

BIG-IQ[®] Cloud: Cloud Administration

Version 4.4



Table of Contents

Legal Notices.....	5
Acknowledgments.....	7
 Chapter 1: BIG-IQ Cloud Overview.....	 15
Additional resources and documentation for BIG-IQ Cloud.....	16
About the BIG-IQ system user interface.....	16
Filtering for associated objects.....	16
Customizing panel order.....	17
 Chapter 2: Configuring BIG-IQ High Availability.....	 19
About a high availability active-active cluster.....	20
Configuring BIG-IQ Cloud in an active-active high availability configuration.....	20
 Chapter 3: Device Resource Management.....	 21
About device discovery and management.....	22
Discovering BIG-IP devices in your network.....	22
Viewing and exporting device inventory details.....	23
 Chapter 4: License Pools.....	 25
About pool licenses.....	26
Automatically activating a pool license.....	26
Manually activating a pool license.....	26
Assigning a pool license to a BIG-IP device.....	27
 Chapter 5: Integrating Amazon Web Services	 29
About Amazon Web Services (AWS) integration.....	30
Network requirements for AWS integration communication	30
Creating an Amazon Identity and Access Management (IAM) user account.....	30
Creating a Virtual Private Cloud.....	31
Launching a virtual server with an Amazon Machine Image (AMI).....	32
Configuring an EC2 cloud connector.....	33
Creating a BIG-IP VE version 11.5 or later in the Amazon EC2 cloud.....	34
Creating a BIG-IP VE version 11.3 or 11.4 in the Amazon EC2 cloud.....	35
Creating a customized application template.....	36
Deploying applications.....	37
Setting up tenant access using IAM.....	38
Viewing activity for cloud resources.....	39
 Chapter 6: Integrating OpenStack	 41

About OpenStack integration.....	42
Network requirements for communication with OpenStack cloud services	42
OpenStack Compute edits required to use BIG-IP VE systems.....	43
Discovering devices located in the OpenStack cloud.....	43
Associating an OpenStack connector with devices.....	44
Chapter 7: Integrating VMware	45
Configuring VMware NSX 6.1 for BIG-IP Cloud.....	46
Network requirements for communication with VMware cloud services	47
Discovering devices located in the VMware cloud.....	47
Create a connection between the BIG-IP device and NSX.....	47
About VMware NSX version 6.0 integration.....	48
Network requirements for communication with VMware cloud services	48
Discovering devices located in the VMware cloud.....	48
Associating a VMware cloud connector with a device.....	49
About vCloud Director integration	50
Network requirements for communication with VMware cloud services	50
Discovering devices located in the VMware cloud.....	50
Associating a VMware cloud connector with a device.....	51
Chapter 8: Local Cloud Integration.....	53
About using a local cloud source.....	54
Discovering BIG-IP devices in your network.....	54
Associating a local cloud connector with a device.....	54
Chapter 9: Cloud Tenant Management.....	57
About creating cloud tenants	58
Creating a tenant.....	58
Creating a cloud user.....	58
Associating a user with a tenant's role.....	59
Chapter 10: iApps Application Template Customization.....	61
About customizing iApp application templates.....	62
Creating a customized application template.....	62
Chapter 11: Glossary.....	65
BIG-IP Cloud terminology.....	66

Legal Notices

Publication Date

This document was published on September 8, 2014.

Publication Number

MAN-0501-02

Copyright

Copyright © 2014, 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, Application Acceleration Manager, Application Security Manager, APM, 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, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, iCall, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, LineRate, LineRate Systems [DESIGN], Local Traffic Manager, LROS, LTM, Message Security Manager, MobileSafe, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, SDAC (except in Japan), SDC, Signalling Delivery Controller, Solutions for an application world, Software Designed Applications Services, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix (except Germany), Traffix [DESIGN] (except Germany), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, 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>

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.

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), which is protected under the GNU Public License.

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

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.

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 (ey@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) 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 jQuery Foundation and other contributors, distributed under the MIT License. Copyright ©2014 jQuery Foundation and other contributors (<http://jquery.com/>).

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.

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 Mojarra project. Source code for the Mojarra software may be obtained at <https://jaserverfaces.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 gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

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

This product includes isc-dhcp software. Copyright © 2004-2013 by Internet Systems Consortium, Inc. ("ISC"); Copyright © 1995-2003 by Internet Software Consortium.

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

This product includes jQuery Sparklines software, developed by Gareth Watts, and distributed under the new BSD license.

This product includes jsdiff software, developed by Chas Emerick, and distributed under the BSD license.

This product includes winston software, copyright © 2010, by Charlie Robbins.

This product includes Q software developed by Kristopher Michael Kowal, and distributed under the MIT license. Copyright © 2009-2013 Kristopher Michael Kowal.

This product includes SlickGrid software developed by Michael Liebman, and distributed under the MIT license.

This product includes JCraft Jsch software developed by Atsuhiko Yamanaka, copyright © 2002-2012 Atsuhiko Yamanaka, JCraft, Inc. All rights reserved.

This product includes DP_DateExtensions software developed by Jim Davis, Copyright © 1996-2004, The Depressed Press of Boston (depressedpres.com). 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.
- Neither the name of the DEPRESSED PRESS OF BOSTON (DEPRESSEDPRESS.COM) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS 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 THE COPYRIGHT OWNER OR CONTRIBUTORS 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.

All code not authored by the Depressed Press is attributed (where possible) to its rightful owners/authors, used with permission and should be assumed to be under copyright restrictions as well.

This product includes Angular software developed by Google, Inc., <http://angularjs.org>, copyright © 2010-2012 Google, Inc., and distributed under the MIT license.

This product includes node.js software, copyright © Joyent, Inc. and other Node contributors. All rights reserved.

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 product includes the epoxy.js library for backbone, copyright © 2012-2013 Greg MacWilliam. (<http://epoxyjs.org>)

This product includes Javamail software, copyright ©1997-2013 Oracle and/or its affiliates, all rights reserved; and copyright © 2009-2013 Jason Mehrens, all rights reserved. This software is distributed under the GPLv2 license.

This product includes underscore software, copyright © 2009-2014 Jeremy Ashkenas, DocumentCloud, and Investigative Reporters & Editors.

This product includes node-static software, copyright © 2010-2014 Alexis Sellier.

This product includes jxrlib software, copyright ©2009 Microsoft Corp. All rights reserved. Distributed under the new BSD license.

This product includes node-uuid software, copyright © 2010-2012, Robert Kieffer, and distributed under the MIT license.

This product includes opensv software, which is distributed under the Apache 2.0 license.

This product includes owasp-jave-encoder software, copyright © 2014, Jeff Ichnowski, and distributed under the New BSD license.

This product includes cookies software, copyright © 2014, Jed Schmidt, <http://jed.is/>, and distributed under the MIT license.

This product includes node-fastcgi software, copyright © 2013, Fabio Massaioli, and distributed under the MIT license.

This product includes socket.io software, copyright © 2013, Guillermo Rauch, and distributed under the MIT license.

This product includes node-querystring software, copyright © 2012. Irakli Gozalishvili. All rights reserved.

This product includes TinyRadius software, copyright © 1991, 1999 Free Software Foundation, Inc., and distributed under the GNU Lesser GPL version 2.1 license.

