

BIG-IP[®] Policy Enforcement Manager[™]: Implementations

Version 11.5



Table of Contents

Legal Notices.....	7
Acknowledgments.....	9
Chapter 1: Overview.....	13
What is Policy Enforcement?.....	14
About enforcement policies.....	15
About enforcement policy rules.....	15
About subscriber provisioning through PCRF.....	16
Best practices for creating enforcement policies.....	16
About sizing considerations.....	17
Chapter 2: Setting Up Application Visibility.....	19
Overview: Setting up application visibility.....	20
What is application visibility?.....	20
Determining and adjusting traffic classifications.....	20
Creating a listener	21
Examining application visibility statistics.....	22
Chapter 3: Configuring Intelligent Traffic Steering.....	25
Overview: Configuring intelligent traffic steering.....	26
What is traffic steering?.....	26
Creating a pool	26
Creating forwarding endpoints.....	27
Creating an enforcement policy.....	28
Creating custom action policies.....	28
Adding rules to an enforcement policy.....	29
Creating a rule using classification criteria.....	30
Creating a rule using URL categorization.....	31
Modifying iRule event for URL categories.....	32
Creating a rule using flow conditions.....	32
Creating a rule for forwarding traffic.....	33
Creating a rule for QoS.....	34
Creating a listener	36
Chapter 4: Configuring Quota Management using Rating Groups.....	37
Overview: Configuring quota management	38
About Gy support and rating groups.....	38
Creating a listener for quota management.....	38
Creating rating groups.....	40

Chapter 5: Reporting Usage Data to an External Analytics Server.....	41
Overview: Reporting usage data to an external analytics server.....	42
Creating a publisher	42
Creating a rule for high-speed logging.....	42
Session-based reporting format.....	43
Flow-based reporting format.....	45
Chapter 6: Usage Monitoring Over a Gx Interface.....	47
Overview: Usage monitoring over a Gx interface	48
Creating a listener for subscriber discovery and policy provisioning.....	48
Creating a rule for usage monitoring.....	49
Chapter 7: Configuring Global Application Policies with Bandwidth Control.....	51
Overview: Global Application Policies with Bandwidth Control.....	52
Creating VLANs.....	52
Creating a static bandwidth control policy.....	53
Creating an enforcement policy.....	53
Creating a rule for bandwidth control.....	54
Creating a listener: example	55
Chapter 8: Enforcing Bandwidth Control Provisioned on PCRF.....	57
Overview: Enforcing bandwidth control provisioned on PCRF.....	58
Creating a dynamic bandwidth control policy for PCRF.....	58
Creating a listener for subscriber discovery and policy provisioning.....	59
Implementation result.....	60
Chapter 9: Configuring Tiered Services with Bandwidth Control.....	61
Overview: Configuring tiered services with bandwidth control.....	62
Creating dynamic bandwidth control policies for tiered services.....	62
Creating enforcement policies for three tiers.....	63
Creating the rules for tiered bandwidth control.....	63
Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF.....	65
Implementation result.....	66
Chapter 10: Configuring Service Chains.....	67
Overview: Configuring service chains.....	68
About services profiles	68
Creating a ICAP profile for policy enforcement.....	69
Creating a Request Adapt profile.....	69
Creating a Response Adapt profile.....	70
Creating an internal virtual server for ICAP server.....	70

Creating a pool	71
Creating endpoints for service chains.....	71
Creating dynamic service chains.....	72
Creating an enforcement policy.....	73
Configuring steering action policy.....	73
Adding rules to an enforcement policy.....	74
Creating a rule for forwarding traffic.....	75
Creating a listener	76
Chapter 11: Provisioning Dynamic Subscribers.....	79
Overview: Provisioning dynamic subscribers	80
Provisioning dynamic subscribers.....	80
Chapter 12: Provisioning Static Subscribers.....	83
Overview: Provisioning static subscribers	84
Provisioning single subscribers.....	84
Subscriber CSV file format.....	85
Provisioning a file of static subscribers.....	85
Chapter 13: Creating Reports using PEM.....	87
Overview: Creating reports using PEM.....	88
Creating format scripts for reports.....	88
Chapter 14: Updating Signatures for Application Recognition.....	91
Overview: Updating classification signatures.....	92
Importing signatures manually.....	92
Scheduling automatic signature updates.....	92
Chapter 15: Creating Custom Classifications.....	95
Overview: Creating custom classifications.....	96
Determining and adjusting traffic classifications.....	96
Creating custom classification categories.....	96
Creating custom classification applications.....	97
Using iRules with classification categories and applications.....	98
Modifying iRule event for URL categories.....	98
Classification iRule commands.....	99
Chapter 16: Configuring PEM with Local Traffic Policies.....	101
Overview: Creating local traffic policy rules for PEM.....	102
Modifying local traffic policy rules for PEM.....	102
Creating local traffic policy rules for PEM.....	103
Creating a virtual server for SSL traffic policy enforcement.....	104

- Chapter 17: Configuring Policy and RADIUS Updates.....107**
 - Overview: Configuring policy and RADIUS updates108
 - Configuring PEM options.....108

- Chapter 18: Enforcing Policy and Classification on IP Protocols.....109**
 - About enforcing policy and classification on IP protocols.....110
 - Creating Any IP profiles for PEM.....110
 - Updating Any IP profile110
 - IPOther filter for current PEM actions.....111

- Chapter 19: Troubleshooting.....113**
 - PEM troubleshooting114
 - Subscriber and policies active sessions.....115
 - Active sessions statistics.....116
 - Configuring subscriber activity log.....116

Legal Notices

Publication Date

This document was published on January 27, 2014.

Publication Number

MAN-0404-04

Copyright

Copyright © 2013-2014, F5 Networks, Inc. All rights reserved.

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

Trademarks

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

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

Patents

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

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

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 software under license from Qosmos (www.qosmos.com).

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

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.

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

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

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

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

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

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

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

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

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

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

Chapter 1

Overview

- *What is Policy Enforcement?*
- *About enforcement policies*
- *About enforcement policy rules*
- *About subscriber provisioning through PCRF*
- *Best practices for creating enforcement policies*
- *About sizing considerations*

What is Policy Enforcement?

BIG-IP® Policy Enforcement Manager™ (PEM) facilitates mobile service providers control subscriber traffic. The system can analyze application traffic and subscriber behavior, and then enforce traffic policing rules that you define. For example, you could have the system drop all web traffic coming from certain IP addresses. You can perform QoS actions on traffic you want to be treated as high priority using DSCP marking, link QoS, or bandwidth control. You could redirect HTTP traffic destined for a particular IP address, and send it to a specific URL. Or, you could send all video traffic from certain subscribers to servers for optimization.

The BIG-IP system is inserted between the subscribers and the network they are trying to access. The system intercepts the traffic that subscribers are trying to send. The goal of the Policy Enforcement Manager is to apply an enforcement policy to a subscriber. To determine what kind of policy to apply to the subscriber, PEM™ needs to understand the subscriber identity.

The system can obtain subscriber identity by looking at RADIUS traffic (if present), or by analyzing subscriber data traffic. RADIUS provides much more information about the subscriber. Although analyzing subscriber traffic is more limited, it does provide the subscriber IP address. The system must have the subscriber IP address in order for PEM to do policy enforcement.

Here is a typical illustration of how policy enforcement works. Traffic from a mobile service provider goes to the BIG-IP system on its way to a network. In order to regulate subscribers, PEM needs to determine the policy to apply. For that reason, PEM collects subscriber identity by intercepting RADIUS traffic and examining (snooping) the RADIUS Authentication and Accounting packets for details about the subscriber. PEM communicates with an external policy server, in this case, PCRF, for dynamic subscriber provisioning. Using the RADIUS information (or the IP address if no RADIUS is present), PEM queries the PCRF for the policy configuration and provisions subscribers dynamically.

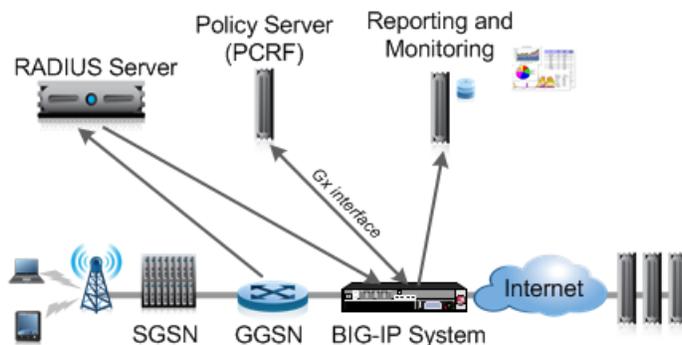


Figure 1: Diagram of policy enforcement overview

Alternatively, you can provision subscribers manually. These subscribers are called *static subscribers*. You use PEM to add static subscribers one at a time, or to import a list of subscribers. Provisioning static subscribers might require the ability to snoop RADIUS traffic but does not require a PCRF connection.

When adding static subscribers on the BIG-IP system, you provide the subscriber ID, subscriber ID type, and one or more policies to apply. You can also specify the IP address, but if it is dynamically assigned, you cannot include it. In this case, you need interception of RADIUS traffic in order to map the subscriber to the IP address. When the subscriber enters the network, the IP address from RADIUS is combined with the information already on PEM. If the static subscriber includes the IP address, no RADIUS interception is required.

About enforcement policies

An *enforcement policy* is a set of rules that determines what to do with specified types of traffic. You can configure policies on the BIG-IP® system using Policy Enforcement Manager™ (PEM), or receive policies from a PCRF.

For a policy to take effect, it needs to be assigned, or provisioned, to a subscriber session (a subscriber and single IP address mapping). A subscriber session is the period of time from when a subscriber logs into the network (authenticated) and logs out, or when the session is terminated by other means. The session is identified by subscriber IP address.

PEM™ supports the following methods for provisioning subscriber policies:

- Subscriber policy provisioning using PCRF
- Subscriber policy provisioning using a subscriber database of static subscribers (up to 100K subscribers)
- Subscriber policy provisioning for unknown subscribers (subscribers that do not currently have policies assigned to them either dynamically or statically)
- Global policy provisioning
- Custom policy provisioning using iRules®

You can use more than one of the subscriber policy provisioning methods. For example, PEM provisions an unknown subscriber policy for a subscriber session, while awaiting a response from PCRF. Or, a global policy might be applied concurrently to other subscriber policies.

As with other BIG-IP modules, like Local Traffic Manager™, you enable PEM functionality by attaching the corresponding profile to one or more virtual servers. To simplify configuration, PEM provides a listener entity that creates the required virtual servers, enables classification, and attaches the policy enforcement profile. When you create a listener, you also define which policies to apply globally or to unknown subscribers.

Advanced users can directly create virtual servers, then configure and attach the Policy Enforcement profile. We recommend that you begin configuring PEM by using listeners instead of using the advanced method. You can get familiar with PEM configuration by examining the virtual servers, settings, and profiles that the listener creates.

About enforcement policy rules

An enforcement policy is made up of a set of rules. In the policy, rules define what to do when the system receives a particular type of traffic. There are many ways you can set up a rule so that you can handle the traffic exactly as you need to. Each rule includes a condition and an action.

A rule defines conditions that the traffic must meet (or not meet) for the rule to apply. The conditions fall into the following criteria:

- *Classification criteria*, such as applications or categories of applications that the system detects. For example, a rule can apply to all webmail traffic or to a specific webmail application.
- *Flow information*, such as traffic associated with specific source and destination IP addresses or ports, or incoming DSCP marking. For example, a rule can apply to all traffic directed to a specific destination port.
- *URL information*, such as URL categories that the system detects. For example, a rule can apply to a specific webpage.
- *Custom criteria*, which are other conditions that you develop using iRules®

If the traffic meets the criteria in the rule, the rule can specify actions to take, such as:

- Dropping traffic
- Forwarding traffic to a specific endpoint or series of endpoints for value-added services
- Redirecting HTTP traffic to a URL
- Generating reporting data for further processing by external analytic systems
- Reporting details about the traffic to the PCRF so it can track mobile usage.
- Setting DSCP bits in the IP header of the traffic by marking all or marking upon the traffic exceeding a threshold
- Setting Layer 2 Quality of Service (QoS) levels for the traffic
- Enforcing rate control using a bandwidth control policy

Because rules provide so much flexibility, you need to plan what you want to do, and consider your options before you add the rules. One option is to simply classify traffic and review reports of the types of traffic your system is receiving to get more information on which to base the rules. This could be the first step when developing enforcement policies using PEM.

About subscriber provisioning through PCRF

When you are provisioning subscriber policies through PCRF, the policies are communicated using Gx protocol in the form of Policy and Charging Control (PCC) rules. A PCC rule can contain the complete rule definition, or it might refer to a predefined or dynamic policy rule, as defined by the Gx protocol specification, Release 9.4. See the 3GPP TS 29.212 specification for details. When the complete rule definition is sent, it is a dynamic PCC rule; when the rule is referenced by name, it is called a predefined rule.

A *predefined PCC rule* on PCRF maps to an enforcement policy in PEM™. For example, a predefined PCC rule with the name `premium-video` on the PCRF applies to video traffic for premium subscribers. In PEM, you can create a policy also called `premium-video` with policy rules that define the enforcement action. The classification criteria for the traffic is video, and the action could be to enforce QoS for the video traffic (for example, specifying a higher bitrate).

A *dynamic PCC rule* is dynamically provisioned by the PCRF over the Gx interface. In this case, the PCC rule contains the rule definition. Therefore, in this case, you do not need to create policies on the BIG-IP® system, since the policy is totally defined on the PCRF.

Best practices for creating enforcement policies

Follow these general recommendations when creating enforcement policies:

- When creating enforcement policies you plan to apply globally or to unknown subscribers, include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.
- Be cautious when developing enforcement policies to be applied globally. The policies affect all the subscribers and are typically applied to subscriber policies in parallel.
- When you remove or disable an enforcement policy, first be sure that it is not currently assigned to any subscribers. At least one policy must be assigned to a subscriber at all times.
- Assign the subscriber IP address when creating static subscribers that include a global or unknown subscriber policy, to ensure that the subscriber gets the entitled service faster.
- If you want to use different types of steering, create separate policies and rules. For example, consider creating a policy that steers traffic from a source VLAN to an endpoint, and another policy to steer VLAN traffic to a service chain.

- Create an empty pool and use it in a forwarding endpoint if you want to route traffic or resolve policy priority conflicts between routing and steering.

These are best practices when writing policy rules:

- Be careful when you mix both L4 and L7 classification criteria in one rule; in most cases, L4 criteria takes precedence. Keep it simple: one rule, one type of criteria.
- Specify different precedence values for the rules that might conflict, to make clear in what order the rules will apply.
- Do not mix different types of policy actions in the same rule; create separate rules for forwarding, reporting, Quality of Service (QoS) actions and finally, for which policy action is implemented.
- A policy (or a rule) should not direct traffic to both a forwarding endpoint and to a service chain. If both are specified, the service chain always takes precedence and is performed first, then traffic is forwarded to the endpoint.
- Dedicate certain bandwidth controllers for use only in PEM™ QoS actions, and do not use them outside PEM.

