

# **BIG-IQ™ Systems and Amazon® EC2®: Setup and Getting Started**

Version 4.1.0





# Table of Contents

<b>Legal Notices.....</b>	<b>5</b>
<b>Acknowledgments.....</b>	<b>7</b>
<b>Chapter 1: Getting Started with BIG-IQ Virtual Edition.....</b>	<b>13</b>
What is BIG-IQ Virtual Edition?.....	14
About BIG-IQ VE compatibility with EC2 hypervisor products.....	14
About the hypervisor guest definition requirements.....	14
<b>Chapter 2: Deploying BIG-IQ Virtual Edition.....</b>	<b>15</b>
About VE EC2 deployment.....	16
Task summary for EC2 deployment.....	16
<b>Chapter 3: BIG-IQ System Overview.....</b>	<b>23</b>
Overview: BIG-IQ system.....	24
Additional resources and documentation for BIG-IQ systems.....	24
<b>Chapter 4: Installation and Initial Configuration.....</b>	<b>25</b>
About licensing and initial configuration.....	26
Automatic license activation and initial configuration.....	26
Manual license activation and initial configuration.....	27
About installing required BIG-IQ system components on managed BIG-IP systems	
.....	28
Installing required BIG-IQ components on BIG-IP systems .....	28
About a high availability configuration.....	29
Configuring a high availability pair.....	29
Replacing a peer in a high availability configuration.....	29
Forcing an active peer BIG-IQ system to standby mode.....	30
About default passwords for pre-defined users.....	30
Changing the default password for the administrator user.....	30
Changing the default password for the root user.....	31



# Legal Notices

---

## Publication Date

This document was published on July 31, 2013.

## Publication Number

MAN-0482-00

## Copyright

Copyright © 2013, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

## Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, Alive With F5, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, Scale<sup>N</sup>, Signalling Delivery Controller, SDC, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TMOS, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, VIPRION, vCMP, VE F5 [DESIGN], Virtual Clustered Multiprocessing, WA, WAN Optimization Manager, WebAccelerator, WOM, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.

## Patents

This product may be protected by one or more patents indicated at:

<http://www.f5.com/about/guidelines-policies/patents>



# Acknowledgments

---

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, <http://www.and.com>.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler ([bazsi@balabit.hu](mailto:bazsi@balabit.hu)), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller ([nisse@lysator.liu.se](mailto:nisse@lysator.liu.se)), which is protected under the GNU Public License.

## Acknowledgments

In the following statement, *This software* refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. *Similar operating systems* includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), [www.gnu.org/copyleft/lgpl.html](http://www.gnu.org/copyleft/lgpl.html).

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at <http://www.perl.com>.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes cryptographic software written by Eric Young ([eay@cryptsoft.com](mailto:eay@cryptsoft.com)).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes RRDtool software developed by Tobi Oetiker (<http://www.rrdtool.com/index.html>) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (<http://www.nominum.com>).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes software developed by Oracle America, Inc. Copyright ©2012.

1. Java Technology Restrictions. Licensee shall not create, modify, change the behavior of, or authorize licensees of licensee to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun" or similar convention as specified by Oracle in any naming convention designation. In the event that Licensee creates an additional API(s) which: (a) extends the functionality of a Java Environment; and (b) is exposed to third party software developers for the purpose of developing additional software which invokes such additional API, Licensee must promptly publish broadly an accurate specification for such API for free use by all developer.
2. Trademarks and Logos. This License does not authorize an end user licensee to use any Oracle America, Inc. name, trademark, service mark, logo or icon. The end user licensee acknowledges that Oracle owns the Java trademark and all Java-related trademarks, logos and icon including the Coffee Cup and Duke ("Java Marks") and agrees to: (a) comply with the Java Trademark Guidelines at <http://www.oracle.com/html/3party.html>; (b) not do anything harmful to or inconsistent with Oracle's



rights in the Java Marks; and (c) assist Oracle in protecting those rights, including assigning to Oracle any rights acquired by Licensee in any Java Mark.

