

BIG-IQ[®] Application Delivery Controller: Administration

Version 4.5



Table of Contents

Legal Notices.....	5
Acknowledgments.....	7
Chapter 1: BIG-IQ Application Delivery Controller: Overview.....	15
About BIG-IQ Application Delivery Controller.....	15
About the BIG-IQ system user interface.....	15
Filtering for associated objects.....	16
Customizing panel order.....	16
Searching for specific objects.....	16
Additional resources and documentation for BIG-IQ systems.....	16
Chapter 2: Device Discovery.....	19
About device discovery and management.....	19
Discovering devices.....	19
Discovering a large group of devices	21
About static and dynamic device groups.....	22
Creating a static group of managed devices.....	22
Creating a dynamic group of managed devices.....	23
Chapter 3: Managing Device Resources.....	25
About device resource management.....	25
Selecting specific devices to manage.....	25
Viewing properties for managed configuration objects.....	26
Overwriting undeployed changes.....	26
Refreshing managed object view.....	27
Changing device local traffic objects.....	28
Creating a new virtual server.....	29
Creating a new pool.....	31
Creating a new node.....	33
About deploying configuration changes.....	35
Chapter 4: Managing Device Permissions.....	37
About permissions management.....	37
Revising managed object permissions.....	37
Chapter 5: Glossary.....	39
BIG-IQ Cloud terminology.....	39
BIG-IQ ADC module terminology.....	40

Legal Notices

Publication Date

This document was published on January 7, 2016.

Publication Number

MAN-0577-00

Copyright

Copyright © 2015-2016, 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 Application Delivery Controller: Overview

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

About BIG-IQ Application Delivery Controller

BIG-IQ[®] Application Delivery Controller (ADC) makes it possible for you to monitor and manage the Local Traffic Manager[™] (LTM) configuration on BIG-IP[®] devices. This module helps the user:

- Create efficient work flows to view the LTM[®] configurations in a relational and dynamic user interface.
- Control access to configuration objects using fine-grained, role-based access control (RBAC). This allows administrators to delegate frequently performed operations (for example, enabling or disabling pool members) to the correct team member.
- Maintain ultimate control of the LTM configuration by providing a staging option. Delegated team members make all relevant changes, then the administrator can apply them after a quick review.

BIG-IQ ADC has two primary interfaces; Configuration and Deployment.

- Use the Configuration interface to work with the settings for the devices the BIG-IQ device manages. The Configuration interface has two interactive modes: On BIG-IQ and On BIG-IP.
 - When **BIG-IP** is selected, the settings that display for the managed devices are from the most recent sync. You cannot make changes to these settings when **BIG-IP** is selected.
 - When **On BIG-IQ** is selected, the settings that display for the managed devices still include the most recent sync settings, but also include any revisions you have made.
- Use the Deployment interface to apply configuration changes, that were made on the BIG-IQ device, to the managed devices.

About the BIG-IQ system user interface

The BIG-IQ[®] system interface is composed of panels. Each panel contains objects that correspond to a BIG-IQ feature. Depending on the number of panels and the resolution of your screen, some panels may be collapsed and show as colored bars on either side of the screen. You can cursor over the collapsed panels

to locate the one you want, and click the panel to open. To associate items from different panels, click an object, and drag and drop it onto the object with which you want to associate it.

Filtering for associated objects

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

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

Customizing panel order

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

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

Searching for specific objects

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

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

Additional resources and documentation for BIG-IQ systems

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® Device: Device Management</i>	This guide provides details about how to deploy software images, licenses, and configurations to managed BIG-IP® devices.
<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® Application Delivery Controller: Administration</i>	This guide provides details about how to centrally manage BIG-IP® Local Traffic Manager™ applications.
<i>BIG-IQ® Security: Administration</i>	This guide contains information used to centrally manage BIG-IP® firewalls, policies, rule lists (as well as other shared objects), and users.
<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.

Chapter 2

Device Discovery

- *About device discovery and management*
- *About static and dynamic device groups*

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 BIG-IQ[®] ADC aware of them through the discovery process. To discover a device, you provide BIG-IQ ADC 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 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.

***Note:** After discovering multiple devices, if there are identical configuration objects that exist on more than one device, those objects appear as unique objects. For example, if you discover ten devices, and a particular object (for example, an iRule, or a VLAN) is defined on each device, that object displays as a unique object for each device on which it is defined.*

Discovering devices

After you license and perform the initial configuration for the BIG-IQ[®] system, you can discover a BIG-IP[®] device running version 11.3.0 or later. For you to discover BIG-IP devices, they must have ports 22 and 443 open.

***Important:** Although the BIG-IQ system can discover BIG-IP devices running version 11.3 or later, the ADC module supports LTM[®] object management only for BIG-IP devices running version 11.5.1, Hot Fix 7 or later.*

Discovering BIG-IP devices is the first step to managing them.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured

on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.

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

The preferred address for discovering a BIG-IP device is its management IP address.

5. (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: `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. (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.

