How to Deploy the F-Series Firewall in the Google Cloud via Google Launcher

https://campus.barracuda.com/doc/66682953/

You can deploy the Barracuda NextGen Firewall F-Series to the Google Cloud as a gateway or remote connectivity device. The firewall is deployed into a dedicated subnet (public subnet) in the Google Cloud network, and the instances for your cloud-based applications are deployed into backend or private subnets of the network. Each subnet is automatically assigned a dedicated gateway IP address and default route that allow the instances to connect to the Internet via the default Google Cloud gateway. An additional tag-based Google Cloud route is introduced to use the firewall as the default gateway. This route is applied automatically to all backend instances with this tag. Google Cloud firewall rules must be created to allow traffic between the firewall and the backend instances, as well as from the Internet to the firewall. By default, the Google Cloud firewall blocks all traffic, even between two instances in a subnet. The firewall has only a single DHCP network interface with a private IP address. Assign a static or ephemeral (dynamic) external IP address to your firewall to be able to connect to the Google Cloud network, even from outside the network.

Before You Begin

- Google Cloud account is required.

Step 1. Create a Network in the Google Cloud

Create the virtual network you are deploying your firewall to.

1. Go to https://console.cloud.google.com.
2. Click the hamburger menu in the upper-left corner.

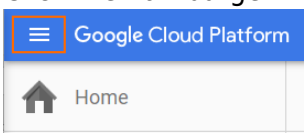

3. In the **Compute** section, click **Networking**.
4. In the main area, click **Create Network**.

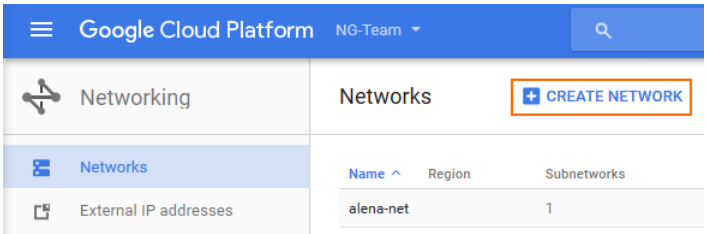

5. Enter the **Name**.
6. In the **Subnetworks** section, click **Custom**.

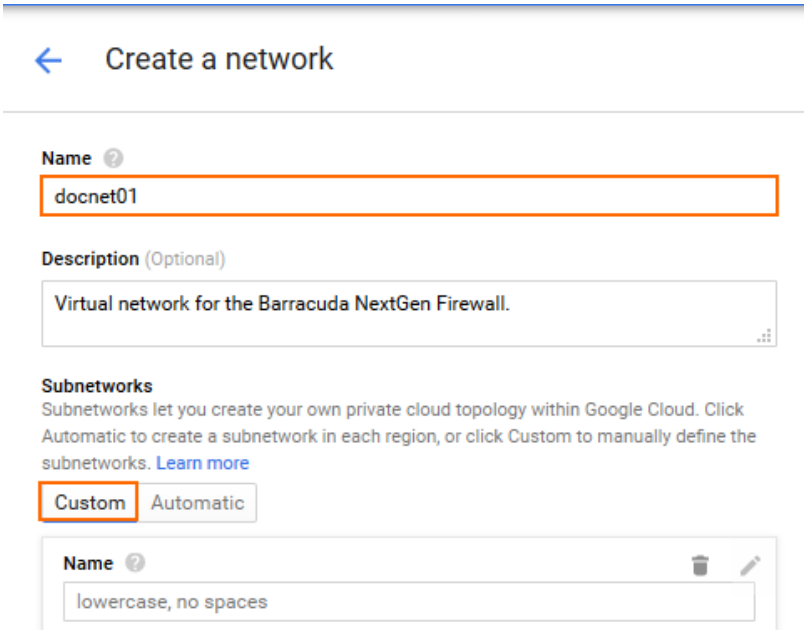

7. Create the public subnet:
   - **Name** – Enter **public-subnet**
   - **Region** – Select your region.
   - **IP address range** – Enter the network in CIDR format. If possible, do not use a network that overlaps with your on-premises network.

