Implementation Guide - NextGen Firewall in Azure

https://campus.barracuda.com/doc/53248931/

This guide contains advanced topics and concepts. Follow the links in each section for step-by-step instructions on how to configure the following features.

Microsoft Azure is one of the big players in the public cloud, offering customers both IaaS and PaaS services. In this guide, we will cover how the Barracuda NextGen Firewall plays a key role in an Azure multi-tiered network and how its connectivity and security features let you access and connect to your cloud resources transparently and securely. We will also look at how the NextGen Control Center appliance helps you manage a large number of firewalls by managing and sharing common configurations and offering the admin a single user interface.

Multi-Tier Network Design in Azure

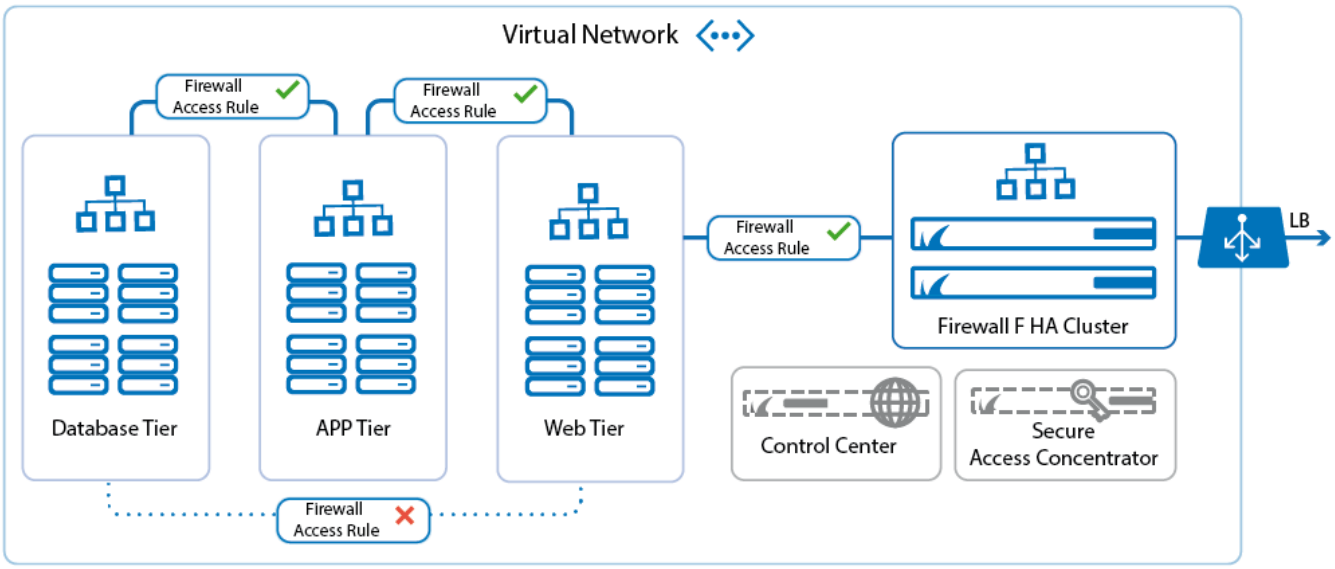

It is a common practice to deploy your applications in multiple tiers. All incoming, outgoing, and internal traffic passes through the firewall, offering you full visibility of the network traffic in your virtual network. Using a firewall allows you to create and enforce policies that separate the application tiers from one another. All network activity is logged and can be streamed to the Control Center or another log collector. The firewall also takes care of connecting your on-premises networks to your virtual networks in Azure. Clients can connect using either client-to-site VPN or SSL VPN. If Secure Connector appliances are also connecting to the application you are running in the cloud, a dedicated Secure Access Concentrator Firewall can terminate and handle the SC traffic.