3. **Source Code.** Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of your license. Source code may not be redistributed unless expressly provided for in the terms of your license.
4. **Third Party Code.** Additional copyright notices and license terms applicable to portion of the Software are set forth in the THIRDPARTYLICENSEREADME.txt file.
5. **Commercial Features.** Use of the Commercial Features for any commercial or production purpose requires a separate license from Oracle. "Commercial Features" means those features identified in Table I-I (Commercial Features In Java SE Product Editions) of the Software documentation accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

This product includes software developed by members of the CentOS Project under the GNU Public License, copyright ©2004-2011 by the CentOS Project.

This product includes software developed by members of the OpenJDK Project under the GNU Public License Version 2, copyright ©2012 by Oracle Corporation.

This product includes software developed by The VMWare Guest Components Team under the GNU Public License Version 2, copyright ©1999-2011 by VMWare, Inc.

This product includes software developed by The Netty Project under the Apache Public License Version 2, copyright ©2008-2012 by The Netty Project.

This product includes software developed by Stephen Colebourne under the Apache Public License Version 2, copyright ©2001-2011 Joda.org.

This product includes software developed by the GlassFish Community under the GNU Public License Version 2 with classpath exception, copyright ©2012 Oracle Corporation.

This product includes software developed by the Mort Bay Consulting under the Apache Public License Version 2, copyright ©1995-2012 Mort Bay Consulting.

This product contains software developed by members of the Jackson Project under the GNU Lesser General Public License Version 2.1, ©2007 – 2012 by the Jackson Project”.

This product contains software developed by QOS.ch under the MIT License, ©2004 – 2011 by QOS.ch.

This product includes software licensed from Gerald Combs ([gerald@wireshark.org](mailto:gerald@wireshark.org)) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

## Acknowledgments

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This software incorporates JFreeChart, ©2000-2007 by Object Refinery Limited and Contributors, which is protected under the GNU Lesser General Public License (LGPL).

This product contains software developed by the Mojarrá project. Source code for the Mojarrá software may be obtained at <https://javaserverfaces.dev.java.net/>.

This product includes JZlib software, Copyright © 2000-2011 ymnk, JCraft, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes Apache Lucene software, distributed by the Apache Software Foundation, under the Apache License version 2.0.

This product includes Apache MINA software, distributed by the Apache Software Foundation, under the Apache License version 2.0.

This product includes OData4J software, distributed under the Apache License version 2.0.

This product includes software developed by the Visigoth Software Society (<http://www.visigoths.org/>).

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes software developed by Addy Osmani, and distributed under the MIT license. Copyright © 2012 Addy Osmani.

This product includes software developed by Charles Davison, and distributed under the MIT license. Copyright © 2013 Charles Davison.

This product includes software developed by The Dojo Foundation, and distributed under the MIT license. Copyright © 2010-2011, The Dojo Foundation.

This product includes google-gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes ec2-tools software, copyright © 2008, Amazon Web Services, and licensed under the Amazon Software License. A copy of the License is located at <http://aws.amazon.com/asl/> .



---

# Chapter 1

---

## Getting Started with BIG-IQ Virtual Edition

---

- *What is 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

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 11.3.0 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 an 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:** There is no longer any limitation on the maximum amount of RAM supported on the hypervisor guest.*

---

***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.*

---

---

# Chapter 2

---

## Deploying BIG-IQ Virtual Edition

---

- *About VE EC2 deployment*
-

## 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 EC2 deployment

To deploy BIG-IQ™ Cloud using the Amazon elastic compute cloud (EC2) public cloud virtual machine management service, you perform a series of tasks using Amazon Web Services (AWS).

#### Task list

#### 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 new virtual private cloud

You need a virtual private cloud (VPC) to deploy the BIG-IQ™ Cloud system because Amazon Web Services (AWS) only provides multiple network interface card (NIC) support for EC2 instances that reside within a VPC.



For the most current instructions for creating a Virtual Private Cloud, 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 availability zone that you choose throughout the configuration process. Objects configured in one zone are not visible within other zones, so they cannot function together.

