AWS Reference Architecture - NextGen Firewall Cold Standby Cluster

https://campus.barracuda.com/doc/70584745/

A NextGen Firewall Cold Standby Cluster is a low-cost architecture for AWS deployments that minimizes downtime to only a few minutes in case of failure of the underlying hypervisor hardware, the firewall instance, or the complete Availability Zone. From a technical standpoint, the Cold Standby Cluster is a NextGen Firewall Auto Scaling Cluster with the size set to one. The firewall configuration is securely stored and synchronized through AWS backend services. Replacing the Elastic Load Balancer used in the NextGen Auto Scaling Cluster with a floating Elastic IP allows the use of both TCP- and UDP-based services on the firewall. The default template uses hourly PAYG licensing, but can be modified to use pool licenses for Control Center-managed instances.

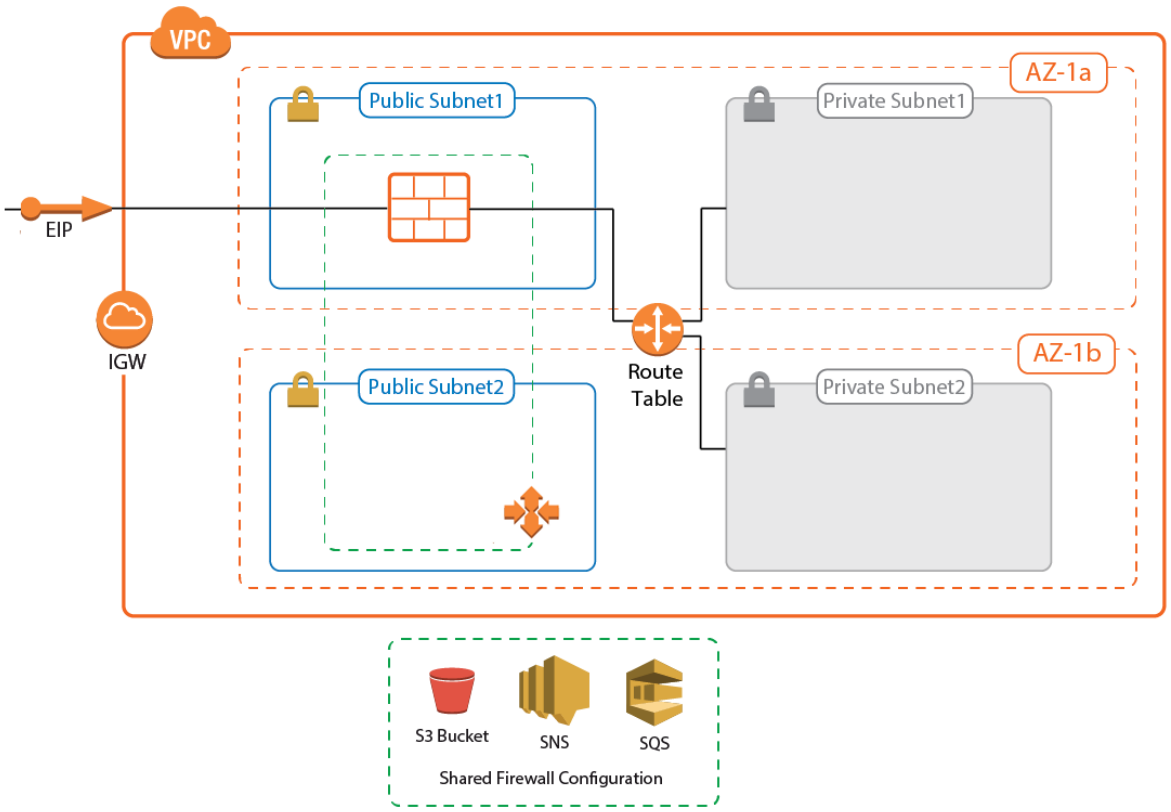

Use Cases for a NextGen Firewall Cold Standby Cluster

The NextGen Firewall Cold Standby Cluster is used to secure access to resources in the private networks of its own VPC, or to deploy in a Transit VPC as a part of a larger cloud infrastructure.

- **Site-to-Site VPN** – One way on-premises to AWS, TINA, and IPsec site-to-site VPN tunnels.

- **Edge Firewall** – Scan for malicious traffic using the built-in IPS and handle access to resources via access rules.
- **Secure Remote Access** – Client-to-site VPN, CudaLaunch, and SSL VPN using TINA, SSL VPN, and IPsec VPN protocols.

Deploying a NextGen Firewall Cold Standby Cluster

The Cold Standby Cluster must be deployed via a CloudFormation template. The template deploys a VPC with public and private subnets in two Availability Zones. The Auto Scaling Cluster is deployed in the public subnets. Instances placed in the private subnets are automatically routed over the active firewall instance.