7. In the **User Name** and **Password** fields, type the administrator user name and password for the managed device.
8. For the **Update Framework** setting, select the **Update Framework On Discovery** 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. Also, in any system upgrade, unexpected errors might occur. Because there is not currently an automatic recovery and rollback feature, if an upgrade fails, a BIG-IP device might 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 practice is to perform a partition restore (restoring both the pre-discovery management components and any related configuration).*

9. Click the **Discover** button.

The BIG-IQ system populates the properties of the device that you added, and displays the device in the ADC panel. If you discover BIG-IP devices configured in a device service clustering, or DSC®, BIG-IQ ADC also populates the DSC Groups panel with the device's details.

Note: *Statistics for managed BIG-IP devices are refreshed every 60 seconds. Consequently, there can be a delay of up to 60 seconds before statistics are available for newly discovered devices. There can be a similar delay before the status icons for the device's managed objects are up to date. For example, you might discover a device, but the status icons for all of its managed objects could display as undefined until the next refresh.*

Discovering a large group of devices

After you license and perform the initial configuration for the BIG-IQ® system, you can discover a BIG-IP® device running version 11.3.0 or later. For you to discover BIG-IP devices, they must have ports 22 and 443 open.

Important: *Although the BIG-IQ system can discover BIG-IP devices running version 11.3 or later, the ADC module supports LTM® object management only for BIG-IP devices running version 11.5.1, Hot Fix 7 or later.*

Before you discover a large group of devices, you must save the information in a .csv file in one of the following formats:

- [address], [userName], [password], [automaticFrameworkUpdate?], [rootUser], [rootPassword], for example: 192.168.2.xxx, admin, password, true, root, password. Use this option if you want BIG-IQ Device to automatically update the framework required to manage the devices.
- [address], [userName], [password], for example: 192.168.2.xxx, admin, password .

If you have a large number of devices to discover, discovering them in a group saves you a significant amount of time, because you are not required to provide the device identification details for each individual device. Instead, you can upload a CSV file that contains the IP address, user name, and password for the devices you want to discover.

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

The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.

3. Hover over the Devices header, click the + icon when it appears, and then select **Import Devices**.
4. From the **Group Name** list select the group to which you want to add the imported devices.
5. Click the **Choose File** button and select the CSV file to which you exported the device list.
Alternatively, you can navigate to the CSV file on your computer and drag and drop it to the Import Devices screen.
6. Click the **Discover** button to complete the discovery process.
If there was a format error for the data in the CSV file, discovery fails and BIG-IQ Device returns an error message.

BIG-IQ software adds devices that it successfully discovers to the list of devices in the group that you specified.

To view or manage the configuration objects on the just added devices, you need to select each device and specify that you want to manage it.

About static and dynamic device groups

To help you manage a large number of BIG-IP® devices, you can organize them into groups. You can create two different types of device groups:

- Static group
- Dynamic group

A *static group* contains a specific set of devices. You may want to create a static group for devices hosting certain applications, in a certain geographical location, or running specific version of BIG-IP software. In contrast, a *dynamic group* is essentially a saved query against a static group. For example, if you create a static group that contained all of the managed BIG-IP devices and you wanted to view only those devices running a specific version of software, you would create a dynamic group with that parameter.

If you delete a managed BIG-IP device from the static group, that change reflects in the dynamic group when you view it.

Creating a static group of managed devices

You must license and discover devices before you can place BIG-IP® devices into a group.

To help you manage a large number of devices, you can organize them into groups. For example, you could group devices by applications, geographical location, or department.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. Hover over the Devices header, click the + icon when it appears, and then select **New Group**.

4. In the **Display Name** field, type the name you want to use to identify this group.
This name is displayed in the Devices panel. You can change this name at any time, after you save this group.
5. In the **Description** field, type a description for this group.
For example, `BIG-IP devices located in Seattle`.
You can change this name at any time, after you save this group.
6. For the **Group Type** setting, select **Static Group**.
7. From the **Parent Group** list, select the source for the group you are creating.
8. Click **Save**.

The associated managed devices now display in the Device panel, within the group you created. If you a saved filter on specific devices within this group, you can create a dynamic group.

Creating a dynamic group of managed devices

You must license and discover devices, and create a static group before you can create a dynamic group.

To filter a static group on specific parameters, you can create a dynamic group. For example, if you have a static group for all devices located in a particular city, you might want to view only those running a specific version of software.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. Hover over the Devices header, click the + icon when it appears, and then select **New Group**.
4. In the **Display Name** field, type the name you want to use to identify this group.
This name is displayed in the Devices panel. You can change this name at any time, after you save this group.
5. In the **Description** field, type a description for this group.
For example, `BIG-IP devices located in Seattle`.
You can change this name at any time, after you save this group.
6. For the **Group Type** setting, select **Dynamic Group**.
7. For the **Source Group** setting, select the static group on which you want to query for results.
8. In the **Search Filter** field, type a term on which you want to filter the group.
You can filter on a single term or, if you want to filter on more than one parameter, use the standard Open Data Protocol (OData) format.
9. Click **Save**.

This dynamic group displays in the ADC panel as a child of the associated static group.

Chapter

3

Managing Device Resources

- *About device resource management*

About device resource management

You can use BIG-IQ[®] ADC 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 device resources is the discovery process. After discovery, you can make revisions, and then deploy the configuration changes to the managed devices for easy asset management. You can make these device configuration modifications without having to log in to each device individually.

