AWS Implementation Guide - Transit VPC using NextGen Firewall

https://campus.barracuda.com/doc/58490910/

Connecting multiple VPCs to multiple locations, such as your datacenter or customer offices, can cause significant configuration overhead, especially if VPCs are frequently added and removed. For example, adding a new VPC requires configuration changes to each on-premises location. A second weak point is the communication between the VPCs. To share common resources, VPCs must be peered if they are in the same region; otherwise, the traffic must be routed through your datacenter.

However, you can reduce the number of VPN connections required by each device participating in the network by using a central VPC as a transit VPC and arranging the VPCs in a hub and spoke topology. The transit VPC uses an active-passive high availability cluster made up of two Barracuda NextGen Firewall F-Series instances in the public subnets.

Since the AWS VPN gateway accepts only IP addresses for customer gateways and NOT hostnames, an Elastic Load Balancer cannot be used for the VPN tunnels connecting the firewall cluster to the spoke VPCs. To work around this, the spoke VPCs use VPN connections configured to connect to both the Elastic IP (EIP) addresses of the primary and secondary firewalls. During normal operations, only the VPN tunnels to the primary firewall are up (the tunnels to the secondary are down). Consequently, all traffic is sent through the primary firewall. If the primary firewall goes down, or needs maintenance, the virtual server fails over to the secondary firewall. Immediately, the VPN tunnels configured to connect to the EIP of the secondary firewall come up, and all traffic is now sent via the secondary firewall. When the secondary firewall is active, the VPN tunnels from the primary firewall to the spoke VPCs are down.

The routing is handled by BGP. The spoke VPCs learn the default route from the firewall and send all traffic through the VPN gateway and the transit VPC high availability firewall cluster. The firewall learns the spoke VPC networks propagated by the VPN Gateway. When a spoke VPC is added or removed, BGP automatically propagates the changes to all connected networks.

Connecting to your on-premises network also works with a combination of BGP over TINA or IPsec tunnels. The VPN tunnels do not have to use the public IP address of the firewall as the destination; instead, they can use an FQDN defined in a Route 53 record set with a failover routing policy. The health check defined for the Route 53 record continuously checks if the primary firewall is active. When the health check fails, the IP address of the secondary firewall is automatically returned. Using an Elastic Load Balancer is not possible because the ELB is TCP only. For TINA tunnels, Route 53 or ELB is optional since TINA also works with two public IP addresses as the destination.

Shared services used by all spoke VPCs can also be located in the transit VPC. For example, if all applications require data from the same on-premises database, creating a read-replica in the transit VPC can reduce the traffic to the primary database server in your datacenter. Another example is when the web applications in the spokes share a Barracuda Web Application Firewall in the transit VPC.

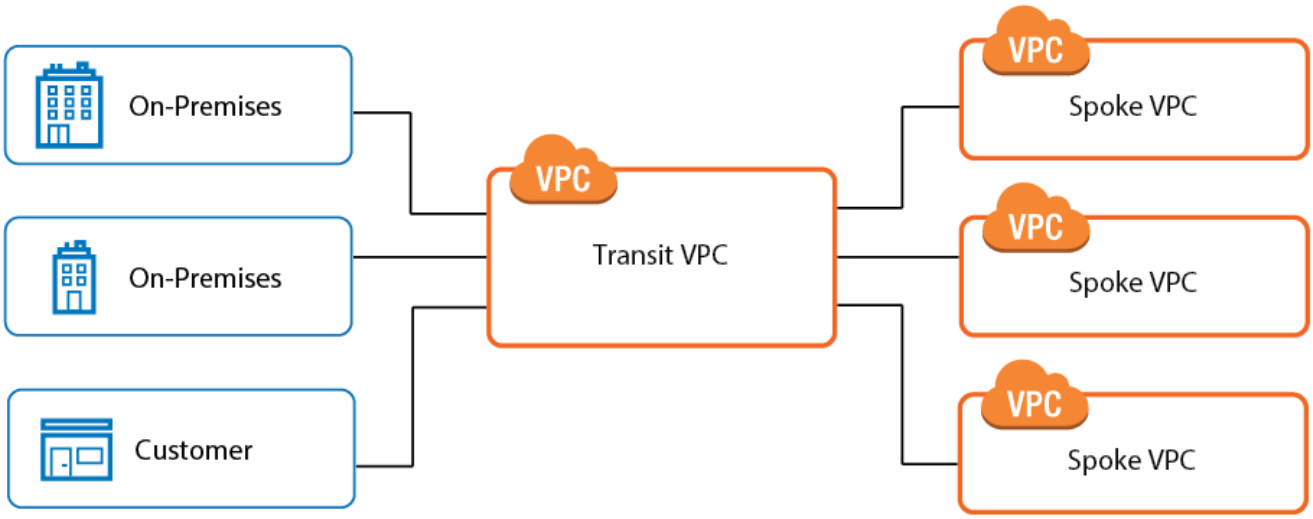

Example CloudFormation template

To deploy the AWS infrastructure of this architecture quickly, use the following CloudFormation template. The NextGen Firewall must be configured manually. This template will deploy the following AWS resources:

- One transit VPC with two Barracuda NextGen Firewall (PAYG) m3.medium instances.
- Three t2.micro Linux instances, one for each of the private subnets in the spokes and one in the private subnet of the transit VPC.
- Two spoke VPCs. Each spoke VPC has an AWS VPN gateway, two customer gateways, and two VPN connections.