1. Create an IAM role for a NextGen Firewall in an Auto Scaling group. For step-by-step instructions, see How to Create an IAM Role for an F-Series Firewall in AWS.
2. Download the **NGF_ColdStandby.json** template and parameter file from the Barracuda Network GitHub account: https://github.com/barracudanetworks/ngf-aws-templates.
3. Accept the Software Terms for the **Barracuda NextGen Firewall PAYG** image in the AWS Marketplace.
4. Create a parameter template file containing your parameters values.
5. Deploy the **coldstandby.json** CloudFormation template via AWS CLI or AWS console.

```
aws cloudformation create-stack  --stack-name "YOUR_STACK_NAME" --
template-body YOUR_S3_BUCKET/NGF_ColdStandby.json
```

During deployment, the following resources are created by the template:

- VPC with private and public subnets in two Availability Zones.
- One S3 bucket.
- Automatically created SNS and SQS queues.
- A Launch Configuration and Auto Scaling group for the firewall. The Barracuda NextGen Firewall PAYG image must be used.

After stack creation is complete, wait for one firewall instance to spin up and finish provisioning.

Control Center-Managed NextGen Firewall Cold Standby Cluster

The NextGen Control Center is a central management appliance for F-Series Firewall, that can be deployed as a virtual appliance on-premises or in the cloud. Managing the Cold Standby Cluster with a NextGen Control Center separates the firewall configuration and monitoring from deployment and integration with other AWS services. This is especially useful for highly specialized or large departments with dedicated network security teams and multiple developer teams using automatic

deployments. Managed firewalls are preconfigured on the Control Center. During provisioning of the firewall instance, the configuration is retrieved from the Control Center. Optionally, pool licenses can be used that are bound to the Control Center license instead of the EC2 instance of the firewall. Pool licenses are available in multiples of 5.

For step-by-step instructions, see How to Modify CloudFormation Templates to Retrieve the PAR File from a Control Center.

Modifying the Default CloudFormation Template

To fetch the configuration from the Control Center, the default template must be edited to invoke the getpar command with the information required to be able to connect to the Control Center. If the PAYG images are used, the licenses are sent to the Control Center before retrieving the firewall configuration. For a Cold Standby Cluster, only one firewall configuration is required on the Control Center.

Managing Firewall Configuration at Scale with the Control Center

Managing multiple similar firewall configurations is greatly simplified by using a Control Center. Configuration nodes and cluster-based services, such as the distributed firewall service, are shared across multiple firewall instances. The Control Center also handles pattern updates and hotfixes centrally.

Login and Default Password

Connect to the firewall cluster via NextGen Admin using the Elastic IP. The default password is the instance ID of the firewall. Go to the **Instance** tab of the Auto Scaling group and use the instance ID from the one instance that is running.

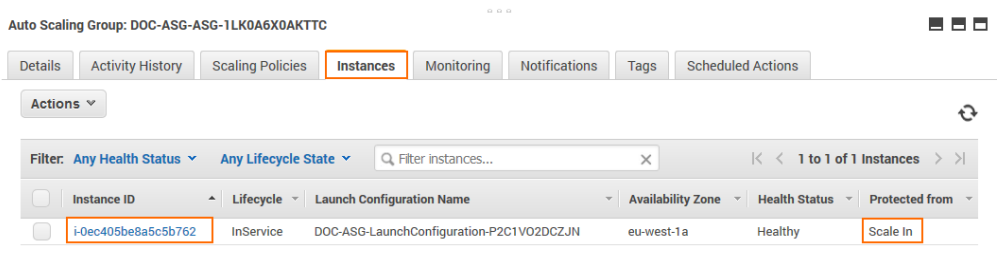

- **IP Address /Name** – Enter the Elastic IP address associated with the firewall.
- **User** – Enter `root`.
- **Password** – The default password is the instance ID of the first instance.

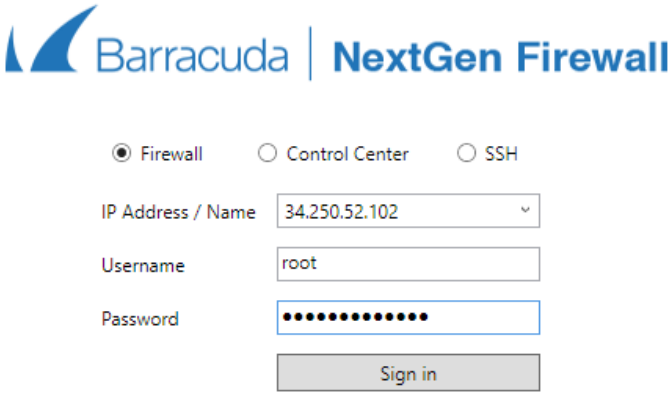

Cold Standby Failover

A failover occurs when the firewall instance is terminated, either due to a cloud-level incident such as emergency maintenance on the hardware, or if the EC2 instance health checks fail for the instance, causing it to be terminated and replaced by the Auto Scaling group. The failover process follows these steps:

1. Health check fails for firewall instance.
2. The unhealthy firewall instance is terminated.
3. The Auto Scaling group launches a replacement firewall instance.
4. Provisioning of the new firewall instance:
    - Configuration from S3 bucket is used.
    - The instance is associated with the Elastic IP.
    - Routes pointing to the firewall instance are rewritten to use the new firewall instance.
    - Hotfixes are installed on the firewall. This may cause the firewall to reboot.
    - (BYOL only) The firewall automatically fetches the configuration from a NextGen Control Center.
5. The new firewall instance is now provisioned.

Note that only routes for which there are AWS CLI commands in the user data section of the templates are changed. The Cold Standby Cluster does not monitor the route tables in the VPC. If this is required, use a NextGen Firewall High Availability Cluster instead.