Chapter 1

BIG-IQ Cloud Overview

- *Additional resources and documentation for BIG-IQ Cloud*
- *About the BIG-IQ system user interface*
- *Filtering for associated objects*
- *Customizing panel order*

Additional resources and documentation for BIG-IQ Cloud

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® Virtual Edition Setup</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.
<i>BIG-IQ® Systems: Licensing and Initial Configuration</i>	This guide provides the network administrator with basic BIG-IQ system concepts and describes the tasks required to license and set up the BIG-IQ system in their network.
<i>BIG-IQ® Cloud: Cloud Administration</i>	This guide contains information to help a cloud administrator manage cloud resources, devices, applications, and tenants (users).
<i>BIG-IQ® Cloud: Tenant User Guide</i>	This guide contains information to help tenants manage applications.
<i>BIG-IQ® Device: Device Management</i>	This guide provides details about how to deploy software images, licenses, and configurations to managed BIG-IP devices.
<i>Platform Guide: BIG-IQ® 7000 Series</i>	This guide provides information about setting up and managing the BIG-IQ 7000 hardware platform.
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.

About the BIG-IQ system user interface

The BIG-IQ® system interface is composed of panels. Each panel contains objects that correspond with a BIG-IQ system feature. Depending on the number of panels and the resolution of your screen, some panels are collapsed on either side of the screen. You can cursor over the collapsed panels to locate the one you want, and click the panel to open. To associate items from different panels, click on an object, and drag and drop it onto the object to which you want to associate it.

Filtering for associated objects

The BIG-IQ system helps you easily see an object's relationship to another object, even if the objects are in different panels.

1. In a panel, click the object on which you want to filter.
The selected object name displays in the Filter field, and the screen refreshes to display unassociated objects as unavailable.

2. To further filter the objects displayed, you can type one additional object in the Filter field, and click the **Apply** button.
3. To display only those objects associated with the object you selected, click the **Apply** button. The screen refreshes and the objects previously displayed in a gray font do not appear. Only objects associated with the object you click display, and the object you selected displays below the Filter field.
4. To remove a filter, click the **x** icon next to the object that you want to remove, below the Filter field.

Customizing panel order

You can customize the BIG-IQ system interface by reordering the panels.

1. Click the header of a panel and drag it to a new location, then release the mouse button. The panel displays in the new location.
2. Repeat step 1 until you are satisfied with the order of the panels.

Chapter

2

Configuring BIG-IQ High Availability

- *About a high availability active-active cluster*
- *Configuring BIG-IQ Cloud in an active-active high availability configuration*

About a high availability active-active cluster

You can ensure that you always have access to managed BIG-IP® devices by installing two or more BIG-IQ® systems in an active-active, high availability (HA) configuration. Any configuration change that occurs on one BIG-IQ system is immediately synchronized with its peer devices. If a BIG-IQ® system in an active-active HA configuration fails, a peer BIG-IQ system takes over the device management.

Configuring BIG-IQ Cloud in an active-active high availability configuration

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the BIG-IQ HA header, and click the + icon when it appears.
3. In the **Peer IP Address** field, type the self IP address (on the internal VLAN) of the peer system.
Do not use the management IP address of the peer system.
4. In the **User name** and **Password** fields, type the administrative user name and password for the system.
5. Click the **Add** button.

If discovery of the newly configured BIG-IQ system 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-type the self IP address, user name, and password.

Chapter

3

Device Resource Management

- *About device discovery and management* |

About device discovery and management

You use BIG-IQ® Device to centrally manage resources located on BIG-IP® devices in your local network, in a public cloud like Amazon EC2, or in a combination of both.

The first step to managing devices is making BIG-IQ Device aware of them through the discovery process. To discover a device, you provide BIG-IQ Device the device IP address, user name, and password. Alternatively, you can upload a CSV file to discover a large number of devices. When you discover a device you place it into a group. These groups help you organize devices with similar features, like those in a particular department or running a certain software version.

After you discover devices, you can view and export inventory details about those devices for easy asset management.

Discovering BIG-IP devices in your network

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

You can discover a device by providing the BIG-IQ system with the device's IP address, user name, and password.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. For devices on the same subnet as the BIG-IQ system, in the **IP Address** field, type the IP address of the device.

You cannot discover a BIG-IP device using its management IP address.

3. When the BIG-IQ system and the BIG-IP device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 - b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

Where `<route name>` is a user-provided name to identify the new route, and `<x.x.x.x>` is the IP address of the default gateway for the internal network.

4. To change the root user name, type a new name in the **Root User Name** field.
5. Type a password for the root user in the **Root Password** field.
6. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
7. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.

8. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

Viewing and exporting device inventory details

You can view detailed data about the managed devices in your network. Information includes associated IP addresses, platform type, license details, software version, and so forth. In addition to viewing this information, you can also export it to a CSV file and edit the data as required to create reports for asset management.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Configuration**.
3. In the Devices panel, click the gear icon next to the device group you want to view, and then click **Inventory**.

The panel expands to display device details.

4. To export the data to a CSV file, click the **Export** button.
You can modify the report as required in Microsoft Excel.

Chapter 4

License Pools

- *About pool licenses*
-

About pool licenses

Pool licenses are purchased for a fixed number of devices, but are not permanently tied to a specific device. As resource demands change, you can revoke and grant those licenses to other resources as required.

Automatically activating a pool license

You must have a base registration key before you can activate the license pool.

If the resources you are licensing are connected to the public internet, you can automatically activate the license pool.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Licenses panel, click the + sign when it appears, and then click **Add New Pool License**. The panel expands to display New License properties.
3. In the **License Name** field, type a name for this license.
4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
5. In the **Add-on Keys** field, paste any additional license key you have.
6. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button. The End User License Agreement (EULA) displays.
7. To accept the EULA, click the **Accept** button.

You can now assign this license to a BIG-IP® device.

Manually activating a pool license

You must have a base registration key before you can activate the pool license.

If the BIG-IQ® Cloud you are licensing is not connected to the public internet, you can still activate the pool license manually.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the **Pool Licenses** panel and click the + sign when it appears.
3. In the **License Name** field, type a name for this license.
4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
5. In the **Add-on Keys** field, paste any additional license key you have.
6. For the **Activation** method setting, select **Manual** and click the **Generate Dossier** button. The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
7. Copy the text displayed in the Device Dossier field, and click the **Access F5 manual activation web portal** link.
Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.
8. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button. The Accept User Legal Agreement displays.
9. To accept the EULA, click the **Accept** button.
10. Copy the license file from the F5 license activation portal to BIG-IQ Cloud.

You can now assign this license to a BIG-IP® device.

Assigning a pool license to a BIG-IP device

Before you can assign a pool license to a device, you must activate a pool license on BIG-IQ® Device and discover at least one device.

Assigning a pool license to a BIG-IP VE device provides you with the flexibility to initiate and remove licenses as needed.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Deployment panel, click the + sign when it appears, and click **New Deployment Job**.
4. In the **Name** field, type a name for this license deployment.
5. From the **Type of Job** list, select **Modify a Virtual Device**.
6. From the **Device** list, select the device you want to license.
7. From the **Licensing** list, select **Use a Pool License**.
8. From the **License Pool** list, select the pool license you want to use for this device.
9. Click the **Deploy** button.
10. To confirm that the license was successfully deployed, click the license pool from which you deployed the license and click the Assignments tab.
The device you licensed displays with the license status and time when the BIG-IQ system last contacted it.

