**DESIGN
GUIDE**

Securing Applications with Cloud NGFW for Azure

AUGUST 2023

Table of Contents

Preface	1
Purpose of This Guide	3
Audience	3
Related Documentation	3
Introduction	4
Public-Cloud Concepts	7
Scaling Methods	7
Reduced Time to Deployment	7
Cloud Infrastructure Protection	8
Azure Concepts and Services	9
Azure Region and Availability Zones	9
Virtual Networks	10
Virtual WAN	16
Palo Alto Networks Azure Design Details	21
Cloud NGFW Components	22
Management	23
Deployment Methods	30
Securing Traffic with the Cloud NGFW	32
Design Model	33
Virtual Network	33
Virtual WAN	41
Summary	49

Preface

GUIDE TYPES



Overview guides provide high-level introductions to technologies or concepts.

Design guides provide an architectural overview for using Palo Alto Networks® technologies to provide visibility, control, and protection to applications built in a specific environment. These guides are required reading prior to using their companion deployment guides.

Deployment guides provide decision criteria for deployment scenarios, as well as procedures for combining Palo Alto Networks technologies with third-party technologies in an integrated design.

Solution guides provide add-on solutions for post-deployment use cases.

DOCUMENT CONVENTIONS



Notes provide additional information.



Cautions warn about possible data loss, hardware damage, or compromise of security.

Blue text indicates a configuration variable for which you need to substitute the correct value for your environment.

In the IP box, enter **10.5.0.4/24**, and then click **OK**.

Bold text denotes:

- Command-line commands.

show device-group branch-offices

- User-interface elements.

In the **Interface Type** list, choose **Layer 3**.

- Navigational paths.

Navigate to **Network > Virtual Routers**.

- A value to be entered.

Enter the password **admin**.

Italic text denotes the introduction of important terminology.

An *external dynamic list* is a file hosted on an external web server so that the firewall can import objects.

Highlighted text denotes emphasis.

Total valid entries: **755**

ABOUT PROCEDURES

These guides sometimes describe other companies' products. Although steps and screen-shots were up-to-date at the time of publication, those companies might have since changed their user interface, processes, or requirements.

GETTING THE LATEST VERSION OF GUIDES

We continually update reference architecture guides. You can access the latest version of this and all guides at this location:

<https://www.paloaltonetworks.com/referencearchitectures>

WHAT'S NEW IN THIS RELEASE

Since the last version of this guide, Palo Alto Networks made the following changes:

- This is a new guide.

Purpose of This Guide

Cloud NGFW for Azure is a Palo Alto Networks next-generation firewall (NGFW) delivered as a cloud-native Microsoft Azure service. This guide describes reference architectures for deploying Cloud NGFW, bringing visibility, control, and protection to applications built on Azure.

This guide:

- Links the technical design aspects of Azure and the Cloud NGFW for Azure solution and then explores several technical design models. The design models include two options that span the scale of enterprise-level operational environments.
- Provides a framework for architectural discussions between Palo Alto Networks and your organization.
- Provides decision criteria for subscription and deployment scenarios, as well as procedures for configuring features of Cloud NGFW for Azure.

AUDIENCE

This guide is for technical readers, including system architects and design engineers, who want to deploy the Cloud NGFW for Azure. This guide assumes the reader is familiar with the basic concepts of applications, networking, virtualization, security, and high availability. The reader should also possess an understanding of Azure network-architecture concepts.

RELATED DOCUMENTATION

The following documents support this guide:

- **Securing Applications with Cloud NGFW for Azure—Virtual WAN Design: Deployment Guide**—Provides deployment guidance for the Cloud NGFW for Azure reference architecture for securing Virtual WAN.
- **Securing Applications with Cloud NGFW for Azure—Virtual Network Design: Deployment Guide**—Provides deployment guidance for the Cloud NGFW for Azure reference architecture for securing Virtual Networks.
- **Panorama on Azure: Deployment Guide**—Provides architectural guidance and deployment details for using a Palo Alto Networks Panorama® virtual appliance, deployed on Microsoft Azure, to provide a single location from which you can create network configurations and security policies that enable visibility, control, and protection to your applications built in an Azure public cloud.
- **Zero Trust Enterprise: Design Guide**—Describes the Zero Trust Enterprise approach to securing users, applications, and infrastructure by eliminating implicit trust and continuously validating every stage of a digital interaction.

Introduction

Organizations are expanding their public-cloud initiatives in a variety of ways, with security remaining top-of-mind. Increased use dictates an effort for more streamlined security workflows and an eye toward cloud-centric architectures that are scalable and resilient.

More workloads reside in public clouds than ever, and the use of the public cloud is increasing dramatically, leading to multi-cloud environments and increased demand for capacity. Security, traditionally viewed as a bottleneck that slows deployment, must more readily support the move toward cloud-centric architectures.

Securing your Azure workloads introduces a range of challenges, including a lack of application visibility, inconsistent security functionality, and difficulty keeping pace with the rate of change commonly found in cloud-computing environments. To be successful, organizations need a public-cloud security solution that does the following:

- Identifies and controls applications based on identity, not the ports and protocols they use
- Stops malware from gaining access to and moving laterally (east-west) within the cloud
- Determines who should be allowed to use the applications and grants access based on need and credentials
- Simplifies management and minimizes the security policy lag when you create, delete, or move VMs within the cloud environment

Palo Alto Networks Cloud NGFW for Azure enables you to protect your public-cloud workloads from cyber threats with our next-generation firewall security and advanced threat prevention features. In addition, when using Panorama network security management, combined with native automation features, allows you to streamline policy management in a manner that minimizes the lag time that may occur when you create, delete, or move workloads.

Applying Palo Alto Networks Next-Generation Security to Public-Cloud Environments

The Cloud NGFW uses the same full-stack traffic classification engine that you can find in our PA-Series and VM-Series firewalls. The Cloud NGFW natively classifies all traffic, inclusive of applications, threats, and content, and then ties that traffic to the user. The app, content, and user—the three elements that run your business—form the basis of your virtualized security policies, resulting in improved security posture and reduced incident response time.

Isolating Mission-Critical Applications and Data by Using Zero Trust Principles

Security best-practices dictate that you should isolate your mission-critical applications and data in secure segments by using Zero Trust (never trust, always verify) principles at each segmentation point. You can deploy the VM-Series throughout your virtualized environment, with the firewall residing as a gateway within your virtual network or between infrastructure VLANs protecting traffic by exerting control based on application and user identity.

Blocking Lateral Movement of Cyberthreats

Cyberthreats commonly compromise an individual workstation or user and then move across the network, looking for a target. Within your virtual network, cyberthreats rapidly move laterally from VM to VM, in an east-west manner, placing your mission-critical applications and data at risk. Exerting application-level control by using Zero Trust principles in between VMs reduces the threat footprint while applying policies to block both known and unknown threats.

Preventing Data Exfiltration

Many cyberattacks are designed to steal customer information or intellectual property that attackers can monetize through corporate blackmail or illicit sale. Just as workloads often must communicate with shared services and each other, some applications (for example, email solutions and collaboration platforms) need to access the public internet, which opens up avenues for data exfiltration. Hybrid cloud architectures pose similar vulnerabilities. To secure traffic between trust zones, the recommended practice is to deploy next-generation firewalls—augmented with capabilities such as DNS Security and URL Filtering—to help guard against data exfiltration.

Benefits of Palo Alto Networks Solutions in Public-Cloud Environments

For organizations that depend upon public and hybrid cloud environments, the Cloud NGFW offers critical business benefits:

- **Consistent security for ongoing compliance**—The Cloud NGFW allows network security teams to manage network security and threat-prevention policies for their Azure environments in the same way that they manage their physical and private cloud environments. Panorama also enables you to manage your Cloud NGFW, VM-Series, and your PA-Series physical security appliances, ensuring policy consistency and cohesiveness. Rich, centralized logging and reporting capabilities provide visibility into virtualized applications, users, and content.
- **Virtual machine monitoring**—As your virtual machines change functions or move from server to server, building security policies based on static data (such as IP address) delivers limited value and can contain outdated information. Security policies must be able to monitor and keep up with changes in virtualization environments, including VM attributes and the addition or removal of VMs. The Panorama Plugin for VMware vCenter automatically polls vCenter for virtual-machine inventory and changes, collecting this data in the form of tags. Dynamic address groups (DAGs) allow you to create policies by using the tags as identifiers for virtual machines instead of a static object definition. Multiple tags representing virtual-machine attributes, such as IP address and operating system, can be resolved within a dynamic address group. DAGs allow you to easily apply policies to new virtual machines or virtual machines moving across hypervisor hosts without administrative intervention.
- **Extended application support**—The Cloud NGFW provides application visibility across all ports, providing relevant information about the network flows in the environment to help security managers make rapid, informed policy decisions. Palo Alto Networks supports more than 3,000 applications and custom application templates that you can easily integrate.
- **Compliance**—Many regulatory standards, such as the Payment Card Industry Data Security Standard, require both segmentation and intrusion detection systems to secure cardholder information from the rest of the environment. The Cloud NGFW built-in intrusion protection system (IPS) capabilities allow network managers to meet these requirements without additional components.

The Cloud NGFW for Azure is ideal for public-cloud deployments where the firewall-as-a-service (FWaaS) form-factor might simplify deployment and provide flexibility while protecting your applications and data. The following sections first describe the required Azure components and then present the solution's necessary Palo Alto Networks components.

Public-Cloud Concepts

Organizations generally move to the public cloud with the goals of increasing scale and reducing time to deployment. Achieving these goals requires application architectures that are built specifically for the public cloud. Before you can architect for the public cloud, you must understand how it is different from traditional on-premises environments.

SCALING METHODS

Traditionally, organizations scale on-premises deployments through the purchase of devices that have increased performance capacity. Scaling up an on-premises deployment in this method makes sense because organizations typically purchase the devices in order to satisfy the performance requirements during the devices' lifetime.

Public-cloud environments focus on scaling out the deployment instead of scaling up. This architectural difference stems primarily from the capability of public-cloud environments to dynamically increase or decrease the number of resources allocated to your environment. In the public cloud, infrastructure used to satisfy performance requirements can have a lifetime in minutes instead of years. Instead of purchasing extra capacity for use at some time in the future, the dynamic nature of the public cloud allows you to allocate just the right amount of resources required to service the application.

In practice, to architect an application for the cloud, you need to distribute functionality, and you should build each functional area to scale out as necessary. Typically, this means a load-balancer distributes traffic across a pool of identically configured resources. When changes occur in the application traffic, the number of resources you have allocated to the pool can dynamically increase or decrease. This design method provides scale and resiliency. However, the application architecture must take into account that the resources are transient. For example, you should not store the application state in the networking infrastructure or in the frontend application servers. Instead, store state information on the client or persistent storage services.

The ability to scale a cloud architecture extends not only to the capacity of an application but also the capacity to deploy applications globally. Scaling an application to a new region in a traditional on-premises deployment requires significant investment and planning. Public-cloud architectures are location-agnostic, and you can deploy them globally in a consistent amount of time.

REDUCED TIME TO DEPLOYMENT

To achieve the goal of reduced time to deployment, you must have a development and deployment process that is repeatable and reacts to changes quickly. DevOps workflows are the primary method for implementing this process. DevOps workflows are highly dependent on the ability to automate, as much as possible, the process of deploying a resource or application. In practice, this means you must be able to programmatically bootstrap, configure, update, and destroy the cloud infrastructure, as well as the

resources running on it. Compared to traditional on-premises deployments where device deployments, configurations, and operations happen manually, automated workflows in a public-cloud environment can significantly reduce time to deployment.

Automation is so core to cloud design that many cloud application architectures deploy new capabilities through the automated build-out of new resources instead of updating the existing ones. This type of cloud architecture provides several benefits, not the least of which is the ability to phase in the changes to a subset of the traffic, as well as the ability to quickly roll back the changes by redirecting traffic from the new resources to the old.

CLOUD INFRASTRUCTURE PROTECTION

Azure provides basic infrastructure components and has a responsibility to ensure that each customer's workloads are appropriately isolated and ensure that the underlying infrastructure and physical environment are secure. However, the customer has the responsibility to securely configure the instances, operating systems, and any necessary applications, as well as maintain the integrity of the data each virtual machine processes and stores. This shared-responsibility model is often a point of confusion for consumers of cloud services.

Services have default configurations that might be secure upon implementation, but to ensure the integrity of the data itself, it is up to the customer to make the assessment and lock those service configurations down.

Security and compliance risks in cloud computing threaten an organization's ability to drive digital business. The dynamic nature of the cloud, the potential complexity of having multiple cloud service providers in the environment, and the massive volume of cloud workloads make security and compliance cumbersome.

Azure Concepts and Services

This section provides an overview of Azure Virtual Networking and Virtual WAN along with some relevant networking services pertaining to building a secure network architecture. For additional information, see the Microsoft Azure documentation as the definitive source of information on these topics.

AZURE REGION AND AVAILABILITY ZONES

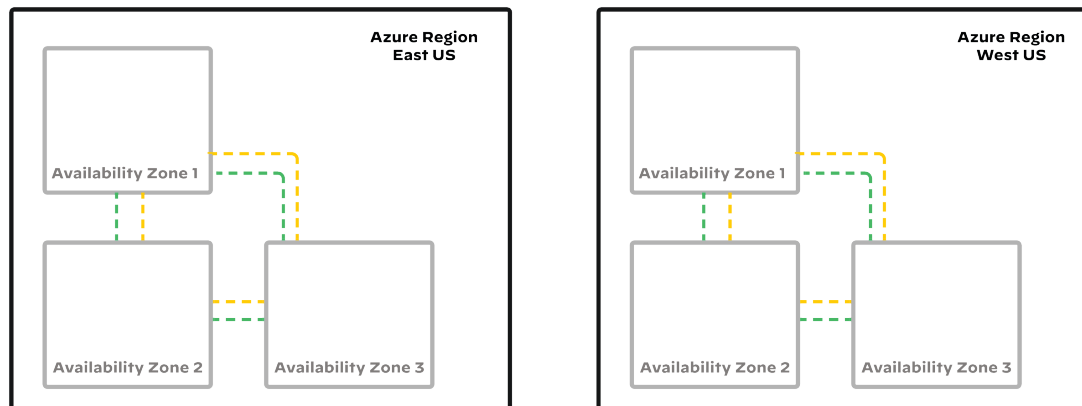
Azure *regions* are strategically located around the world, with each region consisting of multiple data centers. These data centers are deployed within a latency-defined perimeter, meaning they are located in close proximity to each other to minimize network latency. To ensure optimal performance and security, these data centers are connected through a dedicated regional low-latency network. This network allows Azure services within the same region to communicate with each other quickly and efficiently. By having data centers within the same region, Azure can provide high-speed and low-latency connections between its services, resulting in faster response times and improved performance for users. This design can also help with security architecture by keeping data within a specific geographical area and reducing the risk of data breaches or unauthorized access.

