

# **BIG-IQ<sup>®</sup> Cloud: Tenant User Guide**

Version 4.4





# Table of Contents

<b>Legal Notices.....</b>	<b>5</b>
<b>Acknowledgments.....</b>	<b>7</b>
 <b>Chapter 1: BIG-IQ User Interface Basics.....</b>	 <b>15</b>
About the BIG-IQ system user interface.....	16
Filtering for associated objects.....	16
Customizing panel order.....	16
Filtering on multiple objects.....	16
 <b>Chapter 2: Self-Service Application Deployment.....</b>	 <b>19</b>
About self-service application deployment .....	20
Deploying applications.....	20
 <b>Chapter 3: Monitoring.....</b>	 <b>23</b>
About monitoring applications and application servers.....	24
Monitoring applications.....	24
Monitoring application servers.....	24
 <b>Chapter 4: Integrating Amazon Web Services .....</b>	 <b>27</b>
About Amazon Web Services (AWS) integration.....	28
Creating an EC2 node as a tenant.....	28
Viewing activity for cloud resources.....	28
 <b>Chapter 5: Integrating VMware.....</b>	 <b>31</b>
About VMware NSX version 6.1 integration.....	32
Provisioning a BIG-IP VE on NSX.....	32
Using the API to define an NSX runtime deployment specification.....	37
 <b>Chapter 6: Integrating OpenStack .....</b>	 <b>39</b>
About OpenStack integration.....	40
Creating an OpenStack application server .....	40
 <b>Chapter 7: Glossary.....</b>	 <b>41</b>
BIG-IQ Cloud terminology.....	42



# Legal Notices

---

## Publication Date

This document was published on September 8, 2014.

## Publication Number

MAN-0500-02

## Copyright

Copyright © 2014, F5 Networks, Inc. All rights reserved.

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

## Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, Application Acceleration Manager, Application Security Manager, APM, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, iCall, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, LineRate, LineRate Systems [DESIGN], Local Traffic Manager, LROS, LTM, Message Security Manager, MobileSafe, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, SDAC (except in Japan), SDC, Signalling Delivery Controller, Solutions for an application world, Software Designed Applications Services, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix (except Germany), Traffix [DESIGN] (except Germany), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

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

## Patents

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

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

## Export Regulation Notice

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

### **RF Interference Warning**

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

### **FCC Compliance**

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

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

### **Canadian Regulatory Compliance**

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

### **Standards Compliance**

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

# Acknowledgments

---

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

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

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

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

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

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

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

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

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

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

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

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

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

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

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

This product includes software developed by Christos Zoulas.

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

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

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

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

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

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

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

This product includes software developed by Jared Minch.

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

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

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

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

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

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

This product includes Hypersonic SQL.

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

This product includes software developed by the Internet Software Consortium.

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

This product includes software developed by jQuery Foundation and other contributors, distributed under the MIT License. Copyright ©2014 jQuery Foundation and other contributors (<http://jquery.com/>).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

This product includes DP\_DateExtensions software developed by Jim Davis, Copyright © 1996-2004, The Depressed Press of Boston (depressedpres.com). All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the DEPRESSED PRESS OF BOSTON (DEPRESSEDPRESS.COM) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



---

# Chapter

# 1

---

## BIG-IQ User Interface Basics

---

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

## 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.

## Filtering on multiple 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. 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 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.

3. To remove a filter, click the **x** icon next to the object that you want to remove, below the Filter field.



---

# Chapter 2

---

## Self-Service Application Deployment

---

- *About self-service application deployment*
  - *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

---

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

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

1. Log in to the BIG-IQ Cloud with your tenant user name and password.
2. Hover over the Applications header, and click the + icon when it appears.
3. In the **Name** field, type a name for this new application.
4. From the **Application Type** list, select an application.
5. From the **Cloud Connector** list, select the cloud connector associated with where you want to deploy your application.

A *cloud connector* is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.

6. To configure BIG-IQ Cloud to automatically provision additional resources when traffic to your application increases, select **Enable** from the **Server Elasticity** list and specify the settings for the server elasticity options that display.

This option is available only for the EC2 connector. For automatic server provisioning to work, your cloud service provider must have enabled the **Server Elasticity** setting for this EC2 connector.