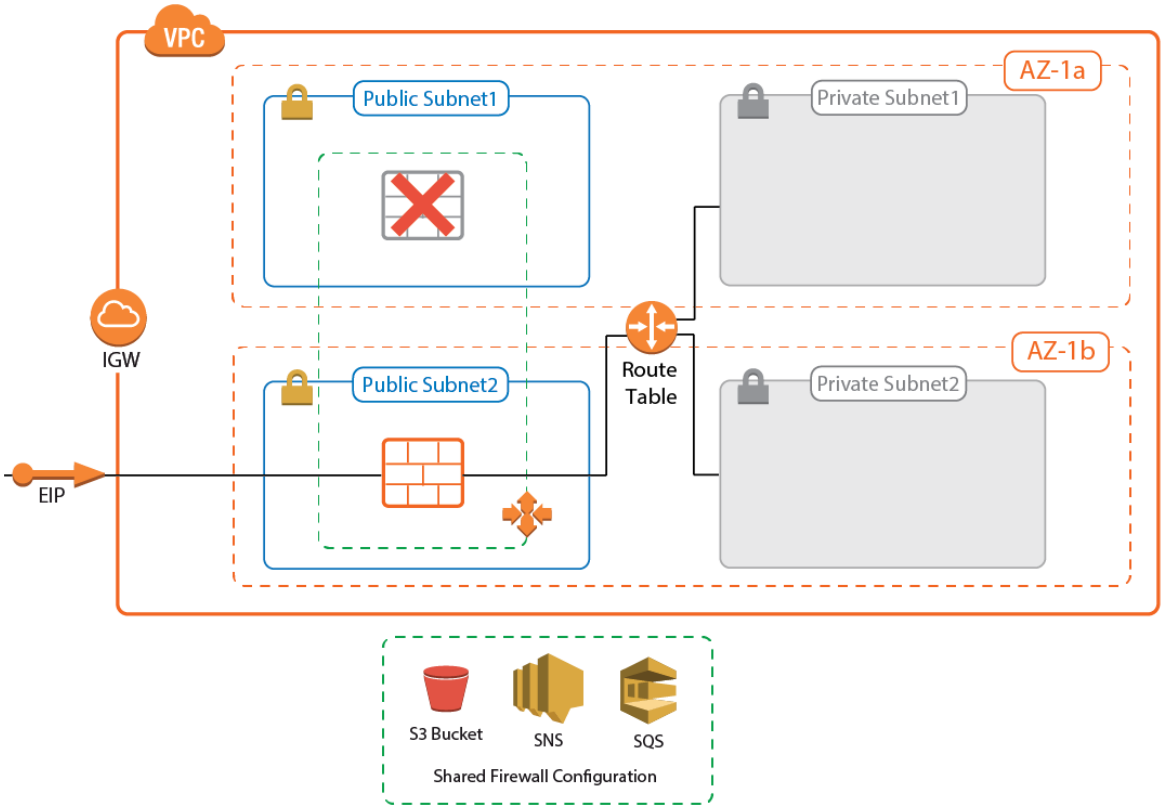

Scaling Up or Scaling Down

Scale the Cold Standby Cluster up by replacing the instance type in the template with a larger instance size and then update the stack. After the stack update is complete, terminate the running instance using the old instance type. The Auto Scaling group now automatically replaces the firewall with an instance using the new instance type.

Installing Hotfixes

Hotfixes are published by Barracuda Networks when an issue requires immediate attention, such as newly discovered security vulnerabilities or critical bugs in the firmware. Hotfixes are installed through the **Firewall** dashboard in NextGen Admin. New instances in the cluster automatically install the same hotfixes when the cluster scales. Most hotfixes reboot the firewall. Consider this when setting the health check grace period because it will take longer for an instance to be ready when multiple hotfixes need to be installed. If the health check is resumed too early, while the firewall is still installing the hotfix or rebooting, the instance is terminated because it is deemed unhealthy. This results in a loop of starting and terminating instances.

To change the health check grace period, modify the **HealthCheckGracePeriod** parameter in the template and update the stack.

```
    "ASG": {
      "Type": "AWS::AutoScaling::AutoScalingGroup",
      "Properties": {
        "VPCZoneIdentifier": [ { "Ref": "PublicSubnetA" }, { "Ref" :"PublicSubnetB" } ],
        "LaunchConfigurationName": { "Ref": "LaunchConfiguration" },
        "MinSize": 1,
        "DesiredCapacity": 1,
        "MaxSize": 1,
        "HealthCheckGracePeriod": 3600,
        "Tags": [ { "Key": "Name", "Value": { "Ref": "ASGName" }, "PropagateAtLaunch": "True" } ]
      },
      "DependsOn": [ "IGWAttachment" ]
    }
},
```

Firmware Update via CloudFormation Stack Update

Although possible, it is not recommended to install firmware updates like hotfixes through the **Update** element on the NextGen Admin dashboard. Instead, replace the AMI in the parameter file of your template and update the CloudFormation stack. The AMI for the new firmware version is listed in the **Manual Launch** tab of the listing for the Barracuda NextGen Firewall PAYG image in the AWS Marketplace.