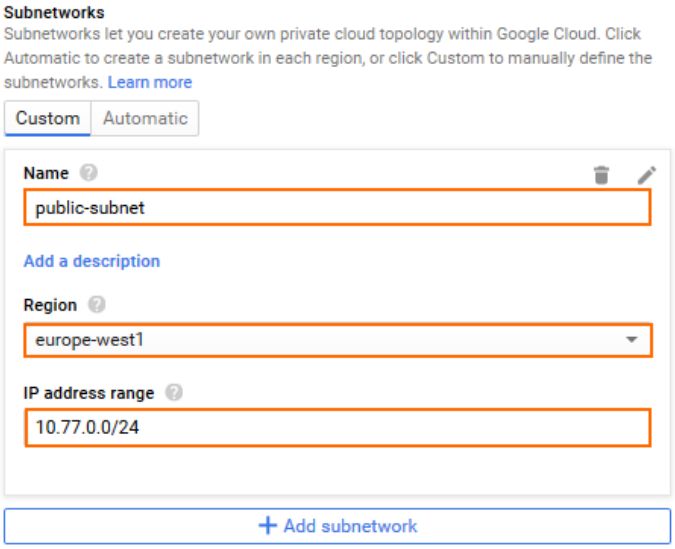

8. Click **Add subnetwork** and create the private subnet:
   - **Name** – Enter **private-subnet**
   - **Region** – Select your region.
   - **IP address range** – Enter the network in CIDR format. If possible, do not use a network that overlaps with your on-premises network.

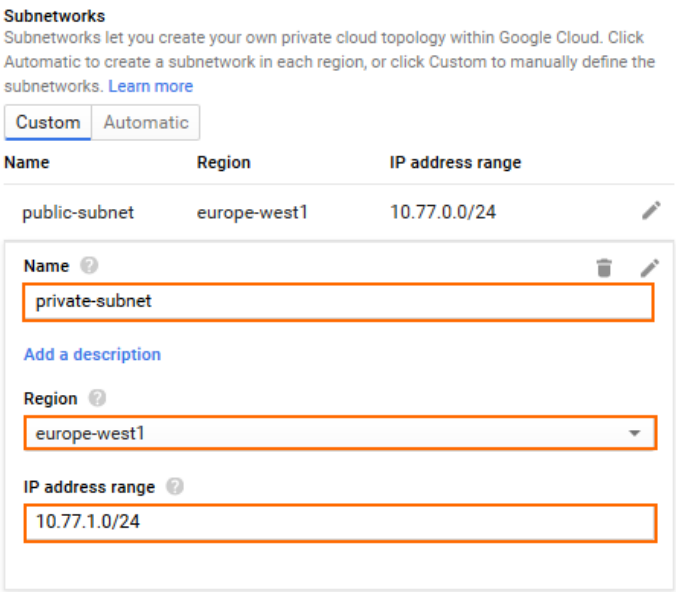

9. Click **Create**.

The network is now listed.

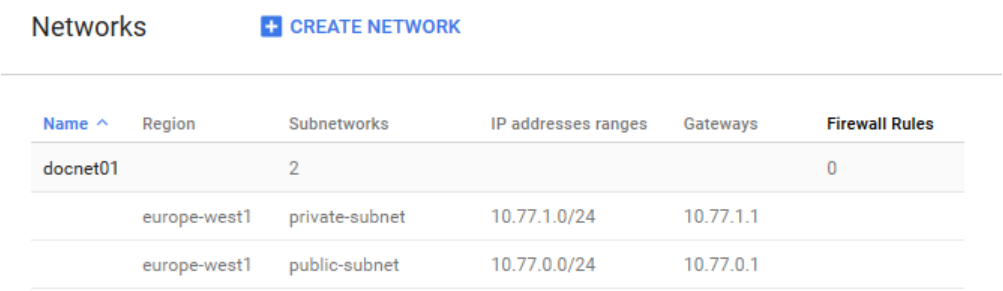

Step 2. Create an External IP Address

Create a static external IP address for your firewall. You can also skip this step and use an ephemeral IP address when creating the firewall instance.

1. Go to https://console.cloud.google.com.
2. Click the hamburger menu in the upper-left corner.
3. In the **Compute** section, click **Networking**.
4. In the left menu, click **External IP addresses**.
5. In the main area, click **Reserve static address**.

