

BIG-IQ[®] Cloud: Cloud Administration

Version 4.5.0



Table of Contents

Legal Notices	7
Acknowledgments	9
Chapter 1: BIG-IQ Cloud Overview	17
Additional resources and documentation for BIG-IQ Cloud.....	18
About the BIG-IQ system user interface.....	18
Filtering for associated objects.....	18
Searching for specific objects.....	19
Customizing panel order.....	19
Chapter 2: BIG-IQ High Availability	21
About a high availability active-active cluster.....	22
Implementing an active-active high availability configuration.....	22
Chapter 3: Device Discovery	23
About device discovery and management.....	24
Discovering BIG-IP devices in your network.....	24
Viewing and exporting device inventory details.....	25
Chapter 4: License Management	27
Overview: Licensing options.....	28
About pool licenses.....	28
Automatically activating a pool license	28
Manually activating a pool license.....	29
Assigning a pool license to a BIG-IP VE.....	30
Revoking a pool license from a BIG-IP VE.....	30
About utility licenses.....	30
Automatically activating a utility license.....	31
Manually activating a utility license.....	31
Assigning a utility license to a BIG-IP device.....	32
Downloading a utility license usage report.....	33
Automatically submitting a utility license usage report to F5.....	33
Revoking a utility license from BIG-IP VE.....	33
About volume licenses.....	34
Automatically activating a volume license.....	34
Manually activating a volume license.....	34
Assigning a volume license to a BIG-IP VE.....	36
Revoking a volume license from a BIG-IP VE.....	36

Chapter 5: Integrating Amazon Web Services	37
About Amazon Web Services (AWS) integration.....	38
Network requirements for AWS integration communication	38
Creating an Amazon Identity and Access Management (IAM) user account.....	38
Creating a Virtual Private Cloud.....	39
Launching a virtual server with an Amazon Machine Image (AMI).....	40
Configuring an EC2 cloud connector.....	41
Creating a BIG-IP VE version 11.5 or later in the Amazon EC2 cloud.....	42
Creating a BIG-IP VE version 11.3 or 11.4 in the Amazon EC2 cloud.....	43
Creating a customized application template.....	45
Deploying applications.....	46
Setting up tenant access using IAM.....	47
Viewing activity for cloud resources.....	47
Chapter 6: Integrating OpenStack	49
About OpenStack integration.....	50
Network requirements for communication with OpenStack cloud services	50
OpenStack Compute edits required to use BIG-IP VE systems.....	51
Discovering devices located in the OpenStack cloud.....	51
Associating an OpenStack connector with devices.....	52
Chapter 7: Integrating with VMware Networking	55
About VMware Networking integration.....	56
Network requirements for communication with VMware cloud services	56
Creating a connection between BIG-IP Cloud and VMware.....	56
Discovering devices located in the VMware cloud.....	56
About vCloud Director integration	58
Network requirements for communication with VMware cloud services	58
Creating a connection between BIG-IP Cloud and VMware.....	58
Discovering devices located in the VMware cloud.....	59
Chapter 8: Integrating with VMware NSX 6.1.....	61
About integrating VMware NSX with a BIG-IP VE.....	62
Network requirements for communication with VMware cloud services	62
Setting up a VMware network for a BIG-IP VE.....	63
Configuring VMware NSX and BIG-IP Cloud for BIG-IP VE systems.....	63
Creating an NSX callback user.....	64
About activating a pool license.....	64
Creating a connection between BIG-IP Cloud and NSX Manager.....	65
Provisioning a BIG-IP VE on NSX version 6.1.....	66
Using the API to define an NSX runtime deployment specification.....	68
About integrating VMware NSX with a BIG-IP device	69

Setting up a VMware network for a VLAN pool bridged to a VXLAN network.....	69
Specifying VLANs on the interfaces to be provisioned.....	70
Creating a connection between BIG-IQ Cloud and NSX Manager.....	71
Creating an application template for NSX.....	72
Confirming that connector is recognized as an NSX service definition.....	72
Creating an NSX Edge Services Gateway for the BIG-IP device.....	73
Creating a load balancing service instance for VLANs bridged to a VXLAN.....	73
Specifying pools for the virtual server.....	74
Specifying virtual servers for the load balancer	75
About integrating VMware NSX with a BIG-IP device using tagged interface VLANs.....	75
Setting up a VMware network for a tagged interface VLAN pool.....	76
Discovering devices located in the VMware cloud.....	77
Specifying VLANs on the interfaces to be provisioned.....	78
Creating a connection between BIG-IQ Cloud and NSX Manager.....	78
Creating an application template for NSX.....	79
Confirming that connector is recognized as an NSX service definition.....	80
Creating an NSX Edge Services Gateway for the BIG-IP device.....	80
Creating a load balancing service instance for tagged VLANs.....	81
Specifying pools for the virtual server.....	82
Specifying virtual servers for the load balancer	82
About integrating VMware NSX with a BIG-IP device using existing VLANs.....	83
Setting up a VMware network for an existing VLAN.....	84
Creating a connection between BIG-IQ Cloud and NSX Manager.....	84
Creating an application template for NSX.....	85
Confirming that connector is recognized as an NSX service definition.....	86
Creating an NSX Edge Services Gateway for the BIG-IP device.....	86
Creating a load balancing service instance for existing VLANs.....	86
Specifying pools for the virtual server.....	87
Specifying virtual servers for the load balancer	88
Chapter 9: Local Cloud Integration.....	91
About using a local cloud source.....	92
Discovering BIG-IP devices in your network.....	92
Associating a local cloud connector with a device.....	93
Chapter 10: Cloud Tenant Management.....	95
About creating cloud tenants	96
Creating a tenant.....	96
Creating a cloud user.....	96
Associating a user with a tenant's role.....	97
Chapter 11: iApps Application Template Customization.....	99
About customizing iApp application templates.....	100

Optional load balancing methods for f5.http, f5.microsoft_sharpoint_2010, and f5.microsoft_iss catalog templates.....	100
Creating a customized application template.....	101
Chapter 12: Glossary.....	103
BIG-IQ Cloud terminology.....	104

Legal Notices

Publication Date

This document was published on June 8, 2015.

Publication Number

MAN-0501-03

Copyright

Copyright © 2015, F5 Networks, Inc. All rights reserved.

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

Trademarks

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

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

Patents

This product may be protected by one or more patents indicated at:
<http://www.f5.com/about/guidelines-policies/patents>

Export Regulation Notice

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

RF Interference Warning

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

FCC Compliance

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

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

Canadian Regulatory Compliance

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

Standards Compliance

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

Acknowledgments

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

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

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

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

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

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

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

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

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

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

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

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

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

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

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

This product includes software developed by Christos Zoulas.

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

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

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

Acknowledgments

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

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

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

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

This product includes software developed by Jared Minch.

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

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

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

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

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

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

This product includes Hypersonic SQL.

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

This product includes software developed by the Internet Software Consortium.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Acknowledgments

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Acknowledgments

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Chapter

1

BIG-IQ Cloud Overview

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

Additional resources and documentation for BIG-IQ Cloud

You can access all of the following BIG-IQ® system documentation from the AskF5™ Knowledge Base located at <http://support.f5.com/>.

Document	Description
BIG-IQ® Systems Virtual Editions Setup guides	BIG-IQ® Virtual Edition (VE) runs as a guest in a virtual environment using supported hypervisors. Each of these guides is specific to one of the hypervisor environments supported for the BIG-IQ system.
<i>BIG-IQ® System: Licensing and Initial Setup</i>	This guide provides the network administrators with basic BIG-IQ system concepts and describes the tasks required to license and set up the BIG-IQ system in their network, including how to add users and assign roles to those users.
<i>BIG-IQ® Cloud: Cloud Administration</i>	This guide contains information to help a cloud administrator manage cloud resources, devices, applications, and tenants (users).
<i>BIG-IQ® Cloud: Tenant User Guide</i>	This guide contains information to help tenants manage applications.
<i>BIG-IQ® Device: Device Management</i>	This guide provides details about how to deploy software images, licenses, and configurations to managed BIG-IP® devices.
<i>Platform Guide: BIG-IQ® 7000 Series</i>	This guide provides information about setting up and managing the BIG-IQ 7000 hardware platform.
Release notes	Release notes contain information about the current software release, including a list of associated documentation, a summary of new features, enhancements, fixes, known issues, and available workarounds.
Solutions and Tech Notes	Solutions are responses and resolutions to known issues. Tech Notes provide additional configuration instructions and how-to information.

About the BIG-IQ system user interface

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

Filtering for associated objects

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

1. To display only items associated with a specific object, hover over the object, click the gear icon, and then select **Show Only Related Items**.
The screen refreshes to display only associated objects in each panel.

2. To highlight only items associated with a specific object, hover over the object, click the gear icon, and then select **Highlight Related Items**.
The screen refreshes, highlighting only associated objects in each panel, and displaying unassociated objects in a gray font.
3. To remove a filter, click the **X** icon next to the filtered object in a panel.

Searching for specific objects

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

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

Customizing panel order

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

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

Chapter 2

BIG-IQ High Availability

- *About a high availability active-active cluster*
- *Implementing an active-active high availability configuration*

About a high availability active-active cluster

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

***Note:** If you are configuring BIG-IQ Security you must configure an active-standby cluster. Refer to the BIG-IQ Security: Administration guide for detailed instructions.*

Implementing an active-active high availability configuration

An active-active, high availability (HA) configuration ensures access to managed BIG-IP® devices in case one BIG-IQ® system fails.

1. Log in to BIG-IQ System with your administrator user name and password.
2. At the top of the screen, click **Configuration**.
3. Hover over the BIG-IQ Systems header, click the + icon when it appears, and then click **Add Device**. The Add Device screen opens.
4. In the **IP Address** field, type the BIG-IQ System's self IP address.
5. In the **User name** and **Password** fields, type the administrative user name and password for the system.
6. From the **Group** list, select **Management Group**.
7. Click the **Save** button.

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

Chapter 3

Device Discovery

- *About device discovery and management* |

About device discovery and management

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

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

After you discover devices, you can view and export inventory details about those devices for easy asset management, and you can modify device configurations as required without having to log in to each device individually.

Discovering BIG-IP devices in your network

After you license and perform the initial configuration for the BIG-IQ[®] system, you can discover BIG-IP[®] devices running version 11.5 or later. For proper communication between the managing BIG-IQ system and the devices it manages, you must configure the BIG-IQ system with a route to each F5 device you want to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

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

1. Log in to BIG-IQ Device with your administrator user name and password.
2. Hover over the Devices header, click the + icon when it appears, and then select **New Device**. The Devices panel expands to show the New Device screen.
3. In the **IP Address** field, type the IP address of the device.
The preferred address for discovering a BIG-IP device is its management IP address.
4. (This step applies only when the BIG-IQ system is hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as an admin user.
 2. Type the following command: `run /util bash`
 3. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as an admin user.
 2. Type the following command: `create net route <route name> {gw <x.x.x.x> network default}`

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

5. (This step applies only if the BIG-IQ system is not hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

6. To change the root user name, type a new name in the **Root User Name** field.
7. Type a password for the root user in the **Root Password** field.
8. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
9. For the **Auto Update Framework** setting, select the **Update Automatically** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

10. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

Viewing and exporting device inventory details

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

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Configuration**.

3. In the Devices panel, click the gear icon next to the device you want to view, and then select **Properties**. The panel expands to display device properties.
4. To export the data to a CSV file, click the **Export** button. You can modify the report as required in Microsoft Excel.

Chapter 4

License Management

- *Overview: Licensing options*
 - *About pool licenses*
 - *About utility licenses*
 - *About volume licenses*
-

Overview: Licensing options

You can centrally manage BIG-IP® virtual edition (VE) licenses for a specific set of F5 offerings (for example, BIG-IP LTM® 25M, BIG-IP LTM 200G, and BIG-IP LTM 1G). When a device is no longer needed, you can revoke the license instance and assign it to another BIG-IP VE device. This flexibility keeps operating costs fixed, and allows for a variety of provisioning options. There are three types of options:

- *Pool licenses* are purchased once, and you assign them to a number of concurrent BIG-IP VE devices, as defined by the license. These licenses do not expire.
- *Utility licenses* are purchased as you need them, and billed at a specific interval (hourly, daily, monthly, or yearly).
- *Volume licenses* are prepaid for a fixed number of concurrent devices, for a set period of time.