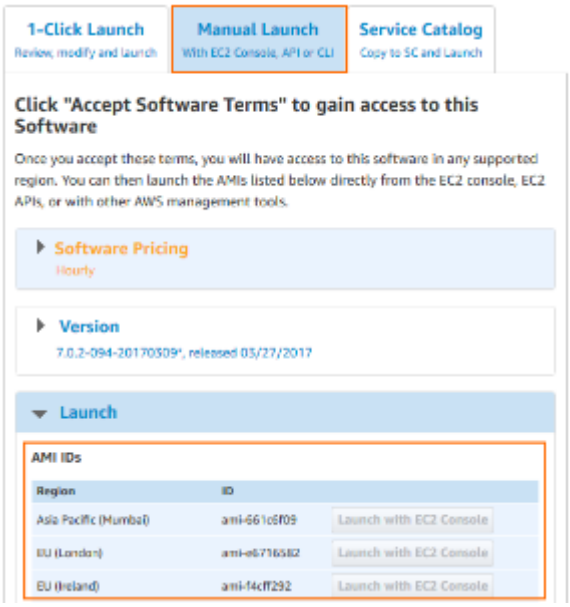

If your templates are stored in an S3 bucket, enter this AWS CLI command to update CloudFormation stack:

```
aws cloudformation update-stack  --stack-name "YOUR_STACK_NAME" --template-
```

```
body YOUR_S3_BUCKET/coldstandby.json --parameter
YOUR_S3_BUCKET/coldstandby_parameters.json
```

After updating the stack, manually terminate the firewall instance. The replacement instance using the new firmware is automatically launched. If the firewall is managed by a Control Center, the cluster the firewall configuration is stored in might have to be migrated after updating the AMI to a new major release, such as 7.1, 7.2, etc...

Site-to-Site VPN Tunnels for Cold Standby Clusters

Site-to-site VPN tunnels transparently connect on-premises networks. The NextGen Firewall supports TINA, IPsec IKEv1, and IKEv2 VPN protocols. Since ESP is not supported, IPsec VPN tunnels must use NAT-T. It is recommended to configure the NextGen Firewall to be the active VPN endpoint. For all instances in the private subnets using the firewall as the default gateway and all instances in VPCs connected via the AWS VPN gateway, the site-to-site VPN tunnel is fully transparent in both directions. If VPC peering is used and the firewall cannot be configured to be the default gateway for the instance, the source IP address for the traffic leaving the tunnel must be rewritten to the address of the DHCP interface of the firewall. Resources in peered VPCs cannot connect directly to the remote networks.

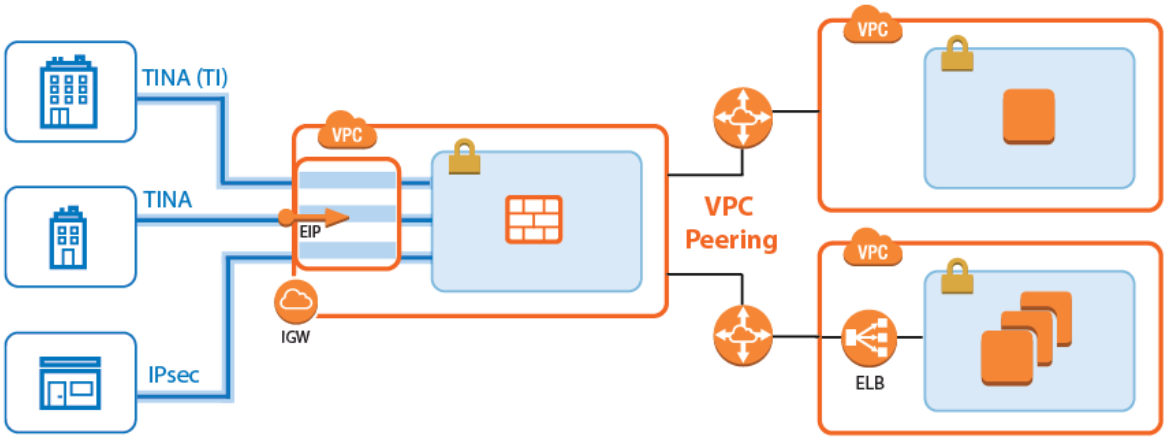

Access Rules

By default, the firewall blocks all traffic. Only traffic matching an access rule with an allowed policy is allowed to pass. If the destination instance is using the firewall as the default gateway, the source IP address does not have to be rewritten; otherwise, the source NAT must be used. Although Dst NAT access rules support basic load balancing, it is recommended to use internal AWS ELBs instead.

For more information, see Access Rules.

Enabling IPS per Access Rule

For the Intrusion Prevention System to scan packets matching an access rule, select the **IPS Policy** in the Pass or Dst NAT access rule. Depending on the configuration of the IPS, malicious traffic patterns are now blocked or reported. For SSL-encrypted traffic, it is recommended to use SSL offloading on the ELB to allow the IPS to analyze the decrypted traffic. This saves processing power on the firewall. If end-to-end encryption is required for regulatory reasons, enable SSL Interception on the access rule and in the IPS configuration to also scan SSL-encrypted traffic.