- a) From the **Node Image** list, select the image from which to create new application servers when capacity is met and additional servers are required.
- b) In the **Min. # of Servers** field, type the minimum number of application servers you want running at any given time.
- c) In the **Max. # of Servers** field, type the maximum number of application servers you want running when additional servers are required.
- d) From the **Monitor By** list, select the category associated with the statistic on which you want to base the threshold value.
- e) For the **When** setting, select a specific statistic, the associated relational operator, and a type a number in the field for the threshold.

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

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

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



---

# Chapter

# 3

---

## Monitoring

---

- *About monitoring applications and application servers*
- *Monitoring applications*
- *Monitoring application servers*

### 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.



---

# Chapter

# 4

---

## Integrating Amazon Web Services

---

- *About Amazon Web Services (AWS) integration*
  - *Creating an EC2 node as a tenant*
  - *Viewing activity for cloud resources*
-

## About Amazon Web Services (AWS) integration

---

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

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

## Creating an EC2 node as a tenant

---

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 nodes using either the user interface or a REST API call. A user with a role of tenant can create an EC2 node 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 node using the `Create Node` API to specify all of the details required for this node.  
`/cm/cloud/tenants/{tenant}/nodes` POST

---

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

---

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

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

## 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 5

---

## Integrating VMware

---

- *About VMware NSX version 6.1 integration* |

## About VMware NSX version 6.1 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 accommodate this requirement, traffic management functions are not permitted on the same network subnet with flowing network traffic. To accomplish this topology, this integration configures a total of four subnets. Two are used for BIG-IQ network management and the other two are for BIG-IP system traffic flow.

### Task summary

## Provisioning a BIG-IP VE on NSX

Configuring and licensing a BIG-IP virtual device on VMware NSX creates a Service Instance Runtime (SIR). Creating this SIR allows vCenter users to deploy load-balanced virtual servers (referred to in the VMware interface as Service Profiles). You can make one virtual device for each defined NSX Edge.

---

**Tip:** The objective of this task is to assemble a data parameter and use it as the payload for a REST API call that updates the configuration of the load balancer and specifies an NSX SIR. Because this parameter can be somewhat complex, you will probably want to use a tool that allows you to assemble multiple lines of content together to include in your REST API call. A script editor or a command line tool that can send files using URL syntax (such as curl) may be very helpful.

---

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

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. Authenticate with the VMware NSX REST API.

---

**Important:** Refer to the *VMware API Reference* documentation for the most authoritative detail about using the VMware APIs required for this task.

---

3. Use a REST API call to the NSX API to retrieve information about the pools and monitors used by the NSX Edge. Copy the REST response to an editor for use in compiling the remainder of the data parameter.

<NSX-Manager-IP-address>/api/4.0/edges/<NSX-Manager-IP-address>loadbalancer/config  
GET

- a) After you copy the content to the editor, edit the value of the `enabled` and `enableServiceInsertion` parameters. Change the value of each parameter to `true`.

The following example illustrates a typical response to this API call.

```
<loadBalancer>
  <version>1</version>
  <enabled>true</enabled>
  <enableServiceInsertion>true</enableServiceInsertion>
  <accelerationEnabled>>false</accelerationEnabled>
```



```

<monitor>
  <monitorId>monitor-1</monitorId>
  <type>tcp</type>
  <interval>5</interval>
  <timeout>15</timeout>
  <maxRetries>3</maxRetries>
  <name>default_tcp_monitor</name>
</monitor>
<monitor>
  <monitorId>monitor-2</monitorId>
  <type>http</type>
  <interval>5</interval>
  <timeout>15</timeout>
  <maxRetries>3</maxRetries>
  <method>GET</method>
  <url>/</url>
  <name>default_http_monitor</name>
</monitor>
<monitor>
  <monitorId>monitor-3</monitorId>
  <type>https</type>
  <interval>5</interval>
  <timeout>15</timeout>
  <maxRetries>3</maxRetries>
  <method>GET</method>
  <url>/</url>
  <name>default_https_monitor</name>
</monitor>
<logging>
  <enable>>false</enable>
  <logLevel>info</logLevel>
</logging>
</loadBalancer>

```

4. Configure the next element in the data parameters for the REST API call that will create your new VMware NSX Service Instance.

---

**Important:** Insert the parameters assembled in this step just before the end of the `</loadBalancer>` element obtained in the previous step.

---

- a) Specify the description and name fields. Specify values that help you identify the service and its purpose.