Download the **NGF_TransitVPC.json** template and parameter file from the Barracuda Network GitHub account: https://github.com/barracudanetworks/ngf-aws-templates and the Network Diagram containing IP addresses.

For step-by-step instructions on how to deploy a CloudFormation template, see How to Deploy an F-Series Firewall in AWS via CloudFormation Template.

Limitations

- VPC networks may not overlap with each other.
- AWS VPN gateways require an IP address as the endpoint. Consequently, the Elastic Load Balancer in front of the firewall HA cluster cannot be used for the VPN connections to the spokes.

- BGP routing must be configured properly on the firewall for the transit VPC to be able to relay traffic between spokes and the on-premises networks.
- The Elastic Load Balancer only performs load balancing.

Deploying a transit VPC

Complete the following configuration steps to deploy two firewalls into a transit VPC. For more detailed descriptions, follow the links for step-by-step instructions.

Create a VPC

Create a VPC in an AWS region of your choice. Make sure the network assigned to the VPC is sized correctly and does not overlap with existing or planned on-premises networks.

- **Public subnets** – These are the subnets for the firewall instances. Only the firewalls have public IP addresses.
- **Private subnets** – These are the subnets for all instances without external connectivity. All traffic is routed over the firewall in the same Availability Zone. To create a robust architecture, create the private subnets in different Availability Zones.

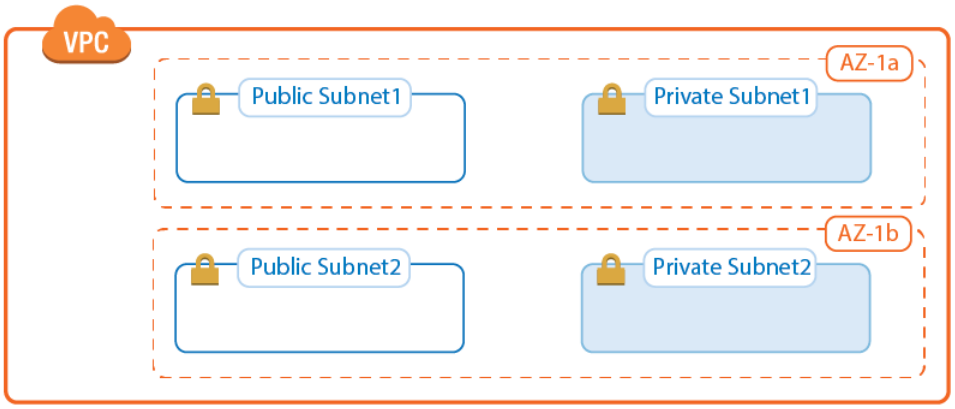

For step-by-step instructions, see Step 1 to 5 in How to Configure a High Availability Cluster in AWS using the Web Portal.

Deploy the firewalls

Launch a firewall instance into each public subnet using either the Barracuda NextGen Firewall BYOL or PAYG images from the AWS Marketplace. Before selecting the instance type, make sure you understand the throughput requirements and size the instances accordingly. Since only one firewall is active at a time, a single instance must be able to handle the highest expected traffic load for your setup. The secondary firewall must use the same instance type.

To avoid the public IP address of the instance from changing, use Elastic IP addresses. Using Elastic IP addresses allows the configuration of the VPN gateway to remain unchanged, even when you redeploy your firewall instances.

If you are using the BYOL image type, activate the license. For both PAYG and BYOL, install the available hotfixes as displayed in the **DASHBOARD > Firmware Update** element when logged in via NextGen Admin. The default username is `root` and the default password is the instance ID.