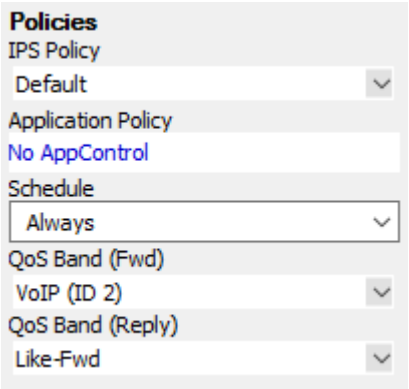

For more information, see Intrusion Prevention System (IPS).

Block Traffic Based on Geographic Location of Source IP Address

Create a network object containing the countries you want to block.

For more information, see How to Create a Geo Location based Network Object.

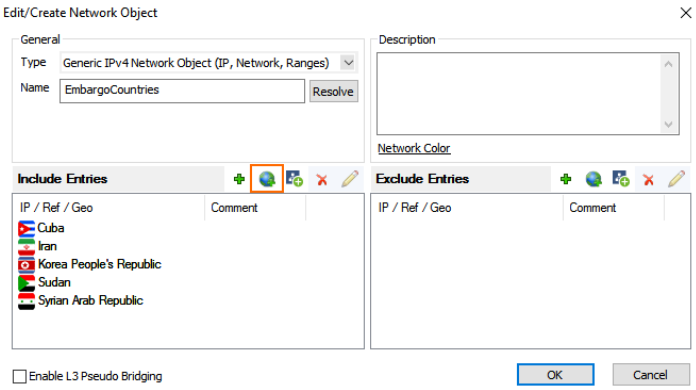

Create a Block access rule using the geolocation network object as the source matching criteria:

- **Action** – Select **Block** or **Deny**.

- **Source** – Select the network object containing the countries you want to block.
- **Service** – Select **Any**.
- **Destination** – Select **DHCP1 Local IP**.

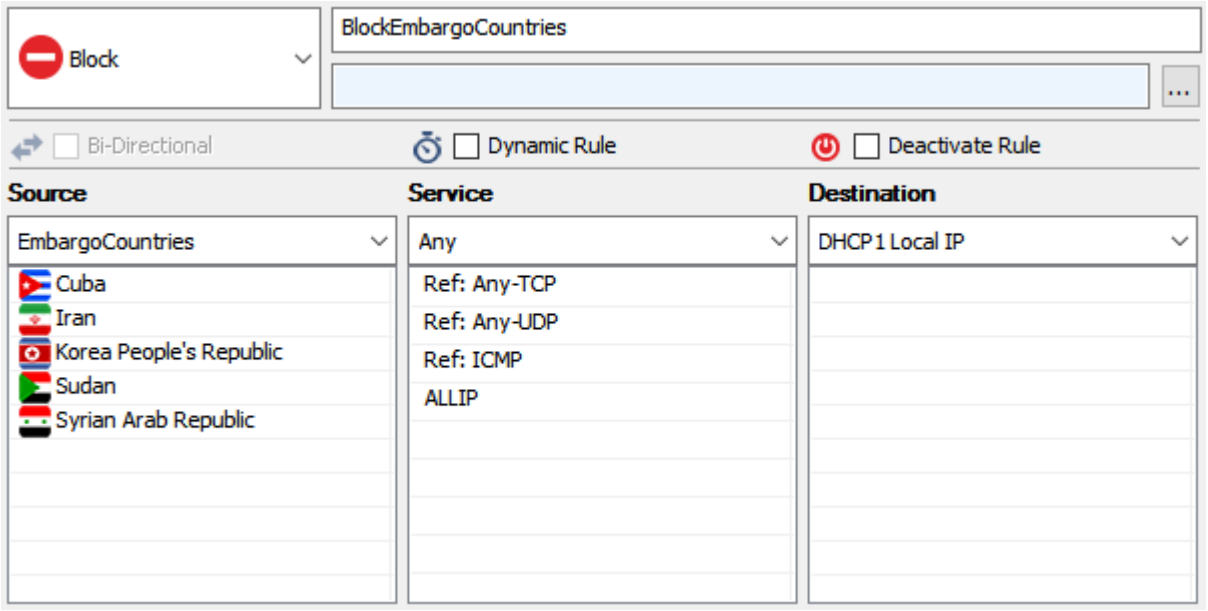

Site-to-Site VPN Tunnel to Backend Services in a Peered VPC

When connecting to services running in a VPC that is peered to the firewall VPC through a site-to-site VPN tunnel, the site-to-site tunnel can only be used one-way from on-premises to the AWS resource. Since the firewall is not the default gateway for the AWS instances running in the private subnets, the source IP address must be rewritten to match the firewall's IP address when exiting the VPN tunnel.

- **Action** – Select **Pass**.
- **Source** – Select a network object containing the on-premises networks. These networks must be configured as the remote network for the site-to-site VPN tunnels.
- **Service** – Select the services, or select **Any** to allow all.
- **Destination** – Select a network object containing the backend networks and/or IP addresses in AWS. These networks must be configured as the local network for the site-to-site VPN tunnels.
- **Connection Method** – **Translated IP from DHCP Interface**.

