Next Generation Firewall

6.4 or higher

How to deploy Forcepoint NGFW in the Azure cloud

Revision G

Contents

Introduction

You can deploy Forcepoint NGFW in the Microsoft Azure cloud to provide VPN connectivity, access control, and inspection for networks in the Azure cloud.

After deployment, you can manage NGFW Engines in the Azure cloud using the Management Client component of the Forcepoint NGFW Security Management Center (SMC) in the same way as other NGFW Engines.

> **Important**
>
> Networking in Microsoft Azure is significantly different compared to traditional networking. Deploying Forcepoint NGFW in the Microsoft Azure cloud requires familiarity with networking in Microsoft Azure. For more information about virtual networks in Microsoft Azure, see https://docs.microsoft.com/en-us/azure/virtual-network/.

For more information about using Azure, see the Microsoft Azure Documentation at https://docs.microsoft.com/en-us/azure/.

Licensing models for Forcepoint NGFW in the Azure cloud

Two licensing models are supported for Forcepoint NGFW in the Azure cloud.

- **Bring Your Own License** — You pay only Microsoft Azure's standard runtime fee for the NGFW Engine instance. You must manually install the license for the NGFW Engine in the Forcepoint NGFW Security Management Center (SMC) is automatically installed.
- **Hourly** (pay as you go license) — You pay Microsoft Azure's standard runtime fee for the NGFW Engine instance plus an hourly license fee based on the runtime of the NGFW Engine. No license installation is needed for the NGFW Engine in the SMC.

For features that require separate licenses, the SMC automatically detects which licensing model the NGFW Engine uses.

Provisioning methods for Forcepoint NGFW in the Azure cloud

Two provisioning methods are supported for Forcepoint NGFW in the Azure cloud.

- Automatic — The NGFW solution template automatically creates the NGFW Engine in the Forcepoint NGFW Security Management Center (SMC) using the SMC API when you deploy Forcepoint NGFW in the Azure cloud. You can use automatic deployment to create single NGFW Engines and Cloud Auto-Scaled Firewalls.

  To use automatic deployment, you must configure the SMC API in the Management Client before you deploy Forcepoint NGFW in the Azure cloud.
- Manual — You must configure the NGFW Engine elements in the SMC before you deploy Forcepoint NGFW in the Azure cloud. When you use manual deployment, you must upload the configuration file generated for the NGFW Engine in the SMC to the Azure cloud.

Limitations of Forcepoint NGFW in the Azure cloud

There are some limitations on features and configuration options when you deploy Forcepoint NGFW in the Microsoft Azure cloud.

> **CAUTION**
>
> To protect the privacy of your data, we recommend using an instance type that runs on dedicated hardware, such as an isolated instance.

- Only single NGFW Engines in the Firewall/VPN role are supported. Clustered NGFW Engines are not supported. Engines in the IPS and Layer 2 Firewall roles are not supported.
- Master NGFW Engines and Virtual Security Engines are not supported.
- The following types of interfaces are not supported: aggregated link interfaces, VLAN interfaces, wireless interfaces, ADSL interfaces.
- FIPS mode is not supported.
- Memory dump diagnostics are not supported.
- Forcepoint NGFW solution templates do not support password authentication for SSH connections to the NGFW Engine command line. You must use SSH keys for authentication of SSH connections to the NGFW Engine command line.

Interfaces and routing for Forcepoint NGFW in the Azure cloud

Interfaces and routing in the Azure cloud work differently than in physical networks. To understand how interfaces and networking work in the Azure cloud, we recommend that you familiarize yourself with the concept of Azure User Defined Routes.

For more information, see the following Microsoft Azure documentation:

- https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-udr-overview

- https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface-vm#constraints

When you deploy Forcepoint NGFW in the Azure cloud, the solution template automatically creates a security subnet, a front end subnet, and a back end subnet. The interface and routing configuration is similar to the configuration that is shown in this example: https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-dmz-nsg-fw-udr-asm

The NGFW Engine has one interface in the security subnet. The route table created by the solution template sends all traffic for the NGFW Engine to the interface in the security subnet. The NGFW Engine does not need to have interfaces in the front end subnet or the back end subnet to communicate with them. Multiple interfaces are not supported.

Configure the SMC

Before you deploy Forcepoint NGFW in the Microsoft Azure cloud, prepare the SMC for the deployment.

These tasks provide an overview of the SMC configuration process. For detailed instructions, see the following documents:

- *Forcepoint Next Generation Firewall Installation Guide*
- *Forcepoint Next Generation Firewall Product Guide*

Configure network connections and contact addresses for the SMC

Make sure that your Management Server and Log Server are reachable from your Forcepoint NGFW instance in Azure.

You must always configure network connections and contact addresses for the SMC, regardless of whether you use automatic deployment or manual deployment.

> **Note**
>
> For NGFW 6.4, the default contact addresses must be the external IP addresses of the Management Server and the Log Server. Make sure that the external IP addresses of the Management Server and the Log Server are reachable from the Internet.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a Location element for elements that contact the SMC servers using a different IP address than the default contact address.
   - a) Select **Configuration**, then browse to **Administration**.
   - b) Browse to **Other Elements** > **Locations**.
   - c) Right-click **Locations**, then select **New Location**.

d) Add the elements to the **Content** pane.

- For NGFW 6.4, add elements that contact the SMC servers using the internal IP addresses of the SMC servers.
- For NGFW 6.5 and higher, add NGFW Engine elements that contact the SMC servers using the external IP addresses of the SMC servers.

e) Click **OK**.

f) In the properties of the NGFW Engine elements or other elements that you added, select this Location.