About pool licenses

Pool licenses are purchased for a particular product offering for a fixed number of devices, but are not permanently tied to a specific device. As resource demands change, you can use BIG-IQ® Device to revoke and reassign those licenses to other BIG-IP® VE devices as required. Pool licenses do not expire.

Automatically activating a pool license

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

Activating a license make it available for assignment to BIG-IP® devices in your network. If the BIG-IQ® system on which you are activating licensing is connected to the public internet, you can automatically activate the pool license.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Pool License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Keys** field, paste any additional license key you have.
7. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button.
The License Agreement displays.
8. To accept the License Agreement, click the **Agree** button.
9. Click the **Activate** button.

If the license does not display as activated in the Licenses panel after several minutes, click the arrow next to the license to contract the list, then click it again to expand. The screen should refresh and display the license as activated.

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

Manually activating a pool license

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

Activating a license make it available for assignment to BIG-IP® devices in your network. If the BIG-IQ® system on which you are activating licensing is not connected to the public internet, you can activate the pool license manually.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Pool License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. In the **Add-on Keys** field, paste any additional license key you have.
7. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button. The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
8. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.

Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.

9. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
10. To accept the License Agreement, click the **Agree** button.
11. Click the **Activate** button.

If the license does not display as activated in the Licenses panel after several minutes, click the arrow next to the license to contract the list, then click it again to expand. The screen should refresh and display the license as activated.

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

Manually activating offering licenses

Before you can activate the individual offering licenses, you must first activate the license itself.

Activating the offering licenses makes them available for assignment.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Licenses panel, click the arrow next to the license you previously activated. The list expands to display the license offerings associated with this license.
4. Hover over an offering license and click the gear icon when it appears.
5. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.

Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.

6. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
7. Copy the license key.
8. On BIG-IQ Device, into the **License Text** field, paste the license key.
9. Click the **Apply** button at the top of the panel.

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

Assigning a pool license to a BIG-IP VE

Before you can assign a pool license to a BIG-IP® VE device, you must activate the license on the BIG-IQ® system and discover the BIG-IP VE device to which you want to assign the license.

Pool licenses provide you with the flexibility to easily manage resources and operating costs. Use this procedure if you have activated a pool license, but have not yet assigned it to a BIG-IP VE.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device you want to license by clicking the arrow next to it.
The panel expands to display the list of devices contained in this group.
4. Click the gear icon next to the device you want to license, and then click **License Device**.
5. In the **Name** field, type a name for this license.
6. From the **Licensing** list, select **Use a Pool License**.
7. From the **Pool License** list, select the pool license you want to assign to this device.
8. Click the **Deploy** button.
9. To confirm that the license was successfully deployed, click the gear icon next to the license you deployed, click **Properties**, and then click **Assignments**.
The device you licensed displays with the license status and the last contact from the BIG-IQ system.

Revoking a pool license from a BIG-IP VE

If traffic decreases to the applications on some of your managed BIG-IP® devices, you can use BIG-IQ® Device to revoke those licenses and assign them to other resources as needed.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device for which you want to revoke a license by clicking the arrow next to it.
4. Click the gear icon next to the device for which you want to revoke a license, and then click **License Device**.
5. From the **Licensing** list, select **Revoke a License**.
6. Click the **Deploy** button.

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

About utility licenses

You are charged for utility licenses only for the duration that the license is activated. You can activate any number of licenses as you need them, specifying the interval (an hour, a day, a month, or a year) at which you want to be billed for each. BIG-IQ® Device tracks license usage in each billing period, and sends that data directly to F5. When a resource is no longer required, you revoke its license and are no longer charged

for that instance until you reassign it to another BIG-IP VE device. Utility licenses can be particularly useful when traffic to certain applications increases for a short period of time, for example, during fiscal year end.

Automatically activating a utility license

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

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

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Utility License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button. The License Agreement displays.
7. To accept the License Agreement, click the **Agree** button.
8. Click the **Activate** button.

If the license does not display as activated in the Licenses panel after several minutes, click the arrow next to the license to contract the list, then click it again to expand. The screen should refresh and display the license as activated.

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

Manually activating a utility license

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

Activating a utility license is the first step to making it available for assignment to BIG-IP® devices in your network. If the BIG-IQ® system on which you are activating licensing is not connected to the public internet, you can activate the utility license manually.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Utility License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button. The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
7. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.

Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.

8. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
9. Select the check box next to the Accept User Legal Agreement to agree to the license terms, and then click the **Next** button.

The license key displays

10. Copy the license key.
11. On BIG-IQ Device, into the **License Text** field, paste the license key.
12. Click the **Apply** button at the top of the panel.

You must now activate each individual utility license offering.

Manually activating offering licenses

Before you can activate the individual offering licenses, you must first activate the license itself.

Activating the offering licenses makes them available for assignment.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Licenses panel, click the arrow next to the license you previously activated.
The list expands to display the license offerings associated with this license.
4. Hover over an offering license and click the gear icon when it appears.
5. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.
Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.
6. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
7. Copy the license key.
8. On BIG-IQ Device, into the **License Text** field, paste the license key.
9. Click the **Apply** button at the top of the panel.

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

Assigning a utility license to a BIG-IP device

Before you can assign a utility pool to a BIG-IP® VE device, you must activate the utility license on the BIG-IQ® system and discover the BIG-IP VE device to which you want to assign a pool license.

Using a utility license for a BIG-IP VE device provides you with the flexibility to easily manage resources and operating costs by choosing a specific billing term for licenses.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device you want to license by clicking the arrow next to it.
The panel expands to display the list of devices contained in this group.
4. Click the gear icon next to the device you want to license, and then click **License Device**.
5. In the **Name** field, type a name for this license.
6. From the **Licensing** list, select **Use a Utility License**.
7. From the **Utility License** list, select the license you want to assign to this device.
8. From the **Offering License** list, select the specific product offering you want to assign to this device.
9. From the **Unit Of Measure** list, select the interval at which you want to be billed for this license.
10. Click the **Deploy** button.

Downloading a utility license usage report

You must assign a utility license to a device before you can create a utility usage report for that license.

You can use this report to augment your internal licensing management and budget planning. You also have the option to submit this report manually to F5 for billing purposes.

Note: If you would like to manually submit this report to F5 for billing purposes instead of automatically, contact F5 Support.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Licenses panel, click the gear icon for the utility license for which you want to download a usage report, and then click **Create Usage Report**.
4. For the **Period** setting, in the **From** and **To** fields, type the date range for the report. Alternatively, click the calendars and navigate to the dates.
5. Select a format option for the report.
6. Click the **Download** button and select an option to open the file, or save the file.

Automatically submitting a utility license usage report to F5

You must assign a utility license to a device before you can submit and save a usage report.

You provide this report to F5 Networks for billing purposes, as per the terms and conditions of your contract.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Licenses panel, click the gear icon for the utility license that you want to submit for billing, and then click **Create Usage Report**.
4. For the usage submission method, select **Automatically submit report to F5**.
5. Click the **Submit** button.
BIG-IQ Device sends a report directly to F5, and saves a copy on BIG-IQ Device.

Revoking a utility license from BIG-IP VE

If traffic decreases to the applications on some of your managed BIG-IP® devices, you can use BIG-IQ® Device to revoke those licenses and assign them to other resources as needed.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device for which you want to revoke a license by clicking the arrow next to it.
4. Click the gear icon next to the device for which you want to revoke a license, and then click **License Device**.
5. From the **Licensing** list, select **Revoke a License**.
6. Click the **Deploy** button.

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

About volume licenses

With volume licenses, you can flexibly manage BIG-IP® VE devices by purchasing a number of prepaid, concurrent licenses. If your needs change throughout the year, you have the option of purchasing more prepaid licenses in increments of 50. BIG-IQ® Device helps you track and distribute these various licenses as required by the applications your customers are using, and notifies you when you reach your prepaid limit. When you revoke a license, you can then assign it to another BIG-IP VE device.

Automatically activating a volume license

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

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

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Volume License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button. The License Agreement displays.
7. To accept the License Agreement, click the **Agree** button.
8. Click the **Activate** button.

If the license does not display as activated in the Licenses panel after several minutes, click the arrow next to the license to contract the list, then click it again to expand. The screen should refresh and display the license as activated.

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

Manually activating a volume license

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

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

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. Hover over the Licenses header, click the + icon when it appears, and then click **Add New Volume License**.
4. In the **License Name** field, type the name you want to use to identify this license.
5. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
6. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button.

The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.

7. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.

Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.

8. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
9. Select the check box next to the Accept User Legal Agreement to agree to the license terms, and then click the **Next** button.

The license key displays

10. Copy the license key.
11. On BIG-IQ Device, into the **License Text** field, paste the license key.
12. Click the **Add** button.
The unactivated volume license displays in the Licenses panel.
13. Click the arrow next to the volume license you created to expand the list of licenses.
14. Click the volume license you want to activate.
15. Copy the license key.
16. On BIG-IQ Device, into the **License Text** field, paste the license key.
17. To accept the License Agreement, click the **Agree** button.

18. Click the **Activate** button.

If the license does not display as activated in the Licenses panel after several minutes, click the arrow next to the license to contract the list, then click it again to expand. The screen should refresh and display the license as activated.

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

Manually activating offering licenses

Before you can activate the individual offering licenses, you must first activate the license itself.

Activating the offering licenses makes them available for assignment.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Licenses panel, click the arrow next to the license you previously activated.
The list expands to display the license offerings associated with this license.
4. Hover over an offering license and click the gear icon when it appears.
5. Copy the text displayed in the **Device Dossier** field, and click the **Access F5 manual activation web portal** link.
Alternatively, you can navigate to the F5 license activation portal at <https://activate.f5.com/license/>.
6. Paste the dossier into the **Enter your dossier** field, and then click the **Next** button.
7. Copy the license key.
8. On BIG-IQ Device, into the **License Text** field, paste the license key.
9. Click the **Apply** button at the top of the panel.

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

Assigning a volume license to a BIG-IP VE

Before you can assign a volume license to a BIG-IP[®] VE device, you must activate the volume license on the BIG-IQ[®] system and discover the BIG-IP VE device to which you want to assign a volume license.

Using a volume license for a BIG-IP VE device provides you with the flexibility to easily manage resources and operating costs by choosing only those features you want to use on the managed BIG-IP VE device.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device you want to license by clicking the arrow next to it.
The panel expands to display the list of devices contained in this group.
4. Click the gear icon next to the device you want to license, and then click **License Device**.
5. In the **Name** field, type a name for this license.
6. From the **Licensing** list, select **Use a Volume License**.
7. Click the **Deploy** button.

Revoking a volume license from a BIG-IP VE

If traffic decreases to the applications on some of your managed BIG-IP[®] devices, you can use BIG-IQ[®] Device to revoke those licenses and assign them to other resources as needed.

1. Log in to BIG-IQ Device with your administrator user name and password.
2. At the top of the screen, click **Provisioning**.
3. On the Devices panel, expand the device group that contains the device for which you want to revoke a license by clicking the arrow next to it.
4. Click the gear icon next to the device for which you want to revoke a license, and then click **License Device**.
5. From the **Licensing** list, select **Revoke a License**.
6. Click the **Deploy** button.

You can now assign this license to another BIG-IP[®] device.

Chapter

5

Integrating Amazon Web Services

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

About Amazon Web Services (AWS) integration

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

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

Network requirements for AWS integration communication

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

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

Creating an Amazon Identity and Access Management (IAM) user account

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

Important: *This task is optional; you can create a virtual machine without creating an IAM user account to control access, but it is best practice to use an IAM account. F5 recommends that you do not use the AWS root account and access keys. Instead, use IAM to create identities you can more easily manage and revoke in the case of a security breach.*

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

For this task, you must create a group and two IAM user accounts. For the most current instructions for performing these steps, refer to the IAM documentation web site, <http://aws.amazon.com/documentation/iam/>.

1. From <https://console.aws.amazon.com/iam>, create a new group named `aws-full-access` and assign the group access rights by attaching the `AdministratorAccess` policy.
2. Create an AWS Admin user, and add that user to the `aws-full-access` group.
3. Create a BIG-IQ[®] Connector user and add that user to the `AmazonEC2FullAccess` group.

Important: The user requires only EC2 Full Access privileges (not AWS Full Access). IAM policy `arn:aws:iam::aws:policy/AmazonEC2FullAccess` represents the EC2 Full Access privileges set.