Chapter

5

Integrating Amazon Web Services

- *About Amazon Web Services (AWS) integration*
- *Creating an Amazon Identity and Access Management (IAM) user account*
- *Creating a Virtual Private Cloud*
- *Launching a virtual server with an Amazon Machine Image (AMI)*
- *Configuring an EC2 cloud connector*
- *Creating a BIG-IP VE version 11.5 or later in the Amazon EC2 cloud*
- *Creating a BIG-IP VE version 11.3 or 11.4 in the Amazon EC2 cloud*
- *Creating a customized application template*
- *Deploying applications*
- *Setting up tenant access using IAM*
- *Viewing activity for cloud resources*

About Amazon Web Services (AWS) integration

BIG-IQ® Cloud provides you with the tools to manage Amazon EC2 and CloudWatch resources required to perform application delivery. Management tasks include discovering and creating BIG-IP® VE virtual machines located in Amazon Virtual Private Cloud (VPC), application pool servers, and deploying applications. You can use these features to accommodate application traffic fluctuations by periodically adding and retracting devices and application servers, as needed. Additionally, you can provide tenants access to self-deployable iApps® through Amazon EC2 integration.

To provide access to these services for Amazon EC2 tenants, you configure communication between Amazon EC2 products, and BIG-IQ Cloud. Then, you associate a Amazon EC2 cloud connector with a device, and create a catalog entry for a corresponding Amazon EC2 service profile. The tenants to whom you give access to the catalog entry see it in their applications panel. From there, they can use it to self-deploy their own iApps.

Network requirements for AWS integration communication

BIG-IQ Cloud integrates with three different Amazon Web Services: Amazon EC2, Amazon CloudWatch, and BIG-IP Virtual Edition deployed in managed Amazon Virtual Private Cloud (VPC).

For proper communication to devices located in an Amazon web service, BIG-IQ® Cloud you must configure an outbound self IP address to DNS and NTP, and you must define a network route between the BIG-IQ Cloud internal VLAN and the public Internet, or the Amazon web services endpoint. For specific instructions, refer to *BIG-IQ® System: Licensing and Initial Setup* and your Amazon documentation .

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 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.

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.
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 of 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.

Configuring an EC2 cloud connector

Before you can create an EC2 cloud connector, you must first discover devices in the Amazon EC2 cloud and create an Amazon Identity and Access Management (IAM) user account. If you want BIG-IQ Cloud to automatically provision additional BIG-IP VE servers and devices for your tenant when more resources are needed, you must also purchase and activate a license pool to associate with this connector.

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Connectors header and click the + icon when it appears.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
4. From the **Cloud Provider** list, select **Amazon EC2**.
5. In the **Region Endpoint** field, type the entry point URL.
For example, `ec2.us-east-1.amazonaws.com` is the region end point for the Amazon EC2 US East (Northern Virginia) Region. Refer to the AWS documentation for a list of all regional end points at http://docs.aws.amazon.com/general/latest/gr/rande.html#ec2_region
6. In the **Key ID** and **Secret Key** fields, type the credentials of the BIG-IQ-Connector IAM user.
For security purposes, it is important to specify a user that has Amazon EC2 Full Control Access.
7. In the **Availability Zone** field, type the location of the region in which the instances are located.
For example, type `us-west-2c` for the availability zone for Oregon state.
8. In the **Virtual Private Cloud** field, you may type the identification for the EC2 Virtual Private Cloud (VCP) network topology inside the Availability Zone.
This step is optional. If you do not specify the identification for a VCP, BIG-IQ Cloud uses the first one it discovers in the Availability Zone.
9. Click the arrow next to **Device & Server Provisioning** to display associated options.
The screen refreshes to display the options.
10. To prompt BIG-IQ Cloud to automatically provision additional BIG-IP VE devices when more resources are needed for application traffic, for the **Device Elasticity** setting, select **Enable**.
11. From the **Device License** list, select a rate at which you want Amazon to direct-bill for additional devices, or select a license pool from which to grant a license.
You must activate a license pool before you can select it.

12. To automatically prompt BIG-IQ Cloud to provision additional servers when more resources are needed to manage an influx in application traffic, for the **Server Elasticity** setting, select **Enable**.
13. Review the network settings populated when you selected a connector, verifying that the proper CIDR blocks display for management, external, and internal.
14. Click the **Save** button.
15. If the system discovered devices, you must expand the device's properties panel, and provide the device's credentials to finalize the discovery process.
16. Review the network settings populated when you selected a connector, verifying that the proper CIDR blocks display for management, external, and internal.

You now create a device associated with this EC2 cloud connector.

Creating a BIG-IP VE version 11.5 or later in the Amazon EC2 cloud

After you license and perform the initial configuration for the BIG-IQ system, you can create devices in the Amazon EC2 cloud. For proper communication, you must configure a route between each instance to the BIG-IQ system. If you do not specify the required network communication route between the devices, then creation fails.

Before you perform this task you must first open specific ports on your EC2 AMI BIG-IQ instance and on any associated EC2 BIG-IP instances. To open these ports, you need additional security group rules in your `allow-only-ssh-https-ping` security group, and you need to associate these rules with the management interface.

You need to create three rules: two outbound rules for the BIG-IQ instance, and one inbound rule for the BIG-IP instance.

Group Name	Group Description	Rule Name	Source	Port
allow-only-ssh-https-ping	Allow only SSH, HTTPS, or PING	Outbound SSH	0.0.0.0/0	22 (SSH)
		Outbound HTTPS	443 0.0.0.0/0	443 (HTTPS)
		Inbound HTTPS	0.0.0.0/0	443 (HTTPS)

To create a BIG-IP VE instance in Amazon EC2 cloud, you associate the EC2 Cloud connector you configured with that device.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Devices header, and click the + icon when it appears.
3. Select the **Create a Device** option.
4. From the Cloud Connector list, select the EC2 cloud connector you created.
5. From the **Device Image** list, select the AMI you created for this device.
6. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you do update the framework or cancel the task.

7. To prompt BIG-IQ Cloud to assign the default user admin and a randomly-selected password, select the **Use "admin"** check box.
8. To assign a specific user name and password, deselect the **Use "admin"** check box. The screen refreshes to display additional settings.
9. In the **User Name** and **Password** fields, type a user name and password for the user of this devices.
10. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

Creating a BIG-IP VE version 11.3 or 11.4 in the Amazon EC2 cloud

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 11.4.

Before you perform this task you must first open specific ports on your EC2 AMI BIG-IQ instance and on any associated EC2 BIG-IP instances. To open these ports, you need additional security group rules in your `allow-only-ssh-https-ping` security group, and you need to associate these rules with the management interface.

You need to create three rules: two outbound rules for the BIG-IQ instance, and one inbound rule for the BIG-IP instance.

Group Name	Group Description	Rule Name	Source	Port
allow-only-ssh-https-ping	Allow only SSH, HTTPS, or PING	Outbound SSH	0.0.0.0/0	22 (SSH)
		Outbound HTTPS	443 0.0.0.0/0	443 (HTTPS)
		Inbound HTTPS	0.0.0.0/0	443 (HTTPS)

To create a BIG-IP VE version 11.3 or 11.4 instance in Amazon EC2 cloud, you must update the BIG-IP VE REST framework that supports the required BIG-IQ Cloud Java-based management services, and then associate the EC2 Cloud connector you configured with that device.

