguard access to actions that can have a big impact.
- **Compliance and governance:** Manage the compliance lifecycle for standards such as PCI DSS, from automated data aggregation and assessment, to remediation and reporting.
- **Cloud security intelligence:** CloudGuard Threat Intelligence is a cloud-native security intelligence technology that delivers cloud intrusion detection, network traffic visualization, and user activity analytics.

Whether public or private clouds are in use, CloudGuard Posture Management facilitates management of server configurations ranging from a few dozen, to hundreds or thousands of instances. Its flexible security management tools ensure compliance while reducing configuration errors and potential breaches.

In the following diagram, the general architecture of the Check Point solution for cloud security posture management can be seen:

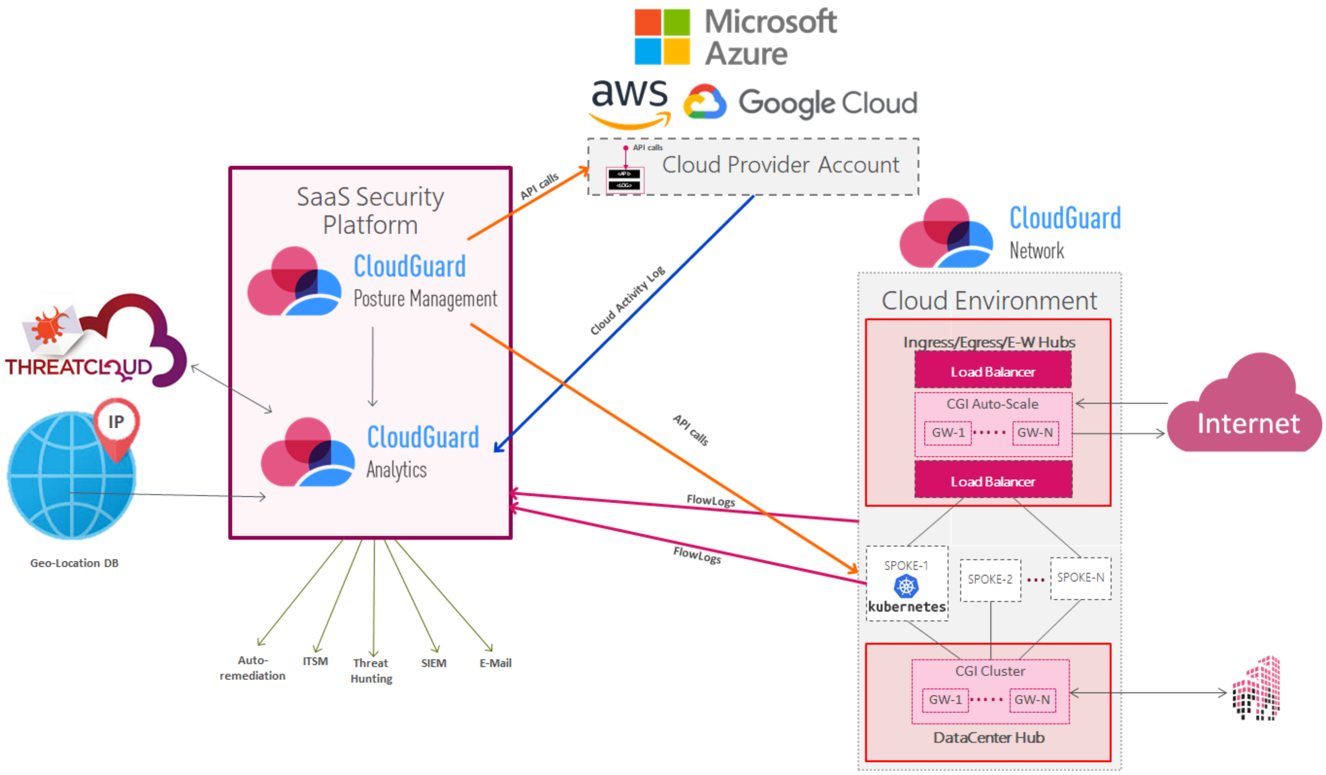

__Figure 58:__ CloudGuard Posture Management Architecture

CloudGuard Posture Management (CPM) can provide the following:

- **Security operations:** Visualize assets, assess security posture, fix misconfigurations and threats, manage the cloud firewall, and enforce security from a single source of network authority.
- **Privileged identity protection:** Protect against compromised credentials and identity theft using a cloud's native IAM capabilities to safeguard access to actions that can have a big impact.
- **Compliance and fovernance:** Manage the compliance lifecycle for standards such as PCI DSS, from automated data aggregation and assessment, to remediation and reporting.
- **Cloud security intelligence:** CloudGuard Intelligence and Threat Hunting is a cloud-native security intelligence technology that delivers cloud intrusion detection, network traffic visualization, and user activity analytics.
- **Workload protection:** Seamlessly integrate protections and controls into your CI/CD tools, like CloudFormation and Terraform, and evaluate security posture pre-deployment; scaling across hundreds of thousands of cloud assets.