Important: For this user, you must download or copy an access key that you use to connect BIG-IQ Cloud to your AWS account

4. From the AWS dashboard, set up an account alias.
Note the IAM user login link. For example,
`https://my-account-alias.signin.aws.amazon.com/console`
5. Log out of the AWS dashboard as the root user.
6. Navigate back to the user login link, and sign in as the AWS-Admin user.

You can now create a new Virtual Private Cloud (VPC).

Creating a Virtual Private Cloud

You need an Amazon Virtual Private Cloud (VPC) to deploy the BIG-IQ® Cloud system, because AWS provides only multiple network interface card (NIC) support for instances that reside within a VPC.

You create a virtual network topology according to your networking needs. The standard network topology used for BIG-IQ Cloud integration includes three subnets. These subnets provide virtual private address spaces used to interconnect your machines and applications. You can use elastic self IP addresses for public internet accessibility.

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

1. Navigate to `https://console.aws.amazon.com/vpc` and select the AWS Region in which you want to manage resources.
For example, Oregon.
2. Click **Start VPC Wizard**.
The Select a VPC Configuration screen opens.
3. Select the VPC Wizard's **VPC with a Single Public Subnet** option, and then click **Select**.
The VPC with a Single Public Subnet screen opens.
4. Set the IP CIDR Block to `10.0.0.0/16`.
5. Set the public subnet to `10.0.0.0/24`.
This is the management network.
6. Select an availability zone.
For example, **us-west-2c**. It is crucial that you use this availability zone throughout the configuration process. Objects configured in one zone are not visible within other zones, so they cannot function together. This availability zone is required when you create a BIG-IQ Cloud connection.
7. Click **Create VPC** to create the new VPC.

Important: Now you need to create another public subnet in the same availability zone as `10.0.1.0/24` and make sure it uses a route table with internet gateway connectivity

8. On the VPC Dashboard, click **Subnets**.

9. Click Create Subnet.

The Create Subnet screen opens.

- 10.** a) For the **Name tag**, type a name for this subnet.
b) For the **VPC**, select the just created VPC.
c) For the **Availability Zone**, select the zone you specified for the public subnet.
d) For the **CIDR block** type `10.0.1.0/24`.
e) Click **Yes, Create**.

AWS creates the new subnet, now you need to make sure it is publically accessible.

11. Select the new subnet, select the Route Table tab on the lower half of the screen, and then click **Edit**.

12. From the Change to list, select the route table that has internet gateway connectivity. and click **Save**.

13. Create a security group named, `allow-all-traffic`, and associate it with the VPC you created.

You must use this exact name.

14. Set the **Inbound Rules ALL Traffic Source** to `0.0.0.0/0`.

15. Set the **Outbound Rules ALL Traffic Destination** to `0.0.0.0/0`.

16. Create a Route Table for the external data network to reach the Internet.

17. Add a route to Destination **0.0.0.0/0** through Target `igw-<xxxx>`.

`<xxxx>` is the Internet Gateway that the VPC Wizard created automatically.

18. Allocate two Elastic IP Addresses.

You now should create a BIG-IQ Cloud connector to associate with this VCP.

Launching a virtual server with an Amazon Machine Image (AMI)

Before you can complete this task, you need to know the name of your key pair and the Availability Zone from which it was created.

You launch an EC2 Amazon Machine Image (AMI) so that you can deploy the virtual machine.

Important: *At publication, this task illustrates the Amazon web interface. However, F5 recommends that you refer to Amazon user documentation for the latest documentation.*

1. Log in to your account on Amazon Web Services (AWS) marketplace.
2. In the Search AWS Marketplace bar, type `F5 BIG-IQ` and then click **GO**.
The F5 BIG-IQ Virtual Edition for AWS option is displayed.
3. Click **F5 BIG-IQ Virtual Edition for AWS** and then click **CONTINUE**.

Tip: *You might want to take a moment here to browse the pricing details to confirm that the region in which you created your security key pair provides the resources you require. If you determine that the resources you need are provided in a region other than the one in which you created your key pair, create a new key pair in the correct region before proceeding.*

The Launch on EC2 page is displayed.

4. Click the **Launch with EC2 Console** tab.

Launching Options for your EC2 AMI are displayed.

5. Select the software version appropriate for your installation, and then click the **Launch with EC2** button that corresponds to the Region that provides the resources you plan to use.

Important: The first time you perform this task, you need to accept the terms of the end user license agreement before you can proceed, so the **Launch with EC2** button reads **Accept Terms and Launch with EC2**.

Important: There are a number of factors that determine which region will best suit your requirements. Refer to Amazon user documentation for additional detail. Bear in mind that the region you choose must match the region in which you created your security key pair.

The Request Instances Wizard opens.

6. Select an **Instance Type** appropriate for your use.
7. From the **Launch Instances** list, select **EC2-VPC**.
8. From the **Subnet** list, select the **10.0.0.0/24** subnet and click **CONTINUE**.
The Advanced Instance Options view of the wizard opens.
9. From the **Number of Network Interfaces** list, select **2**.
10. Click the horizontal **eth1** tab to set values for the second network interface adapter, and then from the **Subnet** list, select the **10.0.1.0/24** subnet and click **CONTINUE**.
The Storage Device Configuration view of the wizard opens.
11. In the **Value** field, type in an intuitive name that identifies this AMI and click **CONTINUE** (for example, `BIG-IQ VE <version>`).
The Create Key Pair view of the wizard opens.
12. From **Your existing Key Pairs**, select the key pair you created for this AMI and click **CONTINUE**.
The Configure Firewall view of the wizard opens.
13. Under Choose one or more of your existing Security Groups, select the **allow-all-traffic** security group, and then click **CONTINUE**.
The Review view of the wizard opens.
14. Confirm that all settings are correct, and then click **Launch**.
The Launch Instance Wizard displays a message to let you know your instance is launching.
15. Click **Close**.

Your new instance appears in the list of instances when it is fully launched.

Configuring an EC2 cloud connector

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

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
4. From the **Cloud Provider** list, select **Amazon EC2**.

5. In the **Region Endpoint** field, type the entry point URL.
For example, `ec2.us-east-1.amazonaws.com` is the region end point for the Amazon EC2 US East (Northern Virginia) Region. Refer to the AWS documentation for a list of all regional end points at http://docs.aws.amazon.com/general/latest/gr/rande.html#ec2_region.
6. In the **Key ID** and **Secret Key** fields, type the credentials of the BIG-IQ-Connector IAM user.
For security purposes, it is important to specify a user that has Amazon EC2 Full Control Access.
7. In the **Availability Zone** field, type the location of the region in which the instances are located.
For example, type `us-west-2c` for the availability zone for Oregon state.
8. In the **Virtual Private Cloud** field, you may type the identification for the EC2 Virtual Private Cloud (VCP) network topology inside the Availability Zone.
This step is optional. If you do not specify the identification for a VCP, BIG-IQ Cloud uses the first one it discovers in the Availability Zone.
9. Click the arrow next to **Device & Server Provisioning** to display associated options.
10. To prompt BIG-IQ Cloud to automatically provision additional BIG-IP VE devices when more resources are needed for application traffic, for the **Device Elasticity** setting, select **Enable**.
11. From the **Device License** list, select a rate at which you want Amazon to direct-bill for additional devices, or select a license pool from which to grant a license.
You must activate a license pool before you can select it.
12. To automatically prompt BIG-IQ Cloud to provision additional servers when more resources are needed to manage an influx in application traffic, for the **Server Elasticity** setting, select **Enable**.
13. Review the network settings populated when you selected a connector, verifying that the proper CIDR blocks display for management, external, and internal.
14. Click the **Save** button.
15. If the system discovered devices, you must expand the device's properties panel, and provide the device's credentials to finalize the discovery process.
16. Review the network settings populated when you selected a connector, verifying that the proper CIDR blocks display for management, external, and internal.

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

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

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

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

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

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

Group Name	Group Description	Rule Name	Source	Port
		Outbound HTTPS	443 0.0.0.0/0	443 (HTTPS)
		Inbound HTTPS	0.0.0.0/0	443 (HTTPS)

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

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

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

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

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

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

You can perform this task only after you have licensed and installed the BIG-IQ® system and at least one BIG-IP® device running version 11.3 or 11.4.

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

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

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

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

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

- Log in to the BIG-IQ system terminal as the root user.
- Optionally, establish SSH trust between the BIG-IQ system and the managed BIG-IP device.


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

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


```
cd /usr/lib/dco/packages/upd-adc
```
- Run the installation script.
 - For devices installed in an Amazon EC2 environment: `./update_bigip.sh -a admin -p <password> -i /<path_to_PEM_file> <BIG-IP Management IP Address>`
 - For devices installed in any other environment: `./update_bigip.sh -a admin -p <password> <BIG-IP Management IP Address>`

Where `<password>` is the administrator password for the BIG-IP device.
- If you established trust in step 2, revoke SSH trust between the BIG-IQ system and the managed BIG-IP device.


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

This step is not required if you did not establish trust in step 2.
- Log in to BIG-IQ® Cloud with your administrator user name and password.
- In the Device panel, click the gear icon next to the legacy device with a yellow triangle next to it and displaying the message, Discovery is incomplete.
- In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
- For the **Auto Update Framework** setting, select the **Update Automatically** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic

recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

10. Click the **Save** button.

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

Creating a customized application template

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

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

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

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

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

6. To allow cloud tenants to specify certificates with SSL encryption when self-deploying applications, select options from the **SSL Cert** and **SSL Key** lists.
BIG-IQ® Cloud uses these options to provide the appropriate certificate and key when the tenant self-deploys this application to a BIG-IP® device. These options are not available for all application templates.
7. Finish making modifications by specifying the Application Properties and Customize Application Template variables.
To allow a tenant to modify a particular setting, select the **Tenant Editable** check box for that setting. For further details about template variables and settings, refer to the *BIG-IP® iApps® Developer's Guide*.
8. If you selected **f5.http**, **f5.microsoft_sharepoint_2010**, or **f5.microsoft_iss** and you want to specify a load balancing option other than the default, Least Connection Member, perform the following steps:
 - a) Click the arrow next to Advanced Properties.

- b) In the **Which load balancing method do you want to use?** field, type the value for the option you want to use.
9. Click the **Save** button.
You can now send the cloud IP addresses to the tenant and use this IP address range in configuring server tiers and pool members, within certain application services. The tenant can self-deploy the application from the catalog.

The customized application displays as an entry in the catalog.

Deploying applications

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

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

1. Log in to the BIG-IQ Cloud with your tenant user name and password.
2. Hover over the Applications header, and click the + icon when it appears.
3. Hover over the Applications header, and click the + icon when it appears.
4. In the **Name** field, type a name for this new application.
5. From the **Application Type** list, select an application.
6. From the **Cloud Connector** list, select the cloud connector associated with where you want to deploy your application.
A cloud connector is a resource that identifies the local or virtual environment in which a tenant deploys applications and, when necessary, adds parameters required by third-party cloud providers.
7. For the **Provision Virtual Server IP** setting, select **Enable** and specify the FQDNs for the virtual servers to prompt BIG-IQ Cloud to automatically provision additional resources when traffic to your application increases.
8. To define a new SSL certificate and private key for this application, for the **SSL Certificate Options**, paste the PEM (CRT or CER) text representation of the certificate and private key.
The SSL certificate and private key must be unbundled Base64 encoded ASCII text with PEM header and footer.
This option is not available for all applications.
9. Alternatively, select the **Use Existing** option to use a SSL certificate and private key already stored on the device.
10. You can further customize this application by specifying an IP address for the virtual server and adding pool hosts.
If your cloud service provider assigned IP addresses for the **Servers**, **Pool Hosts**, and **Pool Members** for this application, the addresses display. If these addresses were specified as not editable, you cannot change them.
11. When you are finished, click the **Deploy** button located at the top of the New Application panel.

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

Setting up tenant access using IAM

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

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

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

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

Viewing activity for cloud resources

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

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

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

Chapter

6

Integrating OpenStack

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

About OpenStack integration

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

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

Network requirements for communication with OpenStack cloud services

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

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

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

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

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

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

OpenStack Compute edits required to use BIG-IP VE systems

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

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

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

2. Restart the OpenStack Compute node services.

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

Discovering devices located in the OpenStack cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.5 or later. For proper communication between the managing BIG-IQ system and the devices it manages, you must configure the BIG-IQ system with a route to each F5 device you want to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

You must know the IP address that the BIG-IQ device will use to access the BIG-IP device.