Warning: When you perform this task, the traffic management interface (TMM) on the BIG-IP VE restarts. Before you perform this task, verify that no critical network traffic is targeted to the BIG-IP VE device.

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.

```
root@<BIG-IP Management IP address>grep -v '<username>@<computername>'
/root/.ssh/authorized_keys > /tmp/authorized_keys.tmp; mv -f
/tmp/authorized_keys.tmp /root/.ssh/authorized_keys
```

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

6. Log in to BIG-IQ® Cloud with your administrator user name and password.
7. In the Device panel, click the gear icon next to the legacy device with a yellow triangle next to it and displaying the message, Discovery is incomplete.
8. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
9. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you do update the framework or cancel the task.

10. Click the **Save** button.

Important: Before you begin using this BIG-IQ system in a production capacity, depending on your security policies, you will likely want to stop using the security group rules that you added as prerequisite to this task.

Creating a customized application template

Before you can customize an application template for a tenant, you must discover at least one F5 device that contains iApps® templates.

As a cloud provider, you modify iApps templates to customize network settings, levels of services, and so forth, for tenants. You can create variations of the same application, offering different types of access (LAN or WAN), or providing a specific limit of connections.

Note: Once you customize and save an application as a catalog entry, you cannot modify it.

1. Hover over the Catalog header, and click the + icon when it appears.
The panel expands to display the Catalog properties.
2. In the **Name** field, type a name for this new application.
3. From the **Application Type** list, select an application.
4. Unless you want to restrict this application template to a specific cloud connector, leave the **Cloud Connector** setting as **Tenant Selectable** so tenants are allowed to select the appropriate cloud connector when they deploy this application.
5. If the **Application Tiers** settings are displayed (expanded), select the options that match the properties for this application; otherwise, keep the default settings.

Important: If you must specify the options for these settings, select the **Tenant Editable** check box for the virtual server and pool members.

6. To allow cloud tenants to specify certificates with SSL encryption when self-deploying applications, select options from the **SSL Cert** and **SSL Key** lists.
BIG-IQ Cloud uses these options to provide the appropriate certificate and key when the tenant self-deploys this application to a BIG-IP® device. These options are not available for all application templates.
7. Finish making modifications by specifying the Application Properties and Customize Application Template variables.
To allow a tenant to modify a particular setting, select the **Tenant Editable** check box for that setting. For further details about template variables and settings, refer to the *BIG-IP® iApps® Developer's Guide*.
8. Click the **Save** button.
You can now send the cloud IP addresses to the tenant and use this IP address range in configuring server tiers and pool members, within certain application services. The tenant can self-deploy the application from the catalog.

The customized application displays as an entry in the catalog.

Deploying applications

Before you can deploy and use an application, your cloud service provider must add you as a user and a tenant, and associate you with at least one cloud connector.

When a cloud administrator adds you as a cloud tenant user, they contact you with the details about the resources to which you have access. These resources are provided to you in the form of an application template. As a cloud tenant user, you can customize these application templates and deploy them.

1. Log in to the BIG-IQ Cloud with your tenant user name and password.
2. Hover over the Applications header, and click the + icon when it appears.
3. In the **Name** field, type a name for this new application.
4. From the **Application Type** list, select an application.
5. From the **Cloud Connector** list, select the cloud connector associated with where you want to deploy your application.
A cloud connector is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.
6. To configure BIG-IQ Cloud to automatically provision additional resources when traffic to your application increases, select **Enable** from the **Server Elasticity** list and specify the settings for the server elasticity options that display.
This option is available only for the EC2 connector. For automatic server provisioning to work, your cloud service provider must have enabled the **Server Elasticity** setting for this EC2 connector.
 - a) From the **Node Image** list, select the image from which to create new application servers when capacity is met and additional servers are required.
 - b) In the **Min. # of Servers** field, type the minimum number of application servers you want running at any given time.
 - c) In the **Max. # of Servers** field, type the maximum number of application servers you want running when additional servers are required.
 - d) From the **Monitor By** list, select the category associated with the statistic on which you want to base the threshold value.
 - e) For the **When** setting, select a specific statistic, the associated relational operator, and a type a number in the field for the threshold.

Base the threshold on the maximum amount of traffic a server can reasonably process for this application to ensure that BIG-IQ Cloud adds additional resources at the right time.

- f) In the **Add Servers** field, type the number of application servers you want BIG-IQ Cloud to add when this threshold is met.
7. To define a new SSL certificate and private key for this application, for the **SSL Certificate Options**, paste the PEM (CRT or CER) text representation of the certificate and private key.
The SSL certificate and private key must be unbundled Base64 encoded ASCII text with PEM header and footer.
This option is not available for all applications.
8. Alternatively, select the **Use Existing** option to use a SSL certificate and private key already stored on the device.
9. You can further customize this application by specifying an IP address for the virtual server and adding pool hosts.
If your cloud service provider assigned IP addresses for the **Servers**, **Pool Hosts**, and **Pool Members** for this application, the addresses display. If these addresses were specified as not editable, you cannot change them.
10. When you are finished, click the **Deploy** button located at the top of the New Application panel.

You can now use this new application, and any application server associated with this new application displays in the Server panel.

Setting up tenant access using IAM

You might want your tenants to have access to all or part of the EC2 cloud you are provisioning so that they are able to configure resources required by their applications. You can provide full access by simply providing the account information (user name and password) that you created previously. More typically, you can provide more limited access by setting up separate user accounts for the tenant, and then configuring the access for those users as best suits your needs.

Important: *If you decide to grant full tenant access to the IAM account, bear in mind that restricting this account to a single tenant becomes even more prudent.*

The following step-sequence provides an outline of the tasks you perform using the AWS EC2 user interface. For the most current instructions for performing each of these tasks, refer to the Amazon Web Services EC2 Management Console web site <https://console.aws.amazon.com/ec2/v2/home>.

1. Log in to the AWS IAM console.
2. Create a user role to encapsulate relevant permissions for this tenant.
If a user needs to create key pairs, make certain that they have sufficient permissions.
3. Configure password policies for this tenant.
4. Create user accounts and set passwords for this tenant.
5. Create the user(s).
6. Specify the IAM AWS Management URL that you will provide to your tenants so that they can log in to this IAM account and directly manage their resources.

Viewing activity for cloud resources

Before you can view dynamic cloud resource activity, you must have an EC2 cloud connector with the **Device Elasticity** setting enabled.

Viewing activity for dynamic cloud resources gives you insight into how cloud resources are expanding to address increased traffic to applications.

1. To view the resource associated with a particular activity, click the activity located on the Activities panel.
The associated objects are highlighted in the relevant panels.
2. To view specific activity details, place your cursor on an activity.
A popup window opens to display further details about the selected activity.

Chapter

6

Integrating OpenStack

- *About OpenStack integration*
- *Network requirements for communication with OpenStack cloud services*
- *OpenStack Compute edits required to use BIG-IP VE systems*
- *Discovering devices located in the OpenStack cloud*
- *Associating an OpenStack connector with devices*

About OpenStack integration

BIG-IQ® Cloud provides you with the tools to manage OpenStack versions 2013.1 (Grizzly) and 2013.2 (Havana) resources required to run applications. Management tasks include discovering BIG-IP® VE virtual machines and discovering, creating, starting, and stopping OpenStack application servers running in the private cloud. You can use this feature to accommodate seasonal traffic fluctuations by periodically adding and retracting devices and application servers as needed. Additionally, you can provide tenants access to self-deployable iApps® through OpenStack integration.