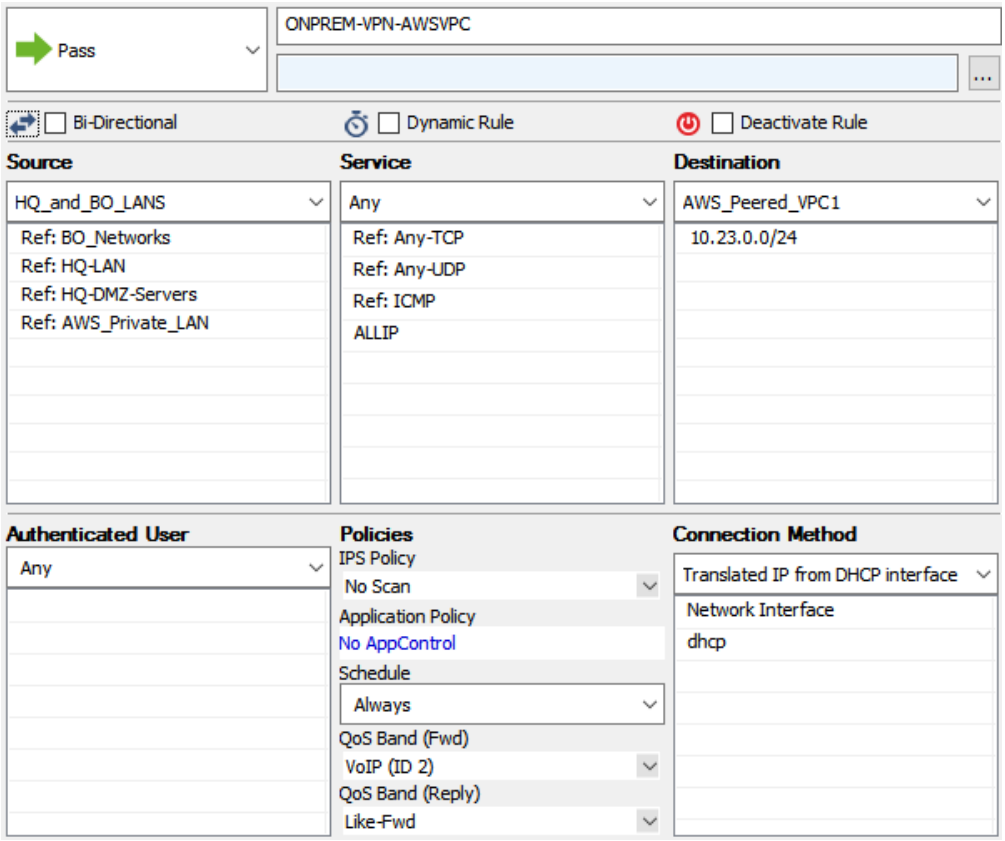

Site-to-Site VPN Tunnel to Backend Services using the Firewall as the Default Gateway

For EC2 instances using the firewall as the default gateway, the site-to-site VPN tunnels do not require source or destination NAT.

- **Action** – Select **Pass**.
- **Source** – Select a network object containing the on-premises networks. These networks must be configured as the remote network for the site-to-site VPN tunnels.
- **Service** – Select the services, or select **Any** to allow all.
- **Destination** – Select a network object containing the backend networks and/or IP addresses in AWS. These networks must be configured as the local network for the site-to-site VPN tunnels.
- **Bi-Directional** – Select to allow traffic in both directions.
- **Connection Method** – **Original Source IP**.