---

**Important:** The first choice you have when creating a VPC is to select a VPC configuration. Choose the VPC with **Public and Private Subnets** option.

---

## Adding an additional subnet

When you create a VPC, Amazon Web Services creates two subnets (Management and External) for it. Many network topologies require three or more subnets (Management, External, and Internal).

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 may 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 dialog box opens.
7. Click **Yes, Associate**.

### Launching a new virtual machine

You need to know the name of your key pair and the Availability Zone from which they were created before you can complete this task.

You need to have an EC2 Amazon Machine Image (AMI) to 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 **VPC**.
8. From the **Subnet** list, select the **10.0.0.0/24** subnet and then 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 then 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 then 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 then 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 AMI 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 dialog box 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 dialog box 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 may 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 may 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 dialog box 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 dialog box 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 may 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.

---



---

# Chapter

# 3

---

## BIG-IQ System Overview

---

- *Overview: BIG-IQ system*
-

## Overview: BIG-IQ system

---

The BIG-IQ™ system is a centralized tool that streamlines the management of devices in your network. The functionality offered is dependent on your software license.

Administrators use BIG-IQ Cloud to provide cloud tenants self-service access to shared computing resources such as networks, servers, storage, applications, and services. Cloud resources can be private or public, depending on the customer's requirements. Each tenant has restricted and dedicated access to cloud resources based on a specific user account or tenant role, ensuring that tenants have access only to their own resources. Cloud resources are easily expanded and reallocated as needed, providing flexible resource balancing.

Firewall managers use BIG-IQ Security to manage security firewalls for multiple devices from one central location. Firewall management includes discovering, editing, and deploying firewall configurations, as well as consolidating shared firewall objects. Once a firewall device is designated for central management, it is no longer managed locally unless there is an exceptional need.

### Additional resources and documentation for BIG-IQ systems

You can access all of the following BIG-IQ™ system documentation from the AskF5™ Knowledge Base located at <http://support.f5.com/>.

Document	Description
<i>BIG-IQ™ Systems: Setup and Getting Started guides</i>	BIG-IQ Virtual Edition (VE) runs as a guest in a virtual environment using supported hypervisors. Each of these guides is specific to one of the hypervisor environments supported for the BIG-IQ system. The guides provide you with the basic concepts and tasks required to set up your virtual environment and perform initial configuration tasks required to start using the BIG-IQ system.
<i>BIG-IQ™ Systems: Cloud Management</i>	This guide contains information to help you manage cloud resources, applications, and tenants.
<i>BIG-IQ™ Security User Guide</i>	This guide contains information to help you manage BIG-IP® firewall resources.
Release notes	Release notes contain information about the current software release, including a list of associated documentation, a summary of new features, enhancements, fixes, known issues, and available workarounds.
Solutions and Tech Notes	Solutions are responses and resolutions to known issues. Tech Notes provide additional configuration instructions and how-to information.



---

# Chapter 4

---

## Installation and Initial Configuration

---

- *About licensing and initial configuration*
- *About installing required BIG-IQ system components on managed BIG-IP systems*
- *About a high availability configuration*
- *About default passwords for pre-defined users*

### About licensing and initial configuration

---

BIG-IQ™ system runs as a virtual machine in specifically-supported hypervisors. After you set up your virtual environment, or your platform, you can license the BIG-IQ system. You initiate the license activation process with the base registration key.

The *base registration key* is a character string that the license server uses to verify the functionality that you are entitled to license. If the system has access to the internet, you select an option to automatically contact the F5 license server and activate the license. If the system is not connected to the internet, you can manually retrieve the activation key from a system that is connected to the internet, and transfer it to the BIG-IQ system.

---

*Note:* If you do not have a base registration key, contact the F5 Networks sales group (<http://www.f5.com>).

---

### Automatic license activation and initial configuration

Before you can activate your F5 product license, you must configure your virtual environment or your platform, define a management IP address, and obtain the base registration key.

If the BIG-IQ™ system is connected to the public internet, use this procedure to activate its license.

