

BIG-IQ™ Cloud: Cloud Administration

Version 4.3



Table of Contents

Legal Notices.....	5
Acknowledgments.....	7
Chapter 1: BIG-IQ Cloud Overview.....	13
Overview: BIG-IQ Cloud.....	14
Additional resources and documentation for BIG-IQ Cloud.....	14
About the BIG-IQ system user interface.....	14
Filtering for associated objects.....	15
Customizing panel order.....	15
Chapter 2: BIG-IQ High Availability.....	17
About a high availability active-active cluster.....	18
Configuring BIG-IQ Cloud in an active-active high availability configuration.....	18
Chapter 3: Device Resource Management.....	19
About device discovery.....	20
Using a script to install required BIG-IQ components on managed BIG-IP devices.....	20
Discovering devices.....	21
Adding devices located in a third-party cloud.....	22
Viewing device inventory details.....	23
Chapter 4: License Pools.....	25
About license pools.....	26
Automatically activating a license pool.....	26
Manually activating a license pool.....	26
Chapter 5: Amazon EC2 Cloud Integration.....	29
About Amazon EC2 integration	30
Network requirements for communication Amazon EC2 cloud services.....	30
Creating a new virtual private cloud.....	30
Launching a new virtual machine.....	31
Creating a new IAM user account.....	32
Associating an EC2 cloud connector with a device and discovering application servers.....	32
Setting up tenant access using IAM.....	33
Viewing activity for elastic cloud resources.....	34
Chapter 6: VMware Cloud Integration.....	35
About VMware integration	36

Network requirements for communication with VMware cloud services36
Associating a VMware cloud connector with a device.....36

Chapter 7: OpenStack Cloud Integration.....37
About OpenStack servers38
Network requirements for communication with OpenStack cloud services38
OpenStack Compute edits required to use BIG-IP VE systems.....39
Associating an OpenStack connector with devices.....39

Chapter 8: Local Cloud Resources.....41
About using a local cloud source.....42
Associating a local cloud connector with a device.....42

Chapter 9: Cloud Tenant Management.....43
About creating cloud tenants44
Creating a tenant.....44
Creating a user.....44
Associating a user with a tenant's role.....45

Chapter 10: iApps Application Template Customization.....47
About customizing iApp application templates.....48
Creating a customized application template.....48

Chapter 11: Glossary.....51
BIG-IP Cloud terminology.....52

Legal Notices

Publication Date

This document was published on March 6, 2014.

Publication Number

MAN-0501-01

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, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, 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, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), 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.

Acknowledgments

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

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

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

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

This product includes software developed by Jared Minch.

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

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

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.

Acknowledgments

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

This product includes 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

Acknowledgments

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 Leaflet software, copyright © 2010-2014, Vladimir Agafonkin, and copyright © 2010-2011, CloudMade; all rights reserved. This software is distributed under the BSD license.

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

Chapter

1

BIG-IQ Cloud Overview

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

Overview: BIG-IQ Cloud

As a cloud administrator, you use BIG-IQ™ Cloud to customize F5 iApps® to supply tenants with self-service access to resources, including servers, storage, applications, and services. You can also easily create one configuration and distribute it to multiple tenants, reducing setup time for your users. Tenants can access these optimized resources without having to perform complicated network tasks, ensuring the accuracy of complex traffic management configurations.

BIG-IQ Cloud also easily integrates with cloud service providers, like Amazon EC2 and VMware products. Depending on the application requirements and traffic volume, you can use your local cloud resources, public cloud resources, or a hybrid of both. In this way, you can easily expand and reallocate resource as needed, without requiring you purchase more equipment.

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.
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

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) cluster. 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 cluster fails, a peer BIG-IQ system takes over the device management that was previously done by the original device.

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

You must install and license at least two BIG-IQ systems before you can configure them in an active-active high availability cluster.

Configuring BIG-IQ Cloud in a high availability cluster means you always have management access to the BIG-IP devices in your network. (Configuring a high availability cluster is optional.)

