BIG-IQ® Centralized Management and Amazon® EC2®: Setup

Version 4.6

Table of Contents

Legal Notices

Legal notices

Publication Date

This document was published on November 23, 2015.

Publication Number

MAN-0512-04

Copyright

Trademarks

Patents

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Getting Started with BIG-IQ Virtual Edition

What is BIG-IQ Virtual Edition?

BIG-IQ® Virtual Edition (VE) is a version of the BIG-IQ system that runs as a guest in specifically-supported hypervisors. BIG-IQ VE emulates a hardware-based BIG-IQ system running a VE-compatible version of BIG-IQ® software.

Note: *The BIG-IQ VE product license determines the maximum allowed throughput rate. To view this rate limit, you can display the BIG-IQ VE licensing page within the BIG-IQ Configuration utility. Lab editions have no guarantee of throughput rate and are not supported for production environments.*

About BIG-IQ VE compatibility with EC2 hypervisor products

Each time there is a new release of BIG-IQ® Virtual Edition (VE) software, it includes support for additional hypervisor management products. The Virtual Edition and Supported Hypervisors Matrix on the AskF5™ website, `http://support.f5.com`, details which hypervisors are supported for each release.

BIG-IQ VE is compatible with the Amazon Web Services (AWS) EC2 hypervisors. This guide documents the AWS interface as it exists just prior to the version 11.6.0 BIG-IP software release.

Important: *Hypervisors other than those identified in this guide are not supported with this BIG-IQ version; any installation attempts on unsupported platforms might not be successful.*

About the hypervisor guest definition requirements

The EC2 virtual machine guest environment for the BIG-IQ® Virtual Edition (VE), at minimum, must include:

- a 64 bit EC2 instance with at least 2 virtual cores (up to 16 are supported in this release)
- at least 4 GB RAM (64GB has been tested, F5® Networks recommends at least 2GB per virtual core)
- 2 x virtual network adapter cards (NICs) (up to 9 are supported)

  Important: *F5® Networks recommends three or more network adapters for most topologies, but the minimum requirement is two (one for management and one for traffic).*

  Important: *To support multiple NICs on Amazon Web Services you must create a virtual private cloud (VPC).*

- 1 x virtual private cloud (VPC).

Important: *Not supplying at least the minimum virtual configuration limits will produce unexpected results.*

Important: *Although you can successfully deploy BIG-IQ software with as few as 2 CPUs and 4 GB RAM, this configuration should only be used for evaluation purposes. For production use, F5 Networks reccomends either 4 CPUs and 16 GB RAM, or (for higher performance) 8 CPUs and 32 GB RAM.*

Note: *Currently, these requirements map to what Amazon designates as an M1 Large instance. Refer to* `http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/instance-types.html#AvailableIpPerENI` *for their most current definition.*

Deploying BIG-IQ Virtual Edition

About VE EC2 deployment

To deploy the BIG-IQ® Virtual Edition (VE) system on EC2, you perform these tasks:

- Verify the host machine requirements.
- Deploy a BIG-IQ® system as a virtual machine.
- Deploy a BIG-IP® system.
- After you have deployed the virtual machines, log in to the BIG-IQ VE system and run the Setup utility. Using the Setup utility, you perform basic network configuration tasks, such as assigning VLANs to interfaces.
- Configure secure communication between the BIG-IQ system and the BIG-IP device.

Task summary for BIG-IQ VE EC2 deployment

To deploy BIG-IQ® Cloud you perform a series of tasks using Amazon Web Services (AWS) to create an elastic compute cloud (EC2) that runs a public cloud virtual machine management service.

When you complete these tasks, your cloud environment will be similar to the basic cloud topology depicted here.