2) Select 🏠 **Home**, then browse to **Others** > **Management Server** > **.**

3) Right-click the Management Server, then select **Properties.**

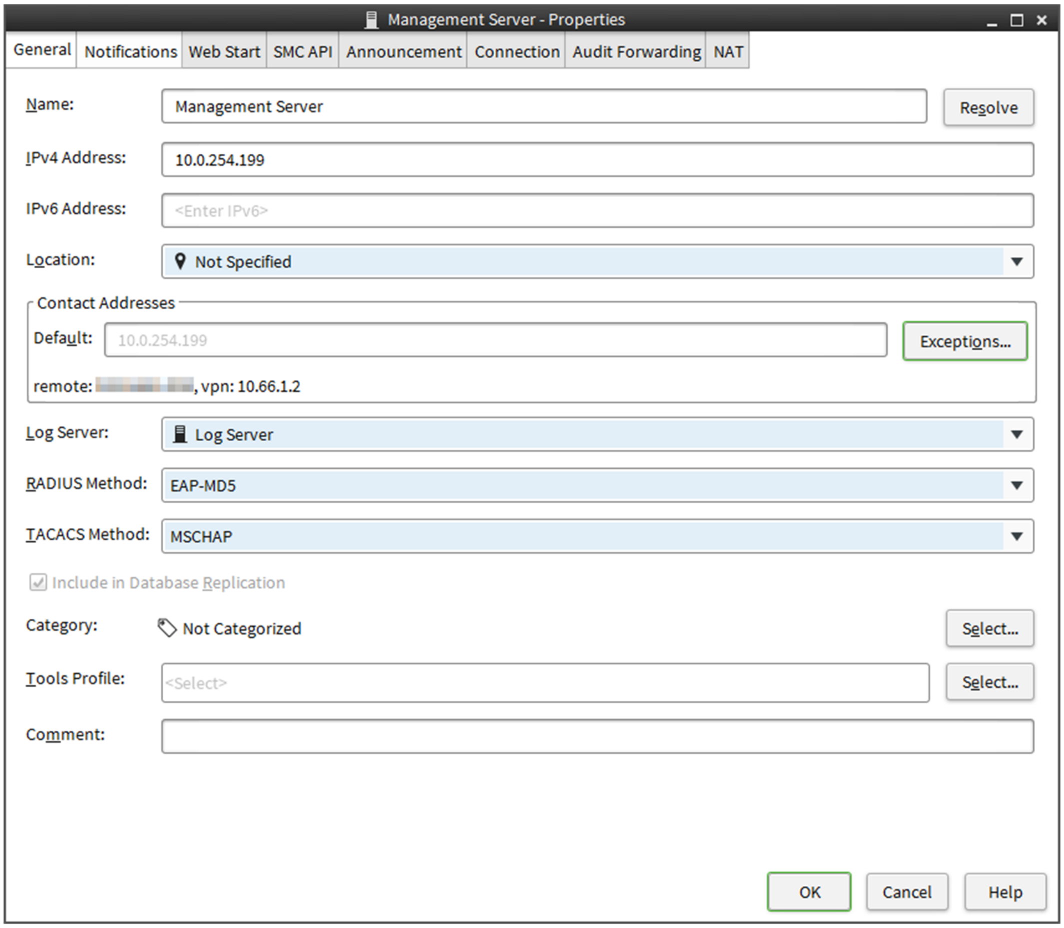

Name: Management Server

IPv4 Address: 10.0.254.199

IPv6 Address: <Enter IPv6>

Location: Not Specified

Contact Addresses

Default: 10.0.254.199

remote: [redacted], vpn: 10.66.1.2

Log Server: Log Server

RADIUS Method: EAP-MD5

TACACS Method: MSCHAP

☑ Include in Database Replication

Category: Not Categorized

Tools Profile: <Select>

Comment:

4) In the **Default** field in the **Contact Addresses** section, enter the default contact address for the Management Server.

- For NGFW 6.4, enter the external IP address of the Management Server.
- For NGFW 6.5 and higher, enter the internal IP address of the Management Server.

5) To configure contact address exceptions for the Management Server, Click **Exceptions**.

6) Click **Add**, select the Location element that you created, then click **Select**.

7) In the **Contact Addresses** cell, enter the IP address, then click **OK**.
   - For NGFW 6.4, enter the internal IP address of the Management Server.
   - For NGFW 6.5 and higher, enter the external IP address of the Management Server.

8) Click **OK** to close the **Management Server Properties** dialog box.

9) Browse to **Others** > **Log Server**, right-click the Log Server, then select **Properties**.

10) In the **Default** field in the **Contact Addresses** section, enter the default contact address for the Log Server.
   - For NGFW 6.4, enter the external IP address of the Log Server.
   - For NGFW 6.5 and higher, enter the internal IP address of the Log Server.

11) To configure contact address exceptions for the Log Server, Click **Exceptions**.

12) Click **Add**, select the Location element that you created, then click **Select**.

13) In the **Contact Addresses** cell, enter the IP address, then click **OK**.
   - For NGFW 6.4, enter the internal IP address of the Log Server.
   - For NGFW 6.5 and higher, enter the external IP address of the Log Server.

14) Click **OK** to close the **Log Server Properties** dialog box.

Create a self-signed certificate for the SMC API

If you plan to deploy Forcepoint NGFW using automatic deployment, create a self-signed certificate for the SMC API.

> **Note**
>
> The certificate for the SMC API must be self-signed. Do not use externally signed certificates for the SMC API.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Configuration**, then browse to **Administration**.

2) Browse to **Certificates** > **TLS Credentials**.

3) Right-click **TLS Credentials**, then select **New TLS Credentials**.

4) In the **Name** field, enter a unique name for the certificate.

5) In the **Common Name** field, enter the fully qualified domain name (FQDN) or IP address that the SMC API service uses.

6) Add the same FQDN or IP address that you entered in the **Common Name** field as the Subject Alternative Name.

   a) Click **Edit** next to the **Subject Alternative Name** field.

   b) Click **Add**, then select **DNS** from the drop-down list in the **Type** cell.

   c) Double-click the **Value** cell, then enter the same FQDN or IP address that you entered in the **Common Name** field.

7) From the signing options, select **Self-Sign**, then click **Finish**.

8) Right-click the certificate element, then select **Properties**.

9) On the **Certificate** tab, click **Export**, then save the certificate file.

Enable the SMC API

If you plan to deploy Forcepoint NGFW using automatic deployment, enable the SMC API and create an SMC API Client account.

These tasks provide an overview of the SMC API configuration process. For more information, see the *Forcepoint NGFW SMC API Reference Guide* ↗.

Steps

❓ For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select 🏠 **Home**, then browse to **Others** > **Management Server** > **.**

2) Right-click the Management Server, then select **Properties.**

3) On the **SMC API** tab, select **Enable**.

4) In the **Host Name** field, enter the same FQDN or IP address that is used in the certificate for the SMC API.

5) (Optional) To use a port other than the default port, enter the port number in the **Port Number** field.

   The default port number is 8082.

   > **Note**
   >
   > If you enter a different port number here, you must also specify the port number when you configure the **NGFW provisioning** settings in the Azure portal.

6) (Optional) If the Management Server has several addresses and you want to restrict access to one address, specify the IP address to use in the **Listen Only on Address** field.