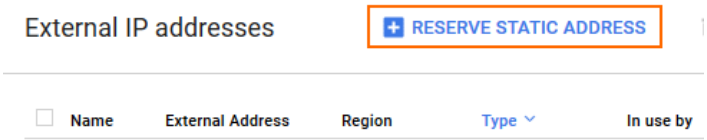

6. Reserve a static address:

- **Name** – Enter a unique name for the external IP address.
- **Type** – Select **Regional**
- **Region** – Select the same region you selected for the public subnet of the network.

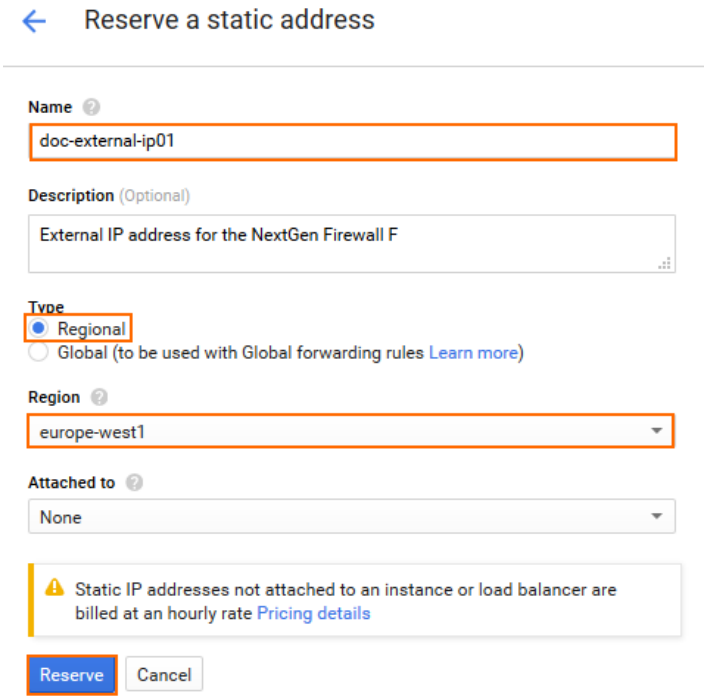

7. Click **Reserve**.

Step 3. Create the Firewall Instance from Cloud Launcher

Deploy a new NextGen Firewall instance from the Cloud Launcher image.

1. Go to the NextGen Firewall solution in Cloud Launcher: https://console.cloud.google.com/launcher/details/barracuda-release/barracuda-nextgen-firewall-f-series
2. Click **Launch on Compute Engine**.

   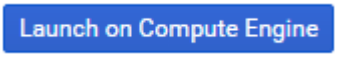

3. Enter the **Deployment name**.
4. From the **Zone** list, select the region for your new firewall instance.

   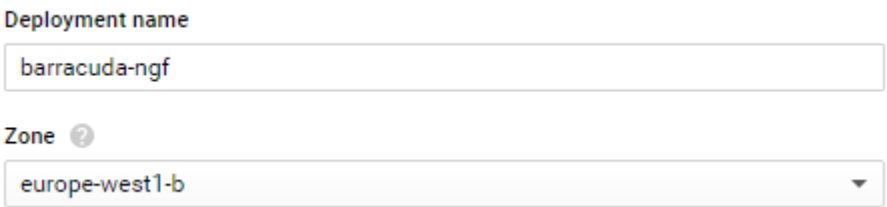

5. Select the **Machine type** with the number of vCPUs corresponding to your NextGen Firewall license and performance needs. For more information, see Public Cloud.

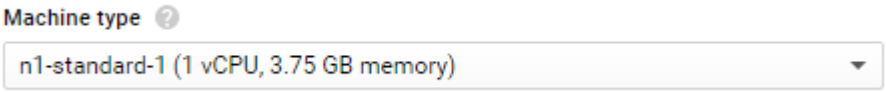

6. Change **Disk type** to SSD if you plan to use IO-intensive features like WAN Opt, Malware Protection, or HTTP Proxy. Otherwise, leave the default setting to **Standard Persistent Disk**.

7. In **Networking**, choose network and subnetwork names for the public subnet you created in Step 1.
8. Leave all default firewall positions checked. You can add more ports, protocols, and IP addresses after deployment.
9. (optional) If you want to use a reserved static address as created in Step 2:
    1. Click **More** to expand the advanced options.
    2. Select your **External IP** from the list.
10. Click **Deploy** to start the deployment.

