

BIG-IQ[®] Cloud: Tenant User Guide

Version 4.5



Table of Contents

Legal Notices	5
Acknowledgments	7
Chapter 1: BIG-IQ User Interface Basics	15
About the BIG-IQ system user interface.....	16
Filtering for associated objects.....	16
Customizing panel order.....	16
Searching for specific objects.....	16
Chapter 2: Self-Service Application Deployment	17
About self-service application deployment	18
Deploying applications with automatic virtual server provisioning.....	18
Deploying applications.....	19
Chapter 3: Monitoring	21
About monitoring applications and application servers.....	22
Monitoring applications.....	22
Monitoring application servers.....	22
Viewing activity for cloud resources.....	23
Chapter 4: Integrating Amazon Web Services	25
About Amazon Web Services (AWS) integration.....	26
Using an API to create a Server.....	26
Creating a new application pool server.....	26
Chapter 5: Integrating VMware	29
About integrating VMware NSX with a BIG-IP VE.....	30
Provisioning a BIG-IP VE on NSX version 6.1.....	30
Using the API to define an NSX runtime deployment specification.....	32
Chapter 6: Integrating OpenStack	35
About OpenStack integration.....	36
Creating an OpenStack application server	36
Chapter 7: Glossary	37
BIG-IQ Cloud terminology.....	38

Legal Notices

Publication Date

This document was published on January 21, 2015.

Publication Number

MAN-0500-03

Copyright

Copyright © 2015, 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.

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 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.

This product includes angular-ui software, which is distributed under the MIT license. Copyright © 2012-2014, AngularUI Team.

This product includes CodeMirror software, which is distributed under the MIT license. Copyright © 2014, Marijn Haverbeke.

This product includes Quartz Scheduler software, which is distributed under the Apache 2.0 license. Copyright © Terracotta, Inc.

Chapter

1

BIG-IQ User Interface Basics

- *About the BIG-IQ system user interface*
- *Filtering for associated objects*
- *Customizing panel order*
- *Searching for specific objects*

About the BIG-IQ system user interface

The BIG-IQ® system interface is composed of panels. Each panel contains objects that correspond to a BIG-IQ feature. Depending on the number of panels and the resolution of your screen, some panels may be collapsed and show as colored bars 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 an object, and drag and drop it onto the object with 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. To display only items associated with a specific object, hover over the object, click the gear icon, and then select **Show Only Related Items**.
The screen refreshes to display only associated objects in each panel.
2. To highlight only items associated with a specific object, hover over the object, click the gear icon, and then select **Highlight Related Items**.
The screen refreshes, highlighting only associated objects in each panel, and displaying unassociated objects in a gray font.
3. To remove a filter, click the **X** icon next to the filtered object in a panel.

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.

Searching for specific objects

The BIG-IQ® system interface makes it easy to search for a specific object. This can be especially helpful as the number of objects increase when you add more users, applications, servers, and so forth.

1. To search for a specific object, in the Filter field at the top of the screen, type all or part of an object's name.
2. Click the **Apply** button.
The screen refreshes to display only the objects associated with the term you typed in the Filter field.
3. To further refine the filter, type another term into the Filter field, and click the **Apply** button again.
4. To remove a filter term, click the **X** icon next to it.

Chapter 2

Self-Service Application Deployment

- *About self-service application deployment*
- *Deploying applications with automatic virtual server provisioning*
- *Deploying applications*

About self-service application deployment

Cloud service providers customize iApps[®] application templates based on your needs as a cloud tenant. For example, they specify such things as an IP address for a virtual server, identify hosts and server pools, set connection limits, and so forth. This customization eliminates the need for you to perform complicated networking tasks, and ensures that your settings are optimized for your resources. When these customized applications are associated with you as a tenant, you have the option to further modify the applications as required, and deploy them as needed.

Additionally, BIG-IQ[®] Cloud provides you with optional self-service access to SSL certificates that you can deploy on demand to your managed BIG-IP[®] devices.

Deploying applications with automatic virtual server provisioning

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. Hover over the Applications header, and click the + icon when it appears.
4. In the **Name** field, type a name for this new application.
5. From the **Application Type** list, select an application.
6. 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.

7. For the **Provision Virtual Server IP** setting, select **Enable** and specify the FQDNs for the virtual servers to prompt BIG-IQ Cloud to automatically provision additional resources when traffic to your application increases.
8. For EC2 connectors, specify the following settings.
 - a) From the **Pool Server 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.
- 9. 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.
- 10. Alternatively, select the **Use Existing** option to use a SSL certificate and private key already stored on the device.
- 11. 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.
- 12. 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.