Important: *To synchronize properly, BIG-IQ systems in a cluster must be running the same version of software. The exact configuration in terms of hardware is not required; however, the systems should have comparable resources. This is required because, in the event of a fail over, a peer must be able to maintain the process requirements. This is especially important in terms of disk space and data collection.*

1. Hover on the **BIG-IQ HA** header and click the + icon when it appears.
2. 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.
3. Click the **Add** button to add this device to this high availability cluster.
4. In the **User Name** and **Password** fields, type the administrative user name and password for the system.

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

Chapter

3

Device Resource Management

- *About device discovery*
- *Using a script to install required BIG-IQ components on managed BIG-IP devices*
- *Discovering devices*
- *Adding devices located in a third-party cloud*
- *Viewing device inventory details*

About device discovery

You use the BIG-IQ™ Cloud to centrally manage cloud resources for tenants. Resources can be located on BIG-IQ devices in your local network, in a public cloud like Amazon EC2, or in a combination of both.

You start managing the resources that are located on devices in your local network by installing on those devices components that are specific to BIG-IQ Cloud, and then making BIG-IQ Cloud aware of them through the discovery process. You accomplish the component installation process from the command line. To discover a device, you provide BIG-IQ Cloud with its IP address, user name, and password.

Alternately, for those devices located in the Amazon EC2 and OpenStack public cloud space, BIG-IQ Cloud automatically detects them when it connects to the Amazon EC2 or OpenStack cloud.

Using a script to install required BIG-IQ components on managed BIG-IP devices

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

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)

Installing requisite BIG-IQ components onto your managed BIG-IP devices results in a REST framework that supports the required Java-based management services. You must perform this installation task on each device before you can discover it.

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

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


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

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

```
cd /usr/lib/dco/packages/upd-adc
```

4. Run the installation script.

- For devices installed in an Amazon EC2 environment: `./update_bigip.sh -a admin -p <password> -i /<path_to_PEM_file> <BIG-IP Management IP Address>`
- For devices installed in any other environment: `./update_bigip.sh -a admin -p <password> <BIG-IP Management IP Address>`

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

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

```
ssh root@<big-ip addr> 'cat /var/ssh/root/authorized_keys' | grep -v -F -f /root/.ssh/identity.pub | ssh root@<big-ip addr> 'cat ->/var/ssh/root/authorized_keys.tmp && mv /var/ssh/root/authorized_keys.tmp /var/ssh/root/authorized_keys'
```

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

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.

Discovering devices

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, device discovery will fail.

For devices located in a third-party cloud (such as EC2, OpenStack, and VMware), you must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. You access this setting by clicking **System > Overview > Services**.

Important: Before you can discover a device, you must first install the required BIG-IQ components on that device.

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

1. Hover on the Devices header, and click the + icon when it appears. The panel expands to display the New Device properties.
2. For devices on the same subnet as the BIG-IQ system, in the **IP Address** field, specify the IP address of the device:
 - For devices in your local network, or located on an OpenStack or VMware cloud device, type the device's internal self IP address.
 - For devices located on Amazon EC2 cloud, type the device's external self IP address.

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

3. If 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: `tmsch 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. For devices on the same subnet as the BIG-IQ system, in the **IP Address** field, specify the IP address of the device:
 - For devices in your local network, or located on an OpenStack or VMware cloud device, type the device's internal self IP address.
 - For devices located on Amazon EC2 cloud, type the device's external self IP address.

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

5. If 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.

6. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
7. Click the **Add** button.

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

Adding devices located in a third-party 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, device discovery will fail.

For devices located in a third-party cloud (such as EC2, OpenStack, and VMware), you must configure BIG-IQ Cloud with DNS so it can resolve the endpoint by name. You access this setting by clicking **System > Overview > Services**.

Important: Before you can discover a device, you must first install the required BIG-IQ components on that device.

You discover a device in a third-party cloud by specifying a connector, selecting a device's image, and providing a user name and password for that device.