7) To enable the use of HTTPS for SMC API connections, click **Select** next to the **Server Credentials**, then select the TLS Credentials element that contains the certificate for the SMC API.

8) Create an SMC API Client.

   a) Select ⚙ **Configuration**, then browse to **Administration**.

   b) Browse to **Access Rights**.

   c) Right-click **Access Rights**, then select **New** > **API Client**.

   d) In the **Name** field, enter a name for the API Client element.

   e) Use the automatically generated authentication key or click **Generate Authentication Key** to generate a new one.

   > ⚠ **Important**
   >
   > This key appears only once, so make sure to record the value in the **Authentication Key** field. You must enter this key when you configure the **NGFW provisioning** settings in the Azure portal.

   f) On the **Permissions** tab, select **Unrestricted Permissions (Superuser)**.

Deploy a single Forcepoint NGFW using automatic deployment

When you use automatic deployment, NGFW Engine elements are automatically created in the SMC.

Before you begin

Configure the SMC and enable the SMC API. The SMC API must be reachable from the Internet.

To use automatic deployment, Forcepoint NGFW 6.4 or higher is required.

The Forcepoint single NGFW solution template includes the NGFW Engine software and the network environment in which it runs. The network environment includes the Security subnet in which the NGFW Engine is deployed, and two protected subnets. The template creates a route from the protected subnets to the Internet through the NGFW Engine. A route is also created between the two protected subnets.

Network environment for single Forcepoint NGFW deployment

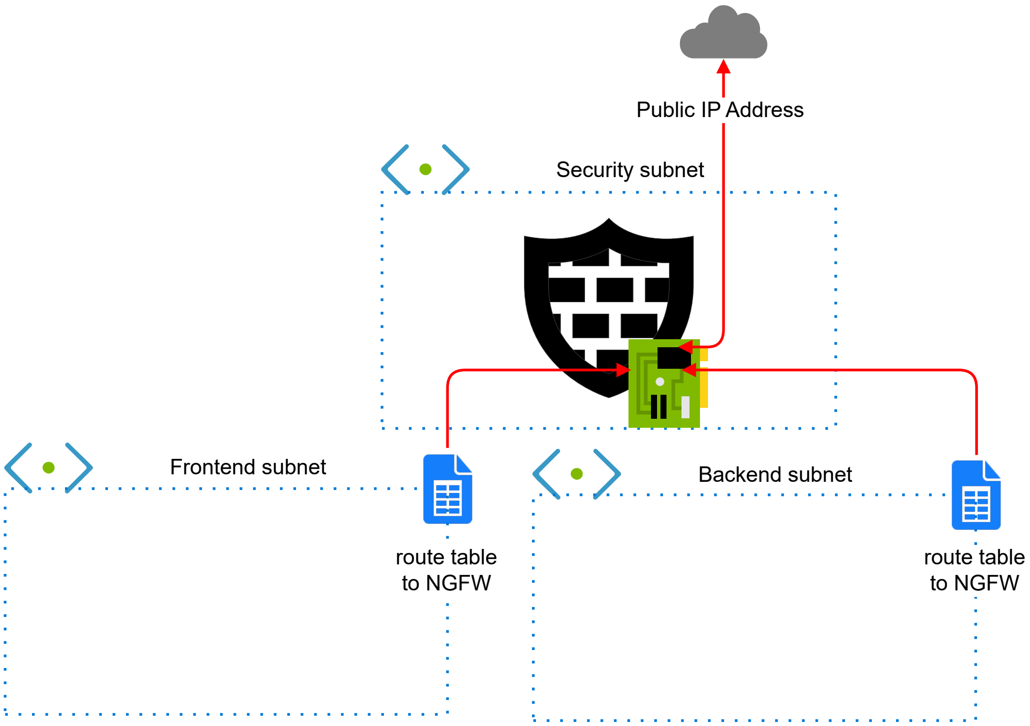

Start the automatic deployment

To start the automatic deployment, create a Forcepoint NGFW instance using an NGFW solution template.

Steps

1) In the Azure dashboard, click **Create a resource.**

2) In the Azure Marketplace pane, search for the NGFW solution template.

   Tip

   Enter `forcepoint LLC` as a search string to find all solution templates published by Forcepoint.

3) Select the Forcepoint single NGFW solution template, then click **Create**.

Configure basic settings for the Forcepoint NGFW instance

Configure settings for how to connect to the Forcepoint NGFW instance, and where to deploy the instance.

Steps

1) In the **NGFW Admin Username** field, enter a user name for the administrator account that is used to make SSH connections to the NGFW Engine.

2) In the **Admin SSH Public Key** field, enter or paste the public key that is used to make SSH connections to the NGFW Engine.

3) From the **Subscription** drop-down list, select the account that is charged for resources.

4) In the **Resource group** options, select **Create new**, then enter a name for the resource group.

   > **Note**
   >
   > You must always create a new resource group for the deployment. It is not possible to use an existing resource group.

5) (Optional) From the **Location** drop-down list, select the region where you want to deploy the NGFW Engine.

   Regions are physical Microsoft Azure data centers. For more information, see https://azure.microsoft.com/en-us/overview/datacenters/how-to-choose/.

   > **Note**
   >
   > When you use an hourly (pay as you go) license, only Microsoft tax remitted countries and regions are supported. For more information, see Knowledge Base article 16011.

6) Click **OK**.

Configure NGFW settings

Configure settings for the virtual machine where the Forcepoint NGFW instance runs.

Steps

1) From the **NGFW licensing model** options, select the licensing model for the NGFW Engine.

2) (Optional) From the **NGFW Version** drop-down list, select the NGFW Engine version.

   Version 6.4 or higher is required.

3) If the default value of the **Virtual network** option does not meet your needs, select a different value.

   You must view and accept the virtual network settings even if you do not change the settings.

4) If the default values of the **NGFW Security Subnet**, **Protected FrontEnd Subnet**, and **Protected BackEnd Subnet** options do not meet your needs, change the settings, then click **OK**.

5) (Optional) If the default value of the **NGFW VM Size** option does not meet your needs, select a different value.

   We recommend selecting a general purpose VM size that has a SKU that starts with the letter D and at least 4 GB of RAM.

6) In the **Resource prefix** field, enter an identifying prefix that is automatically added to the name of the resource.

   The prefix is also added to the name of the automatically created NGFW Engine element in the SMC.

7) From the **VM Zone** options, select the zone to which the NGFW Engine belongs.