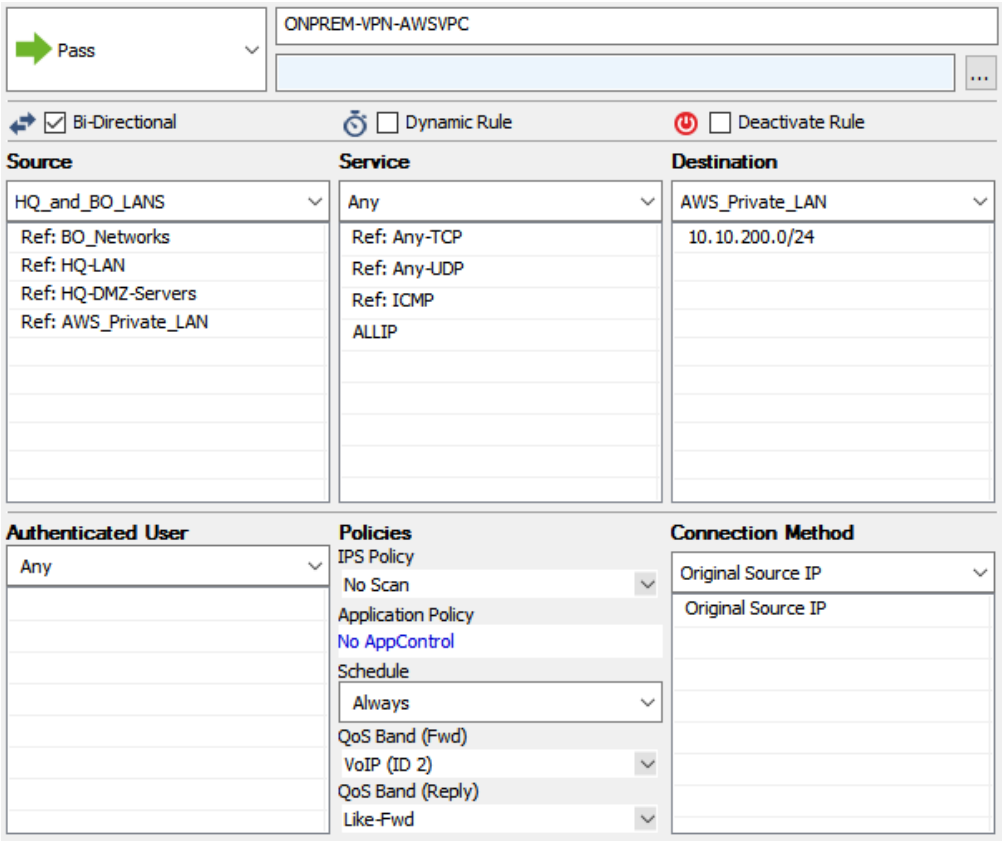

VPN Clients to Backend Services

Each VPN client is assigned an IP address in the VPN client network on the firewall the client is connected to.

- **Action** – Select **Pass**.
- **Source** – Select **Any**.
- **Service** – Select the services remote users are allowed to use, or select **Any** to allow all.
- **Destination** – Select a network object containing the backend services or networks remote users are allowed to access.
- **Connection Method** – **Translated IP from DHCP Interface** or **Original Source IP** depending on the destination.

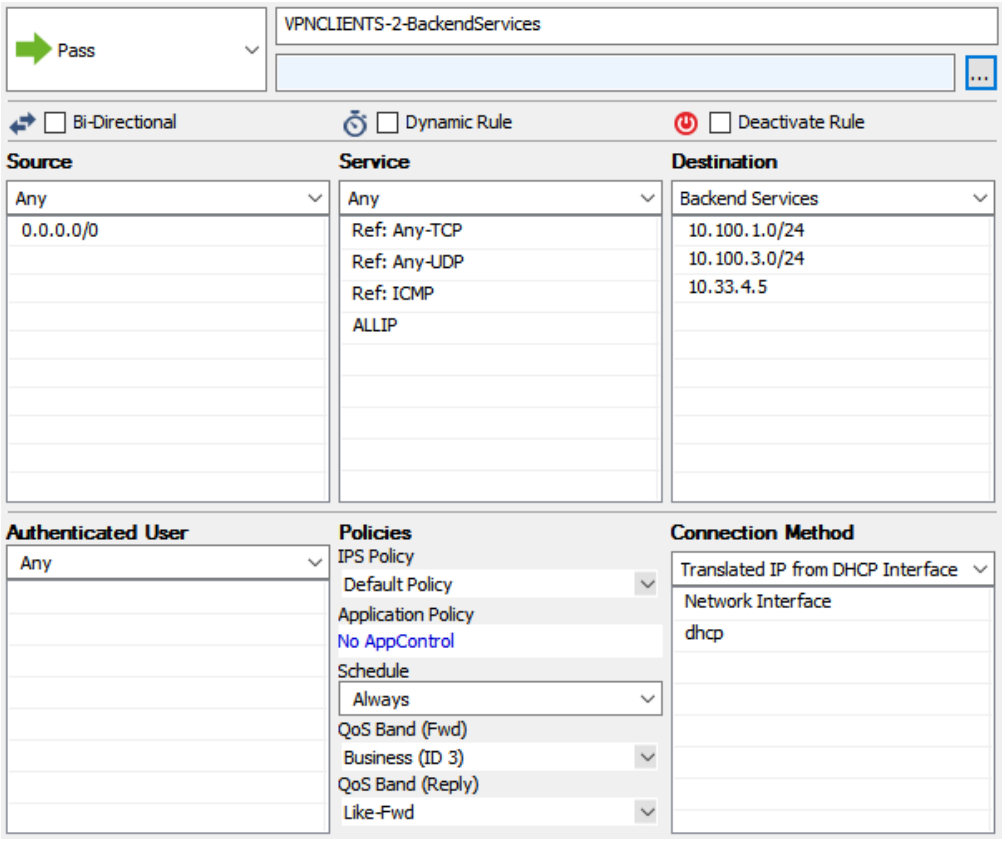

In the **TCP Policy** section of the **Advanced** access rule settings:

- **Syn Flood Protection (Fwd)** – Select **Outbound**.

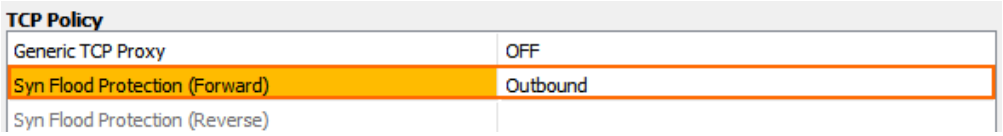

In the **Dynamic Interface Handling** section of the **Advanced** access rule settings:

- **Source Interface** – Select **VPN Clients**.
- **Continue on Source Interface Mismatch** – Select **Yes**.