There are best practices to consider when setting up reporting in enforcement policies:

- Choosing more frequent intervals for generating periodic reporting records (particularly session-based) can greatly increase the amount of reporting data, and could potentially overload the analytics system.
- Flow-based records are generated several times during the flow life and can significantly impact the amount of reporting data sent.

Here are best practices to consider when setting up iRule action:

- If multiple PEM iRules® match a flow, all the iRules are processed. The priority order is as follows:
 - PEM policy priority, and it takes in to consideration if the policy is a global high precedence policy, subscriber policy or low precedence policy.
 - PEM iRule event priority and the default event priority is 500. The event priority can be changed by specifying the priority within the iRule event.
 - Rule precedence.

About sizing considerations

Currently the maximum number of applications or category IDs that PEM™ can store, or report usage statistics for, is limited to 15 per subscriber. This in turn influences the rules, since the traffic statistics for each application or category ID is part of the rule's classification criteria.

When this limitation is exceeded for a given subscriber, an error message is logged into TMM log file. In addition, if the affected rule is installed by PCRF (over Gx connection), a session provisioning failure report is sent back to PCRF. The application or category IDs limitation should be taken into account when designing the subscriber and global policies for a particular PEM deployment.

Note: *If you require granular reporting for large amount of traffic, your performance might be impacted.*

Real performance depends on various factors such as:

- Turning on flow reporting (both performance and memory impacted)
- Frequency of sending session based reporting
- Complexity of classification
- The number of concurrent flows per subscriber
- Over subscription

Chapter 2

Setting Up Application Visibility

- *Overview: Setting up application visibility* |

Overview: Setting up application visibility

This implementation describes how to set up the Policy Enforcement Manager™ (PEM) to analyze application traffic on the network, and provide statistics for application visibility. For example, you can view statistics to see what applications are being used. By monitoring your traffic, you can later create enforcement policies that are tailored for your needs.

Task summary

Determining and adjusting traffic classifications

Creating a listener

Examining application visibility statistics

What is application visibility?

Policy Enforcement Manager™ (PEM™) gives the BIG-IP® system the ability to classify both encrypted and unencrypted traffic into categories for application visibility. You can display statistics about the network traffic in graphical charts, and view classification information by application, category, protocol, virtual server, country, type of device, and so on. In-depth information and application awareness provides visibility into your network infrastructure so you can identify and monitor different types of traffic and resolve performance issues.

Application visibility is particularly useful for service providers. If your organization is using RADIUS protocol for authentication, authorization, and accounting, PEM can intercept accounting messages to retrieve additional information, for example, about mobile devices, subscribers, towers, service plans, and manufacturers.

Charts shown on the **Statistics > Classification** screens display the application visibility data. The classification overview is customizable so you can display the charts or tables that you want. The overview shows top statistics for the categories of which you are most interested.

Determining and adjusting traffic classifications

The BIG-IP® system classifies many categories of traffic and specific applications within those categories. You can determine which categories and applications of traffic the system can classify, and find out information about them such as their application or category ID.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. To view the applications in each category, click the + icon next to the category.
3. To view or edit the properties of the application or category, click the name to open its properties screen.

Tip: Here you can view the application or category ID number.

4. Adjust the properties of the application or category, if necessary.
 - In the **Description** field, you can add text to describe the application or category.
 - Set **State** to **Enabled** to use this classification, or to **Disabled** not to use it.
 - For categories only, set **iRule Event** to **Enabled** if you want the system to trigger an iRule event when it recognizes traffic in this category, or set to **Disabled** if you do not.

- In the **Category** or **Application List** field, you can change which category an application is in, or which applications are in the category.
5. Click **Update** to save any changes.

Creating a listener

If you want to steer specific traffic, or otherwise regulate certain types of traffic, you need to have developed enforcement policies. If using a Gx interface to a PCRF, you need to create a listener that connects to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. Subscriber provisioning using RADIUS is enabled by default. If your system is using RADIUS for snooping subscriber identity, you need to specify VLANs and tunnels. If you are not using RADIUS, you need to disable it.
 - For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
 - If you do not want to use RADIUS, from the **Subscriber Identity Collection** list, select **Disabled**.
7. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

8. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification

and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP® system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Examining application visibility statistics

Before you can look at the application visibility statistics, you must have Adobe® Flash® Player installed on the computer where you plan to view them.

You can review charts that provide application visibility for traffic on your network.

1. On the Main tab, click **Statistics > Classification**.
The Classification Overview screen opens where you can view a summary of the top classification statistics.
2. Review the statistics provided. To quickly change the format of the information, click the icon to the left of the time period.
You can display information in a table, line chart, pie chart, or bar chart.
3. Click the time period (**Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**), to change the interval used for displaying content.
4. To permanently change the format or content of any of the charts, click the cog on the chart, select **Settings**, and adjust the fields in the form.
5. To display additional charts or tables, click the **Add Widget** link and complete the form.
The chart you create becomes a permanent part of the Classification Overview screen.
6. On the menu bar, click **Statistics**.
The Classification Statistics screen opens and the charts display detailed classification statistics by application.
7. Adjust the statistics content in any of the following ways:
 - Use the **View By** setting or Advanced Filters to change the type of classification data shown.
 - Use the **Time Period** setting to change the interval for which statistics are shown.
 - Use the **Measurement to display** setting to change the display of statistics to bits, bytes, or sessions.
 - Use the **Display method** setting to show statistics in different formats.
 - Use the **Expand Advanced Filters** setting to fine-tune even further which types of reports to display.
8. Get detailed information in any of the following ways:
 - Point on the charts to display the details.
 - Review the Details table to see the statistics.
 - In the Details table, click the name of one of the items (application, category, protocol, and so on) to see classification details about that specific item.
 - Use the **Display method** setting to show statistics in different formats.
9. To generate and export a PDF or CSV file of a report to save or email, click **Export**, select the settings, and fill in the appropriate fields.

Note: You must have an SMTP email server configured to use the email option. On the Main tab, click **System > Configuration > Device > SMTP**.

You can use the classification statistics to determine, for example, the types of applications and the specific applications that clients are using. By drilling down into that information, you can find out specifically which applications are being used by a particular IP address.

Chapter

3

Configuring Intelligent Traffic Steering

- *Overview: Configuring intelligent traffic steering*
-

Overview: Configuring intelligent traffic steering

You can use the Policy Enforcement Manager™ to set up the BIG-IP® system to classify and intelligently steer traffic on the network. The system automatically sets up virtual servers for TCP and UDP traffic so that the BIG-IP system can classify the traffic and direct it to one or more steering endpoints based on traffic characteristics.

Task Summary

Creating a pool

Creating forwarding endpoints

Creating an enforcement policy

Creating custom action policies

Adding rules to an enforcement policy

Creating a rule using classification criteria

Creating a rule using URL categorization

Modifying iRule event for URL categories

Creating a rule using flow conditions

Creating a rule for forwarding traffic

Creating a rule for QoS

Creating a listener

What is traffic steering?

Policy Enforcement Manager™ provides the ability to intelligently steer traffic based on policy decision made using classification criteria, URL category, flow information, or custom criteria (iRule events). Steering, also called *traffic forwarding*, can help you police, control and optimize traffic. You can forward a particular type of traffic to a pool of one or more servers designed to handle that type of traffic, or to a location closer to clients requesting a service. For example, you can send HTTP video traffic to a pool of video delivery optimization servers. You can have one policy option to classify each transaction which allows transaction aware steering. This ability to classify traffic for every transaction is called *transactional policy enforcement*. The classification per transaction is for HTTP traffic only.

You set up steering by creating an enforcement policy that defines the traffic that you want to send to a particular location or endpoint. Rules in the enforcement policy specify conditions that the traffic must match, and actions for what to do with that traffic. One of the actions you can take is to forward the traffic to a particular endpoint, called a *forwarding endpoint*.

You can create listeners to set up virtual servers and associate the enforcement policies with the traffic that is sent to them. The system also creates a Policy Enforcement profile that specifies the enforcement policy that the system uses, among other uses, for traffic steering.

Creating a pool

You can create a pool of servers that you can group together to receive and process traffic.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.

The New Pool screen opens.

3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) Type an IP address in the **Address** field.
 - b) Type a port number in the **Service Port** field, or select a service name from the list.
 - c) To specify a priority group, type a priority number in the **Priority Group Activation** field.
 - d) Click **Add**.
5. Click **Finished**.

The new pool appears in the Pools list.

Creating forwarding endpoints

Before you can create an endpoint, you need to create a pool that specifies where you want to direct the classified traffic.

To set up traffic steering, you need to create a forwarding endpoint, which specifies where to send the traffic. If you are configuring w-steering or service chains, you need to create multiple endpoints.

1. On the Main tab, click **Policy Enforcement > Forwarding > Endpoints**.
The Endpoints screen opens.
2. Click **Create**.
The New Endpoint screen opens.
3. In the **Name** field, type a name for the endpoint.
4. From the **Pool** list, select the pool to which you want to steer a particular type of traffic, for example, in a policy rule.
5. If you want to translate the destination address of the virtual server to that of the pool, from the **Address Translation** list, select **Enabled**. Otherwise, leave this setting disabled.
6. If you want to translate the original destination port to another port, from the **Port Translation** list, select **Enabled**. Otherwise, leave this setting disabled.
7. From the **Source Port** list, select the appropriate option for the source port of the connection.

Option	Description
Preserve	Maintains the value configured for the source port, unless the source port from a particular SNAT is already in use.
Preserve Strict	Maintains the value configured for the source port. If the port is in use, the system does not process the connection. Use this setting only when (1) the port is configured for UDP traffic; (2) the system is configured for nPath routing or running in transparent mode; or (3) a one-to-one relationship exists between virtual IP addresses and node addresses, or clustered multi-processing (CMP) is disabled.
Change	Specifies that the system changes the source port.

8. To specify a SNAT pool for address translation, from the **SNAT Pool** list, select the name of an existing SNAT pool.
The steering endpoint uses the SNAT pool to implement selective and intelligent SNATs.
9. If you have multiple pool members and want specific traffic to go to the same pool member every time, from the **Persistence** list, select the appropriate IP address type:

Option	Description
Source Address	Map the source IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.
Destination Address	Map the destination IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.

If you do not need to maintain persistence, leave **Persistence** set to **Disabled**, the default value.

10. Click Finished.

You can direct traffic to the endpoint you created in the policy rules of an enforcement policy.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.

***Tip:** When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.*

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

***Important:** The system performance is significantly affected, depending on complexity of the classification and the type of policy action.*

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Creating custom action policies

In an enforcement policy, custom action can be defined by a Policy Enforcement Manager™ (PEM™) iRule.

1. On the Main tab, click **Policy Enforcement > Policies > iRules**.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a name for the new iRule.
4. In the **Description** field, type a description of the new iRule.
5. In the **Definition** field, specify the TCL syntax that defines a custom iRule action, which can be later attached to a policy enforcement rule.

```
when PEM_POLICY { if {[PEM::policy initial]}
  { /* Commands to run during the first time the policy is evaluated. */
  } else
  { /* Commands to run during policy re-evaluation. */ }
  /* Commands to run during policy eval and re-eval time. */ }
```

There can be two iRule events:

- PEM_POLICY is triggered when a policy evaluation occurs.
- RULE_INIT runs the first time the iRule is loaded or has changed.

The two new PEM iRule commands are PEM::policy initial and PEM::policy name. You can select the **Wrap Text** check box to wrap the definition text, and select the **Extend Text Area** check box to increase the field space of format scripts.

6. Click Finished.

The Policy Enforcement Manager creates a new iRule, and displays the iRule list.

7. To attach a custom action to a specific iRule, follow the steps:

- a) Click **Policy Enforcement > Policies**.
- b) Select a policy name.
- c) Click a policy rule.
- d) From the **Custom Action** list, select a iRule created.

8. Click Update.

You have created a custom action in a policy, using iRules.

Note: The iRule actions are executed at the end of all the other policy actions.

Adding rules to an enforcement policy

Before you can add rules to an enforcement policy, you need to create the policy, then reopen it.

You add rules to an enforcement policy to select the traffic you want to affect, and the actions to take. A *rule* associates an action with a specific type of traffic. So you can, for example, add a rule to select all audio-video traffic and send it to a pool of servers that are optimized to handle that type of traffic.

- 1. On the Main tab, click Policy Enforcement > Policies.**
The Policies screen opens.
- 2. Click the name of the enforcement policy you want to add rules to.**
The properties screen for the policy opens.
- 3. In the Policy Rules area, click Add.**
The New Rule screen opens.
- 4. In the Name field, type a name for the rule.**
- 5. In the Precedence field, type an integer that indicates the precedence for the rule in relation to the other rules.** Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with Gate Status disabled for a search engine and you have rule 2 with precedence 11 with Gate Status enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.
More modify header configuration options display.
8. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.
Other tasks describe how to do this in detail.
If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.
9. Click **Finished**.
10. Repeat steps 3-8 to create as many rules as needed to handle the traffic you are interested in.

The enforcement policy includes the rules with the conditions and actions you added.

Now you need to associate the enforcement policy with the virtual server (or servers) to which traffic is directed.

Creating a rule using classification criteria

You can use Layer 7 classification criteria to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. On the Classification tab, in the **Classification** setting, specify Layer 7 matching criteria for the rule:
 - a) From the **Match Criteria** list, select whether you want perform actions on traffic that matches (select **Match**), or does not match (select **No Match**) the criteria specified.
 - b) From the **Category** list, select the type of traffic this rule applies to, or select **Any** for all traffic.
 - c) Some categories have specific applications associated with them. If this one does, from the **Application** list select the application this rule applies to, or select **Any** for all traffic in this category.
 - d) Click **Add** to add this match criteria to the classification.
Add as many matching criteria as are relevant to this rule.

- Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

- Click **Finished**.

You have created a rule that applies to traffic based on classification criteria.

Creating a rule using URL categorization

You have the ability to enforce policies that are configured as part of the subscriber profile, based on the URL category type. Use Layer 7 criteria to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

- On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
- Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
- In the Policy Rules area, click **Add**.
The New Rule screen opens.
- In the **Name** field, type a name for the rule.
- In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

- On the URL tab, in the **URL** setting, specify Layer 7 matching criteria for the rule :
 - From the **Match Criteria** list, select whether you want perform actions on traffic that matches (select **Match**), or does not match (select **No Match**) the criteria specified.
 - From the **URL Category** list, select the type of traffic this rule applies to.
 - Click **Add** to add this match criteria to the classification.
Add as many matching criteria as are relevant to this rule.

- Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

- Click **Finished**.

You have created a rule that applies to traffic based on URL Category.

Modifying iRule event for URL categories

On the BIG-IP® system, you can modify iRules® Event settings for URL categories.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. Select an URL category.
The URL Properties screen opens.
3. In the **Description** field, type optional descriptive text for the category.
4. In the **iRule Event** field, select the appropriate setting.
 - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
 - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

Note: `CLASSIFICATION::DETECTED` is the only event that is supported.

You have modified an iRule event setting for an existing URL category.

Creating a rule using flow conditions

You can use flow information to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).

6. On the Flow tab, in the **Flow** setting, specify Layer 4 conditions that the traffic must meet (or not meet) for this rule to apply.