1. Hover over the Devices header, click the + icon when it appears, and then select **New Device**. The Devices panel expands to show the New Device screen.
2. In the **IP Address** field, type the device's IP address.

The preferred address for discovering a BIG-IP device is its management IP address.
3. (This step applies only if the BIG-IQ system is not hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

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

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

6. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

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

Associating an OpenStack connector with devices

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

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

1. Log in to BIG-IQ[®] Cloud with your administrator user name and password.
2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
 You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
4. From the **Cloud Provider** list, select **OpenStack**.
5. In the **OpenStack Controller Node URI** field, type the URI for the OpenStack controller node.
6. In the **OpenStack User Name** field, type the user name for the OpenStack administrator.

For example, `https://<IP address>:<Port>` or `http://<IP address>:<Port>`.

Note that default port for `http` is 5000.

7. In the **OpenStack Tenant Name** and **OpenStack Password** fields, type the tenant (also known as, project) name and password.
8. Click the **Save** button.

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

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

Chapter 7

Integrating with VMware Networking

- *About VMware Networking integration*
 - *About vCloud Director integration*
-

About VMware Networking integration

There are two VMware products that you can integrate with the BIG-IQ[®] Networking connector.

- For VMware vShield version 5.1 and 5.5 (also known as VCNS version 5.5), and VMware NSX 6.0, the BIG-IQ software integration provides you with the tools to provide tenants access to self-deployable iApps[®].
- For vCloud Director versions 1.5 and 5.1, BIG-IQ[®] Cloud integration makes it possible for you to use the VCD interface with your cloud applications to manage the F5 cloud applications.

Network requirements for communication with VMware cloud services

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

Creating a connection between BIG-IQ Cloud and VMware

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

1. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **VMware NSX 6.1**.
4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list. BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. In the **VMware Networking Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device's internal VLAN.
7. In the **VMware Networking User Name** and **VMware Networking Password** fields, type the credentials for the VMware administrator.
8. From the **BIG-IQ User Name** list, select the BIG-IQ user the VMware administrator should contact and, in the **BIG-IQ Password** field, type the password for that user.
9. Click the **Save** button.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ[®] system, you can discover BIG-IP[®] devices running version 11.5 or later. For proper communication between the managing BIG-IQ system and the devices it manages, you must configure the BIG-IQ system with a route to each F5 device you want

to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

You must know the IP address that the BIG-IQ device will use to access the BIG-IP device.

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Devices header, click the + icon when it appears, and then select **New Device**. The Devices panel expands to show the New Device screen.
3. In the **IP Address** field, type the device's IP address.
The preferred address for discovering a BIG-IP device is its management IP address.
4. (This step applies only if the BIG-IQ system is not hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

5. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
6. For the **Auto Update Framework** setting, select the **Update Automatically** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

***Important:** When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).*

7. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

If you want to use the BIG-IP device just discovered to host NSX virtual servers, you should now associate it with a VMware cloud connector.

About vCloud Director integration

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

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

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

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

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

Task summary

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

Network requirements for communication with VMware cloud services

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

Creating a connection between BIG-IQ Cloud and VMware

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

1. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **VMware NSX 6.1**.
4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list. BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. In the **VMware Networking Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device's internal VLAN.

7. In the **VMware Networking User Name** and **VMware Networking Password** fields, type the credentials for the VMware administrator.
8. From the **BIG-IQ User Name** list, select the BIG-IQ user the VMware administrator should contact and, in the **BIG-IQ Password** field, type the password for that user.
9. Click the **Save** button.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.5 or later. For proper communication between the managing BIG-IQ system and the devices it manages, you must configure the BIG-IQ system with a route to each F5 device you want to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

You must know the IP address that the BIG-IQ device will use to access the BIG-IP device.

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Devices header, click the + icon when it appears, and then select **New Device**. The Devices panel expands to show the New Device screen.
3. In the **IP Address** field, type the device's IP address.

The preferred address for discovering a BIG-IP device is its management IP address.
4. (This step applies only if the BIG-IQ system is not hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

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

For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

7. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

If you want to use the BIG-IP device just discovered to host NSX virtual servers, you should now associate it with a VMware cloud connector.

Chapter

8

Integrating with VMware NSX 6.1

- *About integrating VMware NSX with a BIG-IP VE*
- *About integrating VMware NSX with a BIG-IP device*
- *About integrating VMware NSX with a BIG-IP device using tagged interface VLANs*
- *About integrating VMware NSX with a BIG-IP device using existing VLANs*

About integrating VMware NSX with a BIG-IP VE

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

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

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

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

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

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

Setting up a VMware network for a BIG-IP VE

Configuring VMware NSX and BIG-IQ Cloud for BIG-IP VE systems

Creating an NSX callback user

Creating a connection between BIG-IQ Cloud and NSX Manager

Provisioning a BIG-IP VE on NSX version 6.1

Network requirements for communication with VMware cloud services

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

Setting up a VMware network for a BIG-IP VE

Before you can begin configuring the BIG-IQ® device integration for a BIG-IP® Virtual Edition (VE), you must perform the following setup tasks.

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

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

Configuring VMware NSX and BIG-IQ Cloud for BIG-IP VE systems

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

1. Log in to vCenter using the vSphere Web Client.
2. In the VMware vSphere Web Client, create additional virtual network connectivity options.

One network must be a management network; typically the BIG-IQ system uses it for provisioning BIG-IQ systems and for discovering BIG-IP devices. You can use an internal network for provisioning and discovering BIG-IP devices as long as that network can be reached by the BIG-IQ device. The other required network is data plane; the BIG-IP device uses it to pass traffic. You need one management network and then you can create up to three data plane networks. You can choose whether each network is a **Logical Switch**, a **Standard Portgroup**, or a **Distributed Portgroup**.

 - a) Define a management network for the BIG-IP device. Use a typical IP address range to refer to this network: 192.168.11.0/24.
 - b) Define a data network. Use a typical IP address range to refer to this network: 10.22.0.0/16.
 - c) Optionally, define another data network. Use a typical IP address range to refer to this network: 10.33.0.0/16.
 - d) Optionally, define another data plane network for the BIG-IP device. Use a typical IP address range to refer to this network: 10.44.0.0/16.
3. In the VMware vSphere Web Client, create two to four IP Pools, one for each network. As you create each pool, you are prompted for a name. Make a note of the names you choose so that when you need to associate each pool to a network interface, you will know which is which.
4. In the VMware vSphere Web Client, set up a web server on the just-created management network.

The NSX Manager uses the URL of this web server to access the installation file (OVA) for the BIG-IP VE you intend to provision.

5. Decompress the OVA file and copy the contents (which include an OVF file) to an accessible location on the just-created web server.

The NSX Manager uses the OVF file to create the BIG-IP VE.

The next tasks to perform are:

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

Creating an NSX callback user

Create an NSX callback user to provide that individual with access to specific NSX resources. You will need the name of this user when you add a VMware NSX 6.1 connector.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the User header, and click the + icon when it appears.
The panel expands to display property fields for the new user.
3. In the **Full Name** field, type a name to identify this user.
The full name can contain a combination of symbols, letters, numbers and spaces.
4. In the **Password** and **Confirm Password** fields, type the password for the new user.
5. Click the **Add** button.

You can now specify this user name as the NSX Callback User Name.

About activating a pool license

When you integrate with VMware NSX to create BIG-IP® VE virtual machines, you can activate a pool license so that BIG-IQ® software can use a license from that pool to license the BIG-IP VE systems that it creates.

You can choose not to use a pool license and skip to discovering devices. If you make this choice, the BIG-IQ device still creates BIG-IP VE systems, but you need to license them before they can be used.

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

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

Automatically activating a pool license

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

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.
3. In the **Add-on Keys** field, paste any additional license key you have.
4. For the **Activation Method** setting, select **Automatic**, and click the **Activate** button. The License Agreement displays.
5. To accept the License Agreement, click the **Agree** button.

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

Manually activating a pool license

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

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

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

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

Creating a connection between BIG-IQ Cloud and NSX Manager

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.

Important: You will need to recall the name you assign to this connector so that you can select it when you are configuring the VMware user interface. The name you specify is used as the service definition name in the VMware user interface.

4. From the **Cloud Provider** list, select **VMware NSX 6.1**.
The screen displays additional settings specific to VMware NSX.
5. From the **Devices** list, select the device you want to associate with this connector.
6. In the **VMware NSX Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device.
7. In the **VMware NSX User Name** and **VMware NSX Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the NSX Manager REST API.
8. In the **VMware vCenter Server Address** field, type the IP address of the vCenter server.
9. In the **VMware vCenter Server User Name** and **VMware vCenter Server Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the vCenter SOAP API.
10. From the **BIG-IQ Callback User Name** list, select the user name that NSX Manager uses to authenticate to the BIG-IQ REST API.
11. In the **BIG-IQ Callback Password** field, type the password that NSX Manager uses to authenticate to the BIG-IQ REST API.
12. From the **BIG-IQ Callback Address** list, select the IP address that this NSX Manager uses to access each BIG-IQ device in the HA cluster.
By default, the management IP address is used, but you can specify a self IP address if you choose.
13. Click the **Save** button.

As part of the connection creation process, the BIG-IQ system does the following:

- Creates a new default tenant for the new connector.
- Verifies connectivity to the NSX Manager and vCenter APIs, and registers the BIG-IQ system as an NSX Partner Service provider.
- Creates a callback user role that enables NSX to access the BIG-IQ software resources necessary for interaction with the BIG-IQ REST API.

Provisioning a BIG-IP VE on NSX version 6.1

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

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

1. In the vSphere web client user interface, create a new NSX Edge in an undeployed state.
If you specify a tenant ID when you create the Edge, BIG-IQ software will create a tenant with that ID when it creates the BIG-IP VE.
Make sure that the NSX Edge you create identifies the Cluster/Resource Pool and the Datastore, but does not identify any interfaces. Otherwise, follow your standard practice for NSX Edge creation.
2. For the just created NSX Edge, navigate to the Manage tab, and then select the Load Balancer tab. Then click the **Edit** button.
The Edit Load balancer global configuration window displays.
3. Select both **Enable Load Balancer** and **Enable Service Insertion**.

4. For the **Service Definition**, select the name of the connector you created for NSX in a previous step.
5. For the **Service Configuration**, select **F5 ADC - Make a BIG-IP VE**.
6. Expand Typed Service Configuration Parameters and then determine which of these optional settings you want to specify.
 - a) In the **Value** field next to **F5-BIG-IP-VE** key, type `yes`.
 - b) If you want to specify a fully qualified host name of the node template to describe the kind of BIG-IP VE this template creates, in the **Value** field next to **F5-BIG-IP-VE-FQ-HOST-NAME** key, type that name in the value box next to **Name of BIG-IP node template?**.

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

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

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

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

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

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

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

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

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

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

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

- The first required vNIC (vNIC0) provides the DHCP-enabled control plane network on which the BIG-IP VE boots. Choose the name specified previously that corresponds to the IP pool `192.168.11.0/24`.
- The next required vNIC (vNIC1) you specify provides the external data network on which the BIG-IP device creates virtual servers. Choose the name specified previously that corresponds to the IP pool `10.22.0.0/16`.
- The first optional vNIC you specify provides the internal data network on which load-balanced pool members are located. Choose the name specified previously that corresponds to the IP pool `10.33.0.0/16`.
- The next optional vNIC you specify provides the data plane network on which the BIG-IQ device discovers and manages BIG-IP devices. Choose the name specified previously that corresponds to the IP pool `192.44.0.0/16`.

8. Click **OK** to close the Edit Load balancer global configuration dialog box.

VMware NSX configures the Edge Gateway based on the settings you specified.

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

Using the API to define an NSX runtime deployment specification

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

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

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

```
{
  "state": "TEMPLATE",
  "properties": [
    {
      "id": "BIG-IP",
      "provider": "true"
    },
    {
      "id": "NodeTemplateName",
      "value": "BIGIP-11.5.0.0.0.221.LTM_1SLOT-scsi.ovf"
    }
  ]
}
```

```

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

```

7. Click **Save**.

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

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

About integrating VMware NSX with a BIG-IP device

The integration between BIG-IQ® Cloud and VMware® NSX makes it possible for you to use existing physical BIG-IP® devices to host NSX virtual servers. Using these servers, you can manage and deploy iApps® on existing VMware NSX environments.

