

BIG-IP[®] Network Firewall: Policies and Implementations

Version 13.1



Table of Contents

About the Network Firewall.....	9
What is the BIG-IP Network Firewall?.....	9
About firewall modes.....	9
Configuring the Network Firewall in ADC mode.....	9
Configuring the Network Firewall to drop or reject traffic that is not specifically allowed.....	10
Configuring the Network Firewall to globally drop or reject traffic.....	10
Firewall Rules and Rule Lists.....	11
About firewall rules.....	11
Firewall actions.....	12
About Network Firewall contexts.....	12
Creating a network firewall management port rule.....	15
About redundant and conflicting rules.....	17
About stale rules.....	18
About firewall rule lists.....	19
Firewall Rule Addresses and Ports.....	25
About firewall rule addresses and ports.....	25
About resolving DNS addresses in Network Firewall rules.....	25
Creating a DNS resolver.....	26
Configuring the Network Firewall to use a DNS resolver.....	26
About address lists.....	26
Creating an address list.....	26
About port lists.....	27
Creating a port list.....	27
Network Firewall Schedules.....	29
About Network Firewall schedules.....	29
Creating a schedule.....	29
Network Firewall Inline Rule Editor.....	31
Using the inline firewall rule editor.....	31
Enabling the Network Firewall inline rule editor.....	32
Creating a rule with the inline editor.....	32
Editing a rule with the inline editor.....	34
Configuring BIG-IP Network Firewall Policies.....	37
About firewall policies.....	37
Creating a Network Firewall policy.....	37
Setting a global firewall policy.....	40
Configuring a route domain with a firewall policy.....	41
Setting network firewall policies for a self IP address.....	41
Creating a virtual server with a firewall policy.....	41
Viewing enforced and staged policy rule logs.....	42
Viewing Network Firewall enforced policy events on the local BIG-IP system	42

Viewing Network Firewall staged policy events on the local BIG-IP system	43
IP Address Intelligence in the Network Firewall.....	45
About IP intelligence policies in the Network Firewall.....	45
Downloading the IP intelligence database.....	46
About IP intelligence blacklist categories.....	47
Blacklist categories.....	47
Creating a blacklist category.....	48
Blacklisting an individual IP address.....	49
Removing an individual IP address from a blacklist.....	49
About IP intelligence feed lists.....	49
Feed list settings.....	50
Creating a feed list.....	51
Configuring and assigning IP intelligence policies.....	51
Configuring a policy to check addresses against IP intelligence.....	51
Assigning a global IP Intelligence policy.....	53
Assigning an IP Intelligence policy to a virtual server.....	53
Assigning an IP Intelligence policy to a route domain.....	53
Inspecting Protocol Anomalies.....	55
About protocol anomaly inspection.....	55
Creating a protocol inspection profile.....	55
Viewing protocol inspection items.....	56
Creating protocol inspection items.....	56
Assigning a protocol inspection profile to a virtual server.....	58
Assigning a protocol inspection profile to a firewall rule.....	59
Viewing protocol security inspection logs.....	60
Deploying the BIG-IP Network Firewall in ADC Mode.....	63
About deploying the network firewall in ADC mode.....	63
Configuration settings for IPv6 pools and ADC mode.....	64
Configuring the Network Firewall in ADC mode.....	65
Creating a VLAN for the network firewall.....	65
Configuring an LTM virtual server with a VLAN for Network Firewall.....	66
Adding a firewall rule to deny ICMP.....	66
Creating an address list.....	67
Denying access with firewall rules on the network virtual server.....	67
Denying access with firewall rules on the application virtual server.....	68
Deploying the BIG-IP Network Firewall in Firewall Mode.....	71
About Firewall mode in the Network Firewall.....	71
Configuring the Network Firewall to drop or reject traffic that is not specifically allowed.....	72
Creating a VLAN for the network firewall.....	73
Configuring an LTM virtual server with a VLAN for Network Firewall.....	74
Creating an address list.....	74
Allowing access from networks on an address list with a firewall rule.....	74
Allowing access from a network to a virtual server with a firewall rule.....	75
Compiling and Deploying Network Firewall rules.....	77
About compiling and deploying rules in the Network Firewall.....	77
Configuring manual or automatic policy compilation for firewall rules.....	77
Configuring manual or automatic policy deployment for firewall rules.....	78

About firewall policy compilation statistics.....	79
Viewing compilation statistics for a firewall rule or policy.....	80
Viewing compilation statistics for all network firewall rules and policies.....	80
Using Firewall NAT for IP and Port Translation.....	83
About using Firewall NAT to translate addresses and ports.....	83
About Firewall NAT and Carrier Grade NAT (CGNAT).....	84
About specifying source translations for Firewall NAT.....	84
Specifying source IP addresses for static NAT.....	86
Specifying source IP addresses for static PAT.....	86
Specifying source IP addresses for deterministic dynamic PAT.....	86
Specifying source IP addresses for dynamic PAT with NAPT.....	88
Specifying source IP addresses for port block allocation mode.....	89
About specifying destination translations for Firewall NAT.....	90
Specifying destination IP addresses for static NAT.....	90
Specifying destination IP addresses for static PAT.....	90
About creating Firewall NAT policies.....	91
Creating a NAT policy.....	91
Creating a NAT match rule.....	91
About specifying NAT context for a Firewall NAT policy.....	92
Adding a global Firewall NAT policy.....	92
Configuring a route domain to use Firewall NAT.....	93
Configuring Firewall NAT on a virtual server.....	95
HTTP Protocol Security.....	97
Overview: Securing HTTP traffic.....	97
Creating an HTTP virtual server to use with HTTP protocol security.....	97
Attaching an HTTP protocol security profile to a virtual server.....	98
Reviewing violation statistics for security profiles.....	98
Overview: Creating a custom HTTP security profile.....	98
Creating a custom HTTP profile.....	99
Creating a security profile for HTTP traffic.....	99
Configuring an HTTP virtual server with an HTTP security profile.....	100
Reviewing violation statistics for security profiles.....	100
Overview: Increasing HTTP traffic security.....	101
About RFC compliance and validation checks.....	101
Modifying HTTP protocol compliance checks.....	101
About evasion techniques checks.....	102
Configuring HTTP protocol evasion techniques blocking policy.....	102
About the types of HTTP request checks.....	102
Configuring length checks for HTTP traffic.....	103
Specifying which HTTP methods to allow.....	103
Including or excluding files by type in HTTP security profiles.....	104
Configuring a mandatory header for an HTTP security profile.....	105
Configuring the blocking response page for HTTP security profiles.....	105
Overview: Configuring Local Protocol Security Event Logging.....	106
Creating a local Protocol Security Logging profile.....	106
Configuring a virtual server for Protocol Security event logging.....	107
Viewing Protocol Security event logs locally on the BIG-IP system.....	107
Disabling logging.....	107
Implementation result.....	108
Overview: Configuring Remote Protocol Security Event Logging.....	108
About the configuration objects of remote protocol security event logging.....	109
Creating a pool of remote logging servers.....	110
Creating a remote high-speed log destination.....	111

Creating a formatted remote high-speed log destination.....	111
Creating a publisher	112
Creating a custom Protocol Security Logging profile	112
Configuring a virtual server for Protocol Security event logging.....	113
Disabling logging	113
Implementation result.....	114
Preventing Attacks with Eviction Policies and Connection Limits.....	115
What are eviction policies and connection limits?.....	115
Creating an eviction policy.....	115
Limiting global connections and flows.....	117
Limiting connections and flows on a virtual server.....	117
Limiting connections and flows on a route domain.....	118
Setting Timers and Preventing Port Misuse with Service Policies.....	119
Creating and Applying Service Policies.....	119
Introduction to service policies.....	119
About service policy types.....	119
Testing Packets with Firewall, IP Intelligence, and DoS Rules.....	123
About packet tracing with the AFM Packet Tester.....	123
Tracing a TCP packet.....	123
Tracing a UDP packet.....	124
Tracing an SCTP packet.....	124
Tracing an ICMP packet.....	125
Packet trace results.....	125
Local Logging with the Network Firewall.....	129
Overview: Configuring local Network Firewall event logging.....	129
Task summary.....	129
Creating a local Network Firewall Logging profile	129
Configuring a virtual server for Network Firewall event logging.....	131
Viewing Network Firewall event logs locally on the BIG-IP system.....	131
Creating a Network Firewall rule from a firewall log entry.....	131
Disabling logging	134
Implementation result.....	135
Remote High-Speed Logging with the Network Firewall.....	137
Overview: Configuring remote high-speed Network Firewall event logging.....	137
About the configuration objects of remote high-speed Network Firewall event logging.....	138
Creating a pool of remote logging servers.....	138
Creating a remote high-speed log destination.....	139
Creating a formatted remote high-speed log destination.....	139
Creating a publisher	140
Creating a custom Network Firewall Logging profile	140
Configuring a virtual server for Network Firewall event logging.....	142
Disabling logging	142
Implementation result.....	143
SNMP Trap Configuration.....	145
Overview: BIG-IP SNMP agent configuration.....	145

Specify SNMP administrator contact information and system location information.....	145
Configure SNMP manager access to the SNMP agent on the BIG-IP system...	145
Grant community access to v1 or v2c SNMP data.....	146
Grant user access to v3 SNMP data.....	146
Overview: SNMP trap configuration.....	147
Enabling traps for specific events.....	147
Setting v1 and v2c trap destinations.....	148
Setting v3 trap destinations.....	148
Viewing pre-configured SNMP traps.....	149
Creating custom SNMP traps.....	149
Overview: About troubleshooting SNMP traps.....	150
AFM-related traps and recommended actions.....	150
AOM-related traps and recommended actions.....	151
ASM-related traps and recommended actions.....	153
Application Visibility and Reporting-related traps and recommended actions....	154
Authentication-related traps and recommended actions.....	155
DDM-related traps and recommended actions.....	155
DoS-related traps and recommended actions.....	157
General traps and recommended actions.....	157
BIG-IP DNS-related traps and recommended actions.....	158
Hardware-related traps and recommended actions.....	162
High-availability system-related traps and recommended actions.....	166
License-related traps and recommended actions.....	167
LTM-related traps and recommended actions.....	167
Logging-related traps and recommended actions.....	169
Network-related traps and recommended actions.....	169
vCMP-related traps and recommended actions.....	170
VIPRION-related traps and recommended actions.....	170
About enterprise MIB files.....	170
Downloading enterprise and NET-SNMP MIBs to the SNMP manager.....	171
Viewing objects in enterprise MIB files.....	171
Viewing SNMP traps in F5-BIGIP-LOCAL-MIB.txt.....	171
Collecting network firewall data using SNMP.....	172
Collecting DoS attack data using SNMP.....	172
About enterprise MIB files.....	172
Downloading enterprise and NET-SNMP MIBs to the SNMP manager.....	173
Viewing objects in enterprise MIB files.....	173
Viewing SNMP traps in F5-BIGIP-COMMON-MIB.txt.....	173
Viewing dynamic routing SNMP traps and associated OIDs.....	173
Monitoring BIG-IP system processes using SNMP.....	174
Collecting BIG-IP system memory usage data using SNMP.....	174
Collecting BIG-IP system data on HTTP requests using SNMP.....	174
Collecting BIG-IP system data on throughput rates using SNMP.....	175
Collecting BIG-IP system data on RAM cache using SNMP.....	176
Collecting BIG-IP system data on SSL transactions using SNMP.....	177
Collecting BIG-IP system data on CPU usage based on a predefined polling interval.....	177
Collecting BIG-IP system data on CPU usage based on a custom polling interval.....	179
Collecting BIG-IP system performance data on new connections using SNMP.	180
Collecting BIG-IP system performance data on active connections using SNMP.....	181
About the RMON MIB file.....	182

Logging Network Firewall Events to IPFIX Collectors.....	183
Overview: Configuring IPFIX logging for AFM.....	183
About the configuration objects of IPFIX logging for AFM.....	183
Assembling a pool of IPFIX collectors.....	183
Creating an IPFIX log destination.....	184
Creating a publisher	184
Creating a custom Network Firewall Logging profile	185
Configuring an LTM virtual server for Network Firewall event logging with IPFIX.....	186
Implementation result.....	187
IPFIX Templates for AFM Events.....	189
Overview: IPFIX Templates for AFM events.....	189
About IPFIX Information Elements for AFM events.....	189
IANA-defined IPFIX Information Elements.....	189
IPFIX enterprise Information Elements.....	189
About individual IPFIX templates for each event.....	191
Network accept or deny.....	191
DoS device.....	192
IP intelligence.....	193
Log Throttle.....	194
Legal Notices.....	197
Legal notices.....	197

About the Network Firewall

What is the BIG-IP Network Firewall?

The BIG-IP® Network Firewall provides policy-based access control to and from address and port pairs, inside and outside of your network. Using a combination of contexts, the network firewall can apply rules in a number of different ways, including: at a global level, on a route domain, on a per-virtual server level, for a self IP address, or for the management port. Firewall rules are combined in firewall policies, which can contain multiple context and address pairs, and can be applied directly to any context except the management port. Rules for the management port context are defined inline, and do not require a separate policy.

By default, the Network Firewall is configured in *ADC mode*, a *default allow* configuration, in which all traffic is allowed through the firewall, and any traffic you want to block must be explicitly specified.

The system is configured in this mode by default so all traffic on your system continues to pass after you provision Advanced Firewall Manager™. You should create appropriate firewall rules to allow necessary traffic to pass before you switch Advanced Firewall Manager to Firewall mode. In *Firewall mode*, a *default deny* configuration, all traffic is blocked through the firewall, and any traffic you want to allow through the firewall must be explicitly specified.

Task list

Configuring the Network Firewall in ADC mode

Configuring the Network Firewall to drop or reject traffic that is not specifically allowed

Configuring the Network Firewall to globally drop or reject traffic

About firewall modes

The BIG-IP® Network Firewall provides policy-based access control to and from address and port pairs, inside and outside of your network. By default, the network firewall is configured in ADC mode. This means it is a *default allow* configuration, in which all traffic is allowed to virtual servers and self IP addresses on the system, and any traffic you want to block must be explicitly specified. This applies only to the virtual server and self IP levels on the system.

Important: *If a packet does not match any rule in any context on the firewall, the Global Reject or Global Drop rule drops the packet (Global Drop) or drops the packet and sends the appropriate reject message (Global Reject) even when the system is in a default allow configuration. In addition, the Global Drop or Global Reject rule does not drop or reject traffic to the management port. Management port rules must be specifically configured and applied.*

Configuring the Network Firewall in ADC mode

If you have changed the firewall setting to Firewall mode, you can configure the BIG-IP® Network Firewall back to ADC mode.

Note: *The firewall is configured in ADC mode, by default.*

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.

2. From the **Virtual Server & Self IP Contexts** list, select the default action **Accept** for the self IP and virtual server contexts.
3. Click **Update**.
The virtual server and self IP contexts for the firewall are changed.

Configuring the Network Firewall to drop or reject traffic that is not specifically allowed

You can configure the BIG-IP® Network Firewall to drop or reject all traffic not explicitly allowed. In Advanced Firewall Manager™, this is called *Firewall mode*, and this is also referred to as a *default deny* policy. Firewall mode applies a default deny policy to all self IP addresses and virtual servers.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Virtual Server & Self IP Contexts** list, select the default action for the self IP and virtual server contexts.
 - Select **Drop** to silently drop all traffic to virtual servers and self IP addresses unless specifically allowed.
 - Select **Reject** to drop all traffic to virtual servers and self IP addresses unless specifically allowed, and to send the appropriate reject message for the protocol.
3. Click **Update**.
The default virtual server and self IP firewall context is changed.

Configuring the Network Firewall to globally drop or reject traffic

If traffic to or from the BIG-IP® Network Firewall does not match a rule, the global rule handles the traffic. You can set the global rule to drop traffic or to reject traffic. The global rule rejects unmatched traffic by default.

Note: Management port traffic is not handled by the global rule. Management port rules must be explicitly defined for the management port context.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Global Context** list, select the default action for the global rule, when the traffic matches no other rule.
 - Select **Drop** to drop traffic silently.
 - Select **Reject** to drop traffic, and send the appropriate reject message for the protocol.
3. Click **Update**.
The global firewall action is changed.

Firewall Rules and Rule Lists

About firewall rules

The BIG-IP® Network Firewall uses rules to specify traffic handling actions. Rules are collected in policies, which are applied at the global context, to a route domain, to a virtual server, or to a self IP address. Rules for the management port do not require a policy, but are defined directly in the management port context.

A rule includes:

Context

The category of object to which the rule applies. Rules can be global and apply to all addresses on the BIG-IP system that match the rule, or they can be specific, applying only to a specific virtual server, self IP address, route domain, or the management port.

Rule or Rule List

Specifies whether the configuration applies to this specific rule, or to a group of rules.

Source Address

One or more addresses, geographic locations, or address lists to which the rule applies. The source address refers to the packet's source.

Source Port

The ports or lists of ports on the system to which the rule applies. The source port refers to the packet's source.

VLAN

Specifies VLANs to which the rule applies. The VLAN source refers to the packet's source.

Destination Address

One or more addresses, geographic locations, or address lists to which the rule applies. The destination address refers to the packet's destination.

Destination Port

The ports or lists of ports to which the rule applies. The destination port refers to the packet's destination.

iRule

Specifies an iRule that is applied to the rule. An iRule can be started when the firewall rule matches traffic.

iRule sampling

When you select an iRule to trigger in a firewall rule, you can select the how frequently the iRule is triggered, for sampling purposes. The value you configure is one out of n times the iRule is triggered. For example, set this field to 5 to trigger the iRule one out of every five times the rule matches a flow.

Protocol

The protocol to which the rule applies. The firewall configuration allows you to select one specific protocol from a list of more than 250 protocols. The list is separated into a set of common protocols, and a longer set of other protocols. To apply a rule to more than one protocol, select **Any**.

Schedule

Specifies a schedule for the firewall rule. You configure schedules to define days and times when the firewall rule is made active.

Action

Specifies the action (accept, accept decisively, drop, or reject) for the firewall rule.

Logging

Specifies whether logging is enabled or disabled for the firewall rule.

Task list

Creating a network firewall management port rule

Firewall actions

These listed actions are available in a firewall rule.

Firewall actions are processed within a context. If traffic matches a firewall rule within a given context, that action is applied to the traffic, and the traffic is processed again at the next context.

Firewall action	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are <code>accepted</code> , traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. For example, if the protocol is TCP, a TCP RST message is sent. One benefit of using Reject is that the sending application is notified, after only one attempt, that the connection cannot be established.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are <code>accepted decisively</code> , traverse the system as if the firewall is not present, and are not processed by rules in any further context after the <code>accept decisively</code> action applies. If you want a packet to be accepted in one context, and not to be processed in any remaining context or by the default firewall rules, specify the <code>accept decisively</code> action. For example, if you want to allow all packets from Network A to reach every server behind your firewall, you can specify a rule that accepts decisively at the global context, from that Network A, to any port and address. Then, you can specify that all traffic is blocked at a specific virtual server, using the virtual server context. Because traffic from Network A is accepted decisively at the global context, that traffic still traverses the virtual server.

About Network Firewall contexts

With the BIG-IP® Network Firewall, you use a context to configure the level of specificity of a firewall policy. For example, you might make a global context rule to block ICMP ping messages, and you might make a virtual server context rule to allow only a specific network to access an application.

Context is processed in this order:

1. Global

2. Route domain
3. Virtual server/self IP
4. Global drop or reject

The firewall processes policies and rules in order, progressing from the global context, to the route domain context, and then to either the virtual server or self IP context. Management port rules are processed separately, and are not processed after previous rules. Rules can be viewed in one list, and viewed and reorganized separately within each context. You can enforce a firewall policy on any context except the management port. You can also stage a firewall policy in any context except management. Management port rules are configured as inline rules specific to the management port.

Important: *You can configure the global drop or reject context. The global drop or reject context is the final context for all traffic, except Management port traffic. Note that even though it is a global context, it is not processed first, like the main global context, but last. If a packet matches no rule in any previous context, the global drop or reject rule drops or rejects the traffic. The default global rule is global reject.*

Notice: *Management port traffic is not affected by the global drop or reject rule, or by global rules in general. Management port rules must be specifically configured and applied.*

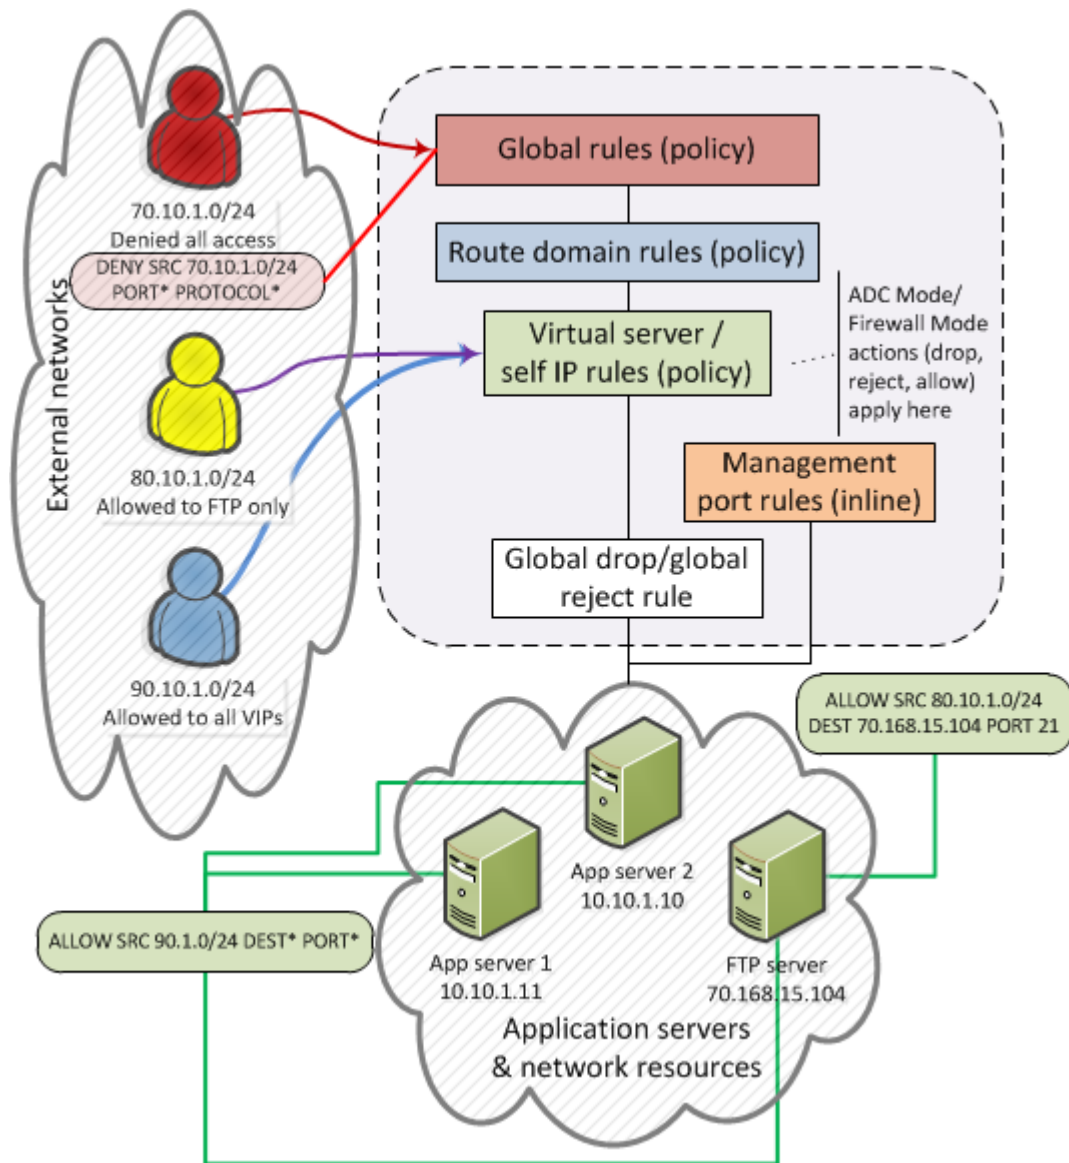


Figure 1: Firewall context processing hierarchy example

Firewall context descriptions

When you create a firewall rule, you can select one of these listed contexts. Each context forms a list of rules. Contexts are processed in hierarchical order, and within each context, rules are processed in numerical order.

Firewall context	Description
Global	Global policy rules are collected in this firewall context. Global rules apply to all traffic that traverses the firewall, and global rules are checked first.
Route Domain	Route domain policy rules are collected in this context. Route domain rules apply to a specific route domain defined on the server. Route domain policy rules are checked after global rules. If you have not configured a route domain, you can apply route domain rules to Route Domain 0, which is effectively the same as the global rule context; however, if you configure another route domain after this, Route Domain 0 is no longer usable as a global context.

Firewall context	Description
Virtual Server	Virtual server policy rules are collected in this context. Virtual server policy rules apply to the selected existing virtual server only. Virtual server rules are checked after route domain rules.
Self IP	Self IP policy rules apply to a specified self IP address on the device. Self IP policy rules are checked after route domain rules.
Management Port	The management port context collects firewall rules that apply to the management port on the BIG-IP® device. Management port rules are checked independently of other rules and are not processed in relation to other contexts.
Global Reject	The Global Reject rule rejects all traffic that does not match any rule in a previous context, excluding Management Port traffic, which is processed independently.

Creating a network firewall management port rule

If you are going to specify address lists, user lists, or port lists with this rule, you must create these lists before creating the firewall rule, or add them after you save the rule.

Create a network firewall management port rule to manage access from an IP or web network address to the BIG-IP® management port.

Note: You cannot add rules created with this task to a rule list at a later time. You must create rules for a rule list from within the rule list.

Important: You can only add management port rules as inline rules. For all other contexts, you must add rules to policies.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
 2. In the Rules area, click **Add** to add a firewall rule to the list.
 3. From the **Context** list, select **Management Port**.
 4. In the **Name** and **Description** fields, type the name and an optional description.
 5. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall rule to the given context and addresses.
 - Select **Disabled** to set the firewall rule to not apply at all.
 - Select **Scheduled** to apply the firewall rule according to the selected schedule.
 6. From the **Schedule** list, select the schedule for the firewall rule.
This schedule is applied when you set the firewall rule state as **Scheduled**.
 7. From the **Protocol** list, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
-

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the global or route domain context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

8. In the **Source** list, specify addresses and geolocated sources to which this rule applies.
 - From the **Address/Region** list, select **Any** to have the rule apply to any packet source IP address or geographic location.
 - From the **Address/Region** list, select **Specify** and click **Address** to specify one or more packet source IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - From the **Address/Region** list, select **Specify** and click **Address List** to select a predefined list of packet source addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - From the **Address/Region** list, select **Specify** and click **Address Range** to specify a contiguous range of packet source IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
 - From the **Address/Region** list, select **Specify** and click **Country/Region** to identify the geographic origin of packet sources, and to apply rules based on selected geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Source address list.
9. From the Source **Port** list, select the type of packet source ports to which this rule applies.
 - Select **Any** to have the rule apply to any packet source port.
 - Select **Specify** and click **Port** to specify one or more packet source ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet source port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
 - Select **Specify** and click **Port List** to select a predefined list of packet source ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
10. From the Destination **Address/Region** list, select the type of packet destination address to which this rule applies.
 - Select **Any** to have the rule apply to any IP packet destination address.
 - Select **Specify** and click **Address** to specify one or more packet destination addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - Select **Specify** and click **Address List** to select a predefined list of packet destination addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - Select **Specify** and click **Address Range** to specify a contiguous range of packet destination IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
11. From the Destination **Port** list, select the type of packet destination ports to which this rule applies.
 - Select **Any** to have the rule apply to any port inside the firewall.
 - Select **Specify** and click **Port** to specify one or more packet destination ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet destination port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.

- Select **Specify** and click **Port List** to select a predefined list of packet destination ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.

12. From the **Action** list, select the firewall action for traffic originating from the specified source address on the specified protocol. Choose from one of the these actions:

Option	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. When a packet is rejected the firewall sends a destination unreachable message to the sender.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall, and does not require any further processing by any of the further firewalls. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.

13. From the **Logging** list, enable or disable logging for the firewall rule.

A logging profile must be enabled to capture logging info for the firewall rule.

14. Click **Finished**.

The list screen and the new items are displayed.

The new firewall rule is created.

About redundant and conflicting rules

When you create rules on the network firewall, it is possible that a rule can either overlap or conflict with an existing rule.

Redundant rule

A rule which has address, user, region, or port information that completely overlaps with another rule, with the same action. In the case of a redundant rule, the rule can be removed with no net change in packet processing because of the overlap with a previous rule or rules.

Conflicting rule

A conflicting rule is a special case of a redundant rule, in which address, user, region or port information overlaps with another rule, but the rules have different actions, and thus conflict.

***Tip:** A rule might be called conflicting even if the result of each rule is the same. For example, a rule that applies to a specific IP address is considered in conflict with another rule that applies to the same IP address, if one has an *Accept* action and the other has an action of *Accept Decisively*, even though the two rules accept packets.*

On a rule list page, redundant or conflicting rules are indicated in the **State** column with either (Redundant) or (Conflicting).

Viewing and removing redundant and conflicting rules

You must have staged or enforced rules configured on your system that are redundant or conflicting.

View and remove redundant or conflicting rules to simplify your configuration and ensure that your system takes the correct actions on packets.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. From the **Type** list, select whether you want to view **Enforced** or **Staged** policies.

*Note: If you select to view **Staged** policies, you can not view management port rules, as they cannot be staged.*

3. View the firewall rule states in the **State** column.

Each rule is listed as Enabled, Disabled, or Scheduled. In addition, a rule can have one of the following states. View and adjust rules with these states, if necessary.

(Redundant)

The rule is enabled, disabled, or scheduled, and redundant. All the functionality of this rule is provided by a previous rule or rules. Hover over the **State** column to see why the rule is considered redundant, and possible solutions. Typically you can disable or delete a redundant rule with no net effect on the system.

(Conflicting)

The rule is enabled, disabled, or scheduled, and conflicting. All the match criteria of this rule is covered by another rule or rules, but this rule has a different action. Hover over the **State** column to see why the rule is considered conflicting, and possible solutions. Typically you should disable or delete a conflicting rule. Because the rule criteria is matched prior to the conflicting rule, there it typically no net change in processing. Note that the **Accept** and **Accept Decisively** actions are treated as conflicting by the system.

(Conflicting & Redundant)

The rule is enabled, disabled, or scheduled, and conflicting or redundant with the actions of more than one other rule. Typically you should disable or delete a conflicting and redundant rule.

4. Resolve conflicting or redundant rules by editing, deleting, or disabling them. Click a rule name to edit, delete, or disable it, and complete the required action.

The firewall rule list is adjusted.

About stale rules

On the rule list page, you can determine whether a rule is stale, infrequently used, or never used. A *stale* rule is one that has not been hit in a long time. In addition, a rule might never be hit, or might be hit infrequently.

Note: Use discretion when tuning rules, and delete rules only when you are sure they are no longer needed.

On the active rules page, or the page of rules for a policy, the **Count** column displays the number of times a rule has been hit. A count of 0 might indicate a rule that will never be hit, and can be removed without changing packet processing. A rule with a low count, when other rules have a high count, might indicate a rule that is stale, and no longer needed.

Use the **Latest Match** column to confirm rule status. A status of **Never** indicates the rule has never been matched, and might be irrelevant. A very long time since the last match indicates a rule that is likely no longer needed.

You can view stale rules from the stale rules reporting page. Go to **Security > Reporting > Network > Stale Rules**.

Viewing and removing unused or infrequently used rules

You must have staged or enforced rules configured on your system, and the system must be processing traffic, to determine whether rules are hit.

View and remove infrequently used or unused rules to reduce firewall processing and simplify your rules, rule lists, and policies.

Caution: Before you remove a rule that is infrequently hit, or never hit, make sure that doing so will not create a security issue. A rule might be hit infrequently, but might still be a required part of your security stance for a specific or rare attack.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. From the **Type** list, select whether you want to view **Enforced** or **Staged** policies.

Note: If you select to view **Staged** policies, you can not view management port rules, as they cannot be staged.

3. View the rule hit count in the **Count** column.
The rule hit count shows how many total times a rule hit has occurred. A very low number indicates that the rule is infrequently hit. A count of 0 indicates the rule has never been hit.
4. View the latest match date in the **Latest Match** column.
The latest match column lists the last time the rule was hit. An old date indicates that the rule has not been hit in a long time. **Never** indicates that the rule has never been hit.
5. Resolve infrequently hit rules by editing, deleting, or disabling them. Click a rule name to edit, delete, or disable it, and complete the required action.

The firewall rule list is adjusted.

About firewall rule lists

The BIG-IP® Network Firewall uses rule lists to collect multiple rules. Rule lists function differently depending on how you create them with Advanced Firewall Manager™ (AFM™).

If you create a rule list with **Security > Network Firewall > Rule Lists > Create**:

This type of rule list is defined with a name and optional description. Once you create a rule list of this type, you can create and add one or more individual firewall rules to it. You can only add firewall rules by creating them from within the rule list. This type of rule list cannot be used on its own, but must be selected in an Active Rules list, or in a Policy Rules list.

If you create a rule list with **Security > Network Firewall > Active Rules > Add and select the Type as Rule List**:

This type of rule list is defined with a name and optional description. You can specify a context (Global, Route Domain, Virtual Server, or Self IP). However, you cannot add individual rules to this rule list. Instead, you select a single rule list you have already created, or one of the predefined rule lists. This type of rule list is used to activate a rule list in the configuration.

If you create a rule list with **Security > Network Firewall > Policies > *policy_name* > Add and select the Type as Rule List**:

This type of rule list is defined with a name and optional description. You cannot specify a context as the context is determined by the policy. You cannot add individual rules to this rule list. Instead, you select a single rule list you have already created, or one of the predefined rule lists. This type of rule list is used to activate a rule list in a policy.

Creating a network firewall rule list

Create a network firewall rule list, to which you can add firewall rules.

1. On the Main tab, click **Security > Network Firewall > Rule Lists**.
The Rule Lists screen opens.
2. Click the **Create** button to create a new rule list.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. Click **Finished**.
The empty firewall rule list is displayed.

Add firewall rules to the rule list to define source, destination, and firewall actions.

Adding a Network Firewall rule to a rule list

Before you add a firewall rule to a rule list, you must create a rule list.

Add a network firewall rule to a rule list so you can collect rules and apply them at once in a policy.

1. On the Main tab, click **Security > Network Firewall > Rule Lists**.
The Rule Lists screen opens.
2. From the list, click the name of a rule list you previously created.
The Rule List properties screen opens.
3. In the Rules area, click **Add** to add a firewall rule to the list.
4. In the **Name** and **Description** fields, type the name and an optional description.
5. From the **Order** list, set the order for the firewall rule.
You can specify that the rule be first or last in the rule list, or before or after a specific rule.
6. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall rule to the given context and addresses.
 - Select **Disabled** to set the firewall rule to not apply at all.
 - Select **Scheduled** to apply the firewall rule according to the selected schedule.
7. From the **Protocol** list, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the *global or route domain* context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

8. If you select ICMP or ICMPv6 as the rule protocol, add ICMP message types and codes in the fields that appear.

If you do not specify specific ICMP/ICMPv6 message types and codes, the rule applies to any ICMP or ICMPv6 message type.
 - In the ICMP/ICMPv6 Message area, select an ICMP message type from the **Type** list, and select an ICMP message code from the **Code** list.
 - Click **Add** to add the message type and code to the firewall rule.
9. From the **Schedule** list, select the schedule for the firewall rule.
This schedule is applied when you set the firewall rule state as **Scheduled**.
10. In the **Source** list, specify addresses and geolocated sources to which this rule applies.

- From the **Address/Region** list, select **Any** to have the rule apply to any packet source IP address or geographic location.
- From the **Address/Region** list, select **Specify** and click **Address** to specify one or more packet source IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
- From the **Address/Region** list, select **Specify** and click **Address List** to select a predefined list of packet source addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
- From the **Address/Region** list, select **Specify** and click **Address Range** to specify a contiguous range of packet source IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
- From the **Address/Region** list, select **Specify** and click **Country/Region** to identify the geographic origin of packet sources, and to apply rules based on selected geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Source address list.

11. From the Source **Port** list, select the type of packet source ports to which this rule applies.

- Select **Any** to have the rule apply to any packet source port.
- Select **Specify** and click **Port** to specify one or more packet source ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
- Select **Specify** and click **Port Range** to specify a list of contiguous packet source port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
- Select **Specify** and click **Port List** to select a predefined list of packet source ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.

12. From the Source **VLAN/Tunnel** list, select the VLAN on which this rule applies.

- Select **Any** to have the rule apply to traffic on any VLAN through which traffic enters the firewall.
- Select **Specify** to specify one or more VLANs on the firewall to which the rule applies. To use a VLAN with this rule, move the VLAN from the **Available** list to the **Selected** list. Similarly, you can remove the VLAN from this rule, by moving the VLAN from the **Selected** list to the **Available** list.

13. In the Destination area and from the **Address/Region** list, select the type of packet destination address to which this rule applies.

- Select **Any** to have the rule apply to any IP packet destination address.
- Select **Specify** and click **Address** to specify one or more packet destination IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
- Select **Specify** and click **Address List** to select a predefined list of packet destination addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
- Select **Specify** and click **Address Range** to specify a contiguous range of packet destination IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
- Select **Specify** and click **Country/Region** to identify the geographic packet destination, and to apply rules based on specific geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in

which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Destination address list.

14. From the Destination **Port** list, select the type of packet destination ports to which this rule applies.

- Select **Any** to have the rule apply to any port inside the firewall.
- Select **Specify** and click **Port** to specify one or more packet destination ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
- Select **Specify** and click **Port Range** to specify a list of contiguous packet destination port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
- Select **Specify** and click **Port List** to select a predefined list of packet destination ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.

15. Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.

16. When you select an iRule to start in a firewall rule, you can enable iRule sampling, and select how frequently the iRule is started, for sampling purposes. The value you configure is one out of n times the iRule is triggered. For example, to trigger the iRule one out of every five times the rule matches a flow, select **Enabled**, then set this field to 5.

17. From the **Action** list, select the firewall action for traffic originating from the specified source address on the specified protocol. Choose from one of the these actions:

Option	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. When a packet is rejected the firewall sends a destination unreachable message to the sender.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall, and does not require any further processing by any of the further firewalls. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.

18. From the **Logging** list, enable or disable logging for the firewall rule.

A logging profile must be enabled to capture logging info for the firewall rule.

19. Click **Finished**.

The list screen and the new items are displayed.

A new firewall rule is created, and appears in the Rules list.

Activating a rule list in a policy

The rule list is a container in which you can select and activate one of the rule lists that you created previously, or one of the predefined system rule lists, to apply a collection of rules at one time, to a policy.

1. On the Main tab, click **Security > Network Firewall > Policies**.

The Policies screen opens.

2. Click the name of a firewall policy to edit that policy.

The Firewall Policy screen opens, or the policy expands on the screen.

3. In the Rules area, click **Add** to add a firewall rule list to the policy.
4. In the **Name** and **Description** fields, type the name and an optional description.
5. From the **Order** list, set the order for the firewall rule.
You can specify that the rule be first or last in the rule list, or before or after a specific rule.
6. From the **Type** list, select whether you are creating a standalone network firewall rule or creating the rule from a predefined rule list.

*Note: If you create a firewall rule from a predefined rule list, only the **Name**, **Description**, **Order**, **Rule List**, and **State** options apply, and you must select or create a rule list to include.*

7. From the **Rule List** setting, select a rule list to activate in the policy or configuration.
8. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall rule to the given context and addresses.
 - Select **Disabled** to set the firewall rule to not apply at all.
 - Select **Scheduled** to apply the firewall rule according to the selected schedule.
9. Click **Finished**.
The list screen and the new items are displayed.

The firewall rule list you selected is activated.

Firewall Rule Addresses and Ports

About firewall rule addresses and ports

In a Network Firewall rule, you have several options for defining addresses and ports. You can use one or more of these options to configure the ports and addresses to which a firewall rule applies.

Note: You can use any combination of inline addresses, ports, address lists, and port lists in a firewall rule.

Users, Groups, or User lists

You can specify predefined user lists, users, and groups. Users and groups must be specified in the form `domain\user_name` or `domain\group_name`. You can select user lists from a list. Users are defined on the BIG-IP® Access Policy Manager®.

Any (address or port)

In both **Source** and **Destination** address and port fields, you can select **Any**. This specifies that the firewall rule applies to any address or port.

Fully qualified domain names

You can specify source or destination addresses as fully qualified domain names. To do this, you must create a DNS resolver cache, and configure the network firewall FQDN Resolver option.

Inline addresses

An inline address is an IP address that you add directly to the network firewall rule, in either the **Source** or **Destination Address** field. You can specify a single IP address, multiple IP addresses, a contiguous range of IP addresses, or you can identify addresses based on their geographic location. IP addresses can be either IPv4 or IPv6, depending on your network configuration.

Address lists

An address list is a preconfigured list of IP addresses that you add directly to the BIG-IP system. You can select this list of addresses to use in either the **Source** or **Destination Address** field. An address list can also contain other address lists, and geographic locations.

Inline ports

An inline port is a port that you add directly to the network firewall rule, in either the **Source** or **Destination Port** field. You can add a single port, or a contiguous port range.

Port lists

A port list is a preconfigured list of ports that you add directly to the BIG-IP system. You can select this list of ports to use in either the **Source** or **Destination Port** field. You can also add port lists to other port lists.

About resolving DNS addresses in Network Firewall rules

You can configure a DNS resolver on the BIG-IP® system to resolve DNS queries and cache the responses, and provide the resolved DNS addresses to network firewall rules that use fully qualified domain names (FQDNs). The next time the system receives a query for a response that exists in the cache, the system returns the response from the cache. The *resolver cache* contains messages, resource records, and the nameservers the system queries to resolve DNS queries.

After you specify a DNS resolver, you specify the DNS resolver in the Network Firewall options, to allow firewall rules to resolve and cache IP addresses from FQDNs.

Creating a DNS resolver

You configure a DNS resolver on the BIG-IP® system to resolve DNS queries and cache the responses. The next time the system receives a query for a response that exists in the cache, the system returns the response from the cache.

1. On the Main tab, click **Network > DNS Resolvers > DNS Resolver List**.
The DNS Resolver List screen opens.
2. Click **Create**.
The New DNS Resolver screen opens.
3. In the **Name** field, type a name for the resolver.
4. Click **Finished**.

Note: When you create an OAuth Server, creating a DNS Resolver with a forward zone named . (period) is mandatory to forward all requests.

Configuring the Network Firewall to use a DNS resolver

You must configure a DNS resolver on the BIG-IP® system before you select the DNS resolver in the firewall options.

The global DNS resolver specifies a DNS resolver for the network firewall to use, when resolving fully qualified domain names (FQDNs) to IP addresses.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. In the FQDN Resolver area, from the **Global Context** list, select the DNS resolver.
3. In the **Refresh Interval** field, specify how often the DNS resolver refreshes the IP addresses associated with fully qualified domain names, in minutes.
The default refresh interval is 60 minutes.
4. Click **Update**.
The DNS resolver is configured for firewall rules.

About address lists

An address list is simply a collection of addresses saved on the server, including fully qualified domain names, IP addresses, contiguous IP address ranges, geographic locations, and other (nested) address lists. You can define one or more address lists, and you can select one or more address lists in a firewall rule. Firewall address lists can be used in addition to inline addresses that are specified within a particular rule.

Creating an address list

Create an address list to apply to a firewall rule, in order to match IP addresses.

1. On the Main tab, click **Security > Network Firewall > Address Lists**.
The Address Lists screen opens.
2. Click **Create** to create a new address list.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. In the Addresses area, add and remove addresses.

- To add an IP address, type the address and press **Enter**.
- To add an IP address range, type the start and end IP addresses, separated by a dash, and press **Enter**.
- To add an existing address list, start typing the name of the address list. A list of items (address lists and geographic locations) will appear. Select the address list and press **Enter**.
- To add a geographic location, start typing the name of the geographic location. A list of items (address lists and geographic locations) will appear. Select the geographic location and press **Enter**.
- To remove an address, select the address in the Addresses list and click the **X**.

Address lists can contain FQDNs, IP addresses, IP address ranges, geographic locations, other address lists, or any combination of these.

5. Click **Finished.**

The list screen and the new items are displayed.

About port lists

A *port list* is simply a collection of ports saved on the server. A port list can also contain other port lists. You can define one or more port lists, and you can specify one or more port lists in a firewall rule. Firewall port lists can be used in addition to inline ports, specified within a particular firewall rule or policy.

Creating a port list

Create a port list to apply to a firewall rule, in order to match ports.

1. On the Main tab, click **Security > Network Firewall > Port Lists.**

The Port Lists screen opens.

2. Click **Create to create a new port list.**

3. In the **Name and **Description** fields, type the name and an optional description.**

4. In the Ports area, add and remove ports.

- To add a single port, type the port number and press the **Enter** key.
- To add a contiguous range of ports, type the first port number, a dash, and the last port number, then press the **Enter** key.
- To add an existing port list to the current port list, start typing the name of the port list. A list of port lists that match the typed input appear on a list in the field. Select the port list you want to add, then press the **Enter** key.
- To remove a port, port range, or port list, select the entry in the Ports area and click the small **X** to the right of the entry.