Option	Description
Match	Select whether you want to perform actions on traffic that matches (select Match) or does not match (select No Match) the criteria specified.
DSCP Marking	To match incoming traffic based on a DSCP value, type an integer from 0 to 63.

Option	Description
Protocol	To specify the applicable traffic by protocol, select UDP , TCP , or leave the default value of Any .
Source Address/Mask	To match incoming traffic based on the address or network it is coming from, type the source IP address/netmask of the network you want the rule to affect. The default value is 0.0.0.0/32.
Source Port	To match incoming traffic based on the port it is coming from, type the port number you want the rule to affect. The default value (empty) matches traffic from all ports.
Source VLAN	To match incoming traffic based on the VLAN, select a previously configured VLAN.
Destination Address/Mask	To match traffic based on the address or network it is directed to, type the source IP address/netmask of the network you want the rule to affect. The default value is 0.0.0.0/32.
Destination Port	To match incoming traffic based on the port it is directed to, type the port number you want the rule to affect. The default value (empty) matches traffic headed to all ports.

- a) Click **Add** to add this match criteria to the classification.

***Tip:** F5 recommends that you keep the matching criteria in a rule simple, adding more rules to specify additional conditions rather than including too many in one rule.*

7. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

8. Click **Finished**.

You have created a rule that classifies traffic.

Creating a rule for forwarding traffic

You can create a rule that forwards traffic to an endpoint. For example, you might want to direct video traffic to a server that is optimized for video viewing.

- On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
- Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
- In the Policy Rules area, click **Add**.
The New Rule screen opens.
- In the **Name** field, type a name for the rule.
- In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. In the Gate area, for **Gate Status**, select **Enabled**.
Options provide several ways to forward the traffic.
 - To redirect traffic to a URL, for **HTTP Redirect**, select **Enabled**, and type the URL.
 - To direct traffic to specific location, from the **Forwarding** list, select an option where you would like to forward the traffic.
 - If you select **Route to Network** then the traffic flow is forwarded to the default destination.
 - If you select **Forwarding to Endpoint**, the flow is steered to a different destination and you can select one of the endpoints.
 - If you select **>Forward to ICAP virtual Server**, the flow is forwarded to the ICAP virtual server.
 - To direct traffic to more than one location (such as value-added services), from the **Service Chain** list, select the name of a service chain that you previously created.
8. Click **Finished**.

You have created a rule that forwards traffic.

Creating a rule for QoS

Before you can create a rule for Quality of Service (QoS), you need to create a bandwidth controller to use rate control.

You can create a rule that results in a QoS action such as DSCP marking, link QoS, or rate limiting.

Note:

In the mobile market, uplink and downlink is sometimes known as forward and reverse respectively.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ.*

For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. For **Gate Status**, select **Enabled**.
If you select **Disabled**, then the corresponding traffic will be dropped.
Forwarding and QoS options are displayed.
8. To set DSCP bits on the downlink traffic, for **IP Marking (DSCP)**, select **Specify**, and type a value between 0 and 63, inclusive.
The traffic that matches this rule is marked with this value.
9. To set DSCP bits on the uplink traffic, for **IP Marking (DSCP)**, select **Specify**, and type a value between 0 and 63, inclusive.
The traffic that matches this rule is marked with this value.
10. To set a Layer 2 Quality of Service (QoS) level in downlink packets, for **L2 Marking (802.1p)**, select **Specify**, and type a value between 0 and 7, inclusive.
Setting a QoS level affects the packet delivery priority.
11. To set a Layer 2 Quality of Service (QoS) level in uplink packets, for **L2 Marking (802.1p)**, select **Specify**, and type a value between 0 and 7, inclusive.
Setting a QoS level affects the packet delivery priority.
12. To apply rate control to downlink traffic, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

Note: You can assign any previously created static or dynamic bandwidth control policies. However, F5® does not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create to enforce dynamic QoS settings provisioned by the PCRF.

Depending on the bandwidth control policy, PEM™ restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, or any combination of these.

13. To apply rate control to uplink traffic and per category of application, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

Note: You can assign any previously created static or dynamic bandwidth control policies. However, we do not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create for communicating with the PCRF.

Depending on the bandwidth control policy, PEM restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, per category of applications or any combination of these.

14. Click **Finished**.

You have created a rule that manages QoS traffic.

Creating a listener

If you want to steer specific traffic, or otherwise regulate certain types of traffic, you need to have developed enforcement policies. If using a Gx interface to a PCRF, you need to create a listener that connects to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup on the BIG-IP[®] system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP[®] system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. Subscriber provisioning using RADIUS is enabled by default. If your system is using RADIUS for snooping subscriber identity, you need to specify VLANs and tunnels. If you are not using RADIUS, you need to disable it.
 - For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
 - If you do not want to use RADIUS, from the **Subscriber Identity Collection** list, select **Disabled**.
7. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

***Note:** For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.*

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

8. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager[™] also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP[®] system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Chapter 4

Configuring Quota Management using Rating Groups

- *Overview: Configuring quota management* |

Overview: Configuring quota management

You can use the Policy Enforcement Manager™ to provision prepaid charging per subscriber or application that communicates with the quota protocol endpoint (QPE), such as online charging system (OCS), over the 3GPP Gy interface. The Gy endpoint allows online credit control for service data flow-based charging. This type of policing is called quota management; this feature ensures that subscribers do not consume resources that are not authorized.

Task summary

Creating a listener for quota management

Creating rating groups

About Gy support and rating groups

The Gy interface in 3GPP architecture facilitates communication between the online charging system (OCS) and the PCEF. In turn, this communication supports the advanced credit authorization and quota-specific reporting. Policy Enforcement Manager™ provides online credit control, through user configuration, for service data flow-based charging. The BIG-IP® system logically links between Gx and Gy.

The subscriber traffic contains allocated quota that is based on applications, category, or a group of them and is measured in terms of volume, time, and events. You can create a rating group, which is the same as a quota bucket. A rating group is identified by a service-identifier AVP that gathers a set of services, which has the same costs and rating type. Once you create a rating group, you can assign it to multiple rules inside the policy. For all the traffic matching the rule, quota is consumed from this bucket to make sure there is no over-subscription of resources. For example, you can have a rating group assigned to managing video traffic of 500 MB. This rating group needs to be assigned to a rule that matches the video traffic, to ensure that there is no over-subscription of subscriber traffic.

***Note:** You need to assign a default rating group on your policy rule or assign a new one. The default rating group is for all traffic that does not belong to another rating group.*

Creating a listener for quota management

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with an online charging system (OCS) over a Gy interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).

The system will create a virtual server for each protocol specified.

7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gy** from the **Supported Apps** options.
13. In the **Product Name** field, type the product name which is used to communicate with the OCS.
14. In the **Origin Host** field, type the fully qualified domain name of the OCS, for example, `ocs.xnet.com`.
15. In the **Origin Realm** field, type the realm name or network in which the OCS resides, for example, `xnet.com`.
16. In the **Destination Host** field, type the destination host name of the OCS, for example, `ocsdest.net.com`.
17. In the **Destination Realm** field, type the realm name or network of the OCS, for example, `net.com`.
18. For the **Pool Member Configuration** setting, add the OCS servers that are to be members of the Gy endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
19. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the OCS over the Gy interface. The default value is 1500.
20. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the OCS. The default value is 2.
21. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
22. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system enables classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gy interface with a diameter endpoint profile.

Creating rating groups

You can assign a rule to a rating group and attach it to a policy. For example, if you want to allocate quota to all the videos a subscriber uses from multiple on-demand Internet streaming media, you can specify a quota bucket that covers all the quota consumption and ensures that the consumption does not exceed the specified time or volume.

1. On the Main tab, click **Policy Enforcement > Rating Groups**.
The Rating Groups List screen opens.
2. Click **Create**.
The New Rating Group screen opens.
3. In the **Name** field, type a name for the rating group.
4. In the **Description** field, type optional descriptive text for the rating group.
5. In the **Rating Group ID** field, type an unique identifier (integer). This Rating Group ID is used by the quota managing endpoint, such as, Gy.
6. In the **Initial Quota** setting, specify **Volume** in octets, the initial quota to receive and send from the OCS, and the total quota volume.
7. In the **Initial Quota** setting, specify **Time** in seconds, the initial time for quota.
8. In the **Default Quota** setting, specify **Volume** in octets, the default quota to receive and send from the OCS, and the total quota volume.
9. In the **Threshold** field, type a default threshold level you want to use for a sending quota replenishment request.
The default value is 0, which indicates that there is no threshold.
10. In the **Usage Time** field, type the quota for how long the traffic can be used.
11. In the **Consumption Time** field, type the maximum idle time that is accounted as quota usage. This is the default value of quota for time and specifies time units for charging as well.
12. In the **Validity Time** field, type the duration for which the quota is used, if the online charging system (OCS) does not specify the validity time.
13. In the **Holding Time** field, type the holding time (in seconds), for which the quota is valid without any usage, if the time is not specified by the OCS.

Note: The default values for consumption time, validity time and holding time are used, when the OCS does not specify them.

14. From the **Breach Action** list, select the appropriate action to be taken when default quota expires or OCS does not provide new quota or breach action.

Breach Action	Description
Terminate	The system stops traffic when quota is breached.
Allow	The system allows traffic to go through even when the quota is breached.
Redirect	The system redirects traffic to the forwarding endpoint, when quota is breached.

15. From the **Request on Install** list, select **Yes** if the quota has to be requested from the Gy, when the policy is installed for a subscriber. Otherwise, select **No** for quota to be requested when one of the applications associated with the rating group is detected.

Chapter 5

Reporting Usage Data to an External Analytics Server

- *Overview: Reporting usage data to an external analytics server*
-

Overview: Reporting usage data to an external analytics server

In Policy Enforcement Manager™, you can create a rule within an enforcement policy that instructs the system to send usage data in high-speed logging (HSL) format to an external analytics server. The rule specifies what type of reporting data you are interested in; one of the actions it can take with the traffic is to send the information collected about it for processing to a centralized analytics server.

The system sends the information as a set of comma-separated values by means of SYSLOG transport. You can choose to use either the session-based or flow-based reporting format, depending on the level of granularity you need.

For example, a rule might collect session-based information about all audio and video traffic. You can specify how often to log the data and set the destination as an HSL server or pool.

Task summary

Creating a publisher

Creating a rule for high-speed logging

Creating a publisher

Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System > Logs > Configuration > Log Publishers**.
The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click << to move the destination to the **Selected** list.

Note: If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.

5. Click **Finished**.

Creating a rule for high-speed logging

Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent.

In an enforcement policy, a rule can specify that statistics about traffic affected by the rule are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.

3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. From the **Reporting** list, select **Enabled**.
8. From the **Report Granularity** list, select the appropriate option:
 - To log details about subscribers and application sessions, select **Session**.
 - For more granular reporting of every TCP connection, select **Flow**.
9. In the **Volume Threshold** setting, specify in octets, the threshold to send HSL reporting records. You can send reporting data from uplink traffic, to downlink traffic and the total traffic volume before logging the information.
10. In the **Destination** setting, specify where to send the usage monitoring data:
 - In the **Gx** field select **Enabled** for the BIG-IP system to send usage monitoring data over a Gx interface. You can then type a string for the **Gx Monitoring Key** that is used for usage monitoring.

***Note:** When you select **Session** in the **Report Granularity** field, the **Gx** field appears.*

- From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs.

***Note:** If you are using a formatted destination, select the publisher that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

11. In the **Interval** field, type an integer that specifies how frequently HSL reporting data is sent.
12. Click **Finished**.

You have created a rule that sends data about the traffic to external high-speed logging servers. The CSV reporting format differs depending on whether the report granularity is flow-based or session-based.

Session-based reporting format

In an enforcement policy, a rule can send session-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
PEM id	Identifies the reporting module (PEM) and the field value is 23003143.

Field	Description
Timestamp seconds	The time the information was logged (along with the timestamp in milliseconds), specifies seconds using UNIX time format.
Timestamp msec	The time the information was logged (along with the timestamp in seconds), specifies milliseconds using UNIX time format.
Report type	The type of report. Always set to 3 for session-based reporting.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber ID type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
3GPP parameters	The list of 3GPP parameters, which can be imsi, imeisv, tower_id, or username.
Application ID	A unique number that represents a particular application, and is used for classifying traffic.
Last Sent	The time, in seconds, since the last log entry was sent.
Bytes in	The number of bytes received during this session.
Bytes out	The number of bytes sent during this session.
Concurrent flows	Always 0 (unsupported).
Opened flows	Always 0 (unsupported).
Terminated flows	Always 0 (unsupported).
Total transactions	Always 0 (unsupported).
Successful transactions	Always 0 (unsupported).
Aggregated category duration	Summary of the duration of all flows for the session.
Reason	The reason for sending the record. It can be 0 - reserved, 1 - volume threshold reached, 2- interval time, 3 - subscriber logout, or 4 - inactivity.

Example session-based reporting format

```
Oct 10 17:19:45 172.31.63.64
23003143,1349914925,546879,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914913,5469633,308908379,
0,0,0,0,0,5052,1
Oct 10 17:19:57 172.31.63.64
23003143,1349914937,546661,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914925,5550857,313317479,
0,0,0,0,0,5063,1
Oct 10 17:20:09 172.31.63.64
23003143,1349914949,546676,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914937,5636605,318053179,
0,0,0,0,0,5074,1
```

Flow-based reporting format

In an enforcement policy, a rule can send flow-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
PEM id	Identifies the reporting module (PEM) and the field value is 2300314.
Timestamp seconds	The time the information was logged in UNIX time format.
Timestamp msec	The msec time value of the timestamp in UNIX time format.
Report type	The type of report; 0 – flow start, 1 – flow interim, 2 – flow end.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber ID type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
Source IP address	The source IP of the subscriber.
Source port	The source port the subscriber.
Destination IP address	The destination IP of the traffic.
Destination port	The destination port for the traffic.
Protocol	The protocol of the traffic for this flow, TCP or UDP.
Application ID	A unique number that represents a particular application in this flow; it is used for classifying traffic.
Urlcat ID	The URL category id that the flow belongs to.
Flow start time seconds	The time, in seconds, the flow started in UNIX time format.
Flow start time msec	The time in milliseconds of the flow start time.
Flow end time seconds	The time the flow ended in UNIX time format.
Flow end time msec	The time in milliseconds of the flow end time.
Transactions count	The count of full transactions seen in the flow.
Bytes in	The number of bytes received during this flow.
Bytes out	The number of bytes sent during this flow.

Example flow-based reporting format

```
Sep 13 13:48:58 172.31.63.60
23003143,1347546777,654398,0,4086007577,E164,2001::10,52784,2001::2,80,6,
67,1347546774,628630,4278124286,4278124286,331,156
Sep 13 13:48:58 172.31.63.60
23003143,1347546777,654398,2,4086007577,E164,2001::10,52784,2001::2,80,6,
67,1347546774,628630,1347546775,382473,547,864
```

Chapter 6

Usage Monitoring Over a Gx Interface

- *Overview: Usage monitoring over a Gx interface*
-

Overview: Usage monitoring over a Gx interface

In Policy Enforcement Manager™, you can create a rule within an enforcement policy that tells the system to send aggregated usage data concerning individual subscribers to a Policy and Charging Rules Function (PCRF). The rule specifies what type of traffic you are interested in, and one of the actions the system can take with the data collected is to send it for processing over a Gx interface to a PCRF.