1. Hover on the Devices header, and click the + icon when it appears.
The panel expands to display the New Device properties.
2. Select the **Create a device** option.
3. From the **Cloud Connector** list, select the connector associated with the device you are adding.
4. From the **Device Image** list, select an image for the device you are creating.
5. If 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.

6. In the **User Name** and **Password** fields, type a user name and password for this device.
7. Click the **Add** button.

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

Viewing 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. To display the details for all managed devices, verify that the filter field at the top of the screen is clear, and then click the show details (|>) button in the Devices panel header.
The panel expands to display the details for all of the managed devices.
2. To view the details for a specific device, click the device listed in the Device panel, and then click the change view button in the panel header.
The screen refreshes to display the details for the selected device.
3. To export the data to a CSV file, click the **Export** button on the device details screen.

You can modify the report as required in Microsoft™ Excel™.

Chapter

4

License Pools

- *About license pools*
 - *Automatically activating a license pool*
 - *Manually activating a license pool*
-

About license pools

The BIG-IQ™ system helps you manage resource usage in a dynamic application environment through the use of license pools. Each license pool is limited to a specific number of licenses. After you perform the initial license activation, BIG-IQ automatically grants and revokes licenses for the BIG-IP® virtual edition systems as resource demands change. This gives you the flexibility to license devices only as needed, keeping your operating costs fixed. You can also use different license pools for different applications, allowing for flexible provisioning options.

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

Note: If you do not have a base registration key, contact your F5 Networks sales representative.

Automatically activating a license pool

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 use this procedure to activate the license pool.

1. Log in to the BIG-IQ system with the administrator user name and password.
2. At the top of the screen, click **Device** and **Virtual**.
3. Hover on the License Pools header, and click the + icon when it appears.
The panel expands to display New License properties.
4. In the **License Name** field, type the name you want to use to identify this license pool.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Keys** field, paste any additional license key you have.
7. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button.
The BIG-IQ system contacts the F5 Networks licensing server and displays the End User License Agreement (EULA).
8. To accept the EULA, click the **Accept** button.
The screen refreshes and displays the license details.

Manually activating a license pool

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

If the resources you are licensing are not connected to the public internet, you can manually activate the license pool.

1. Log in to the BIG-IQ system with the administrator user name and password.

2. At the top of the screen, click **Device** and **Virtual**.
3. Hover on the License Pools header, and click the + icon when it appears.
The panel expands to display New License properties.
4. In the **License Name** field, type the name you want to use to identify this license pool.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Keys** field, paste any additional license key you have.
7. For the **Activation** method setting, select **Manual** and click the **Activate** button.
The BIG-IQ system refreshes and displays the dossier in the **Dossier** field.
8. Copy the displayed dossier and transfer it to a system connected to the internet and navigate to the F5 Licensing Server at <https://activate.f5.com/license/>.
9. Paste the dossier into the **Enter your dossier** text box, or click the **Browse** button to locate it on the system, and click the **Next** button.
10. Copy or save the activation key and transfer it to the BIG-IQ system.
11. The End User License Agreement (EULA) displays.
When you click **Accept**, the screen refreshes to display the license details.
12. To accept the EULA, click the **Accept** button.
The screen refreshes and displays the license details.

Chapter 5

Amazon EC2 Cloud Integration

- *About Amazon EC2 integration*
-

About Amazon EC2 integration

Using Amazon Web Services (AWS) is less expensive and more flexible than building and maintaining a physical computer infrastructure. BIG-IQ™ Cloud provides you with a seamless way to manage Amazon's elastic cloud services (*Amazon EC2*). To support communication between BIG-IQ Cloud and an AWS account, you use the BIG-IQ Cloud's Amazon EC2 Cloud Connector. This EC2 Cloud Connector makes it possible for you to discover BIG-IP® VE virtual machines and application servers running in an AWS account.