5. Click **Finished.**

The list screen and the new items are displayed.

Network Firewall Schedules

About Network Firewall schedules

With a Network Firewall schedule, you can configure date ranges, days of the week, and time ranges for when a firewall rule is applied.

A schedule must be selected in a firewall rule or rule list, to apply to that firewall rule or rule list. The firewall rule or rule list must also be set to the Scheduled state.

When you configure a schedule for a rule list, the rules within the rule list can only be enabled when the rule list is enabled by the schedule. This means that even if the individual rules in a rule list have schedules, the rules are not enabled by their schedules unless the rule list is also enabled by the rule list schedule.

Warning: *You can not specify a schedule for a rule if the system is configured for manual rule compilation or deployment. You also cannot configure the system for manual rule compilation or deployment if there are existing rules on the system that use schedules.*

Task list

[Creating a schedule](#)

Creating a schedule

Create a schedule to define the times, dates, and days of the week for when a firewall rule is applied.

1. On the Main tab, click **Security > Network Firewall > Schedules**.
The Schedules screen opens.
2. Click **Create** to create a new firewall schedule.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. In the **Date Range** area, define the range of dates over which the schedule applies.
 - Select **Indefinite** to have the schedule apply immediately, and run indefinitely. This makes the schedule active until you change the date range, or delete the schedule.
 - Select **Until** to have the schedule apply immediately, and define an end date and ending time. This makes the schedule active now, and disables it when the end date and ending time is reached. Click in the field to choose an end date from a dialog box, and set the ending time with the sliders.
 - Select **After** to have the schedule apply after the specified date and starting time, and run indefinitely. This makes the schedule active starting on the selected date and time, until you change the start date, or delete the schedule. Click in the field to choose a start date from a dialog box, and set the starting time with the sliders.
 - Select **Between** to apply the schedule starting on the specified start date and starting time, and ending on the specified end date and ending time. Click in the fields to choose the start and end dates from a dialog box, and set the starting and ending time with the sliders.
5. In the Time Range area, define the times over which the firewall rule applies.
 - Select **All Day** to have the schedule apply all day, for every day specified in the date range.

- Select **Between** to apply the schedule starting at the specified time, and ending at the specified time each day. Select the start and end hours and minutes from the popup screen, or click **Now** to set the current time.

Note: Specify the hours according to a 24-hour clock. For example, you can specify 3:00 PM with the setting 15.

6. In the Days Valid area, select the days of the week when the schedule is valid. Select check boxes for days of the week when the rule applies, and clear check boxes for days of the week when the schedule does not apply.
7. Click **Finished**.
The list screen and the new items are displayed.

Network Firewall Inline Rule Editor

Using the inline firewall rule editor

The BIG-IP® Network Firewall uses rules to specify traffic handling actions. The inline rule editor provides an alternative way to create and edit rules within a policy, on a single page. The advantage to this type of rule editing is that it provides a simpler and more direct overview of both a rule and the entire policy. You can edit an inline rule for any context. The inline rule editor is available only from the **Active Rules** page.

When using the inline rule editor, the information presented in a firewall rule is simplified to the following categories:

Name

You must specify a name for the rule. You can also specify an optional description.

State

You can enable, disable, or schedule a firewall rule. These states govern whether the rule takes an action, does not take an action, or takes an action only during specific days and times.

Protocol

Specify a protocol to which the firewall rule applies. By default, the rule is TCP.

For ICMP or ICMPv6 protocols, you can specify one or more ICMP types and codes.

Source

A rule can include any number of sources, including IPv4 or IPv6 addresses, IPv4 or IPv6 address ranges, fully qualified domain names, geographic locations, VLANs, address lists, ports, port ranges, port lists, subscribers, subscriber groups, and address lists.

Destination

A rule can include any number of destinations, including IPv4 or IPv6 addresses, IPv4 or IPv6 address ranges, FQDNs, geographic locations, VLANs, address lists, ports, port ranges, port lists, and address lists.

Actions

Specifies an action that applies when traffic matches the rule. The standard rule actions apply (Accept, Drop, Reject, and Accept Decisively). In addition, you can set the rule to start an iRule when the firewall rule matches traffic, and apply timeouts from a service policy to traffic that matches the rule.

Send to Virtual

Specifies a virtual server to which to send traffic that matches the rule. This option is not available for rules that are already at the virtual server context. Traffic that is sent to a virtual server is then evaluated by DDoS rules and firewall rules on that virtual server instead of according to the original rule. Staged rules are also evaluated based on the destination virtual server instead of the originating rule.

Protocol Inspection Profile

Specifies a protocol inspection profile to associate with the firewall rule. Protocol inspection profiles can be configured to run multiple inspections across different protocols.

Classification Policy

Specifies a classification policy to associate with the firewall rule.

Logging

Specifies whether logging is enabled or disabled for the firewall rule.

Task list

Enabling the Network Firewall inline rule editor

Creating a rule with the inline editor

Editing a rule with the inline editor

Enabling the Network Firewall inline rule editor

Enable the inline rule editor to edit rules in place within policies.

***Note:** You can either edit rules with the inline editor or with the standard editor, but not both. You can switch back to the standard rule editor at any time.*

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. Next to **Inline Rule Editor**, select **Enabled**.
3. Click **Update**.
The inline firewall rule editor is enabled.

Creating a rule with the inline editor

The Network Firewall Inline Rule Editor option must be enabled to create a rule with the inline rule editor. If you are going to specify address lists, port lists, custom iRules[®], virtual servers, or service policies to use with this rule, you must create these before you edit the firewall rule, or add them to the rule at a later time.

You edit a Network Firewall policy rule to change course, destination, actions, order, or other items in a firewall rule.

***Note:** You cannot add rules (created with these steps) to a rule list at a later time. You must create rules for a rule list from within the rule list. Similarly, you cannot use the rules created in a policy to apply as inline rules in another context, although you can use rule lists in a policy rule.*

1. On the Main tab, click **Security > Network Firewall > Policies**.
The Policies screen opens.
2. Click the name of the network firewall policy to which you want to add rules. If you want to create a policy, click **Create**, name the policy, and click **Finished**.
3. Click **Add Rule** to add a firewall rule to the policy.
A blank rule appears in the policy.
4. In the **Name** column, type the name and an optional description in the fields.
5. In the **State** column, select the rule state.
 - Select **Enabled** to apply the firewall rule or rule list to the addresses and ports specified.
 - Select **Disabled** to set the firewall rule or rule list to not apply at all.
 - Select **Scheduled** to apply the firewall rule or rule list according to the selected schedule.
6. If you select **Scheduled**, from the **Schedule** list, select the schedule for the firewall policy rule.
This schedule is applied when the firewall policy rule state is set to **Scheduled**.

***Note:** You cannot save a scheduled rule when the firewall compilation or deployment mode is manual.*

7. In the **Protocol** column, select the protocol to which the firewall rule applies.

- Select **Any** to apply the firewall rule to any protocol.
- Select the protocol name to apply the rule to a single protocol.
- Select **Other** and type the port number if the protocol is not listed.

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the global or route domain context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

- 8.** In the **Source** field, begin typing to specify a source address.

As you type, options will appear that match your input. Select the source option you want to use when it appears, or press Return. You can add more addresses by typing in the field labeled **add new source**. A source address can be any of the following:

- Any address
- IPv4 or IPv6 address
- IPv4 or IPv6 address range
- FQDN
- Geographic location
- VLAN
- Address list
- Port
- Port range
- Port list
- Subscriber
- Subscriber group
- Address list

- 9.** In the **Destination** field, begin typing to specify a destination address.

As you type, options will appear that match your input. Select the destination option you want to use when it appears, or press Return. You can add more addresses by typing in the field labeled **add new destination**.

A destination address can be any of the following:

- Any address
- IPv4 or IPv6 address
- IPv4 or IPv6 address range
- FQDN
- Geographic location
- VLAN
- Address list
- Port
- Port range
- Port list
- Address list

- 10.** Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.

- 11.** When you select an iRule to start in a firewall rule, you select how frequently the iRule is started, for sampling purposes. The value you configure is one out of n times the iRule is triggered. For example, to trigger the iRule one out of every five times the rule matches a flow, set this field to 5. To trigger the rule every time the rule matches a flow, set this field to 1.

- Optionally, to send traffic matched by this rule to a specific virtual server, from the **Send to Virtual** list, select the virtual server.

Important: Traffic that is sent to a virtual is processed according to the DDoS rules and firewall rules on that virtual server, not according to the originating context.

- To apply custom timeouts or port misuse profiles to flows that match this rule, from the **Service Policy** field, specify a service policy.
- To apply a protocol inspection profile to check protocol inspection signatures against traffic that matches the rule, select a Protocol Inspection Profile.
- To apply a classification policy to traffic that matches the rule, select a Classification Policy.
- In the **Logging** column, check **Logging** to enable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
- Click **Commit Changes to System**.
The policy with the updated rule is displayed.

The new firewall rule is created and displayed on the firewall policy screen.

Editing a rule with the inline editor

The Network Firewall Inline Rule Editor option must be enabled to edit a rule with the inline rule editor. If you are going to specify address lists, port lists, custom iRules®, or service policies to use with this rule, you must create these before you edit the firewall rule, or add them at a later time.

Edit a network firewall rule to change source or destination components, the rule action, iRules, rule order, and other settings.

- On the Main tab, click **Security > Network Firewall > Policies**.
The Policies screen opens.
- Click the name of the network firewall policy to which you want to add rules.
- To reorder a rule in a policy, click and hold anywhere in the rule row, and drag the rule to a new position within the list.
- To quickly enable or disable a rule in a policy, click the check box next to the rule ID and click the **Enable** or **Disable** button, then click **Commit Changes to System**.
- In the **Description** field, type or change the optional description.
- In the **State** column, select the rule state.
 - Select **Enabled** to apply the firewall rule or rule list to the addresses and ports specified.
 - Select **Disabled** to set the firewall rule or rule list to not apply at all.
 - Select **Scheduled** to apply the firewall rule or rule list according to the selected schedule.
- If you select **Scheduled**, from the **Schedule** list, select the schedule for the firewall policy rule.
This schedule is applied when the firewall policy rule state is set to **Scheduled**.

Note: You cannot save a scheduled rule when the firewall compilation or deployment mode is manual.

- In the **Protocol** column, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
 - Select **Other** and type the port number if the protocol is not listed.

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a

self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the `global` or `route domain` context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

9. In the **Source** field, begin typing to specify a source address.

As you type, options will appear that match your input. Select the source option you want to use when it appears, or press Return. You can add more addresses by typing in the field labeled **add new source**. A source address can be any of the following:

- Any address
- IPv4 or IPv6 address
- IPv4 or IPv6 address range
- FQDN
- Geographic location
- VLAN
- Address list
- Port
- Port range
- Port list
- Subscriber
- Subscriber group
- Address list

10. In the **Destination** field, begin typing to specify a destination address.

As you type, options will appear that match your input. Select the destination option you want to use when it appears, or press Return. You can add more addresses by typing in the field labeled **add new destination**.

A destination address can be any of the following:

- Any address
- IPv4 or IPv6 address
- IPv4 or IPv6 address range
- FQDN
- Geographic location
- VLAN
- Address list
- Port
- Port range
- Port list
- Address list

11. Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.

12. When you select an iRule to start in a firewall rule, you select how frequently the iRule is started, for sampling purposes. The value you configure is one out of *n* times the iRule is triggered. For example, to trigger the iRule one out of every five times the rule matches a flow, set this field to 5. To trigger the rule every time the rule matches a flow, set this field to 1.

13. Optionally, to send traffic matched by this rule to a specific virtual server, from the **Send to Virtual** list, select the virtual server.

Important: Traffic that is sent to a virtual is processed according to the DDoS rules and firewall rules on that virtual server, not according to the originating context.

14. To apply custom timeouts or port misuse profiles to flows that match this rule, from the **Service Policy** field, specify a service policy.
 15. To apply a protocol inspection profile to check protocol inspection signatures against traffic that matches the rule, select a Protocol Inspection Profile.
 16. To apply a classification policy to traffic that matches the rule, select a Classification Policy.
 17. In the **Logging** column, check **Logging** to enable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
 18. Click **Commit Changes to System**.
The policy with the updated rule is displayed.
- The firewall rule is modified.

Configuring BIG-IP Network Firewall Policies

About firewall policies

The BIG-IP® Network Firewall policies combine one or more rules or rule lists, and apply them as a combined policy to one context. You can configure a context to use a specific firewall policy. However, firewall context precedence still applies, so policies applied at the global context still apply, even if they contradict rules applied at a lower precedence context. For example, global policies apply before virtual server policies.

Notice: Global firewall rules are included in an automatic policy called **Global**.

You can apply a Network Firewall policy as a staged policy, while enforcing an existing firewall policy, or no policy. A *staged policy* allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules.

Task list

Creating a Network Firewall policy

Setting a global firewall policy

Configuring a route domain with a firewall policy

Setting network firewall policies for a self IP address

Creating a virtual server with a firewall policy

Creating a Network Firewall policy

Create a BIG-IP® Network Firewall policy to collect and apply one or more firewall rules or rule lists globally, to a virtual server, route domain, or self IP address.

1. On the Main tab, click **Security > Network Firewall > Policies**.
The Policies screen opens.
2. Click **Create** to create a new policy.
3. Type a name and optional description for the firewall policy.
4. Click **Finished**.

The Policies screen shows the new policy in the policy list.

Define firewall rules and rule lists for the policy to affect traffic.

Creating a Network Firewall policy rule

If you are going to specify address lists or port lists to use with this rule, you must create these lists before creating the firewall policy rule, or add them after you save the policy rule.

You create a Network Firewall policy rule to manage access from an IP or web network address to a specified network location, server, or address behind a BIG-IP® system.

Note: You cannot add rules created with this task to a rule list at a later time. You must create rules for a rule list from within the rule list. Similarly, you cannot use the rules created in a policy to apply as inline rules in another context, although you can use rule lists in a policy rule.

1. On the Main tab, click **Security > Network Firewall > Policies**.
The Policies screen opens.
2. Click the name of the network firewall policy to which you want to add rules.
3. In the Rules area, click **Add** to add a firewall rule to the list.
4. In the **Name** and **Description** fields, type the name and an optional description.
5. From the **Type** list, select whether you are creating a standalone network firewall policy rule or creating a rule list.

*Note: If you create a firewall policy rule list, only the **Name**, **Description**, **Order**, **Rule List**, and **State** options apply, and you must select or create a rule list to include.*

6. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall policy rule or rule list to the addresses and ports specified.
 - Select **Disabled** to set the firewall policy rule or rule list to not apply at all.
 - Select **Scheduled** to apply the firewall policy rule or rule list according to the selected schedule.
7. If you select **Scheduled**, from the **Schedule** list, select the schedule for the firewall policy rule.
This schedule is applied when the firewall policy rule state is set to **Scheduled**.

Note: You cannot save a scheduled rule when the firewall compilation or deployment mode is manual.

8. From the **Protocol** list, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the global or route domain context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

9. In the **Source** list, specify addresses and geolocated sources to which this rule applies.
 - From the **Address/Region** list, select **Any** to have the rule apply to any packet source IP address or geographic location.
 - From the **Address/Region** list, select **Specify** and click **Address** to specify one or more packet source IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - From the **Address/Region** list, select **Specify** and click **Address List** to select a predefined list of packet source addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - From the **Address/Region** list, select **Specify** and click **Address Range** to specify a contiguous range of packet source IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
 - From the **Address/Region** list, select **Specify** and click **Country/Region** to identify the geographic origin of packet sources, and to apply rules based on selected geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Source address list.

10. From the Source **Port** list, select the type of packet source ports to which this rule applies.
 - Select **Any** to have the rule apply to any packet source port.
 - Select **Specify** and click **Port** to specify one or more packet source ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet source port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
 - Select **Specify** and click **Port List** to select a predefined list of packet source ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
11. From the Source **VLAN/Tunnel** list, select the VLAN on which this rule applies.
 - Select **Any** to have the rule apply to traffic on any VLAN through which traffic enters the firewall.
 - Select **Specify** to specify one or more VLANs on the firewall to which the rule applies. To use a VLAN with this rule, move the VLAN from the **Available** list to the **Selected** list. Similarly, you can remove the VLAN from this rule, by moving the VLAN from the **Selected** list to the **Available** list.
12. In the Destination area and from the **Address/Region** list, select the type of packet destination address to which this rule applies.
 - Select **Any** to have the rule apply to any IP packet destination address.
 - Select **Specify** and click **Address** to specify one or more packet destination IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - Select **Specify** and click **Address List** to select a predefined list of packet destination addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - Select **Specify** and click **Address Range** to specify a contiguous range of packet destination IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
 - Select **Specify** and click **Country/Region** to identify the geographic packet destination, and to apply rules based on specific geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Destination address list.
13. From the Destination **Port** list, select the type of packet destination ports to which this rule applies.
 - Select **Any** to have the rule apply to any port inside the firewall.
 - Select **Specify** and click **Port** to specify one or more packet destination ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet destination port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
 - Select **Specify** and click **Port List** to select a predefined list of packet destination ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
14. Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.
15. When you select an iRule to start in a firewall rule, you can enable iRule sampling, and select how frequently the iRule is started, for sampling purposes. The value you configure is one out of n times the iRule is triggered. For example, to trigger the iRule one out of every five times the rule matches a flow, select **Enabled**, then set this field to 5.

- Optionally, to send traffic matched by this rule to a specific virtual server, from the **Send to Virtual** list, select the virtual server.

Important: Traffic that is sent to a virtual is processed according to the DDoS rules and firewall rules on that virtual server, not according to the originating context.

- To apply a protocol inspection profile to check protocol inspection signatures against traffic that matches the rule, select a Protocol Inspection Profile.
- To apply a classification policy to traffic that matches the rule, select a Classification Policy.
- From the **Action** list, select the firewall action for traffic originating from the specified source address on the specified protocol. Choose from one of the these actions:

Option	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. When a packet is rejected the firewall sends a destination unreachable message to the sender.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall, and does not require any further processing by any of the further firewalls. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.

- From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
- Click **Finished**.
The list screen and the new items are displayed.
The new firewall policy rule is created.

Setting a global firewall policy

You can assign a firewall policy globally, to provide policy-based network firewall actions at the device level.

Note: To configure a clean-up rule for logging purposes, when IPv6 pool members are configured on the BIG-IP, you must add a specific rule that allows access from the pool member network to the BIG-IP. This rule must allow ICMPv6 protocol 136 (Neighbor Advertisement). This allows the IPv6 pool members to show neighbor table and monitor as available.

- On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
- Under **Active Network Firewall Rules**, click the **Global** link.
The **Global Firewall Rules** screen opens.
- To enforce rules from a firewall policy in the selected context, in the Network Firewall area: from the **Enforcement** list, select **Enabled** and then select the firewall policy to enforce from the **Policy** list.
- To stage rules from a firewall policy in the selected context, in the Network Firewall area: from the **Staging** list, select **Enabled** and then select the firewall policy to stage from the **Policy** list.

The policy rules you selected are enforced at the global level. If you chose to stage policy rules, the results of those rules are logged, but not enforced.

Configuring a route domain with a firewall policy

Before you can configure a route domain with a firewall policy, you need a pre-existing route domain.

Route domains are useful for multi-tenant configurations. You can set firewall policies for enforcement and staging on an existing route domain, and create a route domain on a BIG-IP® system to segment (isolate) traffic on your network.

***Note:** To configure a clean-up rule for logging purposes, when IPv6 pool members are configured on the BIG-IP, you must add a specific rule that allows access from the pool member network to the BIG-IP. This rule must allow ICMPv6 protocol 136 (Neighbor Advertisement). This allows the IPv6 pool members to show the neighbor table and monitor as available.*

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. Click the name of the route domain to show the route domain configuration.
3. Click the Security tab.
4. To enforce rules from a firewall policy on the route domain: in the Network Firewall area: from the **Enforcement** list, select **Enabled** and then select the firewall policy to enforce from the **Policy** list.
5. To stage rules from a firewall policy on the route domain: in the Network Firewall area, from the **Staging** list, select **Enabled** and then select the firewall policy to stage from the **Policy** list.
6. Click **Update** to save the changes to the route domain.

Setting network firewall policies for a self IP address

Ensure that you have created a self IP address.

You can enforce or stage a firewall policy at the self IP context. Stage a firewall policy to verify the results of the firewall policy in the logs without affecting traffic.

1. On the Main tab, click **Network > Self IPs**.
2. Click on the self IP address to which you want to add a network firewall policy.
3. Click the **Security** tab.
4. To enforce rules from a firewall policy on the self IP: In the Network Firewall area, from the **Enforcement** list, select **Enabled**, and then from the **Policy** list, select the firewall policy to enforce.
5. To stage rules from a firewall policy on the self IP: In the Network Firewall area, from the **Staging** list, select **Enabled**, and then from the **Policy** list, select the firewall policy to stage.
6. Click **Update** to save the changes to the self IP.

The selected self IP now enforces or stages rules according to your selections.

Creating a virtual server with a firewall policy

You can create a virtual server with a firewall policy, to provide policy-based network firewall actions at the virtual 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 you want to use for the virtual server.
The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type * or select * **All Ports** from the list.
6. Click **Finished**.
7. Click the name of the virtual server you want to modify.
8. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
9. To enforce rules from a firewall policy on the virtual server, in the Network Firewall area, from the **Enforcement** list, select **Enabled**, then select the firewall policy to enforce from the **Policy** list.
10. To stage rules from a firewall policy on the virtual server, in the Network Firewall area, from the **Staging** list, select **Enabled**, then select the firewall policy to stage from the **Policy** list.
11. Click **Update** to save the changes.

The policy rules you selected are enforced on the virtual server. If you chose to stage policy rules, the results of those rules are logged, but not enforced.

Viewing enforced and staged policy rule logs

With BIG-IP® Advanced Firewall Manager™, you can choose to enforce either inline firewall rules or a firewall policy for a specific context. You can also choose to stage policies for a specific context. *Staged policies* apply all of the specified firewall rules to the policy context, but do not enforce the firewall action. Therefore, the result of a staged policy is informational only, and the result can be analyzed in the firewall logs.

A staged policy on a particular context might not behave the same after you change it to an enforcement policy. Because there can be multiple staged policies on different contexts, the staged policy results you see (in logs and stats) are actually the aggregate of all staged policies on all contexts. Thus, if you enforce a previously staged policy on one or more contexts, but other staged policies remain on other contexts that you do not enforce, the actual enforced results might differ from what you expected from viewing logs and statistics for staged rules.

Important: *You must enable logging for a policy, if you want to view the results of staged or enforced rules in the logs.*

Viewing Network Firewall enforced policy events on the local BIG-IP system

Ensure that the BIG-IP® system is configured to log the types of events you want to view, and to store the log messages locally on the BIG-IP system.

When the BIG-IP system is configured to log events locally, you can view those events using the Configuration utility.

1. On the Main tab, click **Security > Event Logs > Network > Firewall**.
The Network Firewall event log displays.
2. To search for enforced policy events, in the search field, type `Enforced`, then click **Search**.
3. To narrow your search for enforced events, click **Custom Search**. Drag the `Enforced` text from the **Policy Type** column to the custom search table. Narrow your search further by dragging other items from the log display, for example, from the **action**, **policy**, or **rule** columns. the event data that you want to search for from the Event Log table into the Custom Search table, and then click **Search**.

Viewing Network Firewall staged policy events on the local BIG-IP system

Ensure that the BIG-IP® system is configured to log the types of events you want to view, and to store the log messages locally on the BIG-IP system.

When the BIG-IP system is configured to log events locally, you can view those events using the Configuration utility.

Important: *You must enable logging for a policy, if you want to view the results of staged or enforced rules in the logs.*

1. On the Main tab, click **Security > Event Logs > Network > Firewall**.
The Network Firewall event log displays.
2. To search for staged policy events, in the search field, type `Staged`, then click **Search**.
3. To narrow your search for staged policy events, click **Custom Search**. Drag the `Staged` text from the **Policy Type** column to the custom search table. Narrow your search further by dragging other items from the log display. For example, from the **action**, **policy**, or **rule** columns, you can drag event data that you want to search for from the Event Log table into the Custom Search table, and then click **Search**.

IP Address Intelligence in the Network Firewall

About IP intelligence policies in the Network Firewall

In the BIG-IP[®] Network Firewall, you can configure policies to validate traffic against an IP intelligence database. Such traffic can be handled automatically if it originates from known-bad or questionable IP addresses. To use existing lists of known bad IPs, you can configure policies to automatically query *feed lists* that specify blacklist and whitelist IP address entries, and assign default classes and blacklist or whitelist behaviors to those feed lists. In addition, you can manually add an IP address to a blacklist category, or remove an IP address from a blacklist category.

You can control the actions for each blacklist category by specifying those actions in a policy, and you can configure default action and default logging for each policy. Furthermore, you can configure logging and actions per category. You can apply IP Intelligence policies at the global context, to a virtual server, or on a route domain.

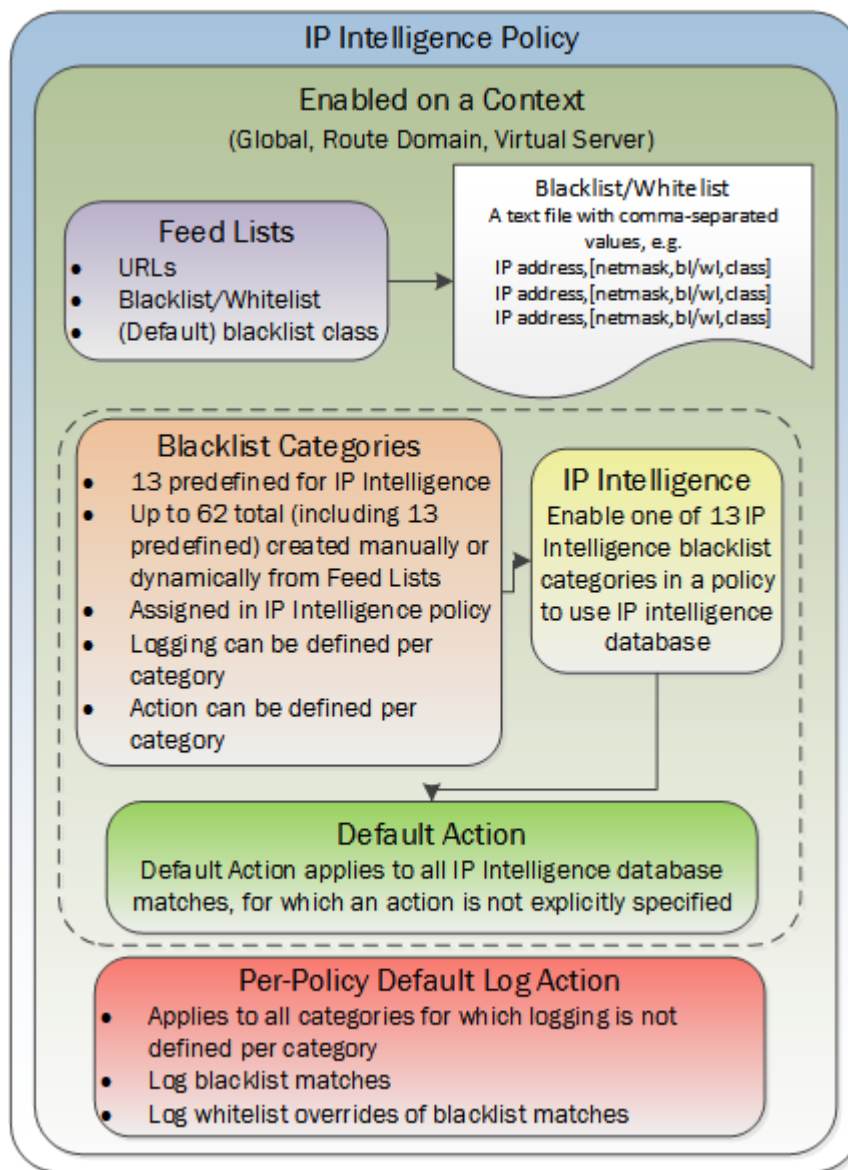


Figure 2: IP Intelligence Policy container and included elements

Downloading the IP intelligence database

Downloading the IP intelligence database

The requirements for using IP Intelligence are:

- The system must have an IP Intelligence license.
- The system must have an Internet connection either directly or through an HTTP proxy server.
- The system must have DNS configured (go to **System > Configuration > Device > DNS**).

Important: IP Intelligence is enabled by default if you have a license for it. You only need to enable it if it was previously disabled.

To enable IP Intelligence on the BIG-IP® system, you enable auto-update to download the IP intelligence database to the system.

1. Log in to the command line for the BIG-IP® system.

2. To determine whether IP intelligence auto-update is enabled, type the following command: `tmsh list sys db iprep.autoupdate`
If the value of the `iprep.autoupdate` variable is `disable`, IP intelligence is not enabled. If it is `enable`, your task is complete. No further steps are necessary.
3. If disabled, at the prompt, type `tmsh modify sys db iprep.autoupdate value enable`
The system downloads the IP intelligence database and stores it in the binary file, `/var/IpRep/F5IpRep.dat`. It is updated every 5 minutes.
4. If the BIG-IP system is behind a firewall, make sure that the BIG-IP system has external access to `vector.brightcloud.com` using port 443.
That is the IP Intelligence server from which the system gets IP Intelligence information.
5. (Optional) If the BIG-IP system connects to the Internet using a forward proxy server, set these system database variables.
 - a) Type `tmsh modify sys db proxy.host value hostname` to specify the host name of the proxy server.
 - b) Type `tmsh modify sys db proxy.port value port_number` to specify the port number of the proxy server.
 - c) Type `tmsh modify sys db proxy.username value username` to specify the user name to log in to the proxy server.
 - d) Type `tmsh modify sys db proxy.password value password` to specify the password to log in to the proxy server.

The IP Intelligence feature remains enabled unless you disable it with the command `tmsh modify sys db iprep.autoupdate value disable`.

You can configure IP intelligence for Advanced Firewall Manager by assigning IP intelligence policies to the global, route domain, or virtual server context.

About IP intelligence blacklist categories

Blacklist categories are categories you can use to differentiate between types of blacklisted URLs. You can specify up to 62 blacklist categories, including 16 that are predefined on the system. A blacklist category definition consists only of a name and description. You can specify actions and logging options for each blacklist category you create, and for predefined categories, in an IP Intelligence policy. The 16 predefined blacklist categories are automatically available for selection in an IP Intelligence policy.

Blacklist categories

Along with the IP address, the IP intelligence database stores the category that explains the reason that the IP address is considered untrustworthy.

Category Name	Description
additional	IP addresses that are added from additional categories not more explicitly defined.
appiq_badactors	IP addresses gathered from AppIQ central management.
application_denial_of_service	IP addresses involved in application DoS Attacks, or anomalous traffic detection.
attacked_ips	Destination IP addresses under attack.
botnets	IP addresses of computers that are infected with malicious software (Botnet Command and Control channels, and infected zombie machines) and are controlled as a group by a Bot master, and are now part of a botnet. Hackers can exploit botnets to send

Category Name	Description
	spam messages, launch various attacks, or cause target systems to behave in other unpredictable ways.
cloud_provider_networks	IP addresses and networks that are used by cloud providers.
denial_of_service	IP addresses that have launched denial-of-service (DoS) attacks, distributed denial-of-service (DDoS) attacks, anomalous SYN flood attacks, or anomalous traffic detection. These attacks are usually requests for legitimate services, but occur at such a fast rate that targeted systems cannot respond quickly enough and become bogged down or unable to service legitimate clients.
infected_sources	Active IP addresses that issue HTTP requests with a low reputation index score, or that are known malicious web sites offering or distributing malware, shell code, rootkits, worms, or viruses.
mobile_threats	IP addresses of malicious and unwanted mobile applications.
phishing	IP addresses that host phishing sites, and other kinds of fraud activities, such as ad click fraud or gaming fraud.
proxy	IP addresses that are associated with web proxies that shield the originator's IP address (such as proxy and anonymization services).
scanners	IP addresses that are involved in reconnaissance, such as probes, host scan, domain scan, and password brute force, typically to identify vulnerabilities for later exploits.
spam_sources	IP addresses tunneling spam messages through proxy, anomalous SMTP activities and forum spam activities.
tor_proxy	IP addresses acting as exit nodes for the Tor Network. Exit nodes are the last point along the proxy chain and make a direct connection to the originator's intended destination.
web_attacks	IP addresses involved in cross site scripting, iFrame injection, SQL injection, cross domain injection, or domain password brute force.
windows_exploits	Active IP addresses that have exercised various exploits against Windows resources by offering or distributing malware, shell code, rootkits, worms, or viruses using browsers, programs, downloaded files, scripts, or operating system vulnerabilities.

Creating a blacklist category

You can create a blacklist category to configure policy-based responses to specific types of addresses. Then you can specify an address as belonging to a blacklist category so you can see the types of categories that are triggered in the logs, and so you can provide unique responses on a per-category basis.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Blacklist Categories**. The Blacklist Categories screen opens.
2. Click **Create** to create a new IP Intelligence blacklist category.
3. In the **Name** field, type a name for the blacklist category.
4. In the **Description** field, type a description for the blacklist category.
5. Select the **Match Type** for the IP intelligence category.

By default, IP intelligence blacklist categories match **Source** only, but you can configure categories to match **Source and Destination** or **Destination** only.

6. Click **Finished**.
The list screen and the new items are displayed.

Blacklisting an individual IP address

You can easily blacklist a single IP address manually. You do this by adding the IP address directly to a blacklist category. The settings for the blacklist category, as defined in an IP intelligence policy, are then applied to the IP address.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Blacklist Categories**.
The Blacklist Categories screen opens.
2. Select the check box next to a blacklist category.
You can select more than one blacklist category.
3. Click the **Add to Category** button.
The **Add Entry** popup screen appears.
4. In the **Insert (IP Address)** field, type an IP address to add to the blacklist category or categories.
5. In the **Seconds** field, specify the duration for which the address should be added to the blacklist category.
6. To allow the IP address to be advertised to edge routers so they will null route the traffic, select **Allow Advertisements**.
7. Click the **Add Address** button.
The IP address is added to the blacklist category or categories.

Removing an individual IP address from a blacklist

You can easily remove single IP address from a blacklist manually. You do this by selecting the blacklist category, and removing the IP address.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Blacklist Categories**.
The Blacklist Categories screen opens.
2. Select the check box next to a blacklist category.
You can select more than one blacklist category.
3. Click the **Delete from Category** button.
The **Delete Entry** popup screen appears.
4. In the **Delete (IP Address)** field, type an IP address to remove from the selected blacklist category or categories.
The **Delete Entry** popup screen appears.
5. In the **Insert (IP Address)** field, type an IP address to add to the blacklist category or categories.
6. Click the **Delete Address** button.
The IP address is removed from the blacklist category or categories.

About IP intelligence feed lists

A *feed list* retrieves blacklists and whitelists from specified URLs. You can use a feed list to dynamically update blacklists and whitelists.

A feed list can retrieve multiple feeds from FTP, HTTP, or HTTPS addresses. You can specify whether a feed is a blacklist or whitelist, and the default category for the feed list. You can also configure a polling interval.

After a blacklist or whitelist is defined in a feed list, you add the feed list to an IP Intelligence policy. The list is then used by the policy to retrieve feeds and dynamically adjust the blacklist and whitelist policy.

Feed list settings

Feed lists dynamically define IP addresses that have been blacklisted or whitelisted. The IP Intelligence policy uses feed lists to dynamically filter traffic.

A feed list defines the feeds that dynamically update the IP address intelligence database for your system.

Feed list setting	Description
URL	Select FTP , HTTP , or HTTPS , then specify the URL for the feed. Feeds are typically text files. An example for a local file might be <code>http://172.10.1.23/feed.txt</code> .
List Type	Whitelist or Blacklist . Specifies the default classification for all URLs in the feed for which a category is not specified.
Blacklist Category	Specifies a default category for the list. This is the default blacklist category for all blacklist URLs in the feed for which a category is not specified. On the BIG-IP [®] system, you can specify a total of 62 categories; however, 9 categories are used by the IP Intelligence database.
Poll Interval	Specifies how often the feed URL is polled for new feeds.
Username	The user name to access the feed list file, if required.
Password	The password to access the feed list file, if required.
Feed URLs	In this area, you can add, replace, or delete feed URLs from the feed list.

A feed is a simple comma-separated value (CSV) file. The file contains four comma-separated values per line.

Position	Value	Definition
1	IP Address	The IP address to be blacklisted or whitelisted. This is the only field that is required in each entry in the file. All other entries are optional. <i>Important: If you append a route domain with a percentage sign and the route domain number, the route domain is not used.</i>
2	Network Mask	(Optional) The network mask for the IP address as a CIDR (such as, 24 for 255.255.255.0). This field is optional.
3	Whitelist/Blacklist	(Optional) Identifies whether the IP address is a whitelist or blacklist address. You can type <code>wl</code> , <code>bl</code> , <code>whitelist</code> , or <code>blacklist</code> , with any capitalization. Leave this field blank to retain the default specified for the feed.

Position	Value	Definition
4	Category	(Optional) Type the category name for the entry. Leave this field blank to retain the default specified for the feed.

In this feed list file example, only the first entry specifies a value for every field. The third and fourth entries, 10.10.0.12 and 10.0.0.12, will be set to blacklist or whitelist entries depending on the setting for the feed. 10.10.0.12 is specified with a category of `botnets`; however, if the default setting for the feed is a whitelist, this is ignored. When an IP address has both a blacklist and a whitelist entry from the configuration, the whitelist entry takes precedence. The more specific entry takes precedence, so if an entry in the feed list file specifies a setting, that setting overrules the default setting for the feed list or category.

```
10.0.0.2,32,bl,spam_sources
10.0.0.3,,wl,
10.10.0.12,,botnets
10.0.0.12,,,
10.0.0.13,,bl,
```

Creating a feed list

You can add whitelist and blacklist IP addresses to your configuration automatically by setting up feeds and capturing them with a feed list.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Feed Lists**.
The Feed Lists screen opens.
2. Click **Create** to create a new IP Intelligence feed list.
3. In the **Name** field, type a name for the feed list.
4. Configure Feed URLs with an HTTP, HTTPS, or FTP URL, the list type, the blacklist category, and the polling interval. Specify a user name and password, if required to access the feed list.
A feed URL includes the actual URL to the text file, and information about the defaults for that file. Within the feed file, however, any URL can be configured to be a whitelist or blacklist entry, and assigned to a blacklist category.
5. Click the **Add** button to add a feed URL to the feed list.
6. Click **Finished**.
The list screen and the new items are displayed.

Configuring and assigning IP intelligence policies

An IP intelligence policy combines combines feed lists, default actions, logging settings, and actions for blacklist categories into a container that you can apply to a virtual server or route domain.

Configuring a policy to check addresses against IP intelligence

You can verify IP addresses against the preconfigured IP Intelligence database, and against IPs from your own feed lists, by creating an IP Intelligence policy.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Policies**.
The IP Intelligence Policies screen opens.
2. Click **Create** to create a new IP Intelligence policy.

3. In the **Name** field, type a name for the IP intelligence policy.
4. To add feed lists to the policy, click the name of an **Available** feed list, and then add it to the **Selected** list.
5. Set the **Default Action** for the policy to Accept or Drop.
 - Select **Accept** to allow packets from categorized addresses that have no action applied on the feed list.
 - Select **Drop** to drop packets from categorized addresses that have no action applied on the feed list.

The default action applies to addresses that are not assigned a blacklist category in the feed list. The IP Intelligence feature uses the action specified in a feed list entry, when available.

6. Set the **Default Log Actions**.
 - **Log Whitelist Overrides** specifies the default log action (**Yes** or **No**) for whitelist matches that override blacklist matches.
 - **Log Blacklist Category Matches** specifies the default log action for blacklist matches, when the category does not have a specified log action in the blacklist matching policy. **Yes** affords visibility into blacklist matches and logs all packets, but provides no hardware acceleration. **Limited** logs statistics every 256 packets and includes hardware acceleration. **No** does not log blacklist matches but provides the highest performance with hardware acceleration.
 - Select both **Log Black List Category Matches** and **Log White List Overrides** to log all black list matches, and all whitelist matches that override blacklist matches.

Note: Whitelist matches always override blacklist matches.

7. To configure matching actions and logging for custom blacklist categories, select the Blacklist Category in the **Blacklist Matching Policy** area. Select a category from the list of predefined and user-defined blacklist categories, and set the default action and default logging action for the category, then click **Add** to add the blacklist category to the policy.
8. In the **Blacklist Matching Policy** area, select default actions for the selected blacklist category.
 - a) For **Blacklist Category**, select the category for which to configure the settings.
 - b) For **Action**, select **Use Policy Default** to use the default policy action for this category. Select **Drop** or **Accept** to override the default action.
 - c) For **Log Blacklist Category Matches**, specify the action for blacklist matches for this category. Select **Use Policy Default** to use the default log action for blacklist matches. **Yes** affords visibility into blacklist matches and logs all packets, but provides no hardware acceleration. **Limited** logs statistics every 256 packets and includes hardware acceleration. **No** does not log blacklist matches but provides the highest performance with hardware acceleration.
 - d) For **Match Override**, select the matching criteria that overrides a blacklist match. You can require a source match, a destination match, or both a source and destination match to override a blacklist match with a whitelist.
 - e) Click **Add** to add the customized blacklist matching policy to the IP intelligence policy. The matching policy for the blacklist category is added to the list.
 - f) To replace a blacklist category matching policy, select the entry and click **Replace**.

When you define a blacklist category and add it to the IP intelligence policy, matches for that category use the default actions specified for the category, and override the policy defaults.

9. To remove a customized category from the policy, select the category in the Blacklist Matching Policy area and click **Delete**.
10. Click **Finished**.

The list screen and the new items are displayed.

Assigning a global IP Intelligence policy

You can assign an IP Intelligence policy globally, to apply blacklist and whitelist matching actions and logging to all traffic.

1. On the Main tab, click **Security > Network Firewall > IP Intelligence > Policies**.
The IP Intelligence Policies screen opens.
2. From the **Global Policy** list, select the IP Intelligence policy to apply to all traffic on the BIG-IP system.
3. Click **Update**.
The list screen and the updated item are displayed.

The specified IP Intelligence policy is applied to all traffic.

Assigning an IP Intelligence policy to a virtual server

You can assign an IP Intelligence policy to a virtual server, to apply blacklist and whitelist matching actions and logging to traffic on that virtual server only.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, from the Security menu, choose Policies.
4. Next to **IP Intelligence**, select **Enabled**, then select the IP intelligence policy to apply to traffic on the virtual server.
5. Click **Update**.
The list screen and the updated item are displayed.

The specified IP Intelligence policy is applied to traffic on the selected virtual server.

Assigning an IP Intelligence policy to a route domain

You can assign an IP Intelligence policy to a route domain, to apply blacklist and whitelist matching actions and logging to route domain traffic.

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. In the Name column, click the name of the relevant route domain.
3. From the **IP Intelligence Policy** list, select an IP Intelligence policy to enforce on this route domain.
4. Click **Update**.
The system displays the list of route domains on the BIG-IP system.

The specified IP Intelligence policy is applied to traffic on the route domain.

Inspecting Protocol Anomalies

About protocol anomaly inspection

In the BIG-IP® Network Firewall, you can configure profiles to inspect traffic against protocol inspection items. Protocol inspection items are arranged in categories by the Service type. You can assign protocol inspection items individually or in groups. You can add a new inspection item by writing a valid Snort rule and defining matching characteristics. You can assign protocol inspection items to a firewall rule, or directly to a virtual server.

Rule precedence applies to protocol inspection profiles. The protocol inspection rules for the most granular context are applied. The only exception is that a virtual server firewall rule takes precedence over a profile applied directly to a virtual server. The order of precedence is:

1. Profile applied to a virtual server firewall rule
2. Profile applied directly to a virtual server
3. Profile applied to a route domain
4. Profile applied to the global context

Task list

Creating a protocol inspection profile

Viewing protocol inspection items

Creating protocol inspection items

Assigning a protocol inspection profile to a virtual server

Assigning a protocol inspection profile to a firewall rule

Viewing protocol security inspection logs

Creating a protocol inspection profile

A protocol inspection profile collects rules for protocol inspection using preinstalled and user-defined inspection items, or signatures. Preinstalled signatures are defined by the Snort project, and new signatures can be defined using Snort syntax. Signatures are selected and added to the profile by Service, and you can narrow the scope of signatures by a number of other characteristics. You can enforce signatures, compliance items, or both.

1. On the Main tab, click **Security > Protocol Security > Inspection Profiles**.
The Inspection Profiles screen opens.
2. Click **Add** and select **New**.
Alternatively, copy an existing inspection profile by selecting the profile and clicking **Add**, then **Clone Existing**.
3. Type a profile name, and optionally add a description.
4. From the **Signatures** menu, select **Enabled** to enforce signatures.
If you are enforcing only Signature items, you can select **Disabled** for compliance items.
5. From the **Compliance** menu, select **Enabled** to enforce compliance items.
If you are enforcing only Compliance items, you can select **Disabled** for signatures.
6. To collect AVR stats, from the **AVR Stats Collect** menu, select **Enabled**.
7. From the Services menu, select the services you want to add to the inspection profile.