In the world of technology and infrastructure, failures can happen unexpectedly. Whether it's a software glitch or a natural disaster like an earthquake or flood, it's crucial to have measures in place to ensure resilience and continuity. Azure *availability zones* provide redundancy and logical isolation of services. In regions where availability zones are enabled, your data and services can have the added protections allowed by using multiple availability zones.

To ensure resiliency, Azure implements a minimum of three separate availability zones in all availability zone-enabled regions. With a round-trip latency of less than 2 ms, a high-performance network interconnects these zones. This network connectivity ensures that your data stays synchronized and accessible even in the face of failures. Each availability zone is made up of one or more data centers, each equipped with independent power, cooling, and networking infrastructure. This level of separation and redundancy means that even if one availability zone is affected, the remaining two zones can continue to support regional services, maintain capacity, and provide high availability.

Many organizations have a critical need for high availability and protection against various types of disasters. Azure regions are specifically designed to address these concerns by offering availability zones to mitigate localized disasters and disaster-recovery options to mitigate regional or large-scale incidents.


Figure 1 Azure regions and availability zones



VIRTUAL NETWORKS

A *virtual network* (VNet) is a logically segmented network within Azure that allows connected resources to communicate with each other. A VNet contains one or more public or private IP address ranges that you can then divide into subnets (/29 or larger). VNet IP address space, both public and private, is reachable only within the VNet or through services connected to the VNet, such as a VPN. Because VNets are isolated from each other, you can overlap IP network definition across VNets.

Figure 2 Single VNet



Virtual machine network interfaces receive IP addresses, default gateways, and DNS servers from the Azure DHCP service. By default, when you start a virtual machine, DHCP dynamically assigns the first available IP address in the subnet to the virtual machine's interface. If you created a VM through the Resource Manager, the IP address does not change when the VM reboots or remains in a stopped state. If you used a classic deployment, then the IP address could change after the VM is restarted from being in a stopped state. In either case, the assigned address is released when the VM is deleted.

Static IP addressing is available when a persistent IP address is required. When there is only one IP on an interface, you do not need to configure static IP addresses in the operating system running on the virtual machine. Instead, you configure the IP address as static in the Azure portal or the template. When you configure a static IP address, the virtual machine still receives the IP address through DHCP.

**Note**

Azure reserves five internal IP addresses from each subnet. You cannot configure these IP addresses on a resource: the first and last addresses of the address space (for the subnet address and multicast) and three addresses reserved for internal use (for DHCP and DNS purposes).

An alternative to static IP addressing for persistent connectivity is the use of name resolution to communicate between resources within a VNet. The Azure DNS servers provide not only public name resolution but also internal name resolution within the VNet. The addition or state change of a virtual machine automatically updates Azure name resolution. If a virtual machine has multiple internal IP addresses, its name resolves to its primary IP address.

Although VNets do not contain publicly routable IP addresses, you can associate resources within a VNet with a publicly routable IP address. You can map public IP addresses to internal IP addresses using a one-to-one assignment, and Azure networking automatically translates the IP addressing of the network flow as it enters and leaves the VNet.

Depending on the configuration, public IP addresses can change as a virtual machine changes state. You can configure a public IP address to be dynamic or static, but in most situations, configuring a DNS name label on the public IP address is the preferred way to ensure persistent connectivity if the underlying IP address can change.

If you need multiple public IP addresses on a single network interface, you can configure one or more secondary IP addresses. You can then associate a public IP address to each secondary IP address on the interface. Because DHCP cannot assign multiple IP addresses to a single interface, you should statically define secondary IP addresses in Azure and configure them in the virtual machine operating system.

**Note**

All resources deployed in a VNet have unfiltered outbound access to the internet, even when they do not have a public IP address directly assigned. Azure automatically translates the IP addresses of outbound traffic to the internet. Assigning a public IP address enables inbound connectivity to the resource using a known IP address or DNS name.

Virtual Network Peering


You can separate workloads across functional environments or administrative domains using multiple VNets. Virtual network peering allows you to logically connect Azure VNets to provide resource connectivity. However, you cannot use overlapping IP address space within the group of peered VNets.

Two types of peering are supported:

- **VNet peering**—The connected VNets are located within the same Azure region (example: West US).
- **Global VNet peering**—The connected VNets are located across multiple Azure regions (example: West US and East US). Some Azure networking capabilities are restricted when using Global VNet peering. For more information, see the Azure documentation for [Virtual Network Peering](#).

You achieve full IP reachability between VNets after you establish the peering connection. All network traffic using VNet peering remains within the Azure backbone. Azure supports 1000 virtual networks within a subscription and supports up to 500 peering connections for each virtual network.

Figure 3 Multiple peered VNets




Traffic Forwarding

Azure uses a Layer 3 overlay network fabric that does not operate in the same fashion as traditional Layer 2 packet forwarding. However, the Azure network infrastructure attempts to make the operational differences transparent. For example, even though you cannot ping the default router or use traceroute during troubleshooting, virtual machines still receive a default route through DHCP.

All resources connected to a VNet communicate directly via the overlay network, even if they are on different IP subnets. When you add or change a subnet, Azure automatically defines system routes to facilitate communication within the VNet as well as to the internet. Traffic-forwarding behavior with peered VNets is essentially the same as within a single VNet. Azure installs system routes for the defined address space of the peered VNets into the active forwarding table for each subnet.

Figure 4 Azure networking



Azure networking includes system routes with the following next-hop types:

- **Virtual network**—System route to a destination prefix within the local VNet address space. Azure automatically creates this route type by default when an address space is defined. You may manually create additional routes of this type as necessary.
- **VNet peering**—System route to a destination prefix within a peered VNet address space. Azure automatically creates this route type by default when a peer connection is established.
- **Internet**—System route to the internet. When you create a VNet, Azure automatically creates a default system route that matches any destination prefix using a wildcard match (0.0.0.0/0). You can manually create additional, more-specific routes to the internet as necessary.
- **None**—Special system route to drop traffic for a specified destination prefix. Azure automatically creates these system routes for RFC-1918 and RFC-6598. You can manually create additional routes of this type in order to drop traffic to other destination prefixes as necessary.
- **Virtual appliance**—Manually created system route for a specified destination prefix with a specified next-hop IP address. You must assign the next-hop address to a virtual device deployed within a VNet (or peered VNet).
- **Virtual network gateway**—System route to a destination prefix assigned to a virtual network gateway connection. Azure automatically creates this system route when you configure the connection.

Table 1 Azure routes example

Address space	Address prefix	Next-hop type
VNet defined	10.110.0.0/16	Virtual network
Peer VNet	10.112.0.0/16	VNet peering
Default (Azure defined)	0.0.0.0/0	Internet
RFC-1918 (Azure defined)	10.0.0.0/8	None
RFC-1918 (Azure defined)	172.16.0.0/12	None
RFC-1918 (Azure defined)	192.168.0.0/16	None
RFC-6598 (Azure defined)	100.64.0.0/10	None

User-Defined Routes

User-defined routes (UDRs) modify the default traffic-forwarding behavior of Azure networking. You primarily use a UDR to direct traffic to a resource, such as a load balancer or a firewall, within the VNet or in a peered VNet. You can also configure them to send traffic to a VPN connection or to be used as a null route to discard unwanted traffic.

Within Virtual WAN, VNets are peered to the Virtual Hub. These VNets that are connected to the Virtual Hub cannot contain VNGs or subnets with VNGs, i.e. gateway subnets. Site to Site, or branch to Azure VPNs will connect directly to the Virtual Hub in each respective region.

You configure a UDR on a per-subnet basis, and it applies to traffic that is sent within the subnet. The destination for a route can be a different subnet in the VNet, a different subnet in another VNet (with an existing peer connection), anywhere on the internet, or a private network connected to the VNet. The next hop for the route can be any resource in the VNet or in a peered VNet in the same region.



Note

The use of UDR summary routes can have unexpected consequences. If you apply a UDR summary route to a subnet that falls within the summary but does not have a more specific UDR or system route, UDR controls traffic within the subnet (host to host).

To view the active system routes for a route table that you have applied to a subnet, you can use the Effective Routes troubleshooting tool. The tool is available under the settings pane of any network interface.

Network Security Groups

Network security groups (NSGs) filter ingress and egress network traffic. You can associate an NSG to a subnet, to the network interface of a virtual machine, or to both. For ease of configuration, when you apply the same policies to more than one resource, you can associate a single network security group with multiple subnets or virtual machine network interfaces. An NSG associated to the subnet with a common policy can be easier to manage than unique NSGs applied at the interface level.



Note

When you are using Standard SKU IP addresses, NSGs are required on the subnet or NIC in order for traffic to reach the resource.

You create a prioritized list of rules that defines the behavior of a network security group. Rules are defined and matched by the traffic source, destination, port, and protocol. In addition to IP addressing, you can set the source and destination of a rule by using Azure service tags or application security groups. There are separate policies for inbound and outbound traffic flows.

Network security groups are pre-configured with default inbound and outbound security rules that perform the following tasks:

- Allowing all traffic within the VNet
- Allowing outbound traffic to the internet
- Allowing inbound traffic that is originating from Azure's load-balancer probe (168.63.129.16/32)
- Denying all other traffic

You cannot modify or remove the default security rules. To override the behavior of a default rule, you must precede it with a custom rule that has a higher priority. The default rules have priority values that begin at 65000, so you must assign a priority value less than 65000 in order to override the default rules. You can view the active rules by using the Effective Security Rules troubleshooting tool in the network security group or network interface settings.

Enterprise Network Connectivity


A *virtual network gateway* (VNG) provides connectivity between an Azure virtual network and your enterprise networks and data centers. Site-to-site connectivity through a VNG can either be through IPsec VPN (also known as a *VPN gateway*) or a dedicated private connection (also known as an *ExpressRoute gateway*). When you deploy a virtual network gateway as a VPN gateway, it supports the configuration of IPsec VPN tunnels to one or more of your locations across the internet. When deployed as an ExpressRoute gateway, the VNG provides connectivity to enterprise locations through a dedicated private circuit facilitated by a service provider.

**Note**

You can only deploy one virtual network gateway of each type in a virtual network, and you deploy them in a dedicated gateway subnet.

A VPN gateway is composed of pair of virtual machines deployed, by default, in an active/standby configuration for resiliency. You may configure IP routing between Azure and the enterprise network to be static or dynamically exchanged through IKEv2 or BGP. A VPN gateway supports configurations with multiple tunnels to a single location, providing additional connection resiliency for deployments with resilient on-premises VPN devices. Active/active configuration is also possible on select VPN gateway sizes.

Figure 5 VPN gateway



Connections through an Azure ExpressRoute gateway do not connect through or traverse the public internet. Instead, Azure resources are accessed directly through colocation facilities, point-to-point Ethernet, or connectivity into your MPLS WAN service. Microsoft recommends ExpressRoute connections for all enterprise customer connectivity and supports a range of bandwidth options from 50 Mbps to 10 Gbps.

For resiliency, Azure terminates ExpressRoute connections on a pair of edge routers. Two BGP connections provide resilient, dynamic IP routing between Azure and your enterprise networks. You can share ExpressRoute circuits across multiple VNETs. In each virtual network that requires backhaul over the ExpressRoute gateway, a VNG connects the VNet to the ExpressRoute circuit.

VIRTUAL WAN

Azure Virtual WAN facilitates central connectivity for users, data centers, remote sites, and branches to Azure. This enables customers to use Microsoft's global network, potentially reducing latency and providing a centralized routing and networking solution contained in each region's virtual hub. Azure Virtual WAN is a hub-and-spoke software technology that is designed to provide a connection point for multiple sites or users into Microsoft Azure. To establish a full mesh connectivity into Azure, remote sites

in Virtual WAN connect to the Microsoft backbone. Each Azure region can serve as a hub. These hubs are a fully connected mesh. A hub enables endpoint connectivity that is transitive and allows for distribution across diverse types of remote-site connections. Azure Virtual WAN allows for users to set up connectivity with existing SD-WAN or VPN-capable devices, or users can manually configure connections in Azure Virtual WAN.

There are two types (SKUs) of Virtual WAN: Basic and Standard. The biggest difference between the Basic and Standard types is that with Basic Virtual WAN, the hubs are not meshed, and with Standard Virtual WAN, all hubs are meshed automatically when you initially set up the Virtual WAN.