To provide access to these services for OpenStack tenants, you configure communication between OpenStack products, and BIG-IQ Cloud. Then, you associate an OpenStack cloud connector with a device, and create a catalog entry for a corresponding OpenStack service profile. The tenants to whom you give access to the catalog entry see it in their applications panel. From there, they can use it to self-deploy their own iApps.

Network requirements for communication with OpenStack cloud services

Before you can manage devices residing in an OpenStack private cloud, you must establish proper communication between the BIG-IQ® Cloud and the OpenStack controller node. Generally, this means defining a network route between the BIG-IQ Cloud internal VLAN and the public Internet, or the OpenStack private cloud endpoint.

The BIG-IQ Cloud connector for OpenStack parses the OpenStack cloud's network naming convention as follows:

- Any network that contains the name `mgmt`, `management`, `internal`, or `external` is assumed to indicate a network type (always-on management network, internal network, and external network, respectively). If there are multiple networks, BIG-IQ Cloud uses the first network it finds with those names to communicate with the OpenStack cloud.
- If there are no networks with those names, BIG-IQ Cloud assigns the network type based on the order in which the network was discovered. For example, if BIG-IQ Cloud discovers networks `10.10.10.0/24`, `20.20.20.0/24`, and `30.30.30.0/24`, it assigns them as follows:
 - Management network `10.10.10.0/24`
 - External network `20.20.20.0/24`
 - Internal network `30.30.30.0/24`

This is important to know, because when you create a new application server in OpenStack through BIG-IQ Cloud, you are allowed to select the internal or external network, but not the management network.

Tip: If you deploy a BIG-IP device in the OpenStack cloud and you want to discover it from BIG-IQ Cloud, you must have an external or interface route from BIG-IQ Cloud to the OpenStack cloud network. If BIG-IQ Cloud is not on same network as OpenStack, you might need to add a floating IP address to the interface to make it accessible. While either external or internal interfaces are acceptable, we recommend using the external interface.

Important: For specific instructions about how to configure your network for OpenStack, refer to the OpenStack documentation.

OpenStack Compute edits required to use BIG-IP VE systems

Before you create BIG-IP VE systems in an OpenStack environment, you must edit a file on each OpenStack Compute node. If you do not edit this file, any BIG-IP VE system you configure fails to start.

1. Log in to the command line of each OpenStack Compute node and edit `/etc/nova/release` to read as follows:

```
[Nova]
vendor = Red Hat
product = Bochs
package = RHEL 6.3.0 PC
```

2. Restart the OpenStack Compute node services.

This edit provides the BIG-IP VE system required access to the OpenStack hypervisor. Any BIG-IP VE systems you create in the OpenStack environment can now properly start.

Discovering devices located in the OpenStack cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

For devices located in a third-party cloud, you must know the internal self IP address (For OpenStack or VMware cloud) or the external self IP address for Amazon EC2. You also must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. To access this setting, log in to BIG-IQ System, select the BIG-IQ system you want to modify, and click the gear icon.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. In the IP Address field, type the device's external self IP address.
You cannot discover a BIG-IP device using its management IP address.

3. When the BIG-IQ system and the BIG-IP device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 - b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

Where `<route name>` is a user-provided name to identify the new route, and `<x.x.x.x>` is the IP address of the default gateway for the internal network.

4. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
5. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.
6. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

You can now associate this device with an OpenStack cloud connector and allocate resources to tenants.

Associating an OpenStack connector with devices

BIG-IQ® Cloud must be able to collect statistics to provide server diagnostics to tenants. By default, most OpenStack deployments are configured to disallow diagnostics collection. For successful deployment, you must change this option by editing the Nova `policy.json` file (typically located in the `/etc/nova/` directory) and changing the following line: `compute_extension:server_diagnostics":`
`"rule:admin_api to compute_extension:server_diagnostics": "rule:admin_or_owner".`

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Connectors header and click the + icon when it appears.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
4. From the **Cloud Provider** list, select **OpenStack**.
5. In the **OpenStack Controller Node URI** field, type the URI for the OpenStack controller node.
6. In the **OpenStack User Name** field, type the user name for the OpenStack administrator.
For example, `https://<IP address>:<Port>` or `http://<IP address>:<Port>`.
Note that default port for `http` is 5000.
7. In the **OpenStack Tenant Name** and **OpenStack Password** fields, type the tenant (also known as, project) name and password.
8. Click the **Save** button.

BIG-IQ Cloud discovers all associated OpenStack servers, and populates them in the Servers panel.

To complete discovery of BIG-IP® devices and populate the Devices panel, provide the administrator user name and password when requested. You can then associate tenants with the OpenStack connector.

Chapter 7

Integrating VMware

- *Configuring VMware NSX 6.1 for BIG-IQ Cloud*
 - *About VMware NSX version 6.0 integration*
 - *About vCloud Director integration*
-

Configuring VMware NSX 6.1 for BIG-IQ Cloud

You must have installed a BIG-IQ® system with two control plane subnets: one to be used for provisioning BIG-IP devices, and the other for BIG-IP® device discovery. These two subnets need to be interconnected.

Additionally, you must configure the following objects in VMware vSphere Web Client before you can perform this task.

- A Datacenter.
- A Datastore for your Datacenter.

Configuring the VMware objects described in this task makes it possible for a BIG-IQ system to configure and license a BIG-IP VE that you can manage with NSX as a load balancing service runtime. Your vCenter users can use this service runtime to deploy load-balanced virtual servers.

1. On the command line for the BIG-IQ system, use the following `tmsh` command to configure the BIG-IQ system to have the default route on the second control plane network.

```
tmsh create net route 0.0.0.0/0 gw 192.168.44.1
```

2. In the VMware vSphere Web Client, create four networks.

Two networks must be control plane networks; the BIG-IQ system uses one for provisioning BIG-IQ systems and the other to discover BIG-IP devices. The other two networks are data plane; the BIG-IP device uses one to pass external traffic and the other to pass internal traffic.

3. In the VMware vSphere Web Client, create four IP Pools, one for each network. As you create each pool, you are prompted for a name. Make a note of the names you choose so that when you need to associate each pool to a network interface, you will know which is which.
 - a) Define the provisioning network for the BIG-IP device. Use a typical IP address range to refer to the first management IP pool: 192.168.11.0/24.
 - b) Define the external data network. Use a typical IP address range to refer to the first data IP pool: 10.22.0.0/16.
 - c) Define the internal data network. Use a typical IP address range to refer to the second data IP pool: 10.33.0.0/16.
 - d) Define the discovery network for the BIG-IP device. Use a typical IP address range to refer to the second management IP pool: 192.168.44.0/24.
4. In the VMware vSphere Web Client, set up a web server on one of the just-created management networks. The NSX Manager uses the URL of this web server to access the installation file (OVF) for the BIG-IP VE you intend to provision.
5. Copy the OVF file that the NSX Manager will use to create the BIG-IP VE to an accessible location on the just-created web server.