You can use this feature to coordinate management-plane changes to a private, public, or hybrid cloud environment. For example, to accommodate seasonal traffic fluctuations, you might need to periodically add devices or application servers in the EC2 environment (referred to as, *cloud bursting*) or retract devices or application servers.

Task summary

Creating a new virtual private cloud

Launching a new virtual machine

Creating a new IAM user account

Associating an EC2 cloud connector with a device and discovering application servers

Setting up tenant access using IAM

Viewing activity for elastic cloud resources

Network requirements for communication Amazon EC2 cloud services

For proper communication to devices located in an Amazon EC2 cloud, BIG-IQ™ Cloud must have network access to those resources. Before you can manage cloud resources, you must define a network route between the BIG-IQ Cloud internal VLAN and the public Internet, or the Amazon EC2 endpoint, for proper communication to devices located in a public cloud. For specific instructions, refer to your Amazon EC2 documentation.

Creating a new virtual private cloud

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

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

Important: *It is crucial to your success that you be consistent in the availability zone that you choose throughout the configuration process. Objects configured in one zone are not visible within other zones, so they cannot function together.*

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

Launching a new virtual machine

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 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.

Creating a new IAM user account

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

Tip: This task is optional; you can create a virtual machine without creating an IAM user account to control access, but using IAM is considered to be best practice.

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

For the most current instructions for creating a new IAM user, refer to the Amazon Virtual Private Cloud (VPC) Documentation web site, <http://aws.amazon.com/documentation/iam/>.

When you complete this task, you will have created a new IAM user and downloaded the credentials (an access key ID and secret access key) that provide access to AWS resources for that new user.

Associating an EC2 cloud connector with a device and discovering application servers

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 first purchase and activate a license pool to associate with this connector.

To provide cloud tenant users with access to resources, you must configure a cloud connector. A *cloud connector* provides two services. First, you can use it to identify a specific set of resources, much like a virtual container, and second, it provides integration with third-party cloud services.

1. Hover on 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 **Amazon EC2**.
4. 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.
5. In the **Key ID** and **Secret Key** fields, type the credentials of an EC2 user that can access your account.
For security purposes, it is important to specify a user that has Amazon EC2 Full Control Access.
6. 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.

7. 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 uses the first one it discovers in the Availability Zone.
8. Click the arrow next to Device & Server Provisioning to display associated options.
9. If you do not want to automatically provision additional BIG-IP VE servers and devices when more resources are needed, for the **Device Elasticity** setting, select **Disable**.
If you do want to provide server elasticity options to a tenant, you must leave the **Cloud Connector** setting as **Tenant Selectable**.
10. For the **Licensing** setting, select the license pool associated with this Amazon EC2 connector.
If there is no license pool associated with this Amazon EC2 connector and you enabled elasticity, Amazon will charge for additional resources on a per-instance basis. Refer to Amazon for their EC2 instance purchasing options.
11. If you do not want to automatically provision additional application servers when more resources are needed, for the **Server Elasticity** setting, select **Disable**.
12. Review network settings populated when you selected a connector and make any necessary corrections.
13. Click the **Save** button at the top of the New Connector panel.
BIG-IQ discovers application servers associated with this connector and populates them in the Server panel. If it discovers F5 devices, it populates them in the Device panel.
14. If the system discovered devices, you must expand the device's properties panel, and provide the device's credentials to finalize the discovery process.
15. Review network settings populated when you selected a connector and make any necessary corrections.

You can now add a cloud tenant using this connector and its associated devices.

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 elastic 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 tenants are expanding cloud resources to address increased traffic to applications.

1. To view the resource and tenant association with a particular activity, click the activity and then click the **Apply** button located next to the filter field.
The screen refreshes to display only items associated with the selected activity.
2. To remove the filter, click the **x** located below the filter field.
3. 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

VMware Cloud Integration

- *About VMware integration*
- *Network requirements for communication with VMware cloud services*
- *Associating a VMware cloud connector with a device*

About VMware integration

To provide services to VMware tenants, you must first configure communication between VMware products, and BIG-IQ™ Cloud. Then, you associate a VMware cloud connector with a device, and create a catalog entry for a corresponding VMware 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 vApps.