There are three connectivity options for this integration.

- One connectivity option uses a pool of VLANs bridged to VXLAN networks.
- One connectivity option uses VLAN trunks that the BIG-IQ device provisions for you. These VLANs must use a tagged interface.
- One connectivity option uses VLANs that are already configured on the BIG-IP device. These VLANs can be either tagged or untagged depending on your network topology.

Most of the API calls for these options are the same; only the network setup varies. When there are specific differences in the API call, it is detailed in the task.

Important: *NSX version 6.1.3 or later is required for this particular connectivity option.*

Task summary

Setting up a VMware network for a VLAN pool bridged to a VXLAN network

Specifying VLANs on the interfaces to be provisioned

Creating a connection between BIG-IQ Cloud and NSX Manager

Creating an application template for NSX

Confirming that connector is recognized as an NSX service definition

Creating an NSX Edge Services Gateway for the BIG-IP device

Creating a load balancing service instance for VLANs bridged to a VXLAN

Specifying pools for the virtual server

Specifying virtual servers for the load balancer

Setting up a VMware network for a VLAN pool bridged to a VXLAN network

Before you can begin configuring the BIG-IQ® device integration for a VLAN pool bridged to a VXLAN network, you must perform the following setup tasks.

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

1. Discover the device on which your VLANs will reside so you can configure the VLANs from the BIG-IQ device.
2. Connect the BIG-IP® device and the ESXi host to a physical switch.
The BIG-IP device and the ESXi host must connect to a port set up to support a trunk interface.
3. Create a vCenter distributed port group for each VLAN in the trunk. Each port group must have a tagged VLAN ID.
This configuration is referred to as virtual switch tagging (VST). It ensures that outbound traffic to the BIG-IP device is tagged.
4. Confirm that the vCenter distributed port group uplinks to the physical switch set up in the previous step.
5. Pre-deploy a Logical Distributed Router (LDR) in the VMware NSX environment.
 - a) Confirm that the VXLAN-transport VLAN uses the same physical switch as the BIG-IP device and the ESXi host.
 - b) Confirm that the logical switches exist on the Distributed vSwitch (DVS).
 - c) Deploy the LDR to the ESXi host on which the physical VLAN used for bridging to VXLAN is connected to the ESXi host.
 - d) Configure the LDR with a unique interface to the management network with a valid IP Address.
The management network LDR port group cannot be part of VXLAN-VLAN mapping.
 - e) Configure the LDR with an interface to any port group on the DVS.
You do not need to specify an IP address. This interface associates the LDR to the correct vSwitch.

Important: There must be at least one LDR per DVS. If VXLAN to VLAN mapping goes above 255 for one DVS, more than one LDR per DVS is required.

6. Define the IP pool for each port group in NSX.
These IP addresses are used as the self IP address for the VLAN.
7. You can also define an IP pool for the NSX virtual server when you deploy it. This step is optional.

Specifying VLANs on the interfaces to be provisioned

Specify VLANs for the interfaces that you plan to use in the VMware NSX integration.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. On the Devices panel, select the device you want to configure, then click the gear icon and select **Properties**.
3. If necessary, under VLAN Connectivity, select the plus sign to add a new entry to the VLAN Pools by Interface list.
4. In the **Interface** field, type in the interface number to which the tagged interfaces are connected.

Important: The BIG-IP device may use untagged, named interfaces in addition to the tagged, numbered interfaces that you configure here. Do not list the untagged interfaces on the BIG-IQ device.

5. In the **VLAN Pool** field, type in the tag numbers assigned to the VLANs connected to this interface.

Entries you make here can be comma-delimited to indicate specific tag numbers, or you can use a dash to indicate a range of tags, or both. (For instance, you could enter 20-60, 90 to indicate the range of tags that exist from 20 to 60, and tag 90.)

6. Click **Update**.

The tagged VLANs you identified will be made available for the integration.

Creating a connection between BIG-IQ Cloud and NSX Manager

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.

Important: You will need to recall the name you assign to this connector so that you can select it when you are configuring the VMware user interface. The name you specify is used as the service definition name in the VMware user interface.

4. From the **Cloud Provider** list, select **VMware NSX 6.1**.
The screen displays additional settings specific to VMware NSX.
5. From the **Devices** list, select the device you want to associate with this connector.
6. In the **VMware NSX Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device.
7. In the **VMware NSX User Name** and **VMware NSX Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the NSX Manager REST API.
8. In the **VMware vCenter Server Address** field, type the IP address of the vCenter server.
9. In the **VMware vCenter Server User Name** and **VMware vCenter Server Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the vCenter SOAP API.
10. From the **BIG-IQ Callback User Name** list, select the user name that NSX Manager uses to authenticate to the BIG-IQ REST API.
11. In the **BIG-IQ Callback Password** field, type the password that NSX Manager uses to authenticate to the BIG-IQ REST API.
12. From the **BIG-IQ Callback Address** list, select the IP address that this NSX Manager uses to access each BIG-IQ device in the HA cluster.
By default, the management IP address is used, but you can specify a self IP address if you choose.
13. Click the **Save** button.

As part of the connection creation process, the BIG-IQ system does the following:

- Creates a new default tenant for the new connector.
- Verifies connectivity to the NSX Manager and vCenter APIs, and registers the BIG-IQ system as an NSX Partner Service provider.
- Creates a callback user role that enables NSX to access the BIG-IQ software resources necessary for interaction with the BIG-IQ REST API.

Creating an application template for NSX

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

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

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

1. Hover over the Catalog header, click the + icon when it appears.
The panel expands to display the application template properties.
2. In the **Name** field, type a name for this new application.
3. From the **Cloud Connector** list select the just-created NSX cloud connector.
4. From the **Application Type** list, select an application.
5. If the **Application Tiers** settings are displayed (expanded), select the options that match the properties for this application; otherwise, keep the default settings.

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

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

The customized application displays as an entry in the catalog. It will also display as an NSX service profile in the VMware NSX 6.1 user interface.

Confirming that connector is recognized as an NSX service definition

The NSX connector you created on BIG-IQ[®] Cloud, must be recognized by vSphere Web Client as a Service Definition.

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

On the Networking and Security panel, select **Service Definitions** and confirm that the NSX connector you created previously appears in the list of recognized service definitions.

Creating an NSX Edge Services Gateway for the BIG-IP device

The NSX Edge Service Gateway establishes the network within which network services such as firewall, NAT, and load balancing are deployed. To integrate a BIG-IP® device with NSX, you must create at least one Edge Service Gateway.

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

In the vSphere web client user interface, create a new NSX Edge.

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.

Both deployed and undeployed modes are supported. But using undeployed mode will simplify implementation.

When you finish editing an Edge, it appears in the list under NSX Edges.

Creating a load balancing service instance for VLANs bridged to a VXLAN

You create an NSX service instance to provide the load balancing service.

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

1. On the Networking and Security pane, select **NSX Edges** and then select the just-created NSX Edge Services Gateway.
2. Ensure that **Manage** is active, and select **Load Balancer**, and then select **Global Configuration**. The Edit Load balancer global configuration screen is displayed.
3. Click the **Edit** button. The Edit Load balancer global configuration screen is displayed.
4. Select both **Enable Load Balancer** and **Enable Service Insertion**.
5. For the **Service Definition**, select the name of the connector you created for NSX previously.
6. For the **Service Configuration**, select **F5 ADC - Use existing BIG-IP VE**.
7. Click the down arrow to expand the **Service Instance Runtime Configuration Parameters** area and then specify settings for the tagged VLANs.
 - a) Select an available vNIC, and then click the edit icon (✎). The Edit Network window is displayed.
 - b) In the **Name** field, type a name to identify the VLAN.
 - c) For **Connectivity Type**, select **Data**.
 - d) For the **Connected To** field, click **Select**. The Select Network screen opens.
 - e) Select **Logical Switch** and then select the VXLAN to which the server pool members are connected, and then click **OK**. The VXLAN you select will be bridged to a VLAN. The Select Network window closes, and the route for this tagged VLAN is created.

- f) Select **IP_POOL** for the **IP allocation Mode**, select the appropriate pool, and then click **OK**.
The Edit Network window closes and the new Service Instance Runtime Configuration that you configured is displayed in place of vNIC0.
 - g) Repeat the preceding six sub-steps, (as needed) until you have specified settings for each tagged VLAN that you want to use.
8. Click the down arrow to expand the **Tyred Service Configuration Parameters** area and then type the IP address of the BIG-IP[®] device in the value field that corresponds to the second entry (F5 BIG-IP address/host).
 9. Click **OK** to close the Edit Load balancer global configuration dialog box.

VMware NSX configures the service instance based on the settings you specified, and associates it with the BIG-IP[®] device.

Specifying pools for the virtual server

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ[®]Cloud.

You specify the virtual resources so that the iApp template has one on which to deploy.

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

1. On the Networking and Security panel, select **NSX Edges** and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Pools**.
The list of existing pools is displayed.
3. Click the add icon (+).
The New Pool screen is displayed.
4. In the **Name** field, type a name for the new server pool.
5. Under Members, click the add icon (+).
The New Member screen is displayed.
6. Select **Enable Member**.
7. In the **Name** field, type a name for the new pool member.
8. In the **IP Address/VC Container** field, type the IP address for the new pool member.

Important: Server pool members must be on the same portgroup or network you used when creating the load balancing service instance. In this case, use the network that you specified when you configured the load balancing service instance.

9. In the **Port** field, type 80.
10. Click **OK**.
The New Member screen closes.
11. Click **OK**.

The New Pool screen closes, and VMware NSX creates the new pool.

Specifying virtual servers for the load balancer

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ® Cloud.

You specify the virtual server on which you want the iApp template to deploy.

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

1. On the Networking and Security panel, select **NSX Edges**, and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Virtual Servers**. The list of existing virtual servers is displayed.
3. Click the add icon (+). The New Virtual Server window is displayed.
4. Select **Enable Virtual Server**.
5. In the **Application Profile** field, select the name of the iApp template you created on the BIG-IQ device for this iApp deployment.
6. In the **Name** field, type a name for this virtual server.
7. In the **IP Address** field, specify the IP address of the virtual server just created.

Important: The Virtual Server must be on the same portgroup or network you used when creating the load balancing service instance.

There are two ways to specify the IP address.

- Type the IP address in the **IP Address** field.
- Click **Select IP Pool** and then choose the IP pool and click **OK**.

8. In the **Port** field, type 80.
9. From the **Default Pool** list select the just-created pool.
10. If you specified Tenant Editable Application Properties when you created the application template, select the Advanced tab to display and revise them as necessary.
11. Click **OK**.
The New Virtual Server screen closes, and VMware NSX creates the new virtual server.

VMware NSX deploys the virtual server. The next time you log in to the BIG-IQ device, the virtual server should appear on the Applications tab.

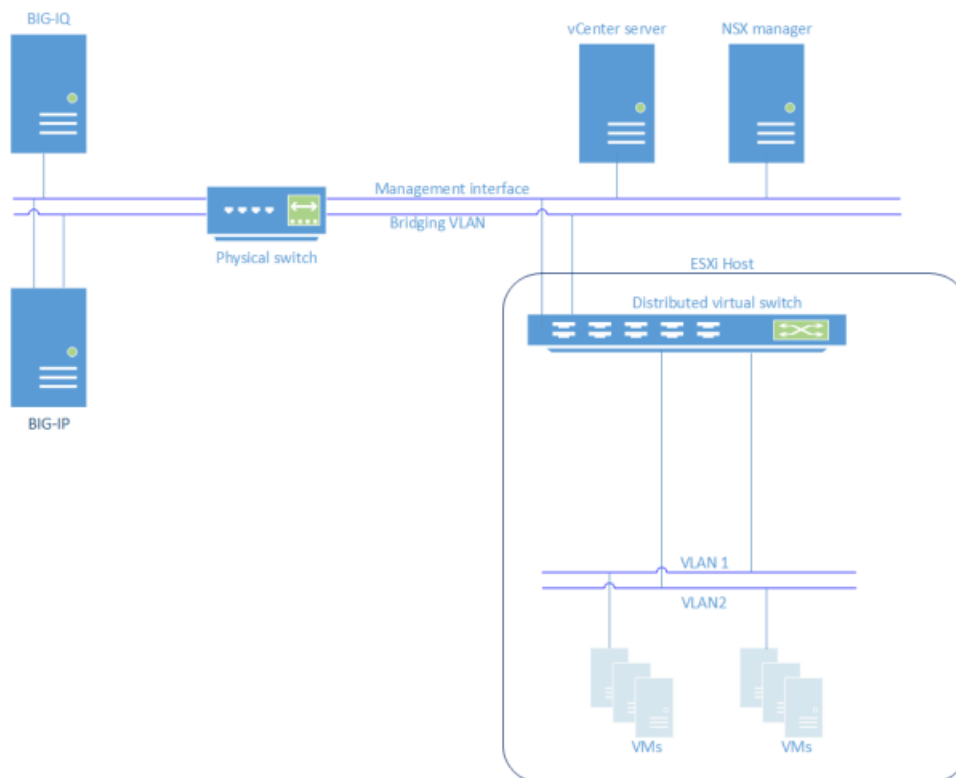
Before you begin using the applications just deployed, you should make sure that the application is healthy in the BIG-IQ Cloud.