The system sends the data in the standard Gx format. The report granularity must be set to session for Gx reporting to be available. The PCRF determines the policies for each subscriber, whether or not reporting is enabled, and how often to send the data.

For example, a rule might collect session-based information about all traffic destined to a particular IP address. The BIG-IP® system communicates with the PCRF and sends information about the subscribers for whom reporting is enabled. You establish the connection to the PCRF by creating a listener with Gx interface enabled.

Task Summary

Creating a listener for subscriber discovery and policy provisioning

Creating a rule for usage monitoring

Creating a listener for subscriber discovery and policy provisioning

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).
The system will create a virtual server for each protocol specified.
7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.

10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** from the **Supported Apps** options.
13. In the **Product Name** field, type the product name which is used to communicate with the PCRF.
14. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
15. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
16. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
17. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
18. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
19. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
20. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
21. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
22. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Creating a rule for usage monitoring

In an enforcement policy, a rule can specify that usage monitoring statistics concerning traffic affected by the rule are sent to a Gx interface.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.

The properties screen for the policy opens.

3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. From the **Reporting** list, select **Enabled**.
8. From the **Report Granularity** list, select **Session**.
You can send only session-based reporting data over the Gx interface.
9. For the **Volume Threshold** setting, specify, in octets, the amount of data to receive from the client, send to the client, and the total traffic volume before logging the information.
10. For the **Destination** setting, specify these values:
 - a) For **Gx**, select **Enabled**.
 - b) In the **Gx Monitoring Key** field, type a string to use for usage monitoring of the service data that the enforcement policy rule or dynamic policy and charging control (PCC) rule controls.
11. Click **Finished**.

You have created a rule that sends data about the traffic to the Gx interface in the standard Gx format.

Chapter

7

Configuring Global Application Policies with Bandwidth Control

- *Overview: Global Application Policies with Bandwidth Control*

Overview: Global Application Policies with Bandwidth Control

You can use bandwidth controllers with Policy Enforcement Manager™ to restrict bandwidth usage per subscriber, group of subscribers, per application, and so on. This implementation uses PEM for global application control to limit overall bandwidth for all P2P (or other application) traffic. For example:

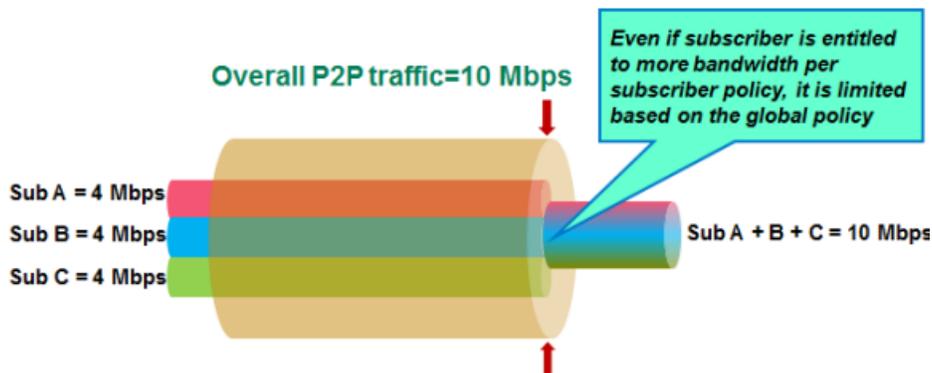


Figure 2: Diagram of bandwidth usage per subscriber

In the figure, three subscribers have individual policies that allow P2P bandwidths of up to 4 Mbps each. The maximum bandwidth for all P2P traffic is limited to 10 Mbps (specified as the maximum rate in a static bandwidth controller). If all were sending P2P traffic, they would all get less bandwidth if you apply a global enforcement policy that enforces bandwidth control.

For this implementation, you create the bandwidth controller and the enforcement policy on the BIG-IP® system. In the enforcement policy, a rule applies bandwidth control to P2P traffic. From the listener, you apply the policy globally to all traffic.

Task Summary

- Creating VLANs*
- Creating a static bandwidth control policy*
- Creating an enforcement policy*
- Creating a rule for bandwidth control*
- Creating a listener: example*

Creating VLANs

VLANs represent a collection of hosts that can share network resources, regardless of their physical location on the network. For Policy Enforcement Manager™, you typically create VLANs for the subscriber traffic coming in to the BIG-IP® system, for traffic going out to the network, and if using w-steering with service chains, you need two VLANs for each value added service to be fully transparent.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.

4. In the **Tag** field, type a numeric tag, from 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting, from the **Available** list, click an interface number or trunk name and add the selected interface or trunk to the **Untagged** list. Repeat this step as necessary.
6. From the **Configuration** list, select **Advanced**.
7. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
8. In the **MTU** field, retain the default number of bytes (**1500**).
9. If you want to base redundant-system failover on VLAN-related events, select the **Fail-safe** box.
10. From the **CMP Hash** list, select the appropriate value depending on the location of the VLAN in the system:
 - On the VLAN coming in to the BIG-IP system (often called *internal*), select **Source Address**.
 - On VLANs going out (often called *external*), leave the value set to **Default**.
 - For traffic returning to the BIG-IP from the Internet, select **Destination Address**.
 - If using w-steering for value-added services, on the VLAN coming back to the BIG-IP system, select **Source Address**.
11. Click **Finished**.
The screen refreshes, and displays the new VLAN from the list.

Create as many VLANs as needed for your configuration.

Creating a static bandwidth control policy

You can create a static bandwidth control policy to limit the bandwidth that traffic uses on the BIG-IP® system.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for the resource you are managing.
The number must be in the range from 1 Mbps to 320 Gbps. This value is the amount of bandwidth available to all the connections going through this static policy.
5. Click **Finished**.

For the bandwidth control policy to take effect, you must apply the policy to traffic, using a virtual server, packet filter, or route domain.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.

The New Policy screen opens.

3. In the **Name** field, type a name for the policy.

***Tip:** When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.*

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

***Important:** The system performance is significantly affected, depending on complexity of the classification and the type of policy action.*

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Creating a rule for bandwidth control

If you want to use rate control, you need to have already created a bandwidth controller.

You can create a rule that provides bandwidth control. For example, the bandwidth controller might limit the total amount of bandwidth that can be used by application traffic, such as P2P.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. In the **Classification** setting, filter the application traffic to which you want to apply bandwidth control.
 - a) For **Match Criteria**, select **Match**.
 - b) For **Category**, select **P2P** (or other application traffic you want to limit on the network).
 - c) Click **Add**.
7. In the Forwarding area, ensure that **Gate Status** is set to **Enabled**.
8. In the **Rate Control** setting, for **Bandwidth Controller**, select the name of the bandwidth controller that you created to limit P2P (or other application) traffic.
9. Click **Finished**.

You have created a rule to restrict the total bandwidth usage for all P2P traffic to the **Maximum Rate** specified in the static bandwidth control policy.

The enforcement policy needs to be associated with the virtual servers required for PEM. You can do this by creating a listener (recommended), or you can edit the virtual servers to specify the enforcement policy as a global policy, and enable classification.

Creating a listener: example

You create a listener to complete the preliminary setup on the BIG-IP® system; in this case, to apply bandwidth management as a global enforcement policy.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. Subscriber provisioning using RADIUS is enabled by default. If your system is using RADIUS for snooping subscriber identity, you need to specify VLANs and tunnels. If you are not using RADIUS, you need to disable it.
 - For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
 - If you do not want to use RADIUS, from the **Subscriber Identity Collection** list, select **Disabled**.
7. In the Policy Provisioning area, for **Global Policy**, move the enforcement policy you created for bandwidth control to **High Precedence**.
The system applies the policy with bandwidth control to all traffic.
8. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. All traffic classified as P2P traffic is limited to the **Maximum Rate** specified in the static bandwidth control policy. Once the maximum rate is reached, no additional P2P traffic is allowed on the network.

Chapter

8

Enforcing Bandwidth Control Provisioned on PCRF

- *Overview: Enforcing bandwidth control provisioned on PCRF*
 - *Implementation result*
-

Overview: Enforcing bandwidth control provisioned on PCRF

Policy Enforcement Manager™ (PEM) can enforce bandwidth limits provisioned by the PCRF using dynamic PCC rules. You do this by creating a dynamic bandwidth controller with the name `dynamic_spm_bwc_policy`. This bandwidth controller must be created on the BIG-IP system® using this predefined name. It does not need to be associated with an enforcement policy in PEM. Subscribers are assigned bandwidth in proportion to their configured rates. So for example, if subscriber A is assigned 4Mbps, and B is assigned 8Mbps, B will always get twice the bandwidth that A gets.

Task Summary

Creating a dynamic bandwidth control policy for PCRF

Creating a listener for subscriber discovery and policy provisioning

Creating a dynamic bandwidth control policy for PCRF

To set up bandwidth control through PCRF, you must have bandwidth control rules configured on the PCRF.

You can create a dynamic bandwidth controller so that PEM™ can enforce the maximum bit rate configured on the PCRF. You must follow the steps exactly as described here using the specified name for the bandwidth controller, and you must create the associated categories.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type the name `dynamic_spm_bwc_policy`.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for all the instances created for this dynamic bandwidth control policy.
For example, use 4Gbps or a value that is appropriate for your hardware platform.

The number must be in the range from 1Mbps to 320Gbps. You can set it to the maximum throughput value for the hardware, or to 90% of that value if you want to retain part of the system bandwidth for other uses.
5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure.
For example, use 50Mbps.

The number must be in the range from 1Mbps to 2Gbps. However, the value you use is just a place holder and is never used by the system. For this example, the value is overridden by the PCRF.
7. Leave the **IP Type of Service** and **Link Quality of Service** values set to **Pass Through**, the default value.
8. In the **Categories** field, add eight categories of traffic with the names `cat1` through `cat8`.
 - a) In the **Category Name** field, type one of the names `cat1` through `cat8`.
 - b) In the **Max Category Rate** field, type 100 and select **%**.
 - c) Click **Add** to add the category to the **Categories** list.
 - d) Repeat to add the remaining seven, to total eight categories.
9. Click **Finished**.

If this is the first bandwidth control policy created on a BIG-IP® device, the system also creates a default static bandwidth control policy named `default-bwc-policy` to handle any traffic that is not included in the policy you created. If you delete all bandwidth controllers, this policy is also deleted.

For PEM to enforce bandwidth control, you need to create a listener with a Gx interface configured.

Creating a listener for subscriber discovery and policy provisioning

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).
The system will create a virtual server for each protocol specified.
7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

***Note:** For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.*

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** from the **Supported Apps** options.

13. In the **Product Name** field, type the product name which is used to communicate with the PCRF.
14. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
15. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
16. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
17. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
18. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
19. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
20. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
21. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
22. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Implementation result

When traffic flows through the BIG-IP® system, the system limits the aggregated bandwidth for all subscribers to the **Maximum Rate** specified in the `dynamic_spm_bwc_policy` bandwidth control policy. The PCRF provides the **Maximum Rate Per User** for each subscriber, overriding the value in the bandwidth control policy. Policy Enforcement Manager™ restricts subscribers to the maximum user rate, and bandwidth is spread among subscribers fairly, on a best effort basis.

Chapter

9

Configuring Tiered Services with Bandwidth Control

- *Overview: Configuring tiered services with bandwidth control*
- *Implementation result*

Overview: Configuring tiered services with bandwidth control

You can set up Policy Enforcement Manager to enforce different levels of bandwidth control on subscribers, providing more bandwidth to subscribers with higher tier subscriptions. Bandwidth control in this case is per subscriber and per application.

This implementation provides three tiers of service: gold (the highest level), silver (the next highest), and bronze (the lowest level). You create three dynamic bandwidth controllers, one for each tier to provide different bandwidth limits for subscribers with different plans. Each tier includes bandwidth control limits for three types of application traffic (P2P, audio-video, and web). You also create three enforcement policies, one for each tier. In the enforcement policies, rules applies bandwidth control to the different types of traffic.

Finally, subscribers are provisioned dynamically through a policy charging and rules function (PCRF) over a Gx interface. On the PCRF, you need to have associated subscribers with one of the subscriber tiers called gold, silver, and bronze.

Task Summary

Creating dynamic bandwidth control policies for tiered services

Creating enforcement policies for three tiers

Creating the rules for tiered bandwidth control

Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF

Creating dynamic bandwidth control policies for tiered services

You can create dynamic bandwidth controllers for tiered services so that PEM can enforce different rates of bandwidth control for subscribers having different policy levels. Use this procedure and the values specified to create three bandwidth controllers, one for each tier of service.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type the name of the bandwidth controller. In this example, name the three bandwidth controllers as follows:
 - Type `gold-bwc` for the premium subscription level.
 - Type `silver-bwc` for the medium subscription level.
 - Type `bronze-bwc` for the lowest subscription level.

There is no requirement to use these names, but it is convenient to use a similar name for the bandwidth controller and the enforcement policy that you will attach it to. Later in this example, you will attach the `gold-bwc` bandwidth controller to the gold enforcement policy.

4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for all the subscribers using each bandwidth controller. For this example, specify **10 Mbps** for all three bandwidth controllers
If you want to use different values, the number must be in the range from 1Mbps to 320Gbps.
5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure relative to the tier of service. For example, use the following values:
 - For `gold-bwc`, specify **8 Mbps**.

- For silver-bwc, specify **4 Mbps**.
 - For bronze-bwc, specify **2 Mbps**.
7. Leave the **IP Type of Service** and **Link Quality of Service** values set to **Pass Through**, the default value.
 8. In the **Categories** field for each bandwidth controller, add three categories of traffic: `P2P`, `Web`, and `Audio-video`.
 - For gold-bwc, set `P2P` to 20%, set `Web` to 70%, and set `Audio-video` to 40%.
 - For silver-bwc, set `P2P` to 20%, set `Web` to 60%, and set `Audio-video` to 30%.
 - For bronze-bwc, set `P2P` to 20%, set `Web` to 50%, and set `Audio-video` to 20%.

In the rule for the policy, different types of traffic, P2P, web, or audio-video traffic is limited to a smaller percentage of the total bandwidth for all subscribers who use that policy.

9. Click **Finished**.

If this is the first bandwidth control policy created on a BIG-IP® device, the system also creates a default static bandwidth control policy named `default-bwc-policy` in the Common partition to handle any traffic that is not included in the policy you created. If you delete all bandwidth controllers, this policy is also deleted.

For PEM to enforce bandwidth control, you need to create enforcement policies with rules that refer to the bandwidth controller.

Creating enforcement policies for three tiers

To impose bandwidth control on multiple subscriber tiers, you need to create an enforcement policy for each tier. Use this procedure and the values specified to create three enforcement policies.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.
 - Type `gold` for the premium subscription level.
 - Type `silver` for the medium subscription level.
 - Type `bronze` for the lowest subscription level.
4. Click **Finished**.
The new enforcement policy is added to the policy list.

You have three enforcement policies that represent the three tiers of subscriber traffic that you are creating.