Organizing Resources in Azure

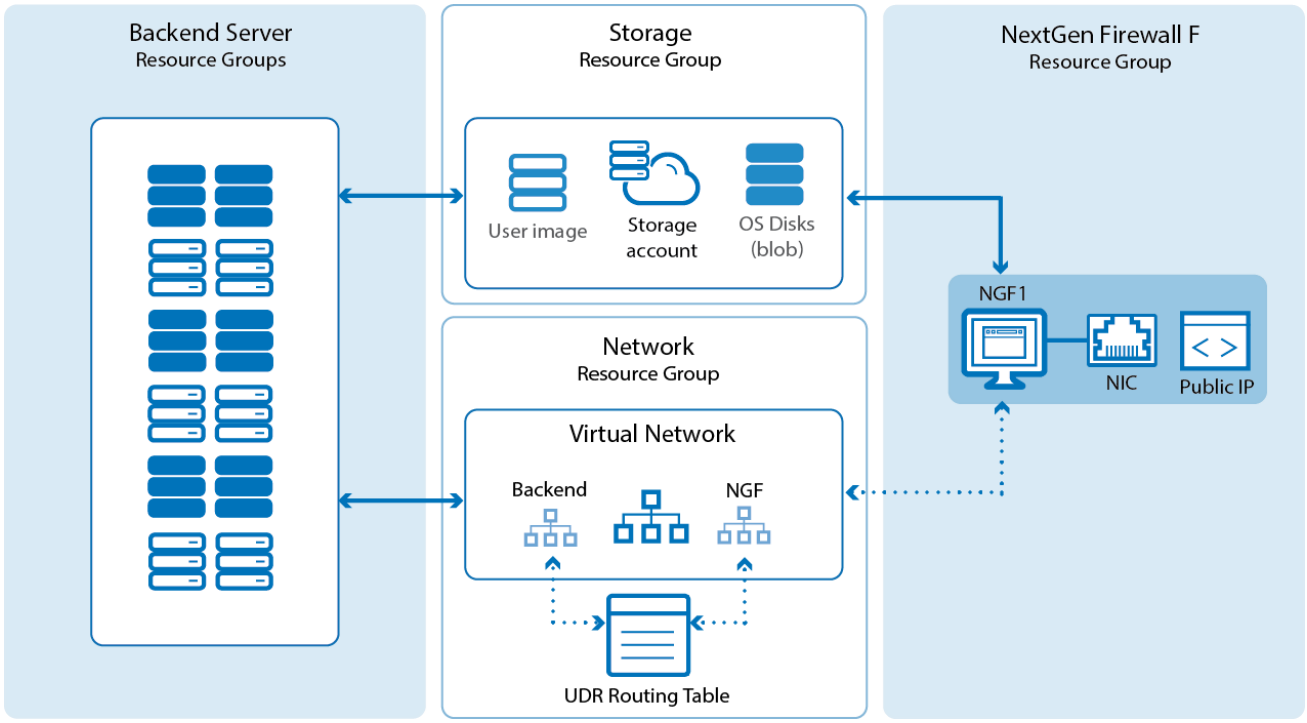

In the Azure Resource Manager deployment model, each resource is placed in a resource group. These logic containers allow you to group your resources to simplify monitoring and to access permissions and billing for your applications. Resource groups also make it easier to redeploy individual parts. For example, if you need to update the application servers, you only redeploy the resource group for these VMs, leaving the rest untouched. A common way to group resources would be:

- **Firewall Resource Group** – This resource group includes the firewall VMs and their network interfaces.
- **Storage Resource Group** – This resource group contains the storage accounts for the OS disks, source images, and all other data your application requires. To satisfy performance requirements, you may have to use multiple storage accounts.
- **Networking Resource Group** – This group contains all Azure networking resources: virtual network, load balancer, public IP address, and Azure route tables.
- **Control Center Resource Group** – This resource group contains the NextGen Control Center VM.
- **Multiple Backend Resource Groups** – These resource groups include VMs actually hosting the application. E.g., Database server resource group, application server resource group, and so on.