The following code snippet illustrates the field values discussed in this step.

```

<globalServiceInstance>
  <description>edge-16__LB__dummy</description>
  <name>edge-16__LB__dummy</name>

```

- b) Use a REST API call to the NSX API to determine the value of the `serviceId` parameter.

You can see all of the NSX connectors using

```
<NSX-Manager-IP-address>/api/2.0/si/servicemanagers GET
```

In the REST response, look for the `<basicinfo>` element in the `<services>` element, and then find the name of the connector you provided when you created the NSX connector. Use the `<objectId>` associated with that connector name as the value for the `serviceId` parameter.

- c) Use the just obtained `serviceId` to make a REST API call to the NSX API to determine the value of the `serviceName` parameter.

```
<NSX-Manager-IP-address>/api/2.0/si/service/<serviceID> GET
```

When you get the REST response, locate the `<name>` parameter and use that value for the `serviceName` parameter.

In the following example, the name parameter value is:

NSX-conn-sample-20140717-113002-22123-0.

```
<objectId>service-446</objectId>
<objectTypeName>Service</objectTypeName>
<vsmUuid>42388336-FE9B-225A-ECAF-1085628911B2</vsmUuid>
<revision>6</revision>
<type>
  <typeName>Service</typeName>
</type>
<name>NSX-conn-sample-20140717-113002-22123-0</name>
```

- d) Use a REST API call to the NSX API to determine the value of the `instanceTemplateUniqueId` and `instanceTemplateId` parameters.

<NSX-Manager-IP-address>/api/2.0/si/service/<serviceID> GET

When you get the REST response, look in the `<serviceInstanceTemplates>` element and locate the name of the service instance template that you plan to use (the default template that is built when you create an NSX connection is F5 ADC - Make a BIG-IP VE. Locate the `<id>` parameter and use that numeric string value for the `instanceTemplateId` parameters.

5. For the `instanceTemplateTypedAttributes` element, specify the type attributes that are required for the service instance template.

There are several attributes to configure. Some are required, some are optional. Each one is identified by the value of the attribute key.

- For the F5-BIG-IP-MAKE-VE key, leave the `<name>` field alone and specify yes for the `<value>` field
- For the F5-BIG-IP-VE-FQ-HOST-NAME key, specify the host name that you want the template to use, or leave it blank to use the default host name.

---

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

---

- For the F5-BIG-IP-VE-OVF key, specify the name you want to specify for the template.

---

**Tip:** If you have already created a node template, then 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.

---

- For the F5-BIG-IP-VE-OVF-URL key, specify the URL that describes the location of the OVF file that the BIG-IP 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.

---

- For the F5-BIG-IP-VE-ADMIN-PASSWORD key, specify the admin password so that you can easily log in as admin to this BIG-IP VE from NSX.

---

**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.

---



---

**Tip:** If you choose to let NSX generate a password, you can view the password by submitting a REST API call: `mgmt/cm/cloud/connectors/vmware-nsx/<connectorId>/nodes<nodeID> GET`. Examine the REST response for the value of the `DeviceMgmtPassword` field.

---

6. For the `runtimeNicInfo` field, specify details for each of the four network interfaces (NICs).

Option	Description
<b>index</b>	The value of this field indicates which of the four NICs you are configuring.
<b>label</b>	Specifies a name for the NIC you are configuring.
<b>objectId field in the network parameter</b>	Specifies which network interface to use for the NIC you are configuring.
<b>connectivityType</b>	Specifies whether the NIC you are configuring is for data (DATA) or control plane (MGMT) use.
<b>ipAllocationType</b>	Specifies the mode to use for IP address allocation. Generally, IP_POOL, is used, but you can also use MANUAL or DHCP.
<b>objectId field in the ipPool parameter</b>	Identifies which IP pool to use for the IP addresses allocated to this NIC.

- The first vNIC you specify 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 vNIC you edit provides the external data network on which the BIG-IP creates virtual servers. Choose the name specified previously that corresponds to the IP pool 10.22.0.0/16.
- The next vNIC you edit 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 last vNIC you edit provides the control 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 192.168.44.0/24.

