F5® BIG-IQ® Centralized Management: Security

Version 5.1

Table of Contents

Overview: BIG-IQ Security 9
Understanding BIG-IQ Network Security and firewall management 9
Understanding Shared Security in BIG-IQ Security 9
Understanding BIG-IQ Web Application Security and application management 10
About the BIG-IQ Security system interface 10
About filtering 10
About browser resolution 10
Setting user preferences 11
About multi-user editing and locking 12
Viewing locks on configuration objects 12
Clearing locks on configuration objects 13
About user roles 13
About BIG-IQ configuration sets 13

Adding BIG-IP Devices to Manage 15
How do I start managing BIG-IP devices from BIG-IQ? 15
Adding devices to the BIG-IQ inventory 15
Importing security service configurations for devices 16
About managing BIG-IP devices 17

Managing Network Security Objects 19
About objects in BIG-IQ Network Security 19
About the policy editor in BIG-IQ Network Security 19
Adding objects 20
Viewing and editing objects 20
Adding objects to firewall contexts and rules 20
Renaming objects 21
Cloning objects 21
Removing objects 21
Filtering content in the policy editor 22
Filtering the policy editor for related objects 24
Filtering the policy editor toolbox frame 24
About address lists 24
Adding address types to address lists 25
Removing entries from address lists 25
Address list properties and addresses 25
About port lists 26
Adding port types to port lists 27
Removing entries from port lists 27
Port list properties and ports 27
About rule schedules 28
Rule schedule properties 28

Managing Firewall Contexts 31
About managing firewall contexts 31
About BIG-IP system firewall contexts 31
About global firewalls 32
About route domain firewalls 32

About virtual server firewalls..... 32
About self IP firewalls..... 33
About management IP firewalls..... 33
About firewall policy types..... 33
Firewall properties..... 33
Adding an enforced firewall policy..... 35
Adding a staged firewall policy..... 35

Managing Rules and Rule Lists..... 37
About rules and rule lists..... 37
Creating rules..... 38
Reordering rules in rule lists..... 38
Removing rules..... 39
Creating and adding rule lists..... 39
Editing rule lists..... 40
Clearing fields in rules..... 41
Cloning rule lists..... 41
Removing rule lists..... 42
Rule properties..... 42

Managing Service, Timer, and Port Misuse Policies..... 47
About service, timer, and port misuse policies..... 47
Create a timer policy..... 47
Create a port misuse policy..... 48
Create a service policy..... 49
Apply a service policy to a firewall rule..... 50
Apply a service policy to a global context..... 50
Apply a service policy to a route domain context..... 51
Apply a service policy to a self IP address context..... 51
Delete a timer policy..... 52
Delete a port misuse policy..... 52
Delete a service policy..... 52

Managing NAT Policies and Translations..... 53
About NAT policies and translations..... 53
Creating a NAT policy..... 53
NAT rule properties..... 53
Cloning a NAT policy..... 55
Deleting a NAT policy..... 55
Creating NAT source translations..... 56
Cloning NAT source translations..... 58
Deleting NAT source translations..... 59
Creating NAT destination translations..... 59
Cloning NAT destination translations..... 60
Deleting NAT destination translations..... 60

Managing FQDN Resolvers..... 61
About FQDN resolvers..... 61
Configuring FQDN resolvers..... 61

Managing Notification Rules..... 63
About notification rules..... 63
Adding and scheduling notification rules..... 63

Editing notification rules....65
Deleting notification rules....65

Managing Change Verifications....67
About change verifications....67
Adding change verifications....67
Viewing change verification properties....68
Change verification properties....68

Managing Firewall Security Locks....71
About firewall security locks....71
Viewing and deleting locks....71

Managing External Logging Devices....73
About external logging devices....73
Add external logging devices....73
Modify external logging devices....74
Remove external logging devices....75
Request authentication token for external logging devices....75
Delete authentication token for external logging devices....76
Access external logging devices....76

Monitoring Firewall Rules....77
About firewall rule monitoring....77
Monitoring firewall rule statistics and hit counts....77
Monitoring firewall rule compilation statistics....78

Managing Firewall Rule Reports....79
About firewall rule reports....79
Creating firewall rule reports....79
Deleting firewall rule reports....80

Managing Firewall Policies in BIG-IQ Network Security....81
About firewall policies in BIG-IQ Network Security....81
Creating firewall policies....81
Managing firewall policies....82
Cloning firewall policies....83
Reordering rules in firewall policies....84
Deleting firewall policies....84
About managing firewall policies using snapshots....85

Managing Security Reports....87
About security reporting....87

Managing Virtual Servers in Shared Security....89
About virtual servers....89
Edit virtual servers....89

Managing DoS Profiles in Shared Security....91
About DoS profiles....91

Creating DoS profiles....91
Configuring for Application Security....92
Configuring for Protocol DNS....95
Configuring for Protocol SIP....96
Configuring for Network Security....97
Editing DoS profiles....97

Managing Device DoS in Shared Security....99
About device DoS....99
Editing device DoS....99

Managing Logging Profiles in Shared Security....101
About logging profiles....101
Create logging profiles....101
Configure for Application Security logging....102
Configure for Protocol Security logging....105
Configure for Network Firewall logging....107
Configure for Network Address Translation logging....109
Configure for DoS Protection logging....110
Editing logging profiles....111
Deleting logging profiles....111

Managing SSH Profiles in Shared Security....113
About SSH profiles....113
Create SSH profiles....113
Configure SSH proxy permissions....113
Configure SSH authentication keys....114
Delete SSH profiles....115

Managing Application Security Policies in BIG-IQ Web Application Security....117
About security policies in BIG-IQ Web Application Security....117
Editing security policies....119
Editing properties settings....119
Editing blocking settings....123
Editing response page settings....134
Editing Data Guard settings....136
Editing Headings and Methods settings....137
Editing IP address settings....138
Adding file types settings....139
Editing parameters settings....140
Editing extractions settings....142
Editing character sets settings....144
Editing attack signatures settings....145
Viewing attack signatures lists....146
Customizing attack signatures lists....147
Adding security policies....148
Importing security policies....149
Exporting security policies....149
Removing security policies....149

Managing Signature Files....151
About signature files in BIG-IQ Web Application Security....151
Viewing signature file properties....151

Signature file properties....151
Updating and pushing signature files....152

Managing Custom Attack Signatures and Signature Sets....153
About custom attack signatures....153
Creating custom attack signatures....153
About signature staging....155
About custom attack signature sets....155
Add custom attack signature sets....155
Edit custom attack signature sets....157
Signatures advanced filter properties....159
Assign custom attack signature sets....162

Managing Virtual Servers in BIG-IQ Web Application Security....163
About virtual servers in the policy editor....163
Displaying virtual server properties and managing policies....163

Managing Event Logs for Web Application Security....165
How do I manage events with a Logging Node?....165
What is a BIG-IQ Logging Node?....166
Discover and activate a logging node....166
Modifying event log indices....167
Define event snapshot storage locations....167
Define Web Application Security snapshot schedules....168
How do I license and do the basic setup to start using a Logging Node?....169
Configuring the BIG-IP logging profile....172
Virtual servers that remote logging uses to route event logs....174
Assigning the logging profile to a virtual server....174
Logging Node management....175
How do I use the event log interface?....177
Viewing event log details....177
Using common filters....178
Filtering the event logs (basic)....178
Filtering (advanced)....179
Filtering by entering query parameters....179
Restore event log snapshots....181
Managing Configuration Snapshots....182
What is snapshot management?....182
Creating a snapshot....182
Comparing snapshots....182
Restoring a snapshot....183

Deploying Changes....185
How do I evaluate changes made to managed objects?....185
Evaluating configuration changes....185
How do I deploy changes made to managed objects?....186
Deploying configuration changes....187
Making an urgent deployment....187
Verifying firewall rules have compiled on all BIG-IP devices....188
Reviewing deployment process states to diagnose problems....189
Device deployment states....189

Managing Configuration Snapshots....191

What is snapshot management?....191
Creating a snapshot....191
Comparing snapshots....191
Restoring a snapshot....192

Managing Audit Logs....193
About audit logs....193
Actions and objects that generate audit log entries....193
Audit log entry properties....194
Viewing audit entry differences....195
Filtering entries in the audit log....195
Customizing the audit log display....197
Managing audit log archive settings....197
About archived audit logs....198
About audit logs in high-availability configurations....198
About the REST API audit log....199
Managing the REST API audit log....199

Legal Notices....201
Legal notices....201

Overview: BIG-IQ Security

Understanding BIG-IQ Network Security and firewall management

BIG-IQ® Network Security is a platform designed for the central management of security firewalls for multiple BIG-IP® systems, where firewall administrators have installed and provisioned the Advanced Firewall Manager™ (AFM™) module.

The BIG-IQ Network Security system provides:

- Device discovery with import of firewalls referenced by discovered devices
- Management of shared objects (address lists, port lists, rule lists, policies, and schedules)
- L3/L4 firewall policy support, including staged and enforced policies
- Firewall audit log used to record every firewall policy change and event
- Role-based access control
- Deployment of configurations from snapshots, and the ability to preview differences between snapshots
- Multi-user editing through a locking mechanism
- Monitoring of rules
- Reports on security

Managing a firewall configuration includes discovering, importing, editing, and deploying changes to the firewall configuration, as well as consolidation of shared firewall objects (policies, rule lists, rules, address lists, port lists, and schedules). BIG-IQ Network Security provides a centralized management platform so you can perform all these tasks from a single location. Rather than log in to each device to manage the security policy locally, it is more expedient to use one interface to manage many devices. Not only does this simplify logistics, but you can maintain a common set of firewall configuration objects and deploy a common set of policies, rule lists, and other shared objects to multiple, similar devices from a central interface.

Bringing a device under central management means that its configuration is stored in the BIG-IQ Network Security database, which is the authoritative source for all firewall configuration entities. This database is also known as the working configuration or working-configuration set.

Once a device is under central management, do not make changes locally (on the BIG-IP device) unless there is an exceptional need. If changes are made locally for any reason, reimport the device to reconcile those changes with the BIG-IQ Network Security working configuration set. Unless local changes are reconciled, the deployment process overwrites any local changes.

In addition, BIG-IQ Network Security is aware of functionality that exists in one BIG-IP system version but not in another. This means, for example, that it prohibits using policies on BIG-IP devices that do not have the software version required to support them.

Understanding Shared Security in BIG-IQ Security

BIG-IQ® Security contains several groups of capabilities. The Shared Security group contains objects that can be used with Network Security objects and with Web Application Security objects.

Understanding BIG-IQ Web Application Security and application management

BIG-IQ® Web Application Security enables enterprise-wide management and configuration of multiple BIG-IP® devices from a central management platform. You can centrally manage BIG-IP devices and security policies, and import policies from files on those devices.

For each device that it discovers, the BIG-IQ system creates a logical container to hold all security policies that are not related to any virtual server on the device. This logical container is called the *inactive virtual server*, and is only used to track policies that are not directly attached to other virtual servers on that device. Policies attached to the inactive virtual server that are distributed are not enforced.

In order for you to deploy a policy to a BIG-IP device, the policy must be attached to one of the device's virtual servers, or to the inactive virtual server. You can deploy policies to a device that already has the policy by overwriting it. If the policy does not yet exist on the device, you can either deploy it as a new policy attached to an available virtual server, or deploy it as a policy attached to the inactive virtual server (which will deploy the policy to the BIG-IP device without attaching it to a virtual server).

From this central management platform, you can perform the following actions:

- Import Application Security Manager™ (ASM) policies from files.
- Import ASM™ policies from discovered devices.
- Distribute policies to BIG-IP devices.
- Export policies, including an option to export policy files in XML format.
- Manage configuration snapshots.
- Edit policy settings. Refer to the table in *About security policies in BIG-IQ Web Application Security* for the supported settings.
- Manage and distribute custom signature sets.
- Manage and distribute custom signatures.
- Distribute signature files to BIG-IP devices.

About the BIG-IQ Security system interface

The BIG-IQ® Security system interface provides many features to assist you in completing tasks.

About filtering

Using filtering, you can rapidly narrow the search scope to more easily locate an entity within the system interface. Each frame in the system interface has its own filter text entry field.

Note: *When you begin typing in the text entry field, you may notice that your browser has cached entries from previous sessions. You can select from the list or continue typing.*

About browser resolution

F5® recommends a minimum screen resolution of 1280 x 1024 to properly display and use the screens efficiently.

It is possible to shrink the browser screen so that system interface elements (screens, scroll bars, icons) no longer appear in the visible screen. Should this occur, use the browser's zoom-out function to shrink the screens and controls.

Setting user preferences

As a firewall policy editor, you can customize the BIG-IQ® Network Security system interface to minimize the information displayed, and to simplify routine editing sessions.

Note: *Setting user preferences is not available through the BIG-IQ Web Application Security system interface.*

For example, you can customize the columns displayed for a particular user in the policy editor.

Note: *This customization does not create an access issue. Users still have access to the resources required by their roles; they just choose not to display them.*

User preference settings persist across sessions. If users log out, they see the same settings when logging back in.

By default, BIG-IQ Network Security replicates user preferences in BIG-IQ high-availability (HA) scenarios.

1. Log in to the BIG-IQ® Security system.
2. At the top-right of the screen, hover over the **admin** icon to display the options.

   The options displayed when you hover over the **admin** vary, depend on whether you have Network Security, Shared Security, or Web Application Security selected.

   - If Network Security is selected, **Global User Settings**, **Security User Settings**, and **Log out** are displayed.
   - If Shared Security is selected, **Global User Settings**, and **Log out** are displayed.
   - If Web Application Security is selected, **Log out** only is displayed and no user settings can be modified.
3. To change Network Security user settings, select **Security User Settings** if it is displayed and select the appropriate options.

Option	Description
Firewall Types	Select or clear the check boxes as required. By default, the interface displays all firewall contexts in the Firewall Contexts screen.
Rule Editor	Select or clear the check boxes as required to modify the policy editor settings. You can set whether to automatically expand rule lists and select which columns to display in the policy editor. By default, the all columns are displayed.

   Changing what columns, or contexts are displayed to the user does not create an access issue. Users still have access to the resources required by their roles; they just choose not to display them.
4. To change global user settings, select **Global User Settings** if it is displayed and select the appropriate options.

Option	Description
Idle Timeout (minutes)	Specify a number indicating how many minutes the BIG-IQ Security user interface can be idle before a user is logged out. The default value is 20.
Default View	Select what part of the BIG-IQ user interface should be initially displayed when a user logs in to the system. The default is **Last Visited** which indicates that the last page used by this user should be displayed when they log in to the system.
5. Click **Save** to save your preferences on either the Global User Settings or Security Settings popup screen. Click **Close** to close the Security Settings popup screen without saving your selections.

Selected preferences are now in effect and persist across user sessions. If you log out, you will see the same settings when you log back in.

About multi-user editing and locking

Within the BIG-IQ® Security system, one or more users may edit firewall security or web application security objects simultaneously. A locking mechanism is used to avoid problems with conflicting changes to objects.

Initially, the user interface displays all objects as read-only. When a user initiates an editing session, the object is locked. Once locked, no one can modify or delete that object except the holder of the lock, or a user with privileges sufficient to break the lock:

- To unlock a locked firewall security object requires the Administrator, Network_Security_Manager, or Security_Manager role.
- To unlock a locked Web application security object requires the Administrator, Web_App_Security_Manager, or Security_Manager role.
- To unlock a locked shared security object, requires the Administrator, Network_Security_Manager, Web_App_Security_Manager, or Security_Manager role.

BIG-IQ Security uses a single repository to hold policy objects and saves each editorial change. With this single-copy design, multiple editors can share the editing task through a locking mechanism.