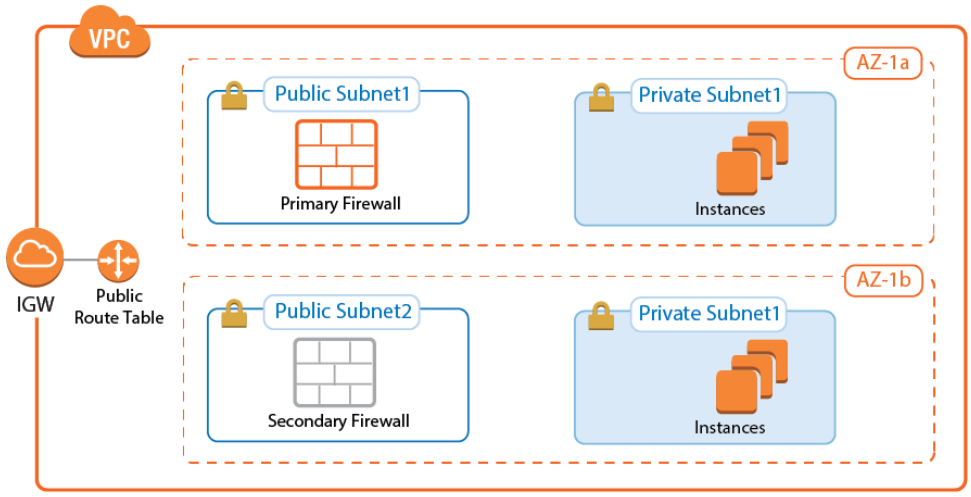

For step-by-step instructions, see Step 6 to 8 in How to Configure a High Availability Cluster in AWS using the Web Portal.

Create AWS route tables

Depending on how you want to access the instances in the private subnets, you must add up to three AWS route tables:

- **Public subnet route table** – Add an AWS Internet gateway and modify the AWS routing table for the public subnets to use the Internet gateway as the target for the default route. Verify that the route table is associated with the public subnets. Instances deployed to the public subnets can now connect to the Internet via the AWS Internet gateway.
- **Two private subnet route tables** – Create a route table for each private subnet and create the default route using the firewall instance in the same Availability Zone as the default gateway.

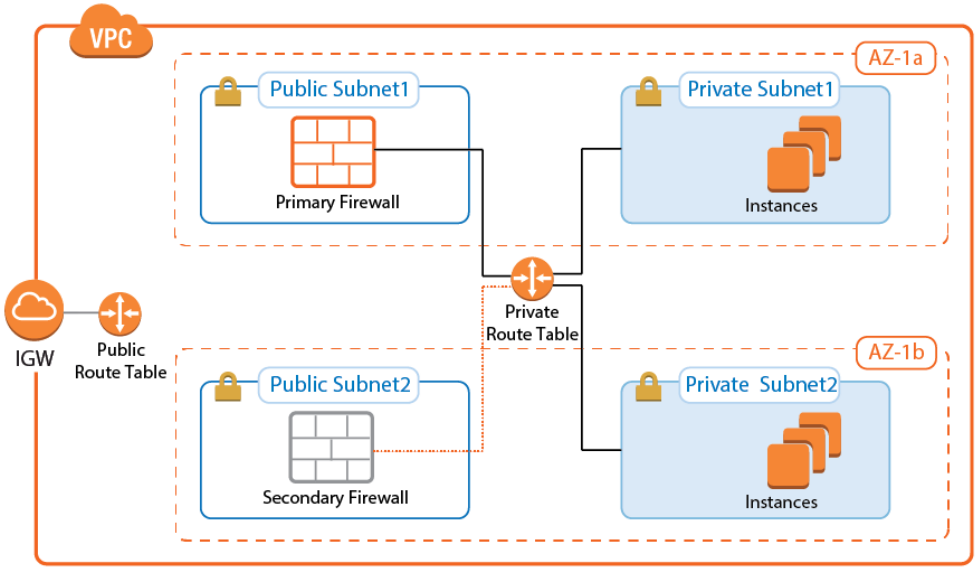

As an alternative to the two private subnet route tables, you can also use rewrite the source and destination address for the connections to the instances in the private subnets. Enter multiple IP addresses as redirection targets of the Dst NAT access rule to load balance traffic between multiple Availability Zones.

For step-by-step instructions, see Step 9 an 10 in How to Configure a High Availability Cluster in AWS using the Web Portal.

Configure Route 53

To be able to reach the cluster through one hostname, configure a health check for the VPN service of the primary firewall and two record sets with failover routing policies. When the virtual server is active on the primary firewall, and the health check is in a healthy state, the FQDN defined in the record set returns the EIP of the primary firewall. When the virtual server fails over to the secondary firewall, the health check fails because the VPN service is no longer running on the primary firewall. The FQDN is now resolved to the EIP of the secondary firewall. To make failover times quicker, use fast 10-second request intervals and low TTL values for the record sets.

For step-by-step instructions, see How to Configure Route 53 for F-Series Firewalls in AWS.

Creating a spoke VPC

Each spoke VPC is made up of one or more private networks connected via the AWS VPN gateway to the transit VPC firewalls. The spoke VPCs do not have an Internet gateway; all traffic will be routed over the firewalls in the transit VPC. The number of spokes supported by the firewall instance depends on how much traffic is sent through the tunnels and which firewall features are enabled.