Each selected service type displays as a new category on the screen. By default, all inspection items are disabled. You must enable items or categories you want to inspect.

8. To enable inspections in the service, click the service category name on the screen. The service category expands to show the inspections.
9. To enable an inspection, select the checkbox for the inspection. The Edit Selected Inspections panel opens on the right of the screen.
10. To enable an inspection, select **Enable**, and click **Apply**.
11. To change the action for the selected inspection, from the Action menu select **Accept**, **Reject**, or **Drop**.
12. To select whether the inspection item is logged, from the Log menu select **Yes** or **No**.

Note: You can select and edit multiple inspections at once. You can select the checkbox at the top of the category to select and edit all inspections in the category.

13. When you have finished adding services and editing inspections, click **Commit Changes to System**. The Inspection Profiles screen appears and the inspection profile you created is displayed in the list. You can attach a protocol inspection profile to a firewall rule or to a virtual server.

Viewing protocol inspection items

View the list of protocol inspection items to see the checks and actions for a category, and to view properties and hits for an inspection item.

1. On the Main tab, click **Security > Protocol Security > Inspection List**. The Inspection List screen opens.
2. To filter the list, select the filter options.

You can type text in the search field to narrow options, and select options from any of the lists. You can add filter lists from the **Add Filter** list. The available category lists are **Service**, **Protocol**, **Inspection Type**, **Direction**, **Risk**, **Accuracy**, **Performance Impact**, **User Defined**, **Action**, and **Log**. You can select multiple items from each list, and you can make selections from multiple lists. For example, to filter items of medium and low accuracy that include the term `blacklist`, select **medium** and **low** from the **Accuracy** list, and type `blacklist` in the search field. The list of inspection items changes to show the results of the filter settings.
3. To view the properties for an inspection item, click the item ID number. The item properties are displayed in the Properties pane on the right.
4. To expand statistics for an inspection item, click the item description. Statistics for hits on the inspection item appear below the item.
5. To change the duration for statistics, click **1 Hour**, **1 Day**, or **1 Week**. You can click **Refresh** to refresh the statistics.
6. To collapse the statistics, click the description again. You can expand multiple items on the same screen.

Creating protocol inspection items

Add inspection items to create new inspections based on Snort signatures. You write signatures in Snort format. For information on writing Snort rules, see <https://www.snort.org/documents>.

1. On the Main tab, click **Security > Protocol Security > Inspection List**. The Inspection List screen opens.
2. Click **New Signature**.
3. In the **Name** field, type a name for the signature.

4. In the **Description** field, type a description.
5. In the **Signature Definition** field, type the valid snort syntax.

***Note:** All remaining fields are optional. However, the default settings accept the signature, and may not be configured correctly for your inspection. Configure settings that are appropriate to your security stance and the detection you want to accomplish.*

6. Specify an action for the signature.
7. Select whether to log the signature.
8. Specify the accuracy for the signature.
9. Specify the direction on which the signature is detected.
10. Specify the performance impact for the signature.
11. Specify the protocol on which the signature acts.
12. Specify the risk level for the attack.
13. In the **Documentation** field, type any documentation for the signature.
14. In the **Attack Type** field, specify the attack type.
15. In the **References** field, type any references for the signature.
16. In the **Reference Links** field, type any reference links.
17. In the **Revision** field, type the revision number.
18. In the **Systems** field, type the systems affected by the signature.
19. Specify the service to which the signature applies.
20. Click **Create** to create the inspection item.

The signature is created and appears in the inspection list.

Assign the inspection item to an inspection profile to enable detection and the action associated with the inspection item.

***Tip:** To view user defined inspection items, you can select **yes** from the **User Defined** list on the **Inspection Profile** or **Inspection List** screens.*

Snort rule reference

This document includes the Snort commands that are currently supported when writing Snort rules.

Snort rule overview

Protocol Anomaly Inspection supports a subset of Snort rules. See the *Snort users manual* for more information. Snort rules can be written as pcre (perl-compatible regular expressions). Negotiation (!) is not supported.

Parameters supported with content and pcre

The following parameters are supported when using the `content` and `pcre` commands. See *content* and *pcre*.

- nocase
- depth
- offset
- distance
- within
- http_client_body
- http_cookie
- http_header

- http_method
- http_uri
- http_stat_code
- http_stat_msg
- fast_pattern

Parameters supported with `byte_test`

All parameters for `byte_test` are supported except `dce` and `bitmask`. See the *byte_test*.

Parameters supported with `byte_jump`

All parameters for `byte_jump` are supported except `dce`, `multiplier`, `align`, `post_offset`, and `bitmask`. See *byte_jump*.

Parameters supported in metadata

The following parameters are supported in `metadata`. See *metadata*.

- service
- policy balanced-ips

The following parameters are supported in `reference`. See *reference*.

- url
- cve
- bugtraq

The following additional commands are supported.

- msg
- classtype
- flow
- rev

The following parameters are added:

- protocol
- accuracy
- risk
- systems
- documentation
- last_updated
- performance_impact

Assigning a protocol inspection profile to a virtual server

Add protocol inspection to a virtual server to detect the configured protocol inspection items on matched traffic on the virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. In the Name column, click the name of the relevant virtual server.
This displays the properties of the virtual server.
3. On the menu bar, from the Security menu, choose Policies.
4. For the Protocol Inspection Profile setting, Select **Enabled**.
5. Select the name of the protocol inspection profile to apply to traffic on the virtual server.

6. Click **Update**.

The protocol inspection profile is enabled on the virtual server.

Assigning a protocol inspection profile to a firewall rule

This task requires an existing network firewall policy.

Assign protocol inspection to a firewall rule to check protocol inspection items on traffic that matches the rule.

1. On the Main tab, click **Security > Network Firewall > Policies**.
The Policies screen opens.
2. Click the name of a firewall policy to edit that policy.
The Firewall Policy screen opens, or the policy expands on the screen.
3. Click **Add Rule** to add a firewall rule to the policy.
A blank rule appears in the policy.
4. In the **Name** column, type the name and an optional description in the fields.
5. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall rule to the given context and addresses.
 - Select **Disabled** to set the firewall rule to not apply at all.
 - Select **Scheduled** to apply the firewall rule according to the selected schedule.
6. From the **Protocol** list, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.

Important: ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the global or route domain context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.

7. In the **Source** field, specify the addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, a geographic location, a subscriber or subscriber group, an address list or port list. After you complete an entry, click **Add**.
8. In the **Destination** field, begin typing to specify a destination address.
As you type, options will appear that match your input. Select the destination option you want to use when it appears, or press Return. You can add more addresses by typing in the field labeled **add new destination**.
A destination address can be any of the following:
 - Any address
 - IPv4 or IPv6 address
 - IPv4 or IPv6 address range
 - FQDN
 - Geographic location
 - VLAN
 - Address list
 - Port

- Port range
- Port list
- Address list

9. From the **Action** list, select the firewall action for traffic originating from the specified source address on the specified protocol. Choose from one of the these actions:

Option	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. When a packet is rejected the firewall sends a destination unreachable message to the sender.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall, and does not require any further processing by any of the further firewalls. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.

10. Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.

11. Optionally, to send traffic matched by this rule to a specific virtual server, from the **Send to Virtual** list, select the virtual server.

***Important:** Traffic that is sent to a virtual is processed according to the DDoS rules and firewall rules on that virtual server, not according to the originating context.*

12. To apply custom timeouts or port misuse profiles to flows that match this rule, from the **Service Policy** field, specify a service policy.

13. To apply a protocol inspection profile to check protocol inspection signatures against traffic that matches the rule, select a Protocol Inspection Profile.

14. To apply a classification policy to traffic that matches the rule, select a Classification Policy.

15. Click **Done Editing**.

16. Click **Commit Changes to System**.

The policy with the updated rule is displayed.

A firewall rule is created with a Protocol Inspection Profile attached.

Viewing protocol security inspection logs

View protocol security inspection logs to check for hits on protocol inspection items.

1. On the Main tab, click **Security > Protocol Security > Inspection Logs**.

The Inspection Logs screen opens.

2. Use the search filters to narrow or expand your search.

For example, you can change the search time range from the default **Last Hour** to see logs for the last week, by selecting **Last Week**. Type text to match in the text field.

3. To customize the search, click **Custom Search**.

With **Custom Search**, you can drag search results directly from the results list to the search table, to provide a narrowed list of results.

4. Click **Search** to search with the custom search parameters.

5. Click **Reset Search** to return to the main search screen.

The search results for the protocol inspection log entries are displayed.

Deploying the BIG-IP Network Firewall in ADC Mode

About deploying the network firewall in ADC mode

The BIG-IP® Network Firewall provides policy-based access control to and from address and port pairs inside and outside of your network. By default, the network firewall is configured in ADC mode, which is a default allow configuration, in which all traffic is allowed through the firewall, and any traffic you want to block must be explicitly specified.

To understand this firewall scenario, imagine that your prerequisite system load-balances all traffic from the Internet to several internal servers. The internal servers are:

Device and location	IP address	Traffic type
Externally accessible FTP server	70.168.15.104	FTP
Application virtual server	192.168.15.101	HTTP, FTP
Server on internal network	10.10.1.10	HTTP, HTTPS
Server on internal network	10.10.1.11	HTTP, HTTPS

The system does not have a separate route domain configured, however you can use Route Domain 0, which is essentially the same as a global rule.

In order for traffic from the internal application virtual server to reach the external network virtual server, you must create a VLAN and enable both internal and external virtual servers on it. In this scenario, these VLANs are specified:

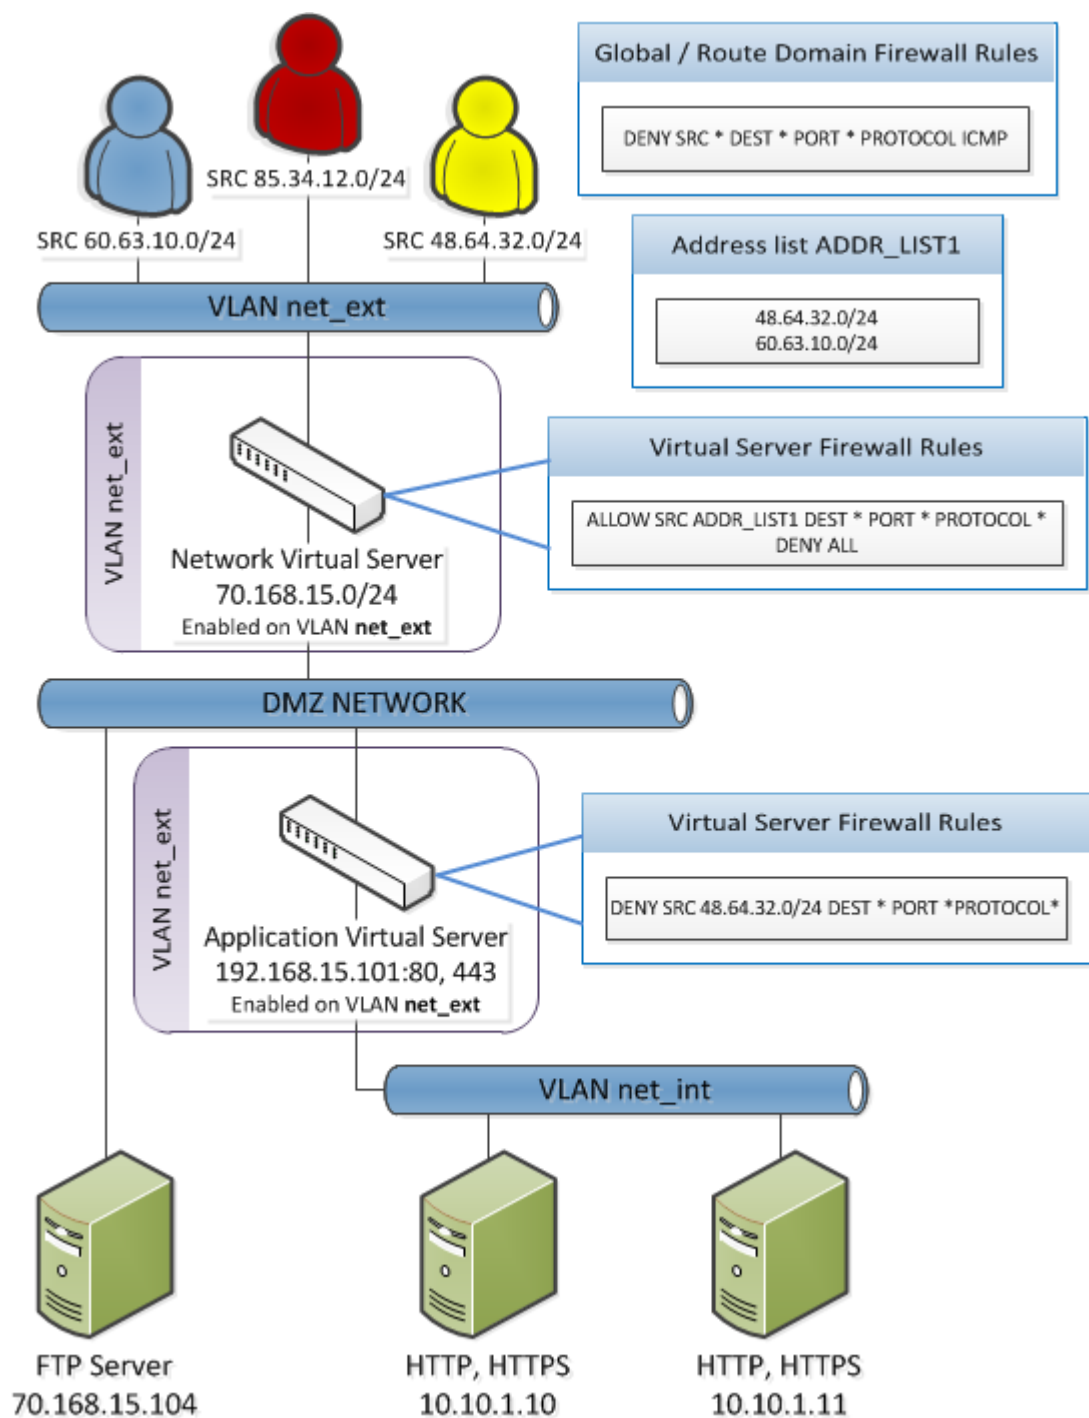
VLAN	Configuration
net_ext	Enabled on 70.168.15.0/24, 192.168.15.101
net_int	Includes pool members 10.10.1.10, 10.10.1.11

In addition, in this firewall configuration, there are three external networks that must be firewalled:

Network	Policy
60.63.10.0/24	Allow all access
85.34.12.0/24	Deny all access
48.64.32.0/24	Allow FTP, deny HTTP and HTTPS

To set up this scenario, you configure addresses, ports, and firewall rules specific to these networks, ports, and addresses. You will also configure a firewall rule that denies all ICMP traffic, to prevent pinging of network devices.

Figure 3: Firewall in ADC mode configuration scenario



Configuration settings for IPv6 pools and ADC mode

In a standard configuration, IPv6 pools work with either ADC mode or Firewall mode without any issues. However, in the specific ADC mode configuration where a Deny All policy is added after any specific Allow rules are configured, IPv6 pools cannot be reached.

If you choose to add a rule to deny all traffic after more specific rules, you must add a preceding rule with the following parameters.

- State: Enabled

- Protocol: ICMPv6 (58)
- Type: Neighbor Advertisement (136)
- Source Address: any affected pool members
- Destination Address: the BIG-IP address, or Any
- Action: Accept
- All other values can be left at their defaults, except the rule name.

Such a rule allows ICMPv6 pools to function, when a rule that denies all traffic is added at the end of the rule list in an ADC mode configuration.

Configuring the Network Firewall in ADC mode

If you have changed the firewall setting to Firewall mode, you can configure the BIG-IP® Network Firewall back to ADC mode.

Note: The firewall is configured in ADC mode, by default.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Virtual Server & Self IP Contexts** list, select the default action **Accept** for the self IP and virtual server contexts.
3. Click **Update**.
The virtual server and self IP contexts for the firewall are changed.

Creating a VLAN for the network firewall

Create a VLAN with tagged interfaces, so that each of the specified interfaces can process traffic destined for that VLAN.

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.
For purposes of this implementation, name the VLAN `net_ext`.
4. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.
5. 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.
6. From the **Configuration** list, select **Advanced**.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. For the **Hardware SYN Cookie** setting, select or clear the check box.
When you enable this setting, the BIG-IP system triggers hardware SYN cookie protection for this VLAN.
Enabling this setting causes additional settings to appear. These settings appear on specific BIG-IP platforms only.
9. For the **Syncache Threshold** setting, retain the default value or change it to suit your needs.

The **Syncache Threshold** value represents the number of outstanding SYN flood packets on the VLAN that will trigger the hardware SYN cookie protection feature.

When the **Hardware SYN Cookie** setting is enabled, the BIG-IP system triggers SYN cookie protection in either of these cases, whichever occurs first:

- The number of TCP half-open connections defined in the LTM[®] setting **Global SYN Check Threshold** is reached.
- The number of SYN flood packets defined in this **Syncache Threshold** setting is reached.

10. For the **SYN Flood Rate Limit** setting, retain the default value or change it to suit your needs.

The **SYN Flood Rate Limit** value represents the maximum number of SYN flood packets per second received on this VLAN before the BIG-IP system triggers hardware SYN cookie protection for the VLAN.

11. Click **Finished**.

The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Enable the new VLAN on both the network virtual server and the application virtual server.

Configuring an LTM virtual server with a VLAN for Network Firewall

For this implementation, at least two virtual servers and one at least one VLAN are assumed, though your configuration might be different.

You enable two virtual servers on the same VLAN to allow traffic from hosts on one virtual server to reach or pass through the other. In the Network Firewall, if you are using multiple virtual servers to allow or deny traffic to and from specific hosts behind different virtual servers, you must enable those virtual servers on the same VLAN.

***Tip:** By default, the virtual server is set to share traffic on **All VLANs and Tunnels**. This configuration will work for your VLANs, but in the firewall context specifying or limiting VLANs that can share traffic provides greater security.*

- 1.** On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
- 2.** Click the name of the virtual server you want to modify.
- 3.** From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
- 4.** Click **Update** to save the changes.
- 5.** Repeat this task for all virtual servers that must share traffic over the VLAN.

The virtual servers on which you enabled the same VLAN can now pass traffic.

Adding a firewall rule to deny ICMP

Use this task to create a firewall rule at the Global context, that denies ICMP packets globally.

- 1.** On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
- 2.** In the Rules area, click **Add** to add a firewall rule to the list.
- 3.** From the **Context** list, select the **Global** context.
- 4.** In the **Name** field, type **deny_icmp**.

5. From the **Type** list, select **Rule**.
6. From the **State** list, select **Enabled**.
7. From the Protocol list, select **ICMP**.
8. In the **ICMP Message** area, from the **Type** list, select `ANY`, and click the **Add** button.

Tip: You can optionally deny only ICMP ping requests, by selecting `Echo (8)` from the **Type** list, and clicking **Add**.

9. Leave the **Source** area configured to allow **Any** address, port, and VLAN.
10. Leave the **Destination** area configured to allow **Any** address or port.
11. From the **Action** list, select **Drop** or **Reject**.
These options either drop ICMP packets from any source and port to any port and address, or send a reject message and reset the the connection.
12. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
13. Click **Finished**.
The list screen and the new items are displayed.

A new firewall rule is created, and appears in the firewall rule list. This firewall rule denies all access to and from all sources and destinations on the ICMP protocol.

Creating an address list

Use this procedure to specify the address list to apply to allow access to specific source addresses.

1. On the Main tab, click **Security > Network Firewall > Address Lists**.
The Address Lists screen opens.
2. Click **Create** to create a new address list.
3. In the name field, type `ADDR_LIST1`.
4. In the Addresses area, add the following addresses: `48.63.32.0/24` and `60.63.10.0/24`. Click **Add** after you type each address.
5. Click **Finished**.
The list screen and the new items are displayed.

Denying access with firewall rules on the network virtual server

The firewall rules in this example apply in the virtual server context. For purposes of this example, the external network-facing virtual server has an IP address of `70.168.15.0/24`. The network virtual server is configured with a pool that includes a publically accessible FTP server at `70.168.15.104`, and an application virtual server at `192.168.15.101`.

Use this task to create a firewall rule that allows all traffic from the networks on the address list `ADDR_LIST1`, and another firewall rule that denies all traffic. This serves the purpose of allowing all traffic from the networks that are allowed access, and denying all other traffic.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. In the Rules area, click **Add** to add a firewall rule to the list.
3. Select the **Virtual Server** context, then select the external network virtual server (in this example, `70.168.15.0/24`).
4. In the **Name** field, type `allow_addr_list`.

5. From the **Type** list, select **Rule**.
6. From the **State** list, select **Enabled**.
7. From the Protocol list, select **Any**.
8. In the **Source** area, from the **Address** list, select **List**.
9. From the **Source Available** list, select ADDR_LIST1, then click the << button to move ADDR_LIST1 to the **Selected** list.
10. Leave the **Destination** area configured with the default **Any / Any** settings.
11. From the **Action** list, select **Accept**.
This allows packets from any source on the address list to any destination and port on any protocol on the DMZ network.
12. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
13. Click the **Repeat** button.
The rule is saved, and a new rule creation page opens, with the same information, so you can create a similar rule.
14. In the **Name** field, type deny_all.
15. In the **Source** area, in the **Address** list, select **Any**.
16. Leave the **Destination** area configured to deny access to **Any** address or port.
17. From the **Action** list, select **Reject**.
This creates a deny all rule for the virtual server.
18. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
19. Click **Finished**.
The list screen and the new items are displayed.
20. From the **Context** list, select **Virtual Server**.
21. From the **Virtual Server** list, select the network virtual server.
22. Click the **Filter** button.
The list screen opens, and all firewall rules that apply to the virtual server are displayed.

Denying access with firewall rules on the application virtual server

The firewall rules in this example apply in the virtual server context. For purposes of this example, the application virtual server on the internal network has an IP address of 192.168.15.101, and is configured to load balance traffic to servers 10.10.1.10 and 10.10.1.11 on ports 80 and 443.

Use this task to create a firewall rule that denies all traffic from the network 48.64.32.0/24 to the internal application servers behind the virtual server **192.168.15.101**.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. In the Rules area, click **Add** to add a firewall rule to the list.
3. Select the **Virtual Server** context, then select the application virtual server (in this example, 192.168.15.101).
4. In the **Name** field, type deny_network_48
5. From the **Type** list, select **Rule**.
6. From the **State** list, select **Enabled**.
7. From the **Schedule** list, select **None**.
8. From the Protocol list, select **Any**.

9. In the **Source** area, from the **Address** list, select **Specify**.
10. In the address field, type `48.64.32.0/24`.
11. Leave the **Destination** area configured to deny access to **Any** address or port.
12. From the **Action** list, select **Drop or Reject**.
This drops packets from the `48.64.32.0` network to any source.
13. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
14. Click **Finished**.
The list screen and the new items are displayed.
15. From the **Context** list, select **Virtual Server**.
16. From the **Virtual Server** list, select the application virtual server.
17. Click the **Filter** button.

The firewall rules are created, and are displayed on the rules list screen for the application virtual server.

Deploying the BIG-IP Network Firewall in Firewall Mode

About Firewall mode in the Network Firewall

The BIG-IP® Advanced Firewall Manager™ (AFM™) provides policy-based access control to and from address and port pairs, inside and outside of your network. In this scenario, the network firewall is configured in *Firewall mode*, a default deny configuration, in which all traffic is blocked through the firewall, and any traffic you want to allow must be explicitly specified.

To understand this firewall scenario, imagine that your prerequisite system load-balances all traffic from the Internet to several internal servers. The internal servers are:

Device and location	IP address	Traffic type
Server on DMZ network	70.168.15.104	FTP
Server on internal network	10.10.1.10	HTTP, HTTPS
Server on internal network	10.10.1.11	HTTP, HTTPS

In order for traffic from the internal application virtual server to reach the external network virtual server, you must create a VLAN and enable both internal and external virtual servers on it. In this scenario, these VLANs are specified:

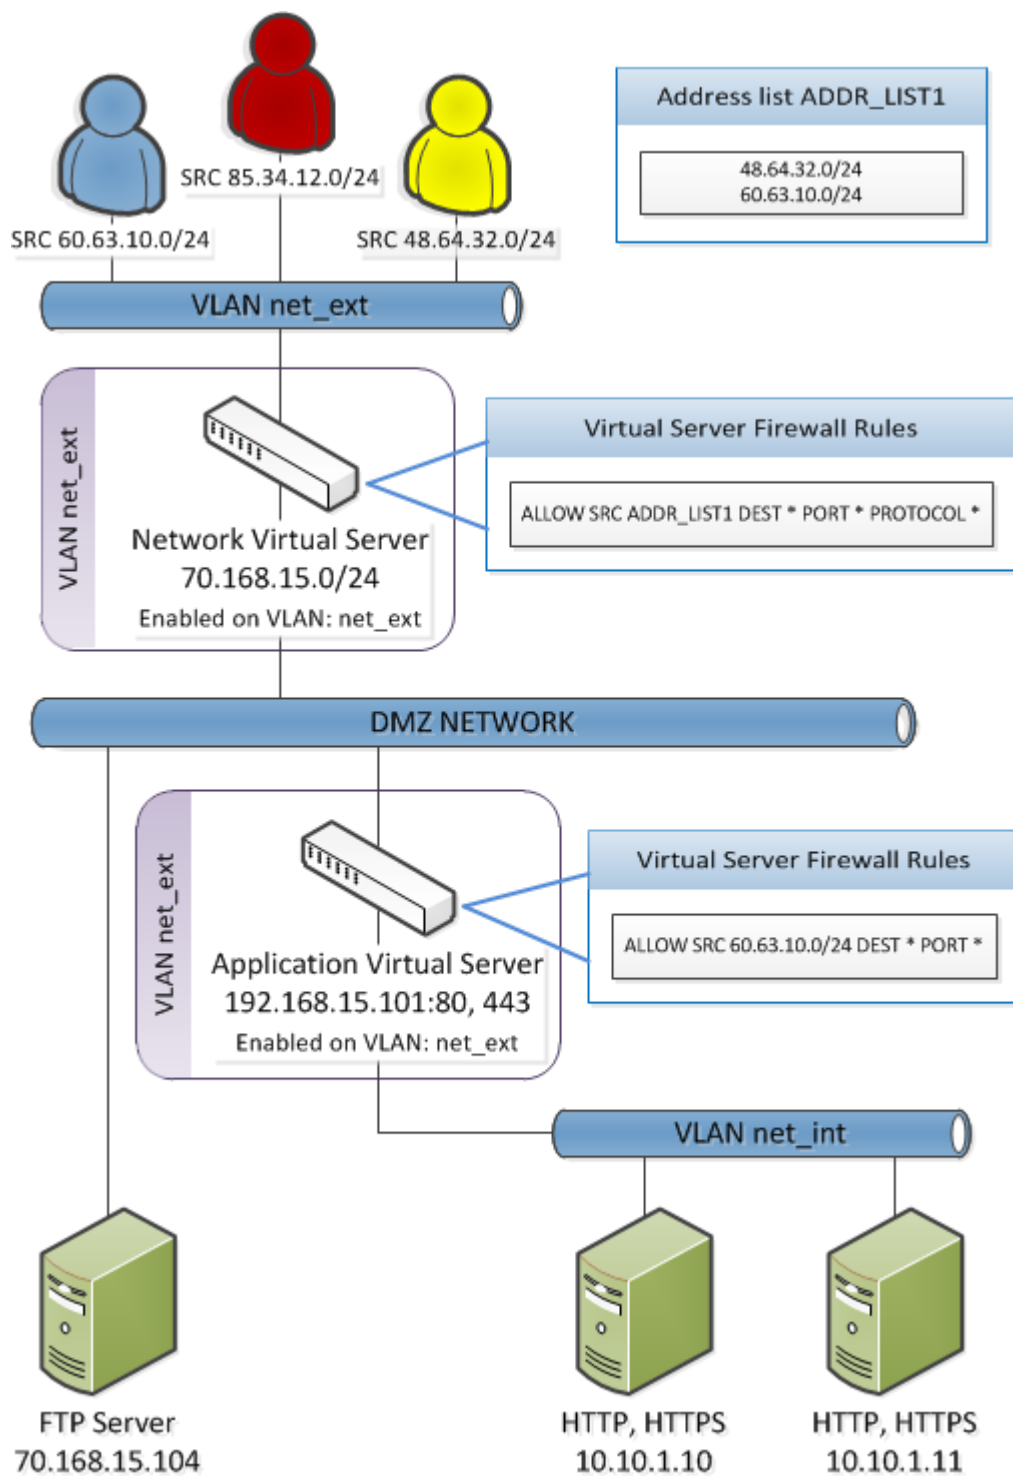
VLAN	Configuration
net_ext	Enabled on 70.168.15.0/24, 192.168.15.101
net_int	Includes pool members 10.10.1.10, 10.10.1.11

In addition, in this firewall configuration, there are three external networks that must be firewalled:

Network	Policy
60.63.10.0/24	Allow all access
85.34.12.0/24	Deny all access
48.64.32.0/24	Allow FTP, deny HTTP and HTTPS

To set up this scenario, you configure addresses, ports, and firewall rules specific to these networks, ports, and addresses.

Figure 4: Firewall configuration scenario



Configuring the Network Firewall to drop or reject traffic that is not specifically allowed

You can configure the BIG-IP® Network Firewall to drop or reject all traffic not explicitly allowed. In Advanced Firewall Manager™, this is called *Firewall mode*, and this is also referred to as a *default deny* policy. Firewall mode applies a default deny policy to all self IP addresses and virtual servers.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Virtual Server & Self IP Contexts** list, select the default action for the self IP and virtual server contexts.
 - Select **Drop** to silently drop all traffic to virtual servers and self IP addresses unless specifically allowed.
 - Select **Reject** to drop all traffic to virtual servers and self IP addresses unless specifically allowed, and to send the appropriate reject message for the protocol.
3. Click **Update**.
The default virtual server and self IP firewall context is changed.

Creating a VLAN for the network firewall

Create a VLAN with tagged interfaces, so that each of the specified interfaces can process traffic destined for that VLAN.

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.
For purposes of this implementation, name the VLAN `net_ext`.
4. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged**.
 - c) Click **Add**.
5. 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.
6. From the **Configuration** list, select **Advanced**.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. For the **Hardware SYN Cookie** setting, select or clear the check box.
When you enable this setting, the BIG-IP system triggers hardware SYN cookie protection for this VLAN.
Enabling this setting causes additional settings to appear. These settings appear on specific BIG-IP platforms only.
9. For the **SynCache Threshold** setting, retain the default value or change it to suit your needs.
The **SynCache Threshold** value represents the number of outstanding SYN flood packets on the VLAN that will trigger the hardware SYN cookie protection feature.
When the **Hardware SYN Cookie** setting is enabled, the BIG-IP system triggers SYN cookie protection in either of these cases, whichever occurs first:
 - The number of TCP half-open connections defined in the LTM[®] setting **Global SYN Check Threshold** is reached.
 - The number of SYN flood packets defined in this **SynCache Threshold** setting is reached.
10. For the **SYN Flood Rate Limit** setting, retain the default value or change it to suit your needs.
The **SYN Flood Rate Limit** value represents the maximum number of SYN flood packets per second received on this VLAN before the BIG-IP system triggers hardware SYN cookie protection for the VLAN.
11. Click **Finished**.

The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Enable the new VLAN on both the network virtual server and the application virtual server.

Configuring an LTM virtual server with a VLAN for Network Firewall

For this implementation, at least two virtual servers and one at least one VLAN are assumed, though your configuration might be different.

You enable two virtual servers on the same VLAN to allow traffic from hosts on one virtual server to reach or pass through the other. In the Network Firewall, if you are using multiple virtual servers to allow or deny traffic to and from specific hosts behind different virtual servers, you must enable those virtual servers on the same VLAN.

***Tip:** By default, the virtual server is set to share traffic on **All VLANs and Tunnels**. This configuration will work for your VLANs, but in the firewall context specifying or limiting VLANs that can share traffic provides greater security.*

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
4. Click **Update** to save the changes.
5. Repeat this task for all virtual servers that must share traffic over the VLAN.

The virtual servers on which you enabled the same VLAN can now pass traffic.

Creating an address list

Use this procedure to specify the address list to apply to allow access to specific source addresses.

1. On the Main tab, click **Security > Network Firewall > Address Lists**.
The Address Lists screen opens.
2. Click **Create** to create a new address list.
3. In the name field, type ADDR_LIST1.
4. In the Addresses area, add the following addresses: 48.63.32.0/24 and 60.63.10.0/24. Click **Add** after you type each address.
5. Click **Finished**.
The list screen and the new items are displayed.

Allowing access from networks on an address list with a firewall rule

The firewall rules in this example apply in the virtual server context. For purposes of this example, the external network-facing virtual server is named ex_vs and has an IP address of 70.168.15.0/24.

Create a firewall rule that allows traffic from the networks on ADDR_LIST1 to the DMZ network, which includes an FTP server that is publicly addressed, and two internal servers on a second virtual server.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.

The Active Rules screen opens.

2. In the Rules area, click **Add** to add a firewall rule to the list.
 3. From the **Context** list, select **Virtual Server**, and then select the external virtual server (in the example, **ex_VS**).
 4. In the **Name** field, type `allow_addr_list`.
 5. From the **Type** list, select **Rule**.
 6. From the **State** list, select **Enabled**.
 7. From the **Protocol** list, select **Any**.
 8. In the Source area, from the **Address** list, select **Specify**, and click **Address List**.
 9. From the list, select **/Common/ADDR_LIST1**, then click **Add** to add **ADDR_LIST1** to the list.
 10. Leave the **Destination** area configured with the default **Any / Any** settings.
 11. From the **Action** list, select **Accept**.
This allows packets from any source on the address list to any destination and port on any protocol on the DMZ network.
 12. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.
 13. Click **Finished**.
The list screen and the new items are displayed.
- A new firewall rule is created, and appears in the firewall rule list.

Allowing access from a network to a virtual server with a firewall rule

The firewall rules in this example apply in the virtual server context. For purposes of this example, the application virtual server is behind the network virtual server with an IP address of `192.168.15.101` and configured for traffic on ports 80 and 443.

Use this procedure to create a firewall rule that allows traffic from a specific external network to the HTTP and HTTPS servers behind an application virtual server.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. In the Rules area, click **Add** to add a firewall rule to the list.
3. In the Context field, select `Virtual Server`, and select the application virtual server (in the example, `192.168.15.101`).
4. In the **Name** field, type `allow_app_vs`.
5. From the **Type** list, select **Rule**.
6. From the **State** list, select **Enabled**.
7. From the Protocol list, select **Any**.
8. In the **Source** area, from the **Address** list, select **Specify**.
9. In the address field, type `60.63.10.0/24`, then click the **Add** button.
10. Leave the **Destination** area configured with the default **Any / Any** settings.
11. From the **Action** list, select **Accept**.
This allows packets from the specified source to any destination and port on any protocol on the internal virtual server. You could specify HTTP and HTTPS protocols, and the internal server addresses, but since these are the only addresses and protocols behind the virtual server, that level of granularity is not necessary.
12. From the **Logging** list, enable or disable logging for the firewall rule.
A logging profile must be enabled to capture logging info for the firewall rule.

13. Click **Finished**.

The list screen and the new items are displayed.

A new firewall rule is created, and appears in the firewall rule list.

Compiling and Deploying Network Firewall rules

About compiling and deploying rules in the Network Firewall

The BIG-IP® Advanced Firewall Manager™ (AFM™) allows you to compile and deploy rules either manually or automatically. Rules are compiled and deployed automatically by default. However, in a large configuration with many rulesets there can be a large number of micro rules created by the compilation process, even when only a small number of rules are added or edited. For such configurations, it might be advantageous to compile all collected rule changes at once, manually. Once rules are compiled, they can be deployed manually or automatically. Deploying manually allows greater control over the rollout of configuration changes. These options provide a more efficient approach to managing large firewall rule sets. When manual rule compilation, manual rule deployment, or both are enabled, the AFM user interface provides feedback about the compilation and deployment status of the current ruleset.

Task list

Configuring manual or automatic policy compilation for firewall rules

Configuring manual or automatic policy deployment for firewall rules

Configuring manual or automatic policy compilation for firewall rules


Set the compilation mode to Manual if you want to collect several rule changes, and then compile them all at one time, or if you want to delay the rule compilation process to another time.

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Firewall Compilation Mode** list, select the compilation mode for the firewall ruleset.
 - Select **Automatic** to compile the firewall ruleset whenever a change is made to any firewall item that is used in the firewall ruleset.
 - Select **Manual** to delay compilation of the firewall ruleset, collect all firewall rule changes, and apply the entire set of changes manually at another time.
3. From the **Log Configuration Changes** list specify the logging option for firewall ruleset compilation and deployment configuration changes.
 - Select **Automatic** to specify that configuration changes are logged only if **Firewall Compilation Mode** or **Firewall Deployment Mode** is set to **Manual**.
 - Select **On** to specify that policy configuration changes are always logged.
 - Select **Off** to specify that policy configuration changes are not logged.
4. Select the log publisher to which to log policy configuration changes.
This field appears only if you specify the **Log Configuration Changes** setting as **Automatic** or **On**.
5. Click **Update**.
The firewall policy compilation mode is configured.

Compiling firewall rules manually

When you have configured the firewall in manual compilation mode, you must manually compile firewall rules after your configuration changes are complete.

1. Look at the status area for Advanced Firewall Manager. If the status shows **Firewall: Pending Rules Compilation**, the rules are ready to be manually compiled.



2. Click the **Firewall: Pending Rules Compilation** link. Alternatively, you can click **Security > Event Logs > Network > Policy Status**.

The Policy Status screen opens, showing the firewall status, an overview of the most recent compilation, and a list of recent configuration changes. If the policy requires compilation, the **Firewall Policy Status** is **Pending Rules Compilation**.

3. Click **Compile**.

The system compiles the collected changes.

After the ruleset is compiled, review the compilation statistics for **Compilation Start Time**, **Compilation End Time**, and **Last Successful Compilation Time**. The status in the **Configuration Change Event** column also shows **Compile Success** after a successful compilation.

If you set the **Firewall Deployment Mode** to automatically deploy after a configuration change, the policies are deployed. If you set the **Firewall Deployment Mode** to manual, you must now deploy the policies.

Configuring manual or automatic policy deployment for firewall rules

Set the deployment mode to Manual if you want to compile rule changes without putting them into effect until a certain time.

Warning: *You can not configure firewall schedules if the firewall deployment mode is manual.*

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Firewall Deployment Mode** list, select the deployment mode for firewall ruleset changes.
 - Select **Automatic** to deploy the firewall ruleset whenever a change is compiled, either manually or automatically.
 - Select **Manual** to delay deployment of the firewall ruleset, collect all compiled firewall ruleset changes, and deploy the entire set of changes manually at another time.
3. From the **Log Configuration Changes** list specify the logging option for firewall ruleset compilation and deployment configuration changes.
 - Select **Automatic** to specify that configuration changes are logged only if **Firewall Compilation Mode** or **Firewall Deployment Mode** is set to **Manual**.
 - Select **On** to specify that policy configuration changes are always logged.
 - Select **Off** to specify that policy configuration changes are not logged.
4. Select the log publisher to which to log policy configuration changes.
This field appears only if you specify the **Log Configuration Changes** setting as **Automatic** or **On**.
5. Click **Update**.
The firewall deployment mode is configured.

Deploying firewall rules manually

When you have configured the firewall in manual deployment mode, you must manually deploy firewall rules after the rules are compiled.

1. Look at the status area for the Advanced Firewall Manager. If the status shows **Firewall: Pending Rules Deployment**, the rules are ready to be manually deployed.



2. Click the **Firewall: Pending Rules Deployment** link. Alternatively, you can click **Security > Event Logs > Network > Policy Status**.

The Policy Status screen opens, showing the firewall status, an overview of the most recent compilation, and a list of the most recent configuration changes. If the policy is compiled, and requires deployment, the **Firewall Policy Status** is **Pending Rules Deployment**.

3. Click **Deploy**.

The system deploys the collected changes.

4. Next to the **Policy Status** setting, select **Advanced** to review additional policy compilation and deployment statistics.

These statistics include the compilation and deployment mode, **Deployment Start Time**, **Deployment End Time**, **Number of Micro Rules**, the **Active BLOB**, and whether the active BLOB is MD5 verified.

After the ruleset is deployed, the status in the **Configuration Change Event** column also shows **Deploy Success** after a successful deployment.

About firewall policy compilation statistics

When firewall rules are recompiled, whether automatically with a rule change, or manually with a manual compile event, the rule list or policy requires some server resources to compile. With large rule sets and deployments, even minor rule changes can cause very large recompilation events. You can view the resources used for policy compilation, either for the entire firewall or by context.

Compiler statistics are displayed on a context for several items.

Activation Time

Displays the time at which firewall policies or rule lists were last activated on this context.

Compilation Duration

Displays the amount of time required to compile the rule sets or policies at the last activation.

Compilation Size

Displays the file size of the compiled rule sets or policies, after the last activation.

Maximum Transient Memory

Displays the maximum memory used to compile the rule sets or policies during the last activation.

Compiler statistics are displayed for several items when displayed for the entire firewall.

Firewall Compilation Mode

Displays whether the firewall is configured to compile ruleset changes manually or automatically.

Firewall Deployment Mode

Displays whether the firewall is configured to deploy ruleset changes manually or automatically.

Firewall Policy Status

Displays whether the firewall ruleset is *Consistent* (all rules are currently compiled and deployed), *Pending Rules Compilation* (some rules have been changed, and the ruleset is not compiled), or *Pending Rules Deployment* (the ruleset is compiled, but not deployed).

Compilation Start Time

Displays the time at which the most recent firewall ruleset compilation event last started.

Compilation End Time

Displays the time at which at which the most recent firewall ruleset compilation event last completed.

Last Successful Compilation Time

Displays the time at which the last successful compilation occurred.

Deployment Start Time

Displays the most recent deployment start time.

Deployment End Time

Displays the most recent deployment end time.

Number of Micro Rules

Displays the number of micro rules compiled in the most recent ruleset compilation event.

Active BLOB

Displays the internal name for the active group of rules to be compiled.

BLOB MD5 Verified

Displays whether the BLOB MD5 is verified.

Viewing compilation statistics for a firewall rule or policy

You can view the most recent compilation statistics for a rule list or policy on the global context, or on a route domain, self IP, or virtual server context.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.

The Active Rules screen opens.

2. From the **Context** list, select **All**.

3. Click on the name of the context for which you want to view statistics.

For example, the global context is always called **Global**. A virtual server or self IP has the name you assigned when you created it; for example, `vs_http_134` or `self_lb_11`. A route domain is identified with a number; for example, `0`.

4. View statistics for rule compilation.

- In the global context, from the **Policy Settings** list, select **Advanced**.
- In a route domain, self IP, or virtual server context, click the Security tab. Then, from the **Policy Settings** list, select **Advanced**.

Statistics are displayed for the most recent rule list and policy compilation on the selected context.

Viewing compilation statistics for all network firewall rules and policies

You can view the most recent compilation statistics for the network firewall.

1. Click **Security > Event Logs > Network > Policy Status**.

The Policy Status screen opens, showing the firewall status, an overview of the most recent compilation, and a list of recent configuration changes.

2. Next to the **Policy Status** setting, select **Advanced** to review additional policy compilation and deployment statistics.

These statistics include the compilation and deployment mode, **Deployment Start Time**, **Deployment End Time**, **Number of Micro Rules**, the **Active BLOB**, and whether the active BLOB is MD5 verified.

Compilation and deployment statistics are displayed for all network firewall policies.

Using Firewall NAT for IP and Port Translation

About using Firewall NAT to translate addresses and ports

Firewall NAT on the BIG-IP[®] Advanced Firewall Manager[™] system supports advanced NAT functionality on the AFM[™] system.

NAT matching policies

NAT policies present a configurable collection of NAT matching rules and NAT translation objects, for inbound and outbound connections. The system matches flows and applies NAT rules, after the matching for firewall rules occurs. Firewall NAT allows you to configure a rule to match traffic, to which NAT source and destination translation rules are applied. Source and destination translation items are configured individually, and can be applied to multiple rules and contexts. Generally, overlapping addresses cannot be configured in NAT source or destination rules. However, you can configure overlapping addresses between two Dynamic PAT items that have the PAT mode set to NAT or Port Block Allocation mode.

Firewall NAT can be used on a system with F5[®] BIG-IP CGNAT (Carrier-Grade NAT). Firewall NAT policies operate with CGNAT policies when applied on the same virtual server.

NAT contexts and precedence

You can configure a firewall NAT policy at the global, virtual server, or route domain context. NAT address and port assignment takes place only at the virtual server level, so a Firewall NAT policy configured at the global context applies on each individual virtual server, and a firewall NAT policy configured at the route domain context applies to all virtual servers on that route domain.

Similarly, NAT policies apply precedence in most-specific to least-specific order. A firewall NAT policy configured on a virtual server takes precedence over a policy configured on the route domain context, or at the global context.