After creating the enforcement policies, you add rules that specify how to treat the subscriber traffic in each tier. In the implementation being developed, subscribers in the different tiers will get different maximum amounts of bandwidth. Further limits will be placed on specific types of traffic (P2P, audio-video, and web).

Creating the rules for tiered bandwidth control

You next add rules to each of the enforcement policies you created (gold, silver, and bronze). The rules limit the amount of bandwidth that P2P, web, and audio-video traffic can use. Create three rules for each enforcement policy.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule. For the first rule, use the name **P2P**.
5. In the **Precedence** field, type an integer that indicates the precedence, 1 being the highest.
In this case, you can use any value, for example, 10, as the precedence for all the rules in all the policies because there is no conflict between the rules you are creating. Each rule applies to a different type of traffic: web, audio-video, and P2P.
6. In the **Classification** setting, specify the type of traffic.
 - a) For the first rule, from the **Category** list, select **P2P**. Use the default values for **Match Criteria (Match)** and **Application (Any)**.
 - b) Click **Add**.
7. In the **Rate Control** setting, for **Bandwidth Controller**, select the name of the bandwidth controller and category. Choose
 - a) For **Bandwidth Controller**, select the name that matches the policy you are working on. For example, if editing the gold policy, select **gold-bwc**.
 - b) For **Category**, select the category that matches the type of traffic specified by the name of the rule. For example, select **P2P**.
8. Click **Finished**.
9. Repeat steps 3-8 to create a second rule for audio-video traffic with these settings.

Option	Values
Name	Audio-video
Precedence	10 (any value is OK)
Classification Category	Audio_video
Rate Control-Bandwidth Controller	Same as the name of the policy you are adding the rule to (gold-bwc, silver-bwc, or bronze-bwc)
Bandwidth Controller-Category	Audio-video

10. Repeat steps 3-8 to create a third rule for web traffic with these settings.

Option	Values
Name	Web
Precedence	10 (any value is OK)
Classification Category	Web
Rate Control-Bandwidth Controller	Same as the name of the policy you are adding the rule to (gold-bwc, silver-bwc, or bronze-bwc)
Bandwidth Controller-Category	Web

The gold, silver, and bronze enforcement policies each have three rules called P2P, Web, and Audio-video. Each of the rules in the gold policy connects to the gold-bwc bandwidth controller; rules in the silver policy connect to the silver-bwc bandwidth controller and; rules in the bronze policy connect to the bronze-bwc policy.

Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF

You create a listener to specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).
The system will create a virtual server for each protocol specified.
7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. For the tiered services example, do not assign global policies.
12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** from the **Supported Apps** options.
13. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
14. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
15. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdst.net.com`.
16. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
17. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
18. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.

19. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
20. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
21. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP system, the system handles policy enforcement.

Implementation result

When traffic flows through a BIG-IP® system, the system limits the aggregated bandwidth for all subscribers having a gold, silver, or bronze policy. Subscribers with a gold policy can use more of the total bandwidth than silver or bronze subscribers. Further, subscriber traffic in any of the tiers that is classified as audio-video, web, or P2P is limited to a percentage of the total bandwidth allowed for that tier.

For example, if a subscriber has a silver subscription level and PEM classifies their traffic as Web, the traffic is limited to 60% of the **Maximum Rate** specified in the `silver-bwc` bandwidth controller (4 Mbps). This leaves 2.4 Mbps as the maximum bandwidth for all web traffic of silver tier subscribers.

Chapter 10

Configuring Service Chains

- *Overview: Configuring service chains* |

Overview: Configuring service chains

You can use the Policy Enforcement Manager™ to create service chains to route traffic to one or more value-added services on the way to its final destination. The *service chains* define the path and order that you want traffic to take. There are several value-added services involved and after each endpoint the traffic comes back to the BIG-IP system. An *endpoint* specifies each place you want to send the traffic, so the service chain is essentially between the value-added services endpoints for traffic to stop at on its way to the server it is headed to. For example, you can forward traffic sequentially for virus scanning, parental control, and caching.

You set up service chains by creating an enforcement policy that defines the traffic that you want to route to the service chain. Rules in the enforcement policy specify conditions that the traffic must match, and actions for what to do with that traffic. One of the actions you can take is to send the traffic to a service chain.

While a static service chain defines fixed value-added services, a dynamic service chain provides service chain action that can dynamically change depending on the flow of parameters and you can attach a steering policy that can override the decision of the next session. You can use dynamic service chain to insert or name header and steer different service. Internet Content Adaptation Protocol (ICAP) is one of the services possible to use in a service chain. Dynamic service chain makes the service chain intelligent and flexible by providing the following support:

- Ability to add or skip different value-added services endpoints by selecting policy based forwarding endpoint.
- Perform header insertion or removal per value-added service chain, depending on the policy.
- Includes one sideband value-added service in the service chain using ICAP as the protocol.

You can create listeners to set up virtual servers and associate enforcement policies with the traffic that is sent to them. The system also creates a Policy Enforcement profile that specifies the enforcement policy that the system uses for the service chain.

Task Summary

Creating a ICAP profile for policy enforcement
Creating a Request Adapt profile
Creating a Response Adapt profile
Creating an internal virtual server for ICAP server
Creating a pool
Creating endpoints for service chains
Creating dynamic service chains
Creating an enforcement policy
Configuring steering action policy
Adding rules to an enforcement policy
Creating a rule for forwarding traffic
Creating a listener

About services profiles

You can configure the Internet Content Adaptation Protocol (ICAP) profile, request adaptation profile, and response adaptation profile for using the dynamic service chain feature in Policy Enforcement Manager™.

The internal virtual server references the pool of content adaptation servers. The internal virtual server also references an ICAP profile, which includes specific instructions for how the BIG-IP® system should modify each request or response. Once the request and response adapt profiles have been created, you can attach the profiles to the HTTP virtual server. The adapt profiles use multiple internal virtual servers for various content types.

The HTTP listener must have adapt profile set. The adapt profiles need to be configured as disabled and are enabled only when the forwarding endpoint is ICAP.

Creating a ICAP profile for policy enforcement

You create this ICAP profile when you want to use an ICAP server to wrap an HTTP request in an ICAP message before the BIG-IP® system sends the request to a pool of web servers. The profile specifies the HTTP request-header values that the ICAP server uses for the ICAP message.

1. On the Main tab, click **Local Traffic > Profiles > Services > ICAP**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. Click **Finished**.

After you create the ICAP profile, you can assign it to an internal virtual server so that the HTTP request that the BIG-IP system sends to an ICAP server is wrapped in an ICAP message, according to the settings you specified in the ICAP profile.

Creating a Request Adapt profile

You create a Request Adapt type of profile when you want a standard HTTP virtual server to forward HTTP requests to an internal virtual server that references a pool of ICAP servers. A Request Adapt type of profile instructs the HTTP virtual server to send an HTTP request to a named internal virtual server for possible request modification.

1. On the Main tab, click **Local Traffic > Profiles > Services > Request Adapt**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. For the **Parent Profile** setting, retain the default value, `requestadapt`.
5. On the right-side of the screen, clear the **Custom** check box.
6. Disable the setting by clearing the **Enabled** check box.
When you clear the **Enabled** check box, Policy Enforcement Manager™ controls this based on the policy.
7. In the **Preview Size** field, type a numeric value.
This specifies the maximum size of the preview buffer. This buffer holds a copy of the HTTP request header and the data sent to the internal virtual server, in case the adaptation server reports that no adaptation is needed. Setting the preview size to 0 disables buffering of the request and should only be done if the adaptation server always returns a modified HTTP request or the original HTTP request.
8. For the **Allow HTTP 1.0** setting, select the **Enabled** check box.
9. Click **Finished**.

After you perform this task, the BIG-IP® system contains a Request Adapt profile that a standard HTTP virtual server can use to forward an HTTP request to an internal virtual server for ICAP traffic.

Creating a Response Adapt profile

You create a Response Adapt type of profile when you want a standard HTTP virtual server to forward HTTP responses to an internal virtual server that references a pool of ICAP servers. A Response Adapt type of profile instructs the HTTP virtual server to send an HTTP response to a named internal virtual server for possible response modification.

1. On the Main tab, click **Local Traffic > Profiles > Services > Response Adapt**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. For the **Parent Profile** setting, retain the default value, `responseadapt`.
5. On the right-side of the screen, select the **Custom** check box.
6. Disable the setting by clearing the **Enabled** check box.
When you clear the **Enabled** check box, Policy Enforcement Manager™ controls the profile based on the policy.
7. In the **Preview Size** field, type a numeric value.
This specifies the maximum size of the preview buffer. This buffer holds a copy of the HTTP response header and the data sent to the internal virtual server, in case the adaptation server reports that no adaptation is needed. Setting the preview size to 0 disables buffering of the response and should only be done if the adaptation server always returns a modified HTTP response or the original HTTP response.
8. For the **Allow HTTP 1.0** setting, check the **Enabled** check box.

After you perform this task, the BIG-IP® system contains a Response Adapt profile that a standard HTTP virtual server can use to forward an HTTP response to an internal virtual server for ICAP traffic.

Creating an internal virtual server for ICAP server

You perform this task to create a standard virtual server that can forward an HTTP request or response to an internal virtual server. The internal virtual server then sends the request or response to a pool of ICAP servers before the BIG-IP® system sends the request or response to the client or web server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address that you want to use as a destination for client traffic destined for a pool of HTTP web servers.
The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type 80, or select **HTTP** from the list.
6. From the **Type** list, select **Internal**.
7. For the **State** setting, verify that the value is set to **Enabled**.
8. From the **HTTP Profile** list, select the name of the HTTP profile that you created previously.
9. From the **Configuration** list, select **Advanced**.
10. From the **Request Adapt Profile** list, select the ICAP profile that you previously created for handling HTTP requests.

11. From the **Response Adapt Profile** list, select the ICAP profile that you previously created for handling HTTP responses.
12. From the **Source Address Translation** list, select **Auto Map**.
The BIG-IP® system uses all of the self IP addresses as the translation addresses for the pool.
13. From the **Default Pool** list, select the pool of ICAP servers that you previously created.
14. Click **Finished**.

After you create the virtual server, the BIG-IP® system can forward an HTTP request or response to a pool of ICAP servers before sending the request or response to the client or web server, respectively.

Creating a pool

You can create a pool of servers that you can group together to receive and process traffic.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) Type an IP address in the **Address** field.
 - b) Type a port number in the **Service Port** field, or select a service name from the list.
 - c) To specify a priority group, type a priority number in the **Priority Group Activation** field.
 - d) Click **Add**.
5. Click **Finished**.

The new pool appears in the Pools list.

Creating endpoints for service chains

Before you can create an endpoint, you need to create a pool that specifies where you want to direct the classified traffic.

If you plan to set up a service chain, you need to create one or more endpoints that specify the locations of the value-added services to which to send the traffic.

1. On the Main tab, click **Policy Enforcement > Forwarding > Endpoints**.
The Endpoints screen opens.
2. Click **Create**.
The New Endpoint screen opens.
3. In the **Name** field, type a name for the endpoint.
4. From the **Pool** list, select the pool to which you want to steer a particular type of traffic.
5. Use the default values for the other fields.
6. Click **Finished**.
The endpoint you created is on the endpoint list.

You link the endpoints together by creating a service chain.

Creating dynamic service chains

Before you can create a service chain, you need to have created endpoints for every service that you want the traffic to be directed to. Set up the servers at those endpoints to handle the traffic and (if conditions are right), return it to the BIG-IP® system. You should have attached the HTTP virtual server to the request adapt profile and response adapt profile. You also need to create VLANs for every traffic entry point.

To send traffic to multiple endpoints, including value-added services, you create service chains that define where to send traffic on the way to its final destination. This way, the system can route traffic to other servers that can handle additional functions. Additionally, you can attach a steering action policy, such as modify headers, when you create a service chain which can be later modified at the other end.

Note: If you want to use steering policy, you must define endpoint in service chain.

1. On the Main tab, click **Policy Enforcement > Forwarding > Service Chains**.
The Service Chains screen opens.
2. Click **Create**
The New Service Chains screen opens.
3. In the **Name** field, type a name for the service chain.
4. In the **Service Chain List** setting, add the endpoints to the service chain. For each place you want to send the traffic, specify the following information:
 - a) From the **Service Endpoint Name** list, type the name of the service endpoint where the traffic is going to.
 - b) From the **VLAN** list, select the name of the VLAN where the traffic is coming from.

Note: Your first service chain should have subscriber VLAN in the VLAN field.

- c) From the **Policy** list, select the name of the steering policy.

Note: If all the service endpoints do not have a steering policy, the service chain is static.

Important: If the policy defining the steering does not match the policy set in the service chain, then the service chain is not processed.

- d) From the **Forwarding Endpoint** list, select the name of the endpoint to which you send traffic.
When you configure a new forwarding endpoint (**Policy Enforcement > Forwarding > Endpoints**), set **Address Translation** and **Port Translation** as **Disabled**.

Note: You need to always configure a default forwarding endpoint or else the flow will exit the service chain and get skipped. If you are in the final leg, then configure without default.

Important: When you use ICAP service, you cannot have a ICAP and a forwarding endpoint on the same service endpoint.

- e) From the **Service Option** list, select the service option in case the service endpoint is not reachable. Select **Optional** if you want to skip the service endpoint. Select **Mandatory** if you want all traffic flows dropped.

Note: To use dynamic service chain, select **Optional**. If service endpoint is not available and set to mandatory, you cannot steer policies.

- f) From the **Internal Virtual** list, select the internal ICAP virtual server.

Important: You cannot have consecutive ICAP on the same VLAN.

- g) Click **Add**.

5. Click **Finished**.

Note: If steering action is applied after the ICAP request, service endpoint with forwarding endpoint should have the same VLAN configured as the service endpoint with ICAP enabled.

You can direct traffic to the service chain you created in the policy rules in an enforcement policy.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.

Tip: When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

Important: The system performance is significantly affected, depending on complexity of the classification and the type of policy action.

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Configuring steering action policy

You can configure HTTP headers of the steering policy in the BIG-IP® system.

Note: If the steering action is enabled, steering policy is evaluated based on the VLAN flow. If no steering policy is configured, then the default endpoint is the next service endpoint.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.

3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.
More modify header configuration options display.
7. To modify the HTTP request header, select the action you want to implement.
 - Select **Insert String Value** to insert a stringvalue that you have specified before.
 - Select **Insert Value from Script** to specify that the BIG-IP system can insert value received from the TCL expression.
 - Select **Remove** to remove the string value that you previously created.
8. In the **Header Name** field, type a header name.
9. In the **String Value** field, type a string value for the header.
10. Click **Finished**.

You can add more rules to an enforcement policy in addition to configuring HTTP header action.

Adding rules to an enforcement policy

Before you can add rules to an enforcement policy, you need to create the policy, then reopen it.

You add rules to an enforcement policy to select the traffic you want to affect, and the actions to take. A *rule* associates an action with a specific type of traffic. So you can, for example, add a rule to select all audio-video traffic and send it to a pool of servers that are optimized to handle that type of traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it*

has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.
More modify header configuration options display.
8. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.
Other tasks describe how to do this in detail.
If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.
9. Click **Finished**.
10. Repeat steps 3-8 to create as many rules as needed to handle the traffic you are interested in.