Figure 6 Virtual WAN overview


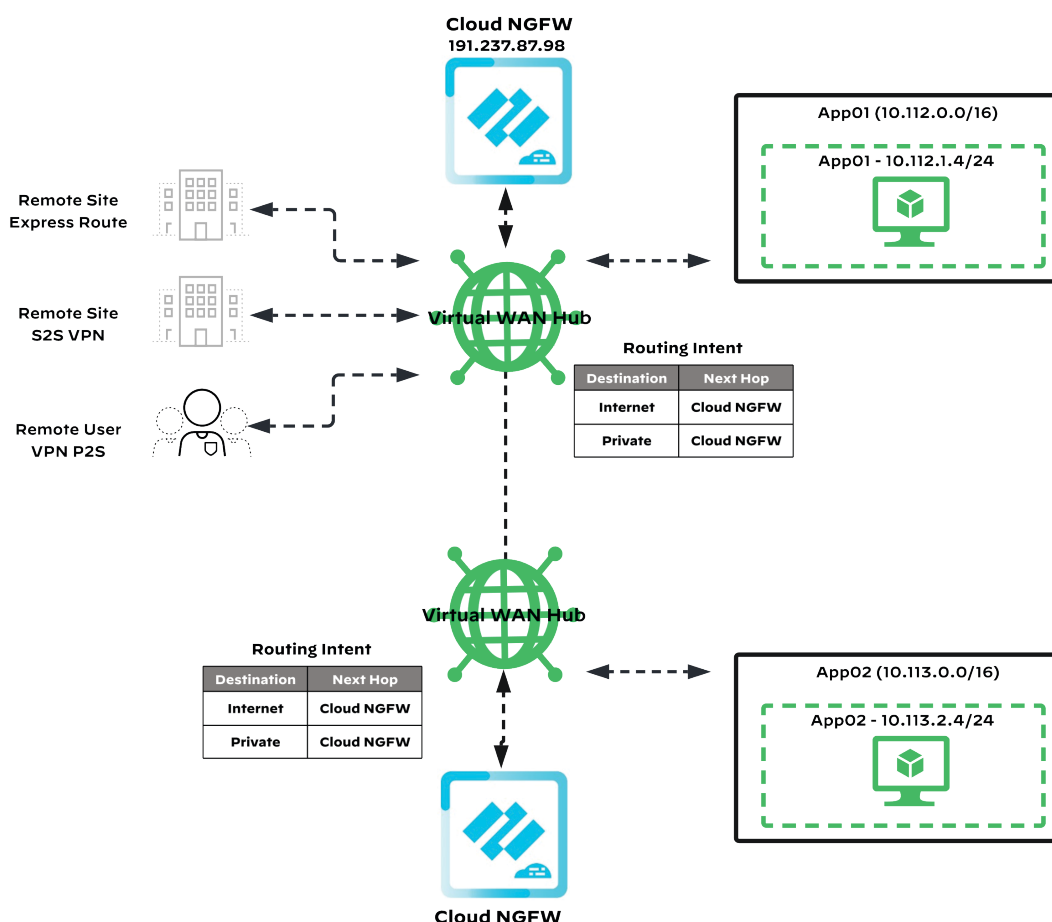


Figure 7 Multi-region Virtual WAN



Routing

To manage the routing needs for Virtual WAN, a router uses BGP to provide the dynamic routing required between gateways. This router also manages transit connectivity between virtual networks that connect to a virtual hub and can manage an aggregate throughput of 50 Gbps. This router is used for customers with the Standard type Virtual WAN, not the Basic type.

Contained within Virtual WAN is a virtual hub router that provides a central location to manage routing. The virtual hub router also provides functionality such as route propagation and association, along with the ability to create route tables. Using a custom algorithm, the virtual router can make path selection decisions for traffic going to Azure and to virtual-hub remote sites. Azure has a feature within Virtual WAN hub called *Hub routing preference*. This feature allows the virtual hub router to use its route-selection algorithm and adapt its path selection based on the hub routing preference configuration for any routes learned from ExpressRoute, site-to-site VPN, SD-WAN, and user VPN connections into Virtual WAN.

The virtual hub router selects paths based on the following criteria, in the order shown:

1. Selecting routes with the longest prefix match
2. Preferring static routes over BGP
3. Using the hub routing preference configuration (ExpressRoute, VPN, or AS Path)
4. Using equal-cost multi-path routing to load-balance the traffic across the multiple routes

Hub

The *Virtual WAN hub* is the regional connection point for Azure's Virtual WAN technology. The virtual hub is the central core for a virtual WAN network in each region. It is possible to connect hubs together between Azure regions and to create multiple virtual hubs per region.

Inside a virtual hub is a hub gateway. The *hub gateway* is different from the virtual network gateway (VNG) that is used for ExpressRoute, and the hub gateway is a separate type of instance from the VPN gateway. Your VNets do not need their own virtual network gateway. Instead, the connectivity is direct through the hub. You can set up the gateway when you create the virtual hub, or you can configure an empty virtual hub (one that does not have any gateways) and add or connect the gateway later.

Upon creation of a virtual hub, a virtual hub router is deployed. This router within the context of Virtual WAN hub is the central piece that manages all routing decisions between VNets and gateways. When you are deploying a virtual hub, sizing is one of your first decisions. You choose the size and capacity of your virtual hub based on two factors:

- How many routing infrastructure units are required
- The required aggregate bandwidth for all VMs that reside in VNets that will be accessed from the virtual hub

If in your Azure environment you have 4000 VMs, you need to size the virtual hub accordingly based on this number and your expected throughput.

The virtual hub has several key components:

- **Hub route table**—You use a *virtual hub route table* to configure path selection and connectivity within, as well as to and from, a Virtual WAN hub. By default, all connections associate and propagate with the default route table.
- **Hub-to-hub connection**—Virtual WAN hubs are regional constructs. It is possible to connect multiple hubs together. You do this with a hub-to-hub connection. Hubs within the same Virtual WAN can be configured with unique regional settings and policies.
- **Hub virtual connection**—You use a *hub virtual connection* to provide connectivity from a Virtual WAN hub to a VNet or virtual network within Azure. At this time, it is only possible for a VNet to be connected to one virtual hub. The hub virtual connection is a key connectivity point leveraged in this solution guide. This is where the transit VNet connects and exchanges routing information with Virtual WAN hub.
- **Association**—Every connection in a virtual hub deployment has an association with one route table. This association allows traffic to be sent to the desired destination indicated by the configured routes in the table. Each connection shows its route table association. It is also typical to have multiple connections associated with the same route table. All VPN, ExpressRoute, and User VPN connections are associated to the same (default) route table. Each branch connection must be associated with the default route table. Each branch or remote site therefore knows about all associated networks. Each ExpressRoute, point-to-site, or site-to-site connection needs an association with the default route table. By default, all connections are associated to a default route table in a virtual hub. Each virtual hub has its own default route table, which you can edit to add a static route(s). Static routes are preferred over dynamically learned routes for the same prefixes.
- **Propagation**—Connections dynamically send routes to a route table. With a VPN connection, ExpressRoute connection, or P2S configuration connection, routes are propagated from the virtual hub to the on-premises router using BGP. Connections can send routes to one or multiple route tables. To ensure that all connections can reach all remote sites and branches, each branch connection needs to propagate its routes to the same set of route tables. If you configure remote sites to propagate their routes to the default route table, then you must do this for all remote sites.
- **Labels**—*Labels* group route tables. This can be especially helpful when propagating routes from connections to multiple route tables. The default route table has a preconfigured label called *Default*. When you propagate connection routes to the Default label, these routes automatically apply to all default route tables across every virtual hub in the Virtual WAN.
- **Static route**—Configuring static routes provides a means of steering traffic through a next-hop IP, which could be of a network virtual appliance (NVA) provisioned in a remote-site VNet attached to a virtual hub. Static routes take precedence over dynamically learned routes for the same prefixes.
- **User-defined routes**—UDRs (or *custom routes*) use connections as the next hop. These connection routes are UDRs assigned to a VNET connection. They direct traffic to an NVA such as the Palo Alto Networks Cloud NGFW or VM-Series firewalls for inline security.

Palo Alto Networks Azure Design Details

Palo Alto Networks Cloud NGFW is an Azure Native ISV Service. This approach allows Palo Alto Networks to develop and manage the FWaaS allowing the user to take advantage of the native Azure UI and APIs. The Cloud NGFW for Azure is accessible in Azure Marketplace. Cloud NGFW provides the best-in-class security capabilities to which Palo Alto Networks customers are accustomed, as a cloud-native service. Many capabilities from PAN-OS® and PA-series and VM-series firewalls are also available, including App-ID™, Threat Prevention, Advanced URL Filtering, and inbound and outbound decryption.

Cloud NGFW is a machine language–powered NGFW delivered in Azure as a fully managed cloud-native service. Cloud NGFW enables advanced application visibility and access control while also including App-ID and URL-filtering technologies. Cloud NGFW includes threat prevention and detection, cloud-delivered security services, and threat-prevention signatures. With this true cloud-native experience, you are able to run an array of applications simultaneously, keeping the applications secure at the speed and scale of the cloud.

To deliver a fully managed, cloud-native service, Cloud NGFW combines best-in-class network security with ease of use. Natively integrated into cloud providers' various service offerings, Cloud NGFW extends Palo Alto Networks threat-prevention capabilities to cloud providers.

Cloud NGFW for Azure provides many uses and key benefits. First and foremost is next-generation firewall security. This includes stopping zero-day attacks and web-based threats in real time, as well as securing applications outbound access to legitimate services. Securing inbound and outbound (north-south) traffic in the cloud is key to providing a secure foundation with which you build a cloud architecture. Cloud NGFW stops web-based attacks, vulnerabilities, exploits, and known evasions—including sophisticated file-based attacks—by using patented, proven App-ID traffic classification technology built upon years of research and development. Cloud NGFW also secures laterally moving traffic, such as traffic crossing trust boundaries within Azure (for example, VNet-to-VNet communications or traffic moving in your Virtual WAN). Cloud NGFW secures this east-west traffic by blocking attackers from gaining access to resources and stopping data exfiltration and command-and-control traffic. Purpose-built to stop unauthorized or east-west lateral movement, Cloud NGFW has the NGFW capabilities that you expect from Palo Alto Networks, including App-ID, URL filtering based on URL categories and geolocations, and SSL/TLS decryption.

Automation is key to building a secure and operational cloud environment. Cloud NGFW integrates security with workflows that are managed by cloud providers. With Cloud NGFW, the first next-generation firewall to integrate with cloud providers, you do not need to have lengthy deployment cycles and are able to get up and running quickly. This is all possible even when setting up required rulestacks and automated security profiles. You can leverage the security model provided by the chosen cloud provider while integrating with their onboarding, monitoring, and logging capabilities. Cloud NGFW provides a unique benefit when integrating with cloud providers. You can take advantage of automatic scaling and high availability while the maintenance occurs, without the need for end-user involvement. This integration enables consistent firewall policy management across multiple cloud-provider accounts.

Palo Alto Networks built Cloud NGFW with automation in mind. With rulestack configuration and automated security profiles, Cloud NGFW meets network security requirements easily, with an intuitive user interface that simplifies the creation of resilient firewall resources that scale with your network traffic.

CLOUD NGFW COMPONENTS

Cloud NGFW Service

To ensure the security of your Azure environment, Cloud NGFW for Azure consists of several components that work together. At the heart of this solution is the Cloud NGFW service, which serves as the firewall and provides robust protection.

When setting up the Cloud NGFW, you have the option to initially provision it as either a VNet or a Virtual WAN network resource type. This flexibility allows you to choose the most suitable option for your specific requirements. The Cloud NGFW is designed with built-in resiliency, scalability, and life-cycle management features, ensuring a reliable and efficient operation. It can span across multiple availability zones and is deployed at a regional level.

Palo Alto Networks integrates the Cloud NGFW as the resource associated with either the VNet or the Virtual WAN hub. During the deployment process, you assign the Cloud NGFW private IP addresses within the specified NGFW subnet. After you have provisioned it, you can configure routing updates by using UDRs in the Virtual Network design and Routing Intent Policies in the Virtual WAN. This allows you to direct traffic through the private IP addresses associated with the Cloud NGFW, ensuring secure and efficient data flow.

Another critical component of the Cloud NGFW is the rulestack. Rulestacks are responsible for defining the traffic-filtering behavior of the NGFW, including advanced access-control features like App-ID and URL Filtering, as well as threat-prevention measures. A rulestack comprises a set of security rules, along with associated objects and security profiles. One or more firewalls can use a rulestack. By associating a rulestack with one or more Cloud NGFW resources, you can enforce consistent security policies across your Azure environment.

Subscription Model

The Cloud NGFW from Palo Alto Networks is delivered as a cloud managed service. This means subscribing to and using the Cloud NGFW is very simple and done via an Azure PAYG subscription. No additional licenses or subscriptions are required.

Availability and Scaling

When you deploy a Cloud NGFW, you are actually deploying a fully redundant solution that scales automatically with your environment. A Cloud NGFW consists of a dedicated subscription and single-tenanted Cloud NGFW resource dedicated to a customer's VNet or Virtual WAN. This resource has built-in

resiliency, scalability, and life-cycle management. Cloud NGFW spans across multiple Azure availability zones. Each Cloud NGFW has the capacity to scale up to 50 Gbps throughput with autoscaling capacity functionality. The Cloud NGFW continuously meters the usage of the Cloud NGFW resource also sending usage records for each Azure subscription to the Azure metering service. This service is responsible for billing.

Logging

A log is an automatically generated, time-stamped file that provides an audit trail for system events on the firewall or for network traffic events that the firewall monitors. Log entries contain artifacts, which are properties, activities, or behaviors associated with the logged event, such as the application type or the IP address of an attacker. Each log type records information for a separate event type. For example, the firewall generates a Threat log to record traffic that matches a spyware, vulnerability, or virus signature or a DoS attack that matches the thresholds configured for a port-scan or host-sweep activity on the firewall.

The Cloud NGFW can send traffic, threat, and decryption logs to either an Azure Log Analytics Workspace (which you can create from the Azure portal) or directly to Panorama. The Log Analytics Workspace is associated with a workspace ID, primary key, and a secondary key that is retrieved through the logging API via the control plane. Azure Log Analytics Workspace is a repository that is configured to store log data from multiple Azure services. You can use a single workspace for all of your data or create multiple workspaces based on your requirements. For Cloud NGFWs managed with Panorama, you can forward logs either to Panorama or Azure.