Each editor has her own copy of a policy (a point-in-time snapshot of the policy managed by BIG-IQ across all devices) and can make changes. When done, an editor can push the changes to the preferred state as one, complete set of changes. Then, an administrator can review a policy change as a single entity before committing it.

For example:

1. If a firewall editor needs to edit `Portlist_1`, `AddressList_2`, and `Rulelist_5`, the editor locks those objects.
2. When the edit pass is complete, the editor saves the object, which clears the lock.

If an editor wants to edit an object that is already locked, the system informs the editor that the object is locked and provides a way to clear the lock if the editor has sufficient privileges. When the lock is cleared, the next firewall editor receives the latest version of the object and any referenced shared objects. Thus, merges and conflicts are avoided. Deleting an object automatically clears all locks associated with it.

BIG-IQ Security supports:

- Multiple, independent locks.
- Locking or unlocking on an object-by-object basis.
- Locks in screens, in the firewall security Policy Editor, and in the Web application security Policy Editor.
- Lock management of firewall security objects using the Locked Objects screen of the firewall security Policy Editor. This screen displays firewall and shared security objects that are locked, the user who locked each object, and when the lock was created. User privileges (assigned by user roles) determine what locks are visible to the user. If you have sufficient privileges, you can use the Locked Objects screen to view and remove multiple firewall and shared security object locks.

Viewing locks on configuration objects

BIG-IQ® Security allows you to view individual locks, and for firewall and shared security objects, allows you to view multiple locks from the Locked Objects screen of the firewall security policy editor.

1. Examine all objects in the BIG-IQ Security screens and policy editors to locate any locked configuration objects.
2. For each locked object, review the lock information on the screen or in the policy editor.

   The displayed lock header displays the owner of the lock and the date and time the lock was created.
3. To view all locked firewall security or shared security objects, use the Locked Objects screen of the firewall security policy editor.

   For each locked object, the Locked Objects screen displays the object name, partition, kind of object, user who locked the object, and when the lock was created.

Clearing locks on configuration objects

The owner of a lock can always clear that lock to enable editing by other users. Other roles (such as Administrator, Network_Security_Manager, Security_Manager, or Web_App_Security_Manager) also carry sufficient privileges to clear locks held by any user. BIG-IQ® Security allows you to clear individual locks, and for firewall and shared security objects, allows you to clear multiple locks from the Locked Objects screen of the firewall security policy editor.

1. Examine all objects in the BIG-IQ Security screens and policy editors to locate any locked configuration objects.
2. For each locked object, review the lock information on the screen or in the policy editor.

   The displayed lock header displays the owner of the lock and the date and time the lock was created. If your role carries sufficient privileges, you will also see a link labeled **Unlock**.
3. In the lock header, click **Unlock**.
4. To clear one or more locked firewall security or shared security objects from a single screen, select the one or more locked objects from the Locked Objects screen of the firewall security policy editor and click **Unlock**.

The lock is cleared; if multiple locks were selected, the locks are cleared.

About user roles

As a security system manager, you need to differentiate between types of users, and to limit user privileges based on user responsibilities. To assist you, the BIG-IQ® system provides a default set of roles. You can associate multiple roles with a given user; for example, you can grant a user the edit (Network_Security_Edit) and the deploy (Network_Security_Deploy) roles for network security functions. Roles persist and are available after a BIG-IQ system failover.

To view the defined roles, both default and locally-defined, log in to BIG-IQ System as administrator, and navigate to the Roles screen.

Select **System Management** from the BIG-IQ menu and then click **USER MANAGEMENT** > **Roles**.

The Roles screen lists each defined role and a description of that role. Refer to the Roles online help or to the *BIG-IQ® Centralized Management: Licensing and Initial Setup* guide for more information on roles and their use.

About BIG-IQ configuration sets

BIG-IQ® system security uses the following terminology to refer to configuration sets for a centrally-managed BIG-IP® device:

Current configuration set

The configuration of the BIG-IP® device as discovered by BIG-IQ. The current configuration is updated during a reimport or rediscovery and before calculating differences during the deployment process.

Working configuration set

The configuration as maintained by the BIG-IQ system. The working configuration is the configuration that is edited on the BIG-IQ system and deployed back to BIG-IP devices. The working configuration for the device is the same as the current configuration when the device is initially managed and when the device is reimported or rediscovered.

The working configuration is created when the administrator first manages the BIG-IP device from the BIG-IQ system. The working configuration is updated when a device is reimported or rediscovered.

If conflicts are observed during reimport or rediscovery, the object in conflict is only updated in the working configuration when the **Use BIG-IP** resolution conflict option is used.

Adding BIG-IP Devices to Manage

How do I start managing BIG-IP devices from BIG-IQ?

To start managing a BIG-IP® device, you must add it to the BIG-IP Devices inventory list on the BIG-IQ® system.

Adding a device to the BIG-IP Devices inventory is a two-stage process.

Stage 1:

- You enter the IP address and credentials of the BIG-IP device you're adding, and associate it with a cluster (if applicable).
- BIG-IQ opens communication (establishes trust) with the BIG-IP device.
- BIG-IQ discovers the current configuration for any selected services you specified are licensed on the BIG-IP system, like LTM® (optional).

Stage 2:

- BIG-IQ imports the licensed services configuration you selected in stage 1 (optional).

Note: *If you only want to do basic management tasks (like software upgrades, license management, and UCS backups) for a BIG-IP device, you do not have to discover and import service configurations.*

Adding devices to the BIG-IQ inventory

Before you can add BIG-IP® devices to the BIG-IQ® inventory:

- The BIG-IP device must be located in your network.
- The BIG-IP device must be running a compatible software version. Refer to *https://support.f5.com/kb/en-us/solutions/public/14000/500/sol14592.html* for more information.
- Port 22 and 443 must be open to the BIG-IQ management address, or any alternative IP address used to add the BIG-IP device to the BIG-IQ inventory. These ports and the management IP address are open by default on BIG-IQ.

If you are running BIG-IP version 11.5.1 up to version 11.6.0, you might need root user credentials to successfully discover and add the device to the BIG-IP devices inventory. Root user credentials are not required for BIG-IP devices running 11.5.0 - 11.5.1 and 11.6.0 - 12.x.

Note: *A BIG-IP device running versions 10.2.0 - 11.4.1 is considered a legacy device and cannot be discovered from BIG-IQ version 5.0. If you were managing a legacy device in previous version of BIG-IQ and upgraded to version 5.0, the legacy device displays as impaired with a yellow triangle next to it in the BIG-IP Devices inventory. To manage it, you must upgrade it to 11.5.0 or later. For instructions, refer to the section titled, Upgrading a Legacy Device.*

You add BIG-IP devices to the BIG-IQ system inventory as the first step to managing them.

Note: *The ADC component is automatically included (first) any time you discover or import services for a device.*

---

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.

3. At the top of the screen, click **Inventory**.
4. Click the **Add Device** button.
5. In the **IP Address** field, type the IPv4 or IPv6 address of the device.
6. In the **User Name** and **Password** fields, type the user name and password for the device.
7. If this device is part of a DSC group, from the **Cluster Display Name** list, select one of the following:
   - For an existing DSC group, select **Use Existing** from the list and select the DSC group from the list.
   - For a new DSC group, select **Create New** from the list and type a name in the field.

   For BIG-IQ to properly associate devices in the same DSC group, the **Cluster Display Name** must be the same for each member in a group.
8. If this device is configured in a DSC group, select an option:
   - **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option if this device is part of a DSC group and you want this device to automatically synchronize configuration changes with other members in the DSC group.
   - **Ignore BIG-IP DSC sync when deploying configuration changes** Select this option if you want to manually synchronize configurations changes between members in the DSC group.
9. Click the **Add** button at the bottom of the screen.
   The BIG-IQ system opens communication to the BIG-IP device, and checks its framework.

   Note: *The BIG-IQ system can properly manage a BIG-IP device only if the BIG-IP device is running a compatible version of the REST framework.*
10. If a framework upgrade is required, in the popup window, in the **Root User Name** and **Root Password** fields, type the root user name and password for the BIG-IP device, and click **Continue**.
11. If in addition to basic management tasks (like software upgrades, license management, and UCS backups) you also want to centrally manage this device's configurations for licensed services, select the check box next to each service you want to discover.

    You can also select these service configuration after you add the BIG-IP device to the inventory.
12. Click the **Add** button at the bottom of the screen.

BIG-IQ displays a discovering message in the Services column of the inventory list.

If you discovered service configurations to manage, you must import them.

Importing security service configurations for devices

Before you can import the security properties defined on a BIG-IP device, the BIG-IQ must discover that device.

Once you import the properties for security configuration objects (virtual servers, firewall policies, signature files, and so on) defined on a BIG-IP device, you can use the BIG-IQ to manage these objects.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. For the device you want to manage, click the link in the Services column.
   The services currently managed for this device are listed. The link text in the Services column varies depend on what services are imported.
5. For the device you want to manage, under Services, click the **Complete import tasks** link.
   The services currently managed for this device are listed.
6. Click **Import** in the Configuration Import row for the security service you want to import.

- Use the Configuration Import row under Advanced Firewall (AFM) to import and manage network firewall security objects.
- Use the Configuration Import row under Application Security (ASM) to import and manage web application security objects.

The BIG-IQ imports the settings for the selected objects defined on the BIG-IP. If the current configuration on the BIG-IQ system is different than the one on the BIG-IP device, BIG-IQ displays a screen for you to resolve the conflicts.

7. If there are conflicts, select one of the following options for each object that is different, and then click the **Continue** button:
   - **Use BIG-IQ** to use the configuration settings stored on BIG-IQ.
   - **Use BIG-IP** to override the configuration setting stored on BIG-IQ with the settings from the BIG-IP device.
8. Click **Close**.

Now you can use this BIG-IQ to manage the security settings on this BIG-IP.

About managing BIG-IP devices

Once you have placed a BIG-IP® device under management by the BIG-IQ® system by discovering and importing that device configuration, you should avoid directly changing the BIG-IP device configuration. All changes to the BIG-IP device configuration should be made using the BIG-IQ system to avoid errors.

During the deployment process, the BIG-IQ system imports the current configuration of the targeted BIG-IP devices. Subsequent changes made directly on the BIG-IP device which add new objects to the configuration will be labeled as being not imported and those objects will not be removed during the next deployment. These objects will continue to be labeled as not imported, until you reimport the configuration using the Device Management BIG-IP Devices screen.

To avoid this situation, when you directly modify a BIG-IP device, you must re-discover and re-import the BIG-IP device from the BIG-IQ system to reconcile the configuration differences.

Managing Network Security Objects

About objects in BIG-IQ Network Security

In BIG-IQ® Network Security, the objects that you can view and manage from the policy editor include:

Contexts (firewall)
Category of object to which a rule applies. In this case, category refers to Global, Route Domain, Self IP, Virtual Server, or Management. Within each context, rules can be viewed and reorganized separately. It is possible to have multiple layers of firewalls on a single BIG-IP® device. These layers constitute the firewall hierarchy. Within the firewall hierarchy, rules progress from Global, to Route Domain, and then to either Virtual Server or Self IP.

Firewall Policies
Set of rules and/or rule lists that specify traffic-handling actions and define the parameters for filtering network traffic. You can assign rule lists or a policy to a firewall. Firewall policies facilitate the assigning of a common collection of rules consistently across multiple firewalls.

Rule lists
Containers for rules; rules are run in the order they appear in their assigned rule list. A rule list can contain thousands of ordered rules, but cannot be nested inside another rule list.

Address lists
Collections of IPv4 or IPv6 addresses, address ranges, nested address lists, geolocations and subnets. These collections are saved on a server and used by policies, rule lists, and rules to allow or deny access to specific IP addresses in IP packets. Firewall rules compare all addresses or address ranges in a given address list to either the source or the destination IP address, depending on how the list is applied. Firewall rules can also compare all geolocations in a given address list to either the source or the destination location, depending on how the list is applied. If there is a match, the rule takes an action, such as accepting or dropping the packet.

Port lists
Collections of ports and port ranges. These collections are saved on a server and used by policies, rule lists, and rules to allow or deny access to specific IP addresses in IP packets. As with address lists, firewall rules compare all ports and port ranges in a given port list to either the source or the destination port, depending on how the list is applied. If there is a match, the rule takes an action, such as accepting or dropping the packet.

Rule Schedules
Rule schedules are assigned to firewall rules, rule lists, and policies to control when rules, rule lists, and firewall policies are active on the firewall. In the Policy Editor, you can hover over schedule names to see the name displayed in a tooltip. This feature is useful if the schedule name is longer than the screen.

About the policy editor in BIG-IQ Network Security

You use the BIG-IQ® Network Security policy editor to rapidly make firewall configuration changes within firewall policies by editing the objects that contain that information. The policy editor provides users with a toolbox that can be used to quickly add objects. The toolbox is located at the bottom of the policy editor.

Adding objects

You add firewall policy objects using the policy editor.

Note: *Address lists and port lists are containers and must contain at least one entry. You cannot create an empty list; you cannot remove an entry from a list if it is the only entry.*

1. Click the type of object you want to add in the navigation list to the left, then click **Create** at the top of the object pane on the right.
2. In the opened screen, populate the property fields as required.
   - All fields that are outlined in gold are required.
   - The **Partition** field is outlined in gold, and although it is pre-populated with `Common`, it is an editable field.
   - You can press Tab to advance from field to field.
3. When you are finished, click **Save** to save your changes, or click **Save & Close** to save and close the current window.

Viewing and editing objects

You use the BIG-IQ® Network Security policy editor to select firewall policy objects for deeper inspection or edit.

Note: *Address lists and port lists are containers, and must contain at least one entry. You cannot create an empty list; you cannot remove an entry in a list if it is the only entry.*

1. Navigate to the object you want to edit.
2. Click the name of the object that you want to edit.
3. Edit the properties and other areas as required.

   You can use the keyboard Tab to advance from field to field.
4. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Adding objects to firewall contexts and rules

BIG-IQ® Network Security enables you to add objects to firewall contexts and rules (used in rule lists and firewall policies).

1. Navigate to the context, rule list or firewall policy to which you want to add an object.
   - If you are editing a firewall context, click the name of the context to add to, click the name of the enforced firewall policy to add to on the Properties tab, and select **Rule & Rule Lists** so that the rules for enforced policies are visible.
   - If you are editing a rule list, click the name of the rule list to add to, and select **Rules** so that the rules are visible.
   - If you are editing a firewall policy, click the name of the policy to add to, and select **Rule & Rule Lists** so that the rules are visible.
2. Click on an existing rule, click **Create Rule** to create a new rule, or click **Add Rule List** to add a rule list as needed.
3. In the toolbox at the bottom, select the type of object you want to add from the drop down list. The objects matching that type are listed.
4. Drag the object you want onto the indicated area.

5. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Renaming objects

BIG-IQ® Network Security does not support renaming an object.

As an alternative to renaming an object, you can create a new object and replace the original object where it is in use.

1. Create the new object. Consider cloning the object as the fastest and most reliable way to create a new object with the same content as the original but with a new name.
2. Locate every instance of the original object by hovering over the object, right-clicking, and selecting **Filter 'related to'**.
   A count is added to the objects in the navigation list on the left, indicating the number of times the object is used in each object.
3. Navigate to each instance where the original object is in use, and replace it with a reference to the newly-created object.
4. Remove the original object.

   Clear the filter by clicking the **X** to the right of the filter text in the field at the top of the navigation list under the filter entry box.

Note: *You cannot remove an object that is still in use.*

Cloning objects

BIG-IQ® Network Security enables you to clone objects to create a copy of that object that is slightly different from the original. You may have an object that serves as a template. You can clone that object, edit it, and then use it in different ways.