Planning your Deployment

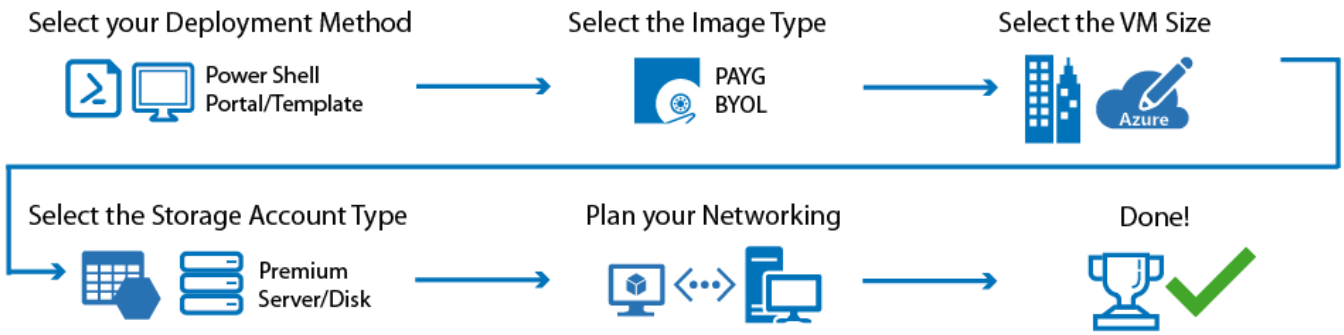

Before deploying the firewall VM, you need to consider several factors because some changes require you to redeploy your setup. The VM pricing tier, storage account type, and network design must all be aligned with each other in order to reach the required level of performance and reliability.

- **Select deployment method**
- **Select image**
- **Select VM pricing tier**
- **Select storage account type**
- **Plan networking**

Select the Deployment Method

Microsoft offers several deployment options: web portal, Azure PowerShell, and Azure templates. Most advanced features require the use of either PowerShell or templates. Resources that must be redeployed often are best deployed via templates because they offer the advantage of deploying as a group. In addition, changes to templates will selectively update only the new parts without requiring a redeployment. On the other hand, Azure PowerShell offers a good way for rapid prototyping while still using the full feature set that Azure provides. The web portal does not offer all the options or the flexibility, but it is easy to use without requiring Azure PowerShell or any knowledge of JSON template syntax.

For more information, see Microsoft Azure Deployment.

Select the Image Type

The Barracuda NextGen Firewall F-Series and Control Center are offered through the Azure Marketplace and as VHD disk images in the Barracuda Download portal. Using the Marketplace images will always give you the latest available firmware release with the option to choose between hourly and BYOL images. However, if you are managing your firewalls with a Control Center and using BYOL images, using the latest firmware may not be an option because the Control Center cannot manage firewalls that have a newer firmware version than it itself is running. In such cases, you need to upload the VHD disk image from the Barracuda Download portal to your storage account and change your deployment to use the source image. With user images, you also do not have to worry about potential migration issues if you use a PAR file created on an older firmware version.

For more information, see How to Upload Azure VHD Images for User Defined Images using ARM and Public Cloud Licensing.

Select the VM Size

You must select the VM size according to the number of CPU cores included in your NextGen Firewall F license. From the VM sizes with the CPU cores, select according to the following criteria:

- **RAM size** – Depending on the features used on the firewall,
- **Max IOPS / Premium disk support** – To increase disk I/O performance, you can either use a premium storage account or use multiple data disks. Data disks can be added from different storage accounts, and if they are present at deployment, a RAID device is automatically created and mounted as /phion0/. Data disks can also be added after the deployment via command-line script on the firewall. Premium storage accounts are only supported by the storage-optimized VM sizes.
- **Load Balancer support** – To create a high availability cluster, the VM tier must support the Azure load balancer.

Test your application deployment to properly select the VM size and the necessary IOPS performance. You can limit the necessary IOPS by configuring the firewall VM to stream them to the Control Center or external log collector, instead of storing logs locally. Local storage is used only for the swap partition. Auto scaling is not supported.

For more information, see Public Cloud. and https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-size-specs/.

Select the Storage Account Type

Create a matching storage account based on the max IOPS supported by the VM tier and series you selected. A fast premium storage account or multiple data disks are recommended for heavy logging and other disk-intensive tasks. Premium storage accounts must also be supported by the VM size.