Create the VPC

Create a VPC in the desired AWS region. Verify that the network assigned to the VPC does not overlap with the transit VPC or the on-premises networks. In the VPC, create at least one private network.

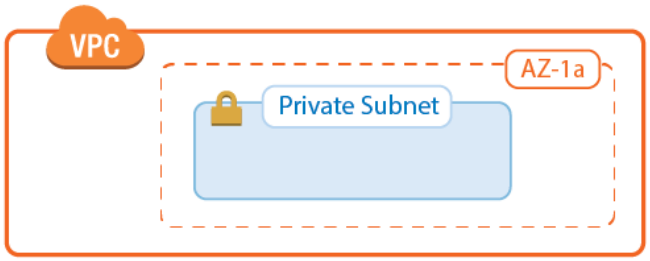

AWS VPN gateway

The AWS VPN gateway connects the EC2 instances in the VPC to the transit VPC via VPN connections. Create a customer gateway for each firewall, and create two VPN connections: one for each transit VPC firewall. Each VPN connection you create for the AWS VPN gateway is made up of two IPsec tunnels. Only one IPsec tunnel pair will be up at the same time. Configure BGP to prefer the first tunnel and to use the secondary tunnel in case the primary is down.

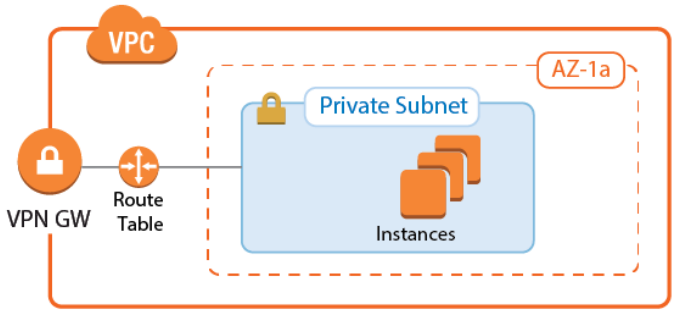

For step-by-step instructions, see Step 1 in How to Configure an IKEv1 IPsec VPN to an AWS VPN Gateway with BGP.

AWS route tables

The AWS route tables can be configured with static routes over the VPN gateway, or they can be configured to learn the routes via BGP. Using BGP has the advantage of being able to control all routing in the firewall's BGP service. However, whether static or dynamic, it is recommended to configure the default route through the VPN gateway. This ensures that all traffic for the VPC passes through the firewalls and that the security policies can be applied in one central location.

Configure the AWS route table for the spoke VPCs to learn the routes propagated by the firewall BGP service. To send all traffic through the transit VPC, propagate the default route to the spoke VPCs. If

propagated routes in the AWS route tables overlap with the local route of the VPC, the local route is always preferred. This applies not only to the local route, but also to all static routes. Static routes are preferred over the learned routes.

Enabling route propagation for AWS route tables

1. Log into the AWS console.
2. Click **Services** and select **VPC**
3. In the **Virtual Private Cloud** section of the left menu, click on **Route Tables**.
4. (optional) Filter the list using the VPC ID.
5. Select the route table for the spoke VPC.
6. In the lower half of the page, click on the **Route Propagation** tab.
7. Click **Edit**.
8. Select the VPN gateway and click **Save**.

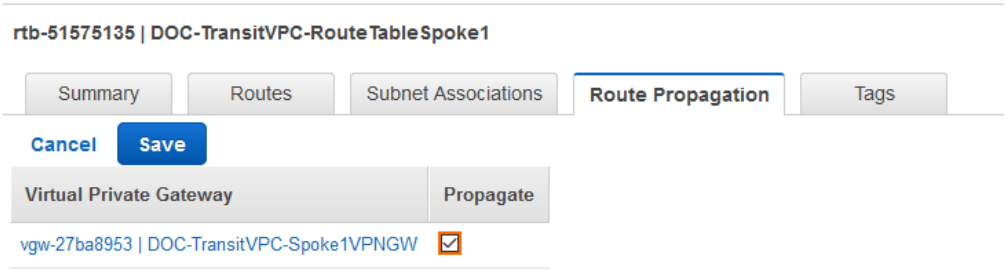

Configure the IPsec tunnels on the transit VPC firewalls

To connect the spoke VPC to the transit VPC, configure four IPsec tunnels: two parallel IPsec tunnels to the EIP of the primary firewall and two for the EIP of the secondary firewall. The AWS VPN gateway treats these connections as if they were separate endpoints. AWS defines a /30 intermediary network for each IPsec tunnel. The IP addresses in this intermediary network are used by BGP. Define BGP neighbors for each next-hop address as per the instructions provided by AWS.