Selecting specific devices to manage

You must have Read permissions to view the configuration objects imported from managed devices, and both Read and Write permissions to manage those objects.

When devices are discovered in other modules (such as BIG-IQ[®] Cloud or BIG-IQ[®] Device), they become visible in the BIG-IQ[®] ADC module's list of managed devices. You can specify whether devices discovered from other modules also use the ADC module's local traffic configuration management capabilities. Opting out of local traffic management for a device allows you to continue Cloud and Device management functions such as license management or backup and restore, without incurring the processing overhead incumbent to local traffic object management.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then, next to the Filter field, click **Monitoring View**. The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view shows the objects and settings most recently imported from the managed BIG-IP device. The list shows only configuration objects for which you have Read or Read/Write permissions.
3. In the Devices panel, expand the group, if necessary, and hover over the device for which you wish to specify management; click the gear icon () , and then select **Properties**. The properties screen for the selected device opens.
4. If you want to change whether the device is being managed, use the **Configuration** setting and select the check box. (There is only one check box, and the text for it reflects whether the device is currently managed.)
 - If the device is currently unmanaged, the check box reads **Manage ADC Configuration**; click it to begin managing the device.

- If the device is currently managed, the check box reads **Unmanage ADC Configuration**; click it to stop managing the device.

5. Click **Save**.

The system makes the changes to management status that you requested.

Viewing properties for managed configuration objects

You must have Read permissions to view the configuration settings imported from managed devices.

Before you can monitor or manage settings for configuration objects on a device, you must be managing that device.

Using BIG-IQ[®] ADC, you can view configuration objects settings for virtual servers, pools, nodes, and iRules[®] that reside on managed BIG-IP[®] devices.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then, next to the Filter field, click **Monitoring View**. The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view shows the objects and settings most recently imported from the managed BIG-IP device. The list shows only configuration objects for which you have Read or Read/Write permissions.
3. On the panel that corresponds to the type of object you want to view, hover over the object you want to view, click the  icon, and then select **Properties** to access the configuration settings that have been imported for this object. The screen displays properties for the selected object.
4. Use the scroll bar to view the entire set of settings defined for the selected configuration.

Important: *If you are viewing settings for a virtual server, do not overlook the two areas at the bottom of the screen (Configuration and Resources) that expand to display additional settings.*

Overwriting undeployed changes

You must have Read permissions to view the configuration settings imported from managed devices.

The default behavior for the BIG-IQ[®] device in its role as manager is to exercise authority over the devices it manages. The settings of the managing BIG-IQ device prevail. That is, if there are differences between the current objects and settings on the managed BIG-IP[®] device, and the objects and settings that the managing BIG-IQ device has for that BIG-IP device, the BIG-IQ device uses the settings it already has.

In situations in which you do not want this to occur, you can overwrite the objects and settings that the BIG-IQ device recognizes for the managed device with the current objects and settings on the managed device. When you do this, settings on the BIG-IQ device (including undeployed configuration revisions) are replaced with the settings from the managed device.

Note: *Overwriting undeployed changes removes all configuration revisions you have made for the selected device. If you have made a significant number of changes on the BIG-IQ device and only want to discard a few of them, it might be better to revert them individually.*

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

The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.

3. In the Devices panel, expand the device group in which your device resides, hover over the device for which you wish to discard changes, click the gear icon (⚙️), and then select **Properties**.
The properties screen for the selected device opens.
4. At the bottom of the Properties screen, click **Overwrite Working Changes** to replace all changes currently pending on the BIG-IQ device for this managed BIG-IP device with the objects and settings as they currently exist on that managed BIG-IP device.
A confirmation popup screen opens.
5. On the confirmation screen, click **Revert Working Changes** to confirm the overwrite of your changes.

The objects and settings for the configuration objects that currently exist on the BIG-IP device overwrite the settings for those objects on the managing BIG-IQ device.

Refreshing managed object view

You must have Read permissions to view the configuration settings imported from managed devices.

Configuration object settings for virtual servers, pools, nodes, and iRules® that reside on managed devices are imported during device discovery. However, if the device administrator makes changes to these settings after device discovery, the settings seen on the BIG-IQ® device may not be completely current. You can refresh the managed object view to make sure that you have the most up to date values for the imported configuration object properties.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then, next to the Filter field, click **Monitoring View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view shows the objects and settings most recently imported from the managed BIG-IP device. The list shows only configuration objects for which you have Read or Read/Write permissions.
3. In the Devices panel, expand the group, if necessary, and hover over the device you want to view, click the gear icon (⚙️), and then select **Properties**.
The properties screen for the selected device opens.
4. At the bottom of the screen, click the **Remote > Local** button.

***Note:** The refresh performed when you click this button does nothing to any configuration changes that you might have made to the managed device, but not yet deployed. To discard those changes, you need to overwrite the undeployed changes.*

The BIG-IQ device refreshes its view of the properties that currently exist on the managed device. The properties of the configuration objects on the managed device are imported again, so that any changes that might have been made on the managed device since the last refresh are recognized by the BIG-IQ device.

The current state of the configuration objects on the managed device is updated. Properties such as statistics, device state, and status are all updated in BIG-IP® view.