Cloud NGFW can capture and save three types of logs:

- **Traffic**—Traffic logs display an entry for the start and end of each session.
- **Threat**—Threat logs display entries when traffic matches one of the security profiles attached to a security rule on the firewall. Each entry includes the following information: date and time, type of threat (such as virus or spyware), threat description or URL (Name column), alarm action (such as allow or block), and severity level.
- **Decryption**—By default, decryption logs display entries for unsuccessful TLS handshakes. These logs can display entries for successful TLS handshakes, if you enable this in decryption policy. If you enable entries for successful handshakes, ensure that you have the system resources (log space) for the logs.

MANAGEMENT

When managing Cloud NGFW, you have the option to use Panorama or the Azure portal. The best method for ensuring up-to-date firewall configuration is to use Panorama for central management of firewall policies. Panorama simplifies consistent policy configuration across multiple independent firewalls through its device group and template stack capabilities. When multiple firewalls are part of the same device group, they receive a common ruleset. Because Panorama enables you to control all of your firewalls—whether they are on-premises or in the public cloud or whether they are a physical appliance

or virtual—device groups also provide configuration hierarchy. With device group hierarchy, lower-level groups include the policies of the higher-level groups. Configuration hierarchy allows you to configure consistent rulesets that apply to all firewalls, as well as consistent rulesets that apply to specific firewall deployment locations such as the public cloud.

You can deploy Panorama in your on-premises data center or in a public-cloud environment such as Azure. When deployed in your on-premises data center, Panorama can manage all the PA-Series, Cloud NGFWs and VM-Series next-generation firewalls in your organization. If you want a dedicated Panorama appliance for the Cloud NGFW and VM-Series firewalls in Azure, deploy Panorama on Azure.

When you have an existing Panorama deployment on-premises for firewalls in your data center and internet edge, you can use it to manage the Cloud NGFW and VM-Series firewalls in Azure. Beyond management, you need to consider your firewall log collection and retention. Log collection, storage, and analysis is an important cybersecurity best practice that organizations perform to correlate potential threats and prevent successful cyber breaches.

The following three deployment mode options are available for Panorama, which, if necessary, allows for the separation of management and log collection:

- **Log Collector mode**—One or more log collectors collect and manage logs from the managed devices. This assumes that another deployment of Panorama is operating in Management-Only mode.
- **Management-Only mode**—Panorama manages configurations for the managed devices but does not collect or manage logs.
- **Panorama mode**—Panorama controls both policy and log management functions for all the managed devices.


Panorama

You can use different methods for deploying Panorama. This section discusses the modes used in this guide.

On Premises Panorama and Dedicated Log Collection

Sending logging data back to the on-premises Panorama can be inefficient, costly, and pose data privacy and residency issues in some regions. An alternative to sending the logging data back to your on-premises Panorama is to deploy Panorama dedicated log collectors on Azure and use the on-premises Panorama for management. Deploying a dedicated log collector on Azure reduces the amount of logging data that leaves the cloud but still allows your on-premises Panorama to manage the Cloud NGFWs and VM-Series firewalls in Azure and have full visibility to the logs as needed.

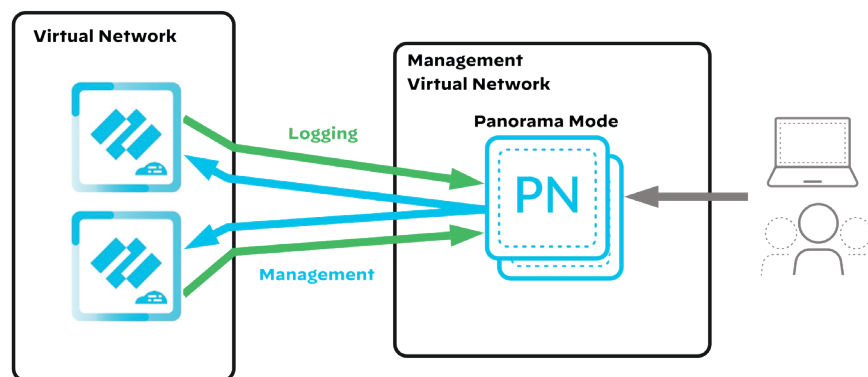
Figure 8 Panorama Log Collector mode on Azure



Panorama Management and Log Collection in the Cloud

In the second design option, you can use Panorama for both management and log collection. You can deploy the management and log collection functionality as a shared virtual appliance or on dedicated virtual appliances. For smaller deployments, you can deploy Panorama and the log collector as a single virtual appliance. For larger deployments a dedicated log collector per region allows traffic to stay within the region and reduce outbound data transfers.

Figure 9 Panorama management and log collection in Azure



Panorama is available as a virtual appliance for deployment on Azure and supports Management-Only mode, Panorama mode, and Log Collector mode with the system requirements defined in Table 2. Panorama on Azure is available with only a BYOL licensing model.

Table 2 Minimum system requirements for the Panorama virtual appliance

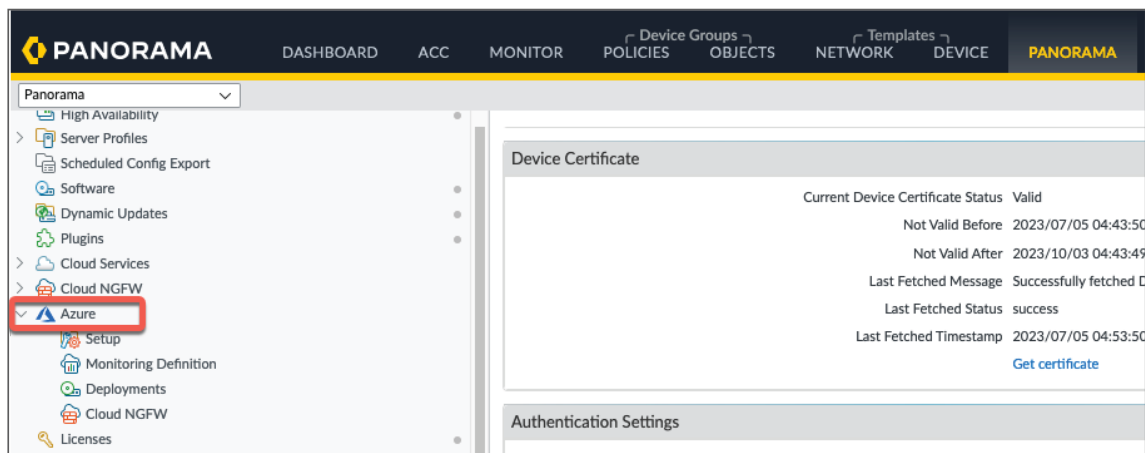
Management-Only mode	Panorama mode	Log Collector mode
16 CPUs 32 GB memory 81 GB system disk	16 CPUs 32 GB memory 2 TB to 24 TB log storage capacity	16 CPUs 32 GB memory 2 TB to 24 TB log storage capacity

Plugins

The Panorama extensible plugin architecture enables support for third-party integration plugins, such as Azure and AWS, along with other Palo Alto Networks products, such as the GlobalProtect® cloud service. With this modular architecture, you are able to take advantage of new capabilities without waiting for a new PAN-OS version.

The Azure plugin enables you to create cloud device groups and Cloud Template stacks, which help you manage policies and objects on Cloud NGFW resources that you manage with Panorama. Version 5.0 is the minimum plugin version supported. The Azure plugin also enables you to monitor your virtual machines on the Azure public cloud. With the plugin, you can enable communication between Panorama and your Azure subscriptions so that Panorama can collect a predefined set of attributes (or metadata elements) as tags for your Azure virtual machines and register the information to your Palo Alto Networks firewalls. When you reference these tags in DAGs and match against them in security policy rules, you can consistently enforce policy across all assets deployed within VNets in your subscriptions.

Figure 10 Azure plugin



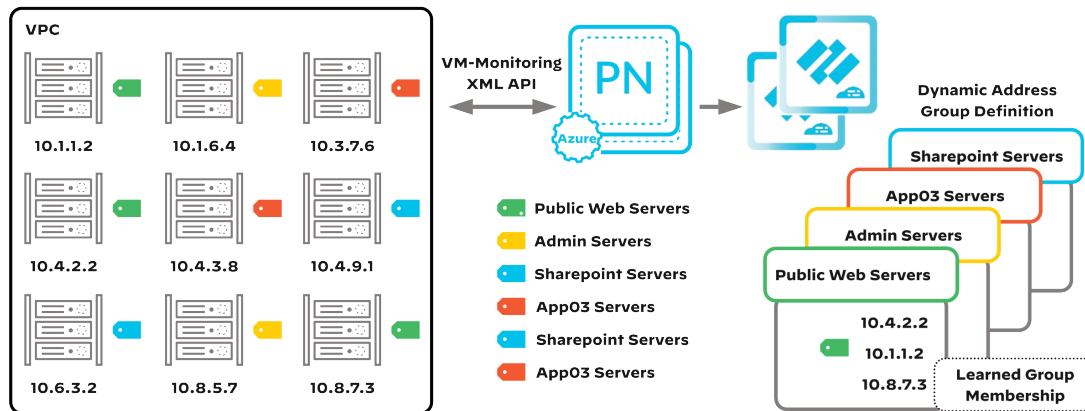
Resource Monitoring

Organizations typically build public-cloud application environments by using a continuous integration/continuous delivery (CI/CD) pipeline. Using CI/CD, you deploy applications and updates quickly and build new infrastructure to accommodate a revised application, as opposed to trying to upgrade the existing operational environment. When the new application or update goes online, you remove the now-unused, older application environment. This amount of change presents a challenge to enforcing security policy unless your security platform is compatible with an agile development and deployment process.

Cloud NGFW supports dynamic address groups. DAGs allow you to create security policy that automatically adapts to compute resource additions, moves, or deletions. DAGs also enable the operational flexibility for applying a security policy to a device based on its role.

A DAG uses tags as a filtering criterion in order to determine its members. You can define tags statically or register them dynamically. You can dynamically register the IP address and associated tags for Azure virtual machines by using the Azure plugin on Panorama.

Figure 11 VM monitoring of Azure-tag-to-DAG mappings



When using Panorama and the Azure plugin, you can centralize the retrieval of tags from Azure and security policy management in order to ensure consistent policies for hybrid and cloud-native architectures. Using a service principal with a built in or custom role that you create, the plugin polls your Azure subscriptions for resource tags and correlates the metadata (IP address-to-tag mapping) into DAGs. Panorama then relays the DAG content to the firewalls, providing scale and flexibility.

The Panorama plugin for Azure allows you to monitor all of your virtual machines, VNets, application gateways, and load-balancers in up to 500 Azure subscriptions. With the plugin, Panorama can retrieve a total of 32 tags for each virtual machine: 11 predefined tags and up to 21 custom tags. The number of tags used impacts the total number of IP addresses you can monitor. For example, Panorama can retrieve 7,000 IP addresses with 10 tags for each, or it can retrieve 6500 IP addresses with 15 tags for each.

Cloud NGFW via Azure Portal

You have numerous options when deciding on a design model for leveraging the Cloud NGFW for securing your applications in Azure. The first choice is to decide how you want to manage the firewall.

An important feature and function of the Cloud NGFW is that you have a choice on how you can configure, manage, and control the firewall. Because the Cloud NGFW is an ISV Service in Azure, fully managed on the backend, you can set up, configure, and manage Cloud NGFW directly from the Azure Portal, if you so choose. This is a common option for organizations seeking to add true NGFW capabilities to their Azure cloud who wish to have a truly cloud integrated and managed solution for securing their Azure traffic. Functioning as part of the Azure cloud service, a native Azure interface is available for those who wish to use Cloud native tooling.

Rulestacks

Rulestacks are a collection of individual rules and security rule objects that you can associate to one or more NGFWs. When building a Cloud NGFW in Azure, you use a local rulestack unless you use Panorama for management. When using the Azure portal or infrastructure-as-code to create an NGFW, you can associate local rulestacks with one or more NGFWs.

Rulestacks define access-control (App-ID, URL Filtering) and threat-prevention behavior of Cloud NGFW resources. To protect the traffic, Cloud NGFW uses your rulestack definitions in two steps. First, it enforces your rules and either allows or denies your traffic, depending on configuration. Second, it performs content inspection on the allowed traffic based on what you specify in the security profiles. A rulestack consists of a set of security rules, associated objects, and profiles.

Rules

Security rules protect networks and devices from threats and disruptions while also helping to optimally allocate network resources for enhancing productivity and efficiency in business processes. On Cloud NGFW for Azure, individual security rules determine whether to block or allow a session, based on configured traffic attributes, such as the source and destination IP address, source and destination FQDNs, or the application.

All traffic passing through the firewall is matched against a session. Each session is matched against a rule. When a session match occurs, the NGFW applies the matching rule to bidirectional traffic in that session (client-to-server and server-to-client). For traffic that doesn't match any defined rules, the default rules apply.

Security policy rules are evaluated left-to-right and from top-to-bottom. A packet is matched against the first rule that meets the defined criteria and, after a match is triggered, subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria.

Rules contain security policy configuration settings. Rules are applied to traffic based on numerous criteria, including the following:

- **Security rule objects**—*Security rule objects* are reusable rulestack settings that you can customize and use within individual rules. A security rule object is a single object, or it could be a collective unit that groups distinct identities such as IP addresses, FQDN, or certificates. As a general best practice, when creating a policy object, you group objects that require similar permissions in the policy. For example, if your organization uses a specific set of server IP addresses for authenticating users, you can group the set of server IP addresses as a prefix list object and reference that prefix list in one or more of your security rules. Grouping objects is beneficial for security operators and administrators because you can greatly reduce the administrative overhead in creating rules.