1. Using a browser on which you have configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you specified for device management:  
`https://<management_IP_address>`
2. Log in to the BIG-IQ system with the default user name `admin` and password `admin`.
3. At the top of the screen, click **System Overview**.
4. In the Setup panel, click the IP address of the BIG-IQ system.  
The panel expands to display additional properties.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Key** field, paste any additional license key you have.  
The options are: `Cloud` or `Security`.
7. For the Activation method setting, select **Automatic** and click the **Activate** button.  
The BIG-IQ system contacts the F5 Networks licensing server and displays The End User License Agreement (EULA) displays.
8. To accept the EULA, click the **Accept** button.  
The screen refreshes and display the license details.
9. Click the **Properties** tab.
10. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.  
The FQDN can consist of letters, numbers, and/or the characters underscore (`_`), dash (`-`), or period (`.`).
11. In the **Self IP Address** field, type the self IP address of your internal VLAN.  
The self IP address must be in Classless InterDomain Routing (CIDR) format. For example:  
`10.10.10.10/24`.
12. To add an additional self IP address, click the + sign and in the new **Self IP Address** field that the system creates, edit the duplicated self IP address to reflect the additional self IP address that you want to add.  
Once you save this self IP address, you cannot change it.

13. Click the **Services** tab.
14. In the **DNS Lookup Servers** field, type the IP address of your DNS Lookup Server.  
The DNS Lookup Server allows you to use IP addresses, host names, or fully-qualified domain names (FQDNs) to access other network objects.
15. In the **DNS Search Domain** field, type the name of your search domain.  
The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
16. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.
17. Click the **Save** button to save your configuration.

## Manual license activation and initial configuration

Before you can activate your F5 product license, you must configure your virtual environment or your platform, define a management IP address, and obtain the base registration key.

If the BIG-IQ™ system is not connected to the public internet, use this procedure to activate its license.

1. Using a browser on which you have configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you specified for device management:  
`https://<management_IP_address>`
2. Log in to the BIG-IQ system with the default user name `admin` and password `admin`.
3. At the top of the screen, click **System Overview**.
4. In the Setup panel, click the IP address of the BIG-IQ system.  
The panel expands to display additional properties.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Key** field, paste any additional license key you have.  
The options are: `Cloud` or `Security`.
7. For the Activation method setting, select **Manual** and click the **Activate** button.  
The BIG-IQ system refreshes and displays the dossier in the **Dossier** field.
8. Copy the displayed dossier and transfer it to a system connected to the internet and navigate to the F5 Licensing Server at `https://activate.f5.com/license/`.
9. Paste the dossier in the **Enter your dossier** text box, or click the **Browse** button to locate it on the system, and click the **Next** button.
10. Copy or save the activation key and transfer it to the BIG-IQ system.
11. Paste the activation key in the **License** field and click the **Activate** button.
12. The End User License Agreement (EULA) displays.  
When you click **Accept**, the screen refreshes to display the license details.
13. Click the **Properties** tab.
14. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.  
The FQDN can consist of letters, numbers, and/or the characters underscore (`_`), dash (`-`), or period (`.`).
15. In the **Self IP Address** field, type the self IP address of your internal VLAN.  
The self IP address must be in Classless InterDomain Routing (CIDR) format. For example:  
`10.10.10.10/24`.
16. To add an additional self IP address, click the **+** sign and in the new **Self IP Address** field that the system creates, edit the duplicated self IP address to reflect the additional self IP address that you want to add.

Once you save this self IP address, you cannot change it.

17. Click the **Services** tab.
18. In the **DNS Lookup Servers** field, type the IP address of your DNS Lookup Server.  
The DNS Lookup Server allows you to use IP addresses, host names, or fully-qualified domain names (FQDNs) to access other network objects.
19. In the **DNS Search Domain** field, type the name of your search domain.  
The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
20. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.
21. Click the **Save** button to save your configuration.