Using the VMware interface, you can manage the virtual machines necessary for application deployments, add or remove a virtual machine, and start or stop virtual machines. You can use your cloud vendor's user interface to duplicate changes.

Important: *To integrate with BIG-IQ Cloud, you must use VMware vShield Manager version 5.1.1 or later, VMware vCNS version 5.5 or later, and VMware NSX version 6.0 or later.*

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's internal VLAN and the management VLAN on the VMware. For specific instructions, refer to your VMware documentation.

Associating a VMware cloud connector with a device

To provide cloud tenant users with access to resources, you must configure a cloud connector. A *cloud connector* provides two services. First, you can use it to identify a specific set of resources, much like a virtual container, and second, it provides integration with third-party cloud services.

1. Hover on 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. In the **VMware Networking Address** field, type the self IP address of the VMware system.
The system administrator VMware IP address must be fully accessible from the internal VLAN.
5. In the **VMware Networking User Name** and **VMware Networking Password** fields, type the credentials for the VMware administrator.
6. 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.
7. Click the **Save** button.

Chapter

7

OpenStack Cloud Integration

- *About OpenStack servers*
-

About OpenStack servers

BIG-IQ™ Cloud includes a feature that provides you with the tools you need to manage OpenStack private cloud resources. These management tasks include discovering, creating, starting, and stopping BIG-IP® VE virtual machines and application servers running in the private cloud.

You can use this feature to accommodate seasonal traffic fluctuations for which you might need to periodically add devices or application servers (referred to as, *cloud bursting*) or retract devices or application servers.

OpenStack Compute edits required to use BIG-IP VE systems

Associating an OpenStack connector with devices

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 requires that you use the OpenStack Grizzly (2013.1) version of the software. Additionally, the OpenStack installation must use Neutron for networking.

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 (management 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.

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 data 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 provide cloud tenant users with access to resources, you must configure a cloud connector. A *cloud connector* provides two services. First, you can use it to identify a specific set of resources, much like a virtual container, and second, it provides integration with third-party cloud services.

1. Log in to BIG-IQ Cloud with your administrator user name and password.
2. Hover on 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.
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 at the top of the New Connector header.

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

8

Local Cloud Resources

- *About using a local cloud source*
 - *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.

Associating a local cloud connector with a device

Before you associate a local cloud connector with a device, you must specify a DNS server for BIG-IQ Cloud and discover one or more devices.

To provide cloud tenant users with access to resources, you must configure a cloud connector. A *cloud connector* provides two services. First, you can use it to identify a specific set of resources, much like a virtual container, and second, it provides integration with third-party cloud services.

1. Hover on 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.
6. Click the **Save** button at the top of the New Connector header.

Chapter 9

Cloud Tenant Management

- *About creating cloud tenants*
 - *Creating a tenant*
 - *Creating a 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, 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.

Creating a tenant

For OpenStack tenants, you must provide admin user privileges to the tenant and the project before they can access to cloud resources.

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

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 it from the additional **Available Connectors** list.
4. If you want to associate this tenant with an existing role, select the role from the **Role** list; otherwise, leave **Create New Role** selected to have BIG-IQ Cloud create a new role associated with this specific tenant.
5. In the **Address**, **Phone**, and **Email** fields, type optional contact information for this tenant.
6. Click the **Save** button.

If you elected to have BIG-IQ Cloud create a new role for this cloud tenant, the new role displays in the Roles panel. You can now associate a user with this tenant to provide access to applications and services.