***Note:** Statistics for managed BIG-IP devices are refreshed every 60 seconds. Consequently, there can be a delay of up to 60 seconds before statistics are available for newly updated devices. There can be a similar delay before the status icons for managed objects are up to date. For example, you might refresh the view*

for a device that has a healthy virtual server, but the status icon for that server could display as unhealthy until the next refresh.

Changing device local traffic objects

Before you make changes to a local traffic object on a managed device, there are two tasks to perform to ensure that you get the expected result.

- Make sure that no undeployed changes exist for the local traffic object on the managing BIG-IQ[®] device. Overwrite undeployed changes before proceeding.
- Make sure you have the most up to date information about the object on the managed BIG-IP[®] device. Refresh the managed object view to update the BIG-IQ device.

You can make revisions to the configuration of local traffic objects (virtual servers, pools, and nodes) on managed devices.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. On the panel that corresponds to the type of object you want to change, hover over the object you want to view, click the  icon, and then select **Properties** to access the configuration settings that have been imported for this object.
The properties screen for the selected object opens.
4. On the Properties screen, make changes to the configuration object you want to modify.
 - a) To enable an iRule on a virtual server, expand **Resources**, then select the iRule from the **Available** list, and use the Move button to move the iRule to the **Enabled** list.
 - b) When you are satisfied with the changes you have made, click **Save**.

The revisions you saved are made, and the Properties screen for the selected object closes.

Changes that you make are made only to the pending version. The *pending version* serves as a repository for changes you stage before deploying them to the managed device. Object settings for the pending version are not the same as the object settings on the actual BIG-IP device until they are deployed or discarded.

Important: *There is an exception to this pattern. When you view properties for a pool and click **Enable**, **Disable**, or **Force Offline**, you can choose whether you want the change to occur immediately (**Change Now**), later (**Change Later**), or not at all (**Cancel**). Changes you decide to make later become part of the pending changes for the managed object.*

To apply the pending version settings to the BIG-IP device, you need to deploy the revisions.

Creating a new virtual server

You can use the BIG-IQ® ADC interface to add a virtual server to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next sync attempt could fail.*

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. Hover over the Virtual Servers panel and click the **+** icon.
The New Virtual Server screen opens.
4. From the **Device** list, select the device on which to create the virtual server.
5. In the **Name** field, type in a name for the virtual server you are creating.
6. In the **Description** field, type in a brief description for the pool you are creating.
7. For the **Source Address**, type an IP address or network from which the virtual server will accept traffic.
For this setting to work, you must specify a value other than 0.0.0.0/0 or ::/0 (that is, any/0, any6/0). In order to maximize the utility of this setting, specify the most specific address prefixes that include your customer addresses, but exclude addresses outside of their range.
8. For the **Destination Address**, type the IP address of the destination you want to add to the Destination list.

The format for an IPv4 address is I<a>.I.I<c>.I<d>:I<port>. For example, 172.16.254.1:443.

The format for an IPv6 address is I<a>:I:I<c>:I<d>:I<e>:I<f>:I<g>:I<h>.I<port>. For example, 2001:db8:85a3:8d3:1319:8a2e:370:7348.443. If the port is well-known to be associated with a specific port number, you can specify it by name. Using the previous example, 2001:db8:85a3:8d3:1319:8a2e:370:7348.443 becomes 2001:db8:85a3:8d3:1319:8a2e:370:7348.https.
9. In the **Service Port** field, type a service port number, or select a type from the list.
When you select a type from the list, the value in the **Service Port** field changes to reflect the associated default, which you can change.
10. To configure the virtual server so that its status contributes to the associated virtual address status, select **Notify Status to Virtual Address**.
When this setting is disabled, the status of the virtual server does not contribute to the associated virtual address status. This status, in turn, affects the behavior of the system when you enable route advertisement of virtual addresses.
11. If you want the pool member and its resources to be available for load balancing, select **State**.
12. To specify configuration parameters for this virtual server, expand **Configuration** and continue with the next thirteen steps. Otherwise, skip to step 25 in this procedure.
13. From the **Source Address Translation** list, select the type of address translation pool used for implementing selective and intelligent source address translation.
 - **None:** The system does not use a source address translation pool for this virtual server.

- **SNAT:** The system uses source network address translation (NAT), as defined in the specified SNAT pool, for address translation.
- **Auto Map:** The system uses all of the self IP addresses as the translation addresses for the pool.

14. In the **Connection Limit** field, type the maximum number of concurrent connections allowed for the virtual server.

15. In the **Connection Rate Limit** field, type the maximum number of connections-per-second allowed for a pool member.

When the number of number of connections-per-second reaches the limit for a given pool member, the system redirects additional connection requests. This helps detect Denial of Service attacks, where connection requests flood a pool member. Setting the limit to 0 turns off connection limits.