The VPN connection information is unique for each VPN connection and can be downloaded by the right-clicking on the VPN connection. In addition to the encryption settings in the AWS configuration file, the following settings are supported:

- **Encryption** – AES, AES256
- **Hash** – SHA1, SHA256
- **Phase 1 DH-Group** – Group 2 and Group 14-18
- **Phase 2 DH-Group** – Group 1, 2, 5 and Group 14-18

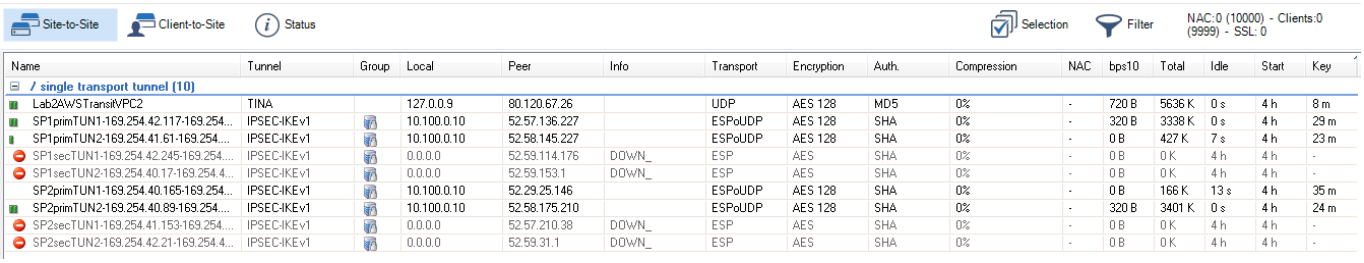

Name	Tunnel	Group	Local	Peer	Info	Transport	Encryption	Auth.	Compression	NAC	bps10	Total	Idle	Start	Key
/ single transport tunnel (10)															
Lab2AWSTransitVPC2	TINA		127.0.0.9	80.120.67.26		UDP	AES 128	MD5	0%	-	720 B	5636 K	0 s	4 h	8 m
SP1primTUN1-169.254.42.117-169.254...	IPSEC-IKEv1		10.100.0.10	52.57.136.227		ESPoUDP	AES 128	SHA	0%	-	320 B	3338 K	0 s	4 h	29 m
SP1primTUN2-169.254.41.61-169.254....	IPSEC-IKEv1		10.100.0.10	52.58.145.227		ESPoUDP	AES 128	SHA	0%	-	0 B	427 K	7 s	4 h	23 m
SP1secTUN1-169.254.42.245-169.254....	IPSEC-IKEv1		0.0.0.0	52.59.114.176	DOWN_	ESP	AES	SHA	0%	-	0 B	0 K	4 h	4 h	-
SP1secTUN2-169.254.40.17-169.254.4...	IPSEC-IKEv1		0.0.0.0	52.59.153.1	DOWN_	ESP	AES	SHA	0%	-	0 B	0 K	4 h	4 h	-
SP2primTUN1-169.254.40.165-169.254...	IPSEC-IKEv1		10.100.0.10	52.29.25.146		ESPoUDP	AES 128	SHA	0%	-	0 B	166 K	13 s	4 h	35 m
SP2primTUN2-169.254.40.89-169.254....	IPSEC-IKEv1		10.100.0.10	52.58.175.210		ESPoUDP	AES 128	SHA	0%	-	320 B	3401 K	0 s	4 h	24 m
SP2secTUN1-169.254.41.153-169.254....	IPSEC-IKEv1		0.0.0.0	52.57.210.38	DOWN_	ESP	AES	SHA	0%	-	0 B	0 K	4 h	4 h	-
SP2secTUN2-169.254.42.21-169.254.4...	IPSEC-IKEv1		0.0.0.0	52.59.31.1	DOWN_	ESP	AES	SHA	0%	-	0 B	0 K	4 h	4 h	-

The IPsec tunnels on the primary firewall are now connected. The tunnels to the EIP of the secondary are down. Since the routing is handled by BGP, no traffic can be sent through the VPN tunnels yet. In addition, access rules in the Forwarding Firewall ruleset must allow the traffic. By default everything is blocked.

For more information, see Step 2 in How to Configure an IKEv1 IPsec VPN to an AWS VPN Gateway with BGP.

Configure BGP on the transit VPC firewalls

The BGP service on the firewall learns and propagates the routes from each location. Create a BGP neighbor configuration for each IPsec tunnel and each on-premises network connected to the transit VPC. If you are not using a static route in the spoke VPCs routing table, propagate the default route to the BGP neighbor for each spoke VPC. The VPN gateway automatically propagates the VPC network via BGP. Since spoke VPCs are always connected by two parallel IPsec tunnels, the route over one IPsec tunnel should be preferred over the other.