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. Hover over the Applications header, and click the + icon when it appears.
4. In the **Name** field, type a name for this new application.
5. From the **Application Type** list, select an application.
6. 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.
7. For the **Provision Virtual Server IP** setting, select **Enable** and specify the FQDNs for the virtual servers to prompt BIG-IQ Cloud to automatically provision additional resources when traffic to your application increases.
8. 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.
9. Alternatively, select the **Use Existing** option to use a SSL certificate and private key already stored on the device.

10. 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.

11. 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.

Chapter

3

Monitoring

- *About monitoring applications and application servers*
- *Monitoring applications*
- *Monitoring application servers*
- *Viewing activity for cloud resources*

About monitoring applications and application servers

As a tenant, you can use BIG-IQ® Cloud to monitor the health statistics and performance of applications and servers. In addition to the application itself, the health of an application is influenced by its associated objects, including:

- Servers that host applications
- Virtual servers that manage traffic to applications
- Connectors

Monitoring applications

Before you can monitor an application, you must first deploy it.

Monitoring statistics and performance details for applications and associated network objects provides you with the information you need to make resource management decisions. The application statistics are collected by the managed BIG-IP® device and include various network statistic, such as connections, bits per second, and so forth. The performance data displays the application performance trend over a period of time.

1. Log in to the BIG-IQ Cloud with your tenant user name and password.
2. On the Applications panel, click the gear icon next to the name of the application that you want to monitor.
The panel expands to display the application's properties.
3. To view the statistics, click **Statistics**.
The statistics display and all of the associated objects for that application are highlighted in the applicable panels.
4. To view the performance, click **Performance**
The performance graph displays.

Monitoring application servers

A cloud provider must have discovered, or you must have added, an application server before you can monitor it.

Monitoring health and performance statistics for your application servers provides you useful information about the health and usage for your resources. This information helps you decide when to increase or decrease resources to support your application requirements.

1. Log in to BIG-IQ Cloud with your tenant user name and password.
2. On the Server panel, click the gear icon next to the name of the application server that you want to monitor.
The panel expands to display the application monitor's properties.
3. To view the statistics, click **Statistics**.
The statistics display and all of the associated objects for that application are highlighted in the applicable panels.

4. To view the performance, click **Performance**
The performance graph displays.

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

4

Integrating Amazon Web Services

- *About Amazon Web Services (AWS) integration*
 - *Using an API to create a Server*
 - *Creating a new application pool server*
-

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.

Using an API to create a Server

Your provider must have already created a tenant user account that includes an EC2 connector, and supplied you with the connector reference.

Both tenants and providers can create EC2 servers using either the user interface or a REST API call. A user with a role of tenant can create an EC2 server using the REST API.

1. Authenticate with the F5 Cloud REST API, specifying a user role of Tenant.

Tip: Refer to *Authentication with the F5 REST API in the BIG-IQ Cloud Overview chapter of this guide for information about authentication strategies.*

Tip: Refer to the *BIG-IQ[®] Cloud Service API Reference Guide for details about using the APIs required for this task.*

2. Create a new EC2 server using the `Create Node` API to specify all of the details required for this server.

```
/cm/cloud/tenants/{tenant}/nodes POST
```

Important: Use the connector reference supplied by your provider when you are specifying the parameters for this server.

The REST response you receive confirms that the new server has the parameter values you specified.

The server you created is now available for your cloud applications to use.

Creating a new application pool server

Before you can create a new server, the cloud administrator must associate you with a cloud connector.

BIG-IQ Cloud provides you with the tools to self-deploy iApps through Amazon EC2 integration. For this integration, you must create an application pool server associated with an EC2 cloud connector.

1. Log in to the BIG-IQ Cloud with your tenant user name and password.
2. Hover on the Servers panel, click the + sign when it appears, and then click **New Server**.
3. From the **Cloud Connector** list, select the EC2 connector you want to associate with this server.
4. From the **Network Interface** list, select the network to associate with this server.
5. From the AMI ID list, select the AMI (Amazon Machine Image) select the associated image.
6. If your EC2 implementation is anything other than m1.small, in the **Instance Type** field, type the name of the instance.
For example, m3.large.
BIG-IQ Cloud supports the following large EC2 instances: m3.large, m3.xlarge, m3.2xlarge, cc2.8xlarge, c3.xlarge, c3.2xlarge, c3.4xlarge, and c3.8xlarge
7. Click the **Save** button.