8) From the **Modify existing vnet to redirect traffic to NGFW** option, select whether to automatically redirect traffic to the NGFW Engine.

   > **Note**
   >
   > The **Modify existing vnet to redirect traffic to NGFW** options are only available if you selected an existing virtual network as the value of the **Virtual network** option. If you created a new virtual network for the NGFW deployment, the new virtual network is automatically configured.

   - **Yes** — Traffic to and from the protect subnets is immediately redirected to the NGFW Engine. Route tables are automatically attached to the virtual networks that are selected for the **Protected FrontEnd Subnet** and **Protected BackEnd Subnet** options. For Cloud Auto-Scaled Firewalls, the mandatory Azure NSG is deployed in the virtual network that you selected for the **NGFW Security Subnet** option.

     > **Note**
     >
     > If additional subnets need to be redirected to the NGFW Engine, you must associate the route table manually with those additional subnets.

   - **No** — You must associate route tables with the virtual networks that are selected for the **Protected FrontEnd Subnet** and **Protected BackEnd Subnet** options to route traffic through the NGFW Engine. For Cloud Auto-Scaled Firewalls, you must manually attach the mandatory Azure NSG to the virtual network that you selected for the **NGFW Security Subnet** option.

9) Click **OK**.

Configure NGFW provisioning and finish the automatic deployment

The NGFW provisioning settings enable the connection between the Forcepoint NGFW instance and the SMC API, and define settings that are applied after the instance starts.

Steps

1) From the **NGFW Deployment model** drop-down list, select **Automatic via SMC REST API**.

2) In the **SMC Contact address (FQDN or IP address)** field, enter the fully qualified domain name (FQDN) or the public IP address of the SMC API.

   Make sure that the information that you enter here matches the common name or subject alternative name in the certificate for the SMC API.

   > **Tip**
   >
   > You can find the FQDN or IP address of the SMC API in the **Host Name** field on the **SMC API** tab of the **Management Server Properties** dialog box in the Management Client.

3) (Optional) If the SMC API uses a port other than the default port, enter the SMC API port number in the **SMC rest API port** field.

   The default port number is 8082.

4) In the **SMC rest API key** field, enter the authentication key of the SMC API Client.

5) Make sure that **Yes** is for **Check REST API TLS certificate options**.

   When **Yes** is selected, the TLS certificate of the SMC API is validated when NGFW Engine elements are automatically created.

   > **Note**
   >
   > The **No** option is intended only for testing purposes. We do not recommend selecting **No** in a production environment.

6) Next to the **Upload SMC rest API certificate** field, click the file browser icon, then select the certificate file.

   > **Tip**
   >
   > To find the certificate in the Management Client, select **Configuration**, then browse to **Administration** > **Certificates** > **TLS Credentials**.

7) (NGFW 6.5 and higher) In the **Engine Location** field, enter the name of the Location element that is selected for the NGFW Engine when the NGFW Engine element is created.

   The Location element must already exist before you deploy the NGFW Engine. The name must match the name of the Location element in the SMC.

   > **Note**
   >
   > Make sure that you have defined contact address exceptions for this location in the properties of the Management Server and the Log Server.

8) (Recommended) In the **Engine policy name** field, enter the name of the Firewall Policy that is uploaded to the NGFW Engine after the NGFW Engine element is created.

   The Firewall Policy must already exist before you deploy the NGFW Engine. The name must match the name of the Firewall Policy element in the SMC.

   > **Note**
   >
   > If you do not specify a Firewall Policy, you must manually install a policy using the Management Client after deploying the NGFW Engine.

9) (Optional) From the **Engine Auto delete when shutting off** options, select **No** if you want the NGFW Engine element to stay in the SMC when the NGFW Engine instance shuts down.

   When **Yes** is selected, the NGFW Engine elements are automatically deleted when the NGFW Engine instances shut down or are restarted in Azure.

10) Click **OK**.

    The deployment continues to a summary and the configuration is validated.

11) When the validation is finished, click **OK**.

12) Review the terms of use, then click **Create**.

Result

The NGFW Engine deployment starts and an NGFW Engine element is automatically created in the SMC. When deployment is finished, you can check the status using the Management Client.

Deploying Cloud Auto-Scaled Firewalls

You can create NGFW Engines that can be scaled manually, on a schedule, or automatically depending on traffic load.

Using scaling features is an advanced task. You must know how scaling works and be aware of the related Azure guidelines. For Forcepoint NGFW, both vertical and horizontal scaling is supported. For more information about scaling in Azure, see https://docs.microsoft.com/en-us/azure/architecture/best-practices/auto-scaling.

The scale set NGFW solution template includes the NGFW Engine software and the network environment in which it runs. The network environment includes the Security subnet in which the NGFW Engines are deployed, and two protected subnets. The template creates a route from the protected subnets to the Internet through the NGFW Engines. A route is also created between the two protected subnets.