Security Rule objects include the following:

- Prefix and FQDN lists allow you to group specific source or destination IP addresses or FQDNs that require the same policy enforcement.
- Custom URL categories allow you to specify exceptions to a URL category's enforcement and to create custom URL categories based on existing categories.
- Intelligent feeds provide an ongoing stream of potential or current threat data. They record and track IP addresses and URLs that are associated with threats.
- **Security services**—Security services are the advanced security capabilities provided by Cloud NGFW for Azure. Network traffic that is allowed, based on the rule settings, to pass through an NGFW is also analyzed to ensure that it meets the requirements of the security services to prevent any disallowed traffic type that the initial rule did not block. You can set security services to the best-practices profiles or customize them to meet the needs of each rulestack. Security services include the following capabilities:
 - IPS examines traffic in order to detect and prevent vulnerability exploits.
 - Anti-spyware examines outbound traffic, especially command-and-control, to detect and prevent an infected client from communicating with a service outside of your network.
 - Malware and file-based threat protection includes antivirus and file blocking, which detect and block viruses, worms, trojans, spyware, and specific file types from being sent through your network.
 - Web-based threat protection includes URL categories and filtering to controls users' access to web content and how they interact with it.
 - Encrypted threat protection sets the TLS certificates that are used for outbound decryption.

- **Prefix and FQDN lists**—Prefix and FQDN lists allow you to group specific source or destination IP addresses or FQDNs that require the same policy enforcement. A prefix list can contain one or more IP addresses or IP netmasks in CIDR notation. An address object with the type IP netmask requires you to enter the IP address or network by using slash notation to indicate the IPv4 network, for example, 192.168.18.0/24. An FQDN (for example, shop.paloaltonetworks.com) object provides further ease of use because DNS provides the FQDN resolution to the IP addresses instead of you needing to know the IP addresses and manually updating them every time the FQDN resolves to new IP addresses.
- **Certificate**—A certificate object is a reference to a TLS certificate stored in the Azure Key Vault in your Azure account and is used in outbound decryption.

**Note**

To use Palo Alto Networks Advanced Cloud-Delivered Security Services (such as Advanced Threat Prevention, Advanced URL Filtering, WildFire®, and DNS Security), you must register your Azure tenant at the Palo Alto Networks Customer Support Portal after the firewall creation.

Without registering your Azure tenant, only the standard Cloud-Delivered Security Services (such as Threat Prevention and URL Filtering) are offered, if enabled.

DEPLOYMENT METHODS

In Azure, there are two deployment methods for using Cloud NGFW: Azure VNets and Azure Virtual WANs. An Azure VNet enables secure communication with other Azure resources, the internet, and on-premises networks. An Azure Virtual WAN represents a networking service that combines networking, security, and routing functionalities together to provide a singular, operational interface.

You can seamlessly deploy the Cloud NGFW in the Virtual WAN hub as a scalable firewall solution in order to secure traffic between critical workloads hosted in a global hybrid network between Azure and on-premises.

**Note**

One private IP address is used for an NGFW resource. For Virtual WAN environments, configure the Virtual WAN hub routing policy to hairpin traffic for the service. That is, the traffic exits an interface and returns before going out to the internet.

The Cloud NGFW for Azure Virtual WAN deployment:

- Is fully integrated into the Azure Virtual WAN by using the software-as-a-service (SaaS) framework.
- Is deployed right into the Virtual WAN virtual hub.
- Uses routing intent and policies to control which traffic gets inspected by the Cloud NGFW service.
- Enables enforcement of consistent security policy for the inter-hub and inter-region traffic

Regardless of whether you use Cloud NGFW to secure VNet traffic, Virtual WAN traffic, or both, deploying and using the service are straightforward. After the Cloud NGFW is created through the Azure Portal, you only need to define security policies in the Azure Portal or in Panorama, then route your traffic to the Cloud NGFW with the use of UDRs. You are free to focus on security. Because there is no infrastructure to manage, there is zero maintenance. Software upgrades, along with built-in scalability and resiliency, are included with the Cloud NGFW service.

Azure Portal

The Cloud NGFW is available through the Azure Marketplace. You can deploy it in just a few simple steps. From the Azure portal, provisioning and setup is simple and intuitive process. You can start your subscription and build your Cloud NGFWs through this process.

You use the Azure Portal to create and manage the Cloud NGFW inside your Virtual WAN hub, then route traffic to the NGFW. The Azure Virtual WAN UI functions in the same way as the Azure Portal or Resource Manager. The Cloud NGFW continuously meters service usage, then forwards it to the Azure metering service for billing purposes.

Infrastructure-As-Code

Infrastructure-as-code is a way of building and maintaining cloud resources and configurations. It ensures that policies and standards are met, maintains state, and reduces drift from intended configuration.

Cloud NGFW for Azure supports leading infrastructure-as-code tools. Azure Resource Manager is a native interface where you can deploy resources via code. Hashicorp Terraform is an independent infrastructure-as-code tool that works for all major public-cloud providers.

If you use infrastructure-as-code for deploying and managing your Cloud NGFW deployment, when you make future changes, you should use only your chosen tool. Although the Cloud NGFW console would still allow administrators to change settings, to maintain the state of the environment, the infrastructure-as-code tool would revert or overwrite the changes.

SECURING TRAFFIC WITH THE CLOUD NGFW

Cloud NGFW provides you with the tools and functionality for securing inbound traffic, outbound traffic, and east-west traffic.

Inbound traffic references any traffic that originates outside of Azure that is destined for resources in the cloud, whether it be virtual machines or applications inside of VNets, including items such as servers and load balancers. Cloud NGFW can prevent malware and vulnerabilities from entering your VNet in the inbound traffic that might be otherwise allowed by Azure Network Security Groups.

Outbound traffic describes traffic that originates from within Azure. This could be traffic coming from an application VNet that has a destination residing outside of Azure. Cloud NGFW protects outbound traffic flows by ensuring that resources in your application VNet connect to allowed services and allowed URLs while preventing exfiltration of sensitive data and information. The Cloud NGFW handles prevents data exfiltration and access of potentially malicious sites.

Another feature of Cloud NGFW is outbound decryption. With outbound decryption enabled, Cloud NGFW behaves like an SSL forward proxy and uses its associated certificates to establish itself as a trusted third party (man-in-the-middle) for the client-server session. However, Cloud NGFW keeps your traffic packet headers and payload intact, providing complete visibility of the source's identity to your destinations.

Outbound decryption uses two certificate objects: Trust and Untrust. The NGFW presents the trust certificate to clients during SSL decryption if the client is attempting to connect to a server that has a certificate signed by a trusted certificate authority (CA). Alternatively, the NGFW presents the untrust certificate to the client attempting to connect to a server that has a certificate signed by a CA that the NGFW does not trust.

You can configure the NGFW resource to decrypt the SSL traffic leaving your VNet or subnet. You can then enforce App-ID and security settings on the plaintext traffic, including Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles. To ensure privacy and security, after decrypting and inspecting traffic, the firewall re-encrypts the plaintext traffic as it exits the firewall. This defines the certificates that the firewall uses for outbound TLS decryption. You must enable outbound TLS decryption during rule creation.

East-west or lateral traffic moves within Azure. Typically, this is seen as VNet-to-VNet traffic or traffic passing between applications but can also include traffic coming from a Virtual WAN spoke that is attempting to reach an application or VNet inside of Azure. Specifically, traffic between source and destination deployed in two different application VNets or in two different subnets in the same VNet is east-west. Cloud NGFW can stop the propagation of malware within your Azure environment by providing inline inspection of all traffic traversing your network. This guide often refers to traffic as *north-south* or *east-west*. This is all traffic that the Cloud NGFW can secure.

Design Model

There are many ways to use the concepts discussed in the previous sections to build a secure architecture for your Azure environment. The design models in this section offer a complete example architecture that leverages the Palo Alto Networks Cloud NGFW to secure both north-south and east-west traffic flows in Azure.

In these design models, Panorama streamlines and consolidates core tasks and capabilities, enabling you to view all your firewall traffic, manage all aspects of device configuration, push global policies, and generate reports on traffic patterns or security incidents. You deploy Panorama in Panorama mode, leveraging the cloud service plugin so that all Cloud NGFW firewall logs are encrypted and sent directly from the firewalls to Panorama or to Azure Log Analytics workspace.

These design models differ slightly in how they create and secure tenant, application, or trust zone boundaries. The two models are not complimentary deployments. The model you choose depends entirely on whether you are using Azure Virtual WAN or running traditional Virtual Networks in Azure.

The design models discussed:

- **Virtual Network integration**—This model showcases protecting north-south and east-west traffic flows, leveraging Palo Alto Networks Cloud NGFW in a proven hub-and-spoke design. Using UDRs, this design steers all traffic so that it passes through the firewall for inspection.
- **Virtual WAN integration**—This model showcases protecting north-south traffic flows as well as east-west traffic flows, leveraging Palo Alto Networks Cloud NGFW's integration in Virtual WAN. By using routing intent and policies, the firewall inspects all traffic passing through the Virtual WAN.

VIRTUAL NETWORK

By using Cloud NGFW to secure VNet communications, you can effectively safeguard your cloud workloads while ensuring consistent and robust security policies, just like you would with an on-premises setup. With Cloud NGFW, you do not have the concern about managing the underlying infrastructure.

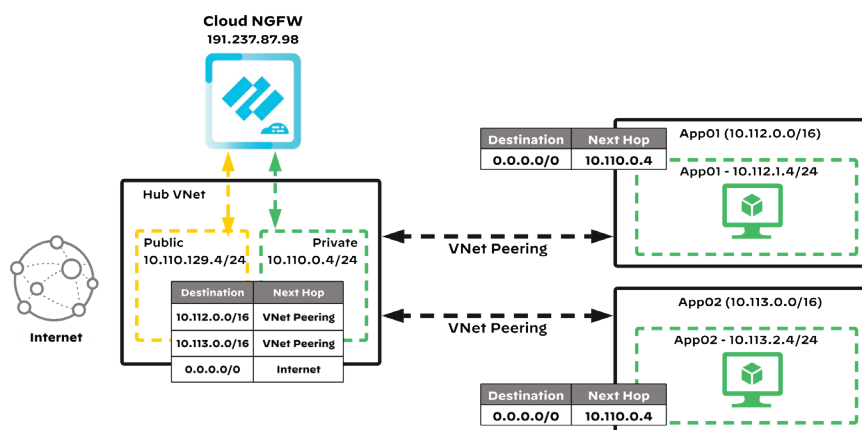
To achieve this, you employ a familiar and straightforward hub-and-spoke design. You create the transit (hub) VNet with a minimum requirement of a /25 network (255.255.255.128). Within the hub VNet, you configure two subnets with a minimum of /26 networks (255.255.255.192). One subnet serves as a public subnet, and you designate the other as a private subnet.

To facilitate inbound and outbound traffic, a public IP address is necessary. You can set it up beforehand or during the Cloud NGFW configuration process.

To connect application or spoke VNets to the hub VNet, you establish peering. Additionally, you configure UDRs and associate them with the spoke subnets, enabling the routing of traffic to the private IP address linked to the Cloud NGFW.

By following this setup, you can effectively leverage Cloud NGFW to protect your cloud workloads while maintaining a high level of security and streamlining your infrastructure management.

Figure 12 VNet design



Management and Configuration

You can use a Panorama appliance to manage security policies centrally on Cloud NGFW resources alongside your physical and virtual firewall appliances. You can also manage all aspects of shared objects and profiles configuration, push these rules, and generate reports on traffic patterns or security incidents of your Cloud NGFW resources—all from a single Panorama console. Organizations who choose the Panorama options are typically those who have hybrid and multi-cloud environments and want to leverage Panorama for a cloud-agnostic, consistent experience. Panorama provides a single location from which you can have centralized policy and firewall management across hardware firewalls, virtual firewalls, and cloud firewalls, which increases operational efficiency in managing and maintaining a hybrid network of firewalls.

Whether you are using Azure or Panorama, the process is the same when building the Cloud NGFW. You make this choice during setup. You continue to subscribe to the Cloud NGFW through Azure Marketplace. During the setup of the Cloud NGFW resource, you choose Panorama for the management option. You can then manage a shared set of security rules centrally on Cloud NGFW resources you create alongside your physical and virtual firewall appliances, and you can use logging, reporting, and log analytics, all from a single Panorama console. Your Panorama appliance can reside in any cloud region or in an on-premises environment. Panorama uses the Azure plugin to push policy and objects to the NGFW resources.

To integrate your Cloud NGFW resource with Panorama, you use the following Palo Alto Networks components.

Palo Alto Networks Policy Management Certificate

You use a Panorama appliance to author and manage policies for your Cloud NGFW resources. The policy-management component also helps to associate your authored policies and objects to multiple Cloud NGFW resources in different Azure regions.

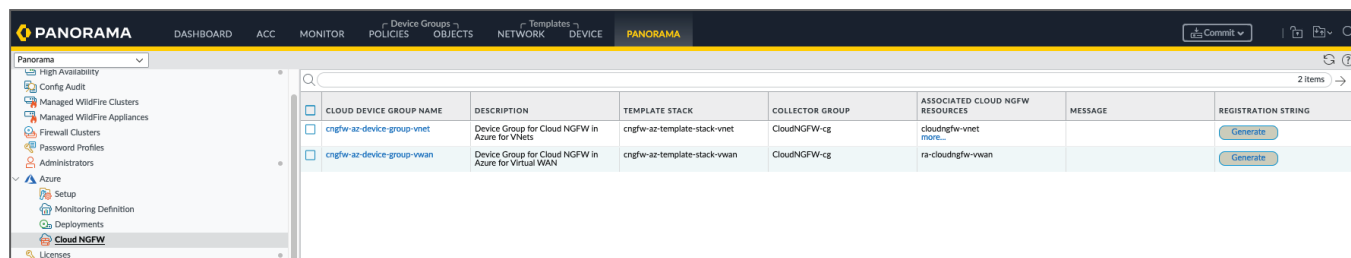
Panorama Azure Plugin Certificate

The Panorama Azure plugin enables you to create cloud device groups and Cloud Template stacks, which help you manage policies and objects on Cloud NGFW resources linked with Panorama.