1. Navigate to the type of object you want to clone.
2. Click the checkbox to the left of object that you want to clone.
3. Click **Clone**.
   The system displays a copy of the object with the original object's name with -CLONE appended to the name and a blank **Description** field.
4. In the opened screen, populate the property fields as required.
   - All fields that are outlined in gold are required.
   - The **Partition** field is outlined in gold, and although it is pre-populated with `Common`, it is an editable field.
   - You can press Tab to advance from field to field.
5. When you are finished, click **Save** to save your changes, or click **Save & Close** to save and close the current window.

The cloned object is added to the existing list in the appropriate section.

Removing objects

From the BIG-IQ® Network Security policy editor, you can remove shared objects.

1. Navigate to the type of object you want to remove.
2. Click the checkbox to the left of object that you want to remove and click **Delete**.
   A popup information screen opens.
3. Respond to the popup screen prompt:

- If the object is being used by another object, policy, rule, or rule list, you cannot remove it; click **Cancel** to not perform the removal.
- If the object can be removed, click **Delete** to confirm the removal.

Filtering content in the policy editor

There are several filter fields you can use to select the data displayed by the Policy Editor. The filter text you enter is used to perform a search of the underlying object's representation in storage (in JSON), which includes not only the name and other displayed data, but also metadata for the object, such as timestamps. Make the text you enter in the filter field specific enough to uniquely identify the one or more objects you want to display.

1. To filter the contents of the Policy Editor frame, log in to BIG-IQ® Security.
2. Navigate to **Network Security** > **Policy Editor**.
3. In the appropriate filter text field, type the text you want to filter on and press **Enter**.

Option	Description
Filter field above navigation list on left	Use the filter field above the navigation list on the left to search objects and list those that match the filter. By default, the filter matches any object that contains the string entered. You select filter options by clicking the arrow to the left of the filter field, and selecting an option. • **Contains** indicates that the filter text matches any object that contains it. This is the default. When searching for times or dates, such as those in a schedule, a partial time, such as September, may be specified. • **Exact** indicates that the filter text matches any object that exactly matches it. When searching for times or dates such as those in a schedule, the complete time and date must be specified. A count of the matching objects appears to the right of each object type in the navigation list. To remove the filter, click the **X** to the right of the filter expression area near the filter field.
Filter field at top right of Policy Editor	Use the filter field at the right top of the Policy Editor to search only the displayed objects for a match to the filter. You select filter options by clicking the arrow to the left of the filter field, and then selecting an option from each option group. The bottom option group in the list controls whether the filter text must be a partial match or an exact match. • **Contains** indicates that the filter text matches any object that contains it. This is the default. When searching for times or dates, such as those in a schedule, a partial time, such as September, may be specified. • **Exact** indicates that the filter text matches any object that exactly matches it. When searching for times or dates, such as those in a schedule, the complete time and date must be specified. The top options group in the list control which objects are filtered. Not all options are displayed on all screens; if none of these options are displayed (**IP Address**, **Name** or **Port**), the default is **All**. • **All** indicates that all objects should be filtered using the filter text. • **IP Address** indicates that only IP address objects should be filtered using the filter text. A complete IPV4 or IPV6 address must be entered as the filter text. • When used with the **Contains** option, the filter text is matched by an IPV4 or IPV6 address that is the same as the filter text, or an IPV4 address range

Option	Description
	or subnet that includes the filter text. IPV6 addresses can not be found within a range or subnet. • When used with the **Exact** option, the filter text is matched by an IPV4 or IPV6 address that is the same as the filter text only. • **Name** indicates that only object names should be filtered using the filter text. • **Port** indicates that only port objects should be filtered using the filter text. A complete port number must be entered as the filter text. • When used with the **Contains** option, the filter text is matched by a port number that is the same as the filter text, or a port number range that includes the filter text. • When used with the **Exact** option, the filter text is matched by a port number that is the same as the filter text only. If the navigation list is displayed, a count of the matching objects appears to the right of each object type in the navigation list. To remove the filter, click the **X** to the right of the filter expression area near the filter field.
Filter field in Policy Editor Toolbox at bottom	Use the filter field in the upper right of the Policy Editor toolbox (displayed at the bottom of the page when active) to search the shared resources list in the toolbox and display only those that have a full or partial match to the filter. To remove the filter, click the **X** to the right of the filter expression area near the filter field.

When specifying a date in a filter, only these date and time formats are supported:

- Sep 1, 2015 2:05:04 PM
- Sep 1, 2015 2:05:04 AM
- Sep 1, 2015 14:05:04
- Sep 1, 2015 2:05
- Sep 1, 2015
- Sep 1 2015
- Sep 1
- September 1
- 2015-09-01T14:05:04
- 2015-09-01T14:05
- 2015-09-01 2015-09
- 2015

You can also use the **Filter 'related to'** option to display objects that are related to that object. Right-click on an object in the initially displayed list of objects and select **Filter 'related to'** to display the objects related to that object. The results of the search only include devices for searches of contexts and firewall policies, all other objects do not include devices in the results. The **Filter 'related to'** option is not available in all policy editor screens.

You clear filter fields by clicking the **X** to the right of the filter field.

Objects are filtered on the text entered and a count for each appears to the right of each object type.

Note: *Filter matches are only displayed for an object and its containing object. For example, when a filter matches a rule name in a rule list within a policy, only the rule and rule list will be shown as matching, but the policy will not.*

Filtering the policy editor for related objects

You can filter contents owithin the Policy Editor frame to show objects related to a selected object.

1. To filter for related objects within the Policy Editor frame, log in to BIG-IQ® Network Security.
2. Navigate to **Network Security** > **Policy Editor**.
3. Locate the object you want to filter on. in either the left frame or in the toolbox at the bottom of the right frame.
4. Right-click the object.
5. Click **Filter 'related to'**.

   You can clear the **Related to** filter by clicking the **X** to the right of the text near the filter field. This option is not available for all objects.

All object types in the left frame are filtered and a count of each **related to** object found appears to the right of each object type.

Filtering the policy editor toolbox frame

You can filter the contents of screens within the policy eitor, such as the toolbox frame, to reduce the amount of data displayed. Filtering techniques can be important for troubleshooting. There are several filter fields you can use within the policy editor. The filter text you enter is used to perform a wildcard search of the underlying object's representation in storage (in JSON), which includes not only the name and other displayed data, but also metadata for the object, such as timestamps. Make the text you enter in the filter field specific enough to uniquely identify the one or more objects you want to display.

1. To filter the contents within the Policy Editor toolbox, log in to BIG-IQ® Security.
2. Navigate to **Network Security** > **Policy Editor**.
3. Use the Filter field in the upper right of the Policy Editor toolbox (the toolbox is displayed at the bottom of the page when active) to search the shared resources list in the toolbox and display only those objects that match the filter.

   You clear the active filter and make all data viewable by clicking the **X** to the right of the filter field.

About address lists

Address lists are collections of IPv4 or IPv6 addresses, address ranges, nested address lists, geolocations and subnets saved on a server and available for use in firewall rules, rule lists and firewall policies

Firewall rules refer to address lists to allow or deny access to specific IP addresses in IP packets. Firewall rules compare all addresses from the list to either the source or the destination IP address (in IP packets), depending on how the list is applied. Firewall rules can also compare all geolocations in a given address list to either the source or the destination location, depending on how the list is applied. If there is a match, the rule takes an action, such as accepting or dropping the packet.

You can see the content of an address list by hovering over its name in the policy editor. If an address list is nested, the tooltip displayed by the hovering will only show the first-level contents. To view address list names that are longer than the display field, hover over the name to see the full name displayed in the tooltip.

Note: *Before nesting an address list inside an address list, check to be sure this option is supported on each BIG-IP® device where you intend to deploy the address list.*

Address lists are containers and must contain at least one entry. You cannot create an empty address list; you cannot remove an entry in an address list if it is the only one.

You can add geolocation awareness to address lists, which enables you to specify source or destination IP addresses by geographic location. Thus, you can specify firewall behavior for traffic to/from entire geographic regions by defining rules based on where the source or destination system is, rather than on its IP address (source or destination). BIG-IQ® Network Security supports specifying geolocation in rules and address lists. The geolocation is validated when the rule or address list is saved.

Note: *If you use a geolocation spec that is valid on the BIG-IQ Network Security system, but not supported on a particular BIG-IP® device because the device has a different geolocation database, it causes a deployment failure for that device. Importing a BIG-IP device with an invalid geolocation spec causes a discovery failure for that device.*

Adding address types to address lists

BIG-IQ® Network Security enables you to add addresses, address ranges, nested address lists, or geolocation to an existing address list.

1. Navigate to the Address Lists area. **Policy Editor** > **Address Lists**
2. Click the name of the address list that you want to edit.
3. Click the Addresses tab and then click the + icon to the right of an address.
   A new row is added to the list of addresses under that row.
4. From the list under the Type column, select **Address**, **Address Range**, **Address List**, **Domain Name** or **Country/Region**.
   - If you select **Address List**, in the **Addresses** field type the first letter of an existing address list. A list of existing address lists appears from which you can select an address.
   - If you select **Address**, in the **Addresses** field supply the address.
   - If you select **Address Range**, in the **Addresses** field supply the beginning range address in the top field and the ending range address in the bottom field.
   - If you select **Domain Name**, in the **Addresses** field supply the domain name.
   - If you select **Country/Region**, in the **Addresses** field select a country from the top list and optionally a region in that country from the bottom list.
5. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Removing entries from address lists

BIG-IQ® Network Security allows you to remove entries from address lists.

1. Navigate to the Address Lists area. **Policy Editor** > **Address Lists**
2. Click the name of the address list that you want to edit.
3. Click the **X** icon to the right of the address list entry to remove.

   The entry is highlighted in red and marked to be deleted. The entry will be deleted when you click **Save** or **Save & Close**.
4. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Address list properties and addresses

Property	Description
Name	Unique, user-provided name for the address list. The text field accepts up to and including 255 characters, including the partition name.
Description	Optional description of the address list.

Property	Description
Partition	Field pre-populated with `Common` (the default). This field is editable when creating or cloning address lists.
Type	After locking the address list for editing, select one of the following: • **Address**. Then, type the address in the **Addresses** field. You can also enter an address range in this field by typing a range in the format: `n.n.n.n-n.n.n.n`. • **Address Range**. The **Addresses** field becomes two fields separated by "to." Type the beginning address and ending addresses in these fields as appropriate. • **Address List**. When you type the first letter of a saved list, the **Addresses** field populates with a picker list that displays saved address lists. You then select from the list. • **Country/Region**. From the first **Addresses** list, select a country. Once you select a country, the second list automatically updates with all available regions for that country. Optionally, select a region from the second list. The wildcard, Unknown, is supported. Note that geolocation is not supported on the management IP context.
Addresses	IPv4 or IPv6 address, address range, or nested address list. There are many ways an IPv4 or IPv6 address or address range can be constructed. The following methods and examples are not meant to be exhaustive. • IPv4 format: *a.b.c.d*[/*prefix*]. For example: `60.63.10.10`. • IPv6 format: *a*:*b*:*c*:*d*:*e*:*f*:*g*:*h*[/*prefix*]. For example: `2001:db7:3f4a:9dd:ca90:ff00:42:8329`. • IPv6 abbreviated form is supported. You can shorten IPv6 addresses as defined in RFC 4291. • You can specify subnets using forward slash (/) notation; for example: `60.63.10.0/24`. Example IPv6 subnet: `2001:db8:a::/64`. • You can append a route domain to an address using the format %RouteDomainID/Mask. For example: `12.2.0.0%44/16`.
Description	Optional text field used to describe the address, address range, or nested address list.

About port lists

Port lists are collections of ports, port ranges, or port lists or nested port lists saved on a server and available for use in firewall rules, rule lists, and firewall policies.

Firewall rules refer to port lists to allow or deny access to specific ports in IP packets. They compare a packet's source port and/or destination port with the ports in a port list. If there is a match, the rule takes an action, such as accepting or dropping the packet. Port lists are containers and must contain at least one entry. You cannot create an empty port list; you cannot remove an entry in a port list if it is the only one.

You can see the content of a port list by hovering over its name in the policy editor. If a port list is nested, the tooltip displayed will only show the first-level contents. To view port list names that are longer than the display field, hover over the name to see the full name displayed in the tooltip.

Note: *Before nesting a port list inside a port list, check to be sure this option is supported on the BIG-IP® device where you intend to deploy the port list.*

Adding port types to port lists

BIG-IQ® Network Security enables you to add ports, port ranges, or nested port lists to an existing port list.

1. Navigate to the Port Lists area. **Policy Editor** > **Port Lists**
2. Click the name of the port list that you want to edit.
3. Click the + icon to the right of a port.
   A new row is added to the Ports table under that row.
4. From the **Type** list, select **Port**, **Port Range**, or **Port List**.
   If you select **Port List**, and type the first letter of an existing port list in the **Ports** field, a list of existing port lists appears from which you can select a port list from the list.
5. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Removing entries from port lists

BIG-IQ® Network Security enables you to remove entries from port lists.

1. Click the name of the port list from which you want to remove an entry. **Policy Editor** > **Port Lists**
2. Click the name of the port list from which to remove the entry.
3. Click the **X** icon to the right of the port list entry to remove.

   The entry is highlighted in red and marked to be deleted. The entry will be deleted when you click **Save** or **Save & Close**.
4. When you are finished, click **Save** to save your edits, or click **Save & Close** to save and release the lock.

Port list properties and ports

Property	Description
Name	Unique name used to identify the port list.
Description	Optional description for the port list.
Partition	Field pre-populated with `Common` (the default). This field is editable when creating or cloning port lists.
Type	Select one of the following: • **Port**. Then, enter the port in the **Ports** field. You can also enter a port range in this field by entering a range in the format: n-n. Valid port numbers are 1-65535. • **Port range**. The **Ports** field becomes two fields separated by "to." Type the beginning port and ending port in these fields as appropriate. • **Port list**. When you type the first letter of a saved list, the **Ports** field is populated with a picker list that displays saved port lists. You then select from the list.
Ports	Port, port range, or port list. Valid port numbers are 1-65535.
Description	Optional text field used to describe the port, port range, or nested port list.

About rule schedules

The Rule Schedules screen displays the defined rule schedules. By default, all rules, rule lists, and policies run continuously. Rule schedules are *continuously active* if created without any scheduling specifics (such as the hour that the rule schedule starts).

You apply a rule schedule to a rule to make that rule active only when needed.

Rule schedule properties

Property	Description
Name	Specifies a unique, user-provided name for the rule schedule.
Description	Specifies an optional description for the rule schedule.
Partition	Displays informational, read-only name of the partition associated with the rule schedule.
Date Range	Specifies the date and time when the rule can be active. Select one of the following: **Indefinite** Specifies that the rule schedule start immediately and run indefinitely. The rule schedule remains active until you change the date range or delete the rule schedule. This is the default. **Until...** Specifies that the rule schedule start immediately and run until a specified end date. The rule schedule is immediately activated and not disabled until the end date and time is reached. Click in the field to choose an end date from a popup calendar. You can specify an end time in the same popup screen. **After...** Specifies that the rule schedule start after the specified date and run indefinitely. The rule schedule is activated starting on the selected date and runs until you change the start date or delete the rule schedule. Click in the field to choose a start date from a popup calendar. You can specify a start time in the same popup. **Between...** Specifies that the rule schedule start on the specified date and run until the specified end date. Click in the fields to choose the start and end dates from a popup calendar. You can specify start and end times in the same popup. ***Note:*** *Using the system interface and popup screens to specify the start and end dates and times is the preferred method. However, if you do specify dates manually, use the format: MMM DD,YYYY HH:MM:SS.*
Time Span	Specifies the time, within the time defined by the Date Range, that the rule schedule can be active. • **All Day** specifies that the rule schedule runs all day. This is the default. • **Between...** specifies the time using the format: HH:MM. You specify this time by typing in the fields.