Step 4. (optional) Create Instances in the Private Subnet

Deploy an instance into the private subnet. The backend instances must be tagged to be able to assign routes and firewall rules to them. Do not assign a public IP address to the backend instances.

Step 5. Create a Default Route for Backend Instances

A default route for each subnet with a metric of 1000 is created for each subnet. For the backend instances to use the firewall as the default gateway, create a default route with a metric lower than 1000. Configure the firewall instance as the next-hop, and add the tags identifying the backend instances. The route is automatically applied to all instances with the same tags as listed in the route.

1. Go to https://console.cloud.google.com.
2. Click the hamburger menu in the upper-left corner.
3. In the **Compute** section, click **Networking**.
4. In the left menu, click **Routes**.

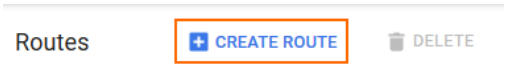

5. Click **Create route** to create the default route for the backend instances:
   - **Name** – Enter a name for the route.
   - **Network** – Select the network created in Step 1.
   - **Destination IP range** – Enter `0.0.0.0/0`.
   - **Priority** – Enter a priority lower than 1000. If two routes for the same destination exist, the route with the lower priority is used.
   - **Instance tags** – Enter the tags used for each instance that should be routed over the NextGen Firewall.
   - **Next hop** – Select **Specify an instance**.
   - **Next hop instance** – Select the firewall instance created in Step 4 from the list.

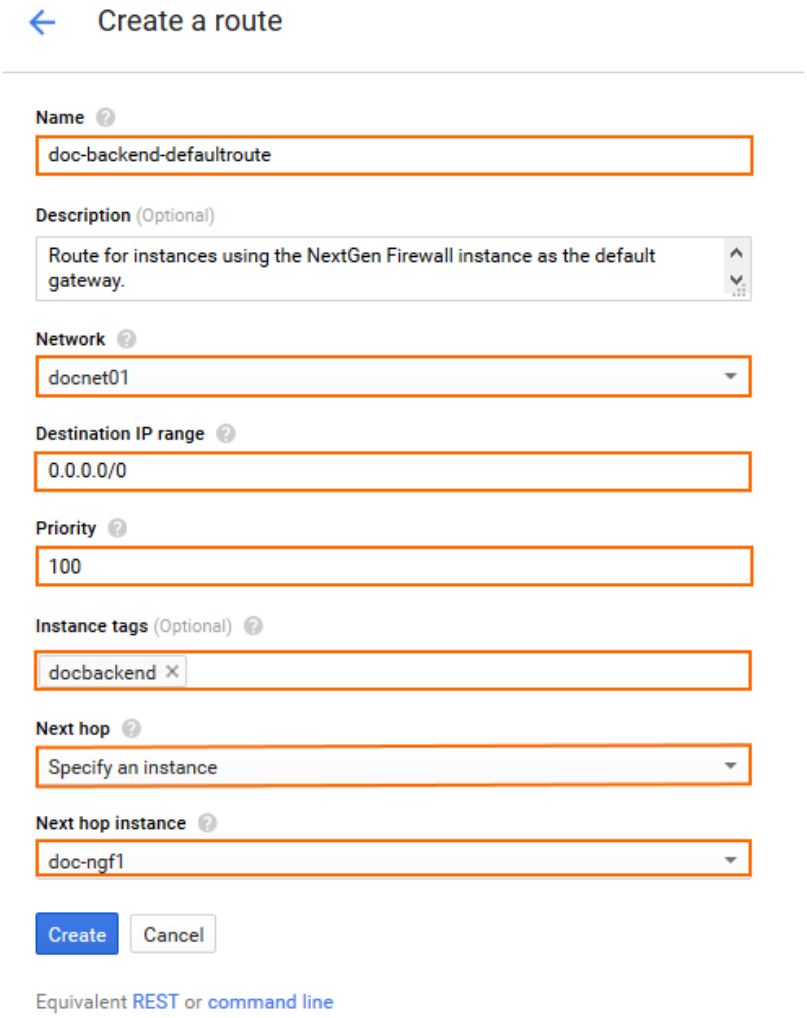

6. Click **Create**.

Step 6. Create Google Cloud Firewall Rules

Create firewall rules to allow traffic into your virtual network and from the firewall to the backend instances. By default, all traffic is blocked.