Cloud Device Groups

Cloud device groups (Cloud DGs) are special-purpose Panorama device groups that allow you to author rules and objects for Cloud NGFW resources. Using the Panorama Azure plugin UI, you create Cloud DGs by specifying the Cloud NGFW and Azure region information. Cloud DG manifests as a global rulestack in that region.

Figure 13 Cloud device groups



	CLOUD DEVICE GROUP NAME	DESCRIPTION	TEMPLATE STACK	COLLECTOR GROUP	ASSOCIATED CLOUD NGFW RESOURCES	MESSAGE	REGISTRATION STRING
<input type="checkbox"/>	cnrgfw-az-device-group-vnet	Device Group for Cloud NGFW in Azure for VNets	cnrgfw-az-template-stack-vnet	CloudNGFW-cg	cloudngfw-vnet more...		Generate
<input type="checkbox"/>	cnrgfw-az-device-group-vwan	Device Group for Cloud NGFW in Azure for Virtual WAN	cnrgfw-az-template-stack-vwan	CloudNGFW-cg	ra-cloudngfw-vwan		Generate

You can create multiple Cloud DGs. You use the Panorama UI's device-group page when managing policy and object configurations in Cloud DGs and their associated objects and security profiles. It is also possible to leverage your existing shared objects and profiles from existing Panorama device groups by referring to them in the security rules you create in your Cloud DGs. Alternatively, if you want to inherit the device group rules and objects, you can add these Cloud DGs to the device-group hierarchy that you manage in Panorama. You can associate the same Cloud DG with multiple regions of the Cloud NGFW resource. This Cloud DG manifests as a dedicated global rulestack in each Azure region of your Cloud NGFW resource.

Cloud Template Stacks

Cloud Template stacks are special-purpose Panorama template stacks that allow your security rules in Cloud DGs to refer to object settings that Panorama allows you to manage by using templates. When creating a Cloud DG, the Panorama Azure plugin enables you to create or specify a Cloud Template stack. The plugin automatically creates this Cloud Template stack and adds it to the cloud device as a reference template stack. From now on, you can use the native Panorama UI's Template Stack page to configure your templates and add them to these Cloud Template stacks.

The Cloud NGFW service manages most device and network configurations in your Cloud NGFW resources. Cloud NGFW therefore ignores infrastructure settings such as interfaces, zones, and routing protocols, if you have configured them in templates added to the Cloud Template stack.

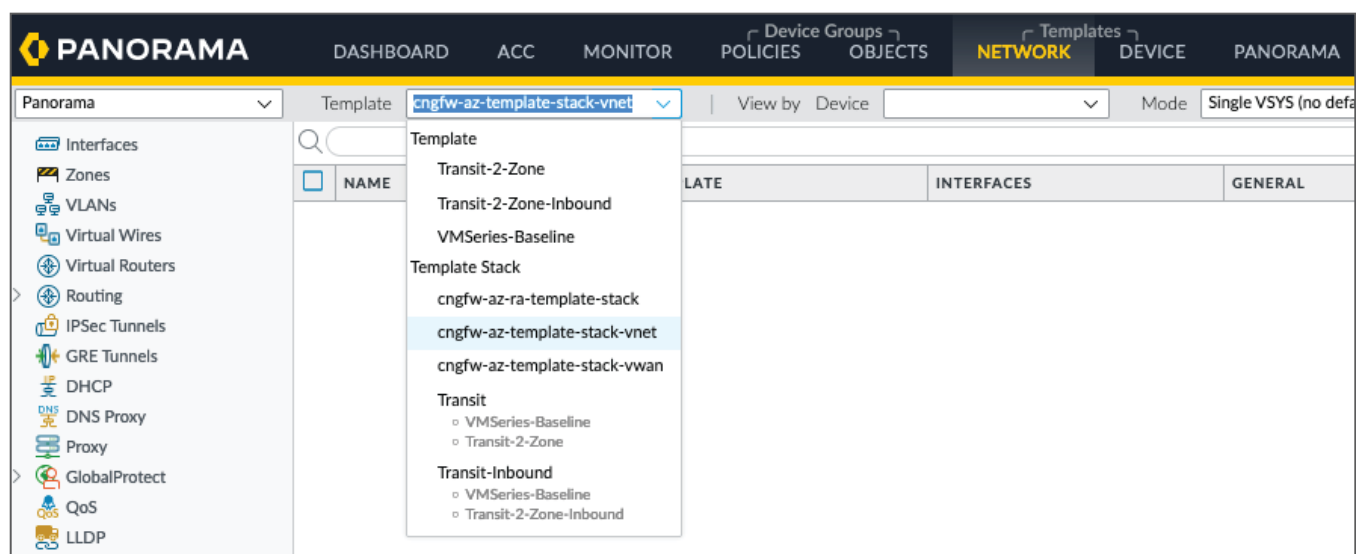
Cloud NGFW currently honors certificate management and log settings in your templates as referenced by the Cloud device group configuration. It ignores all other settings.



Note

You do not assign managed devices to cloud device groups and Cloud Template stacks.

Figure 14 Template stacks



Traffic Flows

The goal behind using any firewall or firewall service is the ability for the firewall (or in the case of this guide, the Cloud NGFW) to inspect any and all traffic you define for securing—as a best practice, all traffic going into Azure, all traffic leaving Azure, and all traffic moving laterally within Azure. This design model assumes no trust of any devices and leverages the Cloud NGFW for the inspection and securing of all Azure traffic.

Inbound

Inbound traffic is any traffic originating outside your Azure region and bound for resources inside your application VNets, such as servers or load balancers. Cloud NGFW can prevent malware and vulnerabilities from entering your VNet in the inbound traffic allowed by Azure security groups.

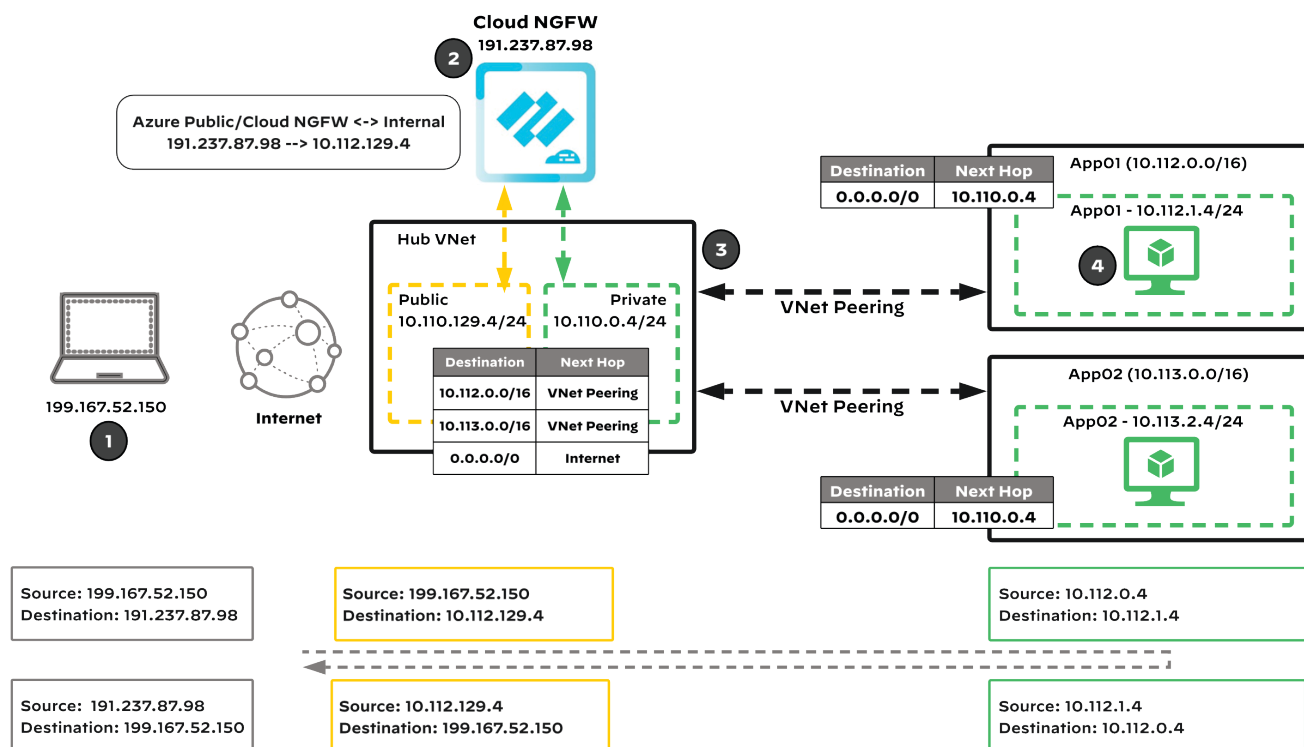
In this design, the Cloud NGFW deployed in the VNet provides inbound protection. You configure ingress traffic from the internet into Azure to pass through the Cloud NGFW for inspection. Inspecting inbound traffic from the internet at the firewall allows for each packet to pass through the deep-inspection engine. An inbound destination could be an application or server residing in a VNet.

Inbound traffic is initiated to the public IP address associated to the Cloud NGFW. This traffic could be destined to a web application but most often is a load balancer, application gateway, or other frontend device to the application or server. The public IP address is associated with the application through a NAT rule. This association ensures that this traffic is transparently forwarded to the Cloud NGFW for inspection. The packet arrives at the Cloud NGFW unchanged. The firewall performs the required translation. The source IP is replaced with the firewall's private IP address, and the destination IP address is replaced with the destination workload's private IP address. The packet is then passed through the Azure network to the destination workload. You can set up multiple translations for use with different applications.

Inbound traffic flow characteristics:

- A device from the outside is accessing the application on its public IP address.
- Public IP address(es) is associated with the public IP of the Cloud NGFW.
- Traffic is transparently forwarded based on the defined NAT rules on the firewall. The packet arrives at the firewall unchanged.
- The firewall performs SNAT and DNAT. The source IP is replaced with the firewall's private interface IP address, and the destination IP address is replaced with the destination workload private IP address.
- The firewall sends the packet out of the private interface.
- The packet is forwarded via the internal subnet to the destination host.

Figure 15 VNet inbound



Outbound

Each Cloud NGFW has an interface in the hub VNet public, private, and management subnet. To reach any private or internal subnets, you configure static routes for all peered application and enterprise network subnets.

User-defined routes in the private subnets of the peered application VNets direct traffic to the Cloud NGFW private interface. This forces traffic destined to the internet to pass through the firewall for inspection.

The Cloud NGFW applies source NAT to outbound internet traffic. The Cloud NGFW translates the original source address to the IP address of its public interface. Azure then automatically translates the interface IP address to the public IP address associated with the firewall's public interface when the traffic egresses to outbound destination.

Outbound traffic flow characteristics:

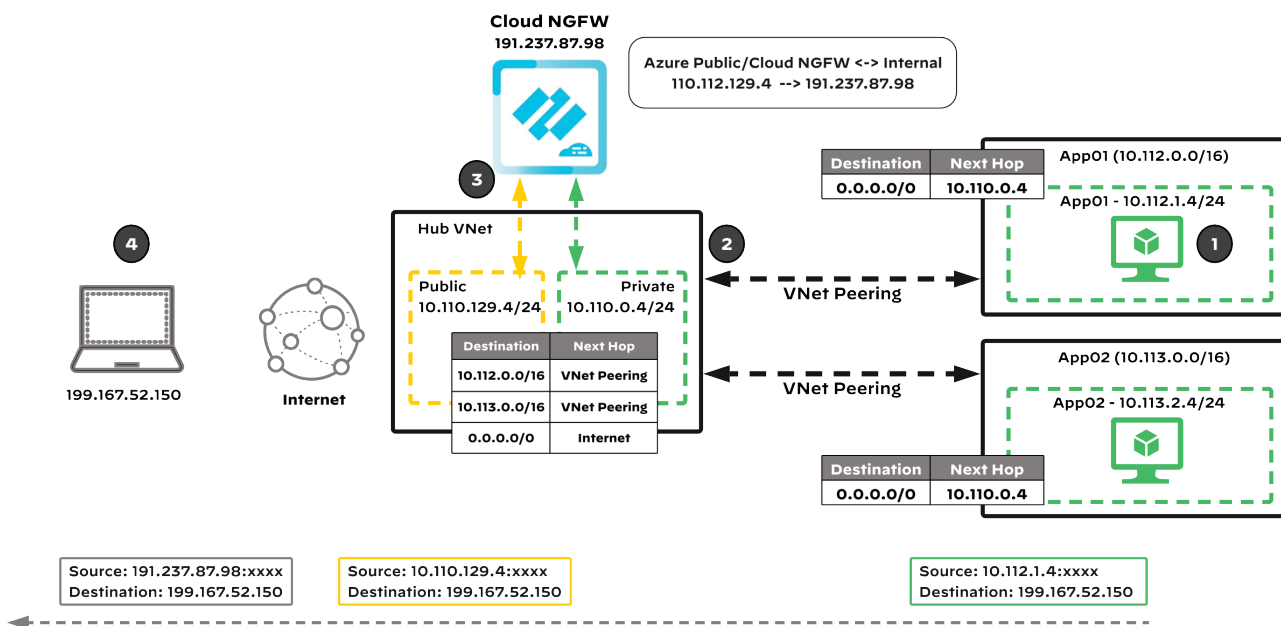
- The firewall security policy allows appropriate application traffic from the resources in the application virtual networks to the internet. Security profiles prevent known malware and vulnerabilities from entering the network in return traffic allowed by the security policy. URL filtering, data loss prevention, file blocking, and data filtering protect against data exfiltration.
- Traffic from source instance is sent to the Cloud NGFW via default route (0.0.0.0/0) defined in the subnet UDR.
- The private IP address is associated with the internal interface of the Cloud NGFW subscription.
- Traffic is forwarded to the private interface of the Cloud NGFW instance.
- The firewall applies a source NAT translation for the traffic using the public interface.