About integrating VMware NSX with a BIG-IP device using tagged interface VLANs

The integration between BIG-IQ® Cloud and VMware NSX makes it possible for you to use existing physical BIG-IP® devices to host NSX virtual servers. Using these servers, you can manage and deploy iApps® on

existing VMware NSX environments. One connectivity option for this integration uses VLAN trunks that the BIG-IQ device provisions for you. These VLANs must use a tagged interface.

This figure illustrates the network topology for this connectivity option.



Task summary

- Setting up a VMware network for a tagged interface VLAN pool
- Discovering devices located in the VMware cloud
- Specifying VLANs on the interfaces to be provisioned
- Creating a connection between BIG-IQ Cloud and NSX Manager
- Creating an application template for NSX
- Confirming that connector is recognized as an NSX service definition
- Creating an NSX Edge Services Gateway for the BIG-IP device
- Creating a load balancing service instance for tagged VLANs
- Specifying pools for the virtual server
- Specifying virtual servers for the load balancer

Setting up a VMware network for a tagged interface VLAN pool

Before you can begin configuring the BIG-IQ[®] device integration for a tagged interface VLAN pool, you must configure the VMware network.

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

1. You must have discovered the device on which your VLANs will reside before you can configure them from the BIG-IQ device.

2. Connect the BIG-IP® device and the ESXi host to a physical switch.
The BIG-IP device and the ESXi host must connect to a port with multiple tagged VLANs.
3. Create a vCenter distributed port group for each VLAN in the trunk. Each port group must have a unique VLAN tag.
This configuration is referred to as *virtual switch tagging* (VST). It ensures that outbound traffic to the BIG-IP device is tagged.
4. Define the IP pool for each port group in NSX.
These IP addresses are used as the self IP address for the VLAN.
5. You also have the option to define an IP pool for the NSX virtual server when you deploy it. This step is not required.

Discovering devices located in the VMware cloud

After you license and perform the initial configuration for the BIG-IQ® system, you can discover BIG-IP® devices running version 11.5 or later. For proper communication between the managing BIG-IQ system and the devices it manages, you must configure the BIG-IQ system with a route to each F5 device you want to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

You must know the IP address that the BIG-IQ device will use to access the BIG-IP device.

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

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. Hover over the Devices header, click the + icon when it appears, and then select **New Device**.
The Devices panel expands to show the New Device screen.
3. In the **IP Address** field, type the device's IP address.
The preferred address for discovering a BIG-IP device is its management IP address.
4. (This step applies only if the BIG-IQ system is not hosted on AWS version 4.4 or later.) If the BIG-IQ system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IQ system using SSH to specify an IP route between them.
 - If the BIG-IQ system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

5. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.

6. For the **Auto Update Framework** setting, select the **Update Automatically** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: *When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).*

7. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

If you want to use the BIG-IP device just discovered to host NSX virtual servers, you should now associate it with a VMware cloud connector.

Specifying VLANs on the interfaces to be provisioned

Specify VLANs for the interfaces that you plan to use in the VMware NSX integration.

1. Log in to BIG-IQ® Cloud with your administrator user name and password.
2. On the Devices panel, select the device you want to configure, then click the gear icon and select **Properties**.
3. If necessary, under VLAN Connectivity, select the plus sign to add a new entry to the VLAN Pools by Interface list.
4. In the **Interface** field, type in the interface number to which the tagged interfaces are connected.

Important: *The BIG-IP device may use untagged, named interfaces in addition to the tagged, numbered interfaces that you configure here. Do not list the untagged interfaces on the BIG-IQ device.*

5. In the **VLAN Pool** field, type in the tag numbers assigned to the VLANs connected to this interface.
Entries you make here can be comma-delimited to indicate specific tag numbers, or you can use a dash to indicate a range of tags, or both. (For instance, you could enter 20-60, 90 to indicate the range of tags that exist from 20 to 60, and tag 90.)
6. Click **Update**.

The tagged VLANs you identified will be made available for the integration.

Creating a connection between BIG-IQ Cloud and NSX Manager

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

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

2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.

Important: *You will need to recall the name you assign to this connector so that you can select it when you are configuring the VMware user interface. The name you specify is used as the service definition name in the VMware user interface.*

4. From the **Cloud Provider** list, select **VMware NSX 6.1**.
The screen displays additional settings specific to VMware NSX.
5. From the **Devices** list, select the device you want to associate with this connector.
6. In the **VMware NSX Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device.
7. In the **VMware NSX User Name** and **VMware NSX Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the NSX Manager REST API.
8. In the **VMware vCenter Server Address** field, type the IP address of the vCenter server.
9. In the **VMware vCenter Server User Name** and **VMware vCenter Server Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the vCenter SOAP API.
10. From the **BIG-IQ Callback User Name** list, select the user name that NSX Manager uses to authenticate to the BIG-IQ REST API.
11. In the **BIG-IQ Callback Password** field, type the password that NSX Manager uses to authenticate to the BIG-IQ REST API.
12. From the **BIG-IQ Callback Address** list, select the IP address that this NSX Manager uses to access each BIG-IQ device in the HA cluster.
By default, the management IP address is used, but you can specify a self IP address if you choose.
13. Click the **Save** button.

As part of the connection creation process, the BIG-IQ system does the following:

- Creates a new default tenant for the new connector.
- Verifies connectivity to the NSX Manager and vCenter APIs, and registers the BIG-IQ system as an NSX Partner Service provider.
- Creates a callback user role that enables NSX to access the BIG-IQ software resources necessary for interaction with the BIG-IQ REST API.

Creating an application template for NSX

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

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

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

1. Hover over the Catalog header, click the + icon when it appears.
The panel expands to display the application template properties.
2. In the **Name** field, type a name for this new application.

3. From the **Cloud Connector** list select the just-created NSX cloud connector.
4. From the **Application Type** list, select an application.
5. If the **Application Tiers** settings are displayed (expanded), select the options that match the properties for this application; otherwise, keep the default settings.

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

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

The customized application displays as an entry in the catalog. It will also display as an NSX service profile in the VMware NSX 6.1 user interface.

Confirming that connector is recognized as an NSX service definition

The NSX connector you created on BIG-IQ[®] Cloud, must be recognized by vSphere Web Client as a Service Definition.

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

On the Networking and Security panel, select **Service Definitions** and confirm that the NSX connector you created previously appears in the list of recognized service definitions.

Creating an NSX Edge Services Gateway for the BIG-IP device

The NSX Edge Service Gateway establishes the network within which network services such as firewall, NAT, and load balancing are deployed. To integrate a BIG-IP[®] device with NSX, you must create at least one Edge Service Gateway.

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

In the vSphere web client user interface, create a new NSX Edge.

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.

Both deployed and undeployed modes are supported. But using undeployed mode will simplify implementation.

When you finish editing an Edge, it appears in the list under NSX Edges.

Creating a load balancing service instance for tagged VLANs

You should create an NSX service instance for each VLAN interface that will provide load balancing services.

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

1. On the Networking and Security pane, select **NSX Edges** and then select the just-created NSX Edge Services Gateway.
2. Ensure that **Manage** is active, and select **Load Balancer**, and then select **Global Configuration**. The Edit Load balancer global configuration window is displayed.
3. Click the **Edit** button. The Edit Load balancer global configuration window is displayed.
4. Select both **Enable Load Balancer** and **Enable Service Insertion**.
5. For the **Service Definition**, select the name of the connector you created for NSX previously.
6. For the **Service Configuration**, select **F5 ADC - Use an existing BIG-IP VE**.
7. Click the down arrow to expand the **Service Instance Runtime Configuration Parameters** area and then specify settings for the tagged VLANs.
 - a) Select an available vNIC, and then click the edit icon (✎). The Edit Network screen opens.
 - b) In the **Name** field, type a name to identify the VLAN.
 - c) For **Connectivity Type**, select **Data**.
 - d) For **Connected To** field click **Select**. The Select Network window is displayed.
 - e) Select one of the distributed port groups associated with the tagged VLANs that are configured on the BIG-IP® device, and then click **OK**. The Select Network window closes, and the route for this tagged VLAN is created.
 - f) Select **IP Pool** for the **IP allocation Mode**, select the appropriate pool, and then click **OK**. The Edit Network window closes, and the new Service Instance Runtime Configuration you configured is displayed in place of vNIC0.
 - g) Repeat the previous six sub-steps, (as needed) until you have specified settings for each tagged VLAN you want to use.
8. Click the down arrow to expand the **Typed Service Configuration Parameters** area, and then type the IP address of the BIG-IP device in the value field that corresponds to the second entry (F5 BIG-IP address/host).
9. Click **OK** to close the Edit Load balancer global configuration dialog box.

VMware NSX configures the service instance based on the settings you specified and associates it with the BIG-IP device.

Specifying pools for the virtual server

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ[®] Cloud.

You specify the virtual resources so that the iApp template has one on which to deploy.

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

1. On the Networking and Security panel, select **NSX Edges** and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Pools**. The list of existing pools is displayed.
3. Click the add icon (+). The New Pool screen is displayed.
4. In the **Name** field, type a name for the new server pool.
5. Under Members, click the add icon (+). The New Member screen is displayed.
6. Select **Enable Member**.
7. In the **Name** field, type a name for the new pool member.
8. In the **IP Address/VC Container** field, type the IP address for the new pool member.

Important: Server pool members must be on the same portgroup or network you used when creating the load balancing service instance. In this case, use the network that you specified when you configured the load balancing service instance.

9. In the **Port** field, type 80.
10. Click **OK**. The New Member screen closes.
11. Click **OK**.

The New Pool screen closes, and VMware NSX creates the new pool.

Specifying virtual servers for the load balancer

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ[®] Cloud.

You specify the virtual server on which you want the iApp template to deploy.

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

1. On the Networking and Security panel, select **NSX Edges**, and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Virtual Servers**. The list of existing virtual servers is displayed.

3. Click the add icon (+).
The New Virtual Server window is displayed.
4. Select **Enable Virtual Server**.
5. In the **Application Profile** field, select the name of the iApp template you created on the BIG-IQ device for this iApp deployment.
6. In the **Name** field, type a name for this virtual server.
7. In the **IP Address** field, specify the IP address of the virtual server just created.

Important: *The Virtual Server must be on the same portgroup or network you used when creating the load balancing service instance.*

There are two ways to specify the IP address.

- Type the IP address in the **IP Address** field.
- Click **Select IP Pool** and then choose the IP pool and click **OK**.

8. In the **Port** field, type 80.
9. From the **Default Pool** list select the just-created pool.
10. If you specified Tenant Editable Application Properties when you created the application template, select the Advanced tab to display and revise them as necessary.
11. Click **OK**.
The New Virtual Server screen closes, and VMware NSX creates the new virtual server.

VMware NSX deploys the virtual server. The next time you log in to the BIG-IQ device, the virtual server should appear on the Applications tab.

Before you begin using the applications just deployed, you should make sure that the application is healthy in the BIG-IQ Cloud.

About integrating VMware NSX with a BIG-IP device using existing VLANs

The integration between BIG-IQ® Cloud and VMware NSX makes it possible for you to use existing physical BIG-IP® devices to host virtual servers. One connectivity option for this integration uses VLANs that you configure on the BIG-IP device. These VLANs can be either tagged or untagged depending on your network topology.