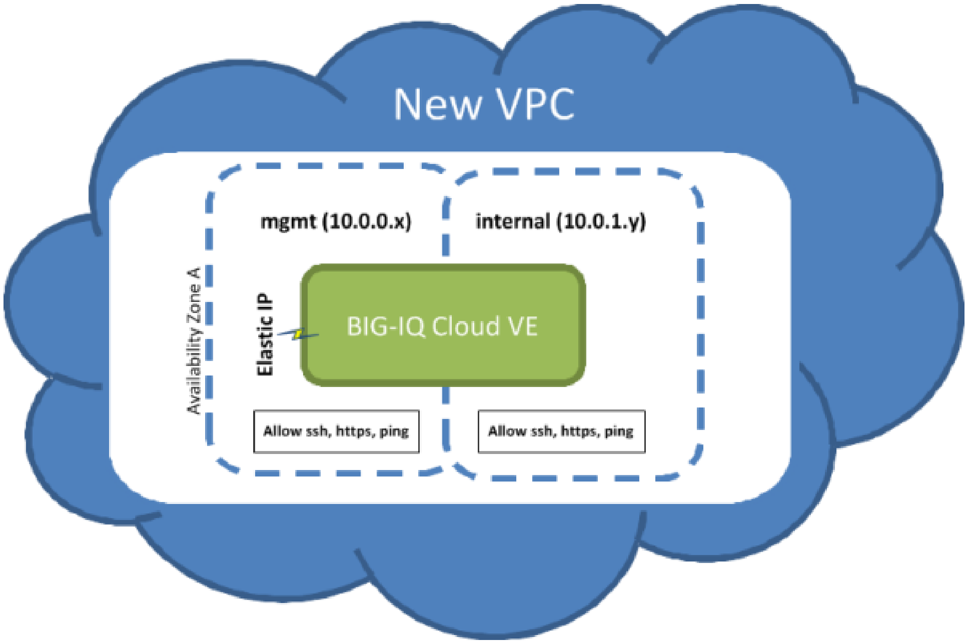

Figure 1: Basic Cloud Topology

Task list

Creating an Amazon Identity and Access Management (IAM) user account
Creating a key pair

Creating a Virtual Private Cloud
Adding an additional subnet
Creating new security groups
Adding a route for external subnet accessibility
Launching a virtual server with an Amazon Machine Image (AMI)
Adding a third network interface
Making the virtual machine management port accessible
Logging in and setting the Admin password

Creating an Amazon Identity and Access Management (IAM) user account

An Amazon Identity and Access Management (IAM) user account provides access to specific Amazon Web Services (AWS) resources. Creating an IAM account provides you with more granular control of the AWS resources your users access.

Important: *This task is optional; you can create a virtual machine without creating an IAM user account to control access, but it is best practice to use an IAM account. F5 recommends that you do not use the AWS root account and access keys. Instead, use IAM to create identities you can more easily manage and revoke in the case of a security breach.*

Tip: *When you manually deploy a virtual machine on AWS EC2, you must create an administrator password in addition to the IAM access keys. If you use the automated process to deploy a virtual server, only the access keys are required.*

For this task, you must create a group and two IAM user accounts. For the most current instructions for performing these steps, refer to the IAM documentation web site, `http://aws.amazon.com/documentation/iam/`.

1. From `https://console.aws.amazon.com/iam`, create a group with aws-full-access (Administrator Access).
2. Create an AWS-Admin user and add that user to the **aws-full-access** group.
3. Create a BIG-IQ Connector user and add that user to the **aws-full-access** group.

   For this user, you must download or copy an access key that you use to connect BIG-IQ Cloud to your AWS account
4. From the AWS dashboard, set up an account alias.

   Note the IAM user login link. For example, `https://my-account-alias.signin.aws.amazon.com/console`
5. Log out of the AWS dashboard as the root user.
6. Navigate back to the user login link and sign in as the **AWS-Admin** user.

You can now create a new Virtual Private Cloud (VPC).

Creating a key pair

Before you can deploy a virtual machine on Amazon Web Services (AWS) Elastic Cloud Computing, you need an AWS account.

To create a virtual private cloud (VPC) on which you can deploy the BIG-IQ® system, you need a (private-public encryption) key pair to authenticate your sessions. Key pairs are reusable, so if you have a key pair, you do not need to repeat this task.

For the most current instructions for creating a key pair, refer to the Amazon Virtual Private Cloud (VPC) Documentation web site, `http://aws.amazon.com/documentation/vpc/`.