1. Go to https://console.cloud.google.com.
2. Click the hamburger menu in the upper-left corner.
3. In the **Compute** section, click **Networking**.
4. In the left menu, click **Firewall rules**.
5. In the main area, click **Create firewall rule**.

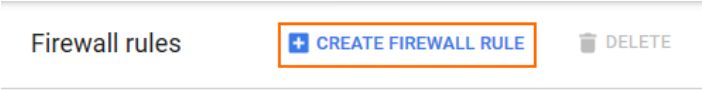

6. Create a firewall rule to allow incoming traffic to your firewall instances:
   - **Name** – Enter the firewall rule name.
   - **Network** – Select the network created in Step 1.
   - **Source filter** – Select **Allow from any source (0.0.0.0/0)**.
   - **Allowed protocols and ports** – Enter a semicolon-delimited, lower-case list of protocols and ports in the following format. tcp:807 is required to be able to connect via NextGen Admin. E.g., Use `tcp:0-65535;udp:0-65535;icmp` to allow all TCP, UDP, and ICMP traffic to the firewall.
   - **Target tags** – Enter the tag assigned to the firewall in Step 3.

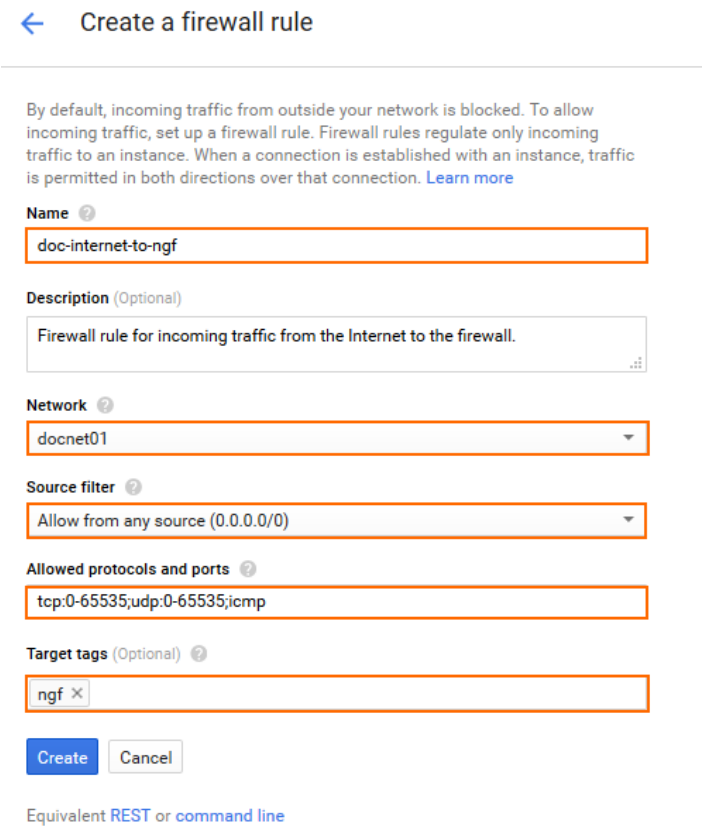

7. Create a firewall rule to allow all traffic from selected subnets to the firewall:
   - **Name** – Enter the firewall rule name.
   - **Network** – Select the network created in Step 1.
   - **Source filter** – Select **Subnetworks**.
   - **Subnetworks** – Select the public subnet and all private subnets with instances that are using the firewall as the default gateway.
   - **Allowed protocols and ports** – Enter a semicolon-delimited, lower-case list of protocols

and ports. E.g., `tcp:0-65535;udp:0-65535;icmp` to allow all TCP, UDP, and ICMP traffic between instances in these subnets.