The enforcement policy includes the rules with the conditions and actions you added.

Now you need to associate the enforcement policy with the virtual server (or servers) to which traffic is directed.

Creating a rule for forwarding traffic

You can create a rule that forwards traffic to an endpoint. For example, you might want to direct video traffic to a server that is optimized for video viewing.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have a rule 1 with precedence 10 with **Gate Status** disabled for a search engine and you have rule 2 with precedence 11 with **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule).*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
Other tasks describe how to do this in detail.
7. In the Gate area, for **Gate Status**, select **Enabled**.

Options provide several ways to forward the traffic.

- To redirect traffic to a URL, for **HTTP Redirect**, select **Enabled**, and type the URL.
- To direct traffic to specific location, from the **Forwarding** list, select an option where you would like to forward the traffic.
 - If you select **Route to Network** then the traffic flow is forwarded to the default destination.
 - If you select **Forwarding to Endpoint**, the flow is steered to a different destination and you can select one of the endpoints.
 - If you select **>Forward to ICAP virtual Server**, the flow is forwarded to the ICAP virtual server.
- To direct traffic to more than one location (such as value-added services), from the **Service Chain** list, select the name of a service chain that you previously created.

8. Click **Finished.**

You have created a rule that forwards traffic.

Creating a listener

If you want to steer specific traffic, or otherwise regulate certain types of traffic, you need to have developed enforcement policies. If using a Gx interface to a PCRF, you need to create a listener that connects to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup on the BIG-IP[®] system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP[®] system.

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. Subscriber provisioning using RADIUS is enabled by default. If your system is using RADIUS for snooping subscriber identity, you need to specify VLANs and tunnels. If you are not using RADIUS, you need to disable it.
 - For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
 - If you do not want to use RADIUS, from the **Subscriber Identity Collection** list, select **Disabled**.
7. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

8. Click Finished.

The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP® system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Chapter 11

Provisioning Dynamic Subscribers

- *Overview: Provisioning dynamic subscribers* |

Overview: Provisioning dynamic subscribers

If you have subscribers that are managed on a separate policy charging and rules function (PCRF), you can connect the BIG-IP® system to that policy server to provision dynamic subscribers. *Dynamic subscribers* are subscribers that are managed by a separate PCRF.

The BIG-IP system receives traffic from GGSN, a gateway between the GPRS mobile network and the Internet. When a subscriber makes a request that is routed to the BIG-IP system, the Policy Enforcement Manager™ queries the PCRF over a Gx interface. The PCRF responds with information about the subscriber. This information is stored on the BIG-IP system, which recognizes the subscriber in future requests.

You can use dynamic subscriber provisioning alone or in combination with static subscribers.

Provisioning dynamic subscribers

If you want to steer specific traffic, or otherwise regulate certain types of traffic, you need to have developed enforcement policies.

To provision subscribers dynamically through a separate PCRF, you create listeners that specify how to handle traffic for policy enforcement. Creating a listener provides preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, reporting, and other features.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Create**.
The New Listener screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Destination** setting, select **Host** or **Network**, and type the IP address or network and netmask to use.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is routed to the BIG-IP® system.

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).
The system will create a virtual server for each protocol specified.
7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** from the **Supported Apps** options.
13. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
14. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
15. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
16. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
17. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
18. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
19. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
20. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
21. Click **Finished**.
The Policy Enforcement Manager creates a listener, and displays the listener list.

When you create a listener, the Policy Enforcement Manager® also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for the Gx interface. The system also creates a virtual server to handle HTTP traffic. The system assigns the appropriate classification and policy enforcement profiles to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Chapter 12

Provisioning Static Subscribers

- *Overview: Provisioning static subscribers* |

Overview: Provisioning static subscribers

If you have subscribers that are not managed on a separate policy charging and rules function (PCRF), you can create static subscribers. *Static subscribers* are individual subscribers of services that are not managed by a separate policy server. You can add static subscribers directly to the Policy Enforcement Manager™, and assign enforcement policies to them. Each subscriber can have one or more enforcement policies associated with them.

You can use static subscribers alone or in combination with those managed by a PCRF.

Provisioning single subscribers

Before you create a subscriber, you need to create the enforcement policy that you want to associate with the subscriber.

You can provision single subscribers to the BIG-IP® Policy Enforcement Manager™ and associate policies with them. This is useful, for example, if you are not using a separate policy charging and rules function (PCRF) for subscribers, or if you need to add a subscriber with a unique enforcement policy.

1. On the Main tab, click **Policy Enforcement > Subscribers**.
The Subscribers screen opens.
2. Click **Create**.
The New Subscriber screen opens.
3. From the **Subscriber ID Type** list, select the format of the subscriber ID.

Option	Description
E.164	A numbering plan that defines the format of an MSISDN international phone number (up to 15 digits). The number typically consists of three fields: country code, national destination code, and subscriber number.
IMSI	International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.
NAI	Network Access Identifier. A fully qualified network name in the form <user>@<realm>; identifies a subscriber and the home network to which the subscriber belongs.
Private	The subscriber ID type is private for the given deployment.

4. In the **Subscriber ID** field, type a unique identifier (up to 64 characters) for the subscriber, such as a phone number. The subscriber ID type determines the format.
5. In the **Subscriber IP Address** field, type the IP address of the subscriber in IPv4 or IPv6-Prefix format. This field is optional but recommended.

Tip: *Assigning the subscriber IP address ensures that the subscriber gets the entitled service faster.*

6. In the **Policies** setting, select at least one enforcement policy from the **Available** list and move it to the **Selected** list.
The selected policy is the one that the system enforces for the subscriber you are adding.
7. Click **Finished**.
The Policy Enforcement Manager creates a static subscriber.

When the subscriber accesses the network through the BIG-IP system, the Policy Enforcement Manager applies the policy you assigned to the subscriber traffic.

Subscriber CSV file format

You can upload a list of static subscribers for policy enforcement on the BIG-IP® system. A CSV file is a text file in a comma-separated-values format. Microsoft® Excel® can be used to view and create such files. The list of subscribers must be in a text file in comma separated value (CSV) format.

Each line of the CSV file represents one subscriber in the format:

```
subscriber_ID,subscriber_ID_type,subscriber_IP,policy_1,policy_2,policy_n
```

subscriber ID

A unique identifier for the subscriber that depends on the subscriber ID type.

subscriber ID type

The format of the subscriber ID. This field is optional, but the comma is required. If you omit subscriber ID type, the system assigns the default value of IMSI (International Mobile Subscriber Identity).

subscriber IP

The IP address of the subscriber. This field is optional, but the comma is required.

policy_1,policy_2,policy_n

One policy or more policies assigned to the subscriber. Multiple policies must be separated by commas.

For example:

```
subscriber_11.1.1.1,private,11.1.1.1,bronze,gold,silver
```

You must include Subscriber ID and at least one policy enforcement policy for each subscriber. You need to include the comma for missing fields. Do not include spaces between values. The policies listed must be included on the policies list in **Policy Enforcement > Policies**, and be provisioned using a listener or a policy enforcement profile.

For example, to specify subscriber 8315555555 in IMSI format with a gold policy, you can leave out the subscriber ID type and subscriber IP address:

```
8315555555,,,gold
```

Provisioning a file of static subscribers

You can upload a CSV file containing a list of static subscribers. The file must be in a specific CSV file format.

1. On the Main tab, click **Policy Enforcement > Subscribers**.
The Subscribers screen opens.
2. Click **Upload**.
3. In the **File Name** setting, click the button to browse to the text file.
The system opens a screen where you can select the text file
4. Click **Upload**.

The Policy Enforcement Manager™ uploads the static subscriber file in chunks of 1000 subscribers. The system performs a validation on each chunk. If a validation fails, the subscribers in the current chunk and

Provisioning Static Subscribers

subsequent chunks are not imported. However, the subscribers loaded in previous chunks are imported onto the system. The Policy Enforcement Manager applies the appropriate policy to traffic from the subscribers in the list.

Chapter 13

Creating Reports using PEM

- *Overview: Creating reports using PEM* |

Overview: Creating reports using PEM

You can view logs of the traffic passing through the BIG-IP® system using the reporting feature in Policy Enforcement Manager™ (PEM™).

The logs can be either session-based or flow-based. Session-based logs provide information about the session, while flow-based logs contain information about the traffic flow.

In Policy Enforcement Manager™, reporting provides a default format. The default format can be changed by defining a format script, which enables you to add your own text and fields, and customize the log messages you would like to see in the logging server.

For example, you can create a format script for session reporting by typing: `return "(session bytes-in:[PEM::session stats reported bytes-in])`. The iRules® command retrieves the value of the attribute, bytes-in, in session record.

Similarly, you can create a format script for flow reporting by typing: `return "(flow bytes-in:[PEM::flow stats reported bytes-in])`. The iRules command retrieves the value of the attribute, bytes-in, in flow record.

Creating format scripts for reports

In an enforcement policy, format scripts can be defined for formatting report messages.

1. On the Main tab, click **Policy Enforcement > Reporting**.
2. Click **Create**.
The Format Scripts screen opens.
3. In the **Name** field, type a name for the custom format script.
4. In the **Description** field, type a description of the format script.
5. In the **Definition** field, specify the format script that defines the log messages you want.

```
return "(session bytes-in:[PEM::session stats reported bytes-in], bytes-out:[PEM::session
stats reported bytes-out],
      subs-id:[PEM::session stats reported subs-id], subs-id-type:[PEM::session stats
reported subs-id-type],
      param-3gpp:[PEM::session stats reported param-3gpp], app-id:[PEM::session stats
reported app-id],
      last-sent-sec:[PEM::session stats reported last-sent-sec], [PEM::session stats reported
last-sent-usec],
      flows-concurrent:[PEM::session stats reported flows-concurrent],
flows-new:[PEM::session stats reported flows-new],
      flows-ended:[PEM::session stats reported flows-ended], flows-duration:[PEM::session
stats reported flows-duration],
      rec-reason:[PEM::session stats report rec-reason])"
```

The iRule expression is in square brackets. You can select the **Wrap Text** check box to wrap the definition text, and select the **Extend Text Area** check box to increase the field space of format scripts.

6. Click **Update** to save any changes.
7. To attach to an HSL endpoint, follow the steps:
 - a) Click **Policy Enforcement > Policies**.
 - b) Select a policy name.
 - c) Click a policy rule.
 - d) From the **Reporting** list, select **Enabled**.

- 8.** From the **Report Granularity** list, select the appropriate option:
 - To log details about subscribers and application sessions, select **Session**.
 - For more granular reporting of every TCP connection, select **Flow**.
- 9.** In the **Volume Threshold** field, specify, in octets, the amount of data to receive from the client, send to the client, and the total traffic volume before logging the information.
- 10.** In the **Destination** setting, for **HSL**, select the HSL endpoint and then select the format script that you created in **Policy Enforcement > Reporting**.
- 11.** Click **Finished**.

You have created a reporting format script that enables you to customize the reporting string to any format.

Chapter 14

Updating Signatures for Application Recognition

- *Overview: Updating classification signatures* |

Overview: Updating classification signatures

Classification signatures define different types of traffic that Policy Enforcement Manager™ (PEM) can recognize. PEM™ recognizes a predefined set of signatures for common applications and application categories that are updated periodically. You can download signature updates from F5 Networks, and schedule the system to automatically update the signatures. You can also manually install the classification signatures and updates, for example, if the BIG-IP® system does not have Internet access.

Task Summary

Importing signatures manually

Scheduling automatic signature updates

Importing signatures manually

If you are uploading classification signatures manually because you do not have Internet access, you must have previously downloaded the signatures file, have it stored locally, and be able to navigate to it. You can obtain the latest signature update file from <http://downloads.f5.com>.

On the BIG-IP® system, you can import a locally stored classification signature update file.

1. On the Main tab, click **Policy Enforcement > Signatures**.
The **Signatures** screen opens.
2. If you have Internet connectivity, click **Check for Updates** to update the signature file if an update is available.
3. To upload signature file update if the BIG-IP system does not have Internet connectivity, under the Signature Definitions area, click **Import Signatures**.
The Signatures screen displays a **Signatures File** field where you can browse to a previously downloaded signature file update.
4. To upload a signature file update, in the Signature Definitions area, click **Import Signatures**.
The Signatures screen displays a **Signatures File** field where you can select the new signature file.
5. In the **Signatures File** field, click **Choose File** to navigate to the previously uploaded signatures file.
6. Click **Upload**.
A message displays indicating whether your upload was successful.

For signature updates to take effect, you must type `bigstart restart tmm` on the `tmsh` command line.

Scheduling automatic signature updates

You can set up the BIG-IP® system to automatically update the classification signatures. This ensures that the system always has the latest classification signature files.

1. On the Main tab, click **Policy Enforcement > Signatures**.
The Signatures screen opens.
2. Click **Check for Updates** to manually upload a signature file update if one is available.
You see the current date and time in the **Latest Update Check** setting of the Signature Definitions area.
3. For the **Automatic Updates** setting, select **Enabled**.

4. From the **Update Schedule** setting, select **Daily**, **Weekly**, or **Monthly** to specify how often you want the system to check for updates.
5. Click **Update** to save your settings.

The signature updates take effect immediately.

Chapter 15

Creating Custom Classifications

- *Overview: Creating custom classifications*
- *Classification iRule commands*

Overview: Creating custom classifications

The Policy Enforcement Manager™ (PEM) includes predefined classification signatures for many standard categories and applications. If the predefined signatures are not sufficient for classifying your traffic, you can create custom categories and applications. To use the custom categories and applications, you need to create iRules® to classify the traffic and act on the traffic.

Task summary

Determining and adjusting traffic classifications

Creating custom classification categories

Creating custom classification applications

Using iRules with classification categories and applications

Modifying iRule event for URL categories

Determining and adjusting traffic classifications

The BIG-IP® system classifies many categories of traffic and specific applications within those categories. You can determine which categories and applications of traffic the system can classify, and find out information about them such as their application or category ID.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. To view the applications in each category, click the + icon next to the category.
3. To view or edit the properties of the application or category, click the name to open its properties screen.

Tip: Here you can view the application or category ID number.

4. Adjust the properties of the application or category, if necessary.
 - In the **Description** field, you can add text to describe the application or category.
 - Set **State** to **Enabled** to use this classification, or to **Disabled** not to use it.
 - For categories only, set **iRule Event** to **Enabled** if you want the system to trigger an iRule event when it recognizes traffic in this category, or set to **Disabled** if you do not.
 - In the **Category** or **Application List** field, you can change which category an application is in, or which applications are in the category.
5. Click **Update** to save any changes.

Creating custom classification categories

On the BIG-IP® system, you can create customized categories for classifying traffic if the predefined categories are not sufficient for your needs. For example, if you plan to create new application types unique to your organization, you can create a category to group them together.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. Click **Create**.

The New Classification screen opens.

3. From the **Type** list, select **Category**.
4. In the **Name** field, type a name for the classification category.
5. In the **Description** field, type optional descriptive text for the category.
6. In the **Category ID** field, type an identifier for this category, a unique number in the range between 20480 and 24576, inclusive.
7. For the **State** setting, select the appropriate value from the list.
 - If you want the system to recognize this classification, select **Enabled**.
 - If you do not need this classification, select **Disabled**.
8. In the **iRule Event** field, select the appropriate setting.
 - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
 - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

Note: CLASSIFICATION::DETECTED is the only event that is supported.