The next tasks to perform are:

- Create a new user
- Activate a pool license
- Create a BIG-IQ software - VMware NSX connector
- Create a BIG-IQ device image (also referred to as an NSX node template)
- Configure your virtual application networks

Network requirements for communication with VMware cloud services

For proper communication, BIG-IQ® Cloud must have network access to the resources on which VMware software is installed. Before you can manage cloud resources, you must define a network route between the BIG-IQ Cloud device's internal VLAN and the management VLAN on the VMware.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

For devices located in a third-party cloud, you must know the internal self IP address (For OpenStack or VMware cloud) or the external self IP address for Amazon EC2. You also must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. To access this setting, log in to BIG-IQ System, select the BIG-IQ system you want to modify, and click the gear icon.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. In the IP Address field, type the device's external self IP address.

You cannot discover a BIG-IP device using its management IP address.

3. When the BIG-IQ system and the BIG-IP device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 - b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

Where `<route name>` is a user-provided name to identify the new route, and `<x.x.x.x>` is the IP address of the default gateway for the internal network.

4. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
5. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.

6. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

You can now associate this device with an VMware cloud connector and allocate resources to tenants.

Create a connection between the BIG-IQ device and NSX

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.

2. Hover over the Connectors header and click the + icon when it appears.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
4. From the **Cloud Provider** list, select **VMware NSX**.
5. In the **VMware NSX Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device's internal VLAN.
6. In the **VMware NSX User Name** and **VMware NSX Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the NSX Manager REST API.
7. In the **VMware vCenter Server Address** field, type the IP address of the vCenter server.
8. In the **VMware vCenter Server User Name** and **VMware vCenter Server Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the vCenter SOAP API.
9. In the **BIG-IQ User Name** and **BIG-IQ Password** fields, type the credentials that NSX Manager uses to authenticate to the BIG-IQ REST API.
10. If you plan to use a pool of licenses, in the **Device License** field, specify the pool of licenses to use when the NSX and BIG-IQ integration provisions a BIG-IP VE.
If you skip this step, you'll need to specify a license each time you add a new device.
11. If you want to specify values for the remaining optional fields (**Timezone**, **NTP Server(s)**, **DNS Servers(s)**, and **DNS Suffix(s)**) so that the NSX and BIG-IQ system integration will use them when it provisions a BIG-IP VE, specify those values next.
12. Click the **Save** button.

About VMware NSX version 6.0 integration

BIG-IQ[®] Cloud provides you with the tools to provide tenants access to self-deployable iApps through VMware NSX 6.0 integration.

Network requirements for communication with VMware cloud services

For proper communication, BIG-IQ[®] Cloud must have network access to the resources on which VMware software is installed. Before you can manage cloud resources, you must define a network route between the BIG-IQ Cloud device's internal VLAN and the management VLAN on the VMware.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ[®] system, you can discover BIG-IP[®] devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

For devices located in a third-party cloud, you must know the internal self IP address (For OpenStack or VMware cloud) or the external self IP address for Amazon EC2. You also must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. To access this setting, log in to BIG-IQ System, select the BIG-IQ system you want to modify, and click the gear icon.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. In the IP Address field, type the device's external self IP address.

You cannot discover a BIG-IP device using its management IP address.

3. When the BIG-IQ system and the BIG-IP device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 - b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

Where `<route name>` is a user-provided name to identify the new route, and `<x.x.x.x>` is the IP address of the default gateway for the internal network.

4. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
5. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.
6. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

You can now associate this device with an VMware cloud connector and allocate resources to tenants.

Associating a VMware cloud connector with a device

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Hover over the Connectors header and click the + icon when it appears.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **VMware Networking**.
4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list. BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. In the **VMware Networking Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device's internal VLAN.
7. In the **VMware Networking User Name** and **VMware Networking Password** fields, type the credentials for the VMware administrator.
8. From the **BIG-IQ User Name** list, select the BIG-IQ user the VMware administrator should contact and, in the **BIG-IQ Password** field, type the password for that user.
9. Click the **Save** button.

About vCloud Director integration

Integrating vCloud Director (VCD) with your cloud applications makes it possible for you to use the VCD interface to manage the F5 cloud applications. The integration process involves tasks using the user interface in both the F5 BIG-IQ® Cloud and the VMware VCD.

After you integrate vCloud Director (VCD) with BIG-IQ Cloud, you can use VCD to manage your cloud applications. After integration, a catalog of BIG-IQ® Cloud applications appears in the VCD user interface.

BIG-IQ Cloud refers to a service provider's customers as *tenants*. The VCD equivalent to a tenant is referred to as an *organization*. BIG-IQ Cloud identifies tenants using a tenant ID. One key to successfully integrating VCD with BIG-IQ Cloud is associating the tenant ID assigned to that catalog with a VCD organization.

To deploy an F5 application catalog in vShield Manager (VSM), you deploy a VSM service profile. While VSM service profiles do not currently recognize F5 tenants, they do recognize VCD organizations. So when your tenant's ID is associated with a VCD organization, you can use VSM and VCD to administer and deploy the tenant's application catalog.

When you create a tenant for VCD integration, make a note of the tenant ID so you can connect it to a VCD organization.

Task summary

When you are integrating vCloud Director (VCD) and BIG-IQ® Cloud, you must configure VCD, then BIG-IQ, then VCD again.

Network requirements for communication with VMware cloud services

For proper communication, BIG-IQ® Cloud must have network access to the resources on which VMware software is installed. Before you can manage cloud resources, you must define a network route between the BIG-IQ Cloud device's internal VLAN and the management VLAN on the VMware.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IQ® devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

For devices located in a third-party cloud, you must know the internal self IP address (For OpenStack or VMware cloud) or the external self IP address for Amazon EC2. You also must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. To access this setting, log in to BIG-IQ System, select the BIG-IQ system you want to modify, and click the gear icon.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. In the IP Address field, type the device's external self IP address.
You cannot discover a BIG-IQ device using its management IP address.
3. When the BIG-IQ system and the BIG-IQ device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.

- b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x>
network default}`

Where `<route name>` is a user-provided name to identify the new route, and `<x.x.x.x>` is the IP address of the default gateway for the internal network.

4. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
5. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.
6. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

You can now associate this device with an VMware cloud connector and allocate resources to tenants.

Associating a VMware cloud connector with a device

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Hover over the Connectors header and click the + icon when it appears.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **VMware Networking**.
4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list. BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. In the **VMware Networking Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device's internal VLAN.
7. In the **VMware Networking User Name** and **VMware Networking Password** fields, type the credentials for the VMware administrator.
8. From the **BIG-IQ User Name** list, select the BIG-IQ user the VMware administrator should contact and, in the **BIG-IQ Password** field, type the password for that user.
9. Click the **Save** button.

Chapter

8

Local Cloud Integration

- *About using a local cloud source*
- *Discovering BIG-IP devices in your network*
- *Associating a local cloud connector with a device*

About using a local cloud source

In addition to providing self-service resources to tenants remotely in a third party cloud, you can also provide them resources to local F5 devices in your network.

Discovering BIG-IP devices in your network

After you license and perform the initial configuration for the BIG-IQ[®] system, you can discover BIG-IP[®] devices running version 11.3 or later. For proper communication, you must configure each F5 device you want to manage with a route to the BIG-IQ system. If you do not specify the required network communication route between the devices, then device discovery fails.

You can discover a device by providing the BIG-IQ system with the device's IP address, user name, and password.

