Implementation Guide - NextGen Firewall in AWS

https://campus.barracuda.com/doc/58490883/

Amazon Web Services follows the shared security responsibility model. Securing and limiting access to the resources deployed on the cloud platform is the responsibility of the cloud architect. This encompasses both frontend access, as well as access to backend EC2 instances. The remote connectivity options of the Barracuda NextGen Firewall F-Series enables you to enforce strong, consistent authentication and encryption of all your traffic. Barracuda Networks supplies several reference architectures as a starting point. Select the reference architecture by use case, the required level of fault tolerance, and acceptable failover times. After you have chosen the architecture, a detailed description and CloudFormation template will help you to understand and quickly deploy the necessary AWS resources.

Use case: North-south firewalls

North-south firewalls secure incoming and outgoing traffic from the private instances in your VPC. All incoming and outgoing traffic for the VPC is routed through the firewall. Incoming traffic can be from the public Internet; however, the firewall also connects remote users via client-to-site VPN and SSL VPN, as well as site-to-site VPN connections to your on-premises datacenters and offices. The policies set through the access rules in the Forwarding Firewall allow you to enable the following advanced features of the NextGen Firewall on a per-access-rule basis:

- **Intrusion Prevention System** – Actively monitor traffic for malicious activities and, if necessary, block suspicious traffic.
- **Virus Scanning / ATP** – Scan incoming files for viruses and advanced malware.
- **DNS Sinkhole** – Use the firewall as the DNS server for the instances in your VPC and intercept traffic to known-bad FQDNs.
- **URL Filtering** – Block access to URL categories.

Select the reference architecture depending on how many VPCs you must protect:

Secure single VPC

To secure a single VPC in one region, use a high availability firewall cluster with route shifting reference architecture. The firewall instances are deployed into two Availability Zones for fault tolerance. All instances in the private subnets send traffic through the active firewall. In case the primary firewall goes down, the secondary immediately takes over and rewrites the AWS routes tables for the private subnets to send traffic over the now-active secondary firewall.

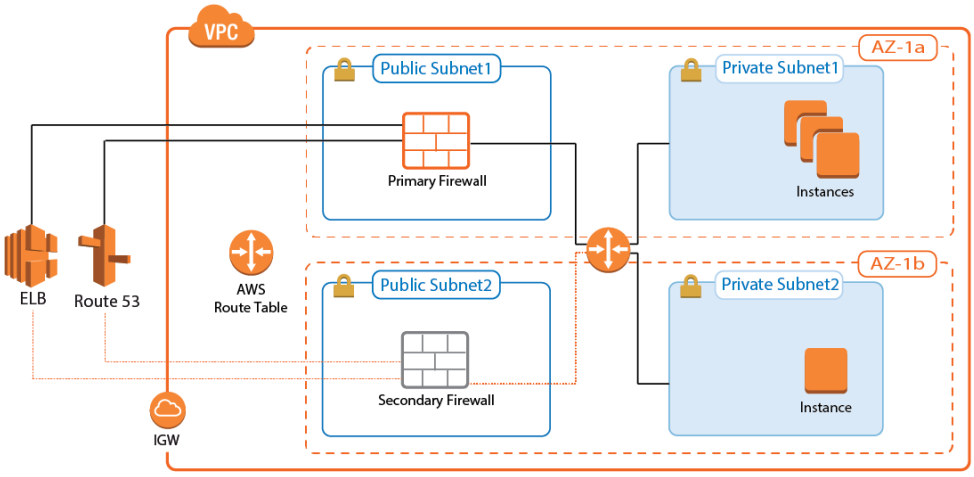

To deploy this solution, go to AWS Implementation Guide - High Availability Firewall Cluster with Route Shifting.

Secure multiple VPCs

If you must secure multiple VPCs distributed over several AWS regions, use the transit VPC reference architecture. Two NextGen Firewall F instances are deployed into a central VPC that acts as a VPN hub for the spoke VPCs. Each spoke VPC is connected to the firewalls in the transit VPC via the AWS VPN Gateway.