Important: *It is crucial to your success that you be consistent in the region that you choose throughout the configuration process. Objects configured in one region are not visible within other regions, so they cannot function together. There are a number of factors that determine which region will best suit your requirements. Refer to Amazon user documentation for additional details.*

The file that downloads from Amazon Web Services uses the extension `.pem`. If you plan to use this key pair with the PuTTY terminal emulator application, you will need to convert the key pair from a `.pem` to a `.ppk` file. At the time of this release, PuTTY does not support the extension `.pem`. PuTTY does have a tool (called PuTTYgen) that converts your key pair to the required PuTTY format.

Creating a Virtual Private Cloud

You need an Amazon Virtual Private Cloud (VPC) to deploy the BIG-IQ® Cloud system, because AWS provides only multiple network interface card (NIC) support for instances that reside within a VPC.

You create a virtual network topology according to your networking needs. The standard network topology used for BIG-IQ Cloud integration includes three subnets. These subnets provide virtual private address spaces used to interconnect your machines and applications. You can use elastic self IP addresses for public internet accessibility.

For the most current instructions for creating a VPC, refer to the VPC Documentation web site, `http://aws.amazon.com/documentation/vpc/`.

1. Navigate to `https://console.aws.amazon.com/vpc` and select the AWS Region in which you want to manage resources.

   For example, Oregon.
2. From the VPC Wizard's **VPC with Public and Private Subnets** option, set the IP CIDR Block to `10.0.0.0/16`.
3. Set the public subnet to `10.0.0.0/24`.

   This is the management network.
4. Select an availability zone.

   For example, **us-west-2c**. It is crucial that you use this availability zone throughout the configuration process. Objects configured in one zone are not visible within other zones, so they cannot function together. This availability zone is required when you create a BIG-IQ Cloud connection.
5. Set the private subnet to `10.0.1.0/24`.

   This is the external data network.
6. Create subnet `10.0.2.0/24`.

   This is the internal network.
7. Create a security group named, `allow-all-traffic`, and associate it with the VPC you created.

   You must use this exact name.
8. Set the **Inbound Rules ALL Traffic Source** to `0.0.0.0/0`.
9. Set the **Outbound Rules ALL Traffic Destination** to `0.0.0.0/0`.
10. Create a Route Table for the external data network to reach the Internet.
11. Add a route to Destination **0.0.0.0/0** through Target `igw-<xxxx>`.

    `<xxxx>` is the Internet Gateway that the VPC Wizard created automatically.

12. Allocate two Elastic IP Addresses.

You now should create a BIG-IQ Cloud connector to associate with this VCP.

Adding an additional subnet

When you create a VPC, Amazon Web Services creates two subnets for it. The first subnet is the management subnet (`10.0.0.0/24`) and the second subnet is external (`10.0.1.0/24`). Many network topologies require three or more subnets (Management, External, and Internal). You can use this task to create an internal subnet (`10.0.2.0/24`).

For the most current instructions for creating an internal subnet, refer to the Amazon Virtual Private Cloud (VPC) Documentation web site `http://aws.amazon.com/documentation/vpc/`.

If you are following a typical deployment strategy, when you finish adding the Internal subnet, your VPC will have three subnets.

- a Management subnet on `10.0.0.0/24`
- an External subnet on `10.0.1.0/24`
- an Internal subnet on `10.0.2.0/24`

Creating new security groups

To use your virtual private cloud (VPC) to deploy your virtual machine, the VPC needs two security groups; each with its own set of rules that govern the security behavior for the traffic that routes through it. The table details the rules required for each group to function properly.

Group Name	Group Description	Rule Name	Source	Rule Type
`allow-only-ssh-https-ping`	Allow only SSH HTTPS or PING	Inbound SSH	0.0.0.0/0	
		Inbound HTTPS	0.0.0.0/0	
		Inbound Custom ICMP	0.0.0.0/0	Echo Request
		Outbound Custom ICMP	0.0.0.0/0	Echo Request
		Outbound Custom ICMP	0.0.0.0/0	Echo Reply
`allow-all-traffic`	Allow all traffic	Inbound All Traffic	0.0.0.0/0	
		Outbound All Traffic	0.0.0.0/0	

1. Create two security groups, one named `allow-only-ssh-https-ping` and the other named `allow-all traffic`.

__Tip:__ For the most current instructions for creating security groups, refer to the Amazon Virtual Private Cloud (VPC) Documentation web site `http://aws.amazon.com/documentation/vpc/`.

