

BIG-IQ™ Cloud: Tenant User Guide

Version 4.2



Table of Contents

Legal Notices	5
Acknowledgments	7
Chapter 1: BIG-IQ User Interface Basics	13
About the BIG-IQ system user interface.....	14
Filtering for associated objects.....	14
Customizing panel order.....	14
Filtering on multiple objects.....	14
Chapter 2: Self-Service Application Deployment	17
About self-service application deployment	18
Deploying applications.....	18
Chapter 3: Monitoring	21
About monitoring applications and application servers.....	22
Monitoring applications.....	22
Monitoring application servers.....	22
Chapter 4: Amazon EC2 Servers	25
About Amazon EC2 servers.....	26
Creating an Amazon EC2 server	26
Chapter 5: OpenStack Cloud Integration	27
About OpenStack servers	28
Creating an OpenStack application server	28
Chapter 6: Glossary	29
BIG-IQ Cloud terminology.....	30

Legal Notices

Publication Date

This document was published on December 20, 2013.

Publication Number

MAN-0500-00

Copyright

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

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

Trademarks

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

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

Patents

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

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

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 Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

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

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

Acknowledgments

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

This product includes 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.

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.

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 on 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.
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 enable 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. 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.

8. 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

Amazon EC2 Servers

- *About Amazon EC2 servers*
 - *Creating an Amazon EC2 server*
-

About Amazon EC2 servers

You can leverage Amazon EC2 cloud servers to host the resources that your applications require. After you create a new server, you can view its statistics and performance to monitor your resource load.

Creating an Amazon EC2 server

Before you can create a new server, your cloud administrator must associate you as a user to an Amazon EC2 cloud connector. You must also have the Amazon Machine Image identification number assigned to you by Amazon.

You can leverage Amazon EC2 cloud 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. In the **AMI ID** field, type the identification for the Amazon Machine Image.
4. In the **Key Pair Name** field, type the name of an existing EC2 key pair.
This step is optional. If you do not provide the name of an existing key pair, BIG-IQ Cloud will create one for you.
5. Click the **Save** button.

You can now deploy applications to the server you created.

Chapter 5

OpenStack Cloud Integration

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

About OpenStack servers

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

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

Task summary

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 6

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 allows efficient management of resources during periods of increased or decreased traffic to applications.
<i>cloud connector</i>	A cloud connector is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.
<i>resources</i>	A resource is any managed object, including devices, web applications, virtual servers, servers, cloud connectors, and so forth.
<i>roles</i>	A role defines specific privileges to which you can associate one or more users. There are two roles for BIG-IQ Cloud: cloud administrator and cloud tenant.
<i>tenant</i>	A tenant is an entity that can consist of one or more users accessing resources provided by a cloud administrator.
<i>user</i>	A user is an individual who has been granted access to specific tenant resources.

Index

A

- adding application server *28*
- Amazon EC2
 - creating servers for *26*
- Amazon EC2 application servers
 - monitoring traffic *22*
- Amazon EC2 cloud servers
 - about creating *26*
- Amazon EC2 servers
 - creating *26*
- Amazon Machine Image
 - for Amazon EC2 servers *26*
- application catalogs
 - about deploying *18*
- applications
 - deploying *18*
 - deploying as a Cloud Tenant user *18*
 - monitoring health and performance *22*
 - viewing statistics for *22*
- application servers
 - monitoring health and performance *22*
- application services
 - about deploying *18*
- application templates
 - defined *30*
 - deploying as a Cloud Tenant user *18*
- application traffic
 - monitoring *22*

B

- BIG-IQ Cloud
 - defined *30*
- BIG-IQ system
 - reordering panel *14*

C

- cloud administrator
 - defined *30*
- cloud bursting
 - defined *30*
- cloud connector
 - defined *30*
- Cloud Tenant users
 - deploying applications *18*
- connecting to servers
 - OpenStack *28*

D

- device elasticity
 - configuring *18*

E

- elasticity
 - configuring for tenants *18*
- elastic resources
 - specifying settings for *18*

F

- filtering
 - objects *14*

G

- glossary *30*

H

- health statistics
 - monitoring *22*

M

- monitoring health and performance *22*

O

- objects
 - finding associations *14*
 - searching for *14*
- OpenStack application server
 - creating *28*
- OpenStack application servers
 - monitoring traffic *22*
- OpenStack integration
 - creating and managing connectors *28*

P

- panels
 - reordering *14*
- performance
 - for application servers *22*
 - viewing for applications *22*

R

- resource elasticity
 - configuring for tenants *18*
- resources
 - defined *30*

S

- search
 - for specific objects *14*

Index

- self-service application deployment [18](#)
- Server [28](#)
- 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 [30](#)
- terms
 - defined [30](#)

U

- user interface
 - customizing [14](#)
 - navigating [14](#)
 - searching for specific objects [14](#)