Translation address persistence

The firewall NAT feature module can assign the same external (translation) address to all connections originated by the same internal client, providing endpoint-independent address mapping.

Efficient logging

Firewall NAT supports log messages that map external addresses and ports back to internal clients for both troubleshooting and compliance with law enforcement/legal constraints.

Network address and port translation

Network address and port translation (NAPT) mode provides standard address and port translation allowing multiple clients in a private network to access remote networks using the single IP address assigned to their router.

Proxy ARP

Currently when using AFM NAT to map a range of client source addresses to a range in the same subnet as the IP address of the egress interface, the BIG-IP system does not proxy ARP for the translated source addresses.

Deterministic assignment of translation addresses

Deterministic mode is an option used to assign translation address, and is port-based on the client address/port and destination address/port. It uses reversible mapping to reduce the amount of log messages, while still maintaining the ability for translated IP address to be discovered for troubleshooting and compliance. Deterministic mode also provides an option to configure backup-members.

Port block allocation of translation addresses

Port block allocation (PBA) mode is an option that reduces logging, by logging only the allocation and release of a block of ports. When a subscriber sends a translation request, the BIG-IP system services the request from a block of ports that is assigned to a single IP address, and only logs the allocation and release of that block of ports. The BIG-IP system applies subsequent requests from the service provider to that block of ports until all ports are used.

Important: To use Firewall NAT, you must create a firewall NAT policy, define a matching rule, attach source or destination translation items, and configure the NAT policy at the device level, on a route domain, or on a virtual server.

About Firewall NAT and Carrier Grade NAT (CGNAT)

Firewall NAT on the BIG-IP® Advanced Firewall Manager™ system can be used with Carrier Grade NAT (CGNAT). Firewall NAT policies are not supported with either LTM® SNAT pools or CGNAT LSN-pool configurations on a virtual server.

- If an LTM SNAT pool or CGNAT LSN-pool is applied to a virtual server, a Firewall NAT policy cannot then be applied to that virtual server.
- If a Firewall NAT policy is applied to a virtual server, an LTM SNAT pool or CGNAT LSN-pool cannot be applied to that virtual server. Note that this extends to all contexts at which the Firewall NAT policy can be applied. For example, if a virtual server uses a Firewall NAT policy that is applied on the route domain, an LTM SNAT pool or CGNAT LSN-pool cannot then be applied to that virtual server.

About specifying source translations for Firewall NAT

Source Translation items

With Firewall NAT, source translation rules are contained in a source translation item. This item contains address and port information for the translation pools, and configuration information for each source translation type.

Static NAT

Static NAT mode provides simple 1:1 mapping between the destination IP address and the router IP address. An equal number of internal and external IP addresses must be specified. Ports are not translated.

Static PAT

Static PAT (also known as NAPT) mode provides standard address and port translation allowing multiple clients in a private network to access remote networks using the single IP address assigned to their router. For outbound packets, Static PAT translates the source IP address and source port. For inbound packets, Static PAT translates the destination IP address, the destination port, and the IP and transport header

checksums. This mode is beneficial for remote access users. NAT is the default mode for firewall NAT source translation items.

Dynamic PAT

Dynamic PAT mode provides inbound connection configuration options and mapping options.

With Dynamic PAT you can configure inbound connections with *endpoint independent filtering*, which specifies that the translation attempts to reuse both the address and port mapping (X:x to X':x') for subsequent packets sent from the same internal IP address and port. The BIG-IP system attempts to map X:x to X':x' in every session. This is called *Endpoint Independent Mapping* in *section 4.1 of RFC 4787*.

Dynamic PAT also allows you to configure the following *mapping modes*.

Address pooling paired

Enables all sessions associated with an internal IP address to map to the same external IP address for the duration of the session.

Endpoint independent mapping

Enables use of the same external address and port for all connections from the host, if it uses the same internal port.

Deterministic Mode

With Dynamic PAT, you can configure a source translation item to use *deterministic* mode. Deterministic mode maps internal addresses to external addresses algorithmically, which significantly reduces the amount of log entries generated, while mapping a subscriber's inside IP address with an outside Internet address and port.

Port Block Allocation Mode

With Dynamic PAT, you can configure a source translation item to use port block allocation (PBA) mode. Port block allocation mode is a translation mode option that reduces logging, by logging only the allocation and release of each block of ports. When a subscriber first establishes a network connection, the BIG-IP system reserves a block of ports on a single IP address for that subscriber. The system releases the block when no more connections are using it. This reduces the logging overhead because the system logs only the allocation and release of each block of ports. When a subscriber first connects, the PBA translation mode applies client port block limits, which the subscriber uses as long as it has addresses allocated. For each subscriber, PBA mode compares the subscriber's allocated number of port blocks to the port block limit for the currently connected pool. If the allocated number of port blocks exceeds the port block limit, then the connection is denied. For example, if a subscriber's allocated number of port blocks is 2, and the port block limit for the currently connected pool is 1, then the connection is denied.

Client Connection Limit

In **Dynamic PAT** modes, you can configure a *client connection limit*. This allows you to specify the maximum number of simultaneous translated connections a client or subscriber is allowed to have.

Hairpin Mode

In **Dynamic PAT** modes, you can configure *hairpin mode*. When a client sends a packet to another client in the same private network, hairpin mode sends the packet directly to the destination client's private address. The BIG-IP system immediately translates the packet's public-side destination address. Rather than going out to the public network and returning later for translation, the packet takes a "hairpin turn" at the BIG-IP device.

Specifying source IP addresses for static NAT

Specify static NAT source IP NAT translations to configure the NAT translation addresses for one-to-one static NAT.

1. On the Main tab, click **Security > Network Address Translation > Source Translation**.
The Source Translation screen opens.
2. Click **Create**.
The New Source Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Static NAT**.
5. In the **Addresses** field, add an address or address range on which source translation is performed. Click **Add** for each address or address range.
6. From the **ICMP Echo** list, select whether to enable or disable ICMP echo on translated addresses.
7. From the **Egress interfaces** area, specify the egress interfaces on which source translation is enabled or disabled. Select **Enabled on** or **Disabled on** to specify the egress interface setting.
Egress interfaces include tunnels and VLANs.

The new source translation item appears on the Source Translation screen.

Associate the source translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

Specifying source IP addresses for static PAT

Specify a static PAT source NAT translation to configure the NAT translation addresses for NAT address and port translation.

1. On the Main tab, click **Security > Network Address Translation > Source Translation**.
The Source Translation screen opens.
2. Click **Create**.
The New Source Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Static PAT**.
5. In the **Addresses** field, add an address or address range on which source translation is performed. Click **Add** for each address or address range.
6. In the **Ports** field, add a port or port range on which source translation is performed. Click **Add** for each port or port range.
7. From the **ICMP Echo** list, select whether to enable or disable ICMP echo on translated addresses.
8. From the **Egress interfaces** area, specify the egress interfaces on which source translation is enabled or disabled. Select **Enabled on** or **Disabled on** to specify the egress interface setting.
Egress interfaces include tunnels and VLANs.

The new source translation item appears on the Source Translation screen.

Associate the source translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

Specifying source IP addresses for deterministic dynamic PAT

Deterministic address translation mode provides address translation that eliminates logging of every address mapping, while still allowing internal client address tracking using only an external address and port, and a destination address and port. Deterministic mode allows unique identification of the internal

client address based on: external address and port (the address and port visible to the destination server), destination address and port (the service accessed by the client), and time. Use Deterministic mode to significantly reduce the logging burden, while mapping a subscriber's inside IP address with an outside Internet address and port.

1. On the Main tab, click **Security > Network Address Translation > Source Translation**.
The Source Translation screen opens.
2. Click **Create**.
The New Source Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Dynamic PAT**.
5. In the **Addresses** field, add an address or address range on which source translation is performed.
Click **Add** for each address or address range.
6. In the **Ports** field, add a port or port range on which source translation is performed. Click **Add** for each port or port range.
7. From the **ICMP Echo** list, select whether to enable or disable ICMP echo on translated addresses.
8. From the **PAT Mode** list, select **Deterministic**.
9. From the Inbound Mode list, select the persistence setting for NAT translation entries.
 - **None** disables persistence. With this setting, the mapping of address X and port x (X:x) to address:port X':x' is never guaranteed to persist from one session to the next.
 - **Endpoint Independent Filtering** specifies that the translation attempts to reuse both the address and port mapping (X:x to X':x') for subsequent packets sent from the same internal IP address and port. The BIG-IP system attempts to map X:x to X':x' in every session. This is called *Endpoint Independent Mapping* in RFC 4787, section 4.1.
10. From the Mapping Mode list, select the mapping mode to determine how dynamic ports are assigned, and specify the timeout in seconds for the mapping mode.
 - Select **Address Pooling Paired** to enable all the sessions associated with an internal IP address to map to the same external IP address for the duration of the session.
 - Select **Endpoint Independent Mapping** to assign the same external address and port for all connections from the host if it uses the same internal port.
 - Select **None** to assign no mapping mode to dynamic port assignments.
11. If required, in the **Client Connection Limit** field, specify the maximum number of simultaneous translated connections a client or subscriber is allowed to have.
The default value of **0** specifies no limit.
12. From the **Hairpin Mode** list, enable or disable hairpin mode.
When a client sends a packet to another client in the same private network, *hairpin mode* sends the packet directly to the destination client's private address; the BIG-IP system immediately translates the packet's public-side destination address. Rather than going out to the public network and coming back later for translation, the packet takes a hairpin turn at the BIG-IP device.
13. From the **Egress interfaces** area, specify the egress interfaces on which source translation is enabled or disabled. Select **Enabled on** or **Disabled on** to specify the egress interface setting.
Egress interfaces include tunnels and VLANs.
14. In the Backup Address field, specify backup IP addresses.
This setting creates a pool of IP addresses available for backup members, which are used if Deterministic mode translation fails and falls back to NAPT mode. This is a collection of IP prefixes with their prefix lengths. You can type backup members in the **Add a Backup IP Address** field, and click **Add**.
15. Click **Submit**.

The new source translation item appears on the Source Translation screen.

Associate the source translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

Specifying source IP addresses for dynamic PAT with NAPT

Specify a dynamic PAT source NAT translation to configure the NAT translation addresses for NAT address and port translation for deterministic mode, which reduces logging of address mapping, while still allowing internal client address tracking using only an external address and port, and a destination address and port.

1. On the Main tab, click **Security > Network Address Translation > Source Translation**.
The Source Translation screen opens.
2. Click **Create**.
The New Source Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Dynamic PAT**.
5. In the **Addresses** field, add an address or address range on which source translation is performed. Click **Add** for each address or address range.
6. In the **Ports** field, add a port or port range on which source translation is performed. Click **Add** for each port or port range.
7. From the **ICMP Echo** list, select whether to enable or disable ICMP echo on translated addresses.
8. From the **PAT Mode** list, select **Deterministic**.
9. From the Inbound Mode list, select the persistence setting for NAT translation entries.
 - **None** disables persistence. With this setting, the mapping of address X and port x (X:x) to address:port X':x' is never guaranteed to persist from one session to the next.
 - **Endpoint Independent Filtering** specifies that the translation attempts to reuse both the address and port mapping (X:x to X':x') for subsequent packets sent from the same internal IP address and port. The BIG-IP system attempts to map X:x to X':x' in every session. This is called *Endpoint Independent Mapping* in *RFC 4787, section 4.1*.
10. From the Mapping Mode list, select the mapping mode to determine how dynamic ports are assigned, and specify the timeout in seconds for the mapping mode.
 - Select **Address Pooling Paired** to enable all the sessions associated with an internal IP address to map to the same external IP address for the duration of the session.
 - Select **Endpoint Independent Mapping** to assign the same external address and port for all connections from the host if it uses the same internal port.
 - Select **None** to assign no mapping mode to dynamic port assignments.
11. If required, in the **Client Connection Limit** field, specify the maximum number of simultaneous translated connections a client or subscriber is allowed to have.
The default value of **0** specifies no limit.
12. From the **Hairpin Mode** list, enable or disable hairpin mode.
When a client sends a packet to another client in the same private network, *hairpin mode* sends the packet directly to the destination client's private address; the BIG-IP system immediately translates the packet's public-side destination address. Rather than going out to the public network and coming back later for translation, the packet takes a hairpin turn at the BIG-IP device.
13. From the **Egress interfaces** area, specify the egress interfaces on which source translation is enabled or disabled. Select **Enabled on** or **Disabled on** to specify the egress interface setting.
Egress interfaces include tunnels and VLANs.
14. Click **Submit**.

The new source translation item appears on the Source Translation screen.

Associate the source translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

Specifying source IP addresses for port block allocation mode

Specify a dynamic PAT source NAT translation to configure the NAT translation addresses for NAT address and port translation for port block allocation (PBA) mode, which reduces logging of address mapping, by assigning a block of ports to a translated address and port.

1. On the Main tab, click **Security > Network Address Translation > Source Translation**.
The Source Translation screen opens.
2. Click **Create**.
The New Source Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Dynamic PAT**.
5. In the **Addresses** field, add an address or address range on which source translation is performed. Click **Add** for each address or address range.
6. In the **Ports** field, add a port or port range on which source translation is performed. Click **Add** for each port or port range.
7. From the **ICMP Echo** list, select whether to enable or disable ICMP echo on translated addresses.
8. From the **PAT Mode** list, select **Port Block Allocation**.
9. From the Inbound Mode list, select the persistence setting for NAT translation entries.
 - **None** disables persistence. With this setting, the mapping of address X and port x (X:x) to address:port X':x' is never guaranteed to persist from one session to the next.
 - **Endpoint Independent Filtering** specifies that the translation attempts to reuse both the address and port mapping (X:x to X':x') for subsequent packets sent from the same internal IP address and port. The BIG-IP system attempts to map X:x to X':x' in every session. This is called *Endpoint Independent Mapping* in *RFC 4787, section 4.1*.
10. From the Mapping Mode list, select the mapping mode to determine how dynamic ports are assigned, and specify the timeout in seconds for the mapping mode.
 - Select **Address Pooling Paired** to enable all the sessions associated with an internal IP address to map to the same external IP address for the duration of the session.
 - Select **Endpoint Independent Mapping** to assign the same external address and port for all connections from the host if it uses the same internal port.
 - Select **None** to assign no mapping mode to dynamic port assignments.
11. If required, in the **Client Connection Limit** field, specify the maximum number of simultaneous translated connections a client or subscriber is allowed to have.
The default value of **0** specifies no limit.
12. From the **Hairpin Mode** list, enable or disable hairpin mode.
When a client sends a packet to another client in the same private network, *hairpin mode* sends the packet directly to the destination client's private address; the BIG-IP system immediately translates the packet's public-side destination address. Rather than going out to the public network and coming back later for translation, the packet takes a hairpin turn at the BIG-IP device.
13. From the **Egress interfaces** area, specify the egress interfaces on which source translation is enabled or disabled. Select **Enabled on** or **Disabled on** to specify the egress interface setting.
Egress interfaces include tunnels and VLANs.

The new source translation item appears on the Source Translation screen.

Associate the source translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

About specifying destination translations for Firewall NAT

Destination Translation items

With Firewall NAT, destination translation rules are contained in a destination translation item. This item contains address and port information for the translation pools, and configuration information for each destination translation type.

Static NAT

Static NAT mode provides simple 1:1 mapping between the destination IP address and the router IP address. An equal number of internal and external IP addresses must be specified. Ports are not translated.

Static PAT

Static PAT (also known as NAPT) mode provides standard address and port translation allowing a single IP address to access remote networks using the single IP address assigned to their router. For outbound packets, Static PAT translates the source IP address and source port. This mode is beneficial for remote access users.

Specifying destination IP addresses for static NAT

Add a static NAT destination translation to a Firewall NAT policy to configure the NAT translation addresses for one-to-one mapping of internal destination addresses to external destination addresses.

1. On the Main tab, click **Security > Network Address Translation > Destination Translation**.
The Destination Translation screen opens.
2. Click **Create**.
The New Destination Translation screen opens.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. From the **Type** list, select **Static NAT**.
5. In the **Addresses** field, add an address or address range on which destination translation is performed.
Click **Add** for each address or address range.
6. Click **Submit**.

The new destination translation item appears on the Destination Translation screen.

Associate the destination translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

Specifying destination IP addresses for static PAT

Define a Static PAT destination NAT translation to define destination addresses and ports to translate from internal to external addresses.

1. On the Main tab, click **Security > Network Address Translation > Destination Translation**.
The Destination Translation screen opens.
2. Click **Create**.
The New Destination Translation screen opens.
3. From the **Type** list, select **Static PAT**.

4. In the **Addresses** field, add an address or address range on which destination translation is performed. Click **Add** for each address or address range.
5. In the **Ports** field, add a port or port range on which destination translation is performed. Click **Add** for each port or port range.
6. Click **Submit**.

The new destination translation item appears on the Destination Translation screen.

Associate the destination translation item to a NAT policy, and associate the policy to a virtual server, route domain, or to the global context.

About creating Firewall NAT policies

Firewall NAT policies collect rules to provide NAT address and port translation for source and destination addresses, including match rules for addresses and protocols, and translation rules for source and destination. You can attach a NAT policy at the device level, a route domain, or to a virtual server.

Creating a NAT policy

Create a NAT policy to attach to the device level, a route domain, or a virtual server, to provide NAT address matching and address and port translation for source and destination addresses.

1. On the Main tab, click **Security > Network Address Translation > Policies**.
The Policies screen opens.
2. Click **Create** to create a new policy.
3. In the **Name** and **Description** fields, type the name and an optional description.
4. Click **Add Rule** to add a NAT rule to the policy.
Click the arrow next to **Add Rule** if you want to choose whether to add the rule at the beginning or end of the list. Note that a rule must be selected to add a rule before or after it.
A blank rule appears in the policy.

You have now configured a NAT policy.

Creating a NAT match rule

You can create a NAT match rule in a NAT policy, to identify traffic flows to which the system applies the NAT source and destination translation items.

1. On the Main tab, click **Security > Network Address Translation > Policies**.
The Policies screen opens.
2. From the policy list, click the name of the NAT policy to which to add the rule.
The NAT policy screen opens.
3. Click **Add Rule** to add a NAT rule to the policy.
Click the arrow next to **Add Rule** if you want to choose whether to add the rule at the beginning or end of the list. Note that a rule must be selected to add a rule before or after it.
A blank rule appears in the policy.
4. In the **Name** and **Description** fields, type the name and an optional description.
5. In the **State** column, select the rule state.
 - Select **Enabled** to apply the rule on the protocol, addresses, and ports specified.
 - Select **Disabled** to disable the rule.
6. In the **Protocol** column, select the protocol to which the NAT rule applies.

- Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
 - Select **Other** and type the port number if the protocol is not listed.
7. In the **Source** field, specify the addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, a geographic location, a subscriber or subscriber group, an address list or port list. After you complete an entry, click **Add**.
 8. In the **Destination** field, specify the destination addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, or an address list or port list. After you complete an entry, click **Add**.
 9. From the **Log Profile** list, select a logging profile to apply to the NAT rule.

***Tip:** You can configure the logging profile on the virtual server security policy, instead of on the match rule.*

10. Click **Commit Changes to System**.
The policy with the updated rule is displayed.

You have now configured a NAT rule to match traffic, and apply NAT translations.

About specifying NAT context for a Firewall NAT policy

You can configure a firewall NAT policy at the global, virtual server, or route domain context. NAT address and port assignment takes place only at the virtual server level, so a Firewall NAT policy configured at the global context applies on each individual virtual server, and a firewall NAT policy configured at the route domain context applies to all virtual servers on that route domain.

NAT policies apply precedence in most-specific to least-specific order. A firewall NAT policy configured on a virtual server takes precedence over a policy configured on the route domain context, or at the global context.

When you specify a NAT policy on a virtual server, you can configure the virtual server to use either the route domain policy, the device policy, or both. Orders of precedence still apply, and the most specific NAT policy is applied.

Adding a global Firewall NAT policy

You can specify a firewall NAT policy at the device level to provide NAT translation for matched traffic on all route domains on the device.

***Note:** Note that you can override the device policy by assigning a policy to a route domain, and by assigning a policy to a specific virtual server.*

1. On the Main tab, click **Security > Options > Network Firewall**.
The Network Firewall screen opens to Firewall Options.
2. From the **Network Address Translation** list, select the NAT policy to use for device-level NAT.
3. Click **Update**.
The options are updated.
4. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
5. Click the name of any virtual server on which you want the Global NAT policy to apply.
6. On the menu bar, from the Security menu, choose Policies.

- To use the global Firewall NAT policy, in the Network Address Translation area, click **Use Device Policy**.

The most specific Firewall NAT policy is applied to the context, so a policy applied at the virtual server takes precedence over a route domain policy, which takes precedence over the global policy.

- Click **Update** to save the changes.

You have now configured a NAT policy for the device.

Adding a NAT match rule to the device policy

You can add a NAT match rule to the device NAT policy, to identify traffic flows to which the system applies the NAT source and destination translation items.

- On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
- From the Context Filter list, select **Global**.
- Click **Add Rule > Add rule to Global** to add a NAT rule to the global policy.
Click the arrow next to **Add Rule** if you want to choose whether to add the rule at the beginning or end of the list. Note that a rule must be selected to add a rule before or after it.
A blank rule appears in the policy.
- In the **Name** and **Description** fields, type the name and an optional description.
- In the **State** column, select the rule state.
 - Select **Enabled** to apply the rule on the protocol, addresses, and ports specified.
 - Select **Disabled** to disable the rule.
- In the **Protocol** column, select the protocol to which the NAT rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
 - Select **Other** and type the port number if the protocol is not listed.
- In the **Source** field, specify the addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, a geographic location, a subscriber or subscriber group, an address list or port list. After you complete an entry, click **Add**.
- In the **Destination** field, specify the destination addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, or an address list or port list. After you complete an entry, click **Add**.
- From the **Log Profile** list, select a logging profile to apply to the NAT rule.

Tip: You can configure the logging profile on the virtual server security policy, instead of on the match rule.

- Click **Commit Changes to System**.
The policy with the updated rule is displayed.

You have now configured a NAT rule in the device policy to match traffic, and apply NAT translations.

Configuring a route domain to use Firewall NAT

Before performing this task, confirm that you have a configured Firewall NAT policy.

Assign a Firewall NAT policy to a route domain to use advanced NAT features for address and port translation on a route domain.

- On the Main tab, click **Network > Route Domains**.

The Route Domain List screen opens.

2. In the Name column, click the name of the relevant route domain.
3. On the Main tab, click **Security**.
The Route Domain Security screen opens.
4. From the Network Address Translation list, select the NAT policy to apply to route domain traffic.

Note: When a NAT policy is specified on a more specific context, that policy is applied. For example, a NAT policy on a route domain takes precedence over a global policy, and a policy on a virtual server takes precedence over a route domain policy.

5. Click **Update**.
The system displays the list of route domains on the BIG-IP system.
6. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
7. Click the name of any virtual server on which you want the Route Domain NAT policy to apply.
8. On the menu bar, from the Security menu, choose Policies.
9. To use the route domain Firewall NAT policy, in the Network address translation area, click **Use Route Domain Policy**.
The most specific Firewall NAT policy is applied to the context, so a policy applied at the virtual server takes precedence over a route domain policy, which takes precedence over the global policy.
10. Click **Update** to save the changes.

The route domain now applies the NAT policy to matching traffic, when the route domain policy takes precedence.

Adding a NAT match rule to a route domain

You can add a NAT match rule to a route domain NAT policy, to identify traffic flows to which the route domain applies the NAT source and destination translation items.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. From the Context Filter list, select **Route Domain**.
3. From the Route Domain list, select the route domain to which you want to add NAT match rules.
4. Click **Add Rule > Add rule to Route Domain** to add a NAT rule to the route domain.
Click the arrow next to **Add Rule** if you want to choose whether to add the rule at the beginning or end of the list. Note that a rule must be selected to add a rule before or after it.
A blank rule appears in the policy.
5. In the **Name** and **Description** fields, type the name and an optional description.
6. In the **State** column, select the rule state.
 - Select **Enabled** to apply the rule on the protocol, addresses, and ports specified.
 - Select **Disabled** to disable the rule.
7. In the **Protocol** column, select the protocol to which the NAT rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
 - Select **Other** and type the port number if the protocol is not listed.
8. In the **Source** field, specify the addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, a geographic location, a subscriber or subscriber group, an address list or port list. After you complete an entry, click **Add**.

9. In the **Destination** field, specify the destination addresses and ports that the rule should match.
You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, or an address list or port list. After you complete an entry, click **Add**.
10. From the **Log Profile** list, select a logging profile to apply to the NAT rule.

Tip: You can configure the logging profile on the virtual server security policy, instead of on the match rule.

11. Click **Commit Changes to System**.
The policy with the updated rule is displayed.

You have now configured a NAT rule in the device policy to match traffic, and apply NAT translations.

Configuring Firewall NAT on a virtual server

After you create a firewall NAT policy, you associate that published policy with the virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. To use the global Firewall NAT policy, in the Network Address Translation area, click **Use Device Policy**.
The most specific Firewall NAT policy is applied to the context, so a policy applied at the virtual server takes precedence over a route domain policy, which takes precedence over the global policy.
5. To use the route domain Firewall NAT policy, in the Network address translation area, click **Use Route Domain Policy**.
The most specific Firewall NAT policy is applied to the context, so a policy applied at the virtual server takes precedence over a route domain policy, which takes precedence over the global policy.
6. From the **Policy** list, select the Firewall NAT policy to apply to the context.
7. Click **Finished**.

The Firewall NAT policy is associated with the virtual server.

Adding a NAT match rule to a virtual server

You can add a NAT match rule to a virtual server policy, to identify traffic flows to which the virtual server applies the NAT source and destination translation items.

1. On the Main tab, click **Security > Network Firewall > Active Rules**.
The Active Rules screen opens.
2. From the Context Filter list, select **Virtual Server**.
3. From the Virtual Server list, select the virtual server to which you want to add NAT match rules.
4. Click **Add Rule > Add rule to Virtual Server** to add a NAT rule to the virtual server.
Click the arrow next to **Add Rule** if you want to choose whether to add the rule at the beginning or end of the list. Note that a rule must be selected to add a rule before or after it.
A blank rule appears in the policy.
5. In the **Name** and **Description** fields, type the name and an optional description.
6. In the **State** column, select the rule state.
 - Select **Enabled** to apply the rule on the protocol, addresses, and ports specified.
 - Select **Disabled** to disable the rule.

7. In the **Protocol** column, select the protocol to which the NAT rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.
 - Select **Other** and type the port number if the protocol is not listed.
8. In the **Source** field, specify the addresses and ports that the rule should match.

You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, a geographic location, a subscriber or subscriber group, an address list or port list. After you complete an entry, click **Add**.
9. In the **Destination** field, specify the destination addresses and ports that the rule should match.

You can type an IP address, a contiguous range of IP addresses, an IP subnet, a port, a range of ports, or an address list or port list. After you complete an entry, click **Add**.
10. From the **Log Profile** list, select a logging profile to apply to the NAT rule.

***Tip:** You can configure the logging profile on the virtual server security policy, instead of on the match rule.*

11. Click **Commit Changes to System**.

The policy with the updated rule is displayed.

You have now configured a NAT rule in the device policy to match traffic, and apply NAT translations.

HTTP Protocol Security

Overview: Securing HTTP traffic

You can secure HTTP traffic by using a default configuration or by customizing the configuration. You can adjust the following security checks in an HTTP security profile:

- HTTP protocol compliance validation
- Evasion technique detection
- Length checking to help avoid buffer overflow attacks
- HTTP method validation
- Inclusion or exclusion of certain files by type
- Mandatory header enforcement

You can also specify how you want the system to respond when it encounters a violation. If the system detects a violation and you enabled the Block flag, instead of forwarding the request, the system can either send a blocking response page or redirect the client to a different location.

Creating an HTTP virtual server to use with HTTP protocol security

When you enable protocol security for an HTTP virtual server, the system scans any incoming HTTP traffic for vulnerabilities before the traffic reaches the HTTP servers.

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. In the **Destination Address** field, type the IP address in CIDR format.

The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Note: The IP address you type must be available and not in the loopback network.

5. In the **Service Port** field, type 80 (for HTTP) or 443 (for HTTPS), or select **HTTP** or **HTTPS** from the list.
6. In the Configuration area, for the **HTTP Profile** setting, select the default profile, `http`.
7. From the **Source Address Translation** list, select **Auto Map**.
8. For the **Default Pool** setting, either select an existing pool from the list, or click the Create (+) button and create a new pool.
9. Click **Finished**.

The HTTP virtual server appears in the Virtual Servers list.

Attaching an HTTP protocol security profile to a virtual server

The easiest method for adding HTTP protocol security to your HTTP virtual server is to use the system default profile. You do this by configuring a virtual server with the **HTTP profile** `http`, and then associating the default HTTP protocol security profile `http_security` with the virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. In the **Name** column, click the virtual server you previously created.
The Properties screen for the virtual server opens.
3. On the menu bar, from the Security menu, choose Policies.
4. From the **Protocol Security** list, select **Enabled**.
5. From the **Profile** list, select `http_security`.
This configures the virtual server with the default HTTP protocol security profile.
6. Click **Update**.

You now have a virtual server configured so that HTTP protocol checks are performed on the traffic that the HTTP virtual server receives.

Reviewing violation statistics for security profiles

You can view statistics and transaction information for each security profile that triggers security violations.

1. On the Main tab, click **Security > Event Logs > Protocol** and click **HTTP, FTP, SMTP, or DNS**.
The appropriate statistics screen opens listing all violations for that protocol, with the number of occurrences.
2. Type a Support ID, if you have one, to filter the violations and view one in particular.
3. Click a violation's hyperlink to see details about the requests causing the violation.
On the Statistics screen, in the left column, you can review information regarding the traffic volume for each security profile configured.

Overview: Creating a custom HTTP security profile

This implementation describes how to set up the BIG-IP[®] system to perform security checks on your HTTP virtual server traffic customized to the needs of your environment. Custom configuration of HTTP security and traffic management requires creating an HTTP security profile, and fine tuning this profile so it protects HTTP traffic the way you want. Once you have all HTTP settings specified, you create a virtual server, attach the custom HTTP security profile, and add a default pool to handle the HTTP traffic.

Task summary

Creating a custom HTTP profile

Creating a security profile for HTTP traffic

Configuring an HTTP virtual server with an HTTP security profile

Reviewing violation statistics for security profiles

Creating a custom HTTP profile

An HTTP profile defines the way that you want the BIG-IP® system to manage HTTP traffic.

1. On the Main tab, click **Local Traffic > Profiles > Services > HTTP**.
The HTTP profile list screen opens.
2. Click **Create**.
The New HTTP Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. From the **Parent Profile** list, select **http**.
5. Select the **Custom** check box.
6. Modify the settings, as required.
7. Click **Finished**.

The custom HTTP profile now appears in the HTTP profile list screen.

Creating a security profile for HTTP traffic

Before performing this procedure, verify that you have installed and provisioned BIG-IP® Advanced Firewall Manager™ (AFM) on the BIG-IP system.

An *HTTP security profile* specifies security checks that apply to HTTP traffic, and that you want the BIG-IP system to enforce. In the security profile, you can also configure remote logging and trusted XFF headers.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. Click the **Create** button.
The New HTTP Security Profile screen opens.
3. In the **Profile Name** field, type a unique name for the profile.
4. If you want the security profile to be case-sensitive, leave the **Profile is case sensitive** check box selected. Otherwise, clear the check box.

Note: You cannot change this setting after you create the security profile.

5. Modify the blocking policy settings by clicking **HTTP Protocol Checks** and **Request Checks**, selecting the appropriate options, and enabling the **Block** or **Alarm** options as needed.

*Note: If you do not enable either **Alarm** or **Block** for a protocol check, the system does not perform the corresponding security verification.*

- **Alarm:** The system logs any requests that trigger the security profile violation.
 - **Block:** The system blocks any requests that trigger the security profile violation.
 - **Alarm and Block:** The system both logs and blocks any requests that trigger the security profile violation.
6. If you want to configure the blocking response page, click **Blocking Page**.
 7. Click **Create**.
The screen refreshes, and you see the new security profile in the list.

The BIG-IP system automatically assigns this service profile to HTTP traffic that a designated virtual server receives.

Configuring an HTTP virtual server with an HTTP security profile

You can configure a local traffic virtual server and a default pool for your network's HTTP servers. When the virtual server receives HTTP traffic, an HTTP security profile can scan for security vulnerabilities, and load balance traffic that passes the scan.

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. In the **Destination Address/Mask** field, type an address, as appropriate for your network.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is ff01::0020/64 or 2001:ed8:77b5:2:10:10:100:42/64. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.
5. In the **Service Port** field, type 443 or select **HTTPS** from the list.
6. From the **HTTP Profile** list, select the `http` profile .
7. From the **Source Address Translation** list, select **Auto Map**.
8. In the Resources area of the screen, for the **Default Pool** setting, click the **Create (+)** button.
The New Pool screen opens.
9. In the **Name** field, type a unique name for the pool.
10. In the Resources area, for the **New Members** setting, select the type of new member you are adding, then type the information in the appropriate fields, and click **Add** to add as many pool members as you need.
11. Click **Finished** to create the pool.
The screen refreshes, and reopens the New Virtual Server screen. The new pool name appears in the **Default Pool** list.
12. Click **Finished** to create the virtual server.
The screen refreshes, and you see the new virtual server in the list.
13. In the Name column, click the name of the relevant virtual server.
This displays the properties of the virtual server.
14. On the menu bar, from the Security menu, choose Policies.
15. From the **Protocol Security** list, select **Enabled**.
16. From the **Protocol Security Profile** list, select your custom HTTP security profile.
17. Click **Update** to save the changes.

Reviewing violation statistics for security profiles

You can view statistics and transaction information for each security profile that triggers security violations.

1. On the Main tab, click **Security > Event Logs > Protocol** and click **HTTP, FTP, SMTP, or DNS**.
The appropriate statistics screen opens listing all violations for that protocol, with the number of occurrences.
2. Type a Support ID, if you have one, to filter the violations and view one in particular.
3. Click a violation's hyperlink to see details about the requests causing the violation.
On the Statistics screen, in the left column, you can review information regarding the traffic volume for each security profile configured.

Overview: Increasing HTTP traffic security

The HTTP security profile consists of many different security checks for the various components of HTTP traffic. This implementation shows you how to fine-tune your HTTP security profile as required by your environment. The custom checks are described under the assumption that you have already created a custom HTTP security profile but have no other prerequisite or special order. You need configure only the custom checks that you are interested in.

You can achieve a greater level of security when you configure the system to perform the following checks:

- HTTP Protocol Checks that are related to RFC compliance and actions to take resulting from a violation
- Request Checks, such as length, allowable HTTP request methods, inclusion or exclusion of file types, and custom headers that must occur in every request
- Blocking Page configuration which describes the page to display in the event of a blocked request when a violation is encountered

About RFC compliance and validation checks

When the BIG-IP® system receives an HTTP request from a client, the first validation check that the system performs is to ensure that it is RFC protocol compliant. If the request passes the compliance checks, the system applies the security profile to the request. So that your system fully validates RFC compliance, keep the following HTTP Protocol Checks enabled (they are enabled by default):

- **Several Content-Length headers:** This security check fails when the incoming request contains more than one content-length header.
- **Null in request:** This security check fails when the incoming request contains a null character.
- **Unparsable request content:** This security check fails when the Advanced Firewall Manager™ is unable to parse the incoming request.

Modifying HTTP protocol compliance checks

F5 Networks® recommends that you retain the default properties for the HTTP protocol security checks. This task allows you to take additional precautions such as enabling the Block flag for the HTTP Protocol Checks setting, even if you enable only the Alarm flag for the other security checks. When you do this, the system blocks all requests that are not compliant with HTTP protocol standards, and performs additional security checks only on valid HTTP traffic.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the HTTP Security Profiles area, in the Profile Name column, click the name of the security profile that you are modifying.
The HTTP Profile Properties screen opens.
3. On the HTTP Protocol Checks tab, for the **HTTP Protocol Checks** setting, select the check boxes for the protocol checks that you want the system to validate.
4. Select **Alarm** or **Block** to indicate how you want the system to respond to a triggered violation.
The default setting is **Alarm**.
 - **Alarm:** The system logs any requests that trigger the violation.
 - **Block:** The system blocks any requests that trigger the violation.
 - **Alarm and Block:** The system both logs and blocks any requests that trigger the violation.

5. Click **Update** to retain changes.

The BIG-IP® system is now enabled for compliance checks on all valid HTTP traffic.

About evasion techniques checks

Advanced Firewall Manager™ can examine HTTP requests for methods of application attack that are designed to avoid detection. When found, these coding methods, called *evasion techniques*, trigger the Evasion technique detected violation. By creating HTTP security profiles, you can detect evasion techniques, such as:

- Directory traversal, for example, `a/b/././c` turns into `a/c`
- Multiple decoding passes
- Multiple backslash characters in a URI, for example, `\\servername`
- Bare byte decoding (higher than ASCII-127) in a URI
- Apache whitespace characters (`0x09`, `0x0b`, or `0x0c`)
- Bad unescape

By default, the system logs requests that contain evasion techniques. You can also block requests that include evasion techniques.

Configuring HTTP protocol evasion techniques blocking policy

You can use HTTP security profiles to detect, log, alarm, and block evasion techniques detected in HTTP traffic.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the HTTP Security Profiles area, in the Profile Name column, click the name of the security profile that you are modifying.
The HTTP Profile Properties screen opens.
3. On the HTTP Protocol Checks tab, for the **Evasion Techniques Checks** setting, select or clear the **Alarm** or **Block** check boxes, as required.

Option	Description
Alarm	The system logs any requests that trigger the violation. This is the default setting.
Block	The system blocks any requests that trigger the violation.
Alarm and Block	The system both logs and blocks any requests that trigger the violation.

4. Click **Update** to retain changes.

About the types of HTTP request checks

By creating HTTP security profiles, you can perform several types of checks on HTTP requests to ensure that the requests are well-formed and protocol-compliant.

Length checks

Specify valid maximum lengths for request components to help prevent buffer overflow attacks.

Method checks

Specify which HTTP methods the system allows in requests.

File type checks

Specify which file types users can or cannot access.

Mandatory headers

Specify custom headers that must occur in every request.

Null in request

This security check fails when the incoming request contains a null character.

Unparsable request content

This security check fails when the system is unable to parse the incoming request.

Configuring length checks for HTTP traffic

Before performing this procedure, verify that you have installed and provisioned BIG-IP® Advanced Firewall Manager™ (AFM) on the BIG-IP system.

You can specify valid maximum lengths for request components in HTTP security profiles to prevent buffer overflow attacks. You can set maximum lengths for URLs, query strings, POST data, and the entire request.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the Profile Name column, click the name of the security profile for which you want to configure length checking.
The Profile Properties screen opens.
3. Click the Request Checks tab.
4. For each option of the **Length Checks** setting, specify **Any** to allow any length or click **Length** and specify the maximum length you want to allow.
5. Select **Alarm** or **Block**, to indicate how you want the system to respond to a triggered violation.
The default setting is **Alarm**.
 - **Alarm**: The system logs any requests that trigger the violation.
 - **Block**: The system blocks any requests that trigger the violation.
 - **Alarm and Block**: The system both logs and blocks any requests that trigger the violation.
6. For the **Request Length Exceeds Defined Buffer Size** setting, select or clear **Alarm** and **Block**, as needed.
 - **Alarm**: The system logs any requests that are longer than allowed by the **long_request_buffer_size** internal parameter (the default is 10,000,000 bytes).
 - **Block**: The system blocks any requests that are longer than allowed by the **long_request_buffer_size** internal parameter (the default is 10,000,000 bytes).
 - **Alarm and Block**: The system both logs and blocks any requests that trigger the violation.
7. Click **Update** to retain changes.

Specifying which HTTP methods to allow

Before performing this procedure, verify that you have installed and provisioned BIG-IP® Advanced Firewall Manager™ (AFM) on the BIG-IP system.

The HTTP security profile accepts certain HTTP methods by default. The default allowed methods are GET, HEAD, and POST. The system treats any incoming HTTP request that includes an HTTP method other than the allowed methods as a violating request. Later, you can decide how to handle each violation.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.

The Security Profiles: HTTP screen opens.

2. In the Profile Name column, click the name of the security profile for which you want to modify allowable HTTP methods.
The Profile Properties screen opens.
3. Click the Request Checks tab.
4. For the **Methods** setting, specify which HTTP methods to allow:
The default allowed methods are GET, HEAD, and POST.
 - From the **Available** list, select the methods you want to allow in a request and move them to the **Allowed** list.
 - To add a new method to the **Available** list: type the name in the **Method** field, click **Add** to add it to the list, and move it to the **Allowed** list.
5. Select **Alarm** or **Block**, to indicate how you want the system to respond to a triggered violation.
The default setting is **Alarm**.
 - **Alarm**: The system logs any requests that trigger the violation.
 - **Block**: The system blocks any requests that trigger the violation.
 - **Alarm and Block**: The system both logs and blocks any requests that trigger the violation.
6. Click **Update** to retain changes.

Including or excluding files by type in HTTP security profiles

Before performing this procedure, verify that you have installed and provisioned BIG-IP® Advanced Firewall Manager™ (AFM) on the BIG-IP system.

By default, an HTTP security profile permits all file types in a request. For tighter security, you can create a list that specifies either all file types you want to allow, or a list specifying all the file types you do not want allowed.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the Profile Name column, click the name of the security profile you want to update.
The Profile Properties screen opens.
3. Click the Request Checks tab.
4. For the **File Types** setting, specify whether you want to create a list of allowed or disallowed file types, and which files you want in the list.
 - To create a list of file types that are permitted in requests, select **Define Allowed**.
 - To create a list of file types not permitted, select **Define Disallowed**.
 - Select file types from the **Available** list, and move them to the **Allowed** or **Disallowed** list.
 - To add a new file type, type the name in the **File Type** field, click **Add** to add it to the **Available** list, and then move it to the **Allowed** or **Disallowed** list.

Important: *If the profile is case-sensitive, the file types are case-sensitive. For example, **jsp** and **JSP** will be treated as separate file types.*

5. Select **Alarm** or **Block**, to indicate how you want the system to respond to a triggered violation.
The default setting is **Alarm**.
 - **Alarm**: The system logs any requests that trigger the violation.
 - **Block**: The system blocks any requests that trigger the violation.
 - **Alarm and Block**: The system both logs and blocks any requests that trigger the violation.

The page you configure is displayed every time one of the security checks set to **Block** is violated.

Configuring a mandatory header for an HTTP security profile

Before performing this procedure, verify that you have installed and provisioned BIG-IP® Advanced Firewall Manager™ (AFM) on the BIG-IP system.

When the BIG-IP® system is managing an application that uses custom headers that must occur in every request, you can specify mandatory HTTP headers in the security profile. The system verifies that all requests contain those headers. If a request does not contain the mandatory header, the system issues the Mandatory HTTP header is missing violation, and takes the action that you configure: Alarm, Block, or both.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the Profile Name column, click the name of the security profile for which you want to configure a Mandatory Header alarm.
The Profile Properties screen opens.
3. Click the Request Checks tab.
4. For the **Mandatory Headers** setting, specify the header that must be in the request:
 - a) In the **Header** field, type the name of the mandatory header, and click the **Add** button to add it to the **Available** list.
 - b) Move the new mandatory header from the **Available** list to the **Mandatory** list.
 - c) Select or clear the **Alarm** or **Block** check boxes as required.

Option	Description
Alarm	The system logs any responses that trigger the Mandatory HTTP header is missing violation. This is the default setting.
Block	The system blocks any requests that trigger the Mandatory HTTP header is missing violation.
Alarm and Block	The system both logs and blocks any requests that trigger the Mandatory HTTP header is missing violation.

5. Click **Update** to retain changes.

All HTTP requests are checked for the mandatory headers you have selected.

Configuring the blocking response page for HTTP security profiles

If your HTTP security profile is set up to block requests that violate one or more of the security checks, the system displays a page, called the blocking response page, on the client's screen. The default blocking response page states that the request was rejected, and provides a support ID. You can also configure the system to redirect the client to a specific web site instead of displaying the blocking response page.

1. On the Main tab, click **Security > Protocol Security > Security Profiles > HTTP**.
The Security Profiles: HTTP screen opens.
2. In the Profile Name column, click the name of the security profile for which you want to configure a blocking page.
The Profile Properties screen opens.
3. Click the Blocking Page tab.
4. For the **Response Type** setting, select one of the options:

- **Default Response:** Specifies that the system returns the system-supplied blocking response page. Though you cannot edit the HTML code on the default blocking page, you can copy it into a custom response and edit it.
- **Custom Response:** Specifies that the system returns a response page that you design or upload.
- **Redirect URL:** Specifies that the system redirects the client to the specified URL.
- **SOAP Fault:** Specifies that the system displays a blocking page in standard SOAP fault message format. Though you cannot edit the SOAP fault code, you can copy it into a custom response and edit it.

The settings on the screen change depending on the selection that you make for the Response Type setting.

5. If you selected the **Custom Response** option, you can either create a new response or upload an HTML file.
 - To create a custom response, make the changes you want to the default responses for the **Response Header** and **Response Body** settings using HTTP syntax for the content, and click **Upload**.
 - To upload an HTML file for the response body, navigate to an existing HTML response page, and click **Upload**.
6. If you selected **Redirect URL**, type the full path of the web page to which the system should redirect the client in the **Redirect URL** field.
7. Click **Update** to retain changes.