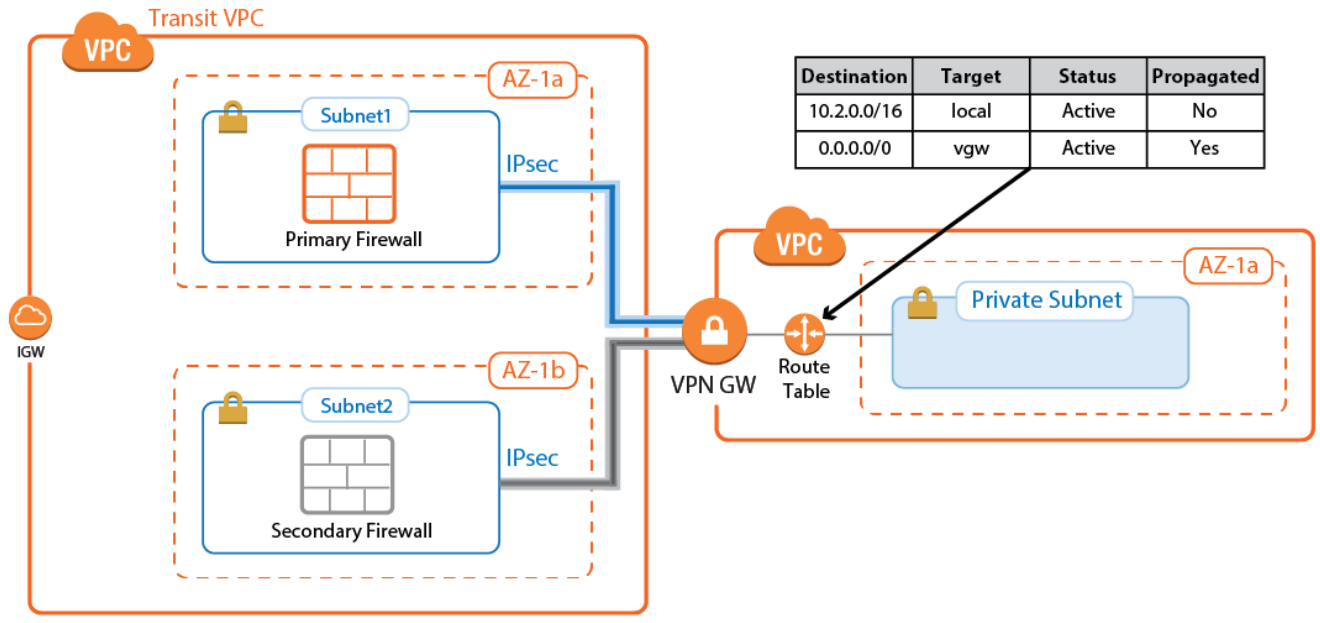

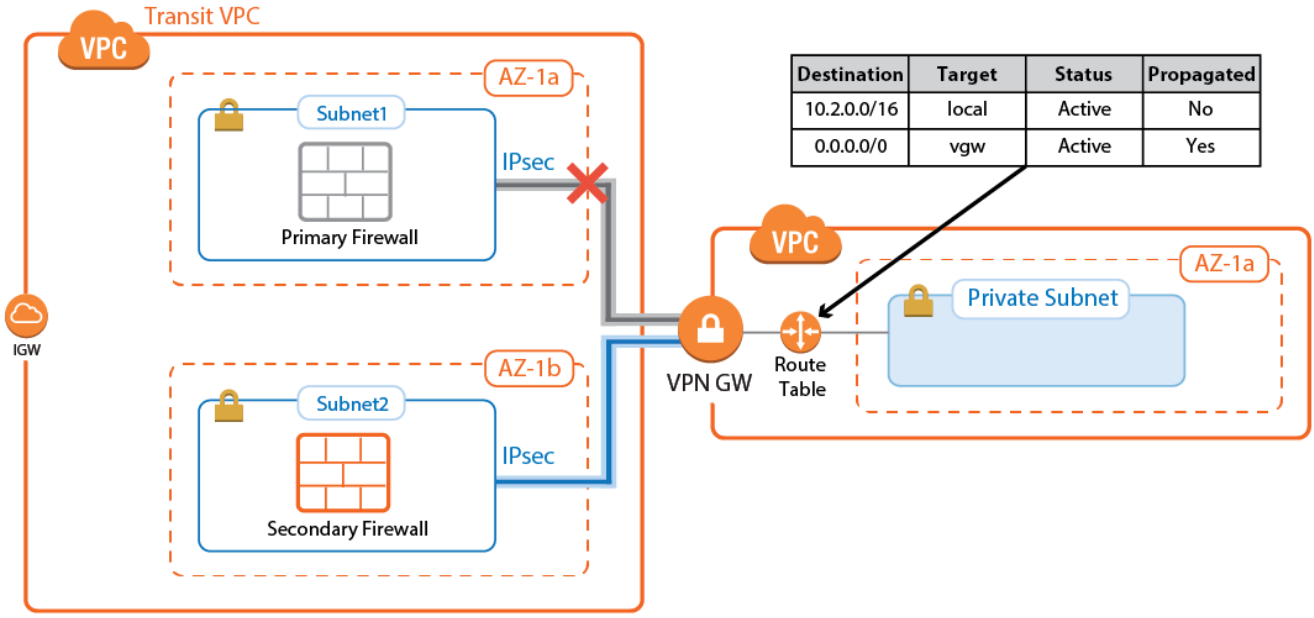

Configure the BGP service on each firewall to exchange information with the BGP service on the other side of the VPN tunnels. Using **Route Maps**, modify the routes learned for the second of the parallel IPsec connections. By lengthening the AS PATH of the IPsec tunnels, traffic is sent through the first tunnel at all times, unless the tunnel is down.

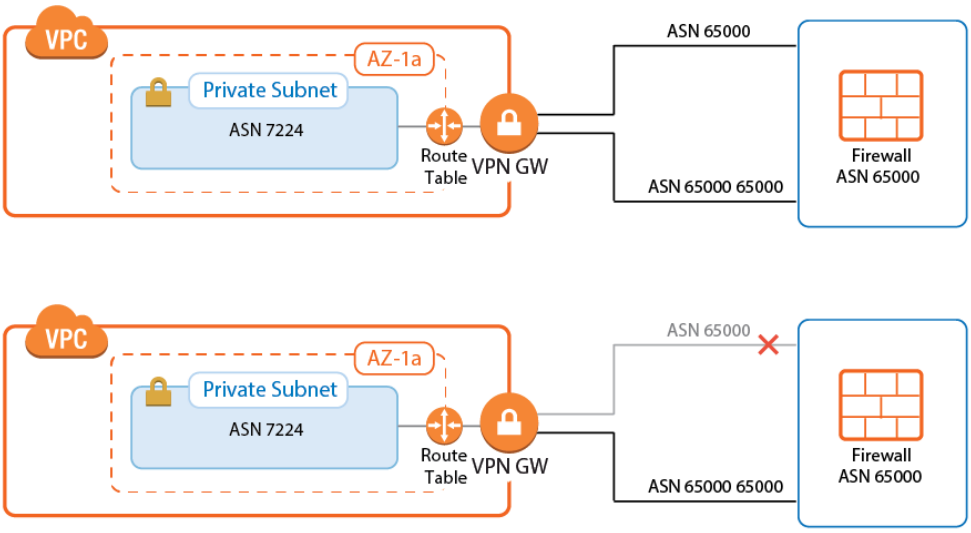

For more information, see Step 3 in [How to Configure an IKEv1 IPsec VPN to an AWS VPN Gateway with BGP](#).

Create access rules to allow traffic

By default, the Forwarding Firewall service blocks all traffic not explicitly allowed by an access rule.

Since all traffic is routed through the transit VPC, create access rules to allow access for individual services and/or entire networks. Access rules allowing traffic through the AWS VPN gateway IPsec tunnels must set the following advanced access settings:

- **Force MSS (Maximum Segment Size)** – Set to 1387.
- **Clear DF Bit** – Set to **yes.**
- **Reverse Interface (Bi-directional)** – If you are using two parallel IPsec tunnels per firewall, set this to **Any**. This allows the traffic to use either IPsec tunnel.

Be sure to sync the access rules on both firewalls to make sure that the behavior is identical no matter which firewall the traffic is sent through.

For more information, see Step 4 in How to Configure an IKEv1 IPsec VPN to an AWS VPN Gateway with BGP.

Launching EC2 instances in spokes

If your transit VPC is created with spokes in a single CloudFormation template, the instances will not have Internet access during launch because the default route over the firewalls is not active yet. To work around this, you can either use multiple templates or, if possible, configure your EC2 instances to use VPC endpoints instead of going over the Internet for the provisioning process.

If your spoke is already connected, verify that access rules are in place that allow the new instance access to all resources required during the provisioning process. If unsure, log into the active firewall and use the **Firewall > History** page in NextGen Admin to check if traffic from the instance was blocked.

For more information, see NextGen Admin History Page.

Connecting to on-premises networks

To be able to forward traffic between your AWS VPC and your on-premises networks, create site-to-site VPN tunnels between the high availability cluster in the transit VPC and the VPN gateways in each remote location.

The networks of the spoke VPC are propagated via BGP over the VPN tunnels. BGP is used to propagate the AWS VPC networks to your on-premises locations. Depending on the remote device, you can use either Barracuda's proprietary TINA VPN or the industry standard IPsec VPN protocol. Failover and preference of the VPN tunnel to the primary firewall is handled by BGP.