Property	Description
Day	Specifies the days the rule schedule is active. Select check boxes for all days that apply. You must select at least one day per week.

Managing Firewall Contexts

About managing firewall contexts

In BIG-IQ® Network Security, a firewall context is a BIG-IP® network object to which a firewall policy can be attached. In BIG-IQ Network Security, these network objects are called Global (global), Route Domain (rd), Virtual Server (vip), Self IP (sip), or Management (mgmt).

Firewall contexts provide policy-based access control to and from address and port pairs, inside and outside the network. Using a combination of contexts, a firewall can apply rules in a number of different ways, including at a global level, per virtual server, per route domain, and even for the management port or a self IP address.

Firewall properties include the firewall name, an (optional) description, its partition, its type, and its parent device on the partition in which it resides. Note that an *administrative partition* is a part of the BIG-IP configuration that is accessible only to a particular group of administrators. The default partition for all BIG-IP configurations, `/Common`, is accessible to all administrators. A sufficiently-privileged administrator can make additional partitions on the BIG-IP device. Each partition corresponds to a folder (with the same name) to hold its configuration objects.

You can use the Policy Editor to view and configure enforced policies or rules whose actions (accept, accept decisively, drop, reject) are in force. You are restricted to a single, enforced policy on any specific firewall. You can edit all other firewall shared objects only from within the object's screen.

Note: *Firewall policies can be enforced in one firewall context and staged in another.*

Considerations when restoring snapshots of BIG-IP devices containing firewall inline rules

If you restore a snapshot of a version 11.5.1 or earlier BIG-IP device that contains inline firewall rules onto a BIG-IP version 11.5.2 or later or BIG-IP version 11.6 or later device, the inline rules are improperly restored to the later version. The inline rules are improperly restored because these later BIG-IP device versions do not support the inline firewall rules that were part of the version 11.5.1 or earlier BIG-IP device snapshot.

When you upgrade a version 11.5.1 or earlier BIG-IP device, the BIG-IP device automatically moves any inline rules into a system-defined policy. The restoration of the version 11.5.1 or earlier snapshot incorrectly writes inline rules back to the configuration of the later version of the BIG-IP device.

To restore a snapshot of a version 11.5.1or earlier BIG-IP device onto a later version BIG-IP device, you must again reimport the upgraded devices after restoring the snapshot. This updates the BIG-IQ system to contain the current policy based firewall configurations and removes the inline rules that were added to the configuration by the restoration of the snapshot for those 11.5.2 or later or 11.6.0 or later devices.

About BIG-IP system firewall contexts

A *firewall context* is the category of object to which a rule applies. In this case, category refers to Global, Route Domain, Virtual Server, Self IP, or Management. Rules can be viewed and reorganized separately within each context.

It is possible to have multiple layers of firewalls on a single BIG-IP® device. These layers constitute the firewall hierarchy. Within the firewall hierarchy, rules progress from Global, to Route Domain, and then to either Virtual Server or Self IP.

If a packet matches a firewall rule within a given context, that action is applied to the packet, and the packet then moves to the next context for further processing. If the packet is accepted, it travels on to the next context. If the packet is accepted decisively, it goes directly to its destination. If the packet is dropped or rejected, all processing stops for that packet; it travels no further.

On each firewall, you can have rules, rule lists, or policies that are enforced or staged. Rules, rule lists, or policies are processed in order within their context and within the context hierarchy.

Rules for the Management interface are processed separately and not as part of the context hierarchy.

About global firewalls

A *global firewall* is an IP packet filter that resides on a global firewall on a BIG-IP® device. Except for packets traveling to the management firewall, it is the first firewall that an IP packet encounters. Any packet reaching a BIG-IP device must pass through the global firewall first.

When you create firewall rules or policies, you can select one of several contexts. Global is one of the contexts you can select. Rules for each context form their own list, and are processed both in the context hierarchy and in the order within each context list.

About route domain firewalls

A *route domain firewall* is an IP packet filter that resides on a route domain firewall on a BIG-IP® device.

A *route domain* is a BIG-IP system object that represents a particular network configuration. After creating a route domain, you can associate various BIG-IP system objects with the domain: unique VLANs, routing table entries such as a default gateway and static routes, self IP addresses, virtual servers, pool members, and firewalls.

When a route domain firewall is configured to apply to one route domain, it means that any IP packet that passes through the route domain is assessed and possibly filtered out by the configured firewall.

When you create firewall rules or policies, you can select one of several contexts. Route domain is one of the contexts you can select. Rules for each context form their own list and are processed both in the context hierarchy and in the order within each context list.

Route domain rules apply to a specific route domain configured on the server. Route domain rules are checked after global rules. Even 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.

Route domain rules are collected in the Route Domain context. Route domain rules apply to a specific route domain defined on the server. Route domain rules are checked after global rules.

About virtual server firewalls

A *virtual server firewall* is an IP packet filter configured on the virtual server and, therefore, designated for client-side traffic. Any IP packet that passes through the virtual server IP address is assessed and possibly filtered out by this firewall.

When you create firewall rules or policies, you can select one of several contexts, including virtual server. Rules for each context form their own list and are processed both in the context hierarchy and in the order within each context list.

Virtual server rules apply to the selected virtual server only. Virtual server rules are checked after route domain rules.

About self IP firewalls

A *self IP firewall* is an IP packet filter configured on the self IP address, a firewall designated for server-side traffic. Any IP packet that passes through the self IP is assessed and possibly filtered out by this firewall.

A self IP address is an IP address on a BIG-IP® system that is associated with a VLAN and used to access hosts in that VLAN. By virtue of its netmask, a self IP address represents an address space; that is, a range of IP addresses spanning the hosts in the VLAN, rather than a single host address.

A static self IP address is an IP address that is assigned to the system and does not migrate between BIG-IP systems. By default, the self IP addresses created with the Configuration utility are static self IP addresses. One self IP address must be defined for each VLAN.

When you create firewall rules or policies, you can select one of several contexts, including self IP. Rules for each context form their own list and are processed both in the context hierarchy and in the order within each context list.

The self IP context collects firewall rules that apply to the self IP address on the BIG-IP device. Self IP rules are checked after route domain rules.

About management IP firewalls

A *management IP firewall* is an IP packet filter configured on the management IP address and, therefore, designated to examine management traffic. Any IP packet that passes through the management IP address is assessed and possibly filtered out by this firewall.

The network software compares IP packets to the criteria specified in management firewall rules. If a packet matches the criteria, then the system takes the action specified by the rule. If a packet does not match a rule, then the software compares the packet against the next rule. If a packet does not match any rule, the packet is accepted.

Management IP firewalls collect firewall rules that apply to the management port on the BIG-IP® device. Management port firewalls are outside the firewall context hierarchy and management port rules are checked independently of other rules.

Note: *Policies and rule lists are not permitted on management IP firewalls. In addition, the management IP firewall context does not support the use of iRules® or geolocation in rules.*

About firewall policy types

In BIG-IQ® Network Security, you can add the following firewall policy types:

Enforced
An enforced firewall policy modifies network traffic based on a set of firewall rules.

Staged
A staged firewall policy allows you to evaluate the effect a policy has on traffic without actually modifying the traffic based on the firewall rules.

Firewall properties

The properties of a firewall context are shown when you select a context type from the list on the left, such as Global or Virtual Server. Some fields are for information purposes only and cannot be edited. Not all columns are shown for each context.

Property	Description
Name	Name as shown in the system interface: `global` for the global firewall; `management-ip` for the management IP firewall; `0` for route domain; the IP address for self-ip; and the firewall name for a virtual server.
Partition	Usually, `Common`. An *administrative partition* is a part of the BIG-IP® configuration that is accessible only to a particular group of administrators. The default partition for all BIG-IP configurations, `Common`, is accessible to all administrators. A sufficiently-privileged administrator can make additional partitions on the BIG-IP device. Each partition corresponds to a folder (with the same name, for instance, `/Common`) to hold its configuration objects.
Firewall Type	One of the following: global (global); route-domain (rd); virtual server (vip); self-ip (self-ip); or management-ip (mgmt).
IP Address	For Virtual server (VIP), self IP, and Management firewall types only; this is an informational, read-only field displaying the IP address retrieved (if available) during DMA.
Description	Optional description for the firewall.
Route Domain ID	Used for Route Domain firewall types only; displays a number that identifies the route domain.
Device	Name of the BIG-IP® device where the firewall resides.
Enforced Policy	Name of the enforced policy assigned to the firewall context. An enforced firewall policy modifies network traffic based on a set of firewall rules. This property is not used for the Management firewall type.
Staged Policy	Name of the staged policy assigned to the firewall context. A staged firewall policy allows you to evaluate the effect a policy has on traffic without actually modifying the traffic based on the firewall rules. This property is not used for the Management firewall type.
Service Policy	Name of the service policy assigned to the firewall context. This property is not used for the Management firewall type.
NAT Policy	Name of the NAT policy assigned to the firewall context.

Adding an enforced firewall policy

You can view and configure firewall policies or rules to force or refine actions (accept, accept decisively, drop, reject) using the Enforced settings. You are restricted to a single, enforced firewall policy on any specific firewall context.

Note: *Policies can be enforced in one firewall context and staged in another.*

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Contexts** in the list on the left to expand the contents and click one of the context types.
4. Click the name of the context to edit. The context properties are displayed.
5. Click **Add Enforced Firewall Policy** in the Enforced Firewall Policy row and in the resulting popup, click the policy to use and click **Add**. Alternatively, drag-and-drop a policy from those listed in the Policy Editor toolbox at the bottom of the page to the Enforced Firewall Policy row.

   Adding an enforced policy results in the removal of all existing rules.
6. Click the name of the enforced policy to display the policy properties.
7. Click **Create Rule** to add a rule by editing the fields in the template.

   You can also add rules by right-clicking in the last rule in the table and selecting **Add rule before** or **Add rule after**. If you right-click after the bottom row in the Rules table, you can select the option **Add rule**. You can then reorder rules by dragging and dropping them until they are in the correct order for execution. You can also reorder rules by right-clicking in the row and selecting among the ordering options.
8. Add a rule list by clicking **Add Rule List**.
9. In the popup screen that opens, select the name of the rule list that you want to add and then click **Add**.
10. Click **Save** to save changes.

    To clear a lock without saving changes, click the **Unlock** link.
11. When finished, click **Save & Close** to save your edits, clear the lock, and exit.

Adding a staged firewall policy

You can stage firewall policies using the Staged settings. Actions (accept, accept decisively, drop, reject) have no effect on network traffic. Rather, they are logged. This gives you the ability to stage a firewall policy first and examine the logs to determine how the firewall policy has affected traffic. Then, you can determine the timing for turning the firewall policy from staged to enforced.

Rule and rule lists are not allowed on staged firewall policies.

Note: *A firewall policy can be staged in one context and enforced in another.*

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Contexts** in the list on the left to expand the contents and click one of the context types.
4. Click the name of the context to edit. The context properties are displayed.

5. Click **Add Staged Firewall Policy** in the Staged Firewall Policy row and in the resulting popup, click the policy to use and click **Add**. Alternatively, drag-and-drop a policy from those listed in the Policy Editor toolbox at the bottom of the page to the Staged Firewall Policy row.

   Adding an enforced policy results in the removal of all existing rules and rule lists.

6. Click **Save** to save changes.

   To clear a lock without saving changes, click the **Unlock** link.

7. When finished, click **Save & Close** to save your edits, clear the lock, and exit.

Managing Rules and Rule Lists

About rules and rule lists

Rule lists are containers for rules, which are run in the order they appear in their assigned rule list. A rule list can contain thousands of ordered rules, but cannot be nested inside another rule list. You can reorder rules in a given rule list at any time.

With BIG-IQ® Network Security, you can manage rules and rule lists from the Rule Lists option (Policy Editor > Rule Lists). You can also create rules and add rule lists from the Contexts and the Policies options. You can import and manage rules (and/or rule lists) from BIG-IP® devices. Furthermore, you can define rules and rule lists within BIG-IQ Network Security, and then deploy back to the BIG-IP device.

You can define a list of rules for a specific firewall and/or refer to one or more shared rule lists by name from other firewalls.

Network firewalls use rules and rule lists to specify traffic-handling actions. The network software compares IP packets to the criteria specified in rules. If a packet matches the criteria, then the system takes the action specified by the rule. If a packet does not match any rule from the list, the software accepts the packet or passes it to the next rule or rule list. For example, the system compares the packet to self IP rules if the packet is destined for a network associated with a self IP address that has firewall rules defined.

A packet must pass all tests to match successfully. For example, to match against a source subnet and several destination ports, a packet must originate from the given subnet and also have one of the specified destination ports.

Rules and rule lists can be applied to all firewall types, such as:

- Global
- Route domain
- Virtual server
- Self IP
- Management IP (rules only, no iRule or geolocation support)

Enabling, disabling and scheduling rules and rule lists

Once a rule or a rule list is created, you can set the state of that rule or rule list to enable it, disable it, or schedule when it is enabled. By default, a rule or rule list is enabled. Settings on a rule list take precedence over those on a rule. For example, if a rule has a state of enabled, but is contained within a rule list that has a state of disabled, the rule used in that rule list will be disabled. The process differs for setting the state of a rule and setting the state of a rule list.

- To set the state for a rule, edit the rule and choose enabled, disabled or scheduled in the State column.
- To set the state for a rule list, edit the rule list, and right click the rule list name and select **Edit Rule List Reference**. The state can now be set by choosing enabled, disabled or scheduled in the State column.

Filtering rule lists

To filter the system interface to display only those objects related to a selected rule list, hover over the rule list name, right-click and then click **Filter 'related to'**. The interface is filtered and a count appears to the right of each object type. The frame to the right provides its own filter field where you can enter text and click on the filter icon to constrain the display to those items that match the filter.

Creating rules

To support a context or policy, you can create specific rules, gather those rules in a rule list, and assign the rule list to the context or policy.

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Select the object to which you want to add the rule:

Option	Description
Rule list	In the left pane, click **Rule Lists** to display the rule lists, then select the rule list to have the rule added.
Context	In the left pane, click **Contexts** to display the contexts, then select the context to have the rule added.
Policy	In the left pane, click **Policies** to display the firewall policies, then select the policy to have the rule added.

4. Add the rule to the object:

Option	Description
Rule list	In the right pane, click **Create Rule**.
Context	In the right pane, click the name of the context staged or enforced policy to which you want to add the rule, then click **Create Rule**.
Policy	In the right pane, click **Create Rule**.

   A new row appears in the table of rules. The row contains a rule template, including defaults, for the new rule.
5. Complete the fields as appropriate.

   You can also add rules by right-clicking in the Rules table, or by right-clicking any row in the Rules table and choosing **Add Rule before** or **Add Rule after**.
6. Click **Save** to save your changes.
7. When you are finished, click **Save & Close** to save your edits, clear the lock, and exit the panel.

Reordering rules in rule lists

You can optimize your network security firewall policy by reordering rules in rule lists.

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Rule Lists** in the left pane and click the specific rule list you want to edit in the right pane.
4. Click the **Rules** tab to ensure it is selected.
5. Drag-and-drop the rules until they are in the correct order.

   If the list of rules expands beyond the editing frame, drag-and-drop will not work. Instead, copy the rule by right-clicking and selecting **Copy Rule**. Then, navigate to the new location for the rule, right-click, and select **Paste Before** or **Paste After** as appropriate. After the copy, delete the rule that you copied. You delete rules by right-clicking on a rule and selecting **Delete Rule**.

   Alternatively, you can reorder rules using the **Cut Rule** option. Right-click on the rule and select **Cut Rule** to select the rule for reordering, then move your cursor to the new position in the rule list, and

select **Paste Before** or **Paste After** as appropriate. The rule is removed from the original position when it is pasted in the new position in the rule list, but not before.

