



# Auto Scaling the Barracuda CloudGen WAF Instances in Microsoft Azure

The Barracuda CloudGen WAF can be deployed on Microsoft Azure using the Azure Resource Manager (ARM) Template. The Barracuda CloudGen WAF integrates with various Microsoft Azure services to provide Auto Scaling capabilities that enable the Barracuda CloudGen WAF deployment to scale up/down based on Azure Monitor autoscaling metrics such as CPU utilization and bandwidth.

Deployment using the Azure template also enables you to bootstrap the configuration of the Barracuda CloudGen WAF. The initial deployment allows you to specify the service configuration during launch. Later, when new instances appear, they will automatically synchronize the configuration from the previously deployed Barracuda CloudGen WAF instances and serve the traffic with complete configuration.

You can define the scaling policies for your instances and set the minimum and maximum number of instances to be used on demand. Auto Scaling can be used for applications that have stable demand as well as for applications that experience hourly, daily, or weekly variability in usage. For more information on Microsoft Azure Auto Scaling, refer to [Autoscaling](#) in the Microsoft Azure Documentation.

The Barracuda Azure Template will deploy the Barracuda CloudGen WAF with the basic service configuration and set up the necessary Azure services (Virtual Machine Scale Sets (VMSS), and Launch Configurations, Blob Storage) for successful Auto Scaling and bootstrapping.

Currently, the Barracuda CloudGen WAF provides the Azure template for Pay-As-You-Go instances only.

Microsoft Azure Services Required for the Auto Scaling Setup

- [Compute Virtual Machines](#)
- [Resource Group](#)
- [Azure Storage](#)
- [Azure Active Directory](#)

## Pay-As-You-Go (PAYG) Auto Scaling

To deploy the Pay-As-You-Go (PAYG) Barracuda CloudGen WAF in the Azure Virtual Machine Scale Sets (VMSS), follow the instructions mentioned in this article.

The Barracuda CloudGen WAF provides two types of VMSS deployments:

- Basic Bootstrapping
- Backup Based Bootstrapping
- No Bootstrapping



## Basic Bootstrapping

In **Basic Bootstrapping**, the ARM template will deploy the Barracuda CloudGen WAF in the Virtual Machine Scale Sets (VMSS) and create one HTTP service with the values provided while deploying the instance using the template. This deployment is recommended if you are starting with your first deployment.

## Backup Based Bootstrapping

In **Backup Based Bootstrapping**, (deployment using the backup file), the service(s) and other configurations are restored from the specified backup file to the virtual machine. This deployment is recommended when you want to replicate the existing backup file configuration.

## No Bootstrapping

In No Bootstrapping, the ARM template will deploy the Barracuda CloudGen WAF in the Virtual Machine Scale Sets (VMSS) without the service configuration. In No Bootstrapping deployment, you can skip the bootstrap configuration and deploy the instance. This deployment is recommended when you want to deploy the virtual machine without the service configuration.

## PAYG Virtual Machine Scale Sets ARM Template

The PAYG Virtual Machine Scale Sets ARM template includes:

- The number of Barracuda CloudGen WAF instances to be deployed in the VMSS.
- Azure Storage Account: The user specified in the ARM template (AD/Service Principal credentials) gains access to the defined Virtual Machine Scale Sets and Azure Storage account.
  - In **Basic Bootstrapping** and **No Bootstrapping**, a new Azure Storage account gets attached to the specified VMSS.
  - In **Backup Bootstrapping**, a new Azure Storage account, and an existing storage account that includes the backup file required for backup bootstrapping gets attached to the specified VMSS.
- VMSS scaling rules created for CPU and network usage (Network In/Out) to determine the scaling up/down of instances.

## Pre-requisites

The following are the prerequisites that you need to have before setting up the VMSS:

- Subnet ID where you want to deploy the Barracuda CloudGen WAF and protect your servers. Ensure the subnet is associated with the resource group where you want to deploy the Barracuda CloudGen WAF.
- Service Principal Credentials generated for the user. To generate the service principal credentials, see “Creating Service Principal Credentials” section in the **Configuring Multiple IP Addresses for the Barracuda CloudGen WAF Instance in Azure Resource Manager** article.

## Default Values of the Barracuda CloudGen WAF PAYG Azure Template

The following are the default values of the Barracuda CloudGen WAF PAYG Azure Template. You can modify the values as needed.



- **Initial Instances:** The number of Barracuda CloudGen WAF instances to be deployed initially to serve the traffic. **Default:** 2
- **Maximum Instances:** The maximum number of instances to be scaled up to handle the traffic whenever required. **Default:** 5
- **Minimum Instances:** The minimum number of instances that needs to be up to handle the traffic during any traffic condition. **Default:** 2

- Ensure that the **Minimum Instances** are lesser than or same as the **Initial Instances**.

- If the **Initial Instances** value is less than the **Minimum Instances**, the deployment of instances will fail.

- **Overprovisioning:** When set to **ENABLE**, the VMSS spins up more number of virtual machines than what is required to handle the traffic. **Default:** *DISABLE*.
- **Scale Up Thresholds:** The instances are scaled up when any of the following threshold value is triggered.

Scale Up Type	Threshold Value	Action	Evaluation Period
CPU	> 85% for 15 minutes	Bring up one instance.	15 minutes
Network In	> 9175040 bytes for 15 minutes	Bring up one instance	15 minutes
Network Out	> 9175040 bytes for 15 minutes	Bring up one instance	15 minutes

- **Scale Down Thresholds:** The instances are scaled down when any of the following threshold is triggered.

Scale Down Type	Threshold Value	Action	Evaluation Period
CPU	< 60% for 60 minutes	Bring down one instance	60 minutes
Network In	< 5242880 bytes for 60 minutes	Bring down one instance	60 minutes
Network Out	< 5242880 bytes for 60 minutes	Bring down one instance	60 minutes

- **Health Probe Settings:** The default values to probe the instances that are in the load-balanced set. Refer to [Azure Load Balancer Overview](#) in the Microsoft Azure Documentation.

**Next Step**

Continue with [Deploying the Barracuda CloudGen WAF Virtual Machine Scale Sets \(VMSS\) - PAYG Instance in Microsoft Azure](#).