Because there are limits on the maximum IOPS a storage account is able to service, you should not have more than 10 OS disks per storage account. If I/O performance issues occur, add additional data disks in a second storage account. Be aware that if a storage account is shared with another resource in an auto-scaling group, these limits can be reached very quickly when the deployment scales up. To avoid this problem, use a separate storage account for the firewall VMs.

Plan Networking

The virtual network is split into one subnet per tier. By default, each VM uses the first IP address in its subnet as the default gateway. The route table in the Azure networking stack contains the following routes:

- Traffic between the subnets in the virtual network
- Traffic to the Internet
- Traffic between different virtual networks using the Azure VPN Gateway
- Traffic from the virtual network to networks connected via the Azure VPN Gateway.

Using the default routing table allows direct access to the Internet for all VMs and direct access from the Internet for all VMs with public IP addresses attached to them. If you need to change that, you must configure a user-defined routing table. User-defined routes are always preferred over the default routes. The route tables are then applied on a per-subnet basis to the subnets. In the multi-tiered setup above, two route tables must be created.

Backend Route Table

This route table is applied to all subnets sending traffic through the firewall, with all internal traffic passing through the firewall. Create the following UDR routes in the Azure route table:

- Traffic to the Internet (0.0.0.0/0) is routed via the firewall VM.
- Traffic to the backend subnets is routed via the firewall VM.
- Traffic to the SC networks is routed via the Secure Access Concentrator VM.

For the firewall VM to forward the traffic, you must enable IP forwarding for the attached NIC in Azure. Create access and application rules in the Forwarding Firewall service to implement your policies.

NGF Route Table

This route table is only needed if a Control Center or Secure Access Concentrator is in the same virtual network. Add routes for the Control Center VIP networks and the SC networks, and then apply the route table to the firewall subnet. When choosing the networks for the VIP and SC networks, make sure you use networks outside of the virtual network address pool.

For more information, see How to Configure Azure Route Tables (UDR) using Azure Portal and ARM and Access Rules.