The system displays the response page when a violation occurs on any of the security checks set to **Block**.

Overview: Configuring Local Protocol Security Event Logging

You can configure the BIG-IP® system to log detailed information about protocol security events and store those logs locally.

Important: The BIG-IP Advanced Firewall Manager™ (AFM) must be licensed and provisioned and DNS Services must be licensed before you can configure Protocol Security event logging.

Creating a local Protocol Security Logging profile

Configuring a virtual server for Protocol Security event logging

Viewing Protocol Security event logs locally on the BIG-IP system

Disabling logging

Creating a local Protocol Security Logging profile

Create a custom Logging profile to log BIG-IP system network firewall events locally on the BIG-IP system.

1. On the Main tab, click **Security > Event Logs > Logging Profiles**.
The Logging Profiles list screen opens.
2. Click **Create**.
The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a unique name for the profile.
4. Select the **Protocol Security** check box.
5. In the HTTP, FTP, and SMTP Security area, from the **Publisher** list, select **local-db-publisher**.
6. In the DNS Security area, from the **Publisher** list, select **local-db-publisher**.

7. Select the **Log Dropped Requests** check box, to enable the BIG-IP system to log dropped DNS requests.
8. Select the **Log Filtered Dropped Requests** check box, to enable the BIG-IP system to log DNS requests dropped due to DNS query/header-opcode filtering.

Note: The system does not log DNS requests that are dropped due to errors in the way the system processes DNS packets.

9. Select the **Log Malformed Requests** check box to enable the BIG-IP system to log malformed DNS requests.
10. Select the **Log Rejected Requests** check box to enable the BIG-IP system to log rejected DNS requests.
11. Select the **Log Malicious Requests** check box to enable the BIG-IP system to log malicious DNS requests.
12. Click **Finished**.

Assign this custom protocol security Logging profile to a virtual server.

Configuring a virtual server for Protocol Security event logging

Ensure that at least one Log Publisher exists on the BIG-IP[®] system.

Assign a custom Protocol Security Logging profile to a virtual server when you want the BIG-IP system to log Protocol Security events on the traffic the virtual server processes.

*Note: This task applies only to systems provisioned at a minimum level (or higher) for **Local Traffic (LTM)**. You can check the provisioning level on the **System > Resource Provisioning** screen.*

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

Viewing Protocol Security event logs locally on the BIG-IP system

Ensure that the BIG-IP[®] system is configured to log the types of events you want to view, and to store the log messages locally on the BIG-IP system.

When the BIG-IP system is configured to log events locally, you can view those events using the Configuration utility.

1. On the Main tab, click **Security > Event Logs > Protocol > DNS**.
The Protocol Security event log displays.
2. To search for specific events, click **Custom Search**. Drag the event data that you want to search for from the Event Log table into the Custom Search table, and then click **Search**.

Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP[®] system to log specific events on the traffic handled by specific resources.

Note: You can disable and re-enable logging for a specific resource based on your network administration needs.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

Implementation result

You now have an implementation in which the BIG-IP[®] system logs specific Protocol Security events locally.

Overview: Configuring Remote Protocol Security Event Logging

You can configure the BIG-IP[®] system to log information about BIG-IP system Protocol Security events and send the log messages to remote high-speed log servers.

***Important:** The Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure Protocol Security event logging.*

This illustration shows the association of the configuration objects for remote high-speed logging.




Figure 5: Association of remote high-speed logging configuration objects

Task summary

Perform these tasks to configure Protocol Security event logging on the BIG-IP® system.

Note: Enabling remote high-speed logging impacts BIG-IP system performance.

- Creating a pool of remote logging servers*
- Creating a remote high-speed log destination*
- Creating a formatted remote high-speed log destination*
- Creating a publisher*
- Creating a custom Protocol Security Logging profile*
- Configuring a virtual server for Protocol Security event logging*
- Disabling logging*

About the configuration objects of remote protocol security event logging

When configuring remote high-speed logging of Protocol Security events, it is helpful to understand the objects you need to create and why, as described here:

Object	Reason	Applies to
Pool of remote log servers	Create a pool of remote log servers to which the BIG-IP® system can send log messages.	Creating a pool of remote logging servers.
Destination (unformatted)	Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers.	Creating a remote high-speed log destination.

Object	Reason	Applies to
Destination (formatted)	If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination.	Creating a formatted remote high-speed log destination.
Publisher	Create a log publisher to send logs to a set of specified log destinations.	Creating a publisher.
DNS Logging profile	Create a custom DNS Logging profile to define the data you want the BIG-IP system to include in the DNS logs and associate a log publisher with the profile.	Creating a custom Protocol Security Logging profile.
LTM [®] virtual server	Associate a custom DNS profile with a virtual server to define how the BIG-IP system logs the DNS traffic that the virtual server processes.	Configuring a virtual server for Protocol Security event logging.

Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP[®] system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
 - **DNS > Delivery > Load Balancing > Pools**
 - **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 the IP address for each remote logging server that you want to include in the pool:
 - a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
 - b) Type a service number in the **Service Port** field, or select a service name from the list.

Note: Typical remote logging servers require port 514.

- c) Click **Add**.
5. Click **Finished**.

Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

Important: If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the **Remote High-Speed Log** type. With this configuration, the BIG-IP system can send data to the servers in the required format.

The BIG-IP system is configured to send an unformatted string of text to the log servers.

5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

Important: ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.

Important: For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

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 custom Protocol Security Logging profile

Create a logging profile to log Protocol Security events for the traffic handled by the virtual server to which the profile is assigned.

Note: You can configure logging profiles for HTTP and DNS security events on Advanced Firewall Manager™, and FTP and SMTP security events on Application Security Manager™.

1. On the Main tab, click **Security > Event Logs > Logging Profiles**.
The Logging Profiles list screen opens.
2. Click **Create**.
The Create New Logging Profile screen opens.
3. Select the **Protocol Security** check box.
4. In the HTTP, FTP, and SMTP Security area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log HTTP, FTP, and SMTP Security events.
5. In the DNS Security area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log DNS Security events.
6. Select the **Log Dropped Requests** check box, to enable the BIG-IP system to log dropped DNS requests.
7. Select the **Log Filtered Dropped Requests** check box, to enable the BIG-IP system to log DNS requests dropped due to DNS query/header-opcode filtering.

Note: The system does not log DNS requests that are dropped due to errors in the way the system processes DNS packets.

8. Select the **Log Malformed Requests** check box to enable the BIG-IP system to log malformed DNS requests.
9. Select the **Log Rejected Requests** check box to enable the BIG-IP system to log rejected DNS requests.
10. Select the **Log Malicious Requests** check box to enable the BIG-IP system to log malicious DNS requests.

11. From the **Storage Format** list, select how the BIG-IP system formats the log.

Option	Description
None	Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example: <pre>"management_ip_address", "bigip_hostname", "context_type", "context_name", "src_ip", "dest_ip", "src_port", "dest_port", "vlan", "protocol", "route_domain", "acl_rule_name", "action", "drop_reason"</pre>
Field-List	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Specify the order the fields display in the log. • Specify the delimiter that separates the content in the log. The default delimiter is the comma character.
User-Defined	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Cut and paste, in a string of text, the order the fields display in the log.

12. Click **Finished**.

Assign this custom Protocol Security Logging profile to a virtual server.

Configuring a virtual server for Protocol Security event logging

Ensure that at least one Log Publisher exists on the BIG-IP® system.

Assign a custom Protocol Security Logging profile to a virtual server when you want the BIG-IP system to log Protocol Security events on the traffic the virtual server processes.

***Note:** This task applies only to systems provisioned at a minimum level (or higher) for **Local Traffic (LTM)**. You can check the provisioning level on the **System > Resource Provisioning** screen.*

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.
5. Click **Update** to save the changes.

Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

***Note:** You can disable and re-enable logging for a specific resource based on your network administration needs.*

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.

3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

Implementation result

You now have an implementation in which the BIG-IP[®] system logs specific Protocol Security events and sends the logs to a specific location.

Preventing Attacks with Eviction Policies and Connection Limits

What are eviction policies and connection limits?

An *eviction policy* provides the system with guidelines for how aggressively it discards flows from the flow table. You can customize the eviction policy to prevent flow table attacks, where a large number of slow flows are used to negatively impact system resources. You can also set how the system responds to such flow problems in an eviction policy, and attach such eviction policies globally, to route domains, and to virtual servers, to protect the system, applications, and network segments with a high level of customization.

A *connection limit* provides a hard limit to the number of connections allowed on a virtual server or on a route domain. If you set such a limit, all connection attempts that exceed this limit are not allowed.

Task list

Creating an eviction policy

Limiting global connections and flows

Limiting connections and flows on a virtual server

Limiting connections and flows on a route domain

Creating an eviction policy

You can create eviction policies to control the granularity and aggressiveness with which the system discards flows.

1. On the Main tab, click **System > Configuration > Local Traffic > Eviction Policy List**.
2. Click **Create**.
The **New Eviction Policy** screen opens.
3. In the **Name** field, type a name for the eviction policy.
4. In the **Trigger** fields, type a high and low water mark for the eviction policy.
This measure specifies the percentage of the quota, for this context, before flow eviction starts (high water mark) and ends (low water mark).
5. Enable **Slow Flow Monitoring** to monitor flows that are considered slow by the system, and specify the slow flow threshold in bytes per second.
This combination of settings monitors the system for flows that fall below the slow flow threshold for more than 30 seconds.
6. In the **Grace Period** field you can set a grace period, in seconds, between the detection of slow flows that meet the threshold requirement, and purging of slow flows according to the **Slow Flow Throttling** settings.
7. In the Slow Flow Throttling area, set the slow flow throttling options.

Option	Description
--------	-------------

Disabled	Slow flows are monitored, but not removed from the system when the threshold requirement is met for 30 seconds.
-----------------	---

Option Description

Absolute Slow flows are removed from the system when the threshold requirement is met for 30 seconds. Setting an absolute limit removes all slow flows beyond the specified absolute number of flows.

Percent Slow flows are removed from the system when the threshold requirement is met for 30 seconds. Setting a percentage limit removes that percentage of slow flows that exceed the specified monitoring setting, so the default value of 100% removes all slow flows that exceed the slow flow threshold, after the grace period.

8. In the Algorithms area, configure the strategies that the eviction policy uses to remove flows by moving algorithms from the **Available** list to the **Selected** list, and configuring applicable settings for the algorithms.

9. Click **Finished**.

The eviction policy appears in the Eviction Policy List.

To use an eviction policy, associate it with a virtual server or a route domain. You can configure a global eviction policy at **System > Configuration > Local Traffic > General**.

Eviction policy strategy algorithms

This table lists the BIG-IP® eviction policy algorithms and associated configuration information.

In an eviction policy, you specify one or more algorithms, or any combination of algorithms, to determine how traffic flows are dropped when the eviction policy threshold limits are reached. Selected algorithms are processed at the same time as a combined strategy, not in a specific order, so the combination of algorithms determines the final strategy used to remove flows. This strategy biases or weights the final algorithm toward the outcomes you have selected, though these choices are not absolute.

***Important:** You must specify at least one algorithm to use to determine how traffic is dropped with an eviction policy, otherwise flows are removed at random when the eviction policy threshold is reached.*

Algorithm	Description
Bias Idle	Biases flow removal toward the existing flows that have been idle, with no payload bytes, for the longest.
Bias Oldest	Biases flow removal toward the oldest existing flows.
Bias Bytes	Biases flow removal toward the flows with the fewest bytes. When this algorithm is selected, add a value to the field Minimum Time Delay in the Strategy Configuration area. This value determines the period of time for which a flow is allowed to exist, at a minimum, before it is subject to removal through the Bias Bytes algorithm.
Bias Fast	Biases flow removal toward the fastest existing flows.
Bias Slow	Biases flow removal toward the slowest existing flows.
Low Priority Route Domains	Biases flow removal toward flows on low priority route domains. When this algorithm is selected, use the Low Priority Route Domains setting in the Strategy Configuration area to move low

Algorithm	Description
Low Priority Virtual Servers	priority route domains from the Available list to the Selected list. Biases flow removal toward flows on low priority virtual servers. When this algorithm is selected, use the Low Priority Virtual Servers setting in the Strategy Configuration area to move low priority virtual servers from the Available list to the Selected list.
Low Priority Countries	Biases flow removal toward flows from lower priority countries. When this algorithm is selected, in the Low Priority Countries setting in the Strategy Configuration area, select low priority countries from the list and click Add to add them to the low priority list.
Low Priority Ports and Protocols	Biases flow removal toward flows on low priority ports and protocols. When this algorithm is selected, use the Low Priority Ports and Protocols setting in the Strategy Configuration area to add ports, protocols, and combinations to the low priority ports and protocols list (you must also specify a name).

Limiting global connections and flows

You must first create an eviction policy before you can assign one globally. The system includes a global eviction policy, by default.

Assign global connection limits and an eviction policy to prevent possible attacks or overflows on system flows.

1. On the Main tab, click **System > Configuration > Local Traffic > General**.
The Local Traffic General settings screen opens.
2. From the **Eviction Policy** list, select the eviction policy to apply globally.

Note: The global context requires an eviction policy. If you do not apply a custom eviction policy, the system default policy, `default-eviction-policy` is applied and selected in this field.

3. Click **Update** to apply the changes.
The eviction policy is applied to the context.

Limiting connections and flows on a virtual server

You must first create an eviction policy before you can assign one to a virtual server.

Assign connection limits and an eviction policy to a virtual server to enact granular control over possible attacks or overflows on system flows.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. In the **Connection Limit** field, type a number that specifies the maximum number of concurrent open connections.

5. From the **Eviction Policy** list, select an eviction policy to apply to the virtual server.
6. Click **Update** to apply the changes.
The eviction policy is applied to the context.

Limiting connections and flows on a route domain

Before performing this task, confirm that you have a configured route domain, or use the common route domain 0. You must add VLANs to a route domain for the route domain to effect traffic.

Assign connection limits and an eviction policy to a route domain to enact granular control over possible attacks or overflows on system flows.

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. In the Name column, click the name of the relevant route domain.
3. In the **Connection Limit** field, type the maximum number of concurrent connections allowed for the route domain. Setting this to 0 turns off connection limits. The default is 0.
4. From the **Eviction Policy** list, select an eviction policy to apply to this route domain.
5. Click **Update**.
The system displays the list of route domains on the BIG-IP system.

The route domain now applies the connection limit and eviction policy to flows and connections.

Setting Timers and Preventing Port Misuse with Service Policies

Creating and Applying Service Policies

Introduction to service policies

A *service policy* collects flow timer and flow timeout features in a policy that can be applied to different contexts, and allows you to configure policies to drop traffic on a specified port when the service does not match.

A service policy can be applied on a route domain, virtual server, self IP, or in a firewall rule.

About service policy types

There are two types of service policies that you can create:

- A *timer policy* allows you to configure a custom idle timeout that differs from the system idle timeout. The timer policy can be applied to the global, virtual server, self IP and route domain contexts, and can also be referenced by firewall rules or rule lists.
- A *port misuse policy* allows you to configure a route domain, firewall rule, or firewall rule list to detect and drop connections that are not using a required application or service for a given port. With a port misuse policy, you can configure ports to allow services, and drop all traffic that does not match the specified service type. You can configure port and service associations without regard for customary port and service pairings.

Creating a timer policy

Create a timer policy to set custom timeouts for self IPs, route domains, firewall rules, or firewall rule lists.

1. Click **Network > Service Policies > Timer Policies**.
2. Click **Create**.
The New Timer Policy screen opens.
3. Type a name for the timer policy.
4. Type an optional description for the timer policy.
5. To save the timer policy and add timer rules, click **Create & Add Rule**.
The New Rule screen opens.
6. Type a name for the rule.
7. From the **Protocol** list, select a protocol.
8. From the **Idle Timeout** list, select the timeout option for the selected protocol.
 - Select **Specify** to specify the timeout for this protocol, in seconds.
 - Select **Immediate** to immediately apply this timeout to the protocol.
 - Select **Indefinite** to specify that this protocol never times out.
 - Select **Unspecified** to specify no timeout for the protocol. When this is selected, the default timeout for the protocol is used.
9. Click **Finished** to save the timer policy rule.

The timer policy is now configured to apply to traffic with this protocol type.

Select the timer policy in a service policy, and apply the service policy to a self IP, route domain, firewall rule, or firewall rule list.

Creating a port misuse policy

Create a port misuse policy to restrict traffic on a port to a specific application. You configure a policy with specific port, protocol, and service rules to specify when port misuse occurs, and what action the policy takes.

1. On the Main tab, click **Network > Service Policies > Port Misuse Policies**.
The Port Misuse screen opens.
2. Click **Create**.
The New Port Misuse Policy screen opens.
3. Type a name for the port misuse policy.
4. Type an optional description for the port misuse policy.
5. Select the **Default Actions** for the port misuse policy.
 - Select **Drop on Service Mismatch** to set a policy default that drops packets when the service does not match the port, as defined in the policy rules.
 - Select **Log on Service Mismatch** to set a policy default that logs service and port mismatches.
6. In the **Rule Name** field, type a name for a policy rule.
7. From the **Port** list, select a port for the port matching rule.
You can select from a list of commonly used ports, or select **Other** and specify a port number.
8. From the **IP Protocol** list, select the IP protocol for the port matching rule.
You can select **TCP**, **UDP**, or **SCTP**.
9. From the **Service** list, select the service.
This setting configures the association between the service and port number. Packets on this port that do not match the specified service type are dropped, if **Drop on Service Mismatch** is applied to this rule.

***Important:** You can specify a service on any port; you are not limited to customary port and service pairings. You can configure any service on any port as a rule in a port misuse policy.*

10. From the **Drop on Service Mismatch** field, select the drop behavior.
 - Select **Use Policy Default** to use the default action for packet drops, when the service does not match the port.
 - Select **Yes** to drop packets when the service does not match the port.
 - Select **No** to allow packets when the service does not match the port.
11. From the **Log on Service Mismatch** field, select the logging behavior.
 - Select **Use Policy Default** to use the default action for logging packet drops, when the service does not match the port.
 - Select **Yes** to log dropped packets when the service does not match the port.
 - Select **No** to not log packet drops when the service does not match the port.
12. Click **Finished** to save the port misuse policy.

The port misuse policy is now configured to drop packets for specified ports, when the service does not match.

Select the port misuse policy in a service policy, and apply the service policy to a self IP, route domain, firewall rule, or firewall rule list.

Creating a service policy

Create a service policy to apply custom timer policies and port misuse settings to self IPs, route domains, firewall rules, or firewall rule lists.

1. Click **Network > Service Policies**.
2. Click **Create**.
The New Service Policy screen opens.
3. Type a name for the service policy.
4. Type an optional description for the service policy.
5. To enable a timer policy in the service policy, in the Timer Policy area, click **Enabled**.
6. From the list, select a timer policy to use in the service policy. The Timer Policy Rules area shows the timer policy rules for the selected timer policy.
7. To enable a port misuse policy in the service policy, in the Port Misuse area, click **Enabled**.
8. From the list, select a port misuse policy to use in the service policy. The Port Misuse Policy Rules area shows the port misuse policy rules for the selected port misuse policy.
9. Click **Finished** to save the service policy and return to the service policies list screen.

The selected self IP now enforces or stages rules according to your selections.

Applying a service policy to a firewall rule

Apply a service policy to a firewall rule to apply custom timers and port misuse settings to traffic matched by the firewall rule.

1. Click **Security > Network Firewall > Active Rules**.
2. Select the service policy.

Option	Description
With the Inline Rules Editor	If you are using the inline rules editor, click in a rule to edit it, and select a service policy in the Action column.
With the standard rules editor	If you are using the standard rule editor, click a rule name and select a service policy from the Service Policy list.

3. Update the rule, or commit your changes.
4. Compile and deploy the changes, if you compile and deploy manually.

When the rule is compiled and deployed, the timeouts and port misuse settings defined in the service policy are applied to the rule.

Applying a service policy to a virtual server

Apply a service policy to a virtual server to use custom timers and port misuse settings on the virtual 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. In the **Destination Address** field, type the IP address in CIDR format.

The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is `ffe1::0020/64` or

Setting Timers and Preventing Port Misuse with Service Policies

2001:ed8:77b5:2:10:10:100:42/64. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. From the **Service policy** list, select the service policy.
7. Configure any other settings that you need.
8. Click **Finished**.

The service policy is now associated with the virtual server, and the timers and port misuse settings are applied to sessions on the virtual server.

Applying a service policy to a route domain

Apply a service policy to a route domain to apply custom timers and port misuse settings to traffic that uses the route domain.

1. On the Main tab, click **Network > Route Domains**.
The Route Domain List screen opens.
2. In the Name column, click the name of the relevant route domain.
3. Click the route domain to which you will apply the service policy.
4. From the **Service Policy** list, select the service policy to apply to the route domain.
5. Click **Update**

Traffic on the route domain that matches the rules defined in the service policy now uses the timeouts and port misuse settings defined in the timer and port misuse policies.

Applying a service policy to a self IP

Apply a service policy to a self IP to apply custom timers and port misuse settings to traffic that uses the self IP address.

1. On the Main tab, click **Network > Self IPs**.
2. In the Name column, click the self IP address that you want to modify.
This displays the properties of the self IP address.
3. Click the self IP to which you will apply the service policy.
4. From the **Service Policy** list, select the service policy to apply to the self IP.
5. Click **Update**

Traffic on the self IP that matches the rules defined in the service policy now uses the timeouts and port misuse settings defined in the timer and port misuse policies.

Testing Packets with Firewall, IP Intelligence, and DoS Rules

About packet tracing with the AFM Packet Tester

The Packet Tester is a troubleshooting tool that allows a user to inject a packet into the traffic processing of BIG-IP® AFM™ and track the resulting processing by the Network Firewall, DoS prevention settings, and IP Intelligence. If the packet hits an Network Firewall, DoS Protection, or IP Intelligence rule, the rule and rule context is displayed. This allows you to troubleshoot packet issues with certain types of packets, and to check that rules for certain packets are correctly configured.

Task list

Tracing a TCP packet

Tracing a UDP packet

Tracing an SCTP packet

Tracing an ICMP packet

Tracing a TCP packet

Before you can trace a TCP packet, you must have BIG-IP® Advanced Firewall Manager™ (AFM) licensed on your system.

You can test a TCP packet to find if it hits a Network Firewall, DoS, or IP Intelligence rule. For purposes of this test, you should test a packet that either represents a type of packet that is currently being dropped (for troubleshooting), or a type of packet you would like to detect with the Network Firewall, DoS rules, or IP Intelligence rules.

1. On the Main tab, click **Network > Network Security > Packet Tester**.
The Packet Tester screen opens.
2. From the **Protocol** list, select **TCP**.
3. Select any TCP flags to set in the TCP packet.
You can select **SYN, ACK, RST, URG, PUSH, FIN**, or a combination.
4. For the **Source** setting, specify the source **IP Address** from which the test packet should appear to originate.
5. Specify the source **Port** from which the test packet should appear to originate.
6. From the list select the source **VLAN** from which the test packet should appear to originate.
7. In the **TTL** field, specify the time to live for the test packet in seconds.
The default setting is 255 seconds.
8. For the **Destination** setting, specify the destination **IP Address** to which the test packet should appear to be sent.
9. In the **Destination** setting, specify the destination **Port** to which the test packet should appear to be sent.
10. In the **Trace Options** setting, specify whether to use the staged network firewall policy for the packet, if one exists.
11. In the **Trace Options** setting, specify whether to trigger logging for the packet, based on the packet test results.

12. Click **Run Trace** to run the packet test.

The packet trace displays the steps in the packet trace process, and the result of the packet trace.

Tracing a UDP packet

Before you can trace a UDP packet, you must have BIG-IP® Advanced Firewall Manager™ (AFM) licensed on your system.

You can test a UDP packet to find if it hits a Network Firewall, DoS, or IP Intelligence rule. For purposes of this test, you should test a packet that either represents a type of packet that is currently being dropped (for troubleshooting), or a type of packet you would like to detect with the Network Firewall, DoS rules, or IP Intelligence rules.

1. On the Main tab, click **Network > Network Security > Packet Tester**.
The Packet Tester screen opens.
2. From the **Protocol** list, select **UDP**.
3. For the **Source** setting, specify the source **IP Address** from which the test packet should appear to originate.
4. Specify the source **Port** from which the test packet should appear to originate.
5. From the list select the source **VLAN** from which the test packet should appear to originate.
6. In the **TTL** field, specify the time to live for the test packet in seconds.
The default setting is 255 seconds.
7. For the **Destination** setting, specify the destination **IP Address** to which the test packet should appear to be sent.
8. In the **Destination** setting, specify the destination **Port** to which the test packet should appear to be sent.
9. In the **Trace Options** setting, specify whether to use the staged network firewall policy for the packet, if one exists.
10. In the **Trace Options** setting, specify whether to trigger logging for the packet, based on the packet test results.
11. Click **Run Trace** to run the packet test.

The packet trace displays the steps in the packet trace process, and the result of the packet trace.

Tracing an SCTP packet

Before you can trace a UDP packet, you must have BIG-IP® Advanced Firewall Manager™ (AFM) licensed on your system.

You can test an SCTP packet to find if it hits a Network Firewall, DoS, or IP Intelligence rule. For purposes of this test, you should test a packet that either represents a type of packet that is currently being dropped (for troubleshooting), or a type of packet you would like to detect with the Network Firewall, DoS rules, or IP Intelligence rules.

1. On the Main tab, click **Network > Network Security > Packet Tester**.
The Packet Tester screen opens.
2. From the **Protocol** list, select **SCTP**.
3. For the **Source** setting, specify the source **IP Address** from which the test packet should appear to originate.
4. Specify the source **Port** from which the test packet should appear to originate.
5. From the list select the source **VLAN** from which the test packet should appear to originate.
6. In the **TTL** field, specify the time to live for the test packet in seconds.

The default setting is 255 seconds.

7. For the **Destination** setting, specify the destination **IP Address** to which the test packet should appear to be sent.
8. In the **Destination** setting, specify the destination **Port** to which the test packet should appear to be sent.
9. In the **Trace Options** setting, specify whether to use the staged network firewall policy for the packet, if one exists.
10. In the **Trace Options** setting, specify whether to trigger logging for the packet, based on the packet test results.
11. Click **Run Trace** to run the packet test.

The packet trace displays the steps in the packet trace process, and the result of the packet trace.

Tracing an ICMP packet

Before you can trace a UDP packet, you must have BIG-IP® Advanced Firewall Manager™ (AFM) licensed on your system.

You can test an ICMP packet to find if it hits a Network Firewall, DoS, or IP Intelligence rule. For purposes of this test, you should test a packet that either represents a type of packet that is currently being dropped (for troubleshooting), or a type of packet you would like to detect with the Network Firewall, DoS rules, or IP Intelligence rules.

1. On the Main tab, click **Network > Network Security > Packet Tester**.
The Packet Tester screen opens.
2. From the **Protocol** list, select **ICMP**.
3. From the **Protocol** list, select **SCTP**.
4. For the **Source** setting, specify the source **IP Address** from which the test packet should appear to originate.
5. From the list select the source **VLAN** from which the test packet should appear to originate.
6. In the **TTL** field, specify the time to live for the test packet in seconds.
The default setting is 255 seconds.
7. For the **Destination** setting, specify the destination **IP Address** to which the test packet should appear to be sent.
8. In the **Destination** setting, specify the destination **Port** to which the test packet should appear to be sent.
9. In the **Trace Options** setting, specify whether to use the staged network firewall policy for the packet, if one exists.
10. In the **Trace Options** setting, specify whether to trigger logging for the packet, based on the packet test results.
11. Click **Run Trace** to run the packet test.

The packet trace displays the steps in the packet trace process, and the result of the packet trace.

Packet trace results

These tables show possible results of an AFM packet trace.

Device DoS results

Device DoS result	Description
Nominal (Green)	The packet matches a vector, but is not categorized as an attack.

Device DoS result	Description
Whitelist (Green)	The packet matches the DoS whitelist and is allowed.
Anomaly (Yellow)	The packet matches an anomaly condition.
Attack (Red)	The packet matches a configured attack condition.

Device IP Intelligence results

Device IP Intelligence result	Description
No match (Green)	The packet does not match an IP Intelligence rule.
Match (Green or Red)	The packet matches an IP Intelligence rule and is either allowed or denied.
Whitelist (Green)	The packet matches the IP Intelligence whitelist and is allowed..
No Policy (Gray)	There is no configured IP intelligence policy for the packet

Device Rules

Device Rules result	Description
Match Allow (Green)	The packet matches a global firewall rule and is allowed.
Match Reject (Red)	The packet matches a global firewall rule and is rejected.
Match Drop (Red)	The packet matches a global firewall rule and is dropped.
Match Decisive (Green)	The packet matches a global firewall rule and is allowed decisively.
No Policy (Gray)	The packet does not match a global firewall rule.

Route Domain IP Intelligence results

Route Domain IP Intelligence result	Description
No match (Green)	The packet does not match a route domain Intelligence rule.
Match (Green or Red)	The packet matches a route domain Intelligence rule and is either allowed or denied.
Whitelist (Green)	The packet matches the route domain Intelligence whitelist and is allowed.
No Policy (Gray)	There is no configured IP intelligence policy for the packet

Route Domain Rules results

Route Domain Rules result	Description
Match Allow (Green)	The packet matches a route domain firewall rule and is allowed.
Match Reject (Red)	The packet matches a route domain firewall rule and is rejected.
Match Drop (Red)	The packet matches a route domain firewall rule and is dropped.
Match Decisive (Green)	The packet matches a route domain firewall rule and is allowed decisively.
No Policy (Gray)	The packet does not match a route domain firewall rule.

Virtual Server DoS results

Virtual Server DoS result	Description
Nominal (Green)	The packet matches a virtual server DoS vector, but is not categorized as an attack.
Whitelist (Green)	The packet matches the virtual server DoS whitelist and is allowed.
Anomaly (Yellow)	The packet matches a virtual server DoS anomaly condition.
Attack (Red)	The packet matches a configured virtual server DoS attack condition.
Prior Whitelist (Gray)	The packet matches a prior whitelist and is allowed.
No Policy (Gray)	No virtual server DoS rule is configured that applies to this packet.

Virtual Server IP Intelligence results

Virtual Server IP Intelligence result	Description
No match (Green)	The packet does not match a virtual server IP Intelligence rule.
Match (Green or Red)	The packet matches a virtual server IP Intelligence rule and is either allowed or denied.
Whitelist (Green)	The packet matches the virtual server IP Intelligence whitelist and is allowed.
No Policy (Gray)	No virtual server IP intelligence policy is configured that applies to this packet.

Virtual Server Rules results

Virtual Server Rules result	Description
Match Allow (Green)	The packet matches a virtual server firewall rule and is allowed.

Virtual Server Rules result	Description
Match Reject (Red)	The packet matches a virtual server firewall rule and is rejected.
Match Drop (Red)	The packet matches a virtual server firewall rule and is dropped.
Match Decisive (Green)	The packet matches a virtual server firewall rule and is allowed decisively.
No Policy (Gray)	The packet does not match a virtual server firewall rule.

Default Rule results

Default Rule result	Description
Allow (Green)	The packet does not match any prior rules, and the default rule is allow, so the packet is allowed.
Reject (Red)	The packet does not match any prior rules, and the default rule is reject, so the packet is rejected.
Drop (Red)	The packet does not match any prior rules, and the default rule is drop, so the packet is dropped.

Local Logging with the Network Firewall

Overview: Configuring local Network Firewall event logging

You can configure the BIG-IP® system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system.

Important: *The BIG-IP system Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure Network Firewall event logging.*

Task summary

Perform these tasks to configure Network Firewall logging locally on the BIG-IP® system.

Note: *Enabling logging and storing the logs locally impacts BIG-IP system performance.*

Creating a local Network Firewall Logging profile

Configuring a virtual server for Network Firewall event logging

Viewing Network Firewall event logs locally on the BIG-IP system

Creating a Network Firewall rule from a firewall log entry

Disabling logging

Creating a local Network Firewall Logging profile

You create a custom Logging profile to log BIG-IP® system Network Firewall events locally on the BIG-IP system.

1. On the Main tab, click **Security > Event Logs > Logging Profiles**.
The Logging Profiles list screen opens.
2. Click **Create**.
The Create New Logging Profile screen opens.
3. In the **Profile Name** field, type a unique name for the profile.
4. Select the **Network Firewall** check box.
5. If you want to enable optional subscriber ID logging:
 - a) Select the **Network Address Translation** check box.
 - b) Then in the Network Address Translation area, select the **Log Subscriber ID** check box.
 - c) Click **Network Firewall**.
6. In the Network Firewall area, from the **Publisher** list, select **local-db-publisher**.
7. Set an **Aggregate Rate Limit** to define a rate limit for all combined network firewall log messages per second.
Beyond this rate limit, log messages are not logged.
8. For the **Log Rule Matches** setting, select how the BIG-IP system logs packets that match ACL rules.
You can select any or all of the options.

Option	Description
---------------	--------------------

Option	Enables or disables logging of packets that match ACL rules configured with:
---------------	--

Option	Description
Accept	action=Accept
Drop	action=Drop
Reject	action=Reject

When an option is selected, you can configure a rate limit for log messages of that type.

9. Select the **Log IP Errors** check box, to enable logging of IP error packets.
When this setting is enabled, you can configure a rate limit for log messages of this type.
10. Select the **Log TCP Errors** check box, to enable logging of TCP error packets.
When this is enabled, you can configure a rate limit for log messages of this type.
11. Select the **Log TCP Events** check box, to enable logging of open and close of TCP sessions.
When this is enabled, you can configure a rate limit for log messages of this type.
12. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for Network Firewall log events.
13. Enable the **Always Log Region** setting to log the geographic location when a geolocation event causes a network firewall event.
14. From the **Storage Format** list, select how the BIG-IP system formats the log.

Option	Description
None	Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example: "management_ip_address", "bigip_hostname", "context_type", "context_name", "src_ip", "dest_ip", "src_port", "dest_port", "vlan", "protocol", "route_domain", "acl_rule_name", "action", "drop_reason"
Field-List	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Specify the order the fields display in the log. • Specify the delimiter that separates the content in the log. The default delimiter is the comma character.
User-Defined	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Cut and paste, in a string of text, the order the fields display in the log.

15. In the IP Intelligence area, from the **Publisher** list, select **local-db-publisher**.

Note: The IP Address Intelligence feature must be enabled and licensed.

16. Set an **Aggregate Rate Limit** to define a rate limit for all combined IP Intelligence log messages per second.
Beyond this rate limit, log messages are not logged.
17. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for IP Intelligence log events.
18. In the Traffic Statistics area, from the **Publisher** list, select **local-db-publisher**.
19. For the **Log Timer Events** setting, enable **Active Flows** to log the number of active flows each second.
20. For the **Log Timer Events** setting, enable **Reaped Flows** to log the number of reaped flows, or connections that are not established because of system resource usage levels.

21. For the **Log Timer Events** setting, enable **Missed Flows** to log the number of packets that were dropped because of a flow table miss. A flow table miss occurs when a TCP non-SYN packet does not match an existing flow.
22. For the **Log Timer Events** setting, enable **SYN Cookie (Per Session Challenge)** to log the number of SYN cookie challenges generated each second.
23. For the **Log Timer Events** setting, enable **SYN Cookie (White-listed Clients)** to log the number of SYN cookie clients whitelisted each second.
24. Click **Finished**.

Now you should assign this custom Network Firewall Logging profile to a virtual server.

Configuring a virtual server for Network Firewall event logging

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom Network Firewall Logging profile to a virtual server when you want the BIG-IP system to log Network Firewall events on the traffic that the virtual server processes.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.

*Note: If you do not have a custom profile configured, select the predefined logging profile **global-network** to log Advanced Firewall Manager™ events. Note that to log global, self IP, and route domain contexts, you must enable a Publisher in the **global-network** profile.*

5. Click **Update** to save the changes.

Viewing Network Firewall event logs locally on the BIG-IP system

Ensure that the BIG-IP® system is configured to log the types of events you want to view, and to store the log messages locally on the BIG-IP system.

When the BIG-IP system is configured to log events locally, you can view those events using the Configuration utility.

1. On the Main tab, click **Security > Event Logs > Network > Firewall**.
The Network Firewall event log displays.
2. To search for specific events, click **Custom Search**. Drag the event data that you want to search for from the Event Log table into the Custom Search table, and then click **Search**.

Creating a Network Firewall rule from a firewall log entry

You must be logging Network Firewall traffic to create a rule from the Network Firewall logs.

You can create a rule from the local log, from an enforced or staged rule or policy. You might use this to change the action taken on specific traffic that is matched by a more general rule. You can also use this to replicate a rule and change some parameter, such as the source or destination ports. Note that the rule you create from a log entry already has some information specified, such as source and destination address and ports, protocol, and VLAN. You can change any of this information as required.

1. On the Main tab, click **Security > Event Logs > Network > Firewall**.
The Network Firewall event log displays.

2. Select the search parameters to show the preferred log results, then click **Search**.
3. Select a log entry, and click **Create Rule**.
4. From the **Context** list, select the context for the firewall rule.
For a firewall rule in a rule list, the context is predefined and cannot be changed.
5. In the **Name** and **Description** fields, type the name and an optional description.
6. From the **Type** list, select whether you are creating a standalone network firewall rule or creating the rule from a predefined rule list.

***Note:** If you create a firewall rule from a predefined rule list, only the **Name**, **Description**, **Order**, **Rule List**, and **State** options apply, and you must select or create a rule list to include.*

7. From the **State** list, select the rule state.
 - Select **Enabled** to apply the firewall rule to the given context and addresses.
 - Select **Disabled** to set the firewall rule to not apply at all.
 - Select **Scheduled** to apply the firewall rule according to the selected schedule.
8. From the **Schedule** list, select the schedule for the firewall rule.
This schedule is applied when you set the firewall rule state as **Scheduled**.
9. From the **Protocol** list, select the protocol to which the firewall rule applies.
 - Select **Any** to apply the firewall rule to any protocol.
 - Select the protocol name to apply the rule to a single protocol.

***Important:** ICMP is handled by the BIG-IP system at the global or route domain level. Because of this, ICMP messages receive a response before they reach the virtual server context. You cannot create rule for ICMP or ICMPv6 on a self IP or virtual server context. You can apply a rule list to a self IP or virtual server that includes a rule for ICMP or ICMPv6; however, such a rule will be ignored. To apply firewall actions to the ICMP protocol, create a rule with the global or route domain context. ICMP rules are evaluated only for ICMP forwarding requests, and not for the IP addresses of the BIG-IP system itself.*

10. In the **Source** list, specify users and groups to which this rule applies.
 - From the **User** list, select **Any** to have the rule apply to any user.
 - From the **User** list, select **Specify** and click **User**, **Group**, or **User List** to specify a user, group, or user list packet source to which the rule applies. When selected, you can type a user or group name in the format `domain\user_name` or `domain\group_name`. You can specify a user list by selecting it from the list. Click **Add** to add a selected user, group, or user list to the packet source list.
11. In the **Source** list, specify addresses and geolocated sources to which this rule applies.
 - From the **Address/Region** list, select **Any** to have the rule apply to any packet source IP address or geographic location.
 - From the **Address/Region** list, select **Specify** and click **Address** to specify one or more packet source IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - From the **Address/Region** list, select **Specify** and click **Address List** to select a predefined list of packet source addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - From the **Address/Region** list, select **Specify** and click **Address Range** to specify a contiguous range of packet source IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.

- From the **Address/Region** list, select **Specify** and click **Country/Region** to identify the geographic origin of packet sources, and to apply rules based on selected geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Source address list.
- 12.** From the Source **Port** list, select the type of packet source ports to which this rule applies.
- Select **Any** to have the rule apply to any packet source port.
 - Select **Specify** and click **Port** to specify one or more packet source ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet source port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.
 - Select **Specify** and click **Port List** to select a predefined list of packet source ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
- 13.** From the Source **VLAN/Tunnel** list, select the VLAN on which this rule applies.
- Select **Any** to have the rule apply to traffic on any VLAN through which traffic enters the firewall.
 - Select **Specify** to specify one or more VLANs on the firewall to which the rule applies. To use a VLAN with this rule, move the VLAN from the **Available** list to the **Selected** list. Similarly, you can remove the VLAN from this rule, by moving the VLAN from the **Selected** list to the **Available** list.
- 14.** In the Destination area and from the **Address/Region** list, select the type of packet destination address to which this rule applies.
- Select **Any** to have the rule apply to any IP packet destination address.
 - Select **Specify** and click **Address** to specify one or more packet destination IP addresses or fully qualified domain names (FQDNs) to which the rule applies. When selected, you can type single IP addresses or FQDNs into the **Address** field, then click **Add** to add them to the address list.
 - Select **Specify** and click **Address List** to select a predefined list of packet destination addresses to which the rule applies. To use an address list with this rule, select the address list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
 - Select **Specify** and click **Address Range** to specify a contiguous range of packet destination IP addresses to which the rule applies. When selected, you can type a start and end IP address in the fields, then click **Add** to add the IP address range to the address list.
 - Select **Specify** and click **Country/Region** to identify the geographic packet destination, and to apply rules based on specific geographic locations. When selected, a field appears in which you can select a country. For many countries, an extra field appears after you select the country, in which you can select a state or province. If you do not select a specific state or province, the entire country is selected. After you select a geographic location, click **Add** to add it to the Destination address list.
- 15.** From the Destination **Port** list, select the type of packet destination ports to which this rule applies.
- Select **Any** to have the rule apply to any port inside the firewall.
 - Select **Specify** and click **Port** to specify one or more packet destination ports to which the rule applies. When selected, you can type single port numbers into the **Port** field, then click **Add** to add them to the port list.
 - Select **Specify** and click **Port Range** to specify a list of contiguous packet destination port numbers to which the rule applies. When selected, you can type the start and end ports into the fields, then click **Add** to add the ports to the port list.

- Select **Specify** and click **Port List** to select a predefined list of packet destination ports to which the rule applies. To use a port list with this rule, select the port list and click the **Add** button. Similarly, to remove the list from this rule, select the list and click the **Delete** button.
16. Optionally, to apply an iRule to traffic matched by this rule, from the **iRule** list, select an iRule.
 17. When you select an iRule to start in a firewall rule, you can enable iRule sampling, and select how frequently the iRule is started, for sampling purposes. The value you configure is one out of n times the iRule is triggered. For example, to trigger the iRule one out of every five times the rule matches a flow, select **Enabled**, then set this field to 5.
 18. From the **Action** list, select the firewall action for traffic originating from the specified source address on the specified protocol. Choose from one of the these actions:

Option	Description
Accept	Allows packets with the specified source, destination, and protocol to pass through the firewall. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached.
Reject	Rejects packets with the specified source, destination, and protocol. When a packet is rejected the firewall sends a destination unreachable message to the sender.
Accept Decisively	Allows packets with the specified source, destination, and protocol to pass through the firewall, and does not require any further processing by any of the further firewalls. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
 19. From the **Logging** list, enable or disable logging for the firewall rule.

A logging profile must be enabled to capture logging info for the firewall rule.
 20. Click **Finished**.

The list screen and the new items are displayed.

The new firewall policy rule is created from the log entry.

Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

Note: You can disable and re-enable logging for a specific resource based on your network administration needs.

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

The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.

The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

Implementation result

You now have an implementation in which the BIG-IP[®] system logs specific Network Firewall events and stores the logs in a local database on the BIG-IP system.

Remote High-Speed Logging with the Network Firewall

Overview: Configuring remote high-speed Network Firewall event logging

You can configure the BIG-IP® system to log information about the BIG-IP system Network Firewall events and send the log messages to remote high-speed log servers.

Important: The BIG-IP system Advanced Firewall Manager™ (AFM™) must be licensed and provisioned before you can configure Network Firewall event logging.

This illustration shows the association of the configuration objects for remote high-speed logging.




Figure 6: Association of remote high-speed logging configuration objects

Task summary

Perform these tasks to configure remote high-speed network firewall logging on the BIG-IP® system.

Note: Enabling remote high-speed logging impacts BIG-IP system performance.

Creating a pool of remote logging servers

Creating a remote high-speed log destination

Creating a formatted remote high-speed log destination

Creating a publisher

Creating a custom Network Firewall Logging profile

Configuring a virtual server for Network Firewall event logging

Disabling logging

About the configuration objects of remote high-speed Network Firewall event logging

When configuring remote high-speed logging of Network Firewall events, it is helpful to understand the objects you need to create and why, as described here:

Object	Reason	Applies to
Pool of remote log servers	Create a pool of remote log servers to which the BIG-IP® system can send log messages.	Creating a pool of remote logging servers.
Destination (unformatted)	Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers.	Creating a remote high-speed log destination.
Destination (formatted)	If your remote log servers are the ArcSight, Splunk, IPFIX, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination.	Creating a formatted remote high-speed log destination.
Publisher	Create a log publisher to send logs to a set of specified log destinations.	Creating a publisher.
DNS Logging profile	Create a custom DNS Logging profile to define the data you want the BIG-IP system to include in the DNS logs and associate a log publisher with the profile.	Creating a custom Network Firewall Logging profile.
LTM® virtual server	Associate a custom DNS profile with a virtual server to define how the BIG-IP system logs the DNS traffic that the virtual server processes.	Creating a virtual server for Network Firewall event logging.