## About installing required BIG-IQ system components on managed BIG-IP systems

---

You must install specific components required by the BIG-IQ™ system on each BIG-IP® device you want to manage. To install these components, you run a series of commands from the command line.

### Installing required BIG-IQ components on BIG-IP systems

You can perform this task only after you have licensed and installed the BIG-IQ™ system and at least one BIG-IP® device running version 11.3 or later.

This task installs, onto your managed BIG-IP devices, a REST framework that supports the required Java-based management services. You must perform this installation task each time you discover a new device.

---

**Important:** *When you run this installation script, the traffic management interface (TMM) on each BIG-IP device restarts. It is important that, before you run this script, you verify that no critical network traffic is targeted to the BIG-IP devices.*

---

1. Log in to the BIG-IQ system terminal as the root user.
2. Establish SSH trust between the BIG-IQ system and the managed BIG-IP device.  

```
ssh-copy-id root@<BIG-IP Management IP Address>
```

This step is optional. If you do not establish trust, you will be required to provide the BIG-IP system's root password multiple times.
3. Navigate to the folder in which the files reside.  

```
cd /usr/lib/dco/packages/upd-adc
```
4. Run the installation script.
  - For devices installed in an Amazon EC2 environment: 

```
./update_bigip.sh -a admin -p <password> -i /<path_to_PEM_file> <BIG-IP Management IP Address>
```
  - For devices installed in any other environment: 

```
./update_bigip.sh -a admin -p <password> <BIG-IP Management IP Address>
```

Where *<password>* is the administrator password for the BIG-IP device.

5. Revoke SSH trust between the BIG-IQ system and the managed BIG-IP device.

```
ssh-keygen -R <BIG-IP Management IP address>
```

This step is not required if you did not establish trust in step 2.

## About a high availability configuration

---

You can ensure that you always have access to your managed BIG-IP devices by installing two BIG-IQ systems in an active/standby high availability (HA) configuration. If the active BIG-IQ system in the HA configuration fails, the standby peer becomes active allowing you to continue to manage devices.

### Configuring a high availability pair

After you install and license two BIG-IQ systems, you can configure them in an active/standby pair.

Having a high availability pair of BIG-IQ systems configured makes it possible for you to always have access to BIG-IP devices in your network. The BIG-IQ systems in a high availability pair synchronize their configurations every 30 minutes. Configuring a high availability pair is optional.

---

**Important:** For a pair of BIG-IQ systems in a high availability configuration to synchronize properly, they must each be running the same version of BIG-IQ. Perform these steps on the active BIG-IQ system.

---

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **System Overview**.
3. Click the **High Availability** tab.
4. In the **Peer IP Address** field, type the self IP address (on the internal VLAN) of the peer BIG-IQ system. Do not use the management IP address of the peer.
5. In the **User Name** and **Password** fields, type the administrative user name and password for the BIG-IQ system you want to add as a peer.
6. Click the **Save** button to save your configuration.

If discovery fails, a delete button displays. Verify the correct self IP address and credentials. Then click the delete button to remove the incorrect information, and re-enter the self IP address, user name, and password.

### Replacing a peer in a high availability configuration

To change the peer BIG-IQ system that you specified in a high availability pair, you must delete the current peer system, and specify a new peer.

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **System Overview**.
3. Click the **High Availability** tab.
4. Click the **Delete** button.
5. Repeat steps 1-4 on the peer system.
6. Log back in to the first BIG-IQ system.
7. At the top of the screen, click **System Overview**.
8. Click the **High Availability** tab.

9. In the **Peer IP Address** field, type the self IP address (on the internal VLAN) of the peer BIG-IQ system.  
Do not use the management IP address of the peer.
10. In the **User Name** and **Password** fields, type the administrative user name and password for the BIG-IQ system you want to add as a peer.
11. Click the **Save** button to save your configuration.

The active BIG-IQ system discovers its peer and displays its status.

If discovery fails, a delete button displays. Verify the correct self IP address and credentials. Then click the delete button to remove the incorrect information, and re-enter the self IP address, user name, and password.