High Availability

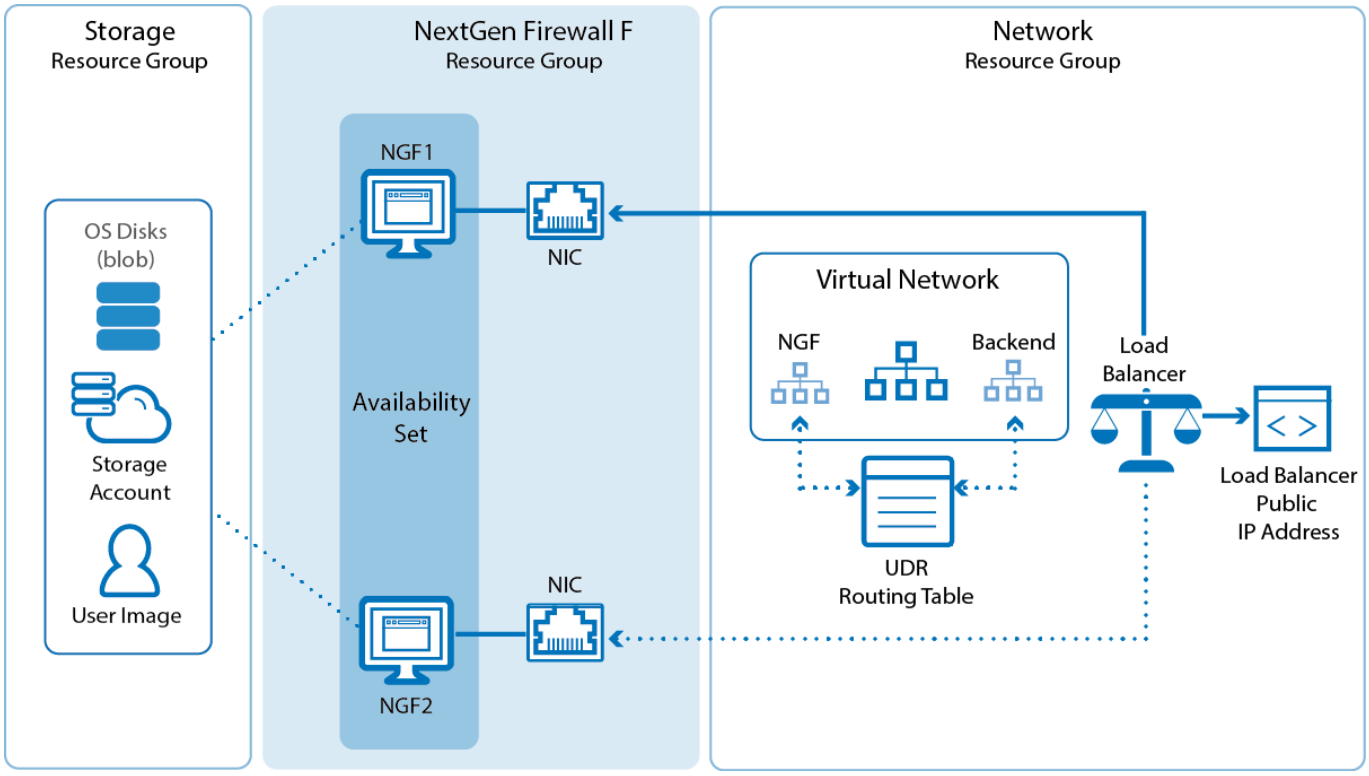

Microsoft requires at least two virtual machines of the same type in an availability set to be included in the Service Level Agreement for Virtual Machines. When deployed as an HA cluster, the firewall VMs continuously sync session information and configuration changes. When one VM fails, or is shut down for maintenance, the secondary firewall takes over the services. Since the IP addresses are fixed and non-transferable between VMs, incoming traffic is handled and delivered via the Azure Load Balancer. The load balancer polls a service running on the virtual server of the active firewall. When the virtual server fails over to the secondary firewall, the load balancer will follow with a latency of a couple of seconds. For backend servers to be able to send traffic through the currently active unit, the F-Series Firewall rewrites and monitors the Azure routing table directly so that routing entries always point to the active unit of the HA cluster. Since Azure does not support floating IP addresses, all active sessions will time out whenever a failover occurs. Barracuda NextGen Control Centers in Azure do not support high availability configurations.

For more information, see High Availability in Azure, How to Configure Azure Route Table Rewriting for HA Clusters using ARM and How to Configure Azure Load Balancer for HA Clusters using PowerShell and ARM.

Connectivity

The F-Series Firewall plays a key role in connecting your users and on-premises networks to your data and applications in the cloud. Remote users connect via client-to-site and SSL VPN, whereas remote

networks use site-to-site tunnels to create a flexible, failsafe hybrid network.