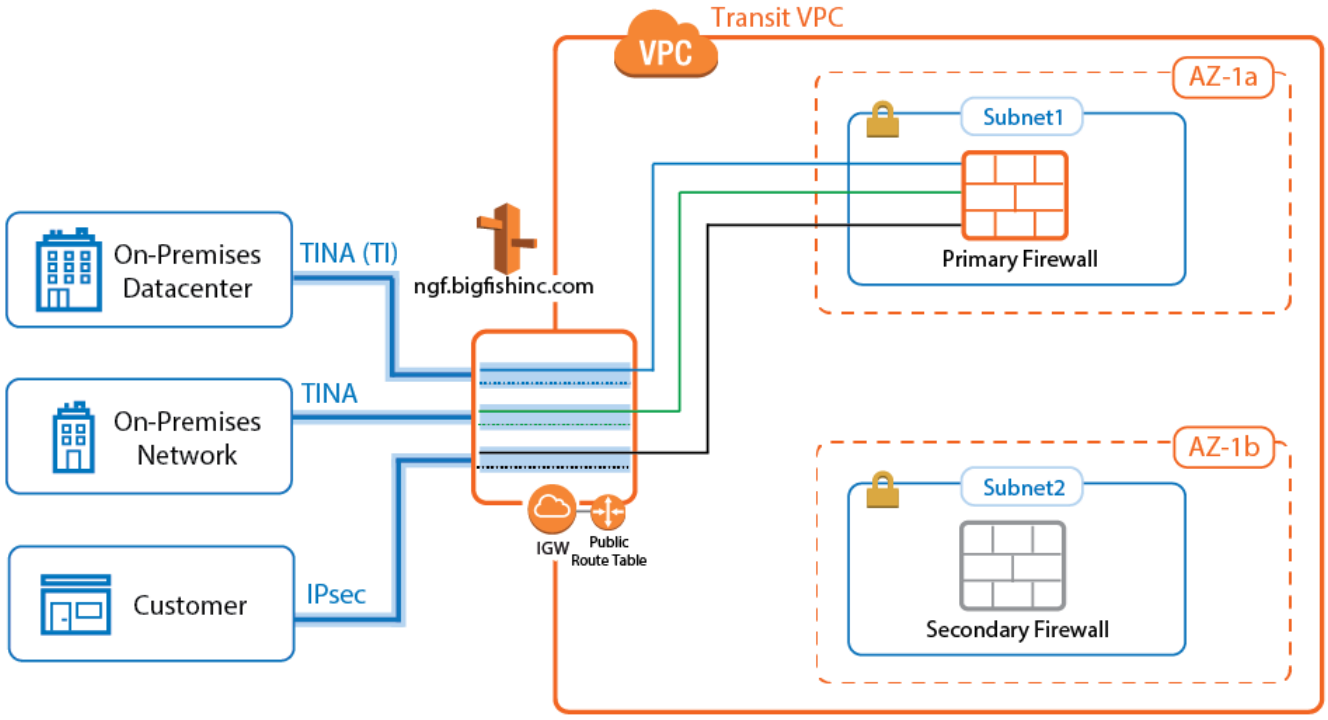

TINA site-to-site VPN tunnels to F-Series Firewalls

If the remote location uses an F-series Firewall, you can take advantage of the TINA VPN protocol. TINA offers many enhancements not featured in the standard IPsec protocol, such as Traffic Intelligence, Traffic Compression, and WAN Optimization. Traffic Intelligence is a logical layer used to manage multiple parallel VPN tunnels (transports) in one VPN tunnel configuration. So if your remote location has multiple Internet connections (perhaps in combination with AWS Direct Connect), all connections can be combined into one VPN tunnel. Traffic intelligence patterns in the connection object of the access rule determine how the traffic is distributed over the VPN transports and failover behavior. WAN Optimization and Compression reduces the amount of traffic sent through the tunnel by using data deduplication.

For step-by-step instructions on how to create a TINA VPN tunnel in combination with BGP, see How to Configure BGP Routing over TINA VPN.

For step-by-step instructions on how to configure Traffic intelligence, see Traffic Intelligence.

For step-by-step instructions on how to configure WAN Optimization, see WAN Optimization.

IPsec site-to-site VPN tunnels to third-party devices

Third-party VPN gateways are connected via IKEv1 IPsec tunnels. The remote device must support routing BGP over IPsec tunnels to be able to learn the routes. Create one IPsec tunnel from the primary firewall and a second IPsec tunnel from the secondary firewall for each location. Using IKEv2 is not supported because it is currently not possible to use BGP over IPsec IKEv2 VPN tunnels with the

F-Series Firewall.

For step-by-step instructions, see How to Configure BGP Routing over IKEv1 IPsec VPN.

Create access rules for on-premises networks

Just like when connecting the spoke VPCs, the firewall blocks all traffic by default. To allow connections to the networks learned via BGP, create pass access rules on both firewalls. These rules must be the same on both firewalls to ensure that if the connection fails over to the secondary firewall, the same policies are applied.

For more information, see Access Rules.

Figures