Network environment for Cloud Auto-Scaled Firewall deployment

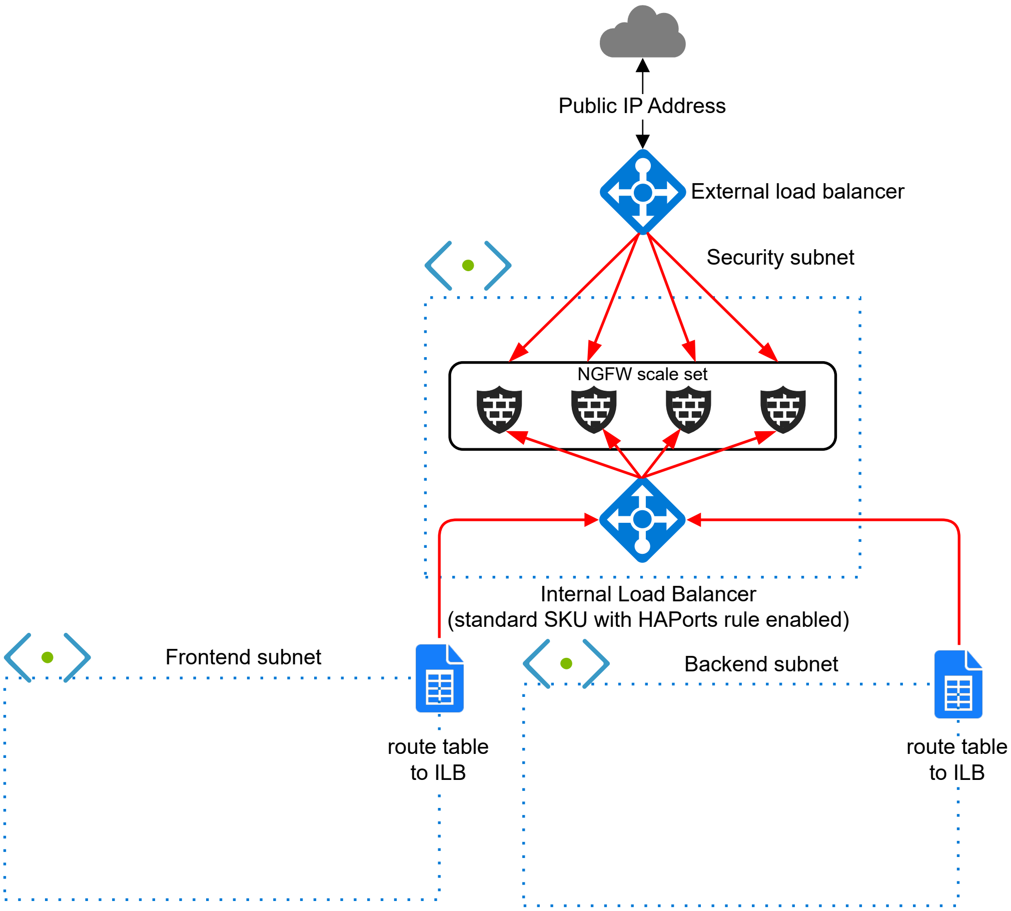

Cloud Auto-Scaled firewalls have the following limitations:

- Cloud Auto-Scaled Firewalls can only be created for Forcepoint NGFW 6.4 or higher.
- The hourly (pay as you go) licensing model is recommended for Cloud Auto-Scaled Firewalls.
- The SMC API is required for Cloud Auto-Scaled Firewalls.
- Because you cannot modify the properties of Cloud Auto-Scaled Firewalls in the SMC, features that require changing the properties of the NGFW Engine elements are not supported.

Monitoring Cloud Auto-Scaled Firewalls in the Management Client

Cloud Auto-Scaled Firewalls are automatically added to Cloud Auto-Scaled Firewall Groups, and the groups are shown in the Home view.

You can preview the Cloud Auto-Scaled Firewalls in the Engine Editor, but you cannot make changes to the configuration.

You can make changes to the installed policy and refresh the policy on the Cloud Auto-Scaled Firewalls, but you cannot install a different policy.

Start the deployment of Cloud Auto-Scaled Firewalls

To deploy a Cloud Auto-Scaled Firewall, use automatic deployment to create a Forcepoint NGFW instance using the scale set NGFW template.

Before you begin

To deploy Cloud Auto-Scaled Firewalls, Forcepoint NGFW 6.4 or higher is required.

Steps

1) In the Azure dashboard, click **Create a resource.**

2) In the Azure Marketplace pane, search for the NGFW solution template.

   Tip

   Enter `forcepoint LLC` as a search string to find all solution templates published by Forcepoint.

3) Select the scale set Forcepoint NGFW solution template, then click **Create**.

Configure basic settings for the Cloud Auto-Scaled Firewalls

Configure settings for how to connect to the Cloud Auto-Scaled Firewalls, and where to deploy the instances.

Steps

1) In the **NGFW Admin Username** field, enter a user name for the administrator account that is used to make SSH connections to the NGFW Engine.

2) In the **Admin SSH Public Key** field, enter or paste the public key that is used to make SSH connections to the NGFW Engine.

3) From the **Subscription** drop-down list, select the account that is charged for resources.

4) In the **Resource group** options, select **Create new**, then enter a name for the resource group.

   Note

   You must always create a new resource group for the deployment. It is not possible to use an existing resource group.

5) (Optional) From the **Location** drop-down list, select the region where you want to deploy the NGFW Engine.

   Regions are physical Microsoft Azure data centers. For more information, see https://azure.microsoft.com/en-us/overview/datacenters/how-to-choose/.

   > **Note**
   >
   > When you use an hourly (pay as you go) license, only Microsoft tax remitted countries and regions are supported. For more information, see Knowledge Base article 16011.

6) Click **OK**.

Configure NGFW settings for the Cloud Auto-Scaled Firewalls

Configure settings for the virtual machines where the Cloud Auto-Scaled Firewalls run.

Steps

1) (Optional) From the **NGFW Version** drop-down list, select the NGFW Engine version.

   Version 6.4 or higher is required.

2) From the **NGFW licensing model** options, select **Pay as you go**.

   > **Note**
   >
   > The hourly (pay as you go) licensing model is recommended for Cloud Auto-Scaled Firewalls. If you use the bring your own license licensing model, you must have an NGFW license in the SMC for each NGFW Engine that is automatically created.

3) If the default value of the **Virtual network** option does not meet your needs, select a different value.

   You must view and accept the virtual network settings even if you do not change the settings.

4) If the default values of the **NGFW Security Subnet**, **Protected FrontEnd Subnet**, and **Protected BackEnd Subnet** options do not meet your needs, change the settings, then click **OK**.

5) In the **Resource prefix** field, enter an identifying prefix that is automatically added to the name of the resource.

   The prefix is also added to the name of the automatically created NGFW Engine elements in the SMC.

6) In the **Initial instance count in the Autoscale group** field, enter the number of NGFW Engine elements that are automatically created when the Cloud Auto-Scaled Firewalls start.

7) Click **OK**.

Configure NGFW provisioning and finish deploying Cloud Auto-Scaled Firewalls

The NGFW provisioning settings enable the connection between the Cloud Auto-Scaled Firewall instances and the SMC API, and define settings that are applied after each instance starts.

Steps

1) In the **SMC Contact address (FQDN or IP address)** field, enter the fully qualified domain name (FQDN) or the public IP address of the SMC API.

   Make sure that the information that you enter here matches the common name or subject alternative name in the certificate for the SMC API.

   > **Tip**
   >
   > You can find the FQDN or IP address of the SMC API in the **Host Name** field on the **SMC API** tab of the **Management Server Properties** dialog box in the Management Client.

2) (Optional) If the SMC API uses a port other than the default port, enter the SMC API port number in the **SMC rest API port** field.

   The default port number is 8082.

3) In the **SMC rest API key** field, enter the authentication key of the SMC API Client.

4) Make sure that **Yes** is for **Check REST API TLS certificate options**.