You can now associate this server with your EC2 implementation.

Chapter 5

Integrating VMware

- *About integrating VMware NSX with a BIG-IP VE*

About integrating VMware NSX with a BIG-IP VE

BIG-IQ[®] Cloud provides you with the tools to manage VMware resources required to deliver highly available applications. Management tasks include discovering and creating BIG-IP[®] devices running in the private cloud. You can use this feature to accommodate seasonal traffic fluctuations by periodically adding and subtracting devices and application servers as needed. Additionally, you can provide NSX users access to self-deployable iApps[®] through VMware integration.

The tasks you perform to set up and configure BIG-IQ devices to manage BIG-IP system traffic in a VMware NSX version 6.1 network, use both the BIG-IQ software user interface and the VMware NSX user interface. There is also a task for which you can have greater control and flexibility using a REST API call to the NSX API. This optional task is included at the end of the task sequence.

In most production environments, data plane and control plane traffic are segregated for security reasons. To accomplish this topology, the network management for all devices is on the control plane subnet.

There are several setup tasks that you must perform before you can begin to configure the BIG-IQ[®] VMware-NSX integration to a BIG-IP VE device.

Important: For the most current instructions for performing these steps, refer to the VMware web site <http://pubs.vmware.com/>.

- You must have installed a BIG-IQ system with a management network subnet. This subnet will be used for provisioning and discovering BIG-IP devices. This subnet must be configured to include DHCP services and the DHCP configuration must include a default gateway.
- The DHCP IP pool must not include the IP address 192.168.1.245. This address is reserved for special use on the BIG-IP device.
- You must set up VMware NSX Manager and VMware vCenter to share the management network subnet that you configured for the BIG-IQ system. When the BIG-IP VE that you configure boots for the first time, it attaches to this shared network.
- You must configure the following objects in VMware vSphere Web Client before you can perform the VMware NSX integration.
 - A data center
 - A data store for your data center
 - A cluster

Provisioning a BIG-IP VE on NSX version 6.1

BIG-IQ[®] software's NSX integration supports provisioning of a BIG-IP[®] VE instance to provide load-balancing services in the context of an NSX Edge.

Important: You perform the following step-sequence using the vSphere Web Client user interface. At time of release, these steps accurately describe the VMware user interface. For the most current instructions for performing these steps, refer to the VMware web site <http://pubs.vmware.com/>.

1. In the vSphere web client user interface, create a new NSX Edge in an undeployed state.

If you specify a tenant ID when you create the Edge, BIG-IQ software will create a tenant with that ID when it creates the BIG-IP VE.

Make sure that the NSX Edge you create identifies the Cluster/Resource Pool and the Datastore, but does not identify any interfaces. Otherwise, follow your standard practice for NSX Edge creation.

2. For the just created NSX Edge, navigate to the Manage tab, and then select the Load Balancer tab. Then click the **Edit** button.
The Edit Load balancer global configuration window displays.
3. Select both **Enable Load Balancer** and **Enable Service Insertion**.
4. For the **Service Definition**, select the name of the connector you created for NSX in a previous step.
5. For the **Service Configuration**, select **F5 ADC - Make a BIG-IP VE**.
6. Expand Typed Service Configuration Parameters and then determine which of these optional settings you want to specify.
 - a) In the **Value** field next to **F5-BIG-IP-VE** key, type `yes`.
 - b) If you want to specify a fully qualified host name of the node template to describe the kind of BIG-IP VE this template creates, in the **Value** field next to **F5-BIG-IP-VE-FQ-HOST-NAME** key, type that name in the value box next to **Name of BIG-IP node template?**.

***Tip:** This step is optional. If you do not specify a host name, the template uses the default host name.*

- c) Specify the name of the node template to describe the kind of BIG-IP VE this template creates; type that name in the **Value** field next to **F5-BIG-IP-VE-OVF-NAME** key.

***Tip:** If you have already created a node template, than specify the name of it here. If you are creating a node template, then specify a name to go along with the URL that you specify in the next step.*

- d) In the **Value** field next to the **F5-BIG-IP-VE-OVF-URL** key, type in the URL that describes the location of the OVF file that the BIG-IQ device uses to create the BIG-IP VE.