9. For the **Application List** setting, move applications that you want to associate with this category from the **Unknown** list to the **Selected** list.
If the applications are not listed yet, you can associate the applications with the category when you create them.
10. Click **Finished**.

You can create new application types to put into this classification category.

Creating custom classification applications

On the BIG-IP® system, you can create customized applications for classifying traffic if the predefined applications are not sufficient for your needs. You can add applications to existing categories or to customized categories that you have previously created.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. Click **Create**.
The New Classification screen opens.
3. From the **Type** list, select **Application**.
4. In the **Name** field, type a name for the classification category.
5. In the **Description** field, type optional descriptive text for the category.
6. In the **Application ID** field, type an identifier for this application, a unique number in the range between 8192 and 16383, inclusive.
7. For the **State** setting, select the appropriate value from the list.
 - If you want the system to recognize this classification, select **Enabled**.
 - If you do not need this classification, select **Disabled**.
8. From the **Category** list, select the category into which to place this application.
9. Click **Finished**.

Using iRules with classification categories and applications

If you are using custom classification categories or applications, you can use iRules[®] to identify the traffic for the custom classifications, or you can initiate an action based on how the traffic is classified.

1. On the Main tab, click **Local Traffic > iRules**.
2. Click **Create**.
3. In the **Name** field, type a 1- to 31-character name.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.

For example, to classify traffic as `xxx_app`, a custom classification application that you created, you can use this iRule:

```
when HTTP_REQUEST {
    if { [HTTP::header "Host"] contains "xxx" } {
        CLASSIFY::application set xxx_app
    }
}
```

For example, to perform an action (in this case, drop) on traffic classified as `xxx_app`, you can use this iRule:

```
when CLASSIFICATION_DETECTED {
    if { [CLASSIFICATION::APP == "xxx_app"]} {
        drop
    }
}
```

For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site <http://devcentral.f5.com>.

5. Click **Finished**.

After creating the iRules, you must assign them as resources for each relevant virtual server on the BIG-IP[®] system.

Modifying iRule event for URL categories

On the BIG-IP[®] system, you can modify iRules[®] Event settings for URL categories.

1. On the Main tab, click **Policy Enforcement > Classification**.
The **Classification** screen opens showing a list of the supported classification categories.
2. Select an URL category.
The URL Properties screen opens.
3. In the **Description** field, type optional descriptive text for the category.
4. In the **iRule Event** field, select the appropriate setting.
 - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
 - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

Note: `CLASSIFICATION::DETECTED` is the only event that is supported.

You have modified an iRule event setting for an existing URL category.

Classification iRule commands

When the BIG-IP® system identifies a specific type of traffic with iRules® enabled, it triggers a `CLASSIFICATION_DETECTED` event. You can use the commands within iRules for additional system flexibility to classify the flow as one or more of the application or category classifications. The `CLASSIFY` commands are available from the `HTTP_REQUEST` or `HTTP_RESPONSE` iRule events.

iRule Command	Description
<code>CLASSIFICATION::app</code>	Gets the name of the classified application (the most explicit classified application).
<code>CLASSIFICATION::category</code>	Gets the category of the application.
<code>CLASSIFICATION::disable</code>	Disables the classification for a flow.
<code>CLASSIFICATION::enable</code>	Enables the classification for a flow.
<code>CLASSIFICATION::protocol</code>	Gets the name of the classified protocol (the least explicit classified application).
<code>CLASSIFY::application set <i>appname</i></code>	Classifies the flow as <i>appname</i> and associates the category that <i>appname</i> belongs to.
<code>CLASSIFY::application set <i>appname</i></code>	Classifies the flow as <i>appname</i> and associates the category that <i>appname</i> belongs to.
<code>CLASSIFY::category set <i>catname</i></code>	Classifies the flow as <i>catname</i> and also associates the flow with the unknown category.
<code>CLASSIFY::application add <i>appname</i></code>	Adds the application <i>appname</i> to the classification statistics.
<code>CLASSIFY::category add <i>catname</i></code>	Adds the category <i>catname</i> to the classification statistics.

Chapter 16

Configuring PEM with Local Traffic Policies

- *Overview: Creating local traffic policy rules for PEM* |

Overview: Creating local traffic policy rules for PEM

Classification signatures are added as rules in the local traffic policy. The classification signatures can be used for many standard categories and applications. In addition, you can create custom categories and applications. The BIG-IP® system automatically creates a local traffic policy that is attached to a virtual server. However, when you use Policy Enforcement Manager™ (PEM™), you can create a policy attached to a virtual server and then the BIG-IP system creates a local traffic policy. You can add an HTTP profile and classification profile in the virtual server. The local traffic policy forms a logical link between the local traffic components and the policy.

When you create a listener, a local traffic policy is attached to the listener HTTP virtual server. If you want to create custom application signatures for certain types of traffic, you can use the local traffic policy to do that, and define the policies that allow you to classify traffic. Some policies can behave like application signatures. An *application signature* is a signature that is assigned to an application (for example, HTTP traffic).

Local traffic policies can include multiple rules. Each rule defines the signature and consists of a condition. Actions are to be performed if the condition holds. Multiple signatures can be assigned to one policy, so you can create a local traffic policy that works with PEM and includes multiple rules that do different things depending on the conditions you set up. In this type of traffic policy, each rule must include one of these PEM actions:

- Enable PEM.
- Attach an application or category ID that you created.

Note: The BIG-IP system does not allow you to attach two classification local traffic policies to the same virtual server.

Task Summary

Modifying local traffic policy rules for PEM

Creating local traffic policy rules for PEM

Creating a virtual server for SSL traffic policy enforcement

Modifying local traffic policy rules for PEM

Before you modify rules on existing policies, you must set up an application or category (**Policy Enforcement > Classification**).

You can add rules to define conditions and run specific actions for different types of application traffic in Policy Enforcement Manager™ (PEM™). For example, if you create an application signature for company A and want to send traffic from company A's website, you can perform actions, such as bandwidth control and disable **Gate status** from PEM. This is a rule that can be assigned to an existing policy.

1. On the Main tab, click **Local Traffic > Policies**.

For more information about local traffic policies, refer to *BIG-IP® Local Traffic Manager™ Implementations*.

The Policy List screen opens.

2. Click **_sys_CEC_video_policy**.

Important: *_sys_CEC_video_policy* is the default local traffic policy that is important for classification; F5® recommends that you keep the policy.

The Policy List screen opens.

3. Click **Add**.
The New Rule screen opens.
 4. In the **Rule Name** field, type a unique name for the policy, for example `companyA`.
 5. In the Rule properties area, define the application traffic to which this rule applies. Specify these values and use default values for the remainder.
 - a) From the **Operand** list, select **http-host**.
 - b) From the **Event** list, select **request**.
 - c) From the **Selector** list, select **all**.
 - d) From the **Condition** list, select **ends-with**.
 - e) Type the value; for example, `companyA.com`.
 - f) Click **Add**.
 6. In the Actions setting, define the action to apply to the traffic. Specify these values and use the default values for the remainder:
 - a) From the **Target** list, select **pem**.

Note: You can specify the application you created; in this example, it is `companyA`.

Event is set to **request** and **Action** is set to **classify**. For **Parameters**, select **application** and specify the application `/common/companyA`; click **Add**.
 - b) In the Actions area, click **Add**.
7. Click **Finished** to add the rule to the local traffic policy.
8. Verify that the rule is added to the policy (**Local Traffic** > **Policies** > `_sys_CEC_video_policy`) and scroll down to view the list of rules.
You should be able to view the rule you just created.

Now you have added a new rule to the existing policy. When you send traffic that matches the rule you defined, you should be able to see the application or category you have configured. You can view the classified traffic, as well (**Statistics** > **Classification** > **Statistics**).

Creating local traffic policy rules for PEM

You can create a new policy with rules in Policy Enforcement Manager™ (PEM™).

1. On the Main tab, click **Local Traffic** > **Policies**.

Important: `_sys_CEC_video_policy` is the default local traffic policy that is important for classification; F5® recommends that you keep the policy.

The Policy List screen opens.

2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a unique name for the policy, for example `f5`.
4. From the **Strategy** list, select **first-match**.
5. For the **Requires** setting, select **http** in the **Available** list, and move it to the **Selected** list using the Move button.
6. For the **Controls** setting, select **classification** in the **Available** list, and move it to the **Selected** list using the Move button.

7. Click **Finished** to add a new policy.
The Policy List screen opens.
8. Click the new policy created. In this example, it is **f5.com**.
9. Click **Add**.
10. In the **Rule Name** field, type the name `f5_web`.
11. In the Rule properties area, define the application traffic to which this rule applies. Specify these values and use default values for the remainder.
 - a) From the **Operand** list, select **http-host**.
 - b) From the **Event** list, select **request**.
 - c) From the **Selector** list, select **all**.
 - d) From the **Condition** list, select **ends-with**.
 - e) Type the value; for example, `f5.com`.
 - f) Click **Add**.
12. In the Actions setting, define the action to apply to the traffic. Specify these values and use the default values for the remainder:
 - a) From the **Target** list, select **pem**.

*Note: You can specify the application created (**Policy Enforcement** > **Classification**); in this example, it is `f5`.*

Event is set to **request**, **Action** is set to **classify**. For **Parameters**, select **application** and specify the application `/common/f5`.

 - b) Click **Add**.
 - c) In the Actions area, click **Add**.
13. Click **Finished** to add the rule to the local traffic policy.
14. Verify that the policy is added to the virtual servers (**Local Traffic** > **Virtual Servers**) and click the HTTP virtual server that you created.
The Virtual Server List screen opens.
15. Click **Resources**.
The listener screen opens.
16. In the Policies area, click the **Manage** button.
The screen for the HTTP virtual server opens.
17. For the **Policies** setting, select the new policy that you created, **f5.com**, from the **Available** list, and move it to the **Enabled** list using the Move button.
18. Click **Finished**.

Now you have created a new policy with an HTTP-based signature (`f5.com`). You can view traffic for `f5.com` (**Statistics** > **Classification** > **Statistics**), and also verify the rule created, by browsing to **f5.com** through the BIG-IP system.

Creating a virtual server for SSL traffic policy enforcement

The BIG-IP® system allows SSL pass through mode to collect certificate information. You have to define a virtual server that references SSL pool and classifies SSL traffic for policy enforcement.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.

The New Virtual Server screen opens.

3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, select **Network**, and type 0.0.0.0 in the **Address** field and 0.0.0.0 in the **Mask** field.
5. In the **Service Port** field, type 443 or select **HTTPS** from the list.
6. From the **Configuration** list, select **Advanced**.
7. From the **Classification** list, select **Enabled** on, for the BIG-IP system to enable classification for virtual servers when a policy enforcement listener is created.
8. From the **Policy Enforcement Profile** list, select the name of the Policy Enforcement Profile that you previously created.
9. Click **Finished**.
10. From the **Default Persistence Profile** list, select **ssl**.
This implements simple persistence, using the default ssl profile.
11. In the Policies area, click the **Manage** button.
12. For the **Policies** setting, from the **Available** list, select the name of the iRule that you want to assign, and use the buttons to move the name into the **Enabled** list.

You have created a virtual server for SSL traffic. The virtual server that references SSL pools appears in the Virtual Servers list.

Chapter 17

Configuring Policy and RADIUS Updates

- *Overview: Configuring policy and RADIUS updates* |

Overview: Configuring policy and RADIUS updates

Policy Enforcement Manager™ (PEM™) enables you to schedule policy reevaluations and radius updates on the BIG-IP® system in the following two ways:

- You can change policies for a subscriber session over a period of time. You can also configure the interval for reevaluation of policies, for a subscriber session, by configuring the re-evaluation interval. The BIG-IP system evaluates changes in the policy for traffic, once the re-evaluation interval is configured.
- The RADIUS traffic contains the subscriber and IP address information that is monitored by the BIG-IP system. If you enable the timeout interval, the BIG-IP system avoids repeated deletion and creation of the subscriber during the configured interval rate.

Configuring PEM options

You can set up the BIG-IP® system to schedule an interval that sets policy reevaluation and RADIUS re-transmission updates periodically.

1. On the Main tab, click **Policy Enforcement > Options**.
The Options screen opens.
2. In the Policy Options area, specify (in seconds) the **Policy Re-evaluation Interval** at which the policy re-evaluation is triggered, to evaluate the flow policy again.
The re-valuation interval is only for active flows.
For example, a subscriber is provisioned over GX which has a policy to allow Netflix with some bandwidth. The subscriber is able to watch a movie using the Netflix service. However, consider that the PCRF installs a policy for this subscriber to block Netflix over the Gx interface. Then, while the subscriber is viewing the content, the Netflix content is blocked for the subscriber after the configured re-evaluation interval.
3. In the RADIUS Options area, for the **Re-Transmit Timeout** setting, select **Enabled** and specify the time in seconds. If you select **Disabled**, each RADIUS message is handled as a new message and this might lead to deletion and creation of sessions even though the radius message is a duplicate. This is the timeout after which the RADIUS message is considered as a new message, by the BIG-IP system.

The policy and RADIUS updates take effect immediately.

Chapter 18

Enforcing Policy and Classification on IP Protocols

- *About enforcing policy and classification on IP protocols*

About enforcing policy and classification on IP protocols

The BIG-IP® system now provides classification and policy enforcement on all non-TCP and non-UDP traffic, which includes IPsec traffic. The Policy Enforcement Manager™ is able to classify and enforce any action on virtually any type of IP traffic. This enables detection of IPsec, ICMP, GRE, and other IP protocols (especially tunneling) for the service providers. For IPsec, Encapsulating Security Payloads (ESP) and Authentication Headers (AH) protocols are used, in both tunnel and transport modes.

A bottom hudfilter forwards non-TCP and non-UDP traffic for both classification and policy enforcement.

Note: HTTP redirect is not supported. Based on the protocol, not all actions work and some traffic is not steered.

Important: You can use SNAT, only when you forward ICMP and ICMPv6 traffic.

Creating Any IP profiles for PEM

Before you create multiple Any IP profiles, you must create a listener in Policy Enforcement Manager™ (PEM™), which creates a virtual server with Any IP profile.

You can create a new Any IP profile through local traffic management in PEM.

1. On the Main tab, click **Local Traffic > Profiles > Any IP**.
The Any IP screen opens.
2. Click **Create**.
The New Any IP Profile screen opens.
3. From the **Parent Profile** list, select the default **ipother** or any other Any IP profile, from where the new profile can inherit the settings.

Note:

You will see multiple Any IP profiles in the list only if you have created the profiles earlier.

4. To specify the idle timeout, click **Customize**, select **Specify**, and type a value (in seconds). The idle time out specifies the number of seconds for which a connection is idle before the connection is eligible for deletion.
5. Click **Finished**.

Now you have created a new Any IP profile. You can view non-TCP and UDP traffic that passes through the BIG-IP system (**Statistics > Classification > Statistics**).