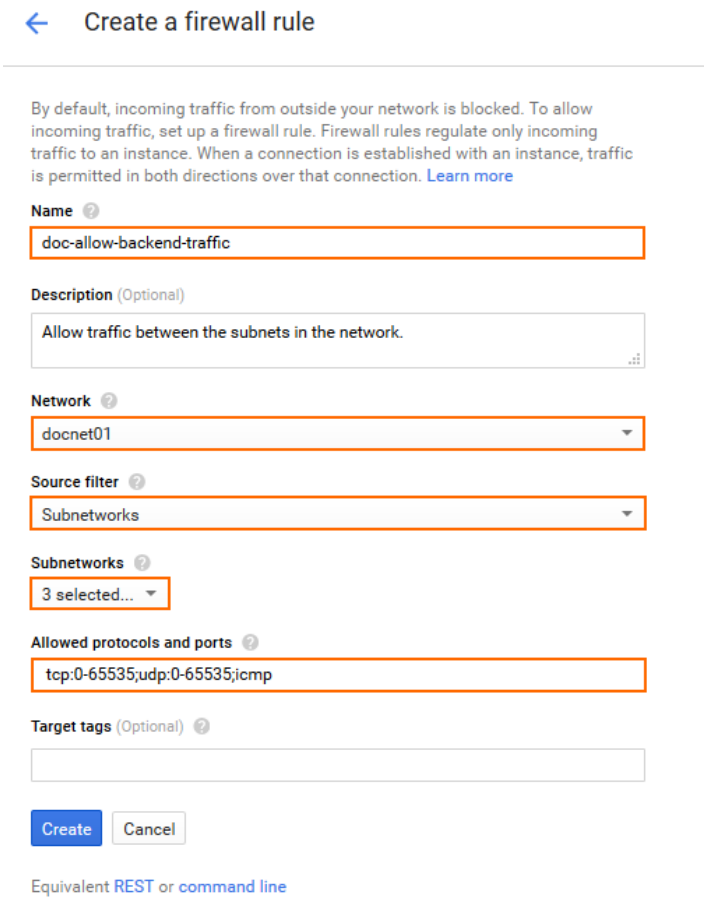

8. Click **Create**.

You can now log in to your firewall instance running in the Google Cloud using NextGen Admin:

- **IP address** – Enter the external IP address created in Step 2.
- **User** – Enter `root`
- **Password** – Enter the instance **Name**.

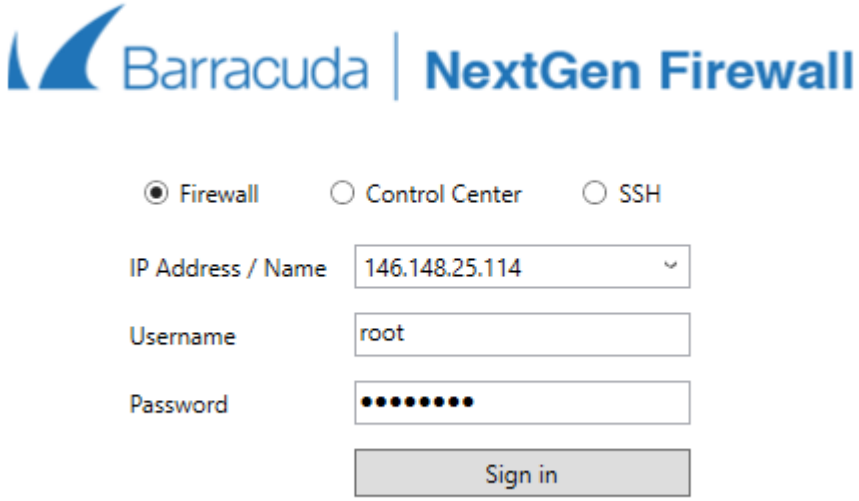

Serial Console

The Google Cloud Platform allows you to enable and connect to the serial port of your firewall instance. This feature allows you to troubleshoot your F-Series Firewall in case of a misconfiguration in a web-based serial console.

For more information, see How to Access the Serial Console on the F-Series Firewall in the Google Cloud.

Next Steps

- (BYOL only) License your firewall. For more information, see Getting Started.
- If DHCP is disabled on the CloudGen Firewall, you must also add network routes for the private IP address of the network interface with a /32 subnet mask and the default subnet gateway assigned by Google. For more information, see Step 2.2 in How to Deploy a NextGen Firewall with Multiple NICs in Google Cloud Using the Command Line.

Figures

© Barracuda Networks Inc., 2019 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.