***Tip:** This step is optional. You only need to specify the URL if the node template has not already been created.*

***Tip:** You can also specify this value using an API call.*

- e) If you want to specify an admin password so that you can easily log in as administrator to this BIG-IP VE from NSX, type the password in the **Value** field next to the **F5-BIG-IP-VE-ADMIN-PASSWORD** key.

***Tip:** If you choose to let the BIG-IQ system generate the password, you can view the password in the BIG-IQ device Servers panel.*

***Tip:** When the BIG-IP VE is initially provisioned, root login is disabled. To access the VE using root login, you must log in as `admin` and set the root password.*

7. Expand Service Instance Runtime Configuration, and then use the controls to specify settings for up to four virtual network interface controllers (vNICs).

For each of the vNICs you specify, the **IP Allocation Mode** must be **IP Pool**.

- The first required vNIC (vNIC0) provides the DHCP-enabled control plane network on which the BIG-IP VE boots. Choose the name specified previously that corresponds to the IP pool `192.168.11.0/24`.
- The next required vNIC (vNIC1) you specify provides the external data network on which the BIG-IP device creates virtual servers. Choose the name specified previously that corresponds to the IP pool `10.22.0.0/16`.
- The first optional vNIC you specify provides the internal data network on which load-balanced pool members are located. Choose the name specified previously that corresponds to the IP pool `10.33.0.0/16`.

- The next optional vNIC you specify provides the data plane network on which the BIG-IQ device discovers and manages BIG-IP devices. Choose the name specified previously that corresponds to the IP pool 192.44.0.0/16.
8. Click **OK** to close the Edit Load balancer global configuration dialog box. VMware NSX configures the Edge Gateway based on the settings you specified.

When you finish editing an Edge with the settings described in this task, BIG-IQ software responds by creating and licensing the BIG-IP VE.

Using the API to define an NSX runtime deployment specification

VMware NSX uses a Runtime Deployment to specify parameters for BIG-IP[®] virtual devices provisioned using a BIG-IQ[®] software connection. Node templates simplify the task of specifying the parameters for the Runtime Deployment. This task uses the `Create node template` API to create a node template. The BIG-IQ system and NSX integration uses this template when it provisions new BIG-IP virtual devices.

Important: *Using an API call to perform this task is optional. If you want to use the NSX user interface to specify the node template, you can do that. However, if you want to create the template in advance or see a list of existing templates before you define a new one, you can use a REST compliant HTTP request to execute an API call. To facilitate the process of submitting REST API calls, F5[®] includes an API management tool called Presentation Manager. This task steps you through its use.*

1. Use a web browser to access and log in to the BIG-IQ device.
`https://<BIG-IQ IP address>`
2. Use the Presentation Manager API tool to access the `Create node template` URL.
`https://<BIG-IQ IP address>/mgmt/cm/cloud/connectors/vmware-nsx/presentation`
The Presentation Manager interface opens for the `Create node template` API.
3. Click **Table of Contents**.
A lengthy list of API endpoints is displayed.
4. From the list of API endpoints, locate the connector just created in the previous task.
The connector will look something like this:
`/mgmt/cm/cloud/connectors/vmware-nsx/<connectorId>/nodes`
5. In the upper right corner, click the plus sign, and then scroll to the very bottom of the page and click the **Advanced** button.
A small field, titled **JSON Input** opens.
6. In the **JSON Input** field, type the values for three property IDs needed to register the node template as a deployment specification.
 - The `OvfUrl` entry identifies the URL specified previously for the OVF file that the BIG-IQ device uses to create the BIG-IP VE.
 - The `BIG-IP` entry set to true indicates that the template specifies provisioning details for a BIG-IP device.
 - The `NodeTemplateName` entry identifies the name you want NSX users to specify when requesting deployment of this type of BIG-IP VE.

```
{
  "state": "TEMPLATE",
  "properties": [
    {
      "id": "BIG-IP",
      "provider": "true"
    }
  ]
}
```



```
    },  
    {  
      "id": "NodeTemplateName",  
      "value": "BIGIP-11.5.0.0.0.221.LTM_1SLOT-scsi.ovf"  
    },  
    {  
      "id": "OvfUrl",  
      "provider":  
"http://server/ovfs/BIGIP-11.5.0.0.0.221.LTM_1SLOT-scsi/BIGIP-11.5.0.0.0.221-scsi.ovf"  
    }  
  ]  
}
```

7. Click **Save**.

Presentation Manager submits the REST API call with the JSON body you specified.