16. From the **Connection Rate Limit Mode** list, select the scope of the rate limit defined for the virtual server.

- **Per Virtual Server:** Applies rate limiting to this virtual server.
- **Per Virtual Server and Source Address:** Applies Connection Rate Limit Source Mask to the source IP address of incoming connections to this virtual server, and applies the rate limit to connections sharing the same subnet. The Connection Rate Limit Source Mask specifies the number of bits in the IP address to use as a limit key.
- **Per Virtual Server and Destination Address:** Applies Connection Rate Limit Destination Mask to the destination IP address of outgoing connections from this virtual server, and applies the rate limit to connections sharing the same subnet. The Connection Rate Limit Destination Mask specifies the number of bits in the IP address to use as a limit key.
- **Per Virtual Server, Destination, and Source Address:** Applies Connection Rate Limit Source Mask and Connection Rate Limit Destination Mask to the source and destination IP address of incoming connections to this virtual server, and applies the rate limit to connections sharing the same subnet. The Connection Rate Limit Source Mask and Connection Rate Limit Destination Mask specify the number of bits in the IP addresses to use as a limit key.
- **Per Source Address (All Rate Limiting Virtual Servers):** Applies rate limiting based on the specified source address for all virtual servers that have rate limits specified.
- **Per Destination Address (All Rate Limiting Virtual Servers):** Applies rate limiting based on the specified destination address for all virtual servers that have rate limits specified.
- **Per Source and Destination Address (All Rate Limiting Virtual Servers):** Applies rate limiting based on the specified source and destination addresses for all virtual servers that have rate limits specified.

17. If you want the system to translate the virtual server address, select **Address Translation**.

This option is useful when the system is load balancing devices that have the same IP address.

18. If you want the system to translate the virtual server port, select **Port Translation**.

This option is useful when you want the virtual server to load balance connections to any service. The default is enabled.

19. From the **Source Port** list, select how you want the system to preserve the connection's source port.

- **Preserve:** Specifies that the system preserves the value configured for the source port, unless the source port from a particular SNAT is already in use, in which case the system uses a different port.
- **Preserve Strict:** Specifies that the system preserves the value configured for the source port. If the port is in use, the system does not process the connection. Restrict the use of this setting to cases that meet at least one of the following conditions:
 - The port is configured for UDP traffic.
 - The system is configured for nPath routing or is running in transparent mode (that is, there is no translation of any other Layer 3 or Layer 4 field).
 - There is a one-to-one relationship between virtual IP addresses and node addresses, or clustered multi-processing (CMP) is disabled.

- **Change:** Specifies that the system changes the source port. This setting is useful for obfuscating internal network addresses.
20. To replicate client-side traffic (that is, prior to address translation) to a member of a specified pool, select that pool from the **Clone Pool (Client)** list.
 21. To replicate server-side traffic (that is, prior to address translation) to a member of a specified pool, select that pool from the **Clone Pool (Server)** list, select the device on which to create the virtual server.
 22. Use the **Auto Last Hop** list to specify whether you want the system to send return traffic to the MAC address that transmitted the request, even if the routing table points to a different network or interface.
 23. From the **Last Hop Pool** list, select the pool the system uses to direct reply traffic to the last hop router.
 24. If you want the system to allow IPv6 hosts to communicate with IPv4 servers, select **NAT64**.
 25. To specify the virtual server score in percent, type that value in the **VS Score** field.
Global Traffic Manager™ (GTM™) uses this value to load balance traffic in a proportional manner.
 26. To specify additional resource details for this virtual server, expand **Resources** and continue with the next two steps. Otherwise, skip to the last step in this procedure.
 27. To specify which iRules® are enabled for this virtual server, use the arrow buttons to move iRules between the **Available** and **Enabled** lists.
iRules are applied in the order in which they are listed.
 28. Use the **Default Pool** list to select the pool name that you want the virtual server to use as the default pool.
A load balancing virtual server sends traffic to this pool automatically, unless an iRule directs the server to send the traffic to another pool.
 29. Click **Save**.
The system creates the new virtual server with the settings you specified.

***Note:** Statistics for managed BIG-IP devices are refreshed every 60 seconds. Consequently, there can be a delay of up to 60 seconds before statistics are available for newly deployed servers. There can be a similar delay before the server status icon display is up to date. For example, you might create a new server, but the status icon could display as undefined until the next refresh.*

Creating a new pool

You can use the BIG-IQ® ADC interface to add a pool to a managed device.

***Important:** When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. Hover over the Pools panel and click the **+** icon.
The New Pool screen opens.
4. In the **Name** field, type in a name for the pool you are creating.

5. From the **Device** list, select the device on which to create the pool.
6. In the **Description** field, type in a brief description for the pool you are creating.
7. In the **Load Balancing Method** field, specify the type of load balancing you want the pool to use. The default is **Round Robin**.
8. In the **Priority Group Activation** setting, specify how the system load balances traffic. The default is **Disabled**.
 - a) To have the system load balance traffic according to the priority number assigned to the pool member, select **Less than**.
 - b) If you use a priority number, from the **Available Member(s)** list, select the minimum number of members that must be available in one priority group before the system directs traffic to members in a lower priority group.