7. In the `objectId` field of the `service` parameter, specify the name of the service manager you want to use.

You can retrieve this name using `<NSX-Manager-IP-address>/api/2.0/si/serviceinstance` GET

8. Once you have completed the data parameter assembly, you can create the SIR by making a REST API call to the NSX API.

`<NSX-Manager-IP-address>api/4.0/edges/<edge objectId>/loadbalancer/config` PUT

The following example illustrates a typical data parameter for this API call.

```
<loadBalancer>
  <version>1</version>
  <enabled>>false</enabled>
  <enableServiceInsertion>>false</enableServiceInsertion>
  <accelerationEnabled>>false</accelerationEnabled>
  <monitor>
    <monitorId>monitor-1</monitorId>
    <type>tcp</type>
    <interval>5</interval>
    <timeout>15</timeout>
    <maxRetries>3</maxRetries>
    <name>default_tcp_monitor</name>
  </monitor>
  <monitor>
    <monitorId>monitor-2</monitorId>
    <type>http</type>
    <interval>5</interval>
    <timeout>15</timeout>
    <maxRetries>3</maxRetries>
    <method>GET</method>
    <url>/</url>
    <name>default_http_monitor</name>
```

```

</monitor>
<monitor>
  <monitorId>monitor-3</monitorId>
  <type>https</type>
  <interval>5</interval>
  <timeout>15</timeout>
  <maxRetries>3</maxRetries>
  <method>GET</method>
  <url>/</url>
  <name>default_https_monitor</name>
</monitor>
<logging>
  <enable>false</enable>
  <logLevel>info</logLevel>
</logging>
<instanceTemplateTypedAttributes>
  <typedAttribute>
    <key>F5-BIG-IP-MAKE-VE</key>
    <name>F5-BIG-IP-MAKE-VE</name>
    <value>yes</value>
    <type>STRING</type>
  </typedAttribute>
  <typedAttribute>
    <key>F5-BIG-IP-VE-OVF</key>
    <name>Name of BIG-IP node template?</name>
    <value />
    <type>STRING</type>
  </typedAttribute>
  <typedAttribute>
    <key>F5-BIG-IP-VE-FQ-HOST-NAME</key>
    <name>Fully qualified host name of BIG-IP VE? Optional - remains default if
unspecified.</name>
    <value />
    <type>STRING</type>
  </typedAttribute>
  <typedAttribute>
    <key>F5-BIG-IP-VE-ADMIN-PASSWORD</key>
    <name>Admin password of BIG-IP VE? Leave blank to have random strong password
generated.</name>
    <value />
    <type>STRING</type>
  </typedAttribute>
</instanceTemplateTypedAttributes>
<baseRuntimeConfig>
  <deploymentScope>
    <resourcePool>domain-c35</resourcePool>
    <datastore>datastore-49</datastore>
    <nics>
      <runtimeNicInfo>
        <index>0</index>
        <label>mgmt</label>
        <network>
          <objectId>network-44</objectId>
        </network>
        <connectivityType>MGMT</connectivityType>
        <ipAllocationType>IP_POOL</ipAllocationType>
        <ipPool>
          <objectId>ipaddresspool-4</objectId>
        </ipPool>
        <ipAddress>10.144.63.213</ipAddress>
        <subnetMask>255.255.240.0</subnetMask>
        <gatewayAddress>10.144.63.203</gatewayAddress>
      </runtimeNicInfo>
      <runtimeNicInfo>
        <index>1</index>
        <label>external</label>
        <network>
          <objectId>network-46</objectId>
        </network>
        <connectivityType>DATA</connectivityType>

```

```

        <ipAllocationType>IP_POOL</ipAllocationType>
        <ipPool>
          <objectId>ipaddresspool-3</objectId>
        </ipPool>
      </runtimeNicInfo>
      <runtimeNicInfo>
        <index>2</index>
        <label>internal</label>
        <network>
          <objectId>network-47</objectId>
        </network>
        <connectivityType>DATA</connectivityType>
        <ipAllocationType>IP_POOL</ipAllocationType>
        <ipPool>
          <objectId>ipaddresspool-2</objectId>
        </ipPool>
      </runtimeNicInfo>
      <runtimeNicInfo>
        <index>3</index>
        <label>ha</label>
        <network>
          <objectId>network-47</objectId>
        </network>
        <connectivityType>MGMT</connectivityType>
        <ipAllocationType>IP_POOL</ipAllocationType>
        <ipPool>
          <objectId>ipaddresspool-2</objectId>
        </ipPool>
      </runtimeNicInfo>
    </nics>
  </deploymentScope>
</baseRuntimeConfig>
</config>
<service>
  <objectId>input-service-id-here</objectId>
</service>
</serviceInstance>
</loadBalancer>

```

After you submit the API call, VMware NSX and the BIG-IQ software engage in a sequence of calls and responses. When this sequence completes successfully, the NSX API status returns an `IN_SERVICE` message to indicate that the device provisioning is complete.