Important: *The* `allow-all-traffic` *security group is critically important for successful operation of the BIG-IP® VE on Amazon EC2.*

2. For each security group, create the rules described in the preceding table. For each rule, define the Group Description, Rule Name, Source, and Rule Type as shown in the table.

   Important: *No punctuation is permitted in the text of the Group Description that you type in.*

When you finish adding the two groups and their associated rules, your VPC should be ready to go with three subnets and two security groups.

It is a good idea to test connectivity before proceeding. You should be able to communicate with your VPC NAT server at this point.

F5 Networks recommends enhancing your security by using the security group source fields to restrict the subnets to allow only management access; however, we recognize that this does not complete your security solution. For enhanced security, you might want to deploy a topology with limited management network access.

Adding a route for external subnet accessibility

Most network topologies require an Amazon Web Services route to the virtual private cloud (VPC) that makes the external subnet used by the virtual machine accessible to the Internet.

1. From the Services tab at the top of the Amazon Web Services Management Console screen, select **VPC**.
2. In the navigation pane, select **Route Tables**.
   The Route Tables screen opens.
3. Select the routing table with one subnet.
4. Click the Associations tab at the bottom of the window.
5. From the **Select a subnet** list, select the **10.0.1.0/24** subnet.
6. Click **Associate**.
   The Associate Route Table popup screen opens.
7. Click **Yes, Associate**.

Launching a virtual server with an Amazon Machine Image (AMI)

Before you can complete this task, you need to know the name of your key pair and the Availability Zone from which it was created.

You launch an EC2 Amazon Machine Image (AMI) so that you can deploy the virtual machine.

Important: *At publication, this task illustrates the Amazon web interface. However, F5 recommends that you refer to Amazon user documentation for the latest documentation.*

1. Log in to your account on Amazon Web Services (AWS) marketplace.
2. In the Search AWS Marketplace bar, type `F5 BIG-IQ` and then click **GO**.
   The F5 BIG-IQ Virtual Edition for AWS option is displayed.
3. Click **F5 BIG-IQ Virtual Edition for AWS** and then click **CONTINUE**.

   Tip: *You might want to take a moment here to browse the pricing details to confirm that the region in which you created your security key pair provides the resources you require. If you determine that the resources you need are provided in a region other than the one in which you created your key pair, create a new key pair in the correct region before proceeding.*

   The Launch on EC2 page is displayed.

4. Click the **Launch with EC2 Console** tab.

   Important: *At the time this was written, the virtual machine must be launched in a VPC so that NICs can be attached. This configuration is supported from the* ***Launch with EC2 Console*** *option, but not the* ***1-Click Launch*** *option.*

   Launching Options for your EC2 AMI are displayed.
5. Select the software version appropriate for your installation, and then click the **Launch with EC2** button that corresponds to the Region that provides the resources you plan to use.

   Important: *The first time you perform this task, you need to accept the terms of the end user license agreement before you can proceed, so the* ***Launch with EC2*** *button reads* ***Accept Terms and Launch with EC2****.*

   Important: *There are a number factors that determine which region will best suit your requirements. Refer to Amazon user documentation for additional detail. Bear in mind that the region you choose must match the region in which you created your security key pair.*

   The Request Instances Wizard opens.
6. Select an **Instance Type** appropriate for your use.
7. From the **Launch Instances** list, select **EC2-VPC**.
8. From the **Subnet** list, select the **10.0.0.0/24** subnet and click **CONTINUE**.
   The Advanced Instance Options view of the wizard opens.
9. From the **Number of Network Interfaces** list, select **2**.
10. Click the horizontal **eth1** tab to set values for the second network interface adapter, and then from the **Subnet** list, select the **10.0.1.0/24** subnet and click **CONTINUE**
    The Storage Device Configuration view of the wizard opens.
11. In the **Value** field, type in an intuitive name that identifies this AMI and click **CONTINUE** (for example, `BIG-IQ VE <version>`).
    The Create Key Pair view of the wizard opens.
12. From **Your existing Key Pairs**, select the key pair you created for this AMI and click **CONTINUE**.
    The Configure Firewall view of the wizard opens.
13. Under Choose one or more of your existing Security Groups, select the **allow-all-traffic** security group, and then click **CONTINUE**.
    The Review view of the wizard opens.
14. Confirm that all settings are correct, and then click **Launch**.
    The Launch Instance Wizard displays a message to let you know your instance is launching.
15. Click **Close**.