When a sufficient number of members become available in the higher priority group, the system again directs traffic to the higher priority group.
9. To specify advanced properties, expand **Advanced Properties** and continue with the next twelve steps. Otherwise, skip to the last step in this procedure.
10. To automatically enable or disable NATs for connections that use this pool, for the **NAT** setting, select **Allow**.
11. To automatically enable or disable SNATs for connections that use this pool, for the **SNAT** setting, select **Allow**.
12. To specify how the system should respond when the target pool member becomes unavailable, select a value from the **Action On Service Down** list.
 - **None**: Specifies that the system takes no action to manage existing connections when a pool member becomes unavailable. The system maintains existing connections, but does not send new traffic to the member.
 - **Reject**: Specifies that, if there are no pool members available, the system resets and clears the active connections from the connection table and sends a reset (RST) or Internet Control Message Protocol (ICMP) message. If there are pool members available, the system resets and clears the active connections, but sends newly arriving connections to the available pool member and does not send RST or ICMP messages.
 - **Drop**: Specifies that the system simply cleans up the connection.
 - **Reselect**: Specifies that the system manages established client connections by moving them to an alternative pool member when monitors mark the original pool member down.
13. To specify the duration during which the system sends less traffic to a newly-enabled pool member, select a value from the **Slow Ramp Time** field.

The amount of traffic is based on the ratio of how long the pool member has been available compared to the slow ramp time, in seconds. Once the pool member has been online for a time greater than the slow ramp time, the pool member receives a full proportion of the incoming traffic. Slow ramp time is particularly useful for the least connections load balancing mode.

***Important:** Setting this to a non-zero value can cause unexpected Priority Group behavior; such as load balancing to a low-priority member even with enough high-priority servers.*

14. To specify whether the system sets a Type of Service (ToS) level within a packet sent to the client, based on the targeted pool, select a value from the **IP ToS to Client** list.

Setting a ToS level affects the packet delivery reliability.

 - **Pass Through**: The system does not change the ToS level within a packet.
 - **Specify**: Provides a field in which you can specify a ToS level to apply. Valid values are from 0 to 255.

- **Mimic:** Specifies that the system sets the ToS level of outgoing packets to the same ToS level of the most-recently received incoming packet. For example, if the most-recently received packet had a ToS level of 3, the system sets the ToS level of the next outgoing packet to 3.
15. To specify whether the system sets a Type of Service (ToS) level within a packet sent to the server, based on the targeted pool, select a value from the **IP ToS to Server** list.
Setting a ToS level affects the packet delivery reliability.
- **Pass Through:** The system does not change the ToS level within a packet.
 - **Specify:** Provides a field in which you can specify a ToS level to apply. Valid values are from 0 to 255.
 - **Mimic:** Specifies that the system sets the ToS level of outgoing packets to the same ToS level of the most-recently received incoming packet. For example, if the most-recently received packet had a ToS level of 3, the system sets the ToS level of the next outgoing packet to 3.
16. To specify whether the system sets a the system sets a Quality of Service (QoS) level within a packet sent to the client, based on the targeted pool, select a value from the **Link QoS to Client** list.
Setting a QoS level affects the packet delivery priority.
- **Pass Through:** The system does not change the QoS level within a packet.
 - **Specify:** Provides a field in which you can specify a QoS level to apply. Valid values are from 0 to 7.
17. To specify whether the system sets a the system sets a Quality of Service (QoS) level within a packet sent to the server, based on the targeted pool, select a value from the **Link QoS to Server** list.
Setting a QoS level affects the packet delivery priority.
- **Pass Through:** The system does not change the QoS level within a packet.
 - **Specify:** Provides a field in which you can specify a QoS level to apply. Valid values are from 0 to 7.
18. To specify the number of times the system tries to contact a new pool member after a passive failure, select a value from the **Reselect Tries** field.
A passive failure consists of a server-connect failure or a failure to receive a data response within a user-specified interval. The default is 0, which indicates no reselects.
19. To enable TCP request queuing, select **Request Queuing**.
20. To specify the maximum number of connection requests allowed in the queue, type an entry in the **Request Queue Depth** field.
The default value of 0 permits unlimited connection requests, constrained only by available memory.
21. To specify the maximum number of milliseconds that a connection request can be queued until capacity becomes available, whereupon the connection request is removed from the queue and reset, type an entry in the **Request Queue Timeout** field.
The default value of 0 permits unlimited time in the queue.
22. Click **Save**.
The system creates the new pool with the settings you specified.

Creating a new node

You can use the BIG-IQ® ADC interface to add a node to a managed device.

Nodes are the basis for creating a load balancing pool. For any server that you want to be part of a load balancing pool, you must first create a node, that is, designate that server as a node. After designating the