Site-to-Site VPN Tunnels

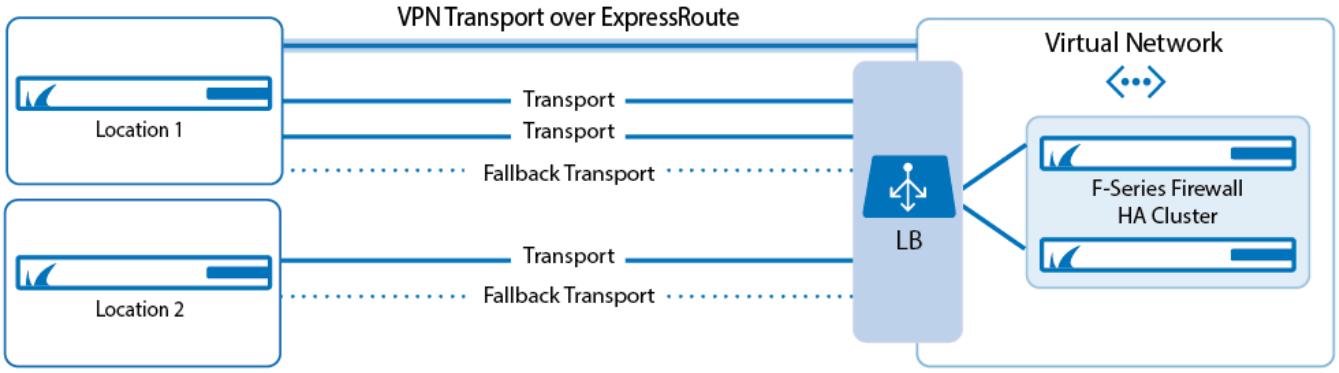

TINA VPN Tunnels

The advanced VPN capabilities of the F-Series Firewall let you optimize your available bandwidth, even when it is spread out over multiple ISP connections and an Express Route MPLS. Traffic Intelligence allows for multiple transports in a VPN tunnel, with each transport using a different connection. The transport used to send the traffic is determined by the settings of the connection object used in the access rule matching the traffic. This lets you select a specific transport or cycle through the available transports to evenly distribute your traffic. When a transport fails, traffic is automatically redirected to the next transport configured in the connection object. The failover from one transport to the other is completely transparent to the user. This means that as long as one transport is operational, the firewall can still send traffic through the VPN tunnel.

For more information, see Site-to-Site VPN, Traffic Intelligence, Application Based Provider Selection, and WAN Optimization.

ExpressRoute

The transport used to transmit data over the ExpressRoute MPLS can be optimized for performance or securely encrypted like the transports transmitted over the public Internet. If you are using ExpressRoute with unlimited data, it can make financial sense to transmit as much data as possible over the MPLS connection since you are not charged per GB. In this case, traffic shaping can help you optimize your traffic flow through the tunnel. As long as bandwidth is available, all traffic is transmitted as fast as possible. However, when the line is at capacity, business-critical traffic is prioritized over other, less important traffic such as file transfers or background tasks. Traffic shaping can be applied on a per-access or application rule basis. You can also reduce the amount of traffic transmitted through the VPN tunnel by enabling VPN compression or, for selected protocols, WAN Optimization.

For more information, see Traffic Shaping, TINA Tunnel Settings, and WAN Optimization.

Third-Party VPN Gateways

To connect to remote third-party VPN gateways, the F-Series Firewall supports IKEv1 and IKEv2 site-to-site IPsec VPN.

For more information, see IPsec VPN Tunnels, How to Configure a Site-to-Site IPsec IKEv1 VPN Tunnel, and How to Configure a Site-to-Site IPsec IKEv2 VPN Tunnel.

App Redirect for the VPN Service in Azure

Because the virtual server listens on a loopback address, typically 127.0.0.2, you must create application redirect rules for the service to be redirected to the listener on the loopback interface. If you are configuring the service on an HA cluster, use both the internal IP address of the primary and secondary firewall as the destination of rule in order to ensure that traffic is forwarded no matter which firewall is currently active. For IPsec VPNs, also verify that **Use dynamic IPs** is disabled in the VPN Settings. Check the listeners by double-clicking on the process on the **CONTROL > Resources** tab.

For more information, see Best Practice - Service Configurations in the Public Cloud.

Remote Access for Users

To connect remote users or mobile clients, the F-Series Firewall offers three options.

- Client-to-site VPN using Barracuda VPN Client, IKEv1, or IKEv2 IPsec clients.
- Web based via SSL VPN
- CudaLaunch

Previously, your remote users would connect to your on-premises firewall using up a lot of bandwidth, especially if your VPN clients are configured to access the Internet through the client-to-site VPN. By using your F-Series firewall in the cloud as the VPN server you can take advantage of the fast datacenter connections. They can directly access the applications and services in Azure behind the firewall and through a site-to-site or ExpressRoute connection also your on-premise resources.

For web-based applications that you may not want to make publicly available, you can allow access through the web-based SSL VPN portal or through CudaLaunch on your mobile devices. CudaLaunch also doubles as a VPN client on your Android phone and will configure your iOS VPN connection. This enables you to use apps on your mobile device to directly access services through the full client-to-site VPN connection.