You can confirm that the provisioning sequence completed successfully, by making a REST API call to the NSX API similar to this one: `<NSX-Manager-IP_address>/api/2.0/si/serviceinstance/<serviceinstance-id>/runtimeinfo/<serviceruntime-id> GET`

The API response `IN_SERVICE` indicates that provisioning completed successfully.

---

**Tip:** If the status does not update to indicate success, you may be able to determine where the sequence broke down by analyzing the call and response sequence. For detail on this sequence, refer to *Provisioning call flow sequence*.

---

## 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 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 your 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 box, titled JSON Input will appear.
6. In the JSON Input box, 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 create a new application server, your cloud administrator must associate you as a user to an OpenStack connector. You must know what the OpenStack virtual machine and flavor are assigned to you by OpenStack.

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 [28](#)
- Amazon CloudWatch
  - about integrating with BIG-IQ Cloud [28](#)
- Amazon EC2
  - about integrating with BIG-IQ Cloud [28](#)
- Amazon EC2 application servers
  - monitoring traffic [24](#)
- Amazon EC2 resources
  - viewing cloud resource activity [28](#)
- Amazon Elastic Computer Cloud, See Amazon EC2
- Amazon Virtual Private Cloud
  - about integrating with BIG-IQ Cloud [28](#)
- application catalogs
  - about deploying [20](#)
- applications
  - deploying [20](#)
  - deploying as a Cloud Tenant user [20](#)
  - monitoring health and performance [24](#)
  - viewing statistics for [24](#)
- application servers
  - monitoring health and performance [24](#)
- application services
  - about deploying [20](#)
- application templates
  - defined [42](#)
  - deploying as a Cloud Tenant user [20](#)
- application traffic
  - monitoring [24](#)

## B

- BIG-IP device
  - provisioning [32](#)
- BIG-IQ Cloud
  - defined [42](#)
- BIG-IQ system
  - reordering panel [16](#)

## C

- cloud administrator
  - defined [42](#)
- cloud bursting
  - defined [42](#)
- cloud connector
  - defined [42](#)
- Cloud Tenant users
  - deploying applications [20](#)

## D

- device elasticity
  - configuring [20](#)

- dynamic cloud resources
  - viewing activity for [28](#)

## E

- elasticity
  - configuring for tenants [20](#)
  - viewing activity for [28](#)
- elastic resources
  - specifying settings for [20](#)

## F

- filtering
  - objects [16](#)

## G

- glossary [42](#)
- Grizzly, See OpenStack

## H

- Havana, See OpenStack
- health statistics
  - monitoring [24](#)

## I

- iApps
  - for OpenStack [40](#)

## M

- monitoring health and performance [24](#)

## N

- node creation
  - as tenant [28](#)
- nodes
  - creating as a tenant [28](#)
- NSX devices
  - connecting to [37](#)
  - provisioning [32](#)
- NSX integration
  - about [32](#)
- NSX runtime deployment specification
  - configuring with API calls [37](#)
  - registering [37](#)

## O

- objects
  - finding associations [16](#)
  - searching for [16](#)

## Index

- OpenStack
  - and iApps [40](#)
- OpenStack application server
  - creating [40](#)
- OpenStack application servers
  - monitoring traffic [24](#)

## P

- panels
  - reordering [16](#)
- performance
  - for application servers [24](#)
  - viewing for applications [24](#)
- provisioning process
  - for NSX device [32](#)

## R

- resource elasticity
  - configuring for tenants [20](#)
- resources
  - defined [42](#)

## S

- search
  - for specific objects [16](#)
- self-service application deployment [20](#)

- server elasticity
  - configuring [20](#)
- server health
  - for Amazon EC2 [24](#)
  - for OpenStack [24](#)
- servers [24](#)
- statistics
  - about [24](#)
  - for Amazon EC2 server traffic [24](#)
  - for application server traffic [24](#)
  - viewing for applications [24](#)

## T

- terminology [42](#)
- terms
  - defined [42](#)

## U

- user interface
  - customizing [16](#)
  - navigating [16](#)
  - searching for specific objects [16](#)

## V

- Virtual Private Cloud [28](#)
  - See also VPC