6. When you are finished, click **Save & Close** to save your edits, clear the lock, and exit the panel.

Removing rules

You can remove specific rules from rule lists, firewalls, or policies, to fine tune security policies.

Note: *You can remove a rule even if it is the only rule in the rule list.*

1. You remove a rule based on the object that you remove it from:

Option	Description
From a rule list	In the left pane, expand **Rules Lists** and click the name of the rule list containing the rule that you want to delete. This opens the Rule List frame that provides access to **Properties** and **Rules** options.
From a firewall context	In the left pane, expand **Contexts**, click the name of the context containing the rule that you want to delete.This opens the Properties frame which contains the Enforced Policy row and the Staged Policy row, either of which may contain a policy. Click the policy name containing the rule to delete and then click **Rules & Rule Lists**.
From a policy	In the left pane, expand **Policies**, click the name of the policy containing the rule that you want to delete. The Policy frame opens and provides access to **Properties** and **Rules & Rule Lists** options. Select **Rules & Rule Lists**.

2. Hover over the row containing the rule, and right-click.
3. Select **Delete rule** and, if prompted, confirm the deletion.
4. Click **Save** to save your changes.

Creating and adding rule lists

To support a specific firewall or policy, you can create a rule list and then assign it to the firewall context or policy.

1. Click **Policy Editor**.
2. Click Rule Lists in the navigation pane on the left.
3. In the Rule Lists pane on the right, click **Create**.
4. Click **Properties** and complete the properties fields as required.

Option	Description
Name	Unique name. The field is read-only field unless creating or cloning the rule list.
Description	Optional description.
Partition	Although pre-populated with `Common` (default), you can set the partition name by typing a unique name for the partition. ***Note:*** *The partition with that name must already exist on the BIG-IP device. No whitespace is allowed in the partition name.* The firewall partition itself is not editable.

5. Click **Rules** and create or add rules to the rule list.

6. Click **Save** to save your changes or **Save & Close** to save your changed and exit the screen.
7. Select the object in the Policy Editor to which you want to add the rule list:

Option	Description
Context	Select Contexts in the navigation frame on the left, and then click the specific firewall context to have a rule list added.
Policy	Select Policies in the navigation frame on the left, and then click the specific firewall policy to have a rule list added.

8. Add the rule list to the selected object:

Option	Description
Context	Click the enforced or staged policy to which the rule list should be added, then click **Add Rule List**, select from the rule lists in the popup dialog, and click **Select**.
Policy	Click **Rules & Rule Lists**, then click **Add Rule List** , then select from the rule lists in the popup dialog, and click **Select**.

   You can add rules by right-clicking in the Rules table, or by right-clicking any row in the Rules table and choosing **add rule before** or **add rule after**.
9. When you are finished, click **Save** or **Save & Close**, as appropriate.

Editing rule lists

You can edit the content of rule lists from Policy Editor Rule Lists, including the order of rules in rule lists.

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Rule Lists** in the left pane and click the specific rule list you want to edit in the right pane.
4. Click **Properties**.

Option	Description
Name	Informational, read-only field set when creating or cloning the rule list.
Description	Optional description.
Partition	Informational, read-only field set when creating or cloning the rule list.

5. Click **Rules**, and click the name of the rule you want to edit.
6. Complete the fields as appropriate.

   You can also add rules by right-clicking in the Rules table, or by right-clicking any row in the Rules table and choosing **Add Rule before** or **Add Rule after**.
7. Complete fields as appropriate.

   To reorder rules, simply drag-and-drop the rules until they are in the correct order. If the list of rules expands beyond the editing frame, drag-and-drop will not work. Instead, copy the rule by right-clicking and selecting **Copy Rule**. Then, navigate to the new location for the rule, right-click, and select **Paste Before** or **Paste After** as appropriate. After the copy, delete the rule that you copied.
8. Click **Save** to save your changes.

Changes made to the rule list are reflected the next time the Contexts or Policies screen is refreshed.

Clearing fields in rules

You can clear the text from fields in rules to fine tune them and, in turn, rule lists and security policies.

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Expand **Rule Lists** and click the name of a rule list that you want to edit.
4. Click the **Rules** tab to ensure it is selected.
5. Click the name of the rule containing the fields whose contents you want to remove.
6. Not all fields can be cleared, but you can remove the contents of these fields as follows:

Option	Description
Address (source or destination)	Click the **X** to the right of the field.
Port (source or destination)	Click the **X** to the right of the field.
VLAN	Click the **X** to the right of the field.
iRule	Click the **X** to the right of the field.
Description	Click the **X** to the right of the field.

7. Click **Save** to save your changes.
8. When you are finished, click **Save & Close** to save your edits, clear the lock, and exit the panel.

Cloning rule lists

Cloning enables you to create and customize rule lists to address unique aspects of your network firewall environment. When you clone a rule list, you create an exact copy of the rule list, which you can then edit to address any special considerations.

Note: *Users with the roles of Network_Security_View or Network_Security_Deploy cannot clone policies.*

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Rule Lists** to display the rule list you want to clone, and then click the checkbox to the left of that rule list.
4. Click **Clone**.
5. Click **Properties** and complete the properties fields as required.

Option	Description
Name	Unique name. The field is read-only field unless creating or cloning the rule list.
Description	Optional description.
Partition	Although pre-populated with `Common` (default), you can set the partition name by typing a unique name for the partition. ***Note:*** *The partition with that name must already exist on the BIG-IP device. No whitespace is allowed in the partition name.* The firewall partition itself is not editable.

6. Click **Rules**, edit the rules as required to configure the clone.

You can also click **Create Rule** to add a new rule.

7. When you are finished, click **Save**.

   If you click **Cancel**, the rule list is not cloned.

The cloned rule list is added alphabetically under **Rule Lists**. In a high-availability configuration, the cloned rule list is replicated on the standby system as soon as it is cloned.

Removing rule lists

You can remove rule lists from firewalls or policies to fine tune security policies.

1. Log in to BIG-IQ® Network Security.
2. Click **Policy Editor**.
3. Click **Rule Lists** to display the rule list you want to remove, and then click the checkbox to the left of that rule list.
4. At the top of the screen, click **Delete**.
5. If it is safe to remove the rule list, a confirmation dialog box opens; click **Delete** to confirm.

   If the rule list is in use, you cannot complete the removal. A popup screen opens informing you that you cannot remove the rule list because it is in use. Click **Close** to acknowledge this message, and then click **Cancel** in the Delete Rule Lists popup screen. To see where a rule list is used, right click on the rule list name and select **Filter 'related to'**. A search is performed and any object using the rule list will have a non-zero number appear next to it in the navigation pane on the left. To clear the search, click the **x** icon to the right of the search string.

Rule properties

This table lists and describes the properties required when you are configuring network firewall rules.

Property	Description
Name	Unique, user-provided name for the rule. If the name is a rule list name, it is preceded by: `referenceTo_` when moved to a firewall or policy. For example: `referenceTo_sys_self_allow_all`.
Address (Source)	There are many ways to construct an IPv4 or IPv6 address, address range, or address list. The following methods and examples are not meant to be exhaustive. • IPv4 format: *a.b.c.d*[/*prefix*]. For example: `60.63.10.10` • IPv6 format: *a*:*b*:*c*:*d*:*e*:*f*:*g*:*h*[/*prefix*]. For example: `2001:db7:3f4a:9dd:ca90:ff00:42:8329` • You can specify subnets using forward slash (/) notation; for example: `60.63.10.0/24`. An example of an IPv6 subnet is as follows: `2001:db8:a::/64`. • You can append a route domain to an address using the format %RouteDomainID/Mask. For example, `12.2.0.0%44/16`. From the list, select: • **Address**. Enter the address in the **Addresses** field. You can also type an address range in the **Addresses** field using the format: n.n.n.n-n.n.n.n. For example: `1.1.1.1-2.2.2.2`. • **Address range**. Type the beginning address in the first **Addresses** field and the ending address in the second **Addresses** field.

Property	Description
	• **Address list**. In the **Addresses** field, type text to display stored address lists. You can select any of the address lists displayed. • **Country/Region**. From the first **Addresses** list, select a country. Once you select a country, the second list automatically updates with all available regions for that country. Optionally, select a region from the second list. The wildcard, Unknown, is supported. Note that geolocation is not supported on the management IP context. Options are provided to add additional addresses, address ranges, address lists, or countries/regions (+) and to delete addresses, address ranges, address lists, or countries/regions (**X**). When you are finished, click **Save** or **Add**.
Port	Ports, port ranges, or port lists. From the list, select: • **Port**. Type the port in the **Ports** field. You can also enter a port range in the port field by typing a range in the format: n-n. For example: `43-44`. • **Port range**. Type the beginning port in the first **Ports** field and the ending port in the second **Ports** field. • **Port list**. In the **Ports** field, type text to display stored port lists. You can select any of the port lists displayed. Options are provided to add additional ports, port ranges, or port lists (+) and to delete ports, port ranges, or port lists (**X**). When you are finished, click **Save** or **Add**.
VLAN	Name of the VLAN physically present on the device (Internal, External, or Any). If you specify a VLAN in a rule without also specifying the VLAN's partition, the deployment task will fail when you attempt to deploy that rule to a firewall. Use the format `partition/VLAN` or `/partition/VLAN`. For example: `Common/external` or `/Common/external`. When you are finished, click **Save** or **Add**.
Address (Destination)	There are many ways to construct an IPv4 or IPv6 address, address range, or address list. The following methods and examples are not meant to be exhaustive. • IPv4 format: *a.b.c.d*[/*prefix*]. For example: `60.63.10.10` • IPv6 format: *a*:*b*:*c*:*d*:*e*:*f*:*g*:*h*[/*prefix*]. For example: `2001:db7:3f4a:9dd:ca90:ff00:42:8329` • You can specify subnets using forward slash (/) notation; for example: `60.63.10.0/24`. An example of an IPv6 subnet is as follows: `2001:db8:a::/64`. • You can append a route domain to an address using the format %RouteDomainID/Mask. For example, `12.2.0.0%44/16`. From the list, select: • **Address**. Type the address in the **Addresses** field. You can also enter an address range in the **Addresses** field using the format: n.n.n.n-n.n.n.n. For example: `1.1.1.1-2.2.2.2`. • **Address range**. Type the beginning address in the first **Addresses** field, and the ending address in the second **Addresses** field. • **Address list**. In the **Addresses** field, type text to display stored address lists. You can select any of the address lists displayed. • **Country/Region**. From the first **Addresses** list, select a country. Once you select a country, the second list automatically updates with all available regions for that country. Optionally, select a region from the second list. The wildcard, Unknown, is supported. Note that geolocation is not supported on the management IP context.

Property	Description
	Options are provided to add additional addresses, address ranges, address lists, or countries/regions (+) and to delete addresses, address ranges, address lists, or countries/regions (**X**). When you are finished, click **Save** or **Add**.
Port	Ports, port ranges, or port lists. From the list, select: • **Port**. Type the port in the **Ports** field. You can also enter a port range in the port field by typing a range in the format: n-n. For example: **43-44**. • **Port range**. Type the beginning port in the first **Ports** field and the ending port in the second **Ports** field. • **Port list**. In the **Ports** field, type text to display stored port lists. You can select any of the port lists displayed. Options are provided to add additional ports, port ranges, or port lists (+) and to delete ports, port ranges, or port lists (**X**). When you are finished, click **Save** or **Add**.
Action	Click in the column and select one of the following: • **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. • **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. If the Rule List is applied to a virtual server, management IP, or self IP firewall rule, then Accept Decisively is equivalent to Accept. • **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. When you are finished, click **Save** or **Add**.
iRule	Click in the column and enter the iRule name, including partition. For example: `/Common/_sys_AXX_Support_OA_BasicAuth`. You can also set sampling rates on iRules® by supplying a number in the **Sampling Rate** field. iRules® use syntax based on the industry-standard Tools Command Language (Tcl). For complete and detailed information on iRules syntax, see the F5 Networks DevCentral web site, `http://devcentral.f5.com`. Note that iRules must conform to standard Tcl grammar rules. For more information on Tcl syntax, see `http://tmml.sourceforge.net/doc/tcl/index.html`. Note that iRules are not supported on the management IP context.
Description	Optional description for the current rule. To add a description, click in the column, type text, and click **Save** or **Add**.
Protocol	IP protocol to compare against the packet. Select the appropriate protocol from the list and click **Save** or **Add**. If you select **ICMP**, **IPv6-ICMP**, or **Other**, a popup dialog box opens where you can specify **Type** and **Code** combinations. The default type is **Any** and the default code is **Any**.

Property	Description
	Note: *The type and code combinations are too numerous to document here. For details, consult the F5 Networks DevCentral site,* `http://devcentral.f5.com` *or the documentation for the specific BIG-IP® platform.*
State	Click in the column and select an option from the list to specify whether the rule is enabled, disabled, or scheduled. The field is updated. Click **Save** or **Add** when you are ready to save your changes. If you select **scheduled** from the list, the **Select Schedule** list is displayed in the screen. Select a schedule and click **OK**. If you have assigned a schedule, then a gear icon appears to the right of the **State** setting in the State column. To make changes to the **State** setting, click the gear icon to open the Select Schedule popup screen. If you have no pre-defined schedules, you cannot assign the scheduled state to the rule.
Log	Click in the column and select an option from the list to specify whether or not the firewall software should write a log entry for any packets that match this rule. From the list, select **true** (log an entry) or **false** (do not log an entry). When your are finished, click **Save** or **Add**. For you to set or edit this setting, the discovered device must be at version 11.3 HF6 or later. The setting is not editable earlier than version 11.3 HF6. When a new rule is added to a firewall through the BIG-IQ® Network Security system interface, editing is enabled for the **Log** setting even for devices with versions earlier than 11.3 HF6.

Managing Service, Timer, and Port Misuse Policies

About service, timer, and port misuse policies

A *service policy* allows you to associate network idle timers (timer policies) or port misuse policies on firewall contexts and rules.

You can discover a service policy on a BIG-IP® device version 12.0, or later. Or you can create one on a BIG-IQ® Centralized Management system using the Network Security policy editor, and then deploy it to a BIG-IP device version 12.0, or later. You can apply a service policy to the global, self IP address, or route domain context. You can also add it to a rule in a rule list, or to a rule on a security policy.

A service policy can contain timer policies or port misuse policies, or both. You create service policies, timer policies, and port misuse policies separately, and then you add the timer policies or port misuse policies to the service policies.

- You use a *timer policy*, also known as a *firewall idle timer*, to configure timer rules that can be associated with firewall contexts and rules. You can discover a timer policy on a BIG-IP device version 12.0, or later, or create one on a BIG-IQ Centralized Management system using the Network Security policy editor and then deploy it to a BIG-IP device version 12.0, or later.
- A *port misuse policy* allows you to configure a firewall context or rule to detect and drop network 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. You can discover a port misuse policy on a BIG-IP device version 12.1, or later, or create one on a BIG-IQ Centralized Management system using the Network Security policy editor, and then deploy it to a BIG-IP device version 12.1, or later.

Create a timer policy

You create a timer policy containing timer rules to add to a service policy that can be applied to the global, self IP address, or route domain contexts.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left, click **Timer Policies**.
4. Click **Create**.
   The Timer Policies - New Item screen opens.
5. In the **Name** field, type a name for the timer policy.
6. In the **Description** field, type an optional description for the timer policy.
7. If needed, change the default Common partition in the **Partition** field.
8. To add timer rules, click the Timer Rules tab and click **Create Timer Rule**.
   A new rule is displayed with default name and values.
9. Click the name of the new rule to enable editing for the rule fields.
10. In the **Name** field, you may specify a more meaningful name than the default.
11. From the **Protocol** list, select the protocol to be used. Select **all-other** as the protocol to have the rule apply to all other protocols not specified in another timer rule in the policy.