Note

You should implement the security policy by using positive security policies (*allow listing*).

Figure 16 VNet outbound



East-West

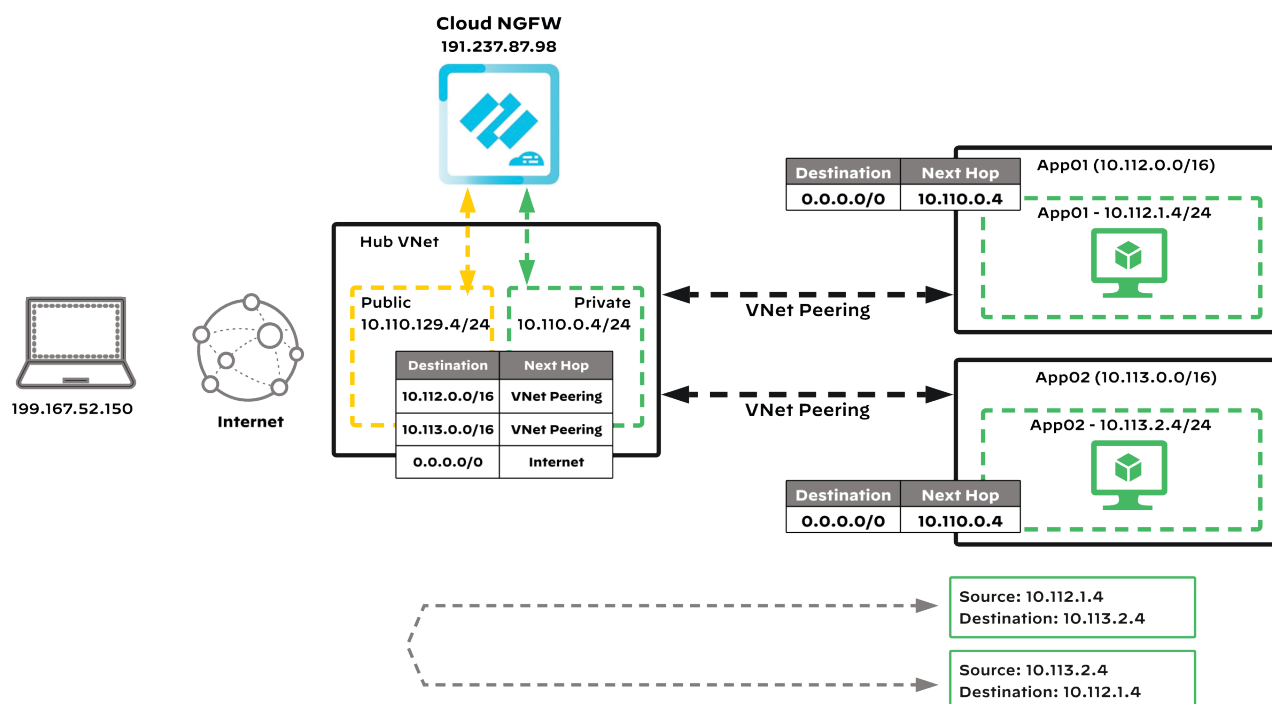
You apply user-defined routes for the private network subnets to the private subnets and direct traffic to the transit VNet's Cloud NGFW private IP address.

For traffic between private application subnets, the firewall does not translate either the source or destination. A positive control security policy should allow only appropriate application traffic between private resources and requires that you create corresponding security policy rules to permit specific traffic. You must then override the default intrazone security policy rule and modify it to deny traffic. You should also enable security profiles in order to prevent known malware and vulnerabilities from moving laterally in the private network through traffic allowed by the security policy.

East-to-west traffic flow characteristics:

- The packet sent from the workload is sent to the private IP by using UDR.
- The private IP address is associated with the Cloud NGFW service.
- The traffic arrives on the private interface of the firewall.
- Upon inspection, the traffic leaves via the same interface to the destination in another VNet.


Figure 17 VNet east-west



VIRTUAL WAN

Azure Virtual WAN provides you with the potential to reduce peering, connect on-premises and branch locations to the cloud, and use the Azure backbone network for communications between sites and VNets. With Cloud NGFW, it is now possible to include best-in-class network security for all traffic flows contained in your Virtual WAN. Setting up Cloud NGFW for use with Virtual WAN is simple and straightforward. You deploy Cloud NGFW directly into your Virtual WAN hub. If you use multiple hubs across different regions, you deploy Cloud NGFW into each. Once set up in Virtual WAN, Cloud NGFW uses routing intent and policies to control which traffic gets inspected. Security best practices and this design call for the inspection of all traffic, regardless of the direction it is traveling. This design assumes no trust (Zero Trust) across network boundaries. This design ensures that you have the capability to enable enforcement with consistent security policies for both inter-hub and inter-region traffic.

Figure 18 Virtual WAN design



Routing Intent Policies

Virtual WAN hub routing intent allows you to set up simple and declarative routing policies to send traffic to a virtual appliance or to a SaaS solution such as Palo Alto Networks Cloud NGFW deployed within a Virtual WAN hub.

When looking at routing intent and routing intent policies, there are two types to choose from: internet traffic and private traffic. Each Virtual WAN hub can have at most one internet-traffic routing policy and one private-traffic routing policy, each with a single next-hop resource. In the case of Azure, the definition of private traffic includes both branch and VNet address prefixes. However, routing policies consider them as one entity within the routing intent concepts.

Figure 19 Virtual WAN routing intent policies

VirtualWAN-hub-East | Routing Intent and Routing Policies

Virtual HUB

Search < Save Cancel Delete

Overview

Configure routing policies for VirtualWAN-hub-East Virtual Hub

Routing Policies for Internet Traffic apply to all connections connected to the Virtual Hub

Routing Policies for Private Traffic apply to all private traffic destined for addresses in the Private Traffic Prefixes below (regardless of the source) that enters the virtual hub

Internet traffic

SaaS solution

Next Hop Resource

RA-CloudNGFW-vWAN

Private traffic

SaaS solution

Next Hop Resource

RA-CloudNGFW-vWAN

Private Traffic: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16,

Connectivity

- VPN (Site to site)
- ExpressRoute
- User VPN (Point to site)

Routing

- Routing Intent and Routing Policies
- BGP Peers
- Route Tables
- Effective Routes

Security

- Azure Firewall and Firewall Manager

Third party providers

- Network Virtual Appliance
- SaaS Solutions

When you configure an internet-traffic routing policy on a Virtual WAN hub, all branch and VNet connections to that virtual WAN hub forward internet-bound traffic to the destination defined in the routing policy: the Cloud NGFW. An internet-traffic routing policy is configured on a Virtual WAN hub, and the Virtual WAN advertises a default (0.0.0.0/0) route to all spokes, gateways, and network virtual appliances (deployed in the hub or spoke).

When you configure a private-traffic routing policy on a Virtual WAN hub, all branch and VNet traffic travelling in and out of the Virtual WAN hub (including any inter-hub traffic) is forwarded to the destination defined in the routing policy: again, the Cloud NGFW. When a private-traffic routing policy is configured on the Virtual WAN hub, all branch-to-branch, branch-to-VNet, VNet-to-branch, and inter-hub traffic is sent through the Cloud NGFW deployed in the Virtual WAN hub.



Note

Routing intent is currently generally available in Azure public cloud. Azure for Government currently has this on the roadmap. For information about current support in your region, see the [Azure documentation](#).

Routing intent simplifies routing by managing route-table associations and propagations for all connections (VNet, site-to-site VPN, point-to-site VPN, and ExpressRoute). You therefore cannot use the routing intent constructs with Virtual WANs that use custom route tables and customized policies.

Management and Configuration

You can use a Panorama appliance to manage security policies centrally on Cloud NGFW resources alongside your physical and virtual firewall appliances. You can also manage all aspects of shared objects and profiles configuration, push these rules, and generate reports on traffic patterns or security incidents of your Cloud NGFW resources—all from a single Panorama console. Organizations who choose the Panorama options are typically those who have hybrid and multi-cloud environments and want to leverage Panorama for a cloud-agnostic, consistent experience. Panorama provides a single location from which you can have centralized policy and firewall management across hardware firewalls, virtual firewalls, and cloud firewalls, which increases operational efficiency in managing and maintaining a hybrid network of firewalls.

Whether you are using Azure or Panorama, the process is the same when building the Cloud NGFW. You make this choice is made during setup. You continue to subscribe to the Cloud NGFW through Azure Marketplace. During the setup of the Cloud NGFW resource, you choose Panorama for the management option. You can then manage a shared set of security rules centrally on Cloud NGFW resources you create alongside your physical and virtual firewall appliances, and you can use logging, reporting, and log analytics, all from a single Panorama console. Your Panorama appliance can reside in any cloud region or in an on-premises environment. Panorama uses the Azure plugin to push policy and objects to the NGFW resources.

To integrate your Cloud NGFW resource with Panorama, you use the following Palo Alto Networks components.

Palo Alto Networks Policy Management Certificate

You use a Panorama appliance to author and manage policies for your Cloud NGFW resources. The policy-management component also helps to associate your authored policies and objects to multiple Cloud NGFW resources in different Azure regions.

Panorama Azure Plugin Certificate

The Panorama Azure plugin enables you to create cloud device groups and Cloud Template stacks, which help you manage policies and objects on Cloud NGFW resources linked with Panorama.

Cloud Device Groups

Cloud device groups (Cloud DGs) are special-purpose Panorama device groups that allow you to author rules and objects for Cloud NGFW resources. Using the Panorama Azure plugin UI, you create Cloud DGs by specifying the Cloud NGFW and Azure region information. Cloud DG manifests as a global rulestack in that region.

Figure 20 Cloud device groups

CLOUD DEVICE GROUP NAME	DESCRIPTION	TEMPLATE STACK	COLLECTOR GROUP	ASSOCIATED CLOUD NGFW RESOURCES	MESSAGE	REGISTRATION STRING
<input type="checkbox"/> cngfw-az-device-group-vnet	Device Group for Cloud NGFW in Azure for VNets	cngfw-az-template-stack-vnet	CloudNGFW-cg	cloudngfw-vnet more...		Generate
<input type="checkbox"/> cngfw-az-device-group-vwan	Device Group for Cloud NGFW in Azure for Virtual WAN	cngfw-az-template-stack-vwan	CloudNGFW-cg	ra-cloudngfw-vwan		Generate

You can create multiple Cloud DGs. You use the Panorama UI's device-group page when managing policy and object configurations in Cloud DGs and their associated objects and security profiles. It is also possible to leverage your existing shared objects and profiles from existing Panorama device groups by referring to them in the security rules you create in your Cloud DGs. Alternatively, if you want to inherit the device group rules and objects, you can add these Cloud DGs to the device-group hierarchy that you manage in Panorama. You can associate the same Cloud DG with multiple regions of the Cloud NGFW resource. This Cloud DG manifests as a dedicated global rulestack in each Azure region of your Cloud NGFW resource.

Cloud Template Stacks

Cloud Template stacks are special-purpose Panorama template stacks that allow your security rules in Cloud DGs to refer to object settings that Panorama allows you to manage by using templates. When creating a Cloud DG, the Panorama Azure plugin enables you to create or specify a Cloud Template stack. The plugin automatically creates this Cloud Template stack and adds it to the cloud device as a reference template stack. From now on, you can use the native Panorama UI's Template Stack page to configure your templates and add them to these Cloud Template stacks.

The Cloud NGFW service manages most device and network configurations in your Cloud NGFW resources. Cloud NGFW therefore ignores infrastructure settings such as interfaces, zones, and routing protocols, if you have configured them in templates added to the Cloud Template stack.

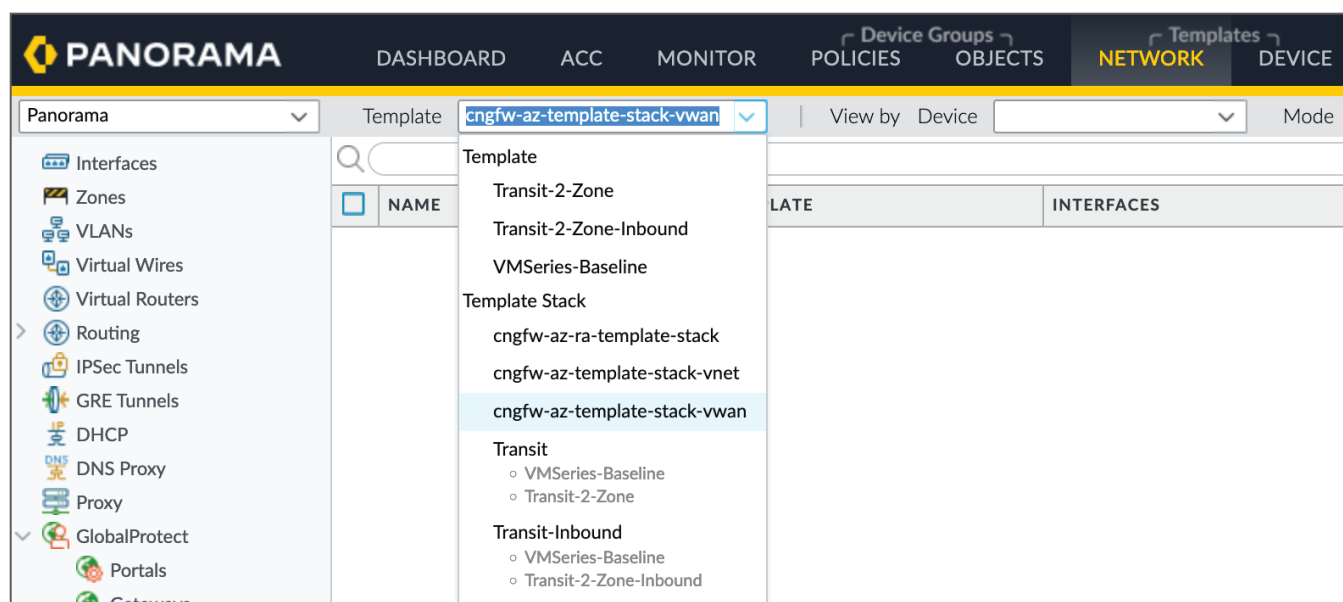
Cloud NGFW currently honors certificate management and log settings in your templates as referenced by the Cloud device group configuration. It ignores all other settings.