CloudGuard Posture Management provides enriched vulnerability management findings to better identify, prioritize, and auto-remediate events based on public exposure, thus minimizing risk. It also prevents critical cloud security misconfigurations and will keep up with evolving posture management security and compliance best practices, including auto-remediation. CloudGuard also complies with regulatory and industry standards, such as HIPAA, CIS BENCHMARKS, NIST CSF/800-53, PCI-DSS, with the most contextual cloud security across 70+ native cloud services.

The following diagram provides the visual representation of all the connections and security relations between the virtual networks and assets.

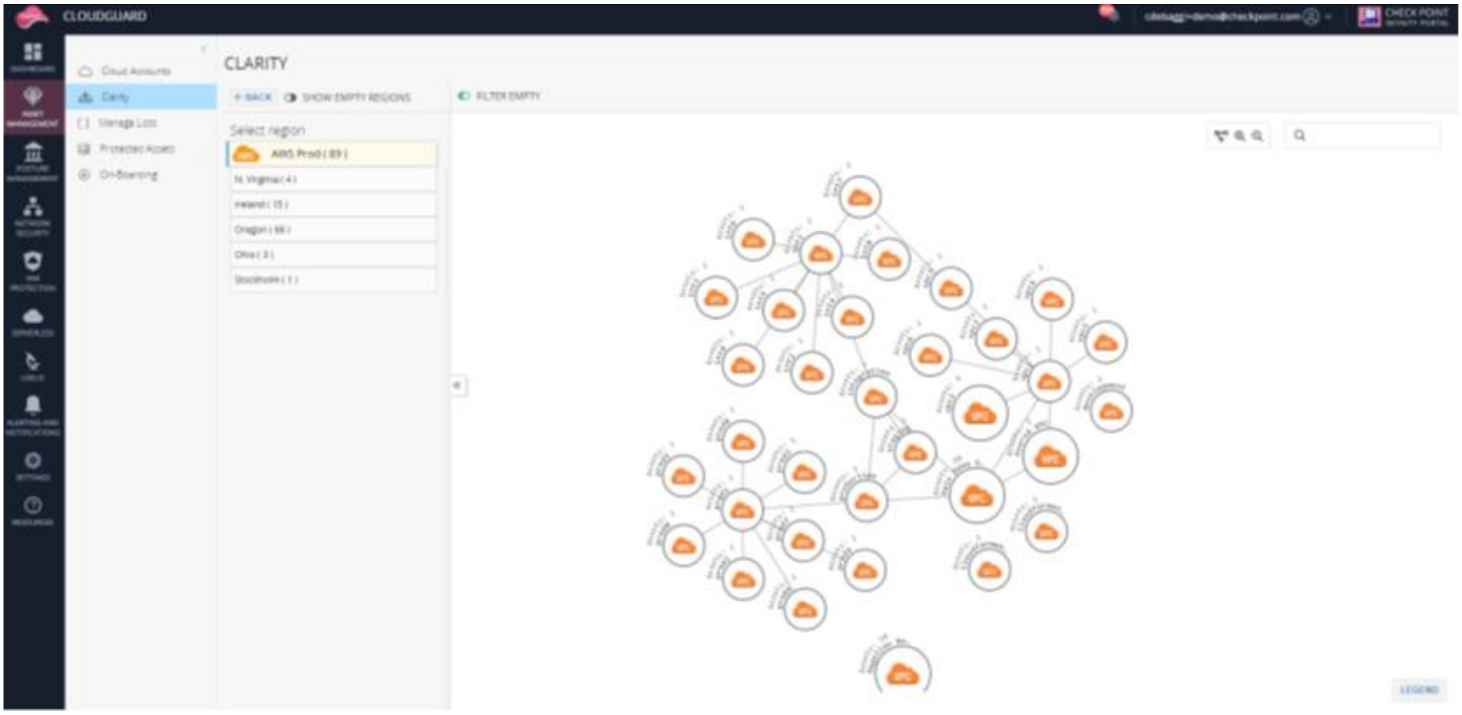

__Figure 59:__ CloudGuard Cloud Security Posture Management: Network Security Flow Analysis

CloudGuard Analytics (Intelligence and Threat Hunting)