   When **Yes** is selected, the TLS certificate of the SMC API is validated when NGFW Engine elements are automatically created.

   > **Note**
   >
   > The **No** option is intended only for testing purposes. We do not recommend selecting **No** in a production environment.

5) Next to the **Upload SMC rest API certificate** field, click the file browser icon, then select the certificate file.

   > **Tip**
   >
   > To find the certificate in the Management Client, select **Configuration**, then browse to **Administration** > **Certificates** > **TLS Credentials**.

6) (NGFW 6.5 and higher) In the **Engine Location** field, enter the name of the Location element that is selected for the NGFW Engine when the NGFW Engine element is created.

   The Location element must already exist before you deploy the NGFW Engine. The name must match the name of the Location element in the SMC.

   > **Note**
   >
   > Make sure that you have defined contact address exceptions for this location in the properties of the Management Server and the Log Server.

7) (Recommended) In the **Engine policy name** field, enter the name of the Firewall Policy that is uploaded to the NGFW Engine after the NGFW Engine element is created.

The Firewall Policy must already exist before you deploy the NGFW Engine. The name must match the name of the Firewall Policy element in the SMC.

> **Note**
>
> If you do not specify a Firewall Policy, you must manually install a policy using the Management Client after deploying the NGFW Engine.

8) (Optional) From the **Engine Auto delete when shutting off** options, select **No** if you want the NGFW Engine elements to stay in the SMC when the NGFW Engine instances shut down or are restarted in Azure.

When **Yes** is selected, the NGFW Engine elements are automatically deleted when the NGFW Engine instances shut down or are restarted in Azure.

If you select **No**, you must manually remove unused Cloud Auto-Scaled Firewalls in the Management Client.

9) (Optional) If the default value of the **NGFW VM Size** option does not meet your needs, select a different value.

We recommend selecting a general purpose VM size that has a SKU that starts with the letter D and at least 4 GB of RAM.

10) Click **OK**.

The deployment continues to a summary and the configuration is validated.

11) When the validation is finished, click **OK**.

12) Review the terms of use, then click **Create**.

13) Add one or more load balancing rules and configure scaling for the Cloud Auto-Scaled Firewalls.

For instructions, see the Microsoft Azure documentation at https://docs.microsoft.com/en-us/azure/.

Result

The NGFW Engine deployment starts and NGFW Engine elements are automatically created in the SMC. When deployment is finished, you can check the status using the Management Client. You can preview the NGFW Engine properties in the Engine Editor, but you cannot make changes to the configuration.

Add NAT rules for Cloud Auto-Scaled Firewalls

To prevent asymmetric routing, add NAT rules in the Management Client.

Steps

1) Select ⚙ **Configuration**.

2) Browse to **Policies** > **Firewall Policies**, then open your Firewall Policy for editing.

3) On the **IPv4 NAT** tab, add the a rule, then define the source, destination, and service:
   - **Source** — ANY
   - **Destination** — $$ DHCP Interface 1.ip Alias element
   - **Service** — Select the service according to the type of traffic that the NGFW Engine handles.

4) To define source and destination translation, double-click the **NAT** cell.

5) On the **Source Translation** tab, configure source NAT.
   a) From the **Translation Type** drop-down menu, select **Dynamic**.

   b) Next to the **IP Address Pool** field, click **Select**.

   c) Browse to the $$ DHCP Interface 1.ip Alias element, then click **Select**.

   d) Deselect **Automatic Proxy ARP**.

6) On the **Destination Translation** tab, configure destination NAT.
   a) Select **Translate Destination**.

   b) Next to the **Translated** field, click **IP Address**, then enter the destination IP address in the protected network.
   For example, if the destination is a web server in the protected network, enter the private IP address of the web server.

   c) Deselect **Automatic Proxy ARP**.

7) Click **OK**.

8) Click **Save and Install**.

Remove unused Cloud Auto-Scaled Firewalls

If you do not set the template solution parameter to automatically remove Cloud Auto-Scaled Firewall instances, you must periodically remove them in the Management Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Configuration**.

2) Browse to **Other Elements** > **Cloud Auto-Scaled Groups**.

3) Right-click a Cloud Auto-Scaled Group, then select **Tools** > **Remove Unused Cloud Auto-Scaled Firewalls**.

4) Click **OK**.

Deploy Forcepoint NGFW using manual deployment

When you use manual deployment, you must create NGFW Engine elements in the SMC before you deploy the NGFW Engine in the Azure cloud environment.

Create the NGFW Engine in the Management Client

If you are deploying using the single NGFW, add and configure a placeholder Single Firewall element for each NGFW Engine that you deploy in the Azure cloud.

These steps provide an overview of the NGFW Engine configuration process. For detailed instructions, see the following documents:

- *Forcepoint Next Generation Firewall Installation Guide*
- *Forcepoint Next Generation Firewall Product Guide*

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add a Single Firewall element.

2) Browse to the **General** branch of the Engine Editor, then select the Location element for elements outside of the local network of the SMC servers from the **Location** drop-down list.

3) Browse to **Interfaces**, then add a layer 3 physical interface and a dynamic IP address.

   a) Add a layer 3 physical interface.

   b) Add an IPv4 address to the interface.

   c) From the IP address type drop-down list, select **Dynamic**.

   d) From the **Dynamic Index** drop-down list, select **First DHCP Interface**.

   e) Select **Automatic Default Route**.

4) Browse to **Interfaces** > **Loopback**, then add the following loopback IP address: 127.0.0.1.

5) Browse to **Interfaces** > **Interface Options**, then make the following selections:

   a) Select Interface ID 0 as the primary control interface.

      The **Node-Initiated Contact to Management Server** option is automatically selected when the control IP address is dynamic. When the option is selected, the NGFW Engine opens a connection to the Management Server and maintains connectivity.

   b) Select the loopback IP address as the identify for authentication requests.

6) Browse to **Routing**, then add a default route through Interface 0.

   a) Right-click the network under Interface 0, then select **Add Router**.

   b) Right-click the Router element, then select **Add**.

   c) Browse to **Networks** > **Any Network**, click **Add**, then click **OK**.

7) Click **Save** to save and validate changes, then close the Engine Editor.

8) (Bring your own license only) Install a license, then bind the license to the Single Firewall element.

9) Save the initial configuration.

   a) Right-click the NGFW Engine, then select **Configuration** > **Save initial Configuration**.