Creating a user

A *user* is an individual who will access 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® 11.3.0 and later. When you discover an F5 device, all iApps® 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 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 pair
 - configuring for the BIG-IQ system 18
- Amazon EC2 33
- See also EC2
 - and network configuration requirements 30
 - See also EC2
- Amazon elastic cloud services 30
- AMI
 - launching new 31
- application catalog 48
- applications
 - customizing for tenants 48
- application templates
 - 48
 - about 48
 - and vApps 36
 - defined 52
 - using 48
- asset management
 - for devices 23
- authentication 32

B

- BIG-IP devices
 - and BIG-IQ components required for management 20
- BIG-IQ Cloud
 - about 14
 - defined 52
 - finding documentation for 14
- BIG-IQ cloud tenant
 - finding documentation for 14
- BIG-IQ system
 - finding documentation for 14
 - reordering panel 15
- BIG-IQ system components
 - installing on managed BIG-IP devices 20

C

- catalog
 - for applications 48
- catalog entries
 - creating for tenants 48
- cloud administrator
 - defined 52
- cloud bursting
 - defined 52
- cloud connector
 - 36
 - defined 52
 - for OpenStack 39
 - local 42
- cloud connector for EC2 32
- cloud resources
 - providing for tenants 44, 48

- cloud resources (*continued*)
 - using locally 42
- cloud tenant
 - adding 44
- cloud tenants
 - about creating 44
- clusters
 - for high availability 18
- cNS, See VMware cNS
- component installation
 - script based 20
- connecting to servers
 - OpenStack 38
- connector
 - local 42
- connectors
 - VMware 36

D

- device discovery
 - about 21–22
- device inventory
 - viewing details 23
- device management
 - about 20
- devices
 - adding 21–22
 - discovering 20
- documentation, finding 14
- dynamic cloud resources
 - viewing activity for 34

E

- EC2 AMI
 - launching 31
- EC2 connector
 - associating with a device 32
- EC2 VPC
 - creating 30
- elasticity
 - viewing activity for 34

F

- failover 18
- filtering
 - objects 15

G

- glossary 52
- guides, finding 14

H

HA, See high availability cluster
high availability
 configuring 18
high availability configuration
 about 18

I

IAM

 creating user account 32

IAM access 33

iApps 48

See also application templates

 customizing for tenants 48

 defined 48

 See also application templates

inventory details

 viewing for devices 23

IP addresses

 for managed devices 20

L

license pool

 automatic activation 26

 manual activation 26

license pools

 about 26

licenses

 managing for devices 26

local cloud connector

 associating with a device 42

local cloud resources

 using 42

M

managed devices

 and BIG-IQ system components required 20

 discovering 20

manuals, finding 14

N

network configuration

 for integrating with OpenStack cloud services 38

 using Amazon EC2 30

 using VMware 36

network configurations

 customizing for tenants 48

network configurations*i*Apps

 customizing for tenants 48

network resources

 using for cloud services 42

O

objects

 finding associations 15

OpenStack

 required configuration 39

 using with BIG-IP VE systems 39

OpenStack cloud services

 and network configuration requirements 38

OpenStack connector

 associating with a device 39

OpenStack integration

 creating and managing connectors 38

P

package endpoint

 installing 20

panels

 reordering 14–15

R

release notes, finding 14

reports

 for asset management 23

resources

 defined 52

 providing access for user 45

RPM endpoint

 installing 20

S

script based

 component installation 20

T

tenant

 adding 44

tenant access

 using IAM 33

tenants

 about creating 44

 associating with a user 45

 creating applications for 48

tenants users

 and tenants 44

 and users 44

terminology 52

terms

 defined 52

U

user interface

 customizing 14–15

 navigating 14

users

 adding 44

 associating with a tenant 45

 defined 44

V

- vApps
 - and application templates 36
 - deploying 36
- virtual machine
 - launching new 31
- virtual private cloud
 - creating 30
- VMware
 - about integration 36
 - and network configuration requirements 36
- VMware (*continued*)
 - supported version 36
- VMware cloud connector
 - associating with a device 36
- VMware NSX
 - integrating with BIG-IQ Cloud 36
- VMware vCNS
 - integrating with BIG-IQ Cloud 36
- VMware vShield
 - integrating with BIG-IQ Cloud 36
- vShield, See VMware vShield