1. Hover over the Devices header, click the + icon when it appears, and then select **Discover Device**.
2. For devices on the same subnet as the BIG-IQ system, in the **IP Address** field, type the IP address of the device.

You cannot discover a BIG-IP device using its management IP address.

3. When the BIG-IQ system and the BIG-IP device are on different subnets, you must create a route:
 - a) Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 - b) Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

Where *<route name>* is a user-provided name to identify the new route, and *<x.x.x.x>* is the IP address of the default gateway for the internal network.

4. To change the root user name, type a new name in the **Root User Name** field.
5. Type a password for the root user in the **Root Password** field.
6. In the **Admin User Name** and **Admin Password** fields, type the administrator user name and password for the managed device.
7. Select the **Auto Update Framework** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework. If you do not select the **Auto Update Framework** check box before you click the **Add** button, a message displays prompting you to update the framework or cancel the task.
8. Click the **Add** button.

BIG-IQ System populates the properties of the device that you added, and displays the device in the Devices panel.

Associating a local cloud connector with a device

Before you associate a local cloud connector with a device, you must discover one or more devices.

To enable integration between a third-party cloud provider and the BIG-IQ device, you must configure a cloud connector. A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

1. Hover over the Connectors header and click the + icon when it appears.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **Local Cloud**.
4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list.
BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. Click the **Save** button.

Chapter

9

Cloud Tenant Management

- *About creating cloud tenants*
 - *Creating a tenant*
 - *Creating a cloud user*
 - *Associating a user with a tenant's role*
-

About creating cloud tenants

As a cloud administrator, you create tenants and allocate resources to them in the form of iApps® application templates. Tenants can then self-deploy the customized application templates to easily define network and application services for several devices, without requiring them to perform complicated networking procedures.

The process of providing resources for a tenant includes these tasks:

- Create a tenant - When you create a tenant, BIG-IQ® Cloud creates a unique role for the tenant and populates it in the Role panel.
- Create a user - When you create a user account, you assign a user name and a password.
- Associate a user with a tenant's role - You associate a user with a tenant to provide that user access to pre-defined cloud resources in the form of self-service customized applications. You can associate multiple users with a single tenant for access to specific resources.

Creating a tenant

You create a tenant to provide access to customized cloud resources and applications.

***Note:** To create an OpenStack tenant, the OpenStack administrator must be a member of the OpenStack tenant/project.*

1. Hover on the Tenants header, and click the + icon when it appears.
The panel expands to display property fields for the new tenant.
2. In the **Name** and **Description** fields, type a name and an optional description for this tenant.
The name can consist of a combination of numbers and symbols, but cannot contain any spaces.
3. From the **Available Connectors** list, select the connector associated with the resources that you are going to provide to this tenant.
To add another connector, click the plus (+) sign and select a connector from the additional **Available Connectors** list.
4. In the **Address**, **Phone**, and **Email** fields, type optional contact information for this tenant.
5. Click the **Save** button.

You can now associate a user with this tenant to provide access to applications and services.

Creating a cloud user

Create a cloud user to provide that individual with access to specific resources.

1. Hover on the User header, and click the + icon when it appears.
The panel expands to display property fields for the new user.
2. In the **Full Name** field, type a name to identify this user.
The full name can contain a combination of symbols, letters, numbers and spaces.

3. In the **Password** and **Confirm Password** fields, type the password for the new user.
4. Click the **Add** button.

You can now associate this user with an existing tenant to provide access to pre-defined cloud resources.

Associating a user with a tenant's role

Before you associate a user with a tenant's role, you must first create the tenant. You can associate multiple users with a tenant's role.

You associate user with a tenant's role to provide that user specific access to cloud resources in the form of self-service applications.

In the Users panel, click the user name that you want to associate with a role and drag and drop it onto that role, in the Roles panel.

This user now has access to all of the resources defined for the associated role.

Chapter 10

iApps Application Template Customization

- *About customizing iApp application templates*
- *Creating a customized application template*

About customizing iApp application templates

An *iApp* is an application template located on F5 devices running TMOS® 11.5.0 and later. When you discover an F5 device, all iApps® templates installed on that device are imported to the BIG-IQ® system.

As a cloud administrator, you can modify imported application templates to offer specific configurations and cloud resource access for tenants. You do this by creating a catalog entry, specifying tenant-specific details such as virtual IP address, application server IP addresses, and so forth. Once saved, these catalog entries are available to tenants for self-deployment from the application panel. This saves tenants time, because it does not require that they perform complex network tasks, and it also makes it possible for you to easily duplicate a configuration for several users simultaneously. You also have the option to allow tenants to further customize the applications as required, and deploy them as needed.

Creating a customized application template

Before you can customize an application template for a tenant, you must discover at least one F5 device that contains iApps® templates.

As a cloud provider, you modify iApps templates to customize network settings, levels of services, and so forth, for tenants. You can create variations of the same application, offering different types of access (LAN or WAN), or providing a specific limit of connections.

Note: *Once you customize and save an application as a catalog entry, you cannot modify it.*

1. Hover over the Catalog header, and click the + icon when it appears.
The panel expands to display the Catalog properties.
2. In the **Name** field, type a name for this new application.
3. From the **Application Type** list, select an application.
4. Unless you want to restrict this application template to a specific cloud connector, leave the **Cloud Connector** setting as **Tenant Selectable** so tenants are allowed to select the appropriate cloud connector when they deploy this application.
5. If the **Application Tiers** settings are displayed (expanded), select the options that match the properties for this application; otherwise, keep the default settings.

Important: *If you must specify the options for these settings, select the **Tenant Editable** check box for the virtual server and pool members.*

6. To allow cloud tenants to specify certificates with SSL encryption when self-deploying applications, select options from the **SSL Cert** and **SSL Key** lists.
BIG-IQ Cloud uses these options to provide the appropriate certificate and key when the tenant self-deploys this application to a BIG-IP® device. These options are not available for all application templates.
7. Finish making modifications by specifying the Application Properties and Customize Application Template variables.
To allow a tenant to modify a particular setting, select the **Tenant Editable** check box for that setting. For further details about template variables and settings, refer to the *BIG-IP® iApps® Developer's Guide*.
8. Click the **Save** button.

You can now send the cloud IP addresses to the tenant and use this IP address range in configuring server tiers and pool members, within certain application services. The tenant can self-deploy the application from the catalog.

The customized application displays as an entry in the catalog.

Chapter 11

Glossary

- *BIG-IQ Cloud terminology*
-

BIG-IQ Cloud terminology

Before you manage cloud resources, it is important that you understand some common terms as they are defined within the context of the BIG-IQ® Cloud.

Term	Definition
<i>application templates</i>	An application template is a collection of parameters (in the form of F5 iApps® templates) that a cloud administrator defines to create a customized configuration for tenants. Cloud administrators add the configured application to a catalog from which a tenant can self-deploy it.
<i>BIG-IQ Cloud</i>	The BIG-IQ® Cloud system is a tool that streamlines management and access for tenants to services and applications hosted by local and/or cloud-based servers.
<i>cloud administrator</i>	Cloud administrators create application templates for tenants to centrally manage access to specific web-based applications and resources. Cloud administrators might also be referred to as cloud providers.
<i>cloud bursting</i>	Cloud bursting is a seamless way to manage an anticipated increase in application traffic by directing some traffic to another cloud resource. When demand falls back into normal parameters, traffic can be directed back to the original cloud resource. This elasticity enables efficient management of resources during periods of increased or decreased traffic to applications.
<i>cloud connector</i>	A cloud connector is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.
<i>resources</i>	A resource is any managed object, including devices, web applications, virtual servers, servers, cloud connectors, and so forth.
<i>roles</i>	A role defines specific privileges to which you can associate one or more users. There are two default roles for BIG-IQ Cloud: cloud administrator and cloud tenant.
<i>tenant</i>	A tenant is an entity that can consist of one or more users accessing resources provided by a cloud administrator.
<i>user</i>	A user is an individual who has been granted access to specific tenant resources.