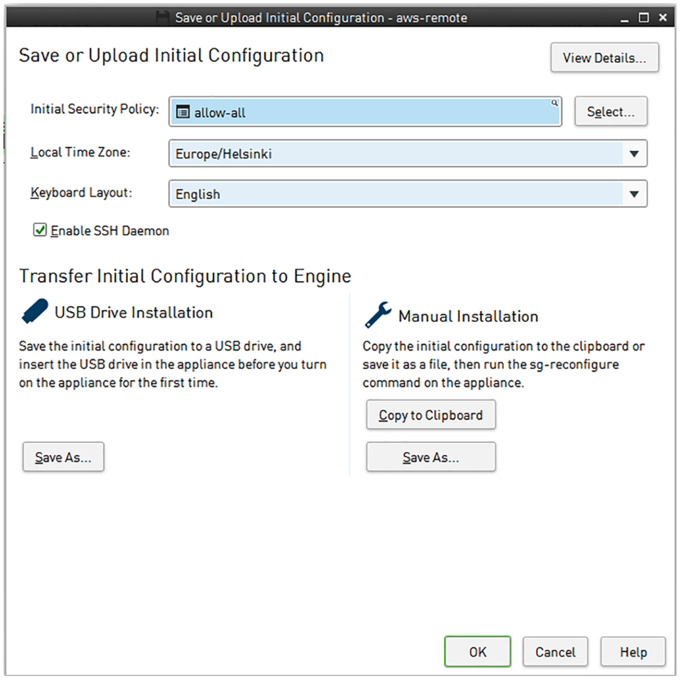

Save or Upload Initial Configuration - aws-remote

Save or Upload Initial Configuration

View Details...

Initial Security Policy: allow-all

Select...

Local Time Zone: Europe/Helsinki

Keyboard Layout: English

☑ Enable SSH Daemon

Transfer Initial Configuration to Engine

USB Drive Installation

Save the initial configuration to a USB drive, and insert the USB drive in the appliance before you turn on the appliance for the first time.

Save As...

Manual Installation

Copy the initial configuration to the clipboard or save it as a file, then run the sg-reconfigure command on the appliance.

Copy to Clipboard

Save As...

OK

Cancel

Help

   b) Next to the **Initial Security Policy** field, click **Select**, then select a policy for the NGFW Engine.

   c) Select **Enable SSH Daemon**.

d) To save the initial configuration file, click **Save As**, then select the location where you want to save the file.

Deploy the NGFW Engine in the Azure cloud environment

The Forcepoint single NGFW solution template includes the NGFW Engine software and the network environment in which it runs.

Before you begin

Configure the SMC and add an NGFW Engine element in the Management Client.

The network environment includes the Security subnet in which the NGFW Engine is deployed, and two protected subnets. The template creates a route from the protected subnets to the Internet through the NGFW Engine. A route is also created between the two protected subnets.

Network environment for single Forcepoint NGFW deployment

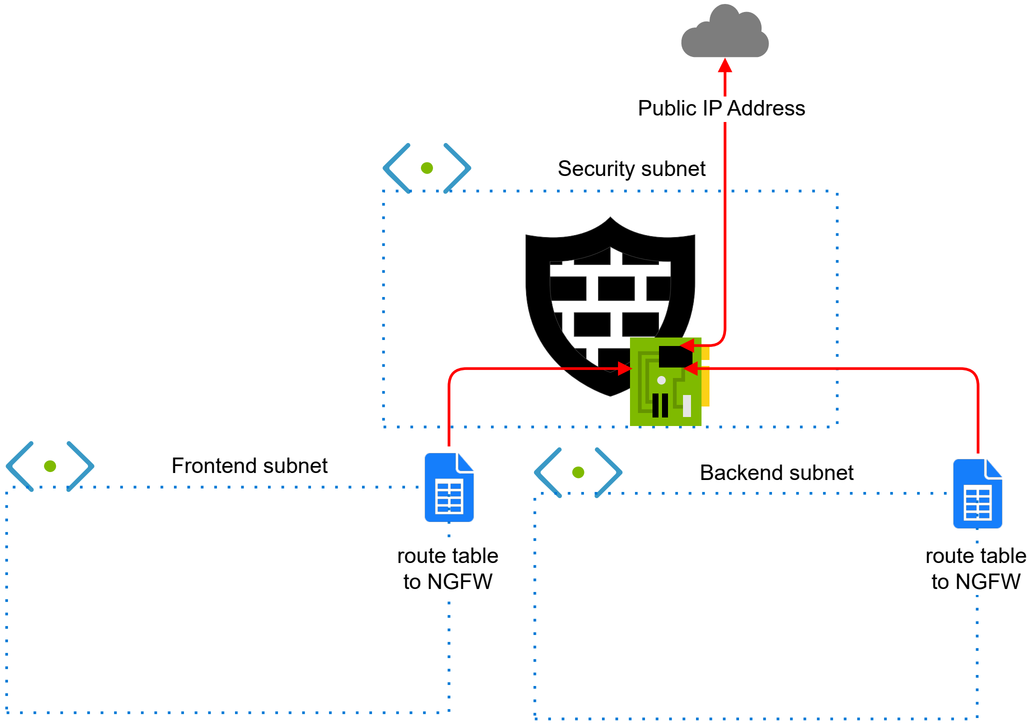

Steps

1) Create a Forcepoint NGFW instance using the Forcepoint single NGFW solution template.

   a) In the Azure dashboard, click **Create a resource.**

b) In the Azure Marketplace pane, search for the NGFW solution template.

> **Tip**
>
> Enter `forcepoint LLC` as a search string to find all solution templates published by Forcepoint.

c) Select the Forcepoint single NGFW solution template, then click **Create**.

2) Configure the options in the **Basics** section.

a) In the **NGFW Admin Username** field, enter a user name for the administrator account that is used to make SSH connections to the NGFW Engine.

b) In the **Admin SSH Public Key** field, enter or paste the public key that is used to make SSH connections to the NGFW Engine.

c) From the **Subscription** drop-down list, select the account that is charged for resources.

d) In the **Resource group** options, select **Create new**, then enter a name for the resource group.

> **Note**
>
> You must always create a new resource group for the deployment. It is not possible to use an existing resource group.

e) (Optional) From the **Location** drop-down list, select the region where you want to deploy the NGFW Engine.

Regions are physical Microsoft Azure data centers. For more information, see https://azure.microsoft.com/en-us/overview/datacenters/how-to-choose/.