Task summary

Setting up a VMware network for an existing VLAN

Creating a connection between BIG-IQ Cloud and NSX Manager

Creating an application template for NSX

Confirming that connector is recognized as an NSX service definition

Creating an NSX Edge Services Gateway for the BIG-IP device

Creating a load balancing service instance for existing VLANs

Specifying pools for the virtual server

Specifying virtual servers for the load balancer

Setting up a VMware network for an existing VLAN

Before you can begin configuring the BIG-IQ[®] device integration for a existing VLAN, you must configure the VMware network.

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

1. You must have discovered the device on which your VLANs will reside before you can configure them from the BIG-IQ device.
2. Make sure that the BIG-IP[®] device has VLANs (internal and external) and self-IPs configured.
3. Create a vCenter distributed port group for the internal and external VLANs.
Each port group must be untagged (that is, the VLAN ID must be set to 0).
4. You can also define an IP pool for the NSX virtual server when you deploy it. This step is optional.

Creating a connection between BIG-IQ Cloud and NSX Manager

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

1. Log in to BIG-IQ[®] Cloud with your administrator user name and password.
2. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
3. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.

Important: You will need to recall the name you assign to this connector so that you can select it when you are configuring the VMware user interface. The name you specify is used as the service definition name in the VMware user interface.

4. From the **Cloud Provider** list, select **VMware NSX 6.1**.
The screen displays additional settings specific to VMware NSX.
5. From the **Devices** list, select the device you want to associate with this connector.
6. In the **VMware NSX Address** field, type the IP address of the VMware system.
The VMware IP address must be fully accessible from the BIG-IQ device.
7. In the **VMware NSX User Name** and **VMware NSX Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the NSX Manager REST API.
8. In the **VMware vCenter Server Address** field, type the IP address of the vCenter server.
9. In the **VMware vCenter Server User Name** and **VMware vCenter Server Password** fields, type the credentials that the BIG-IQ device will use to authenticate to the vCenter SOAP API.
10. From the **BIG-IQ Callback User Name** list, select the user name that NSX Manager uses to authenticate to the BIG-IQ REST API.
11. In the **BIG-IQ Callback Password** field, type the password that NSX Manager uses to authenticate to the BIG-IQ REST API.
12. From the **BIG-IQ Callback Address** list, select the IP address that this NSX Manager uses to access each BIG-IQ device in the HA cluster.

By default, the management IP address is used, but you can specify a self IP address if you choose.

13. Click the **Save** button.

As part of the connection creation process, the BIG-IQ system does the following:

- Creates a new default tenant for the new connector.
- Verifies connectivity to the NSX Manager and vCenter APIs, and registers the BIG-IQ system as an NSX Partner Service provider.
- Creates a callback user role that enables NSX to access the BIG-IQ software resources necessary for interaction with the BIG-IQ REST API.

Creating an application template for NSX

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

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

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

1. Hover over the Catalog header, click the + icon when it appears.
The panel expands to display the application template properties.
2. In the **Name** field, type a name for this new application.
3. From the **Cloud Connector** list select the just-created NSX cloud connector.
4. From the **Application Type** list, select an application.
5. If the **Application Tiers** settings are displayed (expanded), select the options that match the properties for this application; otherwise, keep the default settings.

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

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

The customized application displays as an entry in the catalog. It will also display as an NSX service profile in the VMware NSX 6.1 user interface.

Confirming that connector is recognized as an NSX service definition

The NSX connector you created on BIG-IQ[®] Cloud, must be recognized by vSphere Web Client as a Service Definition.

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

On the Networking and Security panel, select **Service Definitions** and confirm that the NSX connector you created previously appears in the list of recognized service definitions.

Creating an NSX Edge Services Gateway for the BIG-IP device

The NSX Edge Service Gateway establishes the network within which network services such as firewall, NAT, and load balancing are deployed. To integrate a BIG-IP[®] device with NSX, you must create at least one Edge Service Gateway.

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

In the vSphere web client user interface, create a new NSX Edge.

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.

Both deployed and undeployed modes are supported. But using undeployed mode will simplify implementation.

When you finish editing an Edge, it appears in the list under NSX Edges.

Creating a load balancing service instance for existing VLANs

You create an NSX service instance to provide the load balancing service.

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

1. On the Networking and Security pane, select **NSX Edges** and then select the just created NSX Edge Services Gateway.
2. Ensure that **Manage** is active, and select **Load Balancer**, and then select **Global Configuration**. The Edit Load balancer global configuration screen opens.
3. Click the **Edit** button. The Edit Load balancer global configuration screen opens.
4. Select both **Enable Load Balancer** and **Enable Service Insertion**.
5. For the **Service Definition**, select the name of the connector you created for NSX previously.
6. For the **Service Configuration**, select **F5 ADC - Use an existing BIG-IP**.

7. Click the down arrow to expand the **Service Instance Runtime Configuration Parameters** area, and then specify settings for each of the vNICs.
 - a) Select **vnic0**, and then click the edit icon (✎).
The Edit Network screen opens.
 - b) In the **Name** field, type a name that describes the role this subnet plays in your network.
 - c) For **Connectivity Type**, select **Data**.
 - d) For **Connected To**, click **Select**.
The Select Network screen opens.
 - e) Select the distributed port group associated with the VLANs that are configured on the BIG-IP device, and then click **OK**.
The Select Network window closes.
 - f) Select **DHCP** (the default setting) for the **IP allocation Mode**, and then click **OK**.
The Edit Network window closes and the new Service Instance Runtime Configuration you configured is displayed in place of **vnic0**.
 - g) Repeat the previous six sub-steps, but this time configure settings for **vnic1** and specifying the name **internal**.
8. Click the down arrow to expand the **Typed Service Configuration Parameters** area and then type the IP address of the BIG-IP device in the value field that corresponds to the second entry (F5 BIG-IP address/host).
9. Click **OK** to close the Edit Load balancer global configuration dialog box.

VMware NSX configures the service instance based on the settings you specified, and associates it with the BIG-IP device.

Specifying pools for the virtual server

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ® Cloud.

You specify the virtual resources on which you want the iApp template to deploy.

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

1. On the Networking and Security panel, select **NSX Edges** and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Pools**.
The list of existing pools is displayed.
3. Click the add icon (+).
The New Pool screen opens.
4. In the **Name** field, type a name for the new server pool.
5. Under Members, click the add icon (+).
The New Member screen opens.
6. Select **Enable Member**.
7. In the **Name** field, type a name for the new pool member.
8. In the **IP Address/VC Container** field, type the IP address for the new pool member.

Important: Server pool members must be on the same portgroup or network you used when creating the load balancing service instance. In this case, use the network that you specified when you configured NIC1 (internal).

9. In the **Port** field, type 80.
10. Click **OK**.
The New Member screen closes.
11. Click **OK**.

The New Pool screen closes and VMware NSX creates the new pool.

Specifying virtual servers for the load balancer

Before you can perform this task, you must have created an iApp template that uses the NSX connector on BIG-IQ® Cloud.

You specify the virtual server on which you want the iApp template to deploy.

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

1. On the Networking and Security panel, select **NSX Edges**, and then select the just created NSX Edge Services Gateway.
2. Ensure that the **Manage** option is active, select **Load Balancer**, and then select **Virtual Servers**.
The list of existing virtual servers is displayed.
3. Click the add icon (+).
The New Virtual Server window is displayed.
4. Select **Enable Virtual Server**.
5. In the **Application Profile** field, select the name of the iApp template you created on the BIG-IQ device for this iApp deployment.
6. In the **Name** field, type a name for this virtual server.
7. In the **IP Address** field, specify the IP address of the virtual server just created.

Important: The Virtual Server must be on the same portgroup or network you used when creating the load balancing service instance.

There are two ways to specify the IP address.

- Type the IP address in the **IP Address** field.
- Click **Select IP Pool** and then choose the IP pool and click **OK**.

8. In the **Port** field, type 80.
9. From the **Default Pool** list select the just-created pool.
10. If you specified Tenant Editable Application Properties when you created the application template, select the Advanced tab to display and revise them as necessary.
11. Click **OK**.
The New Virtual Server screen closes, and VMware NSX creates the new virtual server.

VMware NSX deploys the virtual server. The next time you log in to the BIG-IQ device, the virtual server should appear on the Applications tab.

Before you begin using the applications just deployed, you should make sure that the application is healthy in the BIG-IQ Cloud.

Chapter

9

Local Cloud Integration

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

About using a local cloud source

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

Discovering BIG-IP devices in your network

After you license and perform the initial configuration for the BIG-IP[®] system, you can discover BIG-IP[®] devices running version 11.5 or later. For proper communication between the managing BIG-IP system and the devices it manages, you must configure the BIG-IP system with a route to each F5 device you want to manage. If you do not specify the required network communication route between the devices, then device discovery fails.

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

1. Log in to BIG-IP Device with your administrator user name and password.
2. Hover over the Devices header, click the + icon when it appears, and then select **New Device**. The Devices panel expands to show the New Device screen.
3. In the **IP Address** field, type the IP address of the device. The preferred address for discovering a BIG-IP device is its management IP address.
4. (This step applies only when the BIG-IP system is hosted on AWS version 4.4 or later.) If the BIG-IP system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IP system using SSH to specify an IP route between them.
 - If the BIG-IP system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.
 1. Use SSH to log in to the BIG-IP system's management IP address as an admin user.
 2. Type the following command: `run /util bash`
 3. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
 - If the BIG-IP system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IP system's management IP address as an admin user.
 2. Type the following command: `create net route <route name> {gw <x.x.x.x> network default}`

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

5. (This step applies only if the BIG-IP system is not hosted on AWS version 4.4 or later.) If the BIG-IP system and the BIG-IP device are on different subnets, then you need to log in to the BIG-IP system using SSH to specify an IP route between them.
 - If the BIG-IP system and the BIG-IP device communicate using the management IP address, then there must be a default route specified. If there is no default route, issue a `route` command.

1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `route <route name> {gw <x.x.x.x> network default}`
- If the BIG-IQ system and the BIG-IP device use something other than the management IP address to communicate, then issue a `tmsh route` command.
 1. Use SSH to log in to the BIG-IQ system's management IP address as the root user.
 2. Type the following command: `tmsh create net route <route name> {gw <x.x.x.x> network default}`

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

6. To change the root user name, type a new name in the **Root User Name** field.
7. Type a password for the root user in the **Root Password** field.
8. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
9. For the **Auto Update Framework** setting, select the **Update Automatically** check box to direct the BIG-IQ system to perform any required REST framework updates on the BIG-IP device.
For the BIG-IQ system to properly manage a BIG-IP device, the BIG-IP device must be running the most recent REST framework.

Important: When you update the REST framework for BIG-IP devices running version 11.6 or earlier, the traffic management interface (TMM) restarts. Before you update the REST framework on a BIG-IP device, verify that no critical network traffic is targeted to that device. Additionally, In any system upgrade scenario, the potential exists for unexpected errors. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, it is conceivable that a BIG-IP device would not be left in the pre-discovery state. If you want to roll back the upgrade due to an error or any other reason, the recommended recovery for this situation is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).

10. Click the **Add** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the Devices panel. Its configuration files display in the Configuration panel.

Associating a local cloud connector with a device

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

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

1. Hover over the Connectors header, click the + icon when it appears, and then click **New Connector**.
2. In the **Name** and **Description** fields, type a name and description.
You can use the name and description to help you organize network resources into logical groups based on certain criteria, such as the location or application.
3. From the **Cloud Provider** list, select **Local Cloud**.

4. From the **Devices** list, select the device you want to associate with this connector.
5. To select additional devices to associate with this connector, click the + icon at the right of the list. BIG-IQ system discovers application servers associated with this connector, and populates them in the Server panel. If the system discovers F5 devices, it populates the Device panel with them.
6. Click the **Save** button.

Chapter 10

Cloud Tenant Management

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

About creating cloud tenants

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

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

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

Creating a tenant

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

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

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

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

Creating a cloud user

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

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

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

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

Associating a user with a tenant's role

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

Attention: *The BIG-IQ system administrator creates roles from the **BIG-IQ System > Access Control** menu. For more information, refer to the **BIG-IQ® System: Licensing and Initial Configuration** guide.*

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

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

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

Chapter

11

iApps Application Template Customization

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

About customizing iApp application templates

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

Note: There are actually two types of iApp application templates.

- *BIG-IP iApp templates* are imported from your managed devices. They are created by F5 developers or third party developers using TCL and APL code. For details on how to create these, refer to the *BIG-IP iApps Developer's Guide* on AskF5, or *iApp Template Development Tips and Techniques* on DevCentral.
- Once imported to the BIG-IQ device, we refer to these templates as *BIG-IQ application templates*. You provide customized versions of these to your tenants.