Index

A

- active-active cluster
 - configuring for the BIG-IQ Cloud [20](#)
- active-active pair
 - configuring for the BIG-IQ system [20](#)
- activities
 - viewing for cloud resource activity [39](#)
- Amazon CloudWatch
 - about integrating with BIG-IQ Cloud [30](#)
- Amazon EC2 [38](#)
- See also EC2
 - about integrating with BIG-IQ Cloud [30](#)
 - See also EC2
- Amazon EC2 cloud
 - creating BIG-P VE version 11.3 or 11.4 devices [35](#)
- Amazon EC2 devices
 - discovering [34](#)
- Amazon EC2 resources
 - viewing cloud resource activity [39](#)
- Amazon Elastic Computer Cloud, See Amazon EC2
- Amazon Machine Images, See AMI
- Amazon Virtual Private Cloud [31](#)
- See also VPC
 - about integrating with BIG-IQ Cloud [30](#)
 - creating [31](#)
 - See also VPC
- Amazon virtual server
 - launching using an AMI [32](#)
- Amazon web services
 - and network configuration requirements [30](#)
- AMI
 - using to launch a virtual server [32](#)
- application catalog [36, 62](#)
- applications
 - customizing for tenants [36, 62](#)
 - deploying as a Cloud Tenant user [37](#)
- application templates
 - about [62](#)
 - and vApps [48](#)
 - defined [66](#)
 - deploying as a Cloud Tenant user [37](#)
 - using [36, 62](#)
- asset management
 - for devices [23](#)
- authentication [30](#)

B

- BIG-IQ Cloud
 - defined [66](#)
 - finding documentation for [16](#)
- BIG-IQ cloud tenant
 - finding documentation for [16](#)
- BIG-IQ system
 - finding documentation for [16](#)
 - reordering panel [17](#)

C

- catalog
 - for applications [36, 62](#)
- catalog entries
 - creating for tenants [62](#)
- cloud administrator
 - defined [66](#)
- cloud bursting
 - defined [66](#)
- cloud connector
 - [49, 51](#)
 - defined [66](#)
 - for OpenStack [44](#)
 - for VMware NSX [47](#)
- cloud connector, local
 - associating with a device [54](#)
- cloud connector for EC2 [33](#)
- cloud resources
 - providing for tenants [58, 62](#)
 - using locally [54](#)
- cloud tenant
 - adding [58](#)
- cloud tenants
 - about creating [58](#)
- Cloud Tenant users
 - deploying applications [37](#)
- clusters
 - for high availability [20](#)
- configure
 - NSX [46](#)
- connector, local
 - associating with a device [54](#)
- connectors
 - VMware [48](#)

D

- device discovery
 - by scanning network [22, 54](#)
- device elasticity
 - configuring [37](#)
- device inventory
 - about [22](#)
 - viewing details [23](#)
- device management
 - about [22](#)
- devices
 - about discovering [22](#)
 - adding [22, 54](#)
 - creating in Amazon EC2 cloud [35](#)
 - discovering Amazon EC2 devices [34](#)
 - discovering OpenStack devices [43](#)
 - discovering VMware devices [47–48, 50](#)
- documentation, finding [16](#)
- dynamic cloud resources
 - viewing activity for [39](#)

E

- EC2 connector
 - associating with a device [33](#)
- elasticity
 - configuring for tenants [37](#)
 - viewing activity for [39](#)
- elastic resources
 - specifying settings for [37](#)
- exportation of inventory details [23](#)

F

- failover [20](#)
- filtering
 - objects [16](#)

G

- glossary [66](#)
- Grizzly, See OpenStack
- guides, finding [16](#)

H

- HA, See high availability cluster
- Havana, See OpenStack
- high availability cluster
 - configuring [20](#)
- high availability configuration
 - about [20](#)

I

- IAM
 - creating user account [30](#)
- IAM access [38](#)
- iApps
 - customizing for tenants [62](#)
 - defined [62](#)
 - for OpenStack [42](#)
- inventory details
 - exporting to CSV file [23](#)
 - viewing for devices [23](#)
- IP addresses
 - for managed devices [22](#)

L

- license pool, See pool license
- licenses
 - for pools [27](#)
- licensing
 - activating pool license automatically [26](#)
 - activating pool license manually [26](#)
 - assigning to BIG-IP devices [27](#)
 - for pool license [26](#)
- local cloud connector
 - associating with a device [54](#)
- local cloud resources
 - using [54](#)

M

- managed devices
 - about discovering [22](#)
- manuals, finding [16](#)

N

- network configuration
 - for integrating with OpenStack cloud services [42](#)
 - using Amazon web services [30](#)
 - using VMware [47–48, 50](#)
- network configurations
 - customizing for tenants [36, 62](#)
- network configurationsiApps
 - customizing for tenants [36, 62](#)
- network resources
 - using for cloud services [54](#)
- NSX
 - configuration [46](#)

O

- objects
 - finding associations [16](#)
- OpenStack
 - and iApps [42](#)
 - required configuration [43](#)
 - using with BIG-IP VE systems [43](#)
- OpenStack cloud services
 - and network configuration requirements [42](#)
- OpenStack connector
 - associating with a device [44](#)
- OpenStack devices
 - discovering [43](#)

P

- panels
 - reordering [16–17](#)
- pool license
 - activating automatically [26](#)
 - activating manually [26](#)
 - assigning to a BIG-IP device [27](#)
- pool licenses
 - about [26](#)

R

- release notes, finding [16](#)
- reports
 - for asset management [23](#)
- resource elasticity
 - configuring for tenants [37](#)
- resources
 - defined [66](#)
 - providing access for user [59](#)

S

server elasticity
 configuring 37

T

tenant
 adding 58
tenant access
 using IAM 38
tenants
 about creating 58
 associating with a user 59
 creating applications for 36, 62
tenants users
 and tenants 58
 and users 58
terminology 66
terms
 defined 66

U

user interface
 customizing 16–17
 navigating 16
users
 adding 58
 associating with a tenant 59
 defined 58

V

vApps
 and application templates 48
 deploying 48
vCloud Director integration
 about 50
vCNS, See VMware vCNS
Virtual Private Cloud 30–31
 See also Amazon Virtual Private Cloud
VMware
 about integration with iApps 48
 and network configuration requirements 47–48, 50
 supported version 48
VMware cloud connector
 associating with a device 49, 51
VMware devices
 discovering 47–48, 50
VMware NSX
 integrating with BIG-IQ Cloud 47, 49, 51
 integration 48
VMware vCNS
 integrating with BIG-IQ Cloud 49, 51
VMware vShield
 integrating with BIG-IQ Cloud 49, 51
VPC
 30–31
 creating 31
vShield, See VMware vShield