CloudGuard Analytics takes vNET/VPC flow logs' enriching data with threat cloud IOC information. As part of the supported services, it shows how the ECS service communicates with other services. CloudGuard Posture Management also provides the user with the ability to visualize the data flows and run GSL queries for immediate incident response and threat hunting purposes. After this is done, you can investigate and identify suspicious ECS related activity, providing investigation and threat analytics capabilities for ECS traffic.

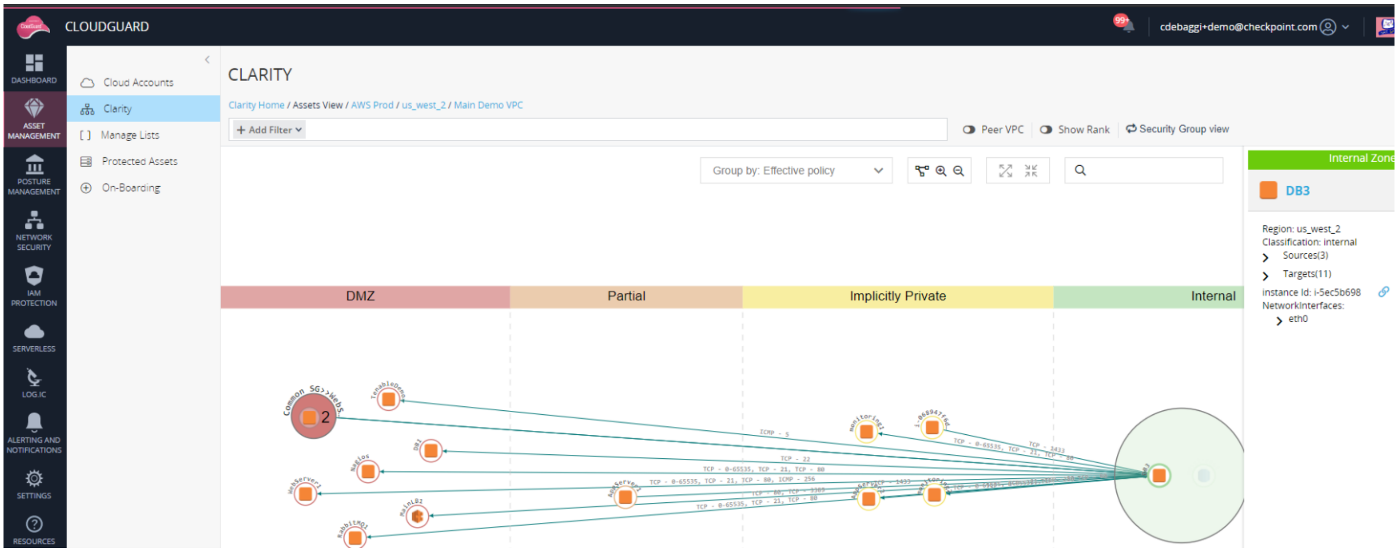

Region: us_west_2
Classification: internal
Sources(3)
Targets(11)
instance Id: i-5ec5b698
NetworkInterfaces:
eth0

__Figure 60__: Flow Visibility for Threat Hunting and Analytics.

Infrastructure as Code Assessment With CloudGuard Posture Management and Terraform

CloudGuard protects AWS, Azure, GCP, and other platforms with continuous security and compliance for your cloud environments. Now, that same security can be integrated into CI/CD pipelines with IaC security capability for DevSecOps. Shifting cloud security will help the CI/CD pipeline to stop misconfigurations and policy violations from ever occurring. Also, instead of being forced to fix post-production deployments, developers are notified of issues immediately.

With IaC security, API calls from the CI/CD pipeline tool insert CloudGuard ShiftLeft to scan the IaC template. CloudGuard will then analyze the template for misconfigurations and policy violations. Afterwards, CloudGuard combines the analysis with any existing violations for a more complete risk understanding if the template connects to existing cloud resources. Finally, the results are quickly returned via API with customizable enforcement levels of pass/fail.

What is Terraform?