server as node, you can add the node to a pool as a pool member. You can also associate a health monitor with the node, to report the status of that server.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. Hover over the Nodes panel and click the **+** icon.
The New Node screen opens.
4. From the **Device** list, select the device on which to create the node.
5. In the **Name** field, type in a name for the node you are creating.
6. For the **Address** field, select either **Address**, or **FQDN**, to specify how to identify the node you are creating.
 - If you select **Address**, type in the IP address that identifies the new node.
 - If you select **FQDN**, type in the fully qualified domain name that identifies the new node.If you select **FQDN**, the screen displays several additional settings.
7. If you chose the FQDN method for identifying this node, specify the **Address Type** for this node by selecting either **IPv4** or **IPv6**.
8. If you chose the FQDN method for identifying this node, specify the **Auto Populate** setting.
When the domain name you specify resolves to multiple IP addresses, you can enable this setting if you want read-only ephemeral nodes to be created for these addresses.
9. If you chose the FQDN method for identifying this node, specify the **Interval** for this node.
This setting specifies the number of seconds that you want the system to spend attempting to resolve a domain name.
10. If you chose the FQDN method for identifying this node, specify the **Down Interval** for this node.
This setting specifies the number of attempts you want the system to make to resolve a domain name.
11. To specify configuration parameters for this node, expand **Configuration** and continue with the next three steps. Otherwise, skip to step fifteen in this procedure.
12. For the **Ratio**, type the ratio weight you want to assign to the new node.
When you are using the Ratio load balancing method, you can assign a ratio weight to each node in a pool. LTM uses this ratio weight to determine the correct node for load balancing. Note that at least one node in the pool must have a ratio value greater than 1. Otherwise, the effect equals that of the Round Robin load balancing method.
13. For the **Connection Limit**, type the maximum number of concurrent connections allowed for this node.
14. For the **Connection Rate Limit**, type the maximum rate of new connections per second allowed for this node.
When you specify this limit, the system controls the number of allowed new connections per second, thus providing a manageable increase in connections without compromising availability. The default value of 0 specifies that there is no limit on the number of connections allowed per second.
15. Click **Save**.
The system creates the new node with the settings you specified.

About deploying configuration changes

Using BIG-IQ® ADC to manage the devices in your network means that you can deploy configuration changes without having to log in to each individual BIG-IP® device. You can review deployment changes before you make them, and then either make the changes, or revert them.

When you deploy changes to a managed device, before the BIG-IQ device applies the configuration changes, it first does a fresh import from the managed device to ensure there are no conflicts. If there are conflicts, the default behavior is to discard any changes made on the managed device before deploying the configuration changes. You can work around this by overwriting undeployed changes. Overwriting undeployed changes performs a fresh import from the managed BIG-IP device and uses those objects and settings to overwrite any revisions performed on the managing BIG-IQ device.

Reviewing and deploying configuration settings

Before you deploy configuration changes, be aware of the following prerequisites:

- You must have a role of Administrator to deploy configuration changes.
- Before you deploy changes to a managed device, make sure that changes have not been made to that device while you were assembling your configuration changes. Deploying changes to a managed device overwrites the objects and settings on the managed device with the settings specified on the BIG-IQ® device. To make sure you are not overwriting settings that you didn't know about, refresh the managed object view before deploying configuration changes.

You must create a deployment job and submit that job before changes to configuration objects you have made are applied to the managed device.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Deployment**.
The list of active deployment jobs is displayed. Jobs are categorized as Pending, Error, or Completed. These are deployments that are already in process. To get your configuration changes applied to the appropriate device, you need to create a new job.
3. In the Deployments panel, click the (+) icon, and then select **Deploy Configuration Changes**.
The New Deployment screen opens.
4. In the **Name** field, type in a name for the deployment task you are creating.
5. In the **Description** field, type in a brief description for the deployment task you are creating.
6. From the list of configuration changes pending deployment, select the device for which you wish to deploy changes.
7. To review the changes before deploying them, select **Review Pending Changes** (to deploy without reviewing, skip this step).
 - a) In the Modified area of the Show Differences popup screen, select each configuration object and scroll through the revisions.

Important: As a prerequisite to this task, make sure that you know the most current configuration settings on the managed device. If you did not perform that refresh, the configuration settings you are comparing your revisions with will be out of sync with any changes made to the BIG-IP device since the last refresh.

Note: If the refresh and review reveals minor changes that have been made on the managed device, and you do not want to lose those changes, consider adding those configuration changes to the property settings on the BIG-IQ device before you deploy the changes.

- b) When you finish reviewing the pending changes, click **Cancel** on the popup screen.
8. To start the task of deploying changes to the managed device, click **Deploy**. The BIG-IQ device starts processing the deployment job. When the job completes successfully, configuration settings on the managed device are overwritten with the settings from the managing BIG-IQ device.

***Note:** To discard the just reviewed changes, overwrite the undeployed changes. The configuration settings currently on the managed device are freshly imported and overwrite the settings on the managing BIG-IQ device. For details, refer to [Overwriting undeployed changes](#).*

When you deploy a configuration job, details display in the Deployment panel's Pending list while the deployment is being processed. These details display until the job either fails or succeeds.

- If the deployment fails, details display in the Deployment panel's Error list.
- If the deployment is successful, details display in the Deployment panel's Completed list.

***Important:** The Completed deployments and Error lists maintain a 7-day history of deployment changes. After a week, these deployment change records are deleted.*

Chapter 4

Managing Device Permissions

- *About permissions management*

About permissions management

The ability to manage resources located on BIG-IP® devices using BIG-IQ® ADC is controlled by the permissions settings associated with your user role. Users with the role of Administrator can set permissions for any role.

Permissions for managing objects follow a fine-grained, role-based access control (RBAC) model. This means that you can grant read, write, create, and delete permissions for a device, a virtual server, a pool, or a node. So for example, a user might be given the ability to make revisions to the settings for a virtual server, but the ability to deploy those changes to the managed device is reserved for the Administrator. Or, you can grant authorization to make changes to one type of managed object (Pools, for instance), but reserve the authorization for other object types. Finally, you might choose to grant authorization to view or make changes on one object (for example, Pool 1), but reserve the authorization for other objects at that same level (for example, Pools 2 - 20).