Updating Any IP profile

If you have created other Any IP profile and you want to attach this profile to the Any IP traffic, then you can attach the profile through local traffic management in Policy Enforcement Manager™ (PEM™).

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Server List**.
The Virtual Server List screen opens.
2. Select any profile.

The profile's properties screen opens.

3. From the **Protocol** list, select ***All Protocols**.
Any IP profile settings displays.
4. From the **Any IP Profile** list, select the default setting **ipother**, or any other Any IP profile from where the new profile can inherit the settings.

Note:

You will see multiple Any IP profiles from the list only if you have created the profiles earlier.

5. Click **Update**.

Now you have updated the Any IP profile and attached it to the Any IP traffic.

IPOther filter for current PEM actions

The policy actions configured in the Policy Enforcement Manager™ can support non-TCP and non-UDP traffic flows. This table contains the information that highlights the actions supported for non-TCP and non-UDP traffic.

Action	All non-TCP and non-UDP flows
Forwarding	Only non-tunnel protocols. <i>Note: ICMP traffic can be steered.</i>
Service-chain	Only non-tunnel protocols. <i>Note: ICMP traffic can be steered.</i>
Cloning	Yes
BWC (both directions)	Yes
L2 QoS markings (both directions)	Yes
Flow Reporting	Yes
Session Reporting	Yes
Gate status drop	Yes
Quota	Yes
HTTP-redirect	No
Modify HTTP headers	No
iRules	CLIENT_DATA and CLIENT_ACCEPTED iRules only (like UDP filter).

Chapter

19

Troubleshooting

- *PEM troubleshooting*
-

PEM troubleshooting

Follow these general troubleshooting suggestions when using Policy Enforcement Manager™ (PEM™):

- If enforcement policies are not enforced as expected, on the VLAN screen for all VLANs set up to receive incoming subscriber traffic, verify that you set **CMP Hash** to **Source Address**.
- If static subscriber policies are not enforced as expected, verify whether you enforced any global high precedence policies with conflicting actions.
- When sending traffic without RADIUS, the unknown subscriber policy (if specified) is assigned to the first flows from dynamic or static subscribers. Subscriber policies are applied to subsequent flows.

Note: An unknown subscriber policy needs to be specified, if there is at least one dynamically provisioned subscriber.

- Policy changes are applied to new and existing flows within a reasonable time.
- For applications with connections initiated from the Internet (FTP, RTSP, TFTP), the BIG-IP® system needs to have **CMP Hash** set to **Destination Address** on the Internet VLAN. In this case, the end-to-end IP addresses have to be preserved; therefore, SNAT should be disabled on all the virtual servers that the applications will use.
- When importing static subscribers, the file is uploaded in chunks of 1000 subscribers. The system performs a validation check on each chunk. If a validation fails, the subscribers in the current chunk and subsequent chunks are not imported. However, the subscribers loaded in previous chunks are imported onto the system.

Steering troubleshooting

- In case of service chains (w-steering), set **CMP Hash** to **Source Address** on all the VLANs for which the w-steering action is to be applied.
- For response-side classification, steering, w-steering, and cloning actions are applied after the results (based on destination IP address and port) are cached in the classification database (srdb). Actions are not applied for the first six flows, by default. (This behavior is configurable by the DB variable `tmm.pem.srdb.entry.step`.)

RADIUS troubleshooting

- If static subscribers are not working as expected with RADIUS, check whether you selected the same **Subscriber ID Type** in the `radiusLB` profile (**Local Traffic > Profiles > Services > RADIUS**) as that assigned when creating the static subscriber. (**IMSI** in the static subscriber corresponds to **3GPP IMSI** in the RADIUS profile; **E164** to **Calling Station ID**, and **NAI** to **User Name**.)
- The RADIUS message also needs to specify the same **Subscriber ID Type** as the RADIUS profile. So make sure that if you select **IMSI**, the IMSI number exists in the RADIUS message. This also applies to the `user-name` for NAI, and `calling station-id` for E164.

Gx interface to PCRF troubleshooting

- If you change the IP address of the Gx server in the listener, the change takes effect after you restart TMM using the command: `bigstart restart tmm`.
- For Gx usage monitoring, the threshold is defined on the Policy and Charging Rules Function (PCRF).

Bandwidth control with PEM troubleshooting

- Do not use dynamic bandwidth control policies in preconfigured enforcement policies (either global or subscriber) when the bitrate is managed by the PCRF through PCC dynamic rules.

- Do not use dynamic bandwidth control policies in global enforcement policies if they are also used in subscriber policies.
- For bandwidth controller to work with PCRF, you need to create a default dynamic bandwidth controller with the name `dynamic_spm_bwc_policy`, with eight categories named `cat1` to `cat8` (all set to 100 percent). You must choose a proper max-rate value for this bandwidth controller (typically, close to network capacity dedicated to subscriber traffic).

Important: *This bandwidth controller is intended for internal usage only and should not be used for other purposes.*

Active sessions troubleshooting (retrieving subscriber data or BIG-IP system information)

- When the BIG-IP system receives policy information from the PCRF for a subscriber, you can verify the active policies on the subscriber session, the subscriber type (static or dynamic) and view subscriber statistics by checking **Active Sessions (Policy Enforcement > Subscribers > Active Sessions)**.
- If you have a static subscriber without an IP address, no active session is created. The incoming RADIUS message has the IP information for the static subscriber and a session is created based on this. When the radius message arrives, verify both the new session and policy attached to the session.
- If your subscriber type is dynamically provisioned, then your assigned policy can be based on a predefined PCC rule or dynamic PCC rule.
- For information about uplink and downlink traffic (byte count and flows), check (**Policy Enforcement > Subscribers > Statistics**).
- You can auto-refresh the subscriber session information for 10 to 300 seconds.

iRules® troubleshooting

- While running the script, if the BIG-IP system receives an error, ignore the error and implement the next custom action script. Although this is the default behaviour, it is possible to change it with the sys db variable `key:pem.tcl.action.error.abort`.
- If policy priority, event priority, and the rule precedence is the same, then there is no guarantee of order of execution.

IPsec® troubleshooting

- For IPsec to work with Policy Enforcement Manager™ (PEM™), disable the DB variable `ipsec.lookupspi`.

Subscriber and policies active sessions

You can view session records based on subscriber ID or session IP. Policy Enforcement Manager™ contains the information presented in this table. You can access this in **Active Sessions (Policy Enforcement > Subscribers > Active Sessions)**.

Field	Description
ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
ID type	The format of the subscriber ID attribute. It can be E.164, IMSI, NAI, or Private (RFC 4006).
Subscriber Type	Specifies a dynamically or statically subscriber.
Calling Station	Radius Attribute Value Pair (AVP) type 31 (3GPP TS 29.061 V9.6.0).

Field	Description
Called Station	Radius Attribute Value Pair (AVP) type 30 (3GPP TS 29.061 V9.6.0).
Tower	Specifies the cell tower where subscriber information goes through.
User Name	Displays the format name name@domain.
IMSI	International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.
IMSEISV	International Mobile Station Equipment Identity Software Version. A globally unique code number that identifies a GSM, UMTS, LTE, or iDEN mobile phone.
Predefined	Specifies the predetermined policy(ies) assigned to the subscriber.
Dynamic	Specifies the dynamic PCC rule applied.
Statistics	Specifies active session statistical information that includes subscriber and session IP identity attributes, assigned policy, and traffic flow information.

Active sessions statistics

You can view subscriber uplink and downlink traffic information. Policy Enforcement Manager™ contains the information presented in this table.

Field	Description
Data Format	Specifies how the system presents the statistics information. The default is Normalized .
Auto Refresh	Automatically updates the screen information at the interval you specify. For example, if you select 60 seconds from the list, the system updates the displayed screen information every 60 seconds. The default is Disabled . When you specify an automatic-refresh interval, the system presents a Stop button for halting the operation, and counts down the seconds to the next update. Select Disabled to turn off automatic refreshing.
Session IP	Specifies the session IP address. The IP address is in either IPv4 or IPv6 format.
Subscriber ID	Specifies a unique identifier subscriber ID.
Uplink	Specifies traffic volume from the subscriber to network.
Downlink	Specifies traffic volume from the network to subscriber.
Current	Specifies current number of flows.
Maximum	Specifies maximum number of open flows.
Total	Specifies accumulated number of flows ever opened by the subscriber.

Configuring subscriber activity log

You can configure the activity logs of the selected subscribers by subscriber or session activity.

1. On the Main tab, click **Policy Enforcement > Subscribers > Activity Log > Configuration**. The Configuration screen opens.

2. From the **Log Publisher** list, select the log publisher that was created. You can create a log publisher in the system at **System > Logs > Configuration > Log Publishers**.
3. From the **Subscriber Type** list, select **Dynamic** (for dynamic provisioning) or **Static** (for static provisioning) subscriber.
4. In the **Subscriber ID** field, type a unique identifier (up to 64 characters) for the subscriber, such as IMSI .
5. Using the Log Subscriber Activity setting, add each subscriber ID to the log settings.
 - a) Type the **Subscriber ID**.
 - b) Click **Add**.
6. To configure settings of the activity logs by sessions, use the Log Session Activity setting to add each session IP to the log settings.
 - a) Type the **Session IP** address.
 - b) Click **Add**.
7. Click **Update**.
Policy Enforcement Manager™ starts generating the subscriber activity logs for the configured subscribers.

You have configured the activity logs settings. Policy Enforcement Manager applies the log settings you assigned and lists subscriber activity and session information.

Index

A

- active sessions
 - records [115](#)
 - statistics [116](#)
- Any IP profile
 - updating [110](#)
- Any IP profiles
 - creating [110](#)
- application rules
 - adding to enforcement policy [29, 74](#)
- applications
 - creating custom classification [97](#)
 - overview custom classification [96](#)
- application visibility
 - defined [20](#)
- application visibility statistics
 - examining [22](#)

B

- bandwidth control
 - creating a rule for [54](#)
- bandwidth control policies
 - creating dynamic for PCRF [58](#)
 - for tiered services [62](#)
 - overview [62](#)
 - static, creating [53](#)
- bandwidth control with PEM
 - overview [52, 58](#)
 - result of [60](#)
- bandwidth control with tiered services
 - result of [66](#)
- best practices, PEM [16](#)

C

- categories
 - creating custom classification [96](#)
 - determining classification [20, 96](#)
- classification
 - using iRules [98](#)
- classification applications
 - creating custom [97](#)
 - overview [96](#)
- classification categories
 - creating custom [96](#)
 - determining [20, 96](#)
- classification data
 - overview [20](#)
- classification iRule commands [99](#)
- classification signatures
 - updating automatically [92](#)
 - updating manually [92](#)
 - updating overview [92](#)
- classification statistics
 - examining [22](#)
- creating [69–70](#)

- custom action
 - iRules [28](#)

D

- dynamic service chains
 - creating [72](#)
- dynamic subscribers
 - overview [80](#)
 - provisioning [80](#)

E

- endpoints
 - creating [27](#)
 - creating for service chains [71](#)
- enforcement policy
 - about rules [15](#)
 - adding Gx reporting rules [49](#)
 - adding HSL reporting rules [42](#)
 - adding rules to [29, 74](#)
 - best practices [16](#)
 - classifying traffic rules [30, 32](#)
 - controlling bandwidth [54, 63](#)
 - creating [28, 53, 63, 73](#)
 - creating listeners [21, 36, 55, 76](#)
 - overview [14–15](#)
 - QoS traffic rules [34](#)
 - troubleshooting [114](#)
 - URL categorization [31](#)
- enforcement rules
 - overview [15](#)

F

- forwarding traffic
 - creating rules [33, 75](#)
 - overview [26](#)

G

- global application policies
 - overview [52](#)
- Gx reporting
 - overview [48](#)
- Gy support
 - about [38](#)

H

- header values
 - for HTTP requests [69](#)
- HSL reporting
 - overview [42](#)
- HSL reporting format
 - flow-based [45](#)
 - session-based [43](#)

- HTTP request-header values *69*
- HTTP virtual server
 - send HTTP requests *69*

I

- ICAP content adaptation *68*
- ICAP profiles
 - 69*
 - assigning *70*
- internal virtual server
 - for response modification *70*
- internal virtual servers
 - creating *70*
- internal virtual server type
 - defined *68*
- IPOther filter
 - PEM actions *111*
- IPsec traffic
 - classification *110*
- iRule commands, classification *99*
- iRules
 - using with traffic classification *98*

L

- limitation, PEM *17*
- listener
 - creating *38, 48, 59, 65*
- listeners
 - connecting to a PCRF *21, 36, 55, 76*
 - creating *21, 36, 55, 76*
- local traffic policy
 - classification *102*
 - creating PEM rules *103*
 - modifying PEM rules *102*
 - SSL *104*
- logging
 - and publishers *42*
- logging format
 - HSL flow-based *45*
 - HSL session-based *43*

M

- maximum rate of throughput, See bandwidth control policies
- modify header
 - creating *73*

O

- online charging system (OCS)
 - connecting to *38*

P

- PCRF
 - enforcing bandwidth control *58*
 - provisioning subscribers *16*
- performance, PEM *17*

- policy
 - creating enforcement *28, 53, 63, 73*
 - provisioning dynamic subscribers *80*
- Policy and Charging Rules Function (PCRF)
 - connecting to *48, 59, 65*
- policy enforcement
 - overview *14*
- policy re-evaluation
 - interval *108*
- policy updates
 - overview *108*
- pools
 - creating *26, 71*
- publishers
 - creating for logging *42*

Q

- quota management with PEM
 - overview *38*

R

- RADIUS re-transmission
 - overview *108*
 - re-transmit timeout *108*
- rating groups
 - about *38*
 - creating *40*
- recommendations, PEM *16*
- reporting
 - creating format scripts *88*
 - overview *88*
- reporting format
 - HSL flow-based *45*
 - HSL session-based *43*
- request adapt profile *69*
- response adapt profile *70*
- rules
 - about enforcement *15*
 - for bandwidth control *54, 63*
 - for classifying traffic *30, 32*
 - for QoS *34*
 - local traffic policy *102*
 - setting up Gx reporting *49*
 - setting up HSL reporting *42*

S

- servers
 - and publishers for log messages *42*
- service chain
 - creating endpoints for *71*
- service chains
 - overview *68*
- signatures
 - updating automatically *92*
 - updating manually *92*
- static bandwidth control policies, See bandwidth control policies
- static subscribers
 - overview *84*

- statistics
 - examining for classification *22*
- steering
 - adding classification rules to policy *29, 74*
 - creating a policy *28, 53, 73*
 - creating endpoints *27*
 - creating service chains *72*
- steering policy
 - adding *73*
- subscriber activity
 - debuggability *116*
- subscriber CSV file format *85*
- subscribers
 - adding *84*
 - overview *80, 84*
 - provisioning dynamic *80*
 - provisioning with PCRF *16*
 - uploading *85*

T

- tiered services
 - overview *62*

- traffic classification
 - defined *20*
 - overview *20*
 - using iRules *98*
- traffic steering
 - overview *26*
- transactional policy enforcement
 - overview *26*
- troubleshooting, PEM *114*

U

- URL categories
 - creating iRule Events *32, 98*

V

- virtual servers
 - and internal type *68*
 - creating *70*
- VLANs
 - creating *52*