For more information, see Client-to-Site VPN, SSL VPN, CudaLaunch, and Application Control

NextGen FSC-Series

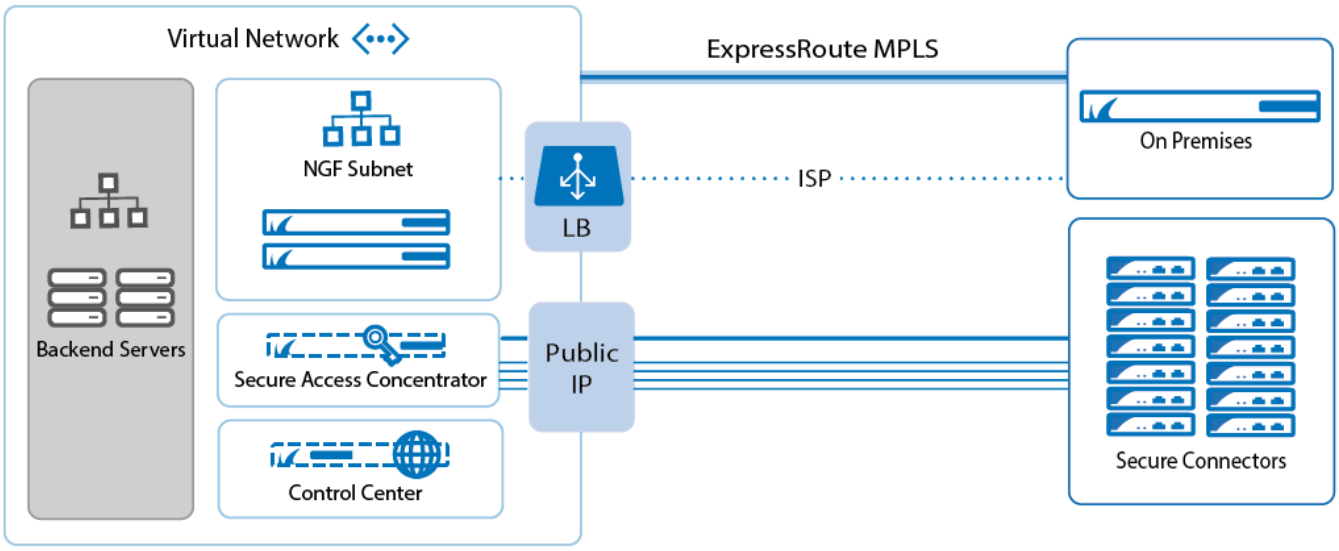

The Barracuda NextGen Firewall offers large-scale remote access capabilities for IoT devices, machine-to-machine communication and remote micro-networks. A large number of small Secure Connector (SC) appliances connect via TINA VPN to the Secure Access Concentrator (SAC). The SAC is deployed similarly to the Control Center in a separate subnet. The SAC applies your application control policies to the traffic exiting the SC VPN tunnels and then forwards the management traffic to the Control Center. Before creating the SC networks, verify that the SC networks do not overlap with the address space used for your virtual network. You can assign a public IP address to the SAC, or use the F-Series Firewall VM to forward all the traffic to the SAC. Using a separate public IP address frees up bandwidth on the F-Series Firewall VM. To correctly route traffic for the SCs, create and assign a route table to the SAC subnet and create a default route using the firewall VM as the gateway. Add routes for the SC networks to both the backend and the firewall subnet route tables in order to be able to access the SCs through a site-to-site VPN and the backend networks. Use access rules on the firewall and SAC to restrict access and to apply application control policies to the traffic exiting and entering the SC VPN tunnels.

For more information, see NextGen FSC-Series.

Monitoring and Logging

The firewall not only applies the security policies to your traffic, but also offers and monitors real-time visibility and comprehensive reporting. The admin thus always knows how traffic is flowing through the virtual network and where potential bottlenecks or misconfigurations lie. By configuring FW Audit, all connections are logged and can either be stored on the firewall or streamed to the Control Center or to an external log collector, such as Splunk, for further processing. If your firewall is acting as a VPN hub with all traffic from the client-to-site connections passing through the firewall monitor, dashboards and Report Creator help you analyze which resources are being accessed and how much

bandwidth is being used for a particular service.

For more information, see FW Audit, Logging, Monitor Page, Barracuda Report Creator, Splunk Integration.

Internal Protection - Creating and Enforcing a DMZ

After routing is in place, access rules on the firewall control which connections are allowed, blocked, or redirected between the different tiers. By default, all connections are blocked. To allow, for example, the web server tier to access the application servers, you need to create an access rule with the source set to the subnet the webservers are in and the destination the application servers are in. Select the bidirectional check box to allow connections to be initiated both ways. Otherwise, the source networks can create a connection that matches this rule. To be more restrictive, you can also limit access based on the service type.