### Forcing an active peer BIG-IQ system to standby mode

In the event that both BIG-IQ systems in a high availability pair become active, the system displays a warning message at the top of every screen. If this occurs, you can use this procedure to put the standby system back into standby mode.

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **System Overview**.
3. Click the **High Availability** tab.
4. Click the **Force Standby** button.  
The BIG-IQ system is forced into standby mode.

### About default passwords for pre-defined users

---

When you initially license the BIG-IQ™ system, it creates the following administrative roles with a default password.

- administrator
- root

### Changing the default password for the administrator user

You must specify the management IP address settings for the BIG-IQ™ system to prompt the system automatically create the administrator user.

After you initially install and configure the BIG-IQ system, it is important to change the password for the administrator password user from the default password, `admin`.

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **Users**.
3. On the Users panel, click **Admin User**
4. Click the properties gear.  
The screen refreshes to display the properties for this user.
5. In the **Password** and **Confirm Password** fields, type a new password.
6. Click the **Save** button.

## Changing the default password for the root user

You must specify the management IP address settings for the BIG-IQ™ system to prompt the system automatically create the root user.

After you initially install and configure the BIG-IQ system, it is important to change the password for the root user from the default password, `default`.

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **Users**.
3. On the Users panel, click **root** user.
4. Click the properties gear.  
The screen refreshes to display the properties for this user.
5. In the **Password** and **Confirm Password** fields, type a new password.
6. Click the **Save** button.





# Index

## A

- active active mode
  - for high availability pair 30
- admin, See administrator
- administrator user
  - changing password for 30
  - default password 30
- administrator user password
  - changing 30
- Admin password
  - changing 20
  - setting 20
- AMI
  - launching new 18
- authentication 16

## B

- base registration key
  - about 27
- BIG-IP devices
  - and BIG-IQ components required for management 28
- BIG-IQ Cloud
  - about 24
  - finding documentation for 24
- BIG-IQ Security
  - about 24
  - finding documentation for 24
- BIG-IQ system
  - about 24
- BIG-IQ system components
  - installing on BIG-IP devices 28
  - installing on managed BIG-IP devices 28

## C

- configuration
  - initial setup 26–27
- CPU
  - and guest definition 14

## D

- deployment
  - for EC2 16
- deployment overview 16
- documentation, finding 24
- dossier
  - providing 26–27

## E

- EC2 AMI
  - launching 18
- EC2 VPC
  - creating 16

- Elastic Compute Cloud
  - and compatible versions 14
- environment, for guest 14
- external subnet
  - adding route for accessibility 18

## F

- failover 29
  - See also high availability

## G

- guest environment 14
- guides, finding 24

## H

- high availability
  - 29
  - configuring 29
  - replacing a peer 29
- high availability configuration
  - about 29
- high availability pair
  - forcing peer to standby 30
- hypervisor, See guest environment.
- hypervisor guest definition 14

## K

- key pairs
  - creating 16

## L

- license
  - activating automatically 26
  - activating manually 27
- licensing
  - BIG-IQ system 26–27

## M

- managed devices
  - and BIG-IQ system components required 28
- management port
  - making it accessible 20
- manuals, finding 24
- maximum allowed throughput rate 14

## N

- network
  - incorporating BIG-IQ systems 26
  - port 443 26

## Index

network interface  
adding 20

## P

package endpoint  
installing 28  
password  
changing Admin 20  
changing for administrator user 30  
changing for root user 31  
setting Admin 20  
pre-defined users  
administrator  
root 30  
product license 14

## R

release notes, finding 24  
required port, for network communication 26  
root user  
changing password for 31  
default password 30  
root user password  
changing 31  
route  
adding for external subnet 18

RPM endpoint  
installing 28

## S

security groups  
creating 17  
Setup utility 16  
subnet  
adding additional 17  
adding route for external 18

## T

task list  
for deploying on EC2 16  
for deploying on virtual machine 16

## V

virtual configuration, and hypervisor guest definition 14  
virtual machine  
launching new 18  
virtual machine settings 14  
virtual private cloud  
creating 16