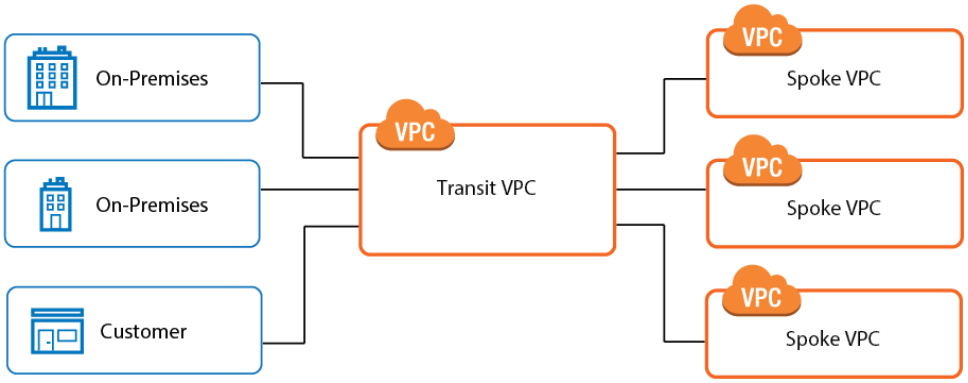

To deploy this solution, go to AWS Implementation Guide - Transit VPC using NextGen Firewall.

Use case: East-west firewall

To secure east-west traffic in your VPC, deploy the NextGen Firewall F instance as a segmentation firewall.

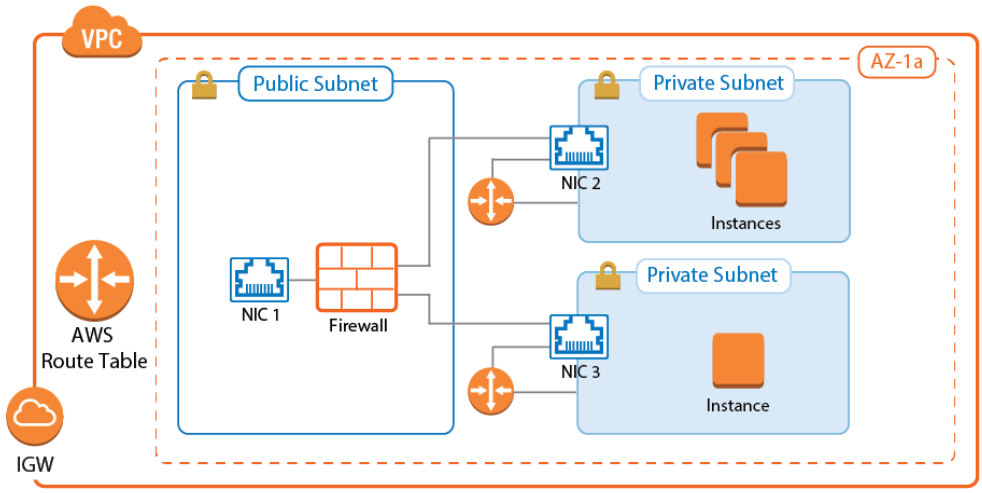

By default, all traffic within the VPC is routed over the AWS default gateway. To be able to route the traffic over the firewall, the local route for internal VPC traffic must be circumvented. This route cannot be overridden by other more specific routes, nor can it be changed to use the firewall as the gateway instead. Using a combination of a single firewall instance with multiple network interfaces and adding a route on the client instances allows you to sent traffic from one private subnet to the other over the firewall instance. This allows you to apply security policies and gain real-time visibility into the connections between instances in your private subnets. Since all network interfaces of an EC2 instance must be in the same Availability Zone, it is not possible to use a high availability cluster as a segmentation firewall. All subnets in the VPC must be in the same Availability Zone. If you are using multiple Availability Zones in your VPC, a firewall must be deployed to each AZ.

To deploy this solution, go to AWS Implementation Guide - Segmentation Firewall for Single AZ VPCs.

Figures

1. multi_AZ_routeshifting_ha0.png
2. transit_vpc_overview.png
3. segmentation.png