Creating a pool of remote logging servers

Before creating a pool of log servers, gather the IP addresses of the servers that you want to include in the pool. Ensure that the remote log servers are configured to listen to and receive log messages from the BIG-IP® system.

Create a pool of remote log servers to which the BIG-IP system can send log messages.

1. On the Main tab, click the applicable path.
 - **DNS > Delivery > Load Balancing > Pools**
 - **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 the IP address for each remote logging server that you want to include in the pool:
 - a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
 - b) Type a service number in the **Service Port** field, or select a service name from the list.

Note: Typical remote logging servers require port 514.

- c) Click **Add**.
5. Click **Finished**.

Creating a remote high-speed log destination

Before creating a remote high-speed log destination, ensure that at least one pool of remote log servers exists on the BIG-IP® system.

Create a log destination of the **Remote High-Speed Log** type to specify that log messages are sent to a pool of remote log servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **Remote High-Speed Log**.

***Important:** If you use log servers such as Remote Syslog, Splunk, or ArcSight, which require data be sent to the servers in a specific format, you must create an additional log destination of the required type, and associate it with a log destination of the **Remote High-Speed Log** type. With this configuration, the BIG-IP system can send data to the servers in the required format.*

The BIG-IP system is configured to send an unformatted string of text to the log servers.

5. From the **Pool Name** list, select the pool of remote log servers to which you want the BIG-IP system to send log messages.
6. From the **Protocol** list, select the protocol used by the high-speed logging pool members.
7. Click **Finished**.

Creating a formatted remote high-speed log destination

Ensure that at least one remote high-speed log destination exists on the BIG-IP® system.

Create a formatted logging destination to specify that log messages are sent to a pool of remote log servers, such as Remote Syslog, Splunk, or ArcSight servers.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select a formatted logging destination, such as **IPFIX**, **Remote Syslog**, **Splunk**, or **ArcSight**.

***Important:** ArcSight formatting is only available for logs coming from Advanced Firewall Manager™ (AFM™), Application Security Manager™ (ASM™), and the Secure Web Gateway component of Access Policy Manager® (APM®). IPFIX is not available for Secure Web Gateway. Remote Syslog formatting*

is the only type supported for logs coming from APM. The Splunk format is a predefined format of key value pairs.

The BIG-IP system is configured to send a formatted string of text to the log servers.

5. If you selected **Remote Syslog**, then from the **Syslog Format** list select a format for the logs, and then from the **High-Speed Log Destination** list, select the destination that points to a pool of remote Syslog servers to which you want the BIG-IP system to send log messages.
-

Important: For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.

6. If you selected **Splunk** or **IPFIX**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.
7. Click **Finished**.

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 custom Network Firewall Logging profile

Create a custom Logging profile to log messages about BIG-IP® system Network Firewall events.

1. On the Main tab, click **Security > Event Logs > Logging Profiles**.
The Logging Profiles list screen opens.
2. Click **Create**.
The Create New Logging Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select the **Network Firewall** check box.
5. In the Network Firewall area, from the **Publisher** list, select the publisher the BIG-IP system uses to log Network Firewall events.
6. Set an **Aggregate Rate Limit** to define a rate limit for all combined network firewall log messages per second.
Beyond this rate limit, log messages are not logged.
7. For the **Log Rule Matches** setting, select how the BIG-IP system logs packets that match ACL rules. You can select any or all of the options.

Option	Description
---------------	--------------------

Option	Enables or disables logging of packets that match ACL rules configured with:
---------------	--

Option	Description
Accept	action=Accept
Drop	action=Drop
Reject	action=Reject

When an option is selected, you can configure a rate limit for log messages of that type.

8. Select the **Log IP Errors** check box, to enable logging of IP error packets.
When this setting is enabled, you can configure a rate limit for log messages of this type.
9. Select the **Log TCP Errors** check box, to enable logging of TCP error packets.
When this is enabled, you can configure a rate limit for log messages of this type.
10. Select the **Log TCP Events** check box, to enable logging of open and close of TCP sessions.
When this is enabled, you can configure a rate limit for log messages of this type.
11. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for Network Firewall log events.
12. Enable the **Log Geolocation IP Address** setting to specify that when a geolocation event causes a network firewall action, the associated IP address is logged.
13. From the **Storage Format** list, select how the BIG-IP system formats the log.

Option	Description
None	Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example: "management_ip_address", "bigip_hostname", "context_type", "context_name", "src_ip", "dest_ip", "src_port", "dest_port", "vlan", "protocol", "route_domain", "acl_rule_name", "action", "drop_reason"
Field-List	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Specify the order the fields display in the log. • Specify the delimiter that separates the content in the log. The default delimiter is the comma character.
User-Defined	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Cut and paste, in a string of text, the order the fields display in the log.

14. In the IP Intelligence area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log source IP addresses, which are identified and configured for logging by an IP Intelligence policy.

Note: The IP Address Intelligence feature must be enabled and licensed.

15. Set an **Aggregate Rate Limit** to define a rate limit for all combined IP Intelligence log messages per second.
Beyond this rate limit, log messages are not logged.
16. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for IP Intelligence log events.
17. In the Traffic Statistics area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log traffic statistics.
18. For the **Log Timer Events** setting, enable **Active Flows** to log the number of active flows each second.

19. For the **Log Timer Events** setting, enable **Reaped Flows** to log the number of reaped flows, or connections that are not established because of system resource usage levels.
20. For the **Log Timer Events** setting, enable **Missed Flows** to log the number of packets that were dropped because of a flow table miss. A flow table miss occurs when a TCP non-SYN packet does not match an existing flow.
21. For the **Log Timer Events** setting, enable **SYN Cookie (Per Session Challenge)** to log the number of SYN cookie challenges generated each second.
22. For the **Log Timer Events** setting, enable **SYN Cookie (White-listed Clients)** to log the number of SYN cookie clients whitelisted each second.
23. Click **Finished**.

Assign this custom network firewall Logging profile to a virtual server.

Configuring a virtual server for Network Firewall event logging

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom Network Firewall Logging profile to a virtual server when you want the BIG-IP system to log Network Firewall events on the traffic that the virtual server processes.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles, and move them from the **Available** list to the **Selected** list.

*Note: If you do not have a custom profile configured, select the predefined logging profile **global-network** to log Advanced Firewall Manager™ events. Note that to log global, self IP, and route domain contexts, you must enable a Publisher in the **global-network** profile.*

5. Click **Update** to save the changes.

Disabling logging

Disable Network Firewall, Protocol Security, or DoS Protection event logging when you no longer want the BIG-IP® system to log specific events on the traffic handled by specific resources.

Note: You can disable and re-enable logging for a specific resource based on your network administration needs.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Disabled**.
5. Click **Update** to save the changes.

The BIG-IP system does not log the events specified in this profile for the resources to which this profile is assigned.

Implementation result

You now have an implementation in which the BIG-IP[®] system logs specific Network Firewall events and sends the logs to a remote log server.

SNMP Trap Configuration

Overview: BIG-IP SNMP agent configuration

You can use the industry-standard SNMP protocol to manage BIG-IP® devices on a network. To do this, you must configure the SNMP agent on the BIG-IP system. The primary tasks in configuring the SNMP agent are configuring client access to the SNMP agent, and controlling access to SNMP data.

Task summary

Perform these tasks to configure SNMP on the BIG-IP system.

Specify SNMP administrator contact information and system location information

Configure SNMP manager access to the SNMP agent on the BIG-IP system

Grant community access to v1 or v2c SNMP data

Grant user access to v3 SNMP data

Specify SNMP administrator contact information and system location information

You specify contact information for the SNMP administrator, as well as the physical location of the BIG-IP® system running an SNMP agent.

1. On the Main tab, click **System > SNMP > Agent > Configuration**.
2. In the Global Setup area, in the **Contact Information** field, type contact information for the SNMP administrator for this BIG-IP system.
The contact information is a MIB-II simple string variable. The contact information usually includes both a user name and an email address.
3. In the **Machine Location** field, type the location of the system, such as `Network Closet 1`.
The machine location is a MIB-II simple string variable.
4. Click **Update**.

Configure SNMP manager access to the SNMP agent on the BIG-IP system

Before you start this task, you should gather the IP addresses of the SNMP managers that you want to have access to the SNMP agent on this BIG-IP® system.

You configure the SNMP agent on the BIG-IP system so that a client running the SNMP manager can access the SNMP agent to remotely manage the BIG-IP system.

1. On the Main tab, click **System > SNMP > Agent > Configuration**.
2. In the **Client Allow List** area, for the **Type** setting, select either **Host** or **Network**, depending on whether the IP address you specify is a host system or a subnet.

Note: By default, SNMP is enabled only for the BIG-IP system loopback interface (127.0.0.1).

3. In the **Address** field, type either an IP address or network address from which the SNMP agent can accept requests.
4. If you selected **Network** in step 2, type the netmask in the **Mask** field.
5. Click **Add**.
6. Click **Update**.

The BIG-IP system now contains a list of IP addresses for SNMP managers from which SNMP requests are accepted.

Grant community access to v1 or v2c SNMP data

To better control access to SNMP data, you can assign an access level to an SNMP v1 or v2c community.

Note: SNMPv1 does not support Counter64 OIDs, which are used for accessing most statistics. Therefore, for SNMPv1 clients, an `snmp walk` command skips any OIDs of type Counter64. We recommend that you use only clients that support SNMPv2 or later.

1. On the Main tab, click **System > SNMP > Agent > Access (v1, v2c)**.
2. Click **Create**.
3. From the **Type** list, select either **IPv4** or **IPv6**.
4. In the **Community** field, type the name of the SNMP community for which you are assigning an access level.
5. From the **Source** list, select **All**, or select **Select** and type the source IP address in the field that displays.
6. In the **OID** field, type the OID for the top-most node of the SNMP tree to which the access applies.
7. From the **Access** list, select an access level, either **Read Only** or **Read/Write**.

Note: When you set the access level of a community or user to read/write, and an individual data object has a read-only access type, access to the object remains read-only. In short, the access level or type that is the most secure takes precedence when there is a conflict.

8. Click **Finished**.

The BIG-IP® system updates the `snmpd.conf` file, assigning only a single access setting to the community as shown in this sample `snmpd.conf` file.

Example `snmpd.conf` file

In the following sample code from an `snmpd.conf` file, string `rocommunity public default` identifies a community named `public` that has the default read-only access-level. This access-level prevents any allowed SNMP manager in community `public` from modifying a data object, even if the object has an access type of read/write. The string `rwcommunity public1` identifies a community named `public1` as having a read/write access-level. This access-level allows any allowed SNMP manager in community `public1` to modify a data object under the tree node `.1.3.6.1.4.1.3375.2.2.10.1 (ltmVirtualServ)` on the local host `127.0.0.1`, if that data object has an access type of read/write.

```
rocommunity public default
rwcommunity public1 127.0.0.1 .1.3.6.1.4.1.3375.2.2.10.1
```

Grant user access to v3 SNMP data

To better control access to SNMP data, you can assign an access level to an SNMP v3 user.

1. On the Main tab, click **System > SNMP > Agent > Access (v3)**.
2. Click **Create**.
3. In the **User Name** field, type the name of the user for which you are assigning an access level.

4. In the Authentication area, from the **Type** list, select a type of authentication to use, and then type and confirm the user's password.
5. In the Privacy area, from the **Protocol** list, select a privacy protocol, and either type and confirm the user's password, or select the **Use Authentication Password** check box.
6. In the **OID** field, type the OID for the top-most node of the SNMP tree to which the access applies.
7. From the **Access** list, select an access level, either **Read Only** or **Read/Write**.

Note: When you set the access level of a community or user to read/write, and an individual data object has a read-only access type, access to the object remains read-only. In short, the access level or type that is the most secure takes precedence when there is a conflict.

8. Click **Finished**.

The BIG-IP® system updates the `snmpd.conf` file, assigning only a single access setting to the user.

Overview: SNMP trap configuration

SNMP *traps* are definitions of unsolicited notification messages that the BIG-IP® alert system and the SNMP agent send to the SNMP manager when certain events occur on the BIG-IP system. Configuring SNMP traps on a BIG-IP system means configuring how the BIG-IP system handles traps, as well as setting the destination to which the notifications are sent.

The BIG-IP system stores SNMP traps in two specific files:

`/etc/alertd/alert.conf`
Contains default SNMP traps.

Important: Do not add or remove traps from the `/etc/alertd/alert.conf` file.

`/config/user_alert.conf`
Contains user-defined SNMP traps.

Task summary

Perform these tasks to configure SNMP traps for certain events and set trap destinations.

- Enabling traps for specific events*
- Setting v1 and v2c trap destinations*
- Setting v3 trap destinations*
- Viewing pre-configured SNMP traps*
- Creating custom SNMP traps*

Enabling traps for specific events

You can configure the SNMP agent on the BIG-IP® system to send, or refrain from sending, notifications to the traps destinations.

1. On the Main tab, click **System > SNMP > Traps > Configuration**.
2. To send traps when an administrator starts or stops the SNMP agent, verify that the **Enabled** check box for the **Agent Start/Stop** setting is selected.
3. To send notifications when authentication warnings occur, select the **Enabled** check box for the **Agent Authentication** setting.
4. To send notifications when certain warnings occur, verify that the **Enabled** check box for the **Device** setting is selected.

5. Click **Update**.

The BIG-IP system automatically updates the `alert.conf` file.

Setting v1 and v2c trap destinations

You specify the IP address of the SNMP manager in order for the BIG-IP® system to send notifications.

1. On the Main tab, click **System > SNMP > Traps > Destination**.
2. Click **Create**.
3. For the **Version** setting, select either `v1` or `v2c`.
4. In the **Community** field, type the community name for the SNMP agent running on the BIG-IP system.
5. In the **Destination** field, type the IP address of the SNMP manager.
6. In the **Port** field, type the port number on the SNMP manager that is assigned to receive the traps.
7. For the **Network** setting, select a trap network.
The BIG-IP system sends the SNMP trap out of the network you select.
8. Click **Finished**.

Setting v3 trap destinations

You specify the destination SNMP manager to which the BIG-IP® system sends notifications.

1. On the Main tab, click **System > SNMP > Traps > Destination**.
2. Click **Create**.
3. For the **Version** setting, select `v3`.
4. In the **Destination** field, type the IP address of the SNMP manager.
5. In the **Port** field, type the port number on the SNMP manager that is assigned to receive the traps.
6. For the **Network** setting, select a trap network.
The BIG-IP system sends the SNMP trap out of the network you select.
7. From the **Security Level** list, select the level of security at which you want SNMP messages processed.

Option	Description
Auth, No Privacy	Process SNMP messages using authentication but without encryption. When you use this value, you must also provide values for the Security Name , Authentication Protocol , and Authentication Password settings.
Auth and Privacy	Process SNMP messages using authentication and encryption. When you use this value, you must also provide values for the Security Name , Authentication Protocol , Authentication Password , Privacy Protocol , and Privacy Password settings.

8. In the **Security Name** field, type the user name the system uses to handle SNMP v3 traps.
9. In the **Engine ID** field, type an administratively unique identifier for an SNMP engine. (This setting is optional.) You can find the engine ID in the `/config/net-snmp/snmpd.conf` file on the BIG-IP system.
Note that this ID is identified in the file as the value of the `oldEngineID` token.
10. From the **Authentication Protocol** list, select the algorithm the system uses to authenticate SNMP v3 traps.
When you set this value, you must also enter a value in the **Authentication Password** field.

11. In the **Authentication Password** field, type the password the system uses to handle an SNMP v3 trap.

When you set this value, you must also select a value from the **Authentication Protocol** list.

Note: The authentication password must be at least 8 characters long.

12. If you selected **Auth and Privacy** from the **Security Level** list, from the **Privacy Protocol** list, select the algorithm the system uses to encrypt SNMP v3 traps.

When you set this value, you must also enter a value in the **Privacy Password** field.

13. If you selected **Auth and Privacy** from the **Security Level** list, in the **Privacy Password** field, type the password the system uses to handle an encrypted SNMP v3 trap.

When you set this value, you must also select a value from the **Privacy Protocol** list.

Note: The authentication password must be at least 8 characters long.

14. Click **Finished**.

Viewing pre-configured SNMP traps

Verify that your user account grants you access to the advanced shell.

Pre-configured traps are stored in the `/etc/alertd/alert.conf` file. View these SNMP traps to understand the data that the SNMP manager can use.

Use this command to view the SNMP traps that are pre-configured on the BIG-IP® system:

```
cat /etc/alertd/alert.conf.
```

Creating custom SNMP traps

Verify that your user account grants you access to tmsh.

Create custom SNMP traps that alert the SNMP manager to specific SNMP events that occur on the network when the pre-configured traps do not meet all of your needs.

1. Log in to the command line.
2. Create a backup copy of the file `/config/user_alert.conf`, by typing this command: `cp /config/user_alert.conf backup_file_name`
For example, type: `cp /config/user_alert.conf /config/user_alert.conf.backup`
3. With a text editor, open the file `/config/user_alert.conf`.
4. Add a new SNMP trap.

The required format is:

```
alert alert_name "matched message" {
    snmptrap OID=".1.3.6.1.4.1.3375.2.4.0.XXX"
}
```

- `alert_name` represents a descriptive name. The `alert_name` or `matched_message` value cannot match the corresponding value in any of the SNMP traps defined in the `/etc/alertd/alert.conf` or `/config/user_alert.conf` file.
- `matched_message` represents the text that matches the Syslog message that triggers the custom trap. You can specify either a portion of the Syslog message text or use a regular expression. Do not include the Syslog prefix information, such as the date stamp and process ID, in the match string.

- The *xxx* portion of the OID value represents a number that is unique to this OID. Specify any OID that meets all of these criteria:
 - Is in standard OID format and within the range .1.3.6.1.4.1.3375.2.4.0.300 through .1.3.6.1.4.1.3375.2.4.0.999.
 - Is in a numeric range that can be processed by your trap receiving tool.
 - Does not exist in the MIB file `/usr/share/snmp/mibs/F5-BIGIP-COMMON-MIB.txt`.
 - Is not used in another custom trap.

As an example, to create a custom SNMP trap that is triggered whenever the system logs switchboard failsafe status changes, add the following trap definition to `/config/user_alert.conf`.

```
alert SWITCHBOARD_FAILSAFE_STATUS "Switchboard Failsafe (.*)" {
    snmptrap OID=".1.3.6.1.4.1.3375.2.4.0.500"
}
```

This trap definition causes the system to log the following message to the file `/var/log/ltn`, when switchboard failsafe is enabled: `Sep 23 11:51:40 bigip1.askf5.com lacpd[27753]: 01160016:6: Switchboard Failsafe enabled.`

5. Save the file.
6. Close the text editor.
7. Restart the `alertd` daemon by typing this command: `bigstart restart alertd`
If the `alertd` daemon fails to start, examine the newly-added trap entry to ensure that the format is correct.

Overview: About troubleshooting SNMP traps

When the BIG-IP® alert system and the SNMP agent send traps to the SNMP manager, you can respond to the alert using the recommended actions for each SNMP trap.

AFM-related traps and recommended actions

This table provides information about the AFM™-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
BIGIP_TMM_TMMERR_DOS_ATTACK_START (.1.3.6.1.4.1.3375.2.4.0.133)	The start of a possible DoS attack was registered.	Determine your response to this type of DoS attack, if required.
BIGIP_TMM_TMMERR_DOS_ATTACK_STOP (.1.3.6.1.4.1.3375.2.4.0.134)	The end of a possible DoS attack was detected.	None, informational.
BIGIP_DOSPROTECT_DOSPROTECT_AGGREAPEROID (.1.3.6.1.4.1.3375.2.4.0.22)	The flow sweeper started or stopped.	None, informational.

AOM-related traps and recommended actions

This table provides information about the Always-On Management (AOM)-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipLibhalAomEventWarning (.1.3.6.1.4.1.3375.2.4.0.167)	AOM has issued a warning event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomEventError (.1.3.6.1.4.1.3375.2.4.0.168)	AOM has issued an error event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomEventAlert (.1.3.6.1.4.1.3375.2.4.0.169)	AOM has issued an alert event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomEventCritical (.1.3.6.1.4.1.3375.2.4.0.170)	AOM has issued a critical event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomEventEmergency (.1.3.6.1.4.1.3375.2.4.0.171)	AOM has issued an emergency event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomEventInfo (.1.3.6.1.4.1.3375.2.4.0.172)	AOM has issued an information event.	Inspect the <code>/var/log/ltn</code> file for additional messages that might provide further clarity on why the warning was raised.
bigipLibhalAomSensorTempWarning (.1.3.6.1.4.1.3375.2.4.0.173)	AOM has issued a temperature sensor warning level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorTempError (.1.3.6.1.4.1.3375.2.4.0.174)	AOM has issued a temperature sensor warning level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorTempAlert (.1.3.6.1.4.1.3375.2.4.0.175)	AOM has issued a temperature sensor alert level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper

Trap name	Description	Recommended action
		airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorTempCritical (. 1.3.6.1.4.1.3375.2.4.0.176)	AOM has issued a temperature sensor critical level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorTempEmergency (. 1.3.6.1.4.1.3375.2.4.0.177)	AOM has issued a temperature sensor emergency level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorTempInfo (. 1.3.6.1.4.1.3375.2.4.0.178)	AOM has issued a temperature sensor information level event.	Check the fan status from the output of your <code>tmsh show sys hardware query</code> , and see if any are down. Make sure the system has proper airflow. Verify that the unit has a sufficiently cool ambient room temperature.
bigipLibhalAomSensorFanWarning (. 1.3.6.1.4.1.3375.2.4.0.179)	AOM has issued a fan sensor warning level event.	Inspect the system for anything obstructing the system fans. Ensure that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorFanError (. 1.3.6.1.4.1.3375.2.4.0.180)	AOM has issued a fan sensor error level event.	Inspect the system for anything obstructing the system fans. Ensure that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorFanAlert (. 1.3.6.1.4.1.3375.2.4.0.181)	AOM has issued a fan sensor alert level event.	Inspect the system for anything obstructing the system fans. Ensure that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorFanCritical (. 1.3.6.1.4.1.3375.2.4.0.182)	AOM has issued a fan sensor critical level event.	Inspect the system for anything obstructing the system fans. Ensure that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorFanEmergency (. 1.3.6.1.4.1.3375.2.4.0.183)	AOM has issued a fan sensor emergency level event.	Inspect the system for anything obstructing the system fans. Ensure that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorFanInfo (. 1.3.6.1.4.1.3375.2.4.0.184)	AOM has issued a fan sensor information level event.	Inspect the system for anything obstructing the system fans. Ensure

Trap name	Description	Recommended action
		that the system fan tray is fully seated using the supplied screws.
bigipLibhalAomSensorPwrWarning (. 1.3.6.1.4.1.3375.2.4.0.185)	AOM has issued a power sensor warning level event.	Ensure that the power supply unit (PSU) is properly seated. Ensure that the PSU has an appropriate power feed.
bigipLibhalAomSensorPwrError (. 1.3.6.1.4.1.3375.2.4.0.186)	AOM has issued a power sensor error level event.	Ensure that the PSU is properly seated. Ensure that the PSU has an appropriate power feed.
bigipLibhalAomSensorPwrAlert (. 1.3.6.1.4.1.3375.2.4.0.187)	AOM has issued a power sensor alert level event.	Ensure that the power supply unit PSU is properly seated. Ensure that the PSU has an appropriate power feed.
bigipLibhalAomSensorPwrCritical (. 1.3.6.1.4.1.3375.2.4.0.188)	AOM has issued a power sensor critical level event.	Ensure that the power supply unit PSU is properly seated. Ensure that the PSU has an appropriate power feed.
bigipLibhalAomSensorPwrEmergency (. 1.3.6.1.4.1.3375.2.4.0.189)	AOM has issued a power sensor emergency level event.	Ensure that the power supply unit PSU is properly seated. Ensure that the PSU has an appropriate power feed.
bigipLibhalAomSensorPwrInfo (. 1.3.6.1.4.1.3375.2.4.0.190)	AOM has issued a power sensor information level event.	Ensure that the power supply unit PSU is properly seated. Ensure that the PSU has an appropriate power feed.

ASM-related traps and recommended actions

This table provides information about the ASM™-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipAsmRequestBlocked (. 1.3.6.1.4.1.3375.2.4.0.38)	The BIG-IP® system blocked an HTTP request because the request contained at least one violation to the active security policy.	Check the HTTP request to determine the cause of the violation.
bigipAsmRequestViolation (. 1.3.6.1.4.1.3375.2.4.0.39)	The BIG-IP system issued an alert because an HTTP request violated the active security policy.	Check the HTTP request to determine the cause of the violation.
bigipAsmFtpRequestBlocked (. 1.3.6.1.4.1.3375.2.4.0.79)	The BIG-IP system blocked an FTP request because the request contained at least one violation to the active security policy.	Check the FTP request to determine the cause of the violation.

Trap name	Description	Recommended action
bigipAsmFtpRequestViolation (. 1.3.6.1.4.1.3375.2.4.0.80)	The BIG-IP system issued an alert because an FTP request violated the active security policy.	Check the FTP request to determine the cause of the violation.
bigipAsmSmtpRequestBlocked (. 1.3.6.1.4.1.3375.2.4.0.85)	The BIG-IP system blocked an SMTP request because the request contained at least one violation to the active security policy.	Check the SMTP request to determine the cause of the violation.
bigipAsmSmtpRequestViolation (. 1.3.6.1.4.1.3375.2.4.0.86)	The BIG-IP system issued an alert because an SMTP request violated the active security policy.	Check the SMTP request to determine the cause of the violation.
bigipAsmDosAttackDetected (. 1.3.6.1.4.1.3375.2.4.0.91)	The BIG-IP system detected a denial-of-service (DoS) attack.	Determine the availability of the application by checking the response time of the site. Check the BIG-IP ASM logs: <ul style="list-style-type: none"> Identify the source IP of the attack and observe other violations from the same source. Determine if the source IP is attacking other resources. Consider blocking the source IP in the ACL. Identify the URL that is under attack. Consider disabling the URL, if the attack is not mitigated quickly.
bigipAsmBruteForceAttackDetected (. 1.3.6.1.4.1.3375.2.4.0.92)	The BIG-IP system detected a brute force attack.	Check the BIG-IP ASM logs: <ul style="list-style-type: none"> Identify the source IP of the attack and observe other violations from the same source. Determine if the source IP is attacking other resources. Consider blocking the source IP in the ACL. Identify the user name that is under attack. Consider contacting the user and locking their account.

Application Visibility and Reporting-related traps and recommended actions

This table provides information about the Application Visibility and Reporting (AVR) notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipAvrAlertsMetricSnmp (. 1.3.6.1.4.1.3375.2.4.0.105)	A BIG-IP system AVR SNMP metric changed.	Information only, no action required.
bigipAvrAlertsMetricSntp (. 1.3.6.1.4.1.3375.2.4.0.106)	A BIG-IP system AVR SMTP metric changed.	Information only, no action required.

Authentication-related traps and recommended actions

This table provides information about the authentication-related notifications that an SNMP manager can receive.

Trap Name	Description	Recommended Action
bigipTamdAlert (. 1.3.6.1.4.1.3375.2.4.0.21)	More than 60 authentication attempts have failed within one second, for a given virtual server.	Investigate for a possible intruder.
bigipAuthFailed (. 1.3.6.1.4.1.3375.2.4.0.27)	A login attempt failed.	Check the user name and password.

DDM-related traps and recommended actions

This table provides information about the Digital Diagnostic Monitoring (DDM)-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipDDMPowerAlarm (. 1.3.6.1.4.1.3375.2.4.0.158)	The Digital Diagnostic Monitoring (DDM) on a pluggable optical transceiver detected an alarm condition. DDM monitors both transmit and receive optical power to ensure the laser power is between vendor-specified power thresholds for pluggable optical modules such as SFP/SFP+/QSFP+/QSFP28. An alarm can occur when a cable is removed from a plugged port, or when the front panel port or the transceiver is not configured or operating properly.	<p>Refer to the text of the alert: is it a low or high alarm? Is it a transmit or receive alarm? The action to take for F5 branded optics (the following troubleshooting steps) depends on a condition derived from the two states (low/high and transmit/receive):</p> <ul style="list-style-type: none"> • Low (Alarm)/Transmit (Alarm): See if the BCM port is enabled. If not, then enable it. • High (Alarm)/Transmit (Alarm): Hot swap extract and insert F5 Optics multiple times. Check to see if a link comes up without a DDM error after each insertion. If a problem persists, then it is a bad F5 Optic. • Low (Alarm)/Receive (Alarm): Verify F5 optics module with local loopback cable. Verify that the transmission power on the other end of the cable is correct. Recheck the optical link budget calculations. Clean the optical cables, connectors, and/or lens. For any receive problem, look at

Trap name	Description	Recommended action
		<p>the transmitter to make sure it is okay and the correct protocol.</p> <ul style="list-style-type: none"> High (Alarm)/Receive (Alarm): Check the protocol setting on both link partners and make sure they are compatible. Verify that the transmission power on the other end is okay. Recheck the optical link budget calculations. For any receive problem, look at the transmitter to make sure it is okay and the correct protocol.
<p>bigipDDMPowerWarn (.1.3.6.1.4.1.3375.2.4.0.159)</p>	<p>The DDM on a pluggable optical transceiver detected a warning condition. DDM monitors both transmit and receive optical power to ensure the laser power is between vendor-specified power thresholds for pluggable optical modules such as SFP/SFP+/QSFP+/QSFP28. A warning can occur when a cable is removed from a plugged port, or when the front panel port or the transceiver is not configured or operating properly.</p>	<p>Refer to the text of the alert: is it a low or high alarm? Is it a transmit or receive alarm? The action to take for F5 branded optics (the following troubleshooting steps) depends on a condition derived from the two states (low/high and transmit/receive):</p> <ul style="list-style-type: none"> Low (Alarm)/Transmit (Alarm): See if the BCM port is enabled. If not, then enable it. High (Alarm)/Transmit (Alarm): Hot swap extract and insert F5 Optics multiple times. Check to see if a link comes up without a DDM error after each insertion. If a problem persists, then it is a bad F5 Optic. Low (Alarm)/Receive (Alarm): Verify F5 optics module with local loopback cable. Verify that the transmission power on the other end of the cable is correct. Recheck the optical link budget calculations. Clean the optical cables, connectors, and/or lens. For any receive problem, look at the transmitter to make sure it is okay and the correct protocol. High (Alarm)/Receive (Alarm): Check the protocol setting on both link partners and make sure they are compatible. Verify the transmission power on the other end is okay. Recheck the optical link budget calculations. For any receive problem, look at the transmitter to make sure it is okay and the correct protocol.

Trap name	Description	Recommended action
bigipDDMPowerAlarmClear (.1.3.6.1.4.1.3375.2.4.0.160)	The DDM on a pluggable optical transceiver no longer detects an alarm condition. DDM monitors both transmit and receive optical power to ensure the laser power is between vendor-specified power thresholds.	Depending on the state of the network, action might or might not be required. The previous alarm has cleared.
bigipDDMPowerWarnClear (.1.3.6.1.4.1.3375.2.4.0.161)	The DDM on a pluggable optical transceiver no longer detects a warning condition. DDM monitors both transmit and receive optical power to ensure the laser power is between vendor-specified power thresholds.	Depending on the state of the network, action might or might not be required. The previous alarm has cleared.
bigipDDMNonF5Optics (.1.3.6.1.4.1.3375.2.4.0.162)	A non-F5 pluggable optical transceiver is present in an interface. See K8153 at http://support.f5.com for restrictions on third-party hardware components with F5 products.	Might need to replace with an F5 branded optic.

DoS-related traps and recommended actions

This table provides information about the denial-of-service (DoS)-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipAggrReaperStateChange (.1.3.6.1.4.1.3375.2.4.0.22)	The state of the aggressive reaper has changed, indicating that the BIG-IP® system is moving to a distress mode.	Use the default denial-of-service (DoS) settings. You can also add rate filters to survive the attack.
bigipDosAttackStart (.1.3.6.1.4.1.3375.2.4.0.133)	The BIG-IP system detected a DoS attack start.	Check the attack name in the notification to determine the kind of attack that is detected.
bigipDosAttackStop (.1.3.6.1.4.1.3375.2.4.0.134)	The BIG-IP system detected a DoS attack stop.	Information only, no action required.

General traps and recommended actions

This table provides information about the general notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipDiskPartitionWarn (.1.3.6.1.4.1.3375.2.4.0.25)	Free space on the disk partition is less than the specified limit. By default, the limit is 30% of total disk space.	Increase the available disk space.

Trap name	Description	Recommended action
bigipDiskPartitionGrowth (.1.3.6.1.4.1.3375.2.4.0.26)	The disk partition use exceeds the specified growth limit. By default, the limit is 5% of total disk space.	Increase the available disk space.
bigipUpdatePriority (.1.3.6.1.4.1.3375.2.4.0.153)	There is a high priority software update available.	Download and install the software update.
bigipUpdateServer (.1.3.6.1.4.1.3375.2.4.0.154)	Unable to connect to the F5 server running update checks.	Verify the server connection settings.
bigipUpdateError (.1.3.6.1.4.1.3375.2.4.0.155)	There was an error checking for updates.	Investigate the error.
bigipAgentStart (.1.3.6.1.4.1.3375.2.4.0.1)	The SNMP agent on the BIG-IP® system has been started.	For your information only. No action required.
bigipAgentShutdown (.1.3.6.1.4.1.3375.2.4.0.2)	The SNMP agent on the BIG-IP system is in the process of being shut down.	For your information only. No action required.
bigipAgentRestart (.1.3.6.1.4.1.3375.2.4.0.3)	The SNMP agent on the BIG-IP system has been restarted.	This trap is for future use only.

BIG-IP DNS-related traps and recommended actions

This table provides information about the DNS-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipGtmBoxAvail (.1.3.6.1.4.1.3375.2.4.0.77)	The BIG-IP® system has come UP.	Information only, no action required.
bigipGtmBoxNotAvail (.1.3.6.1.4.1.3375.2.4.0.78)	The BIG-IP system has gone DOWN.	Information only, no action required.
bigipGtmBig3dSslCertExpired (.1.3.6.1.4.1.3375.2.4.0.81)	The certificate <code>/config/big3d/client.crt</code> has expired.	Replace the certificate.
bigipGtmBig3dSslCertWillExpire (.1.3.6.1.4.1.3375.2.4.0.82)	The certificate <code>/config/big3d/client.crt</code> will expire soon.	Replace the certificate.
bigipGtmSslCertExpired (.1.3.6.1.4.1.3375.2.4.0.83)	The certificate <code>/config/gtm/server.crt</code> has expired.	Replace the certificate.
bigipGtmSslCertWillExpire (.1.3.6.1.4.1.3375.2.4.0.84)	The certificate <code>/config/gtm/server.crt</code> will expire soon.	Replace the certificate.
bigipGtmPoolAvail (.1.3.6.1.4.1.3375.2.4.0.40)	A global traffic management pool is available.	Information only, no action required.
bigipGtmPoolNotAvail (.1.3.6.1.4.1.3375.2.4.0.41)	A global traffic management pool is not available.	Information only, no action required.

Trap name	Description	Recommended action
bigipGtmPoolDisabled (.1.3.6.1.4.1.3375.2.4.0.42)	A global traffic management pool is disabled.	Check the status of the pool.
bigipGtmPoolEnabled (.1.3.6.1.4.1.3375.2.4.0.43)	A global traffic management pool is enabled.	Information only, no action required.
bigipGtmLinkAvail (.1.3.6.1.4.1.3375.2.4.0.44)	A global traffic management link is available.	Information only, no action required.
bigipGtmLinkNotAvail (.1.3.6.1.4.1.3375.2.4.0.45)	A global traffic management link is not available.	Check the status of the link, as well as the relevant detailed log message.
bigipGtmLinkDisabled (.1.3.6.1.4.1.3375.2.4.0.46)	A global traffic management link is disabled.	Check the status of the link.
bigipGtmLinkEnabled (.1.3.6.1.4.1.3375.2.4.0.47)	A global traffic management link is enabled.	Information only, no action required.
bigipGtmWideIpAvail (.1.3.6.1.4.1.3375.2.4.0.48)	A global traffic management wide IP is available.	Information only, no action required.
bigipGtmWideIpNotAvail (.1.3.6.1.4.1.3375.2.4.0.49)	A global traffic management wide IP is unavailable.	Check the status of the wide IP, as well as the relevant detailed log message.
bigipGtmWideIpDisabled (.1.3.6.1.4.1.3375.2.4.0.50)	A global traffic management wide IP is disabled.	Check the status of the wide IP.
bigipGtmWideIpEnabled (.1.3.6.1.4.1.3375.2.4.0.51)	A global traffic management wide IP is enabled.	Information only, no action required.
bigipGtmPoolMbrAvail (.1.3.6.1.4.1.3375.2.4.0.52)	A global traffic management pool member is available.	Information only, no action required.
bigipGtmPoolMbrNotAvail (.1.3.6.1.4.1.3375.2.4.0.53)	A global traffic management pool member is not available.	Check the status of the pool member, as well as the relevant detailed log message.
bigipGtmPoolMbrDisabled (.1.3.6.1.4.1.3375.2.4.0.54)	A global traffic management pool member is disabled.	Check the status of the pool member.
bigipGtmPoolMbrEnabled (.1.3.6.1.4.1.3375.2.4.0.55)	A global traffic management pool member is enabled.	Information only, no action required.

Trap name	Description	Recommended action
bigipGtmServerAvail (.1.3.6.1.4.1.3375.2.4.0.56)	A global traffic management server is available.	Information only, no action required.
bigipGtmServerNotAvail (.1.3.6.1.4.1.3375.2.4.0.57)	A global traffic management server is unavailable.	Check the status of the server, as well as the relevant detailed log message.
bigipGtmServerDisabled (.1.3.6.1.4.1.3375.2.4.0.58)	A global traffic management server is disabled.	Check the status of the server.
bigipGtmServerEnabled (.1.3.6.1.4.1.3375.2.4.0.59)	A global traffic management server is enabled.	Information only, no action required.
bigipGtmVsAvail (.1.3.6.1.4.1.3375.2.4.0.60)	A global traffic management virtual server is available.	Information only, no action required.
bigipGtmVsNotAvail (.1.3.6.1.4.1.3375.2.4.0.61)	A global traffic management virtual server is unavailable.	Check the status of the virtual server, as well as the relevant detailed log message.
bigipGtmVsDisabled (.1.3.6.1.4.1.3375.2.4.0.62)	A global traffic management virtual server is disabled.	Check the status of the virtual server.
bigipGtmVsEnabled (.1.3.6.1.4.1.3375.2.4.0.63)	A global traffic management virtual server is enabled.	Information only, no action required.
bigipGtmDcAvail (.1.3.6.1.4.1.3375.2.4.0.64)	A global traffic management data center is available.	Information only, no action required.
bigipGtmDcNotAvail (.1.3.6.1.4.1.3375.2.4.0.65)	A global traffic management data center is unavailable.	Check the status of the data center, as well as the relevant detailed log message.
bigipGtmDcDisabled (.1.3.6.1.4.1.3375.2.4.0.66)	A global traffic management data center is disabled.	Check the status of the data center.
bigipGtmDcEnabled (.1.3.6.1.4.1.3375.2.4.0.67)	A global traffic management data center is enabled.	Information only, no action required.
bigipGtmAppObjAvail (.1.3.6.1.4.1.3375.2.4.0.69)	A global traffic management application object is available.	Information only, no action required.
bigipGtmAppObjNotAvail (.1.3.6.1.4.1.3375.2.4.0.70)	A global traffic management application object is unavailable.	Check the status of the application object, as well as the relevant detailed log message.

Trap name	Description	Recommended action
bigipGtmAppAvail (.1.3.6.1.4.1.3375.2.4.0.71)	A global traffic management application is available.	Information only, no action required.
bigipGtmAppNotAvail (.1.3.6.1.4.1.3375.2.4.0.72)	A global traffic management application is unavailable.	Check the status of the application, as well as the relevant detailed log message.
bigipGtmJoinedGroup (.1.3.6.1.4.1.3375.2.4.0.73)	The BIG-IP system joined a global traffic management synchronization group.	Information only, no action required.
bigipGtmLeftGroup (.1.3.6.1.4.1.3375.2.4.0.74)	The BIG-IP system left a global traffic management synchronization group.	Information only, no action required.
bigipGtmKeyGenerationExpiration (.1.3.6.1.4.1.3375.2.4.0.95)	A generation of a DNSSEC key expired.	Information only, no action required.
bigipGtmKeyGenerationRollover (.1.3.6.1.4.1.3375.2.4.0.94)	A generation of a DNSSEC key rolled over.	Information only, no action required.
bigipGtmProberPoolDisabled (.1.3.6.1.4.1.3375.2.4.0.99)	A global traffic management prober pool is disabled.	Check the status of the prober pool.
bigipGtmProberPoolEnabled (.1.3.6.1.4.1.3375.2.4.0.100)	A global traffic management prober pool is enabled.	Information only, no action required.
bigipGtmProberPoolStatusChange (.1.3.6.1.4.1.3375.2.4.0.97)	The status of a global traffic management prober pool has changed.	Check the status of the prober pool.
bigipGtmProberPoolStatusChangeReason (.1.3.6.1.4.1.3375.2.4.0.98)	The reason the status of a global traffic management prober pool has changed.	The action required is based on the reason given.
bigipGtmProberPoolMbrDisabled (.1.3.6.1.4.1.3375.2.4.0.103)	A global traffic management prober pool member is disabled.	Check the status of the prober pool member.
bigipGtmProberPoolMbrEnabled (.1.3.6.1.4.1.3375.2.4.0.104)	A global traffic management prober pool member is enabled.	Information only, no action required.
bigipGtmProberPoolMbrStatusChange (.1.3.6.1.4.1.3375.2.4.0.101)	The status of a global traffic management prober pool member has changed.	Check the status of the prober pool member.
bigipGtmProberPoolMbrStatusChangeReason (.1.3.6.1.4.1.3375.2.4.0.102)	The reason the status of a global traffic management prober pool member has changed.	The action required is based on the reason given.

Hardware-related traps and recommended actions

This table provides information about hardware-related notifications that an SNMP manager can receive. If you receive any of these alerts, contact F5® Networks technical support.

Trap name and Associated OID	Description	Recommended action
bigipAomCpuTempTooHigh (. 1.3.6.1.4.1.3375.2.4.0.93)	The AOM is reporting that the air temperature near the CPU is too high.	Check the input and output air temperatures. Run an iHealth® report and troubleshoot based on the results. If the condition persists, contact F5 Networks technical support.
bigipBladeNoPower (. 1.3.6.1.4.1.3375.2.4.0.88)	A blade lost power.	Contact F5 Networks technical support.
bigipBladeTempHigh (. 1.3.6.1.4.1.3375.2.4.0.87)	The temperature of a blade is too high.	This trap might be spurious. If the condition persists, contact F5 Networks technical support.
bigipBladeOffline (. 1.3.6.1.4.1.3375.2.4.0.90)	A blade has failed.	Remove the blade. Contact F5 Networks technical support.
bigipChmandAlertFanTrayBad (. 1.3.6.1.4.1.3375.2.4.0.121)	A fan tray in a chassis is bad or was removed.	Replace the fan tray. If the condition persists, contact F5 Networks technical support.
bigipCpuTempHigh	The CPU temperature is too high.	Check the input and output air temperatures. Run an iHealth report and troubleshoot based on the results. If the condition persists, contact F5 Networks technical support.
bigipCpuFanSpeedLow (. 1.3.6.1.4.1.3375.2.4.0.5)	The CPU fan speed is too low.	Check the CPU temperature. If the CPU temperature is normal, the condition is not critical. If the condition persists, contact F5 Networks technical support.
bigipCpuFanSpeedBad (. 1.3.6.1.4.1.3375.2.4.0.6)	The CPU fan is not receiving a signal.	Check the CPU temperature. If the CPU temperature is normal, the condition is not critical. If the condition persists, contact F5 Networks technical support.
bigipSystemCheckAlertFanSpeedLow (. 1.3.6.1.4.1.3375.2.4.0.115)	The system fan speed is too low.	This condition is critical. Replace the fan tray. These appliances do not have fan trays: 1600, 3600, 3900, EM4000, 2000, 4000. If the condition persists, contact F5 Networks technical support.
bigipSystemCheckAlertVoltageHigh (. 1.3.6.1.4.1.3375.2.4.0.114)	The system voltage is too high.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support.

Trap name and Associated OID	Description	Recommended action
		<i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertVoltageLow (. 1.3.6.1.4.1.3375.2.4.0.123)	The system voltage is too low.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertMilliVoltageHigh (. 1.3.6.1.4.1.3375.2.4.0.124)	The system milli-voltage is too high.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertMilliVoltageLow (. 1.3.6.1.4.1.3375.2.4.0.127)	The system milli-voltage is too low.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertTempHigh (. 1.3.6.1.4.1.3375.2.4.0.113)	The system temperature is too high.	Check the system and air temperatures. If the condition persists, contact F5 Networks technical support.
bigipSystemCheckAlertCurrentHigh (. 1.3.6.1.4.1.3375.2.4.0.125)	The system current is too high.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertCurrentLow (. 1.3.6.1.4.1.3375.2.4.0.128)	The system current is too low.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertPowerHigh (. 1.3.6.1.4.1.3375.2.4.0.126)	The system power is too high.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support.