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

Optional load balancing methods for f5.http, f5.microsoft_sharpoint_2010, and f5.microsoft_iss catalog templates

For **f5.http**, **f5.microsoft_sharepoint_2010**, or **f5.microsoft_iss** catalog templates (only), you can specify a load balancing method other than the default, `least connection member`, by typing a value for the **Which load balancing method do you want to use?** setting.

Load Balancing Method	Value
Dynamic ratio member	dynamic-ratio-member
Dynamic ratio node	dynamic-ratio-node
Fastest application response	fastest-app-response
Fastest node	fastest-node
Least connections node	least-connections-node
Least sessions	least-sessions
Observed member	observed-member
Observed node	observed-node
Round robin	round-robin
Ratio member	ratio-member
Ratio node	ratio-node
Ratio session	ratio-session
Ratio least connections member	ratio-least-connections-member
Ratio least connections node	ratio-least-connections-node

Load Balancing Method	Value
Weighted least connections member	weighted-least-connections-member

Creating a customized application template

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

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

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

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

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

6. To allow cloud tenants to specify certificates with SSL encryption when self-deploying applications, select options from the **SSL Cert** and **SSL Key** lists.
BIG-IQ® Cloud uses these options to provide the appropriate certificate and key when the tenant self-deploys this application to a BIG-IP® device. These options are not available for all application templates.
7. Finish making modifications by specifying the Application Properties and Customize Application Template variables.
To allow a tenant to modify a particular setting, select the **Tenant Editable** check box for that setting. For further details about template variables and settings, refer to the *BIG-IP® iApps® Developer's Guide*.
8. If you selected **f5.http**, **f5.microsoft_sharepoint_2010**, or **f5.microsoft_iss** and you want to specify a load balancing option other than the default, Least Connection Member, perform the following steps:
 - a) Click the arrow next to Advanced Properties.
 - b) In the **Which load balancing method do you want to use?** field, type the value for the option you want to use.
9. Click the **Save** button.
You can now send the cloud IP addresses to the tenant and use this IP address range in configuring server tiers and pool members, within certain application services. The tenant can self-deploy the application from the catalog.

The customized application displays as an entry in the catalog.

Chapter 12

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>provider</i>	The terms provider or (cloud provider) are used interchangeably with the term cloud administrator/administrator.
<i>resources</i>	A resource is any managed object, including devices, web applications, virtual servers, servers, cloud connectors, and so forth.
<i>roles</i>	A role defines specific privileges to which you can associate one or more users. There are two default roles for BIG-IQ Cloud: cloud administrator and cloud tenant.
<i>tenant</i>	A tenant is an entity that can consist of one or more users accessing resources provided by a cloud administrator.
<i>user</i>	A user is an individual who has been granted access to specific tenant resources.

Index

A

- active-active cluster
 - configuring for the BIG-IQ system 22
 - active-active pair
 - configuring for the BIG-IQ system 22
 - activities
 - viewing for cloud resource activity 47
 - Amazon CloudWatch
 - about integrating with BIG-IQ Cloud 38
 - Amazon EC2 47
 - See also EC2
 - about integrating with BIG-IQ Cloud 38
 - See also EC2
 - Amazon EC2 cloud
 - creating BIG-P VE version 11.3 or 11.4 devices 43
 - Amazon EC2 devices
 - discovering 42
 - Amazon EC2 resources
 - viewing cloud resource activity 47
 - Amazon Elastic Compute Cloud, See Amazon EC2
 - Amazon Machine Images, See AMI
 - Amazon Virtual Private Cloud 39
 - See also VPC
 - about integrating with BIG-IQ Cloud 38
 - creating 39
 - See also VPC
 - Amazon virtual server
 - launching using an AMI 40
 - Amazon web services
 - and network configuration requirements 38
 - AMI
 - using to launch a virtual server 40
 - application catalog 45, 72, 79, 85, 101
 - applications
 - customizing for tenants 45, 72, 79, 85, 101
 - deploying as a Cloud Tenant user 46
 - application templates
 - about 100
 - and vApps 56
 - defined 104
 - deploying as a Cloud Tenant user 46
 - using 45, 72, 79, 85, 101
 - asset management
 - for devices 25
 - authentication 38
- ## B
- BIG-IP device
 - and service definition 84
 - service definition 69, 72, 76, 80, 86
 - BIG-IP devices
 - provisioning and discovering 62–63
 - BIG-IP VE
 - configuring NSX for 63
 - provisioning 66
 - setting up VMware network 62–63

- BIG-IP virtual resources
 - specifying 74, 82, 87
- BIG-IQ Cloud
 - configuring for BIG-IP VE 63
 - defined 104
 - finding documentation for 18
- BIG-IQ cloud tenant
 - finding documentation for 18
- BIG-IQ system
 - finding documentation for 18
 - reordering panels 19
- billing
 - for utility licenses 33

C

- callback users
 - adding 64
- catalog
 - for applications 45, 72, 79, 85, 101
- catalog entries
 - creating for tenants 100
- Catalog entries
 - for f5.http, f5.microsoft_sharepoint_2010, or f5.microsoft_iss iApp 100
- cloud administrator
 - defined 104
- cloud bursting
 - defined 104
- cloud connector
 - 56, 58
 - defined 104
 - for OpenStack 52
 - for VMware NSX 65, 71, 78, 84
- cloud connector, local
 - associating with a device 93
- cloud connector for EC2 41
- cloud resources
 - providing for tenants 96, 100
 - using locally 92
- cloud tenant
 - adding 96
- cloud tenants
 - about creating 96
- Cloud Tenant users
 - deploying applications 46
- clusters
 - for high availability 22
- connector, local
 - associating with a device 93
- connectors
 - VMware 56

D

- device discovery
 - by scanning network 24, 92

- device inventory
 - about 24
 - viewing details 25
- device management
 - about 24
- devices
 - about discovering 24
 - adding 24, 92
 - creating in Amazon EC2 cloud 43
 - discovering Amazon EC2 devices 42
 - discovering OpenStack devices 51
 - discovering VMware devices 56, 59, 70, 77–78
- documentation, finding 18
- dynamic cloud resources
 - viewing activity for 47

E

- EC2 connector
 - associating with a device 41
- Edge Services Gateway
 - creating for NSX 73, 80, 86
- elasticity
 - viewing activity for 47
- exportation of inventory details 25

F

- f5.http iApp
 - specifying load balance option 100
- f5.microsoft_iss iApp
 - specifying load balance option 100
- f5.microsoft_sharepoint_2010 iApp
 - specifying load balance option 100
- failover 22
- filtering process
 - finding associated objects 18

G

- glossary 104
- Grizzly, See OpenStack
- guides, finding 18

H

- HA, See high availability cluster
- Havana, See OpenStack
- high availability cluster
 - configuring 22
- high availability configuration
 - about 22

I

- IAM
 - creating user account 38
- IAM access 47
- iApps
 - customizing for tenants 100
 - defined 100

- iApps (*continued*)
 - for OpenStack 50
- iApp templates
 - specifying load balancing option for f5.http, f5.microsoft_sharepoint_2010, or f5.microsoft_iss 100
- iApp virtual servers
 - specifying 75, 82, 88
- inventory details
 - exporting to CSV file 25
 - viewing for devices 25
- IP addresses
 - for managed devices 24

L

- license
 - activating pool license 28
 - adding pool license 28
 - manually activate a pool license 29
- licenses
 - about managing for devices 28
 - about pool licenses 28
 - about utility licenses 30
 - about volume licenses 34
 - assigning a volume license 36
 - assigning utility 32
 - for pools 30
 - revoking a volume license for managed device 36
 - revoking for managed device 30, 33

licensing

- activating a utility license automatically 31
- activating a volume license automatically 34
- activating a volume license manually 34
- activating pool license automatically 28, 64
- activating pool license manually 65
- activating utility license manually 31
- assigning a volume license to BIG-IP devices 36
- assigning utility license to BIG-IP devices 32
- for managed devices 28, 30, 34
- for pool license 28, 64–65
- for pools for BIG-IP devices 30
- manually activating pool license manually 29

licensing process

- for managed devices 64
- load balancing service instance
 - creating 73, 81, 86
- local cloud connector
 - associating with a device 93
- local cloud resources
 - using 92

M

- managed devices
 - about discovering 24
- manual activation
 - for pool license 29
- manuals, finding 18

N

- network configuration
 - and requirements for using VMware 56, 58, 62
 - for integrating with OpenStack cloud services 50
 - using Amazon web services 38
- network configurations
 - customizing for tenants 45, 72, 79, 85, 100–101
- network configurations iApps
 - customizing for tenants 45, 72, 79, 85, 100–101
- network resources
 - using for cloud services 92
- NSX devices
 - about provisioning 62, 69, 75, 83
 - connecting to 68
- NSX Edge pools
 - specifying 74, 82, 87
- NSX Edge Services Gateway
 - creating 73, 80, 86
- NSX integration
 - about 62, 69, 75, 83
- NSX runtime deployment specification
 - configuring with API calls 68
 - registering 68
- NSX service
 - about 62–63
- NSX service definition
 - provisioning 84
- NSX service instance
 - creating 73, 81, 86
- NSX VE
 - configuring 63

O

- objects
 - finding associations 18
 - searching for 19
- offering licenses
 - activating for a license 29, 32, 35
- OpenStack
 - and iApps 50
 - required configuration 51
 - using with BIG-IP VE systems 51
- OpenStack cloud services
 - and network configuration requirements 50
- OpenStack connector
 - associating with a device 52
- OpenStack devices
 - discovering 51

P

- panels
 - reordering 18–19
- pool license
 - about activating 64
 - activating automatically 28, 64
 - activating manually 29, 65
 - adding 28
 - revoking for a BIG-IP device 30, 33

- pool licenses
 - about 28
 - assigning to a BIG-IP device 30
- provisioning
 - NSX service definition 76
- provisioning process
 - for NSX device 66
 - for VXLAN network 69
 - NSX service definition 69, 72, 80, 86

R

- release notes, finding 18
- reports
 - for asset management 25
 - for utility license 33
 - for utility license billing 33
- resources
 - defined 104
 - providing access for user 97

S

- search function
 - finding specific objects 19
- service instance
 - creating for NSX 73

T

- tenant
 - adding 96
- tenant access
 - using IAM 47
- tenants
 - about creating 96
 - and users 64, 96
 - associating with a user 97
 - creating applications for 45, 72, 79, 85, 101
- tenants users
 - and tenants 96
 - and users 64, 96
- terminology 104
- terms
 - defined 104

U

- user interface
 - and searching for specific objects 19
 - customizing 18–19
 - navigating 18
- users
 - adding 96
 - and tenants 64
 - associating with a tenant 97
 - defined 64, 96
- utility license
 - activating an offering license for 29, 32, 35
 - activating automatically 31
 - activating manually 31

Index

- utility license (*continued*)
 - assigning to a BIG-IP device 32
 - submitting usage report 33
- utility license reports
 - downloading 33
- utility licenses
 - about 30
- V**
- vApps
 - and application templates 56
 - deploying 56
- vCloud Director integration
 - about 58
- vCNS, See VMware vCNS
- Virtual Private Cloud 39
 - See also VPC
 - See also Amazon Virtual Private Cloud
- virtual servers
 - specifying for iApp 75, 82, 88
- VLAN interfaces
 - discovering 70, 78
- VMware
 - about integration with iApps 56
 - and network configuration requirements 56, 58, 62

- VMware (*continued*)
 - supported version 56
- VMware cloud connector
 - associating with a device 56, 58
- VMware devices
 - discovering 56, 59, 77
- VMware network
 - for a BIG-IP VE 62–63
 - setting up for an existing VLAN 84
- VMware NSX
 - integrating with BIG-IQ Cloud 56, 58, 65, 71, 78, 84
 - integration 56
- VMware vCNS
 - integrating with BIG-IQ Cloud 56, 58
- VMware vShield
 - integrating with BIG-IQ Cloud 56, 58
- volume license
 - activating automatically 34
 - activating manually 34
 - assigning to a BIG-IP device 36
 - revoking for a BIG-IP device 36
- volume licenses
 - about 34
- VPC
 - 39
 - creating 39
- vShield, See VMware vShield