Terraform is an open source “Infrastructure as Code” tool, created by HashiCorp.

As a declarative coding tool, Terraform enables developers to use a high-level configuration language called HCL (HashiCorp Configuration Language) to describe the desired “end-state” cloud or on-premises infrastructure for running an application. It then generates a plan for reaching that end-state and executes the plan to provision the infrastructure.

Because Terraform uses a simple syntax, it is possible to provision infrastructure across multiple cloud and on-premises data centers, and can safely and efficiently re-provision infrastructure in response to configuration changes, it is currently one of the most popular infrastructure automation tools available.

Note: CloudGuard ShiftLeft can effectively define, provision, and manage the resources needed without needing the IT member interference.

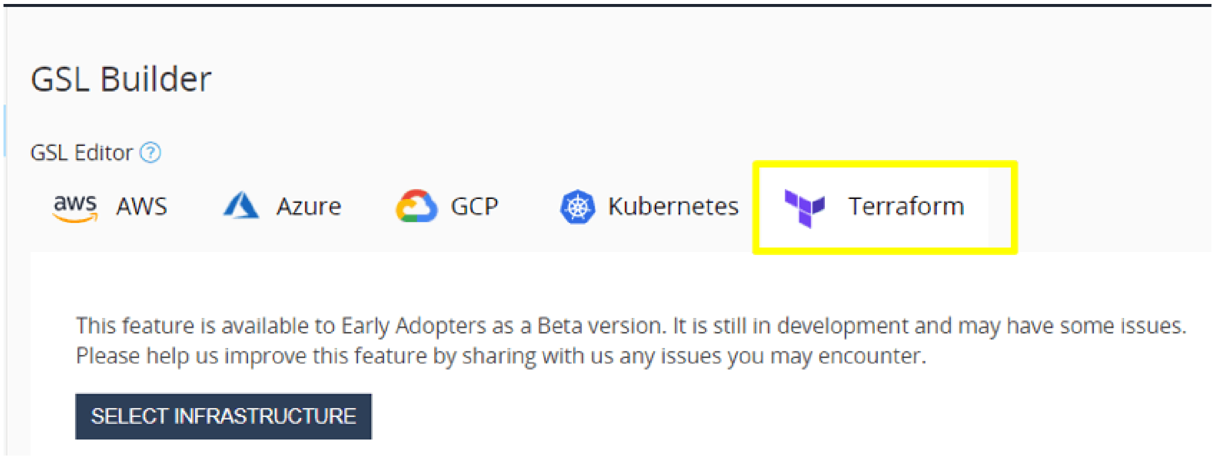

This feature is available to Early Adopters as a Beta version. It is still in development and may have some issues. Please help us improve this feature by sharing with us any issues you may encounter.

Figure 61*: CloudGuard CPSM Integrated with Terraform*

The CloudGuard Terraform rulesets, based on various compliance frameworks, are applied to infrastructure-as-code plans. The plans are evaluated for compliance before being created and deployed in cloud accounts. Misconfigurations and other compliance issues are eliminated at source.