> **Note**
>
> When you use an hourly (pay as you go) license, only Microsoft tax remitted countries and regions are supported. For more information, see Knowledge Base article 16011.

f) Click **OK**.

3) Configure the options in the **NGFW configuration** section.

a) From the **NGFW licensing model** options, select the licensing model for the NGFW Engine.

b) (Optional) From the **NGFW Version** drop-down list, select the NGFW Engine version.

Version 6.4 or higher is required.

c) If the default value of the **Virtual network** option does not meet your needs, select a different value.

d) If the default values of the **NGFW Security Subnet**, **Protected FrontEnd Subnet**, and **Protected BackEnd Subnet** options do not meet your needs, change the settings, then click **OK**.

e) (Optional) If the default value of the **NGFW VM Size** option does not meet your needs, select a different value.

   We recommend selecting a general purpose VM size that has a SKU that starts with the letter D and at least 4 GB of RAM.

f) From the **VM Zone** options, select the zone to which the NGFW Engine belongs.

g) From the **Modify existing vnet to redirect traffic to NGFW** option, select whether to automatically redirect traffic to the NGFW Engine.

   > **Note**
   >
   > The **Modify existing vnet to redirect traffic to NGFW** options are only available if you selected an existing virtual network as the value of the **Virtual network** option. If you created a new virtual network for the NGFW deployment, the new virtual network is automatically configured.

   - **Yes** — Traffic to and from the protect subnets is immediately redirected to the NGFW Engine. Route tables are automatically attached to the virtual networks that are selected for the **Protected FrontEnd Subnet** and **Protected BackEnd Subnet** options. For Cloud Auto-Scaled Firewalls, the mandatory Azure NSG is deployed in the virtual network that you selected for the **NGFW Security Subnet** option.
   - **No** — You must associate route tables with the virtual networks that are selected for the **Protected FrontEnd Subnet** and **Protected BackEnd Subnet** options to route traffic through the NGFW Engine. For Cloud Auto-Scaled Firewalls, you must manually attach the mandatory Azure NSG to the virtual network that you selected for the **NGFW Security Subnet** option.

h) Click **OK**.

4) Configure the options in the **NGFW provisioning** section.

   a) From the **NGFW Deployment model** drop-down list, select **Manual via engine.cfg**.

   b) Next to the **Upload initial contact file engine.cfg** field, click the file browser icon, then select the engine.cfg file that contains the initial configuration for the NGFW Engine.

5) Click **OK**.

   The deployment continues to a summary and the configuration is validated.

6) When the validation is finished, click **OK**.

7) Review the terms of use, then click **Create**.

Result

The NGFW Engine deployment starts. When deployment is finished, you can check the status and manage the NGFW Engine using the Management Client.

Configure a VPN with an NGFW Engine in Azure

When you have deployed an NGFW Engine in Azure, you can use it as an endpoint in VPNs with other NGFW Engines in your network.

Note

You cannot use Cloud Auto-Scaled Firewalls in VPNs.

Configuring a VPN between NGFW Engines that are managed by the same SMC has the following advantages compared to using Azure's native VPN tools:

- Access control for VPN traffic
- Centralized management of the NGFW Engines that act as VPN gateways

Because the public IP addresses of NGFW Engines deployed in Azure are dynamic, the following restrictions apply when you use an NGFW Engine deployed in Azure as a VPN gateway:

- The VPN gateway must use the fully qualified domain name (FQDN) of your NGFW Engine as the phase-1 ID.
- IKEv1 main mode with pre-shared key authentication is not supported. Aggressive mode allows the use of pre-shared keys, but for security reasons certificate-based authentication is also recommended when IKEv1 is set in aggressive mode.

Configure the NGFW Engine deployed in Azure as a VPN gateway

Configure settings for the NGFW Engine deployed in Azure that allow you to use it as a VPN gateway.

Steps ❓ For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Azure portal, select your virtual machine, then select **Overview** to find the FQDN of your NGFW Engine.

2) (NGFW 6.4 only) In the Management Client, add the FQDN of the NGFW Engine to the dynamic IP address under interface 0.

   For NGFW 6.5 or higher, the FQDN is entered automatically.

   a) Right-click the NGFW Engine, then select **Edit <element type>**.

   b) In the navigation pane on the left, select **Interfaces**.

   c) Right-click the IP address, then select **Edit IP Address.**

   d) In the **Contact Addresses** options, enter the FQDN of your NGFW Engine in the **Default** field.

e) Click **OK**.

3) Configure the phase-1 ID of the VPN endpoint.

   a) In the navigation pane on the left, select **VPN** > **End-Points**.

   b) Right-click the internal endpoint, then select **Properties**.

   c) In the **Phase-1 ID** settings, select **DNS Name** from the **ID Type** drop-down list, then enter the FQDN of your NGFW Engine in the **ID Value** field.

   d) Click **OK**.

4) Click **Save**.

Define a policy-based VPN

To a policy-based VPN, first you define some basic properties for the VPN, then you add gateways.

These steps provide an overview of the VPN configuration process. For detailed instructions, see the *Forcepoint Next Generation Firewall Product Guide*.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Configuration**, then browse to **SD-WAN**.

2) Browse to **Policy-Based VPNs**.

3) Right-click **Policy-Based VPNs**, then select **New Policy-Based VPN**.

4) In the **Name** field, enter a name for the VPN.

5) (Optional) From the **Default VPN Profile** drop-down list, select the VPN Profile element that defines the settings for authentication, integrity checking, and encryption.

6) Click **OK**.
   The Policy-Based VPN opens for editing.

7) On the **Site-to-Site VPN** tab, drag and drop the gateways that you want to include in this VPN into either of the two panes for the VPN topology.

   - To allow a gateway to establish a VPN tunnel with any other gateway in the VPN, add it to the **Central Gateways** pane.
   - To allow a gateway to establish a VPN tunnel only with central gateways in this VPN, add it to the **Satellite Gateways** pane.

8) Click 💾 **Save**.

9) Add Access rules and possibly also NAT rules to direct outgoing traffic to the VPN and allow incoming traffic from the VPN.

Find product documentation

On the Forcepoint support website, you can find information about a released product, including product documentation, technical articles, and more.

You can get additional information and support for your product on the Forcepoint support website at https://support.forcepoint.com. There, you can access product documentation, release notes, Knowledge Base articles, downloads, cases, and contact information.

You might need to log on to access the Forcepoint support website. If you do not yet have credentials, create a customer account. See https://support.forcepoint.com/CreateAccount.

Published 29 May 2020