12. From the **Destination Ports** list, specify the one or more ports to use, if necessary. The default is to use any port.
    - Select **Port** to specify an individual port: type the port in the field provided, and then click +. You can enter multiple individual ports, one at a time.

      Enter 0 as the port value to specify all other ports that have not been specified using **Port** or **Port Range**.
    - Select **Port Range** to specify a range of ports: type the beginning port in the first field, and the ending port of the range in the second field provided, and then click +. You can enter multiple ports ranges, one at a time.
    - Select **All Other** to specify all other ports that have not been specified using **Port** or **Port Range**.
13. From the **Idle Timeout** list, select the timeout option for the selected protocol.
    - Select **Specify** to specify the timeout for this protocol, in seconds. Type the number of seconds in the field provided.
    - 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 system uses the default timeout for the protocol.
14. Save your changes in one of two ways:
    - Click **Save** to save the timer policy rule.
    - Click **Save & Close** to save the timer policy rule and return to the Timer Policies screen.

The timer policy is now configured and can be added to a service policy.

You now need to add the timer policy to a service policy, and apply the service policy to a global, self IP address, or route domain context. You can also add it to a firewall rule on a policy, or in a rule list.

Create a port misuse policy

You create a port misuse policy containing port misuse rules to add to a service policy that can be applied to the global, self IP address, or route domain contexts.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left, click **Port Misuse Policies**.
4. Click **Create**.
   The Port Misuse Policies - New Item screen opens.
5. Type a name and an optional description for the port misuse policy.
6. If needed, change the default Common partition in the **Partition** field.
7. In the **Default Actions** row, select the default actions to occur when port misuse is detected. You can select none, one, or both options.
   - 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.
8. To add port misuse rules, click the Port Misuse Policy Rules tab and click **Create Port Misuse Rule**. The screen displays a new port misuse rule with default name and values.
9. Click the name of the new rule to enable editing for the rule fields.
10. In the **Name** field, you may specify a more meaningful name than the default.
11. In the **Port** field, 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. The default port number is automatically supplied for the common ports.

12. In the **IP Protocol** field, select the IP protocol for the port matching rule.
13. In the **Service** field, select the service to use. 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.

    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.
14. In the **Drop on Service Mismatch** list, select the drop behavior.
    - 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.
    - Select **Use Policy Default** to use the default action for packet drops, when the service does not match the port.
15. In the **Log on Service Mismatch** list, select the behavior for logging packet drops.
    - 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.
    - Select **Use Policy Default** to use the default action for logging packet drops, when the service does not match the port.
16. Save your changes.

You have configured the port misuse policy.

You now can add the port misuse policy to a service policy, and apply the service policy to a global, self IP address, or route domain context. You can also add it to a firewall rule on a policy, or in a rule list.

Create a service policy

You create a service policy to contain timer policies that can be applied to the global, self IP address, or route domain contexts. Service policies can also be added to a rule in a rule list or a rule on a security policy.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left click **Service Policies**.
4. Click **Create**.
   The Service Policies - New Item screen opens.
5. In the **Name** field type a name for the service policy.
6. If needed, change the default Common partition in the **Partition** field.
7. In the **Description** field, type an optional description for the service policy.
8. Select a timer policy from those listed in the **Timer Policy** list.

   If no timer policy is listed, you need to create one and then assign it to the service policy.
9. In the **Pin Policy to Device(s)** area, select the BIG-IP devices to be pinned to this policy, if needed. Pinning a BIG-IP device to a policy enables the policy to be deployed even if it is not associated with a firewall context for that device. You select the BIG-IP device to use by moving it from the **Available** list to the **Selected** list using the arrow buttons. You can filter the list of available BIG-IP devices using the filter field at the top of the **Available** list. Moving a BIG-IP device that is part of a cluster to the **Selected** list will cause the other member of the cluster to move to that list as well.

   If you have a self IP context with a static (non-floating) IP address, you may be required to assign the device depending on you cluster deployment settings. For example, this property must be set for a

peer BIG-IP device that is part of a DSC cluster managed by the BIG-IQ Centralized Management system. You may be directed to set this property as a result of an evaluation critical error.

10. Save your changes.

You have defined the service policy. You can now assign it to a global, self IP address, or route domain context. You can also add it to a rule in a rule list, or a rule on a security policy.

Apply a service policy to a firewall rule

You apply a service policy to a firewall rule to apply timer policies or port misuse policies to traffic that is matched by the firewall rule. The rule can be associated with a rule list or with a firewall security policy.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**.
4. Display the list of rules from a rule list or from a firewall security policy in the policy editor.

Option	Description
If the rule is in a rule list:	On the left, click **Rule Lists**, and then click the name of the rule list containing the rule. The rules are listed on the Rules tab.
If the rule is associated with a policy:	On the left, click **Firewall Policies**, and then click the name of the policy containing the rule. The rules are listed on the Rules & Rule Lists tab.

5. To make it editable, click the name of the rule to which you want to add the service policy.
6. Add the service policy to the rule.

Option	Description
Add the service policy by typing.	Type the name of the service policy in the Service Policy column for the rule. The system completes name of the service policy once you begin typing the name.
Add the service policy by drag and drop.	In the Shared Resources area, select **Service Policies**, and then drag the service policy from that list and drop it into the Service Policy column for the rule.

7. Save your changes.

The service policy is added to the rule.

Apply a service policy to a global context

You apply a service policy to a global firewall context to use a timer or port misuse policy with that context.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **Global** .
4. Click the name of the global context to open it for editing.
5. Add the service policy to the Service Policy row:
   a) Click **Add Service Policy**.

b) From the popup screen select the service policy to add.
c) Click **Select**.

You can also add a service policy by selecting **Service Policies** in the Shared Resources list, and then dragging one of the displayed service policies and dropping it onto the Service Policy row. To remove a service policy, click the **X** to the right of the service policy name in the Service Policy row.

6. Save your changes.

The service policy is now associated with the global context.

Apply a service policy to a route domain context

You apply a service policy to a route domain firewall context in order to use a timer policy.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor** and then **Route Domain** from the list on the left.

   If the Route Domain context is not displayed, click **Contexts** in the list to expand the list of contexts and display it.
4. Click the name of the route domain to open it for editing.
5. Add the service policy to the Service Policy row:
   a) Click **Add Service Policy**.
   b) From the popup screen select the service policy to add.
   c) Click **Select**.

   You can also add a service policy by selecting **Service Policies** in the Shared Resources list, and then dragging one of the displayed service policies onto the Service Policy row. To remove a service policy, click the **X** to the right of the service policy name in the Service Policy row.
6. Save your changes.

The service policy is now associated with the route domain context.

Apply a service policy to a self IP address context

You apply a service policy to a self IP address firewall context so you can use a timer policy.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left click **Self IP**.
4. Click the name of the self IP address to open it for editing.
5. Add the service policy to the Service Policy row:
   a) Click **Add Service Policy**.
   b) From the popup screen select the service policy to add.
   c) Click **Select**.

   You can also add a service policy by selecting **Service Policies** in the Shared Resources list, and then drag one of the displayed service policies and drop it onto the Service Policy row. To remove a service policy, click the **X** to the right of the service policy name in the Service Policy row.
6. Save your changes.

The service policy is now associated with the self IP address context.

Delete a timer policy

You can delete obsolete timer policies that are no longer used by a service policy to avoid clutter in the user interface.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left, click **Timer Policies**.
4. Select the check box to the left of any timer policy that you want to remove.
5. Click **Delete**.
6. Confirm that you want to remove the timer policy by clicking **Delete** in the confirmation dialog box.

The system removes the selected timer policies.

Delete a port misuse policy

You can delete obsolete port misuse policies that are no longer used by a service policy to avoid clutter in the user interface.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left, click **Port Misuse Policies**.
4. Select the check box to the left of any port misuse policy that you want to remove.
5. Click **Delete**.
6. Confirm that you want to remove the port misuse policy by clicking **Delete** in the confirmation dialog box.

The system removes the selected port misuse policy.

Delete a service policy

You should delete service policies that are no longer used, to simplify your view.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then in the list on the left click **Service Policies**.
4. Select the check box to the left of any service policy you want to remove.
5. Click **Delete**.
6. Confirm that you want to remove the service policy by clicking **Delete** in the confirmation dialog box.

The system removes the selected service policies.

Managing NAT Policies and Translations

About NAT policies and translations

You can use network translation address (NAT) policies to translate network addresses. These NAT policies contain rules that contain NAT source translations and NAT destination translations.

You associate a NAT policy with a global context, route domain context, or virtual server firewall context by adding it to the NAT Policy property of the firewall context.

You can discover a NAT policy on a BIG-IP® device version 12.1 or later, or create one on a BIG-IQ® system using the Network Security Policy Editor and then deploy it to a BIG-IP device version 12.1 or later.

Creating a NAT policy

You create a NAT policy to contain rules that contain NAT source translations and NAT destination translations.

1. Log in to the BIG-IQ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Policies**.
4. Click **Create**.
   The NAT Policies - New Item screen opens with the Properties tab displayed.
5. Type a name for the NAT policy in the **Name** field.
6. Type an optional description for the NAT policy in the **Description** field.
7. If needed, change the default `Common` partition in the **Partition** field.
8. Click the Rules tab and the click **Create Rule**.
   A new row appears in the table of rules. The row contains a rule template, including defaults, for the new rule.
9. Click the name of the rule to edit the default rule properties.
10. Complete the rule fields as appropriate.

    You can also add rules by right-clicking in the Rules table, or by right-clicking any row in the Rules table and choosing one of the options available.
11. Click **Save** to save the NAT policy, or click **Save & Close** to save the NAT policy and return to the NAT Policies page.

The NAT policy is now defined and can be assigned to a firewall context.

NAT rule properties

The following table lists and describes the properties required when configuring NAT policy rules. These rules are similar to rules used in firewall policies, but have a different set of properties.

Property	Description
Name	Unique, user-provided name for the rule.

Property	Description
Address (Source)	Source address or addresses. Select the type of source address from the list: • **Address**. Type a single address in the **Address** field and then click + to the right of the address field to add it. • **Address List**. In the **Address** field, type the name of the address list. Alternatively, from the **Shared Resources** list at the bottom, you can select **Address Lists** to list those available, and then drag and drop it into the **Address** field. • **Address Range**. Type the beginning address in the first **Address Range** field and the ending address in the second **Address Range** field. Then click + to the right of the address field to add it. When you are finished, click **Save** or **Save & Close**.
Port (Source)	Source port or ports. Select the type of port from the list: • **Port**. Type the port in the **Port** field. • **Port Range**. Type the beginning port in the first **Port** field and the ending port in the second **Port** field. Then click + to the right of the address field to add it. • **Port List**. In the **Port** field, type the name of the port list. Alternatively, from the **Shared Resources** list at the bottom, you can select **Port Lists** to list those available and then drag and drop it into the **Port** field. When you are finished, click **Save** or **Save & Close**.
VLAN (Source)	Name of the VLAN physically present on the device (Internal, External, or Any). If you specify a VLAN in a rule without also specifying the VLAN's partition, the deployment task will fail when you attempt to deploy that rule to a firewall. Use the format `partition/VLAN` or `/partition/VLAN`. For example: `Common/external` or `/Common/external`. When you are finished, click **Save** or **Save & Close**.
Address (Destination)	Select the type of destination address from the list: • **Address**. Type a single address in the **Address** field and then click + to the right of the address field to add it. • **Address List**. In the **Address** field, type the name of the address list. Alternatively, from the **Shared Resources** list at the bottom, you can select **Address Lists** to list those available and then drag and drop it into the **Address** field. • **Address Range**. Type the beginning address in the first **Address Range** field and the ending address in the second **Address Range** field. When you are finished, click **Save** or **Save & Close**.
Port (Destination)	Destination port or ports. Select the type of port from the list: • **Port**. Type the port in the **Port** field. • **Port Range**. Type the beginning port in the first **Port** field and the ending port in the second **Port** field. • **Port List**. In the **Port** field, type the name of the port list. Alternatively, from the **Shared Resources** list at the bottom, you can select **Port Lists** to list those available and then drag and drop it into the **Port** field. When you are finished, click **Save** or **Save & Close**.
Description	Optional description for the current rule. To add a description, click in the column, type text, and click **Save** or **Add**.

Property	Description
Protocol	IP protocol to compare against the packet. Select the appropriate protocol from the list and click **Save** or **Save & Close**. The default type is **Any** and the default code is **Any**. ***Note:*** *The type and code combinations are too numerous to document here. For details, consult the F5 Networks DevCentral site,* `http://devcentral.f5.com`*, or the documentation for the specific BIG-IP® platform.*
State	Select whether the rule is enabled or disabled. The field is updated. Click **Save** or **Save & Close** to save your changes.
Translated Source	Type the name of a NAT Source Translation in the field. Alternatively, from the **Shared Resources** list at the bottom, you can select NAT Source Translations to list those available and then drag and drop it into the **Translated Source** field.
Translated Destination	Enter the name of a NAT Destination Translations in the field. Alternatively, from the **Shared Resources** list at the bottom, you can select NAT Destination Translations to list those available and then drag and drop it into the **Translated Destination** field.
Log Profile	Type the name of a logging profile in the field. This logging profile must already be defined using Logging Profiles in Shared Security.

Cloning a NAT policy

Cloning enables you to create a copy of the NAT policy, which you can then edit to address any special considerations.

1. Log in to the BIG-IQ ®system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left click **NAT Policies**.
4. Select the NAT policy to clone by selecting the check box for it, and then clicking **Clone**.
   A copy of that NAT policy is created with the same name, but with -CLONE appended to the name.
5. Change the NAT policy as needed.
6. Click **Save** to save the NAT policy, or click **Save & Close** to save the NAT policy and return to the NAT Policies page.

The NAT policy is now defined and can be assigned to a firewall context.

Deleting a NAT policy

You delete NAT policies that are no longer used.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Policies**.
4. Select the one or more NAT policies to be removed by selecting the check box for the appropriate NAT policy.
5. Click **Delete**.
6. Confirm that you want to remove the NAT policy by clicking **Delete** in the confirmation dialog box.

The selected NAT policies are removed.

Creating NAT source translations

Create NAT source translations to use within a NAT policy rule.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Source Translations**.
4. Click **Create**.
   The NAT Source Translations - New Item screen opens.
5. Type a name for the NAT source translations in the **Name** field.
6. In the **Description** field, type an optional description for the NAT source translations.
7. If needed, change the default `Common` partition in the **Partition** field.
8. From the **Type** list, select the type of address translation to use. The type of address translation you select determines what additional properties are available.
   - Select **Static NAT** for static network address translation.
   - Select **Static PAT** for static network port and address translation.
   - Select **Dynamic PAT** for dynamic network port and address translation.
9. If you selected **Static NAT** for the value of the **Type** list, supply values for the following settings.

Property	Description
Addresses	Add one or more addresses or address ranges by entering them and then clicking the + button. Remove them by clicking the **X** button next to the address or address range.
ICMP Echo	Select whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes. • Select **disabled** to disable ICMP echoes.
Egress Interfaces	Select whether the source address is translated for egressing network traffic, and on what interfaces, such as the `/Common/http-tunnel` interface. • Select **Disabled on** to disable source address translation for the specified interfaces, and then select the check box for the interfaces to be disabled. • Select **Enabled on** to enable source address translation for the specified interfaces and then select the check box for the interfaces to be enabled.

10. If you selected **Static PAT** for the value of the **Type** list, fill in the following settings.

Property	Description
Addresses	Add one or more addresses or address ranges by typing them and then clicking the + button. Remove them by clicking the **X** button next to the address or address range.