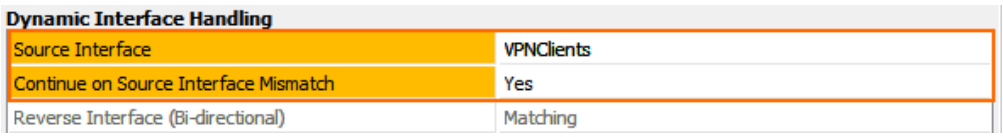

Internet to Backend Services not Using the Firewall as the Default Gateway

Create the following access rule to forward traffic from the Internet to an internal web server.

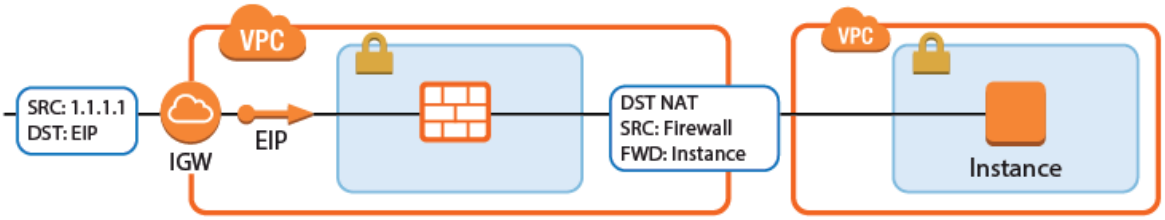

- **Action** – Select **Dst NAT**.
- **Source** – Select **Any** or a network object containing the networks the ELB is deployed in.
- **Service** – Select the service. E.g., **HTTP+S**.
- **Destination** – Select **DHCP1 Local IP**.
- **Connection Method** – Select **Dynamic NAT** or **Translated from DHCP Interface**.
- **Redirection Target** – Enter the IP address of the backend service. Optionally, append the port number to redirect to a different port. E.g, `10.100.1.2` or `10.100.1.2:8080`

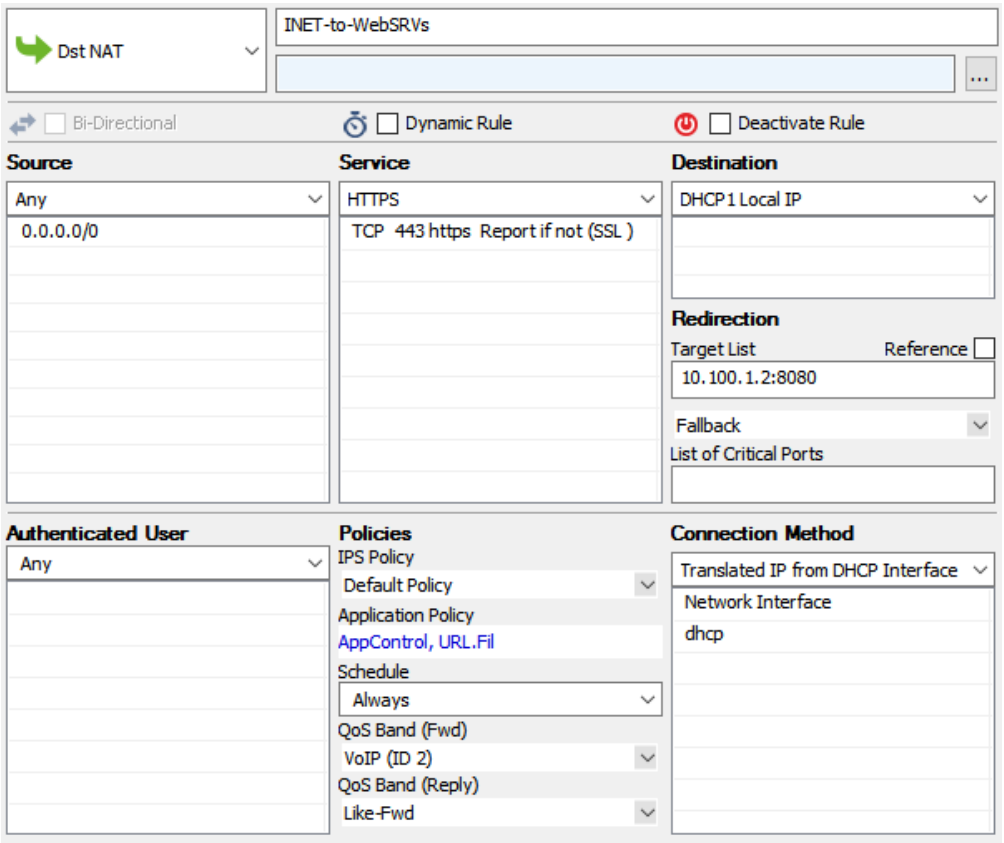

Figures

1. cold_standby_01.png
2. asg_default_pwd01.png
3. asg_default_pwd02.png
4. cold_standby_failover.png
5. awsIG_health_check_grace_period.png
6. awsIG_list_AMIs.png
7. aws_cold_standby_site2site.png
8. ips.png
9. awsIG_geolocation_network_object.png
10. awsIG_geolocation_access_rule.png
11. awsIG_access_rule_s2s_SNAT.png
12. awsIG_access_rule_s2s_noSNAT.png
13. awsIG_access_rule_vpnclients_01.png
14. awsIG_access_rule_vpnclients_02.png
15. awsIG_access_rule_vpnclients_03.png
16. aws_cold_standby_access_rule3.png
17. awsIG_dstnat_websrv01.png