Revising managed object permissions

You must have Read/Write permissions to make revisions to a configuration object. If you only have Read permissions, you can still view the configuration settings imported from managed devices.

You can revise the permissions for any configuration object (virtual servers, pools, and nodes) based on the role assigned to a user's login credentials.

1. Log in to BIG-IQ ADC with the administrator user name and password.
2. At the top of the screen, click **Configuration**, and then click **Editing View**.
The screen displays the list of devices that the BIG-IQ device is currently managing, along with the configuration objects on those devices. This view displays the objects and settings currently configured on the managing BIG-IQ device. Only configuration objects for which you have Read or Read/Write permissions are displayed.
3. On the panel that corresponds to the type of object you want to change, hover over the object you want to view, click the  icon, and then select **Properties** to access the configuration settings that have been imported for this object.
The properties screen for the selected object opens.
4. Click the  icon, and then select **Properties** to access the configuration settings that have been imported for this object.
The properties for the selected object are displayed.

Managing Device Permissions

5. Click **Permissions** to access the permissions settings that have been imported for this object.
6. In the **Role** field, type the name of the role to which you want to assign permissions, and then click **Read** or **Read/Write** as appropriate.
7. To grant permissions to another role, click the add (+) icon. To remove a role to which you have granted permissions, click the remove (x) icon.
8. When you are satisfied with the changes you have made, click **Save**.
The permissions changes are made, and the screen for the selected object closes.

Chapter 5

Glossary

- *BIG-IQ Cloud terminology*
- *BIG-IQ ADC module 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.

Term	Definition
<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.

BIG-IQ ADC module terminology

Before you manage device resources, it is important that you understand some common terms as they are defined within the context of the BIG-IQ[®] ADC module.

Term	Definition
<i>deployment</i>	In the ADC module, you can make configuration changes to a managed device, but the changes are not actually applied to the device until you deploy them.
<i>BIG-IP</i>	In the ADC module, there are two primary modes for interacting with the managed device. In this read-only mode, settings that display are from the most recent refresh of the managed device.
<i>BIG-IQ</i>	In the ADC module, there are two primary modes for interacting with the managed device. In this mode, you can edit any setting for which you have permissions; settings that display include any modifications you have made to the managed object configuration.
<i>refresh configuration</i>	In the ADC module, you can refresh the view of the configuration objects for a managed object and get an update of the device statistics.
<i>overwrite configuration</i>	In the ADC module, you can discard any undeployed changes you have made to the configuration objects for a managed object.

Index

A

- ADC
 - about 15
- ADC Configuration function
 - turning off 25
- application templates
 - defined 39–40

B

- BIG-IQ Cloud
 - defined 39–40
 - finding documentation for 16
- BIG-IQ Device
 - about ADC 15
 - finding documentation for 16
- BIG-IQ Security
 - finding documentation for 16
- BIG-IQ system
 - reordering panels 16

C

- cloud administrator
 - defined 39–40
- cloud bursting
 - defined 39–40
- cloud connector
 - defined 39–40
- configuration changes
 - about deploying 35
 - deploying to a device 35
 - discarding 26
 - reviewing 35
- configuration deployment 35
- configuration objects
 - viewing properties 26
- CSV file
 - uploading for bulk device discovery 21

D

- deployment
 - defined 40
 - discarding configuration changes 26
 - of configuration changes 35
- device configuration settings
 - refreshing the view 27
- device discovery
 - using a CSV file for bulk discovery 21
- device groups
 - about dynamic 22
 - about static 22
- device inventory
 - about 19

- device management
 - about 19
- device properties
 - refreshing the view 27
- device resource management
 - about 25, 37
- devices
 - about adding 19
 - about discovering 19
 - discovering 19
- discovery
 - using a CSV file for bulk device discovery 21
- documentation, finding 16
- dynamic device groups
 - about 22
- dynamic group
 - creating 23

F

- filtering process
 - finding associated objects 16

G

- glossary 39–40
- groups
 - about dynamic device groups 22
 - about static device groups 22
 - creating dynamic 23
 - creating static 22
- guides, finding 16

I

- IP addresses
 - for managed devices 19

M

- managed device
 - changing objects 28
 - changing to unmanaged 25
 - identifying 25
 - making revision 25
 - revising 37
- managed device configuration settings
 - viewing 27
- managed device properties
 - refreshing the view 27
 - viewing 27
- managed devices
 - about discovering 19
 - revising 26
 - viewing properties 26
- managed object
 - revising permissions 37

Index

manuals, finding 16

N

nodes
creating 33

O

objects
finding associations 16
searching for 16
on BIG-IP
defined 40
on BIG-IQ
defined 40
overwrite configuration
defined 40

P

panels
reordering 15–16
permissions
about managing 37
pools
creating 31

R

refresh configuration
defined 40

release notes, finding 16

resources
defined 39–40

S

search function
finding specific objects 16
static device groups
about 22
static group
creating 22

T

terminology 39–40
terms
defined 39–40

U

undeployed changes
overwriting 26
user interface
and searching for specific objects 16
customizing 15–16
navigating 15

V

virtual servers
creating 29