Property	Description
Ports	Add one or more ports or port ranges by typing them and then clicking the + button. Remove them by clicking the **X** button next to the port or port range.
ICMP Echo	Select whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes. • Select **disabled** to disable ICMP echoes.
Egress Interfaces	Select whether egress interfaces are available. • Select **Disabled on** to disable egress filtering interfaces. • Select **Enabled on** to disable egress filtering interfaces.

11. If you selected **Dynamic PAT** for the value of the **Type** list, supply values for the following settings.

Property	Description
Addresses	Add one or more addresses or address ranges by typing them and then clicking the + button. Remove them by clicking the **X** button next to the address or address range.
Ports	Add one or more ports or port ranges by typing them and then clicking the + button. Remove them by clicking the **X** button next to the port or port range.
ICMP Echo	Select whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes. • Select **disabled** to disable ICMP echoes.
PAT Mode	Select the port address translation mode. The mode you select determines what additional properties are available. • Select **NAPT** (default) • Select **Deterministic** • Select **Port Block Allocation**
Inbound Mode	Select the inbound mode. • Select **None** to disable inbound mode. • Select **Endpoint Independent Filtering** to use endpoint independent filtering. This property is available for all PAT modes.
Mapping	Select the mapping to use. For all mappings, the default timeout value is 300 seconds, and can be modified. The range is 0 to 31536000 seconds. • Select **None** to disable inbound mode. • Select **Endpoint Independent Mapping** to use endpoint independent filtering.

Property	Description
	• Select **Address Pooling Paired** to use paired address pooling. This property is available for all PAT modes.
Client Connection Limit	Enter a number as the maximum number of client connections allowed. The default is 0, which indicates no connection limit. This property is available for all PAT modes.
Hairpin mode	Select the hairpin mode. • Select **enabled** to enable hairpin mode. • Select **disabled** to not enable hairpin mode. This property is available for all PAT modes.
Backup Addresses	Add one or more backup IP addresses by typing them and then clicking the + button. Remove them by clicking the **X** button next to the address This property is available when the deterministic PAT mode is set.
Port Block Allocation	Select numeric values for one or more of the following fields; the default is to not have a value set: • **Block Idle Timeout**. The range is 30 31536000 seconds. • **Block Life Time**. The range is 0 to 31536000 seconds. • **Block Size**. Must be 1 or greater, and less than or equal to the number of ports in the port range. • **Client Block Limit**. Must be 1 or greater. • **Zombie Timeout**. Must be 0 to 31536000 seconds. This property is available when the port block allocation PAT mode is set.
Egress Interfaces	Select whether egress interfaces are available. • Select **Disabled on** to disable egress filtering interfaces. • Select **Enabled on** to disable egress filtering interfaces.

12. Click **Save** to save the NAT source translations, or click **Save & Close** to save the NAT source translations and return to the NAT Source Translations page.

The NAT source translations are now defined and can be assigned to a rule used by a NAT policy.

Cloning NAT source translations

Cloning enables you to create an exact copy of the NAT source translations, which you can then edit.

1. Log in to the BIG-IQ ® system with your user name and password.

2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Source Translations**.
4. Select the NAT source translations to clone by selecting the check box, and then clicking **Clone**.
   A copy of the NAT source translations is created with the same name, but with -CLONE appended to the name.
5. Change the NAT source translations as needed.
6. Click **Save** to save the NAT source translations, or click **Save & Close** to save the NAT source translations and return to the NAT Source Translations page.

The cloned NAT source translations can now be assigned to a rule in a NAT policy.

Deleting NAT source translations

You delete NAT source translations that are no longer used.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Source Translations**.
4. Select check box for one or more NAT source translations to remove.
5. Click **Delete**.
6. Confirm that you want to remove the NAT source translations by clicking **Delete** in the confirmation dialog box.

The selected NAT source translations are removed.

Creating NAT destination translations

You create NAT source translations to use within a NAT policy rule.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Destination Translations**.
4. Click **Create**.
   The NAT Destination Translations - New Item screen opens.
5. Type a name for the NAT destination translations in the **Name** field.
6. In the **Description** field, type an optional description for the NAT destination translations.
7. If needed, in the **Partition** field change the default `Common` partition.
8. From the **Type** list, select the type of address translation to use. The type of address translation you select determines what additional properties are available.
   - Select **Static NAT** for static network address translation.
   - Select **Static PAT** for static network port and address translation.
9. If you selected **Static NAT** or **Static PAT** for the **Type** setting, supply values for the **Addresses** setting.
   - Add one or more addresses or address ranges by typing them in, and then clicking the + button.
   - Remove the address or address range by clicking the **X** button next to it.
10. If you selected **Static PAT** from the **Type** list, supply values for the **Ports** setting.
    - Add one or more ports or port ranges by typing them in and then clicking the + button.
    - Remove the port or port range by clicking the **X** button next to it.

11. Click **Save** to save the NAT destination translations, or click **Save & Close** to save the NAT destination translations and return to the NAT Destination Translations page.

The NAT destination translations are now defined and can be assigned to a rule used by a NAT policy.

Cloning NAT destination translations

With cloning, you create an exact copy of the NAT source translations, which you can then edit.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Destination Translations**.
4. Select the check box for the NAT destination translations to clone and then click **Clone**.
   The system creates a copy of the NAT destination translations with the same name, but with -CLONE appended to the name.
5. Change the NAT source translations as needed.
6. Click **Save** to save the NAT destination translations, or click **Save & Close** to save the NAT destination translations and return to the NAT Destination Translations page.

The cloned NAT destination translations can now be assigned to a rule in a NAT policy.

Deleting NAT destination translations

You delete NAT destination translations that are no longer used.

1. Log in to the BIG-IQ ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **NAT Destination Translations**.
4. Select one or more NAT destination translations to remove by selecting the check box for the appropriate NAT destination translations.
5. Click **Delete**.
6. Confirm that you want to remove the NAT destination translations by clicking **Delete** in the confirmation dialog box.

The system removes the selected NAT destination translations.

Managing FQDN Resolvers

About FQDN resolvers

FQDN is an acronym for a fully qualified domain name. The FQDN resolver in the Network Security Policy Editor works with the ADC DNS resolver to allow you to use fully qualified domain names where you would otherwise only be able to enter IP addresses.

You configure an FQDN resolver by clicking the device name of the FQDN resolver on the FQDN Resolvers page.

You access the DNS resolver by selecting **ADC** from the BIG-IQ menu, and then clicking **DNS Resolvers** on the left.

The BIG-IQ® system can discover FQDN support on a BIG-IP ®device version 12.0 or later, or created on a BIG-IQ system using the Network Security Policy Editor and then deployed to a BIG-IP device version 12.0 or later.

Configuring FQDN resolvers

You configure FQDN resolvers for use in your environment, including associating them with a DNS resolver.

1. Log in to the BIG-IQ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **FQDN Resolvers**.

   A list of the FQDN resolvers displays, one listed for each discovered BIG-IP® device.
4. Click the name of the BIG-IP device with an FQDN resolver to configure.
   The FQDN Resolvers - global-fqdn-policy screen opens for that BIG-IP device. Note that the device, name, and partition used by the FQDN resolver cannot be changed.
5. If needed, change the minimum refresh interval value in the **Min Refresh Interval** field.

   By default, the value of the **Min Refresh Interval** field is 60 minutes. The interval is given as the number of minutes, expressed as an integer from 10 to 46080, inclusive.
6. Select a DNS resolver from those listed in the **DNS Resolver** field.

   If no DNS resolver is listed, create one and then select it from the **DNS Resolver** field. You create DNS resolvers separately by selecting **ADC** from the BIG-IQ menu and then **DNS Resolvers**. You can have different DNS resolvers for different BIG-IP devices, unless those BIG-IP devices are clustered, in which case the DNS resolver should be the same.
7. Click **Save** to save the FQDN resolver changes, or click **Save & Close** to save the FQDN resolver changes and return to the FQDN Resolvers screen.

The FQDN resolver is now defined and can be used to resolve fully qualified domain names on the BIG-IP device.

Managing Notification Rules

About notification rules

Notification rules are accessed from within the Policy Editor and are used to notify users when a policy (firewall policy or NAT policy) is changed or when a percentage of the maximum supported configuration objects is reached. The notifications are configured using notification rules and are delivered through email, such an email is referred to as a notification email. Notification rules can be useful for administrators who wish to be notified when policies are changed, or who wish to notify others of such changes. Notifications can also be sent based on shared resources used by policies.

About Notification Email

There are two kinds of notification email that can be sent using notification rules:

- If a Policy Notify rule type is selected, the notification email lists what specified policies have changed, possibly including any changed shared resources.
- If a Limit Notify rule type is selected, the notification email lists what specified limits have been reached. The limits can include device limits, object limits or both.
  - The limit for a device is determined by the license for that device. The email contains the number of devices and the percentage of devices used, based on the maximum number of devices.
  - The limit for objects is a total number of objects for all devices being managed by the BIG-IQ system. Shared objects, such as virtual servers, only count as a single object even if they are used multiple times. The email contains the number of objects for all devices being managed and the percentage of objects used, based on the maximum number of objects.

Adding and scheduling notification rules

Use the Notification Rules screen of the Policy Editor to add and schedule a new notification rule.

Creating notification rules

1. Select Notification Rules from the navigation list to display the Notification Rules screen.
2. Click **Create** and the Notification Rules - New Item screen is displayed.
3. On the Properties tab, specify the appropriate values for the following fields.

Property	Description
Name	Specify a name for the notification. This is required.
Description	Specify a description for the notification
Email Comment	Specify the content of the email for this notification
Format	Select the format of the notification to be either **Plain Text** or **CSV**.
Rule Type	Select the type of notification rule to use. • **Policy Notify** indicates that the notification is triggered when the policy has changed. You specify the policy on the Policy Notify tab. • **Limit Notify** indicates that the notification is triggered when a limit has been reached. You specify the limit on the Limit Notify tab.

Property	Description
Email Recipients	Specify information about one or more email recipients. • In the **Name** field, specify a name for the recipient. • In the **Email Address** field, specify the email address of the recipient. To add another recipient, click the (+) plus sign and supply the **Name** and **Email Address** fields for that recipient. To remove a recipient, click the (**X**) to the right of the email recipient.

4. If you specified the **Rule Type** as **Policy Notify**, specify the appropriate values for the following fields on the Policy Notify tab.

Field	Description
Available Firewall Policies	Select the firewall policy the rule should monitor and notify you when it changes, then click **Add** . The selected policy is added to the list of firewall policies below the **Available Firewall Policies** field.
Notify on Dependent Objects	Determines whether or not dependent objects, such as shared resources, are also monitored by the rule. By default this option is selected, indicating that shared resources should also be monitored.
Available NAT Policies	Select the NAT policy the rule should monitor and notify you when it changes, then click **Add**. The selected NAT policy is added to the list of policies below the **Available NAT Policies** field.
Notify on Dependent Objects	Determines whether or not dependent objects, such as shared resources, are also monitored by the rule. By default this option is selected, indicating that shared resources should also be monitored.

You can also add firewall policies or NAT policies by dragging them from the Shared Resources area and dropping them on to the Drop Firewall Policies here to add to the list area or dropping them on the Drop NAT Policies here to add to the list area.

To delete a policy from the list, click the **X** to the right of the **Notify on Dependent Objects** option.

5. If you specified the **Rule Type** as **Limit Notify**, specify the appropriate values for the following fields on the Limit Notify tab.

Field	Description
Device Limit Notification	Select this check box to be notified when the BIG-IQ system reaches a specified limit.
Device Limit Thresholds	Specify the device limit thresholds at which a notification email is sent. A device limit is a percentage of the number of BIG-IP devices your BIG-IQ system is managing. You can set up to three limits that are each a percentage of the limit amount by modifying the percentage amount in each of the three device limit threshold fields. For example, if your BIG-IQ system is licensed to handle 10 BIG-IP devices, you would go over the threshold of 49% when 5 BIG-IP devices were being managed.
Object Limit Notification	Select this check box to be notified when a specified object limit is exceeded.
Object Limit Thresholds	Specify the object limit thresholds at which a notification email is sent. You can set up to three limits that are each a percentage of the limit amount by modifying the percentage amount in each of the three object limit threshold fields.

Field	Description
	As more operations occur with more BIG-IP devices, the number of objects in use by the BIG-IQ system grows. The maximum number of objects supported varies depending on the BIG-IP device configuration.

6. Click **Save** to save the information you have entered and to verify it is syntactically correct.
7. When finished, click **Save & Close** to save changes, release the lock, and exit the screen.

Scheduling notification rules

Once a notification rule has been created it can be scheduled. To schedule a notification rule, click the check box to the left of the rule to select it, and then click **Edit Schedule**. The Notification Rules Evaluation Interval dialog is displayed. In this dialog, specify the interval at which the schedule should run.

Editing notification rules

From the Notification Rules screen in the Policy Editor, you can edit notification rules.

1. Click the name of the notification rule to edit that rule and display the editing screen.
2. The rule is locked for editing.
3. Change the property and field values that you need to modify.
4. Click **Save** to save changes.
5. When finished, click **Save & Close** to save changes, release the lock, and exit the screen.

Deleting notification rules

From the Notification Rules screen in the Policy Editor, you can remove notification rules.

1. Select the rule to be removed by clicking the check box to the left of the rule.
2. Click **Delete**.

Managing Change Verifications

About change verifications

Use change verifications to ensure that the changes you have made to a firewall security policy in BIG-IQ® Network Security are compatible with the specified BIG-IP® devices before attempting to deploy those changes.

In some environments, the person who edits the firewall policy is not the same person as the one who deploys that policy. The person who edits the firewall policy can use the change verifications feature to make sure their changes to the firewall are compatible with the BIG-IP devices before someone else deploys those policy changes.

Firewall policy changes can be verified against either the working configuration or a configuration snapshot. In either case, the entire configuration is verified, not just the latest changes to that configuration. If the working configuration is used, make sure that while the verification is processing, other users are not changing the working configuration by changing address lists, rule lists and so on.

You create, view, and delete change verifications in the Policy Editor by selecting **Change Verifications** from the navigation list on the left. This displays the list of change verifications, including these details:

- The name of the change verification.
- The status of the change verification.
- When the change verification was created.
- What BIG-IQ system user created the change verification.
- What non-critical and critical errors were encountered during the change verification. If the number of errors is not zero, the number of errors are links that you can click for more detailed error information.

To view the properties of a change verification, click the change verification name.

To create a new change verification, click **Create**.

To delete one more change verifications, select the check box to the left of one or more change verifications and click **Delete**.

To filter which change verifications are displayed, use the Policy Editor filter fields.

Adding change verifications

You add change verifications in the BIG-IQ® Network Security Policy Editor to ensure that the changes you have made to a firewall security policy are compatible with the specified BIG-IP® devices before attempting to deploy those changes.

1. Log in to BIG-IQ Network Security with an account with the appropriate role assigned to it.

   Valid roles for adding change verifications are: Administrator, Security_Manager, Network_Security_Deploy, Network_Security_Edit, Network_Security_Manager, and Network_Security_View.
2. From the main BIG-IQ list, select **Security**, then click **Network Security** > **Policy Editor**.
3. From the Policy Editor navigation list, select **Change Verifications** to display the Change Verifications screen.
4. Click **Create**.

The Change Verifications - New Item screen opens.