Trap name and Associated OID	Description	Recommended action
		<i>Note: This alert does not happen for standby power.</i>
bigipSystemCheckAlertPowerLow (. 1.3.6.1.4.1.3375.2.4.0.129)	The system power is too low.	Review additional error messages in the log files. Unplug the system. Contact F5 Networks technical support. <i>Note: This alert does not happen for standby power.</i>
bigipChassisTempHigh (. 1.3.6.1.4.1.3375.2.4.0.7)	The temperature of the chassis is too high.	Contact F5 Networks technical support.
bigipChassisFanBad (. 1.3.6.1.4.1.3375.2.4.0.8)	The chassis fan is not operating properly.	Replace the fan tray. If the condition persists, contact F5 Networks technical support.
bigipChassisPowerSupplyBad (. 1.3.6.1.4.1.3375.2.4.0.9)	The chassis power supply is not functioning properly.	Verify that the power supply is plugged in. In the case of a dual-power-supply system, verify that both power supplies are plugged in. Contact F5 Networks technical support.
bigipLibhalBladePoweredOff (. 1.3.6.1.4.1.3375.2.4.0.119)	A blade is powered off.	Contact F5 Networks technical support.
bigipLibhalSensorAlarmCritical (. 1.3.6.1.4.1.3375.2.4.0.120)	The hardware sensor on a blade indicates a critical alarm.	Review any additional error messages that you receive, and troubleshoot accordingly. If the condition persists, contact F5 Networks technical support.
bigipLibhalDiskBayRemoved (. 1.3.6.1.4.1.3375.2.4.0.118)	A disk sled was removed from a bay.	Information only, no action required.
bigipLibhalSsdLogicalDiskRemoved (. 1.3.6.1.4.1.3375.2.4.0.117)	An SSD logical disk was removed from the BIG-IP® system.	Information only, no action required.
bigipLibhalSsdPhysicalDiskRemoved (. 1.3.6.1.4.1.3375.2.4.0.116)	An SSD physical disk was removed from the BIG-IP system.	Information only, no action required.
bigipRaidDiskFailure (. 1.3.6.1.4.1.3375.2.4.0.96)	An disk in a RAID disk array failed.	On www.askf5.com , see <i>SOL10856: Overview of hard drive mirroring</i> . If the problem persists, contact F5 Networks technical support.
bigipSsdMwiNearThreshold (. 1.3.6.1.4.1.3375.2.4.0.111)	An SSD disk is reaching a known wear threshold.	Contact F5 Networks technical support.
bigipSsdMwiReachedThreshold (. 1.3.6.1.4.1.3375.2.4.0.112)	An SSD disk is worn out.	If this is the first alert, the disk might continue to operate for a

Trap name and Associated OID	Description	Recommended action
		short time. Contact F5 Networks technical support.
bigipNetLinkDown (. 1.3.6.1.4.1.3375.2.4.0.24)	An interface link is down.	This alert applies to L1 and L2, which are internal links within the device connecting the CPU and Switch subsystems. These links should never be down. If this occurs, the condition is serious. Contact F5 Networks technical support.
bigipExternalLinkChange (. 1.3.6.1.4.1.3375.2.4.0.37)	The status of an external interface link has changed to either UP, DOWN, or UNPOPULATED.	This occurs when network cables are added or removed, and the network is reconfigured. Determine whether the link should be down or up, and then take the appropriate action.
bigipPsPowerOn (. 1.3.6.1.4.1.3375.2.4.0.147)	The power supply for the BIG-IP system was powered on.	Information only, no action required, unless this trap is unexpected. In that case, verify that the power supply is working and that system has not rebooted.
bigipPsPowerOff (. 1.3.6.1.4.1.3375.2.4.0.148)	The power supply for the BIG-IP system was powered off.	Information only, no action required, unless power off was unexpected. In that case, verify that the power supply is working and that system has not rebooted.
bigipPsAbsent (.1.3.6.1.4.1.3375.2.4.0.149)	The power supply for the BIG-IP system cannot be detected.	Information only, no action required when the BIG-IP device is operating with one power supply. For BIG-IP devices with two power supplies installed, verify that both power supplies are functioning correctly and evaluate symptoms.
bigipSystemShutdown (. 1.3.6.1.4.1.3375.2.4.0.151)	The BIG-IP system has shut down.	Information only, no action required when the shut down was expected. Otherwise, investigate the cause of the unexpected reboot.
bigipFipsDeviceError (. 1.3.6.1.4.1.3375.2.4.0.152)	The FIPS card in the BIG-IP system has encountered a problem.	Contact F5 Networks technical support.

High-availability system-related traps and recommended actions

This table provides information about the high-availability system-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipStandby (. 1.3.6.1.4.1.3375.2.4.0.14)	The BIG-IP® system has switched to standby mode.	Review the log files in the <code>/var/log</code> directory and then search for core files in the <code>/var/core</code> directory. If you find a core file, or find text similar to <code>fault at location xxxx stack trace:</code> , contact F5® Networks technical support.
bigipStandByFail (. 1.3.6.1.4.1.3375.2.4.0.75)	In failover condition, this standby system cannot become active.	Investigate failover condition on the standby system.
bigipActive (. 1.3.6.1.4.1.3375.2.4.0.15)	The BIG-IP system has switched to active mode.	Information only, no action required.
bigipActiveActive (. 1.3.6.1.4.1.3375.2.4.0.16)	The BIG-IP system is in active-active mode.	Information only, no action required.
bigipFeatureFailed (. 1.3.6.1.4.1.3375.2.4.0.17)	A high-availability feature has failed.	View high-availability processes and their current status.
bigipFeatureOnline (. 1.3.6.1.4.1.3375.2.4.0.18)	A high-availability feature is responding.	View high-availability processes and their current status.
bigipTrafficGroupStandby (. 1.3.6.1.4.1.3375.2.4.0.141)	The status of a traffic group has changed to stand by.	Information only, no action required. To determine the reason for the failover, review the LTM® log <code>/var/log/ltm</code> and search for keywords <code>active</code> or <code>standby</code> . Additionally, you can run the <code>tmsh</code> command <code>tmsh show sys ha-status</code> to view the failover conditions.
bigipTrafficGroupActive (. 1.3.6.1.4.1.3375.2.4.0.142)	The status of a traffic group has changed to active.	Information only, no action required. To determine the reason for the failover, review the LTM log <code>/var/log/ltm</code> and search for keywords <code>active</code> or <code>standby</code> . Additionally, you can run the <code>tmsh</code> command <code>tmsh show sys ha-status</code> to view the failover conditions.
bigipTrafficGroupOffline (. 1.3.6.1.4.1.3375.2.4.0.143)	The status of a traffic group has changed to offline.	Information only, no action required.
bigipTrafficGroupForcedOffline (. 1.3.6.1.4.1.3375.2.4.0.144)	The status of a traffic group has changed to forced offline.	Information only, no action required.
bigipTrafficGroupDeactivate (. 1.3.6.1.4.1.3375.2.4.0.145)	A traffic group was deactivated.	Information only, no action required. To determine the reason for the deactivation, review the LTM

Trap name	Description	Recommended action
		log /var/log/ltm and search for the keyword deactivate.
bigipTrafficGroupActivate (. 1.3.6.1.4.1.3375.2.4.0.146)	A traffic group was activated.	Information only, no action required. To determine the reason for the deactivation, review the LTM log /var/log/ltm and search for the keyword activate.

License-related traps and recommended actions

This table provides information about the license-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipLicenseFailed (. 1.3.6.1.4.1.3375.2.4.0.19)	Validation of a BIG-IP® system license has failed, or the dossier has errors.	Occurs only when first licensing the system or adding a module key (such as HTTP compression) to an existing system. If using automatic licensing, verify connectivity to the outside world, fix the dossier if needed, and try again.
bigipLicenseExpired (. 1.3.6.1.4.1.3375.2.4.0.20)	The BIG-IP license has expired.	Call F5® Networks technical support.
bigipDnsRequestRateLimiterEngaged (. 1.3.6.1.4.1.3375.2.4.0.139)	The BIG-IP DNS Services license is rate-limited and the system has reached the rate limit.	Call F5 Networks technical support to upgrade your license.
bigipGtmRequestRateLimiterEngaged (. 1.3.6.1.4.1.3375.2.4.0.140)	The BIG-IP DNS license is rate-limited and the system has reached the rate limit.	Call F5 Networks technical support to upgrade your license.
bigipCompLimitExceeded (. 1.3.6.1.4.1.3375.2.4.0.35)	The compression license limit is exceeded.	Purchase additional compression licensing from F5 Networks.
bigipSslLimitExceeded (. 1.3.6.1.4.1.3375.2.4.0.36)	The SSL license limit is exceeded, either for transactions per second (TPS) or for megabits per second (MPS).	Purchase additional SSL licensing from F5 Networks.

LTM-related traps and recommended actions

This table provides information about the LTM®-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipUnsolicitedRepliesExceededThreshold (. 1.3.6.1.4.1.3375.2.4.0.122)	The BIG-IP® system DNS cache received unsolicited	Check the BIG-IP system logs to determine if the system

Trap name	Description	Recommended action
	query replies exceeding the configured threshold.	is experiencing a distributed denial-of-service (DDoS) attack.
bigipNodeRate (.1.3.6.1.4.1.3375.2.4.0.130)	A local traffic management node has received connections exceeding the configured rate-limit.	Consider provisioning more resources on the BIG-IP system for this virtual server.
bigipNodeDown (.1.3.6.1.4.1.3375.2.4.0.12)	A BIG-IP system health monitor has marked a node as down.	Check the node and the cable connection.
bigipNodeUp (.1.3.6.1.4.1.3375.2.4.0.13)	A BIG-IP system health monitor has marked a node as up.	Information, no action required.
bigipMemberRate (.1.3.6.1.4.1.3375.2.4.0.131)	A local traffic management pool member has received connections exceeding the configured rate-limit.	Consider provisioning more resources on the BIG-IP system for this virtual server.
bigipVirtualRate (.1.3.6.1.4.1.3375.2.4.0.132)	A local traffic management virtual server has received connections exceeding the configured rate-limit.	Consider provisioning more resources on the BIG-IP system for this virtual server.
bigipLtmVsAvail (.1.3.6.1.4.1.3375.2.4.0.135)	A local traffic management virtual server is available to receive connections.	Information only, no action required.
bigipLtmVsUnavail (.1.3.6.1.4.1.3375.2.4.0.136)	A local traffic management virtual server is not available to receive connections.	Check the virtual server.
bigipLtmVsEnabled (.1.3.6.1.4.1.3375.2.4.0.137)	A local traffic management virtual server is enabled.	Information only, no action required.
bigipLtmVsDisabled (.1.3.6.1.4.1.3375.2.4.0.138)	A local traffic management virtual server is disabled.	Information only, no action required.
bigipServiceDown (.1.3.6.1.4.1.3375.2.4.0.10)	A BIG-IP system health monitor has detected a service on a node to be stopped and thus marked the node as down.	Restart the service on the node.
bigipServiceUp (.1.3.6.1.4.1.3375.2.4.0.11)	A BIG-IP system health monitor has detected a service on a node to be running and has therefore marked the node as up.	Information only, no action required.
bigipPacketRejected (.1.3.6.1.4.1.3375.2.4.0.34)	The BIG-IP system has rejected some packets.	Check the detailed message within this trap and act accordingly.
bigipInetPortExhaustion (.1.3.6.1.4.1.3375.2.4.0.76)	The TMM has run out of source ports and cannot	Either increase the number of addresses

Trap name	Description	Recommended action
	open new communications channels with other machines.	available for SNAT automapping or SNAT pools, or lower the idle timeout value if the value is excessively high.

Logging-related traps and recommended actions

This table provides information about the logging-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipLogEmerg (.1.3.6.1.4.1.3375.2.4.0.29)	The BIG-IP® system is unusable. This notification occurs when the system logs a message with the log level LOG_EMERG.	Check the detailed message within this trap and within the <code>/var/log</code> files to determine which process has the emergency. Then act accordingly.
bigipLogAlert (.1.3.6.1.4.1.3375.2.4.0.30)	The BIG-IP system requires immediate action to function properly. This notification occurs when the system logs a message with the log level LOG_ALERT.	Check the detailed message within this trap and within the <code>/var/log</code> files to determine which process has the alert situation. Then act accordingly.
bigipLogCrit (.1.3.6.1.4.1.3375.2.4.0.31)	The BIG-IP system is in critical condition. This notification occurs when the system logs a message with the log level LOG_CRIT.	Check the detailed message within this trap and within the <code>/var/log</code> files to determine which process has the critical situation. Then act accordingly.
bigipLogErr (.1.3.6.1.4.1.3375.2.4.0.32)	The BIG-IP system has some error conditions. This notification occurs when the system logs a message with the log level LOG_ERR.	Check the detailed message within this trap and within the <code>/var/log</code> files to determine which processes have the error conditions. Then act accordingly.
bigipLogWarning (.1.3.6.1.4.1.3375.2.4.0.33)	The BIG-IP system is experiencing some warning conditions. This notification occurs when the system logs a message with the log level LOG_WARNING.	Check the detailed message within this trap and within the <code>/var/log</code> files to determine which processes have the warning conditions. Then act accordingly.

Network-related traps and recommended actions

This table provides information about the network-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipARPConflict (.1.3.6.1.4.1.3375.2.4.0.23)	The BIG-IP® system has detected an ARP advertisement for any of its own ARP-enabled addresses. This can occur for a virtual server address or a self IP address.	Check IP addresses and routes.

vCMP-related traps and recommended actions

This table provides information about the virtual clustered multiprocessing (vCMP®)-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipVcmpAlertsVcmpPowerOn (. 1.3.6.1.4.1.3375.2.4.0.107)	The BIG-IP® system powered on a vCMP guest from a suspended or powered-off state.	Information only, no action required.
bigipVcmpAlertsVcmpPowerOff (. 1.3.6.1.4.1.3375.2.4.0.108)	The BIG-IP system powered off a vCMP guest.	Information only, no action required.
bigipVcmpAlertsVcmpHBLost (. 1.3.6.1.4.1.3375.2.4.0.109)	The BIG-IP system cannot detect a heartbeat from a vCMP guest.	Check the guest and restart, if necessary.
bigipVcmpAlertsVcmpHBDetected (. 1.3.6.1.4.1.3375.2.4.0.110)	The BIG-IP system detected a heartbeat from a new or returning vCMP guest.	Information only, no action required.

VIPRION-related traps and recommended actions

This table provides information about the VIPRION®-related notifications that an SNMP manager can receive.

Trap name	Description	Recommended action
bigipClusterdNoResponse (. 1.3.6.1.4.1.3375.2.4.0.89)	The cluster daemon failed to respond for 10 seconds or more.	Start the cluster daemon.
bigipClusterPrimaryChanged (. 1.3.6.1.4.1.3375.2.4.0.150)	The primary cluster has changed.	Information only, no action required.

About enterprise MIB files

The enterprise MIB files contain F5® Networks specific information. All OIDs for the BIG-IP® system data are contained in the F5 enterprise MIB files, including all interface statistics (1.3.6.1.4.1.3375.2.1.2.4 (sysNetwork.sysInterfaces)). These enterprise MIB files reside on the BIG-IP system:

F5-BIGIP-COMMON-MIB.txt

Contains information that the SNMP manager can use to help manage F5-specific notifications (SNMP traps) that all other BIG-IP MIB files reference.

F5-BIGIP-SYSTEM-MIB.txt

Contains information that the SNMP manager can use to help manage BIG-IP system objects, such as global statistic data, network information, and platform information.

F5-BIGIP-LOCAL-MIB.txt

Contains information that the SNMP manager can use to help manage BIG-IP local traffic objects, such as virtual servers, pools, nodes, profiles, health monitors, iRules®, and SNATs. Also contains information on AFM™ objects, such as firewall rules and DoS vectors.

F5-BIGIP-GLOBAL-MIB.txt

Contains information that the SNMP manager can use to help manage global traffic objects, such as wide IPs, virtual servers, pools, links, servers, and data centers.

F5-BIGIP-APM-MIB.txt

Contains information that the SNMP manager can use to help manage access policy objects, such as profiles, statistics, lease pools, and ACLs.

F5-BIGIP-WAM-MIB.txt

Contains information that the SNMP manager can use to help manage traffic acceleration objects, such as applications, profiles, and statistics.

Task summary

Perform these tasks when working with MIB files.

Downloading enterprise and NET-SNMP MIBs to the SNMP manager

Viewing objects in enterprise MIB files

Viewing SNMP traps in F5-BIGIP-LOCAL-MIB.txt

Collecting network firewall data using SNMP

Collecting DoS attack data using SNMP

Downloading enterprise and NET-SNMP MIBs to the SNMP manager

View the set of standard SNMP MIB files that you can download to the SNMP manager, by listing the contents of the BIG-IP® system directory `/usr/share/snmp/mibs`.

Download compressed files that contain the enterprise and NET-SNMP MIBs.

1. Click the **About** tab.
2. Click **Downloads**.
3. Click **Download F5 MIBs (mibs_f5.tar.gz)** or **Download NET-SNMP MIBs (mibs_netsnmp.tar.gz)**.
4. Follow the instructions on the screen to complete the download.

Viewing objects in enterprise MIB files

You must have the `Administrator` user role assigned to your user account.

View information about a BIG-IP system object by listing the contents of an enterprise MIB file.

1. Access a console window on the BIG-IP system.
2. At the command prompt, list the contents of the directory `/usr/share/snmp/mibs`.
3. View available objects in the relevant MIB file.

Viewing SNMP traps in F5-BIGIP-LOCAL-MIB.txt

Verify that you have the `Administrator` user role assigned to your user account.

When an F5-specific trap sends a notification to the SNMP manager, the SNMP manager receives a text message describing the event or problem that has occurred. You can identify the traps specified in the `F5-BIGIP-LOCAL-MIB.txt` file by viewing the file.

1. Access a console window on the BIG-IP system.
2. At the command prompt, list the contents of the directory `/usr/share/snmp/mibs`.
3. View the `F5-BIGIP-LOCAL-MIB.txt` file.

Look for objects with the prefix `ltmFw` for firewall rules, and `ltmDos` for DoS attacks.

Collecting network firewall data using SNMP

You can use SNMP commands to collect firewall rule data.

Write an SNMP command to gather data on firewall rules, contexts, and rule hits.

For example, this SNMP command collects data on firewall rules memory usage, where `public` is the community name and you are logged in to the BIG-IP system: `snmpwalk -c public localhost ltmFwRuleStat`

The SNMP manager now queries the system about firewall rules.

Collecting DoS attack data using SNMP

You can use SNMP commands to gather DoS attack data.

Write an SNMP command to gather DoS attack data from the BIG-IP system.

For example, this SNMP command collects DoS attack data, where `public` is the community name and you are logged in locally to the BIG-IP® system: `snmpwalk -c public localhost ltmDosAttackDataStat`

The SNMP manager displays a list of all the DoS attack types and hits on those attack types.

About enterprise MIB files

The enterprise MIB files contain F5® Networks specific information. All OIDS for the BIG-IP® system data are contained in the F5 enterprise MIB files, including all interface statistics (1.3.6.1.4.1.3375.2.1.2.4 (`sysNetwork.sysInterfaces`)). These enterprise MIB files reside on the BIG-IP system:

F5-BIGIP-COMMON-MIB.txt

Contains information that the SNMP manager can use to help manage F5-specific notifications (SNMP traps) that all other BIG-IP MIB files reference.

F5-BIGIP-SYSTEM-MIB.txt

Contains information that the SNMP manager can use to help manage BIG-IP system objects, such as global statistic data, network information, and platform information.

F5-BIGIP-LOCAL-MIB.txt

Contains information that the SNMP manager can use to help manage BIG-IP local traffic objects, such as virtual servers, pools, nodes, profiles, health monitors, iRules®, and SNATs. Also contains information on AFM™ objects, such as firewall rules and DoS vectors.

F5-BIGIP-GLOBAL-MIB.txt

Contains information that the SNMP manager can use to help manage global traffic objects, such as wide IPs, virtual servers, pools, links, servers, and data centers.

F5-BIGIP-APM-MIB.txt

Contains information that the SNMP manager can use to help manage access policy objects, such as profiles, statistics, lease pools, and ACLs.

F5-BIGIP-WAM-MIB.txt

Contains information that the SNMP manager can use to help manage traffic acceleration objects, such as applications, profiles, and statistics.

Task summary

Perform these tasks when working with MIB files.

Downloading enterprise and NET-SNMP MIBs to the SNMP manager

Viewing objects in enterprise MIB files

Viewing SNMP traps in F5-BIGIP-LOCAL-MIB.txt

Collecting network firewall data using SNMP

Collecting DoS attack data using SNMP

Downloading enterprise and NET-SNMP MIBs to the SNMP manager

View the set of standard SNMP MIB files that you can download to the SNMP manager, by listing the contents of the BIG-IP® system directory `/usr/share/snmp/mibs`.

Download compressed files that contain the enterprise and NET-SNMP MIBs.

1. Click the **About** tab.
2. Click **Downloads**.
3. Click **Download F5 MIBs (mibs_f5.tar.gz)** or **Download NET-SNMP MIBs (mibs_netsnmp.tar.gz)**.
4. Follow the instructions on the screen to complete the download.

Viewing objects in enterprise MIB files

You must have the `Administrator` user role assigned to your user account.

View information about a BIG-IP system object by listing the contents of an enterprise MIB file.

1. Access a console window on the BIG-IP system.
2. At the command prompt, list the contents of the directory `/usr/share/snmp/mibs`.
3. View available objects in the relevant MIB file.

Viewing SNMP traps in F5-BIGIP-COMMON-MIB.txt

Verify that you have the `Administrator` user role assigned to your user account.

When an F5-specific trap sends a notification to the SNMP manager, the SNMP manager receives a text message describing the event or problem that has occurred. You can identify the traps specified in the F5-BIGIP-COMMON-MIB.txt file by viewing the file.

1. Access a console window on the BIG-IP system.
2. At the command prompt, list the contents of the directory `/usr/share/snmp/mibs`.
3. View the F5-BIGIP-COMMON-MIB.txt file. Look for object names with the designation NOTIFICATION-TYPE.

Viewing dynamic routing SNMP traps and associated OIDs

Verify that you have the `Administrator` user role assigned to your user account.

When you want to set up your network management systems to watch for problems with dynamic routing, you can view SNMP MIB files to discover the SNMP traps that the dynamic routing protocols send, and to find the OIDs that are associated with those traps.

1. Access a console window on the BIG-IP system.
2. At the command prompt, list the contents of the directory `/usr/share/snmp/mibs`.

3. View the following dynamic routing MIB files:

- BGP4-MIB.txt
- ISIS-MIB.txt
- OSPF6-MIB.txt
- OSPF-MIB.txt
- OSPF-TRAP-MIB.txt
- RIPv2-MIB.txt

Monitoring BIG-IP system processes using SNMP

Ensure that your SNMP manager is running either SNMP v2c or SNMP v3, because all BIG-IP® system statistics are defined by 64-bit counters, and only SNMP v2c and SNMP v3 support 64-bit counters. Ensure that you have downloaded the F-5 Networks enterprise and NET-SNMP MIBs to the SNMP manager.

You can monitor a specific process on the BIG-IP system using SNMP. To do this you can use the `HOST-RESOURCES` MIB and write a script to monitor the process.

Write a script to monitor a BIG-IP system process using the `HOST-RESOURCES` MIB.

For example, this command determines the number of TMM processes currently running on the system: `snmpwalk -v2c -c public localhost hrSWRunName | egrep "\"tmm([0-9]+)?\" | wc -l`

The script can now query the BIG-IP system about the status of processes.

Collecting BIG-IP system memory usage data using SNMP

You can use an SNMP command with OIDs to gather data on the number of bytes of memory currently being used on the BIG-IP® system.

Note: To interpret data on memory use, you do not need to perform a calculation on the collected data.

Write an SNMP command to gather data on the number of bytes of memory currently being used on the BIG-IP system.

For example, this SNMP command collects data on current memory usage, where `public` is the community name and `bigip` is the host name of the BIG-IP system: `snmpget -c public bigip sysGlobalStat.sysStatMemoryUsed.0`

The SNMP manager can now query the BIG-IP system about CPU and memory usage.

Collecting BIG-IP system data on HTTP requests using SNMP

You can use SNMP commands with an OID to gather and interpret data on the number of current HTTP requests on the BIG-IP® system. The following table shows the required OIDs for polling data on HTTP requests.

Performance Graph	Graph Metrics	Required SNMP OIDs
HTTP Requests	HTTP Requests	sysStatHttpRequests (.1.3.6.1.4.1.3375.2.1.1.2.1.56)

The following table shows the required calculations for interpreting metrics on HTTP requests.

Performance Graph	Graph Metric	Required calculations for HTTP requests
HTTP Requests	HTTP Requests	$\langle \text{DeltaStatHttpRequests} \rangle / \langle \text{interval} \rangle$

1. For each OID, perform two separate polls, at an interval of your choice. For example, poll OID `sysStatHttpRequests (.1.3.6.1.4.1.3375.2.1.1.2.1.56)` twice, at a 10-second interval. This results in two values, `<sysStatHttpRequests1>` and `<sysStatHttpRequests2>`.
2. Calculate the delta of the two poll values. For example:

```
 $\langle \text{DeltaStatHttpRequests} \rangle = \langle \text{sysStatHttpRequests2} \rangle - \langle \text{sysStatHttpRequests1} \rangle$ 
```

3. Perform the calculation on the OID deltas. The value for `interval` is 10. For example, to calculate the value of the HTTP Requests graph metric:

```
 $(\langle \text{DeltaStatHttpRequests} \rangle) / \langle \text{interval} \rangle$ 
```

Collecting BIG-IP system data on throughput rates using SNMP

You can use SNMP commands with various OIDs to gather and interpret data on the throughput rate on the BIG-IP® system. The following table shows the individual OIDs that you must poll, retrieving two separate poll values for each OID.

Performance Graph	Graph Metrics	Required SNMP OIDs
Throughput (summary graph)	Client Bits	<code>sysStatClientBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.3)</code>
	Client Bits	<code>sysStatClientBytesOut (.1.3.6.1.4.1.3375.2.1.1.2.1.5)</code>
	Server Bits	<code>sysStatServerBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.10)</code>
	Server Bits	<code>sysStatServerBytesOut (.1.3.6.1.4.1.3375.2.1.1.2.1.12)</code>
Client-side Throughput (detailed graph)	Client Bits In	<code>sysStatClientBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.3)</code>
	Client Bits Out	<code>sysStatClientBytesOut (.1.3.6.1.4.1.3375.2.1.1.2.1.5)</code>
Server-side Throughput (detailed graph)	Server Bits In	<code>sysStatServerBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.10)</code>
	Server Bits Out	<code>sysStatServerBytesOut (.1.3.6.1.4.1.3375.2.1.1.2.1.12)</code>
HTTP Compression Rate (detailed graph)	Compression	<code>sysHttpCompressionStatPrecompressBytes (.1.3.6.1.4.1.3375.2.1.1.2.22.2)</code>

The following table shows the required calculations for interpreting metrics on throughput rates.

Performance Graph	Graph Metrics	Required calculations for throughput rates
Throughput (summary graph)	Client Bits	$((\langle \text{DeltaStatClientBytesIn} \rangle + \langle \text{DeltaStatClientBytesOut} \rangle) * 8) / \langle \text{interval} \rangle$
	Server Bits	$((\langle \text{DeltaStatServerBytesIn} \rangle + \langle \text{DeltaStatServerBytesOut} \rangle) * 8) / \langle \text{interval} \rangle$
	Compression	$(\langle \text{DeltaHttpStatPrecompressBytes} \rangle) * 8 / \langle \text{interval} \rangle$

Performance Graph	Graph Metrics	Required calculations for throughput rates
Throughput (detailed graph)	Client Bits In	$(\langle \Delta \text{StatClientBytesIn} \rangle * 8) / \langle \text{interval} \rangle$
	Client Bits Out	$(\langle \Delta \text{StatClientBytesOut} \rangle * 8) / \langle \text{interval} \rangle$
	Server Bits In	$(\langle \Delta \text{StatServerBytesIn} \rangle * 8) / \langle \text{interval} \rangle$
	Server Bits Out	$(\langle \Delta \text{StatServerBytesOut} \rangle * 8) / \langle \text{interval} \rangle$
	Compression	$(\langle \Delta \text{HttpStatPrecompressBytes} \rangle * 8) / \langle \text{interval} \rangle$

- For each OID, perform two separate polls, at an interval of your choice. For example, poll OID `sysStatServerBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.10)` twice, at a 10-second interval. This results in two values, `<sysStatServerBytesIn1>` and `<sysStatServerBytesIn2>`.
- Calculate the delta of the two poll values. For example, for the Server Bits In graphic metric, perform this calculation:

$$\langle \Delta \text{StatServerBytesIn} \rangle = \langle \text{sysStatServerBytesIn2} \rangle - \langle \text{sysStatServerBytesIn1} \rangle$$

- Perform the calculation on the OID deltas. For this calculation, it is the average per second in the last `<interval>`. The value for `interval` is 10. For example, to calculate the value of the Server Bits In graph metric:

$$(\langle \Delta \text{StatServerBytesIn} \rangle) / \langle \text{interval} \rangle$$

Collecting BIG-IP system data on RAM cache using SNMP

You can use an SNMP command with various OIDs to gather and interpret data on RAM cache use. The following table shows the required OIDs for polling for data on RAM Cache use.

Performance Graph	Graph Metric	Required SNMP OIDs
RAM Cache Utilization	Hit Rate	<code>sysWebAccelerationStatCacheHits (.1.3.6.1.4.1.3375.2.1.1.2.23.2)</code> <code>sysWebAccelerationStatCacheMisses (.1.3.6.1.4.1.3375.2.1.1.2.23.3)</code>
CPU Cache Utilization	Byte Rate	<code>sysWebAccelerationStatCacheHitBytes (.1.3.6.1.4.1.3375.2.1.1.2.23.5)</code> <code>sysWebAccelerationStatCacheMissBytes (.1.3.6.1.4.1.3375.2.1.1.2.23.6)</code>
RAM Cache Utilization	Eviction Rate	<code>sysWebAccelerationStatCacheEvictions (.1.3.6.1.4.1.3375.2.1.1.2.23.10)</code> , <code>sysWebAccelerationStatCacheHits (.1.3.6.1.4.1.3375.2.1.1.2.23.2)</code> <code>sysWebAccelerationStatCacheMisses (.1.3.6.1.4.1.3375.2.1.1.2.23.3)</code>

The following table shows the required calculations for interpreting metrics on RAM Cache use.

Performance Graph	Graph Metric	Required SNMP OIDs
RAM cache Utilization	Hit Rate	$\langle \text{sysWebAccelerationStatCacheHits1} \rangle / (\langle \text{sysWebAccelerationStatCacheHits1} \rangle + \langle \text{sysWebAccelerationStatCacheMisses1} \rangle) * 100$
RAM cache Utilization	Byte Rate	$\langle \text{sysWebAccelerationStatCacheHitBytes1} \rangle / (\langle \text{sysWebAccelerationStatCacheHitBytes1} \rangle + \langle \text{sysWebAccelerationStatCacheMissBytes1} \rangle) * 100$
RAM cache Utilization	Eviction Rate	$\langle \text{sysWebAccelerationStatCacheEvictions1} \rangle / (\langle \text{sysWebAccelerationStatCacheHits1} \rangle + \langle \text{sysWebAccelerationStatCacheMisses1} \rangle) * 100$

1. For each OID, poll for data. For example, poll OID `sysWebAccelerationStatCacheHits (.1.3.6.1.4.1.3375.2.1.1.2.23.2)`. This results in a value `<sysWebAccelerationStatCacheHits>`.
2. Poll OID `sysWebAccelerationStatCacheHits (.1.3.6.1.4.1.3375.2.1.1.2.23.2)`. This results in a value `<sysWebAccelerationStatCacheMisses>`.
3. Perform the calculation using the OID data. For example, to calculate the value of the Hit Rate graphic metric:

```
<sysWebAccelerationStatCacheHits> / (<sysWebAccelerationStatCacheHits> + <>) *100).
```

Collecting BIG-IP system data on SSL transactions using SNMP

You can use SNMP commands with an OID to gather and interpret data on SSL performance. The following table shows the individual OIDs that you must use to poll for SSL transactions using SNMP.

Performance Graph	Graph Metrics	Required SNMP OIDs
SSL TPS	SSL TPS	<code>sysClientsslStatToNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.9.6)</code>
SSL TPS	SSL TPS	<code>sysClientsslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.9.9)</code>
SSL TPS	SSL TPS	<code>sysServersslStatTotNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.10.6)</code>
SSL TPS	SSL TPS	<code>sysServersslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.10.9)</code>

The following table shows the required calculations for interpreting metrics on SSL transactions using SNMP.

Performance Graph	Graph Metric	Required calculations for SSL TPS
SSL TPS	SSL TPS	<code><DeltaClientsslStatClientTotConns> / (<interval></code>

1. For each OID, poll for data. For example, poll OID `sysClientsslStatToNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.23.2)` and `sysClientsslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.9.9)`.
2. Add the two values together. This results in the value `sysClientsslStatTotConns1`.
3. Poll the two OIDs again, within ten seconds of the previous polls.
4. Again, add the two values together. This results in the value `sysClientsslStatToComms2`.
5. Calculate the delta of the two sums:

```
<DeltaClientsslStatTotConns> = <sysClientsslStatTotConns2> - <sysClientsslStatTotConns1>.
```

6. Perform the calculation on the OID deltas. The value for interval is 10. For example, to calculate the value of the SSL transactions using SNMP:

```
(<DeltaClientsslStatClientTotConns>) / <interval>
```

Collecting BIG-IP system data on CPU usage based on a predefined polling interval

For the CPU[0-n] and Global Host CPU Usage graph metrics, you can instruct the BIG-IP® system to gather and collect CPU usage data automatically, based on a predefined polling interval. Use the `sysMultiHostCpu` and `sysGlobalHostCpu` MIBs.

The following table shows the required OIDs for automatic collection of CPU[0-n] graphic metrics.

Performance Graph	Graph Metric	Required SNMP OIDs
CPU Usage	CPU[0-n]	5-second Polling Interval sysMultiHostCpuUser5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.12) sysMultiHostCpuNice5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.13) sysMultiHostCpuSystem5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.14) sysMultiHostCpuIdle5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.15) sysMultiHostCpuIrq5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.16) sysMultiHostCpuSoftirq5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.17) sysMultiHostCpuLowait5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.18) sysMultiHostCpuUsageRatio5s (.1.3.6.1.4.1.3375.2.1.7.5.2.1.19) sysMultiHostCpuUsageRatio (.1.3.6.1.4.1.3375.2.1.7.5.2.1.11)
CPU Usage	CPU[0-n]	1-minute Polling Interval sysMultiHostCpuUser1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.20) sysMultiHostCpuNice1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.21) sysMultiHostCpuSystem1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.22) sysMultiHostCpuIdle1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.23) sysMultiHostCpuIrq1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.24) sysMultiHostCpuSoftirq1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.25) sysMultiHostCpuLowait1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.26) sysMultiHostCpuUsageRatio1m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.26)
CPU Usage	CPU[0-n]	5-minute Polling Interval sysMultiHostCpuUse5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.28) sysMultiHostCpuNice5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.29) sysMultiHostCpuSystem5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.30) sysMultiHostCpuIdle5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.31) sysMultiHostCpuIrq5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.32) sysMultiHostCpuSoftirq5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.33) sysMultiHostCpuLowait5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.34) sysMultiHostCpuUsageRatio5m (.1.3.6.1.4.1.3375.2.1.7.5.2.1.35)

The following table shows the required OIDs for automatic collection of Global Host CPU Usage graph metrics.

Performance Graph	Graph Metric	Required SNMP OIDs
CPU Usage	Global Host CPU Usage	5-second Polling Interval sysGlobalHostCpuUser5s (.1.3.6.1.4.1.3375.2.1.1.2.20.14) sysGlobalHostCpuNice5s (.1.3.6.1.4.1.3375.2.1.1.2.20.15) sysGlobalHostCpuSystem5s (.1.3.6.1.4.1.3375.2.1.1.2.20.16) sysGlobalHostCpuIdle5s (.1.3.6.1.4.1.3375.2.1.1.2.20.17) sysGlobalHostCpuIrq5s (.1.3.6.1.4.1.3375.2.1.1.2.20.18) sysGlobalHostCpuSoftirq5s (.1.3.6.1.4.1.3375.2.1.1.2.20.19) sysGlobalHostCpuLowait5s (.1.3.6.1.4.1.3375.2.1.1.2.20.20) sysGlobalHostCpuUsageRatio5s (.1.3.6.1.4.1.3375.2.1.1.2.20.21) sysGlobalHostCpuUsageRatio (.1.3.6.1.4.1.3375.2.1.1.2.20.13)
CPU Usage	Global Host CPU Usage	1-minute Polling Interval sysGlobalHostCpuUser1m (.1.3.6.1.4.1.3375.2.1.1.2.20.22) sysGlobalHostCpuNice1m (.1.3.6.1.4.1.3375.2.1.1.2.20.23) sysGlobalHostCpuSystem1m (.1.3.6.1.4.1.3375.2.1.1.2.20.24)

Performance Graph	Graph Metric	Required SNMP OIDs
		sysGlobalHostCpuIdle1m (.1.3.6.1.4.1.3375.2.1.1.2.20.25) sysGlobalHostCpuIrq1m (.1.3.6.1.4.1.3375.2.1.1.2.20.26) sysGlobalHostCpuSoftirq1m (.1.3.6.1.4.1.3375.2.1.1.2.20.27) sysGlobalHostCpuLowait1m (.1.3.6.1.4.1.3375.2.1.1.2.20.28) sysGlobalHostCpuUsageRatio1m (.1.3.6.1.4.1.3375.2.1.1.2.20.29)
CPU Usage	Global Host CPU Usage	5-minute Polling Interval sysGlobalHostCpuUse5m (.1.3.6.1.4.1.3375.2.1.1.2.20.30) sysGlobalHostCpuNice5m (.1.3.6.1.4.1.3375.2.1.1.2.20.31) sysGlobalHostCpuSystem5m (.1.3.6.1.4.1.3375.2.1.1.2.20.32) sysGlobalHostCpuIdle5m (.1.3.6.1.4.1.3375.2.1.1.2.20.33) sysGlobalHostCpuIrq5m (.1.3.6.1.4.1.3375.2.1.1.2.20.34) sysGlobalHostCpuSoftirq5m (.1.3.6.1.4.1.3375.2.1.1.2.20.35) sysGlobalHostCpuLowait5m (.1.3.6.1.4.1.3375.2.1.1.2.20.36) sysGlobalHostCpuUsageRatio5m (.1.3.6.1.4.1.3375.2.1.1.2.20.37)

Collecting BIG-IP system data on CPU usage based on a custom polling interval

For the CPU[0-n], Global Host CPU, and TMM CPU Usage graph metrics, an alternative to instructing the BIG-IP® system to collect CPU usage data automatically, is to do it manually, based on a custom polling interval. For the CPU[0-n] and Global Host CPU graph metrics, use the sysMultiHostCpu and sysGlobalHostCpu MIBs. For the TMM CPU Usage graphic metric, use the sysStatTm MIB.

The following table shows the required SNMP OIDs for collecting CPU data manually.

Performance Graph	Graph Metric	Required SNMP OIDs
CPU Usage	CPU[0-n]	sysMultiHostCpuUser (.1.3.6.1.4.1.3375.2.1.7.5.2.1.4) sysMultiHostCpuNice (.1.3.6.1.4.1.3375.2.1.7.5.2.1.5) sysMultiHostCpuSystem (.1.3.6.1.4.1.3375.2.1.7.5.2.1.6) sysMultiHostCpuIdle (.1.3.6.1.4.1.3375.2.1.7.5.2.1.7) sysMultiHostCpuIrq (.1.3.6.1.4.1.3375.2.1.7.5.2.1.8) sysMultiHostCpuSoftirq (.1.3.6.1.4.1.3375.2.1.7.5.2.1.9) sysMultiHostCpuLowait (.1.3.6.1.4.1.3375.2.1.7.5.2.1.10)
CPU Usage	TMM[0-m]	sysTmmStatTmUsageRatio5s (.1.3.6.1.4.1.3375.2.1.8.2.3.1.37.[tmm_id]) sysTmmStatTmUsageRatio1m (.1.3.6.1.4.1.3375.2.1.8.2.3.1.38.[tmm_id]) sysTmmStatTmUsageRatio5m (.1.3.6.1.4.1.3375.2.1.8.2.3.1.39.[tmm_id])
CPU Usage	Global Host CPU Usage	sysGlobalHostCpuCount (.1.3.6.1.4.1.3375.2.1.1.2.20.4) sysGlobalHostActiveCpu (.1.3.6.1.4.1.3375.2.1.1.2.20.5) sysGlobalHostCpuUser (.1.3.6.1.4.1.3375.2.1.1.2.20.6) sysGlobalHostCpuNice (.1.3.6.1.4.1.3375.2.1.1.2.20.7) sysGlobalHostCpuSystem (.1.3.6.1.4.1.3375.2.1.1.2.20.8) sysGlobalHostCpuIdle (.1.3.6.1.4.1.3375.2.1.1.2.20.9) sysGlobalHostCpuIrq (.1.3.6.1.4.1.3375.2.1.1.2.20.10) sysGlobalHostCpuSoftirq (.1.3.6.1.4.1.3375.2.1.1.2.20.11) sysGlobalHostCpuLowait (.1.3.6.1.4.1.3375.2.1.1.2.20.12)
CPU Usage	Global TMM CPU Usage	sysGlobalTmmStatTmUsageRatio5s (.1.3.6.1.4.1.3375.2.1.1.2.21.34) sysGlobalTmmStatTmUsageRatio1m (.1.3.6.1.4.1.3375.2.1.1.2.21.35) sysGlobalTmmStatTmUsageRatio5m (.1.3.6.1.4.1.3375.2.1.1.2.21.36)

Performance Graph	Graph Metric	Required SNMP OIDs
CPU Usage	TMM CPU Usage	sysStatTmTotalCycles (.1.3.6.1.4.1.3375.2.1.1.2.1.41) sysStatTmIdleCycles (.1.3.6.1.4.1.3375.2.1.1.2.1.42) sysStatTmSleepCycles (.1.3.6.1.4.1.3375.2.1.1.2.1.43)

The following table shows the formulas for calculating metrics on CPU use.

Performance Graph	Graph Metric	Required calculations for CPU use
CPU Usage	CPU[0-n]	(<DeltaCpuUsers>) + (<DeltaCpuNice> + <DeltaCpuSystem> / (<DeltaCpuUsers>) + <DeltaCpuNice> + <DeltaCpuIdle> + <DeltaCpuSystem> + <DeltaCpulrq> + <DeltaCpuSoftirq> + <DeltaCpulowait>) *100
CPU Usage	Global Host CPU Usage	(<DeltaCpuUsers> + <DeltaCpuNice> + <DeltaCpuSystem>) / (<DeltaCpuUsers> + <DeltaCpuNice> + <DeltaCpuIdle> + <DeltaCpuSystem> + <DeltaCpulrq> + <DeltaCpuSoftirq> + <DeltaCpulowait>) *100

1. Poll the OID `sysMultiHostCpuUser` (.1.3.6.1.4.1.3375.2.1.7.5.2.1.4) twice, at a 10-second interval. This results in two values, `sysMultiHostCpuUser1` and `sysMultiHostCpuUser2`.
2. Calculate the delta of the two poll values. For example:

$$\langle \text{DeltaCpuUser} \rangle = \langle \text{sysMultiHostCpuUser2} \rangle - \langle \text{sysMultiHostCpuUser1} \rangle.$$

3. Repeat steps 1 and 2 for each OID pertaining to the **CPU[0-n]** graph metric.
4. Repeat steps 1 and 2 again, using the OIDs from the MIBs **sysStatTm** and **sysGlobalHostCpu**.
5. Calculate the values of the graphic metrics using the formulas in the table above.

Collecting BIG-IP system performance data on new connections using SNMP

You can use SNMP commands with various OIDs to gather and interpret data on the number of new connections on the BIG-IP® system. The following table shows the required OIDs for the Performance graphs in the Configuration utility.

Performance Graph	Graph Metrics	Required SNMP OIDs
New Connections Summary	Client Accepts Server Connects	sysTcpStatAccepts (.1.3.6.1.4.1.3375.2.1.1.2.12.6) sysStatServerTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.14)
Total New Connections	Client Accepts Server Connects	sysStatClientTotConns(.1.3.6.1.4.1.3375.2.1.1.2.1.7) sysStatServerTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.14)
New SSL Profile Connections	SSL Client SSL Server	sysClientsslStatTotNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.9.6), sysClientsslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.9.9) sysServersslStatTotNativeConns(.1.3.6.1.4.1.3375.2.1.1.2.10.6), sysServersslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.10.9)
New Accepts/Connects	Client Accepts Server Connects	sysTcpStatAccepts (.1.3.6.1.4.1.3375.2.1.1.2.12.6) sysTcpStatConnects (.1.3.6.1.4.1.3375.2.1.1.2.12.8)

The following table shows the required calculations for interpreting metrics on new connections.