When creating the access rules, use the firewall Live and History pages in NextGen Admin to see which connections are allowed or blocked.

For more information, see Access Rules, Live Page and History Page.

Publish Cloud-Hosted Applications

Similar to the access rules used to control the internal traffic between the subnets/VMs of the virtual network, access rules secure and manage access to your public-facing services. In addition to limiting which services are allowed to connect via access rules, the Intrusion Prevention System (IPS) actively monitors traffic matching the access rule for malicious activities and can also block suspicious traffic. You can create, edit, and override default and custom IPS signature handling policies. You can customize the IPS policies and apply them on a per-access-rule basis.

You can also limit access to your services based on the geographic location of the source IP address. Create a block access rule and use a geo location network object containing all countries you wish to deny access.

For more information, see Access Rules, Geo Location based Network Object and Intrusion Prevention System (IPS).

NextGen Control Center – Central Management for Firewalls in Azure

Since cloud deployments can change and iterate very quickly, working with stand-alone firewalls can result in repetitive configuration tasks to replicate the firewall setup for a new deployment. The NextGen Control Center addresses this problem by using customizable default settings for new firewalls and a repository that lets you share configurations between firewalls. When you create a new firewall configuration in a Control Center cluster, the box level settings of the **Default Box** in the cluster are used to create the new configuration. Changing the configuration of the default box lets you reuse the configurations. If you modify the Default Box configuration directly, the scope of these changes is limited to the cluster. To be able to use the same settings over multiple clusters, create repository entries and link them to the default box configuration. New configurations will then automatically use the same links.

For more information, see NextGen Control Center Central Management

Preparing the Configuration

Configure all settings in the **Default Box** node of the cluster that you want to deploy your Azure firewalls to. If you are using only one cluster, you can continue by creating a new firewall configuration in that cluster that will use the new default settings. If you want to use these settings in multiple clusters, copy the configuration to the repository and link each node to the configuration from the repository. Repeat this for each cluster containing an F-Series Firewall in Azure. Be careful if you are mixing cloud, virtual, and hardware firewalls in one cluster: these default settings are model agnostic and are also used when you reset your configuration by selecting **copy from default**.

The configuration for services in the virtual server you assign to the firewall can also be linked to the repository wherever it makes sense. You can also copy the configuration from the repository to get a starting point that requires you only to change the firewall-specific settings, but saves you the work of having to configure the entire service again. If your firewall VMs are all in a cluster and use the same or very similar rulesets, use the distributed firewall service to only administer one ruleset without the need to synchronize changes between multiple firewall services.

For more information, see Repositories.

Deploying a New Firewall

Before you deploy the F-Series Firewall VM, create a new firewall configuration and select the appropriate cluster. Since you have already modified the default box, you need to set only the configurations.

- Box Properties
- Network
- Identity
- Box Licenses

Create a virtual server and, where possible, link to preconfigured repository entries. Otherwise, copy from the repository and customize the configuration for the specific firewall. After completing the configuration, you can now export it in a PAR file and apply it to the new firewall.

For more information, see Central Management, How to Add a new F-Series Firewall to the Control Center, and How to Configure a Remote Management Tunnel for an F-Series Firewall.

Special Considerations for PAYG Licenses

If you are using PAYG (hourly) images to deploy your firewall VM, special consideration must be given to the license of the firewall. Licenses for PAYG images are generated on first boot of the VM and cannot be recreated afterwards. You must export the licenses after starting the firewall VM and before using a PAR file to import the configuration from the Control Center for the firewall you have prepared. Import the license files into the firewall configuration on the Control Center and then create the PAR file for the managed unit.

For more information, see Public Cloud Licensing.

Figures

1. Azure_IG_MultiTier_01-01.png
2. Azure_IG_ResourceGroups.png
3. Azure_IG_Deployment.png
4. Azure_IG_ResourceGroups_ha.png
5. Azure_IG_Connectivity.png
6. Azure_IG_S-Series.png