Your new instance appears in the list of instances when it is fully launched.

Adding a third network interface

When you first create a virtual private cloud (VPC), there are typically only two network interfaces associated with it. F5 Networks recommends adding a third network interface to the VPC before you use it to deploy the virtual machine.

1. From the Services tab at the top of the Amazon Web Services (AWS) Management Console screen, select **EC2**.
2. In the navigation pane, select **Network Interfaces**.
   The Network Interfaces screen opens.

3. Click the **Create Network Interface** button (at top left).
   The Create Network Interface popup screen opens.
4. In the **Description** field, type `Internal 10.0.2.0-24` (or a similarly mnemonic name).
5. In the **Subnet** field, select **10.0.2.0/24**.
6. From the **Security Groups** list, select **allow-all-traffic**.
7. Click **Yes, Create**
   AWS adds your network interface to the list.
8. Right-click the new network interface, and then select **Attach**.
   The Attach Network Interface popup screen opens.
9. From the **Instance** list, select the VE AMI that you created.

Making the virtual machine management port accessible

The management port for your virtual machine might require accessibility over the Internet. However, there are alternative topologies that do not require exposing the management port to the Internet.

F5 Networks recommends, at a minimum, adding restrictions to your source addresses in the `allow-only-ssh-https-ping` security group.

Alternatively, you might find the Amazon Web Services EC2 VPN sufficiently effective so that you do not need to associate an Internet-accessible Elastic IP with the management port.

1. From the Services tab at the top of the Amazon Web Services Management Console screen, select **EC2**.
2. In the navigation pane, select **Elastic IPs**.
   The Addresses screen opens.
3. Click **Allocate New Address**.
   The Allocate New Address popup screen opens.
4. From the **EIP used in** list, select **VPC**.
5. Click **Yes, Allocate**.
6. In the Address column, right-click the newly created Elastic IP and select **Associate** from the popup menu.
   The Associate Address popup screen opens.
7. From the **Instance** list, select the VE AMI that you created as an EC2 hypervisor.
8. From the **Private IP Address** list, select **10.0.0.0/24** (the Management subnet).
9. Click **Yes, Associate**.

Logging in and setting the Admin password

To perform this task, you must have completed the following tasks:

- Created a key pair
- Created and configured a VPC
- Instantiated and launched a BIG-IQ® Virtual Edition (VE) AMI
- Made the virtual machine management port accessible through the Internet

To maintain security, the first time you log in to your EC2 AMI, you should log in as root, and change the Admin password.

1. Log in to the new AMI that you just launched.

   Use the name of the key pair (`.pem` file), and the elastic IP address of your EC2 instance. `$ ssh -i <`*`username`*`>-aws-keypair.pem root@<`*`elastic IP address of EC2 instance`*`>`

Tip: You can also use a terminal emulator such as PuTTY to test your connectivity. At publication, PuTTY does not support the extension `.pem`, so remember that you will also need to convert the key pair `.pem` file to a `.ppk` file before you can use it with PuTTY.

2. At the command prompt, type `tmsh modify auth password admin`.

   Warning: *Because this login is visible externally, make sure to use a strong, secure password.*

   The terminal window displays the message: `changing password for admin`, and then prompts: `new password`.
3. Type in your new password and then press **Enter**.
   The terminal window displays the message: `confirm password`.
4. Re-type the new password and then press Enter.
5. To ensure that the system retains the password change, type `tmsh save sys config`, and then press Enter.

   Important: *Without your security key pair, you cannot access this AMI. Once you log in with your key pair, you could create a root password. However, if you decide to do this, choose the root password wisely, bearing in mind that depending on your Security Group policies, this login might provide external SSH access.*

The Admin password is now changed.

Tip: *If at some point you determine you want to restore your virtual machine to its default state, use the following `tmsh` command: `tmsh load sys config default; bigstart restart dhclient` Using a less explicit `tmsh` command has been shown to produce undesirable results.*

Index