Performance Graph	Graph Metrics	Required SNMP OIDs
New Connections Summary	Client Accepts Server Connects	<DeltaTcpStatAccept> / <interval> <DeltaStatServerTotConns> / <interval>
Total New Connections	Client Connects Server Connects	<DeltaStatClientTotConns> / <interval> <DeltaStatServerTotConns> / <interval>
New Client SSL Profile Connections	SSL Client SSL Server	(<DeltaClientsslStatTotNativeConns> + <DeltaClientsslStatTotCompatConns>) / <interval> (<DeltaServersslStatTotNativeConns> + <DeltaServersslStatTotCompatConns>) / <interval>
New Accepts/Connects	Client Accepts Server Connects	<DeltaTcpStatAccepts> / <interval> <DeltaTcpStatConnects> / <interval>

- For each OID, perform two separate polls, at an interval of your choice.
For example, for the client accepts metric, poll OID `sysTcpStatAccepts` (.1.3.6.1.4.1.3375.2.1.1.2.12.6) twice, at a 10-second interval. This results in two values, `<sysTcpStatAccepts1>` and `<sysTcpStatAccepts2>`.
- Calculate the delta of the two poll values.
For example, for the client accepts metric, perform this calculation:

```
<DeltaTcpStatAccepts> = <sysTcpStatAccepts2> - <sysTcpStatAccepts1>
```

- Perform a calculation on the OID deltas. The value for `interval` is the polling interval. For example, to calculate the value of the client accepts metric:

```
<DeltaTcpStatAccepts> / <interval>
```

Collecting BIG-IP system performance data on active connections using SNMP

Write an SNMP command with the various OIDs shown in the table to gather and interpret data on the number of active connections on the BIG-IP® system.

Note: To interpret data on active connections, you do not need to perform any calculations on the collected data.

Performance Graph	Graph Metrics	Required SNMP OIDs
Active Connections Summary	Connections	sysStatClientCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.8)
Active Connections Detailed	Client Server SSL Client SSL Server	sysStatClientCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.8) sysStatServerCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.15) sysClientsslStatCurConns (.1.3.6.1.4.1.3375.2.1.1.2.9.2) sysServersslStatCurConns (.1.3.6.1.4.1.3375.2.1.1.2.10.2)

About the RMON MIB file

The BIG-IP[®] system provides the remote network monitoring (RMON) MIB file, RMON-MIB.txt. This file contains remote network monitoring information. The implementation of RMON on the BIG-IP system differs slightly from the standard RMON implementation, in the following ways:

- The BIG-IP system implementation of RMON supports only these four of the nine RMON groups: statistics, history, alarms, and events.
- The RMON-MIB.txt file monitors the BIG-IP system interfaces (that is, sysIfIndex), and not the standard Linux interfaces.
- For hardware reasons, the packet-length-specific statistics in the RMON statistics group offer combined transmission and receiving statistics only. This behavior differs from the behavior described in the definitions of the corresponding OIDs.

Logging Network Firewall Events to IPFIX Collectors

Overview: Configuring IPFIX logging for AFM

You can configure the BIG-IP® system to log information about Advanced Firewall Manager™ (AFM™) processes and send the log messages to remote IPFIX collectors.

The BIG-IP system supports logging of AFM events over the IPFIX protocol. IPFIX logs are raw, binary-encoded strings with their fields and field lengths defined by IPFIX templates. *IPFIX collectors* are external devices that can receive IPFIX templates and use them to interpret IPFIX logs.

Task summary

Perform these tasks to configure IPFIX logging of AFM processes on the BIG-IP® system.

Note: Enabling IPFIX logging impacts BIG-IP system performance.

Assembling a pool of IPFIX collectors

Creating an IPFIX log destination

Creating a publisher

Creating a custom Network Firewall Logging profile

Configuring an LTM virtual server for Network Firewall event logging with IPFIX

About the configuration objects of IPFIX logging for AFM

The configuration process involves creating and connecting the following configuration objects:

Object	Reason	Applies to
Pool of IPFIX collectors	Create a pool of IPFIX collectors to which the BIG-IP system can send IPFIX log messages.	Assembling a pool of IPFIX collectors.
Destination	Create a log destination to format the logs in IPFIX templates, and forward the logs to the IPFIX collectors.	Creating an IPFIX log destination.
Publisher	Create a log publisher to send logs to a set of specified log destinations.	Creating a publisher.

Assembling a pool of IPFIX collectors

Before creating a pool of IPFIX collectors, gather the IP addresses of the collectors that you want to include in the pool. Ensure that the remote IPFIX collectors are configured to listen to and receive log messages from the BIG-IP® system.

These are the steps for creating a pool of IPFIX collectors. The BIG-IP system can send IPFIX log messages to this pool.

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 the IP address for each IPFIX collector that you want to include in the pool:
 - a) Type the collector's IP address in the **Address** field, or select a node address from the **Node List**.
 - b) Type a port number in the **Service Port** field.
By default, IPFIX collectors listen on UDP or TCP port 4739 and Netflow V9 devices listen on port 2055, though the port is configurable at each collector.
- c) Click **Add**.
5. Click **Finished**.

Creating an IPFIX log destination

A log destination of the **IPFIX** type specifies that log messages are sent to a pool of IPFIX collectors. Use these steps to create a log destination for IPFIX collectors.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **IPFIX**.
5. From the **Protocol** list, select **IPFIX** or **Netflow V9**, depending on the type of collectors you have in the pool.
6. From the **Pool Name** list, select an LTM[®] pool of IPFIX collectors.
7. From the **Transport Profile** list, select **TCP**, **UDP**, or any customized profile derived from TCP or UDP.
8. The **Template Retransmit Interval** is the time between transmissions of IPFIX templates to the pool of collectors. The BIG-IP system only retransmits its templates if the **Transport Profile** is a **UDP** profile.
An IPFIX template defines the field types and byte lengths of the binary IPFIX log messages. The logging destination sends the template for a given log type (for example, NAT44 logs or customized logs from an iRule) before sending any of those logs, so that the IPFIX collector can read the logs of that type. The logging destination assigns a template ID to each template, and places the template ID into each log that uses that template.

The log destination periodically retransmits all of its IPFIX templates over a UDP connection. The retransmissions are helpful for UDP connections, which are lossy.
9. The **Template Delete Delay** is the time that the BIG-IP device should pause between deleting an obsolete template and re-using its template ID. This feature is helpful for systems that can create custom IPFIX templates with iRules.
10. The **Server SSL Profile** applies Secure Socket Layer (SSL) or Transport Layer Security (TLS) to TCP connections. You can only choose an SSL profile if the **Transport Profile** is a **TCP** profile. Choose an SSL profile that is appropriate for the IPFIX collectors' SSL/TLS configuration.

SSL or TLS requires extra processing and therefore slows the connection, so we only recommend this for sites where the connections to the IPFIX collectors have a potential security risk.
11. Click **Finished**.

Creating a publisher

A publisher specifies where the BIG-IP[®] system sends log messages for IPFIX logs.

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. Use the **Log Destinations** setting to select an existing IPFIX destination (perhaps along with other destinations for your logs): click any destination name in the **Available** list, and click << to move it to the **Selected** list.
5. Click **Finished**.

Creating a custom Network Firewall Logging profile

You create a custom Logging profile to log messages about BIG-IP® system Network Firewall events.

1. On the Main tab, click **Security > Event Logs > Logging Profiles**.
The Logging Profiles list screen opens.
2. Click **Create**.
The Create New Logging Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select the **Network Firewall** check box.
5. If you want to enable optional subscriber ID logging:
 - a) Select the **Network Address Translation** check box.
 - b) Then in the Network Address Translation area, select the **Log Subscriber ID** check box.
 - c) Click **Network Firewall**.
6. In the Network Firewall area, from the **Publisher** list, select the IPFIX publisher the BIG-IP system uses to log Network Firewall events.
7. Set an **Aggregate Rate Limit** to define a rate limit for all combined network firewall log messages per second.
Beyond this rate limit, log messages are not logged.
8. For the **Log Rule Matches** setting, select how the BIG-IP system logs packets that match ACL rules.
You can select any or all of the options.

Option	Description
---------------	--------------------

Option	Enables or disables logging of packets that match ACL rules configured with:
---------------	--

Accept	<code>action=Accept</code>
---------------	----------------------------

Drop	<code>action=Drop</code>
-------------	--------------------------

Reject	<code>action=Reject</code>
---------------	----------------------------

When an option is selected, you can configure a rate limit for log messages of that type.

9. Select the **Log IP Errors** check box, to enable logging of IP error packets.
When this setting is enabled, you can configure a rate limit for log messages of this type.
10. Select the **Log TCP Errors** check box, to enable logging of TCP error packets.
When this is enabled, you can configure a rate limit for log messages of this type.
11. Select the **Log TCP Events** check box, to enable logging of open and close of TCP sessions.
When this is enabled, you can configure a rate limit for log messages of this type.
12. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for Network Firewall log events.
13. Enable the **Log Geolocation IP Address** setting to specify that when a geolocation event causes a network firewall action, the associated IP address is logged.

14. From the **Storage Format** list, select how the BIG-IP system formats the log.

Option	Description
None	Specifies the default format type in which the BIG-IP system logs messages to a remote Syslog server, for example: <pre>"management_ip_address", "bigip_hostname", "context_type", "context_name", "src_ip", "dest_ip", "src_port", "dest_port", "vlan", "protocol", "route_domain", "acl_rule_name", "action", "drop_reason"</pre>
Field-List	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Specify the order the fields display in the log. • Specify the delimiter that separates the content in the log. The default delimiter is the comma character.
User-Defined	Allows you to: <ul style="list-style-type: none"> • Select, from a list, the fields to be included in the log. • Cut and paste, in a string of text, the order the fields display in the log.

15. In the IP Intelligence area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log source IP addresses, which are identified and configured for logging by an IP Intelligence policy.

Note: The IP Address Intelligence feature must be enabled and licensed.

16. Set an **Aggregate Rate Limit** to define a rate limit for all combined IP Intelligence log messages per second.
 Beyond this rate limit, log messages are not logged.
17. Enable the **Log Translation Fields** setting to log both the original IP address and the NAT-translated IP address for IP Intelligence log events.
18. In the Traffic Statistics area, from the **Publisher** list, select the publisher that the BIG-IP system uses to log traffic statistics.
19. For the **Log Timer Events** setting, enable **Active Flows** to log the number of active flows each second.
20. For the **Log Timer Events** setting, enable **Reaped Flows** to log the number of reaped flows, or connections that are not established because of system resource usage levels.
21. For the **Log Timer Events** setting, enable **Missed Flows** to log the number of packets that were dropped because of a flow table miss. A flow table miss occurs when a TCP non-SYN packet does not match an existing flow.
22. For the **Log Timer Events** setting, enable **SYN Cookie (Per Session Challenge)** to log the number of SYN cookie challenges generated each second.
23. For the **Log Timer Events** setting, enable **SYN Cookie (White-listed Clients)** to log the number of SYN cookie clients whitelisted each second.
24. Click **Finished**.

Now you can assign this custom network firewall Logging profile to a virtual server.

Configuring an LTM virtual server for Network Firewall event logging with IPFIX

Ensure that at least one log publisher exists on the BIG-IP® system.

Assign a custom Network Firewall Logging profile to a virtual server when you want the BIG-IP system to log Network Firewall events to IPFIX collectors on the traffic that the virtual server processes.

Note: This task applies only to LTM[®]-provisioned systems.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Security > Policies**.
The screen displays policy settings for the virtual server.
4. In the **Log Profile** setting, select **Enabled**. Then, select one or more profiles that log specific events to IPFIX collectors, and move them from the **Available** list to the **Selected** list.

*Note: To log global, self IP, and route domain contexts, you must enable a Publisher in the **global-network** profile.*

5. Click **Update** to save the changes.

Implementation result

Now you have an implementation in which the BIG-IP[®] system logs messages about AFM[™] events and sends the log messages to a pool of IPFIX collectors.

Note: Network firewall events are logged only for rules or policies for which logging is enabled.

IPFIX Templates for AFM Events

Overview: IPFIX Templates for AFM events

The IP Flow Information Export (IPFIX) Protocol is a logging mechanism for IP events. This appendix defines the IPFIX Information Elements (IEs) and Templates used to log the F5[®] Application Firewall Manager[™] (AFM[™]) events. An *IE* is the smallest form of useful information in an IPFIX log message, such as an IP address or a timestamp for the event. An *IPFIX template* is an ordered collection of specific IEs used to record one IP event, such as the acceptance of a network packet.

About IPFIX Information Elements for AFM events

Information Elements (IEs) are individual fields in an IPFIX template. An IPFIX template describes a single Advanced Firewall Manager[™] (AFM[™]) event.

IANA-defined IPFIX Information Elements

IANA maintains a list of standard IPFIX Information Elements (IEs), each with a unique Element Identifier. The F5[®] AFM[™] IPFIX implementation uses a subset of these IEs to publish AFM events. This subset is summarized in the table.

Information Element (IE)	ID	Size (Bytes)
destinationIPv4Address	12	4
destinationIPv6Address	28	16
destinationTransportPort	11	2
ingressVRFID	234	4
observationTimeMilliseconds	323	8
protocolIdentifier	4	1
sourceIPv4Address	8	4
sourceIPv6Address	27	16
sourceTransportPort	7	2

IPFIX enterprise Information Elements

IPFIX provides for enterprises to define their own Information Elements. F5[®] currently uses the following non-standard IEs for AFM[™] events:

Information Element (IE)	ID	Size (Bytes)
aclPolicyName	12276 - 26	Variable
aclPolicyType	12276 - 25	Variable
aclRuleName	12276 - 38	Variable
action	12276 - 39	Variable

Information Element (IE)	ID	Size (Bytes)
attackType	12276 - 46	Variable
bigipHostName	12276 - 10	Variable
bigipMgmtIPv4Address	12276 - 5	4
bigipMgmtIPv6Address	12276 - 6	16
contextName	12276 - 9	Variable
contextType	12276 - 24	Variable
destinationFqdn	12276 - 99	Variable
destinationGeo	12276 - 43	Variable
deviceProduct	12276 - 12	Variable
deviceVendor	12276 - 11	Variable
deviceVersion	12276 - 13	Variable
dosAttackEvent	12276 - 41	Variable
dosAttackId	12276 - 20	4
dosAttackName	12276 - 21	Variable
dosPacketsDropped	12276 - 23	4
dosPacketsReceived	12276 - 22	4
dropReason	12276 - 40	Variable
errdefsMsgNo	12276 - 4	4
flowId	12276 - 3	8
ipfixMsgNo	12276 - 16	4
ipintelligencePolicyName	12276 - 45	Variable
ipintelligenceThreatName	12276 - 42	Variable
logMsgDrops	12276 - 96	4
logMsgName	12276 - 97	Variable
logprofileName	12276 - 95	Variable
messageSeverity	12276 - 1	1
msgName	12276 - 14	Variable
partitionName	12276 - 2	Variable
saTransPool	12276 - 37	Variable
saTransType	12276 - 36	Variable
sourceFqdn	12276 - 98	Variable
sourceGeo	12276 - 44	Variable
sourceUser	12276 - 93	Variable
transDestinationIPv4Address	12276 - 31	4
transDestinationIPv6Address	12276 - 32	16
transDestinationPort	12276 - 33	2

Information Element (IE)	ID	Size (Bytes)
transIpProtocol	12276 - 27	1
transRouteDomain	12276 - 35	4
transSourceIPv4Address	12276 - 28	4
transSourceIPv6Address	12276 - 29	16
transSourcePort	12276 - 30	2
transVlanName	12276 - 34	Variable
vlanName	12276 - 15	Variable

Note: IPFIX, unlike NetFlow v9, supports variable-length IEs, where the length is encoded within the field in the Data Record. NetFlow v9 collectors (and their variants) cannot correctly process variable-length IEs, so they are omitted from logs sent to those collector types.

About individual IPFIX templates for each event

F5® uses IPFIX templates to publish AFM™ events.

Network accept or deny

This IPFIX template is used whenever a network packet is accepted or denied by an AFM™ firewall.

Information Element (IE)	ID	Size (Bytes)	Notes
aclPolicyName	12276 - 26	Variable	This IE is omitted for NetFlow v9.
aclPolicyType	12276 - 25	Variable	This IE is omitted for NetFlow v9.
aclRuleName	12276 - 38	Variable	This IE is omitted for NetFlow v9.
action	12276 - 39	Variable	This IE is omitted for NetFlow v9.
bigipHostName	12276 - 10	Variable	This IE is omitted for NetFlow v9.
bigipMgmtIPv4Address	12276 - 5	4	
bigipMgmtIPv6Address	12276 - 6	16	
contextName	12276 - 9	Variable	This IE is omitted for NetFlow v9.
contextType	12276 - 24	Variable	This IE is omitted for NetFlow v9.
observationTimeMilliseconds	323	8	
destinationFqdn	12276 - 99	Variable	This IE is omitted for NetFlow v9.
destinationGeo	12276 - 43	Variable	This IE is omitted for NetFlow v9.
destinationIPv4Address	12	4	
destinationIPv6Address	28	16	
destinationTransportPort	11	2	
deviceProduct	12276 - 12	Variable	This IE is omitted for NetFlow v9.
deviceVendor	12276 - 11	Variable	This IE is omitted for NetFlow v9.

Information Element (IE)	ID	Size (Bytes)	Notes
deviceVersion	12276 - 13	Variable	This IE is omitted for NetFlow v9.
dropReason	12276 - 40	Variable	This IE is omitted for NetFlow v9.
msgName	12276 - 14	Variable	This IE is omitted for NetFlow v9.
errdefsMsgNo	12276 - 4	4	
flowId	12276 - 3	8	
ipfixMsgNo	12276 - 16	4	
protocolIdentifier	4	1	
messageSeverity	12276 - 1	1	
partitionName	12276 - 2	Variable	This IE is omitted for NetFlow v9.
ingressVRFID	234	4	
saTransPool	12276 - 37	Variable	This IE is omitted for NetFlow v9.
saTransType	12276 - 36	Variable	This IE is omitted for NetFlow v9.
sourceFqdn	12276 - 98	Variable	This IE is omitted for NetFlow v9.
sourceGeo	12276 - 44	Variable	This IE is omitted for NetFlow v9.
sourceIPv4Address	8	4	
sourceIPv6Address	27	16	
sourceTransportPort	7	2	
sourceUser	12276 - 93	Variable	This IE is omitted for NetFlow v9.
transDestinationIPv4Address	12276 - 31	4	
transDestinationIPv6Address	12276 - 32	16	
transDestinationPort	12276 - 33	2	
transIpProtocol	12276 - 27	1	
transRouteDomain	12276 - 35	4	
transSourceIPv4Address	12276 - 28	4	
transSourceIPv6Address	12276 - 29	16	
transSourcePort	12276 - 30	2	
transVlanName	12276 - 34	Variable	This IE is omitted for NetFlow v9.
vlanName	12276 - 15	Variable	This IE is omitted for NetFlow v9.

DoS device

Information Element (IE)	ID	Size (Bytes)	Notes
action	12276 - 39	Variable	This IE is omitted for NetFlow v9.
bigipHostName	12276 - 10	Variable	This IE is omitted for NetFlow v9.
bigipMgmtIPv4Address	12276 - 5	4	
bigipMgmtIPv6Address	12276 - 6	16	

Information Element (IE)	ID	Size (Bytes)	Notes
contextName	12276 - 9	Variable	This IE is omitted for NetFlow v9.
observationTimeMilliseconds	323	8	
destinationIPv4Address	12	4	
destinationIPv6Address	28	16	
destinationTransportPort	11	2	
deviceProduct	12276 - 12	Variable	This IE is omitted for NetFlow v9.
deviceVendor	12276 - 11	Variable	This IE is omitted for NetFlow v9.
deviceVersion	12276 - 13	Variable	This IE is omitted for NetFlow v9.
dosAttackEvent	12276 - 41	Variable	This IE is omitted for NetFlow v9.
dosAttackId	12276 - 20	4	
dosAttackName	12276 - 21	Variable	This IE is omitted for NetFlow v9.
dosPacketsDropped	12276 - 23	4	
dosPacketsReceived	12276 - 22	4	
msgName	12276 - 14	Variable	This IE is omitted for NetFlow v9.
errdefsMsgNo	12276 - 4	4	
flowId	12276 - 3	8	
ipfixMsgNo	12276 - 16	4	
messageSeverity	12276 - 1	1	
partitionName	12276 - 2	Variable	This IE is omitted for NetFlow v9.
ingressVRFID	234	4	
sourceIPv4Address	8	4	
sourceIPv6Address	27	16	
sourceTransportPort	7	2	
vlanName	12276 - 15	Variable	This IE is omitted for NetFlow v9.

IP intelligence

Information Element (IE)	ID	Size (Bytes)	Notes
action	12276 - 39	Variable	This IE is omitted for NetFlow v9.
attackType	12276 - 46	Variable	This IE is omitted for NetFlow v9.
bigipHostName	12276 - 10	Variable	This IE is omitted for NetFlow v9.
bigipMgmtIPv4Address	12276 - 5	4	
bigipMgmtIPv6Address	12276 - 6	16	
contextName	12276 - 9	Variable	This IE is omitted for NetFlow v9.
contextType	12276 - 24	Variable	This IE is omitted for NetFlow v9.
observationTimeMilliseconds	323	8	

Information Element (IE)	ID	Size (Bytes)	Notes
destinationIPv4Address	12	4	
destinationIPv6Address	28	16	
destinationTransportPort	11	2	
deviceProduct	12276 - 12	Variable	This IE is omitted for NetFlow v9.
deviceVendor	12276 - 11	Variable	This IE is omitted for NetFlow v9.
deviceVersion	12276 - 13	Variable	This IE is omitted for NetFlow v9.
msgName	12276 - 14	Variable	This IE is omitted for NetFlow v9.
errdefsMsgNo	12276 - 4	4	
flowId	12276 - 3	8	
ipfixMsgNo	12276 - 16	4	
ipintelligencePolicyName	12276 - 45	Variable	This IE is omitted for NetFlow v9.
ipintelligenceThreatName	12276 - 42	Variable	This IE is omitted for NetFlow v9.
protocolIdentifier	4	1	
messageSeverity	12276 - 1	1	
partitionName	12276 - 2	Variable	This IE is omitted for NetFlow v9.
ingressVRFID	234	4	
saTransPool	12276 - 37	Variable	This IE is omitted for NetFlow v9.
saTransType	12276 - 36	Variable	This IE is omitted for NetFlow v9.
sourceIPv4Address	8	4	
sourceIPv6Address	27	16	
sourceTransportPort	7	2	
transDestinationIPv4Address	12276 - 31	4	
transDestinationIPv6Address	12276 - 32	16	
transDestinationPort	12276 - 33	2	
transIpProtocol	12276 - 27	1	
transRouteDomain	12276 - 35	4	
transSourceIPv4Address	12276 - 28	4	
transSourceIPv6Address	12276 - 29	16	
transSourcePort	12276 - 30	2	
transVlanName	12276 - 34	Variable	This IE is omitted for NetFlow v9.
vlanName	12276 - 15	Variable	This IE is omitted for NetFlow v9.

Log Throttle

Information Element (IE)	ID	Size (Bytes)	Notes
bigipHostName	12276 - 10	Variable	This IE is omitted for NetFlow v9.

Information Element (IE)	ID	Size (Bytes)	Notes
bigipMgmtIPv4Address	12276 - 5	4	
bigipMgmtIPv6Address	12276 - 6	16	
observationTimeMilliseconds	323	8	
deviceProduct	12276 - 12	Variable	This IE is omitted for NetFlow v9.
deviceVendor	12276 - 11	Variable	This IE is omitted for NetFlow v9.
deviceVersion	12276 - 13	Variable	This IE is omitted for NetFlow v9.
msgName	12276 - 14	Variable	This IE is omitted for NetFlow v9.
errdefsMsgNo	12276 - 4	4	
ipfixMsgNo	12276 - 16	4	
messageSeverity	12276 - 1	1	
contextType	12276 - 24	Variable	This IE is omitted for NetFlow v9.
contextName	12276 - 9	Variable	This IE is omitted for NetFlow v9.
logprofileName	12276 - 95	Variable	This IE is omitted for NetFlow v9.
logMsgName	12276 - 97	Variable	This IE is omitted for NetFlow v9.
logMsgDrops	12276 - 96	4	

Legal Notices

Legal notices

Publication Date

This document was published on February 5, 2019.

Publication Number

MAN-0439-08

Copyright

Copyright © 2018, 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

For a current list of F5 trademarks and service marks, see <http://www.f5.com/about/guidelines-policies/trademarks>.

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: <https://f5.com/about-us/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.

Legal Notices

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.

Index

A

- access control
 - and SNMP data *146*
- access levels
 - assigning *146*
- access levels, assigning *145*
- actions
 - firewall rule *12*
- active connections data, collecting using SNMP commands *181*
- ADC mode
 - cofngi for firewall *64*
 - network firewall configuration *63*
 - setting for firewall *9, 65*
- adding a firewall rule in a list *20*
- address list
 - creating *26, 67, 74*
- addresses
 - lists *26*
- AFM
 - IANA IPFIX IEs for *189*
 - IPFIX template for DoS device events *192*
 - IPFIX template for IP intelligence events *193*
 - IPFIX template for log throttle events *194*
 - IPFIX template for network session *191*
- AFM-related SNMP traps, defined *150*
- algorithms
 - for eviction policy *116*
- allowing access
 - with a firewall rule *75*
- AOM traps
 - table of always-on management-related *151*
- AOM-related SNMP traps, defined *151*
- application virtual server
 - denying access with firewall rules *68*
- ASM-related SNMP traps, defined *153*
- authentication-related SNMP traps, defined *155*
- automatic compilation
 - of firewall rules *77*
- automatic deployment
 - of firewall rules *78*
- AVR-related SNMP traps, defined *154*

B

- BIG-IP DNS-related SNMP traps, defined *158*
- BIG-IP system information
 - specifying *145*
- BIG-IP system processes, monitoring using SNMP *174*
- blacklisting an IP address *49*
- blacklist
 - removing IP address *49*
- blacklist categories *47*
- blacklist category
 - adding an IP address to *49*
 - defining *48*
 - removing an IP address *49*

- blacklist matching policy *51*
- blocking response page
 - configuring in HTTP profile *105*

C

- CGNAT
 - using with Firewall NAT *84*
- checking IP address reputation
 - for a route domain *53*
 - globally *53*
- client access
 - allowing *145*
- collectors
 - for IPFIX *183*
- compilation statistics
 - viewing for firewall *80*
- compiling rules
 - manual or automatic *77*
 - manually *77*
- conflicting rules
 - resolving *17*
 - viewing *17*
- connection limits
 - about *115*
- connections
 - collecting data about active *181*
 - collecting data about HTTP *174*
 - collecting data about new *180*
 - collecting data about RAM *176*
 - collecting data about SSL *177*
 - collecting data about throughput *175*
- context
 - for network firewall rule *14*
- CPU usage
 - collecting based on a custom polling interval *179*
 - collecting based on a predefined polling interval *177*
- creating a firewall policy *37*
- creating a firewall rule
 - to deny ICMP packets *66*
- creating a firewall rule list *20*
- creating a list of addresses *26, 67, 74*
- creating a list of ports *27*
- creating a network firewall rule *15*
- creating a rule from a log entry *131*
- creating a rule list in a policy *22*
- creating a schedule *29*
- custom profiles
 - and Network Firewall Logging *129, 140, 185*
 - and Protocol Security logging *112*
 - and Protocol Security Logging *106*

D

- DDM-related SNMP traps, defined *155*
- default access levels, modifying *145*
- default deny policy *10, 72*
- default dropping of traffic *10, 72*

- default rejection of traffic *10, 72*
- denying access
 - with a firewall rule *74*
 - with firewall rules *67*
- denying all access
 - with a firewall rule *68*
- deploying rules
 - manual or automatic *78*
 - manually *78*
- destination IP addresses
 - translating with Firewall NAT *90*
- destination SNMP managers, specifying *148*
- destination translation
 - static PAT *90*
- destinations
 - for IPFIX logging *184*
 - for logging *111, 139*
 - for remote high-speed logging *111, 139*
- Device NAT
 - adding a match rule *93*
- DNS cache
 - and DNS resolver *25*
- DNS resolver
 - configuring for network firewall *26*
 - creating *26*
- DNS resolver
 - and DNS cache *25*
- DoS attack data, collecting using SNMP commands *172*
- DoS-related SNMP traps, defined *157*
- dynamic PAT
 - source translation *86, 88*
- dynamic routing, and viewing SNMP traps *173*

E

- editing a rule with the inline editor *34*
- editing rules
 - with inline editor *32*
- enterprise MIB files
 - and SNMP *170, 172*
 - and viewing objects *171, 173*
 - downloading *171, 173*
- evasion techniques checks *102*
- event logs
 - viewing *107, 131*
 - viewing enforced events *42*
 - viewing staged events *43*
- events
 - setting SNMP traps *147*
- eviction policies
 - about *115*
- eviction policy
 - algorithms *116*
 - assigning globally *117*
 - creating *115*

F

- F5-BIGIP-COMMON-MIB.txt, and viewing SNMP traps *173*
- F5-BIGIP-LOCAL-MIB.txt, and viewing SNMP traps *171*
- feed list
 - defining *51*

- feed list settings *50*
- feed lists *49, 51*
- firewall
 - configuring firewall mode *10, 72*
 - dropping traffic not explicitly allowed *10, 72*
 - dropping traffic that matches no rule *10*
 - NAT, about *83*
 - rejecting traffic not explicitly allowed *10, 72*
 - rejecting traffic that matches no rule *10*
 - setting ADC mode *9, 64, 65*
- firewall contexts *14*
- firewall data, collecting using SNMP commands *172*
- firewall mode
 - setting for firewall *10, 72*
- Firewall mode
 - network firewall configuration *71*
- firewall NAT policy
 - associating with virtual servers *95*
- firewall policies
 - enforcing *37*
 - evaluating *37*
 - resources to compile *79*
 - staging *37*
- firewall policy
 - adding to a virtual server *41*
 - assigning globally *40*
 - creating *37*
 - defining *37*
 - viewing compilation statistics *80*
- firewall policy rule
 - creating *37*
 - creating with inline rule editor *32*
 - editing with inline rule editor *34*
- firewall rule
 - adding to a rule list *20*
 - allow access to a single network *75*
 - allow access to an address list *74*
 - assigning protocol inspection *59*
 - creating for management port *15*
 - creating from a log entry *131*
 - creating in a policy *37*
 - creating with inline rule editor *32*
 - denying access to specific servers *67*
 - denying ICMP packets *66*
 - editing with inline rule editor *34*
 - resolving conflicting *17*
 - resolving infrequently used *19*
 - resolving never used *19*
 - resolving redundant *17*
 - viewing conflicting *17*
 - viewing infrequently used *19*
 - viewing never used *19*
 - viewing redundant *17*
- firewall rule list
 - creating *20*
- firewall rules
 - actions *12*
 - conflicting *17*
 - context ordering *12*
 - denying access to specific networks *68*
 - overlapping *17*
 - redundant *17*

- firewall rules (*continued*)
 - unused or infrequently used rules 18
- flow table overflows
 - preventing globally 117
- flows
 - preventing overflow attacks on a route domain 118
 - preventing overflow attacks on a virtual server 117
- FQDNs
 - about resolving in firewall rules 25
 - configuring DNS resolver 26
- fully qualified domain names
 - configuring DNS resolver 26

G

- general SNMP traps, defined 157
- global actions
 - about dropping traffic 9
 - allowing traffic 9
 - rejecting traffic 9
- global context
 - assigning IP intelligence policy 53
 - creating a firewall policy 40
 - viewing compilation statistics 80
- global drop rule 10
- global eviction policy
 - assigning 117
- global NAT policy 92
- global reject rule 10

H

- hardware-related SNMP traps, defined 162
- high-availability system-related SNMP traps, defined 166
- high-speed logging
 - and server pools 110, 138
- HOST-RESOURCES MIB, using in a script 174
- HTTP
 - and evasion techniques checks 102
 - configuring request checks 102
- HTTP profiles
 - attaching security profile 98
 - configuring mandatory headers 105
 - configuring the blocking response page 105
 - creating 99
- HTTP protocol validation
 - checking, importance 101
- HTTP rates data, collecting using SNMP commands 174
- HTTP request checks
 - allowing or disallowing files by type 104
 - configuring length checks 103
 - specifying HTTP methods to allow 103
- HTTP RFC compliance
 - ensuring in HTTP traffic 101
- HTTP security
 - fine-tuning profile settings 101
 - increasing 101
- HTTP security profiles
 - allowing files by type 104
 - configuring length checks 103
 - creating 98
 - disallowing files by type 104

- HTTP security profiles (*continued*)
 - fine-tuning 101
 - specifying allowable methods 103
- HTTP traffic
 - blocking evasion techniques 102
 - configuring protocol compliance checks 101
 - creating security profile 99
 - securing 97

I

- ICMP packet
 - tracing 125
- infrequently used rules
 - resolving 19
 - viewing 19
- inline rule editor
 - enabling 32
- inspecting protocol anomalies
 - creating a profile 55
- interfaces
 - tagging 65, 73
- IP address intelligence
 - assigning globally 53
 - assigning to a route domain 53
 - assigning to a virtual server 53
 - blacklisting an IP address 49
 - creating a blacklist category 48
 - creating a feed list 51
 - feed lists 50
 - removing an IP address from blacklist 49
- IP intelligence 45
- IP Intelligence
 - downloading the database 46
 - enabling 46
- IP intelligence policy 51
- IPFIX
 - AFM template overview 189
 - and server pools 183
 - template for accept or deny through AFM firewall session 191
 - template for DoS device events 192
 - template for IP intelligence events 193
 - template for log throttle events 194
- IPFIX collectors
 - and destinations for log messages 184
 - and publishers for log messages 184
- IPFIX logging
 - and AFM 183
 - configuring 183
 - creating a destination 184
 - overview 183
- iprep.autoupdate command 46

L

- license-related SNMP traps, defined 167
- lists of addresses 25
- lists of ports 25
- log entry
 - using to create a firewall rule 131
- logging

logging (*continued*)

- and destinations 111, 139, 184
- and network firewall 129, 137
- and Network Firewall profiles 129, 140, 185
- and pools 110, 138, 183
- and protocol security 106
- and Protocol Security 108
- and Protocol Security profiles 106, 112
- and publishers 112, 140, 184
- protocol inspection items 60

Logging profile

- and network firewalls 66, 74, 131, 142
- and Protocol Security events 107, 113
- and the network firewall 186

Logging profiles, disabling 107, 113, 134, 142

logging-related SNMP traps, defined 169

LTM-related SNMP traps, defined 167

M

manual compilation

- of firewall rules 77

manual deployment

- of firewall rules 78

memory usage data, collecting using SNMP commands 174

MIB files

- about enterprise 170, 172
- about RMON 182
- and viewing enterprise objects 171, 173

N

name resolution

- using the BIG-IP system 26

NAT

- adding a global match rule 93
- adding a route domain match rule 94
- adding a virtual server match rule 95
- and proxy ARP 83
- configuring on a route domain 93
- creating a match rule 91
- creating a policy 91
- creating policies 91
- global policy 92
- specifying a policy for the device 92
- specifying destination translations 90
- specifying source translations 84
- translating destination addresses and ports 90
- translating source addresses and ports 84
- using with CGNAT 84
- using with network firewall 83

NAT policy

- context 92
- for device 92
- on route domain 92
- on virtual server 92

NET-SNMP MIB files, downloading 171, 173

network firewall

- about address lists 26
- about inline rule editing 31
- about modes 9
- about rule lists 19

network firewall (*continued*)

- about rules 11
- and logging 42
- compiler statistics 79
- compiling rules 77
- context 12
- deploying in ADC mode 63
- deploying in Firewall mode 71
- IP Intelligence 49, 51
- policy compilation 79
- port lists 27
- viewing compilation statistics 80

Network Firewall

- about 9
- about policies 37
- addresses 25
- blacklist categories 47
- blacklists 45
- enabling a VLAN on a virtual server 66, 74
- feed lists 45
- IP intelligence 45
- IP Intelligence 47
- policy rule precedence 37
- ports 25
- protocol inspection 55
- schedules 29
- whitelists 45

network firewall logging

- overview of local 129

Network Firewall logging

- disabling 107, 113, 134, 142

Network Firewall Logging

- customizing profiles 129, 140, 185

Network Firewall Logging profile, assigning to virtual server 131, 142, 186

network firewall logging, configuring of high-speed remote 138

network firewall logging, overview of high-speed remote 137

network firewall policy

- and self IP addresses 41

network virtual server

- denying access with firewall rules 67

network-related SNMP traps, defined 169

new connections data, collecting using SNMP commands 180

notifications, sending 148

O

options

- specifying a device firewall policy 92

P

packet

- tracing a TCP 123
- tracing a UDP 124
- tracing an ICMP 125
- tracing an SCTP 124

packet tester

- tracing a TCP packet 123
- tracing a UDP packet 124

- packet tester (*continued*)
 - tracing an ICMP packet 125
 - tracing an SCTP packet 124
- packet testing 123
- packet trace results 125
- packet tracing 123
- packets
 - about testing through the firewall 123
- ping
 - preventing with a firewall rule 66
- policies
 - creating for Firewall NAT 91
- policy logging
 - enforced policies 42
 - staged policies 42
- pools
 - for high-speed logging 110, 138
 - for IPFIX 183
- port block allocation mode
 - source translation 89
- port list
 - creating 27
- port lists 27
- port misuse policy
 - creating 120
- preventing port misuse 120
- profiles
 - and disabling Network Firewall logging 107, 113, 134, 142
 - creating for HTTP 99
 - creating for HTTP security 98
 - creating for Network Firewall Logging 129, 140, 185
 - creating for Protocol Security logging 112
 - creating for Protocol Security Logging 106
 - customizing settings for HTTP 103–105
 - HTTP security, attaching 98
- protocol anomalies
 - adding inspection items 56
 - inspecting with a profile 55
 - viewing inspection items 56
- protocol anomaly inspection 55
- protocol anomaly rules 57
- protocol inspection
 - adding items 56
 - assigning to a firewall rule 59
 - assigning to a virtual server 58
 - creating a profile 55
 - viewing items 56
 - viewing logs 60
- protocol security
 - configuring for HTTP traffic 97
- Protocol Security
 - viewing event logs locally 107
- protocol security logging
 - overview of local 106
- Protocol Security logging
 - configuring 109
 - customizing profiles 112
 - overview 108
- Protocol Security Logging
 - customizing profiles 106

- Protocol Security Logging profile, assigning to virtual server 107, 113
- publishers
 - and logging 184
 - creating for logging 112, 140

R

- RAM cache data, collecting using SNMP commands 176
- redundant rules
 - resolving 17
 - viewing 17
- remote servers
 - and destinations for log messages 111, 139
 - for high-speed logging 110, 138
- request checks
 - configuring for HTTP protocol 102
- requests
 - accepting 145
- resolver DNS cache
 - about 25
- resolving DNS addresses
 - in network firewall rules 26
- resolving FQDNs
 - with DNS resolver 25
- resolving rule conflicts 17
- resolving rule redundancy 17
- RFC compliance
 - ensuring in HTTP traffic 101
- RMON MIB file, and SNMP 182
- route domain
 - adding a NAT match rule 94
 - applying a port misuse policy 122
 - applying a timer policy 122
 - assigning IP intelligence policy 53
- route domain NAT policy 92
- route domains
 - adding an eviction policy 118
 - configuring firewall NAT 93
 - configuring for firewall policy 41
 - controlling flows 118
 - limiting flows 118
 - setting a firewall policy 41
 - viewing compilation statistics 80
- rule
 - creating in a firewall policy 37
- rule compilation mode
 - automatic 77
 - manual 77
- rule deployment mode
 - automatic 77
 - manual 77
- rule list
 - activating in a policy 22
 - viewing compilation statistics 80
- rule lists 19
- rules
 - conflicting 17
 - creating for firewall with the inline editor 32
 - editing inline 31
 - redundant 17
 - stale 18

S

- schedule
 - creating 29
- scheduling
 - firewall rules 29
- SCTP packet
 - tracing 124
- security profiles
 - creating for HTTP 99
 - viewing statistics 98, 100
- self IP addresses
 - enforcing a firewall policy 41
 - setting firewall policies 41
 - staging a firewall policy 41
- self IPs
 - viewing compilation statistics 80
- servers
 - and destinations for log messages 111, 139, 184
 - and publishers for IPFIX logs 184
 - and publishers for log messages 112, 140
 - for high-speed logging 110, 138
- service policies
 - types defined 119
- service policy
 - applying to firewall rule 121
 - applying to route domain 122
 - applying to self IP 122
 - associating port misuse policy 121
 - associating timer policy 121
 - creating 121
 - definition 119
- setting ADC mode 9, 64, 65
- setting the compilation mode 77
- setting the deployment mode 78
- setting the global drop or reject rule 10
- SNMP
 - and enterprise MIB files 170, 172
 - and monitoring BIG-IP system processes 174
 - and the RMON MIB file 182
- SNMP access levels, assigning 145
- SNMP agent configuration
 - overview of 145
- SNMP agents
 - allowing access to 145
- SNMP alerts, sending 147
- SNMP commands
 - collecting active connections data 181
 - collecting DoS attack data 172
 - collecting firewall rule data 172
 - collecting HTTP rates data 174
 - collecting memory usage data 174
 - collecting network firewall data 172
 - collecting new connections data 180
 - collecting RAM cache data 176
 - collecting SSL transactions 177
 - collecting throughput rates data 175
- SNMP data
 - and controlling access 146
 - controlling access to 146
- SNMP events, setting traps 147
- SNMP manager, and downloading MIB files 171, 173

- SNMP notifications, sending 148
- SNMP protocol, managing 145
- SNMP traps
 - about troubleshooting 150
 - and dynamic routing 173
 - creating 149
 - defined 147
 - enabling 147
 - table of advanced firewall manager-related 150
 - table of application security management-related 153
 - table of authentication-related 155
 - table of AVR-related 154
 - table of digital diagnostic monitoring-related 155
 - table of DoS-related 157
 - table of general 157
 - table of global traffic management-related 158
 - table of hardware-related 162
 - table of high-availability system-related 166
 - table of license-related 167
 - table of local traffic management-related 167
 - table of logging-related 169
 - table of network-related 169
 - table of vCMP-related 170
 - table of VIPRION-related 170
 - viewing 149, 171, 173
- SNMP v1 and v2c traps, setting destination 148
- SNMP v3 traps, setting destination 148
- Snort
 - supported commands 57
- source 86
- source IP addresses
 - translating with Firewall NAT 84
- source translation
 - dynamic PAT 86, 88
 - port block allocation mode 89
 - static PAT 86
- specifying for the device 92
- SSL transactions, collecting using SNMP commands 177
- stale rules
 - resolving 19
 - viewing 19
- static PAT
 - destination translation 90
 - source translation 86
- statistics
 - viewing for security profiles 98, 100
- system information
 - specifying 145

T

- tagged interfaces
 - configuring 65, 73
- TCP packet
 - tracing 123
- throughput rates data, collecting using SNMP commands 175
- timer policy
 - creating 119
- traffic
 - dropping by default 10, 72
 - rejecting by default 10, 72

- traffic flows
 - controlling with eviction policy *115*
- traps
 - about troubleshooting SNMP *150*
 - defined *147*
 - table of advanced firewall manager-related SNMP *150*
 - table of always-on management-related SNMP *151*
 - table of application security management-related SNMP *153*
 - table of authentication-related SNMP *155*
 - table of AVR-related SNMP *154*
 - table of digital diagnostic monitoring-related SNMP *155*
 - table of DoS-related SNMP *157*
 - table of general SNMP *157*
 - table of global traffic management-related SNMP *158*
 - table of hardware-related SNMP *162*
 - table of high-availability system-related SNMP *166*
 - table of license-related SNMP *167*
 - table of local traffic management-related SNMP *167*
 - table of logging-related SNMP *169*
 - table of network-related SNMP *169*
 - table of vCMP-related SNMP *170*
 - table of VIPRION-related SNMP *170*
- troubleshooting SNMP traps *150*

U

- UDP packet
 - tracing *124*

V

- vCMP-related SNMP traps, defined *170*
- violations statistics
 - viewing *98, 100*
- VIPRION-related SNMP traps, defined *170*
- virtual server
 - adding a NAT match rule *95*
 - adding an eviction policy *117*
 - applying a port misuse policy *121*
 - applying a service policy *121*
 - applying a timer policy *121*
 - assigning eviction policy *117*
 - assigning Network Firewall Logging profile *131, 142, 186*
 - assigning protocol inspection *58*
 - assigning Protocol Security Logging profile *107, 113*
 - controlling flows *117*
 - enabling on a VLAN *66, 74*
 - limiting flows *117*
- virtual server NAT policy *92*
- virtual servers
 - associating firewall NAT policy *95*
 - creating for HTTP traffic *97, 100*
 - creating with a firewall policy *41*
 - viewing compilation statistics *80*
- VLANs
 - creating for network firewall *65, 73*