The API call registers the deployment specification received from the NSX API with the BIG-IQ software's NSX Partner Service. The REST API response includes the property ID `ImageId`. This value identifies the just-created deployment specification that confirms that the connection between the BIG-IQ system and the NSX device is established.

Chapter 6

Integrating OpenStack

- *About OpenStack integration*
 - *Creating an OpenStack application server*
-

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.

Creating an OpenStack application server

Before you can integrate BIG-IQ Cloud with OpenStack, you must deploy an application and associate it with a OpenStack connector.

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`.

You can leverage OpenStack servers to host the resources that your applications require. This eliminates the need for you to purchase and manage hardware and allows you to expand and retract virtual resources as needed.

1. Log in to BIG-IQ Cloud with your tenant user name and password.
2. Hover on the Servers header, and click the + icon when it appears.
The panel expands to display fields for the new server details.
3. From the **Cloud Connector** list, select OpenStack.
4. From the **Network Interface** list, select the interface from which you are accessing the server.
The **Address** field populates with the IP address of the virtual machine in the OpenStack network.
5. In the **Name** field, type a name for this application server.
6. From **Image ID** list, select the OpenStack virtual machine image.
7. In the **OpenStack Flavor** field, type the name of the OpenStack flavor that you received from OpenStack.
A flavor is an available hardware configuration for this server. Each flavor has a unique combination of disk space and memory capacity.
8. Click the **Save** button.

You can now deploy applications to the server you created.

Chapter

7

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

- activities
 - viewing for cloud resource activity 23
- Amazon CloudWatch
 - about integrating with BIG-IQ Cloud 26
- Amazon EC2
 - about integrating with BIG-IQ Cloud 26
 - deploying applications to 18
- Amazon EC2 application servers
 - monitoring traffic 22
- Amazon EC2 resources
 - viewing cloud resource activity 23
- Amazon Elastic Compute Cloud, *See* Amazon EC2
- Amazon Virtual Private Cloud
 - about integrating with BIG-IQ Cloud 26
- application catalogs
 - about deploying 18
- applications
 - deploying 18
 - deploying as a Cloud Tenant user 18–19
 - monitoring health and performance 22
 - viewing statistics for 22
- application servers
 - monitoring health and performance 22
- application services
 - about deploying 18
- application templates
 - defined 38
 - deploying as a Cloud Tenant user 18–19
- application traffic
 - monitoring 22

B

- BIG-IP devices
 - provisioning and discovering 30
- BIG-IP VE
 - provisioning 30
 - setting up VMware network 30
- BIG-IQ Cloud
 - defined 38
- BIG-IQ system
 - reordering panels 16

C

- cloud administrator
 - defined 38
- cloud bursting
 - defined 38
- cloud connector
 - defined 38
- Cloud Tenant users
 - deploying applications 18–19

D

- device elasticity
 - configuring 18
- dynamic cloud resources
 - viewing activity for 23

E

- EC2 integration
 - using a large instance 26
- elasticity
 - configuring for tenants 18
 - viewing activity for 23
- elastic resources
 - specifying settings for 18

F

- filtering process
 - finding associated objects 16

G

- glossary 38
- Grizzly, *See* OpenStack

H

- Havana, *See* OpenStack
- health statistics
 - monitoring 22

I

- iApps
 - for OpenStack 36

M

- monitoring health and performance 22

N

- node creation
 - as tenant 26
- nodes
 - creating as a tenant 26
- NSX devices
 - about provisioning 30
 - connecting to 32
- NSX integration
 - about 30
- NSX runtime deployment specification
 - configuring with API calls 32
 - registering 32

Index

NSX service
about 30

O

objects
finding associations 16
searching for 16
OpenStack
and iApps 36
OpenStack application server
creating 36
OpenStack application servers
monitoring traffic 22

P

panels
reordering 16
performance
for application servers 22
viewing for applications 22
provisioning process
for NSX device 30

R

resource elasticity
configuring for tenants 18
resources
defined 38

S

search function
finding specific objects 16
self-service application deployment 18
server elasticity
configuring 18
server health
for Amazon EC2 22
for OpenStack 22
servers 22
statistics
about 22
for Amazon EC2 server traffic 22
for application server traffic 22
viewing for applications 22

T

terminology 38
terms
defined 38

U

user interface
and searching for specific objects 16
customizing 16
navigating 16

V

Virtual Private Cloud, See VPC
VMware network
for a BIG-IP VE 30