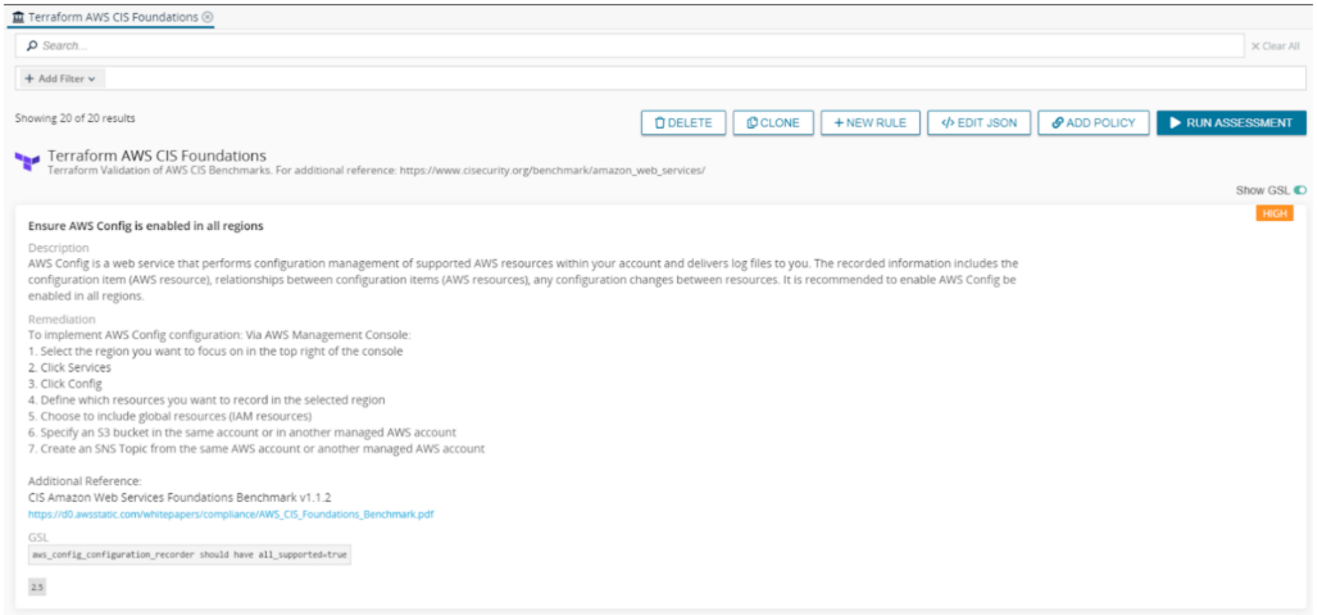

Showing 20 of 20 results

Terraform AWS CIS Foundations

Terraform Validation of AWS CIS Benchmarks. For additional reference: https://www.cisecurity.org/benchmark/amazon_web_services/

Ensure AWS Config is enabled in all regions

Description

AWS Config is a web service that performs configuration management of supported AWS resources within your account and delivers log files to you. The recorded information includes the configuration item (AWS resource), relationships between configuration items (AWS resources), any configuration changes between resources. It is recommended to enable AWS Config be enabled in all regions.

Remediation

To implement AWS Config configuration: Via AWS Management Console:
1. Select the region you want to focus on in the top right of the console
2. Click Services
3. Click Config
4. Define which resources you want to record in the selected region
5. Choose to include global resources (IAM resources)
6. Specify an S3 bucket in the same account or in another managed AWS account
7. Create an SNS Topic from the same AWS account or another managed AWS account

Additional Reference:
CIS Amazon Web Services Foundations Benchmark v1.1.2
https://d0.awsstatic.com/whitepapers/compliance/AWS_CIS_Foundations_Benchmark.pdf

GSL

```
aws_config_configuration_recorder should have all_supported=true
```

2.5

Figure 62*: Terraform AWS – CIS foundations*

The list of security risks for IaC platforms are extensive and includes network exposure, configuration drifts, and ghost resources. These risks must be considered as part of the development and testing of lifecycles. IaC scanning removes barriers for security by improving efficiency in software development, configurations, and speed.

CONCLUSION

For organizations looking to optimize and improve their business processes, the cloud transformation journey begins with the public IaaS transition. The Massachusetts Institute of Technology describes this need for transition well: *“the cloud native approach focuses on building, running and deploying applications by integrating all the new advantages provided by cloud technologies, such as on-demand delivery, flexibility, global deployment, productivity, agility, scalability, and cost savings, among others.”*

Cloud-native Zero Trust principles are shining a light on the path of this journey, which aims to protect the data and information of the organizations. This is why Check Point CloudGuard provides unified cloud native security for all your assets and workloads, giving to you the confidence to automate security, prevent threats, and manage posture across your cloud deployments.

In this whitepaper, we discussed how cloud network security is one of CloudGuard’s key capabilities, as it provides advanced threat prevention and automated network security through a virtual security gateway. This is done using a unified security management system across all your cloud and on-premises deployments, while promoting the automation of processes using APIs, and supporting Infrastructure as Code (IaC) practices.

CloudGuard can provide to security teams a unified security management system for hybrid and on-premises deployments with a single pane-of glass, while promoting a consistency of security policies across public IaaS and hybrid-clouds. This, is a much more efficient way of handling security than using a single management console for each deployed cloud.

CONTACT US

Worldwide Headquarters | 5 Shlomo Kaplan Street, Tel Aviv 67897, Israel | Tel: 972-3-753-4555 | Fax: 972-3-624-1100 | Email: info@checkpoint.com **U.S. Headquarters** | 959 Skyway Road, Suite 300, San Carlos, CA 94070 | Tel: 800-429-4391; 650-628-2117 | Fax: 650-654-4233 | www.checkpoint.com