Note

You do not assign managed devices to cloud device groups and Cloud Template stacks.

Figure 21 Template Stacks



Traffic Flows

The goal behind using any firewall or firewall service is the ability for the firewall (or in the case of this guide, the Cloud NGFW) to inspect any and all traffic you define for securing—as a best practice, all traffic going into Azure, all traffic leaving Azure, and all traffic moving laterally within Azure. This design model assumes no trust of any devices and leverages the Cloud NGFW for the inspection and securing of all Azure traffic.

Inbound

Inbound traffic is any traffic originating outside your Azure region and bound for resources inside your application VNETs, such as servers or load balancers. Cloud NGFW can prevent malware and vulnerabilities from entering your VNet in the inbound traffic allowed by Azure security groups.

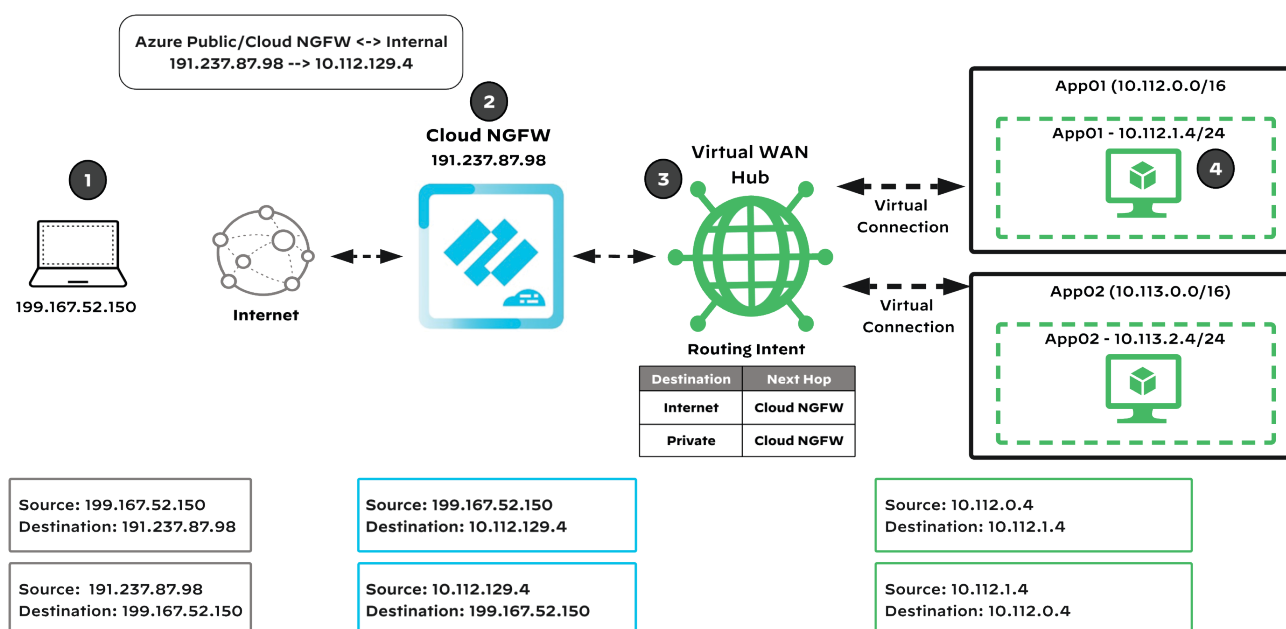
When you deploy Cloud NGFW to secure Virtual WAN, you are able to secure each Virtual WAN hub and all connections into and out of the hub.

Inbound protection is provided through the Cloud NGFW deployed in the VNet. You should configure ingress traffic from the internet into Azure to pass through the Cloud NGFW for inspection. Inspecting inbound traffic from the internet at the firewall allows for each packet to pass through the deep inspection engine. Once processed through the firewall, traffic egresses the trust side interface of the Cloud NGFW before following the Azure routing, which sends the request in this case to the Virtual WAN hub. This inbound destination could be an application or server residing in a VNet, or it could be destined towards a resource residing in a Virtual WAN remote site, as covered in this guide.

Inbound traffic flow characteristics:

- A client from the outside is accessing the workload on its public IP address.
- Public IP address(es) is associated with the Cloud NGFW.
- Traffic is transparently forwarded based on the defined Frontend rules to the firewall. The packet arrives at the firewall unchanged.
- The firewall performs a NAT translation. The source IP is replaced with the firewall's private interface IP address, and the destination IP address is replaced with the destination workload private IP address.
- The firewall sends the packet out of the firewall's private interface.
- The packet is forwarded via the private interface in the internal subnet to the destination host.

Figure 22 Virtual WAN inbound



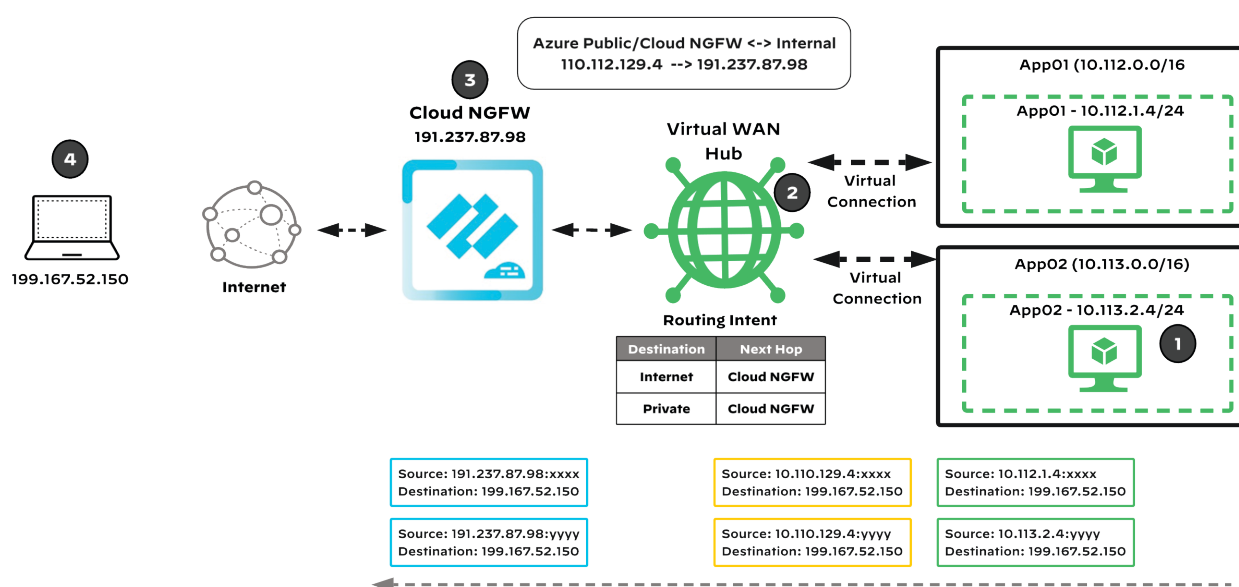
Outbound

Cloud NGFW secures traffic that is leaving from Azure and destined to the internet. In most cases, Virtual WAN remote sites access the internet directly and not need to go through Azure. However, there could be cases where these Virtual WAN remote sites are sending and exchanging traffic through Azure to the internet. In this situation, the Cloud NGFW associated with the Virtual WAN hub inspects and secures these flows.

Outbound traffic flow characteristics:

- Traffic from the source application is sent to the virtual hub.
- Traffic is then sent to the Private IP address associated with the Cloud NGFW interface on the private subnet.
- The firewall applies a source NAT to the traffic using the public interface.
- Next, another source NAT is applied using the public IP addresses Cloud NGFW, before being sent to the destination IP address.

Figure 23 Virtual WAN outbound



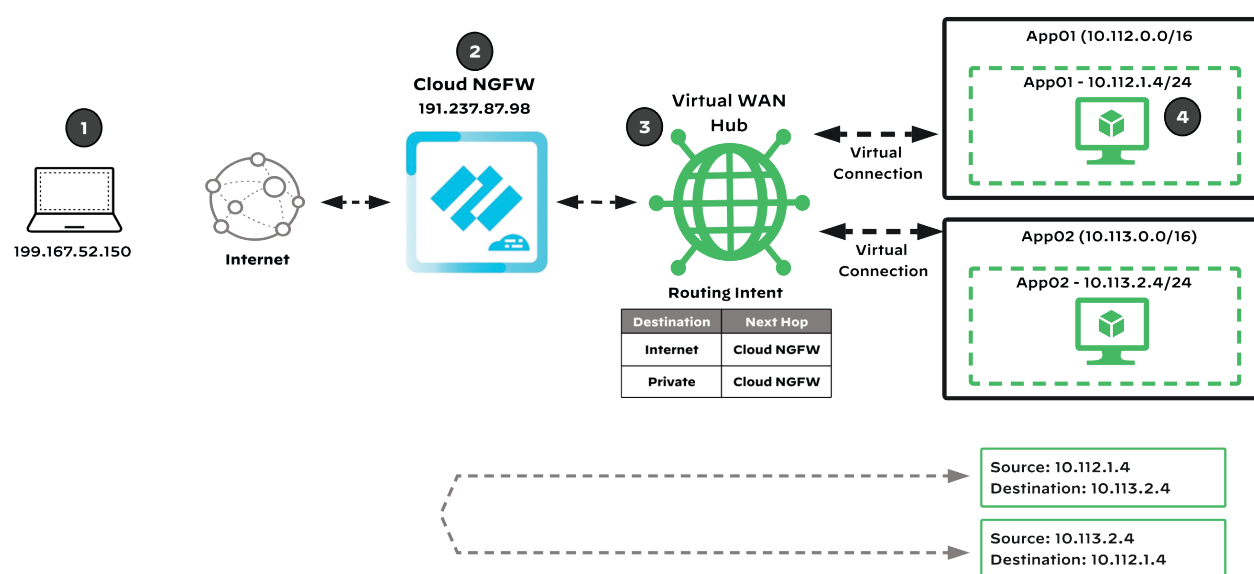
East-West

Although this design allows for the inline inspection and protection of multiple directional flows or traffic, the east-west flows are among the most common requirements for providing inline security between Virtual WAN remote sites and resources contained in Azure. Typically, remote sites are connecting to Virtual WAN in order to access resources contained in Azure. These resources could be an application, virtual machine, or other Azure service. Accessing this data from outside of Azure means it is important for this traffic to be inspected, logged, and secured. In this guide, *east-west traffic* refers to any VNet-to-VNet traffic, VNet-to-remote-sites and remote-sites-to-VNet traffic, and remote-site-to-remote-site traffic through Virtual WAN.

East-to-west traffic flow characteristics:

- The Virtual WAN hub advertises a route for RFC 1918 addresses space via the Cloud NGFW with the private IP as the next hop.
- The packet arrives on the private interface of the Cloud NGFW VNet.
- The firewall inspects the traffic and applies any relevant security policy rules.
- The traffic leaves via private interface. The Virtual WAN hub routes the traffic to the destination VNet.

Figure 24 Virtual WAN east-west



Summary

Moving applications to the cloud requires the same enterprise-class security as your private network. The shared-security model in cloud deployments places the responsibility of protecting applications and data on your organization. Deploying Palo Alto Networks Cloud NGFW for Azure in your Azure infrastructure provides a scalable infrastructure with protections from known and unknown threats, complete application visibility, a common security policy, and native cloud automation support. Your ability to move applications to the cloud securely helps you to meet challenging business requirements.

Securing network traffic in Azure is a key component and requirement for protecting your business and technical resources. With multiple configurations available, you can choose the design model that best suits your needs. Whether you prefer managing your Cloud NGFW through the Azure portal or using Panorama, our best practice design ensures that your inbound, outbound, and east-west traffic is secure. Cloud NGFW provides inline security for all Azure network traffic, ensuring that your data is protected at all times. You use centralized policy and rule configuration to easily configure and manage your security policies and rules from a single location, whether it's through the Azure portal or Panorama. You have the flexibility to choose the management option that works best for you. If you prefer a truly cloud-integrated and managed solution, you can configure and manage Cloud NGFW directly from the Azure portal. Alternatively, if you have hybrid and multi-cloud environments, are an existing Palo Alto Networks customer, or seek additional controls and features for your environment, Panorama offers a cloud-agnostic, consistent experience. Panorama allows you to have centralized policy and firewall management across hardware firewalls, virtual firewalls, and cloud firewalls, increasing operational efficiency in managing and maintaining your network of firewalls.

The designs in this guide provide you the details needed to design and secure your VNet communications by leveraging Cloud NGFW. With a hub-and-spoke design, you protect your cloud workloads while maintaining strong and consistent security policies. If you are using (or plan to use) Azure Virtual WAN to connect your on-premises and branch locations to the cloud along with your VNets residing in Azure, you can use Cloud NGFW directly on your Virtual WAN hub to ensure best-in-class network security for all traffic flows, controlling your traffic flows with routing intent and policies to ensure traffic gets inspected.

By deploying Palo Alto Networks Cloud NGFW in your Azure environment, you can ensure the comprehensive inspection and securing of all of your Azure traffic. Whether it's inbound, outbound, or east-west traffic, Cloud NGFW provides the necessary protection to safeguard your resources.

HEADQUARTERS

Palo Alto Networks	Phone: +1 (408) 753-4000
3000 Tannery Way	Sales: +1 (866) 320-4788
Santa Clara, CA 95054, USA	Fax: +1 (408) 753-4001
https://www.paloaltonetworks.com	info@paloaltonetworks.com

© 2023 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies. Palo Alto Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.



You can use the [feedback form](#) to send comments about this guide.