5. In the **Name** field type a name for the change verification.
6. In the **Description** field type a description of the change verification.
7. Specify a source for the change verification.
   - Select **Working Config** to use the current working configuration as the source. Be sure that the working configuration does not change while the change verification process is occurring. There could be unexpected results in the verification if other users are editing and changing any part of the current configuration, including address lists, rule lists and so on.
   - Select **Snapshot** to use a specified snapshot as the source. Click **Select Snapshot** to display the list of available snapshots, click the name of the snapshot to use, and then click **Select**. The selected snapshot is displayed.
8. From **Available Devices**, select one or more devices to verify the source against.
   - Choose devices by selecting the check box to the left of each device to use for verification.
   - Choose a group of devices by selecting the check box to the left of **View by groups** to display devices organized by group, and then selecting the check box to the left of the group name to choose all devices in that group for verification.
9. Click **Verify**.
   The selected source is verified against each selected device and the change verification is shown in a list with the results. If there are errors in the verification, the number of errors are shown as links that can be clicked for more detail.

Viewing change verification properties

You view change verifications in the BIG-IQ® Network Security Policy Editor to ensure that the changes you have made to a firewall security policy are compatible with the specified BIG-IP® devices before attempting to deploy those changes.

1. Log in to BIG-IQ Network Security with an account with the appropriate role assigned to it.
   Valid roles for viewing change verifications are: Administrator, Security_Manager, Network_Security_Deploy, Network_Security_Edit, Network_Security_Manager, and Network_Security_View.
2. From the main BIG-IQ list, select **Security**, then click **Network Security** > **Policy Editor**.
3. From the Policy Editor navigation list, select **Change Verifications** to display the list of change verifications on the Change Verifications screen.
4. Click the name of a change verification to view the properties, the device used, and the number of errors.
5. Click **Cancel** to exit the Change Verifications properties screen.
   If there are errors in the change verification, the number of errors are shown as links that you can click for more detail on the error.

Change verification properties

This table lists the properties of a change verification and any associated devices.

Table 1: Change verification properties

Property	Description
Name	Name of the change verification.

Property	Description
Description	Optional description of the change verification.
User	The BIG-IQ® system user who performed the change verification.
Snapshot Name	The name of the snapshot used. If the working configuration was used instead of a snapshot, this field is blank.
Task Status	The status of the change verification task.
Start Time	When the change verification process started.
End Time	When the change verification process completed.

Table 2: Change verification device properties

Property	Description
Device	Name of the BIG-IQ device.
Verification Errors	The number of non-critical verification errors. If this number is greater than zero, it is a link which can be clicked to get more details on the errors.
Critical Errors	The number of critical errors. If this number is greater than zero, it is a link which can be clicked to get more details on the errors.
Status	The status of the change verification.

Managing Firewall Security Locks

About firewall security locks

The Locked Objects screen of the firewall security Policy Editor displays firewall and shared security objects that are locked on the BIG-IQ® system, the user who locked each object and when the lock was created. User privileges (assigned by user roles) determine what locks are visible to the user. If you have sufficient privileges, you can use the Locked Objects screen to view and remove locks.

Viewing and deleting locks

If the role associated with your account has sufficient privileges, you can display and unlock firewall security and shared security objects that are locked on the BIG-IQ® system, see which user locked the object and when the lock was created.

1. Log in to BIG-IQ Security.
2. At the top of the screen, click **Network Security**, then **Policy Editor**, and on the left, click **Locked Objects**.
3. Review the locked objects.
4. To unlock an object, select the check box to the left of that object and click **Unlock**.
5. Confirm the unlock by clicking **Delete** in the Delete Locks dialog box.

If you decide not to unlock the object, click **Cancel** instead of **Delete** in the Delete Lock dialog box.

Managing External Logging Devices

About external logging devices

You can use external logging devices to collect log files from BIG-IP® devices for viewing from a web interface launched from the F5® BIG-IQ® Centralized Management system. Currently, only the SevOne Performance Log Appliance (SevOne PLA) is supported. As part of supporting external logging devices, the BIG-IQ Centralized Management system pushes (transfers) BIG-IP device management IP- and self IP address information to the external logging device.

You need to add an external logging device to the management system before you can use it. You add, modify, or remove external logging devices from the management system by selecting **System Management** > **BIG-IQ LOGGING** > **External Logging Devices**.

Once you have added the external logging device, you need to set up authentication with the device so that you can then launch the user interface to manage the device. You perform these tasks by selecting **Network Security** > **Monitoring** > **External Logging Devices**.

Add external logging devices

You can add an external logging device to the F5® BIG-IQ® Centralized Management system to collect log files from BIG-IP devices. Currently only the SevOne Performance Log Appliance (SevOne PLA) is supported.

1. Log in to the BIG-IQ Centralized Management system with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. In the list on the left, click **BIG-IQ LOGGING** and then click **External Logging Devices**.
4. Click **Add**.
   The Add External Logging Device screen opens.
5. Type a name and an optional description for the external logging device.
6. Type the connection information for the SevOne PLA device.
   a) Specify the IP address of the external logging device in the **IP Address** field. If the external logging device is also a query server, select the **Use as query server** check box to the right of the **IP Address** field. Query servers are members of the SevOne PLA logging device cluster that process user queries.
   b) In the **User Name** field, specify the user name of the account used to securely access the SevOne PLA device.
   c) In the**Password** field, specify the password for the selected account .
   d) Once the connection information is complete, click **Test** to verify that the connection information can be used to access the SevOne PLA device. If the test is successful, a green check mark and the text `Connection Established` are shown in the Test Connection row.
7. Type the query server information for the device.
   a) Add the first query server IP address in the **IP Address** field.
   b) Add additional query servers by clicking the + to the right of the description and then supplying an IP address for each in the created **IP Address**fields.
   c) Remove query servers by clicking the **X** to the right of the query server to remove.
8. Optionally, select the **Enabled** check box in the **Status** setting to allow scheduling for data pushes from the BIG-IQ Centralized Management system to the SevOne PLA device, and to enable editing for the other fields in this area.

- To have the data pushed every day,for the **Push Frequency** setting, select **Daily**.
- To push the data once a month, in the **Push Frequency** setting, select **Monthly**, and then select the day of the month.
- To push the data each week, in the **Push Frequency** setting, select **Weekly** and then select one or more days of the week.

9. Specify when to begin, and optionally end, the scheduled data push.

   Select the date and time using the calendar tool, or type the date, hour, and minute to use. By default, **No End Date** is selected and must be cleared to specify an end time.

10. Click **Push Now** in the **Last Push Date** setting to manually push the data to the external logging device.
11. Click **Add** to save your changes.

    After you successfully add your changes, the screen shows the **Last Push Date** setting. You can use the **Push Now** button in that row to manually push data to the external logging device.

The external logging device is added to the BIG-IQ Centralized Management system.

You can now configure authentication with the external logging device so that you can manage logs using the SevOne PLA user interface.

Modify external logging devices

You can modify settings for external logging devices used with the F5® BIG-IQ® Centralized Management system when the computing environment changes.

1. Log in to the BIG-IQ Centralized Management system with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. In the list on the left, expand **BIG-IQ LOGGING** and then click **External Logging Devices**.
4. Click the name of the external logging device to modify.
   The External Logging Device (*device-name*) screen opens.
5. Modify the description for the external logging device, as needed.
6. Modify the connection information for the SevOne PLA device, as needed.
   a) Specify the IP address of the external logging device in the **IP Address** field. If the external logging device is also a query server, select the **Use as query server** check box to the right of the **IP Address** field. Query servers are members of the SevOne PLA logging device cluster that process user queries.
   b) Specify the user name of the account used to securely access the SevOne PLA device in the **User Name** field.
   c) Specify the password for the selected account in the**Password** field.
   d) Click **Test** once the connection information is complete to verify that the connection information can be used to access the SevOne PLA device. If the test is successful, a green check mark and the text `Connection Established` are shown in the **Test Connection** setting.
7. Modify the query server information for the device, as needed.
   a) Add the first query server IP address in the **IP Address** field.
   b) Add additional query servers by clicking the + to the right of the description and then supplying an IP address for each in the created **IP Address**fields.
   c) Remove query servers by clicking the **X** to the right of the query server to be removed.
8. Optionally, select the **Enabled** check box in the **Status** setting to allow data pushes from the BIG-IQ Centralized Management system to the SevOne PLA device to be scheduled and to enable the other fields in this area for editing.

   - Select **Daily** in the **Push Frequency** setting to have the data pushed every day.

- Select **Monthly** in the **Push Frequency** setting and then select the day of the month to push the data once a month.
- Select **Weekly** in the **Push Frequency** setting and then select one or more days of the week to push the data each week.

9. Specify when to begin, and optionally end, the scheduled data push.

   Select the date and time using the calendar tool, or type the date, hour, and minute to use. By default, **No End Date** is selected and must be cleared to specify an end time.

10. Click **Save** to save your changes.

    You can use the **Push Now** button in the **Last Push Date** setting to manually push the data to the external logging device.

The external logging device is updated in the BIG-IQ Centralized Management system.

Depending on the changes you made, you may need to configure authentication with the external logging device so that you can manage logs using the SevOne PLA user interface.

Remove external logging devices

You can remove external logging devices from the F5® BIG-IQ® Centralized Management system when those devices are no longer needed.

1. Log in to the BIG-IQ Centralized Management system with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. In the list on the left, click **BIG-IQ LOGGING** and then click **External Logging Devices**.
4. Click the check box to the left of the name of the external logging device to remove.
5. Click **Remove**.
   A dialog box opens, asking you to confirm that you want to remove the device.
6. Click **Delete** in the dialog box.

The device is removed and is no longer seen in the list of devices.

Request authentication token for external logging devices

You must add the external logging device to theF5® BIG-IQ® Centralized Management system before you can request an authentication token for it.

You request an authentication token for the external logging device so you can access the web interface to the external logging device and the BIG-IP® device logs stored there.

1. Log in to the BIG-IQ Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. From the menu at the top click **Monitoring**, and then on the left click **External Logging Devices**.
4. Request an authentication token by clicking the link displayed in the Auth token column for the appropriate external logging device.

   - If you do not have an authentication token to access the SevOne PLA, you request it by clicking **Request Auth Token**.
   - If you want to replace an existing authentication token, you request a new one by clicking **Manage Auth Token**.

   The Manage SevOne PLA Authentication Token dialog box opens.
5. Type the user name and password for an account with privileges to use the web interface to the SevOne PLA device.

6. Click **Request Token**.
   A new authentication token is sent from the SevOne PLA device, and the new token string is displayed.
7. Click **Save** to save the token and close the dialog box.

The authentication token is created and is used when you launch the SevOne PLA user interface.

Delete authentication token for external logging devices

You must have an authentication token before you can delete it.

You delete authentication tokens with an external logging device when then are no longer needed or should be replaced.

1. Log in to the F5® BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. From the menu at the top click **Monitoring**, and then on the left click **External Logging Devices**.
4. Click the check box to the left of the external logging device that should have the authentication token removed.
5. Click **Delete Auth Token**.

You no longer have an authentication token for the external logging device. You will need to request a new authentication token to communicate with the external logging device user interface.

Access external logging devices

You must add the external logging device to the F5® BIG-IQ® Centralized Management system before you can access it. You must also have authenticated your account with the external logging device.

You access and launch the web interface to the external logging device to access the BIG-IP® device logs stored there.

1. Log in to the BIG-IQ Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. At the top click **Monitoring**, and then on the left click **External Logging Devices**.
4. In the External Device column, click **Launch** for the appropriate external logging device.
   The SevOne PLA user interface opens.

You can now use SevOne PLA to review the BIG-IP device logs.

Monitoring Firewall Rules

About firewall rule monitoring

In BIG-IQ™ Network Security, you can monitor:

- Firewall rule statistics, such as the number of times inbound network traffic matches a firewall rule on a BIG-IP™ device (also referred to as a firewall rule hit count) as well as the rule overlap status.
- Firewall rule compilation statistics for a set of rules associated with a firewall context on a BIG-IP device.

You access this firewall rule monitoring by selecting **Network Security** from the BIG-IQ menu and then clicking **Monitoring**.

You can generate reports about firewall rules by selecting **Network Security** from the BIG-IQ menu and then clicking **Policy Editor**, and then selecting **Firewall Rule Reports**.

Monitoring firewall rule statistics and hit counts

You can monitor firewall rule statistics and hit counts on one or more BIG-IP™ devices using Network Security monitoring.

Note: *Firewall rule statistics are collected for the rules in the enforced policy associated with a firewall, but not the rules in a staged policy.*

Note: *If a virtual server, route domain or self IP is created using the BIG-IQ™ system, firewall statistics cannot be collected until the changes are deployed to the device and reimported.*

1. Log in to the BIG-IQ system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Monitoring**.
4. Click **Firewall Rule Statistics**.
   The Firewall Rule Statistics screen opens and displays a list of firewall contexts, including their name, partition, type, and on what BIG-IP device they occur.
5. Click the name of the firewall context to monitor.
6. The Firewall Rule Statistics page for that firewall context displays.

   The following information is listed in the named columns for each firewall rule on the BIG-IP device:

   - Rule Name specifies the name of the rule used in the policy. If not listed, the rule is not running.
   - Rule List Name specifies the name of the rule list if the rule is in a rule list.
   - Rule specifies the name of the rule within a rule list. If the rule is not in a rule list, this field is blank.
   - Overlap Status specifies whether the rule overlaps with another rule.
   - Hit Count specifies the number of times the rule has been matched.
   - Last Hit Time specifies when the rule was last matched.

Monitoring firewall rule compilation statistics

You can monitor rule compilation statistics on one or more BIG-IP™ devices using Network Security monitoring. This information is similar to what is displayed when using the `tmsh show security firewall container-stat` command.

Note: *If a firewall context references a policy that is both staged and enforced, there will be two entries in the compilation statistics: one for the enforced policy and one for the staged policy.*

1. Log in to the BIG-IQ system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Monitoring**.
4. Click **Firewall Compilation Statistics**.
   The Firewall Compilation Statistics screen opens and displays the list of BIG-IP devices managed by the BIG-IQ system, including their network name, IP address, and BIG-IP device version.
5. Click the name of the BIG-IP device to monitor.
6. The Firewall Compilation Statistics page for that BIG-IP device displays.

   Depending on the version of the BIG-IP device, the following information, or a subset of this information, may be listed in the named columns for the one or more firewall rules within the specified firewall context on the BIG-IP device:

   - **Context Name** specifies the context name associated with the one or more rules, such as `/Common/global-firewall-rules`.
   - **Context Type** specifies the firewall context type associated with the one or more rules, such as global or self IP.
   - **Policy Name** specifies the name of the policy associated with the one or more rules.
   - **Policy Type** specifies type of policy associated with the one or more rules, such as enforced or staged.
   - **Rule Count**Specifies the number of rules compiled for this BIG-IP device context, such as 30. This count includes rules in rule lists as well as rules that are not in rule lists.
   - **Compile Duration** specifies the amount of time required to compile the rules, expressed as `hours:minutes:seconds`.
   - **Overlap Check Duration** specifies the amount of time required to check overlapping rules, expressed as `hours:minutes:seconds`.
   - **Size** specifies the size of the compiled rules in bytes.
   - **Max Memory** specifies the maximum amount of memory consumed by the rules in bytes.
   - **Activation Time** specifies when the rules are activated and available for use.

Managing Firewall Rule Reports

About firewall rule reports

You can generate different types of firewall rule reports for selected BIG-IP® devices in either CSV or HTML format. These reports capture information similar to that gathered using the firewall rule monitoring. The types of reports you can generate include:

- Stale Rule Report. Creates a report on firewall rules that are not being used on the BIG-IP device.
- Overlap Status Stats Report. Creates a report on firewall rules that are overlapping on the BIG-IP device.
- Compilation Status Report. Creates a report on the compilation of firewall rules on the BIG-IP device.

Creating firewall rule reports

You create firewall rule reports to capture statistics about firewall rules in a report format.

1. Log in to the BIG-IQ® system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **Firewall Rule Reports**.
4. Click **Create**.
   The Firewall Rule Reports - New Item screen opens.
5. Type a name for the report in the **Name** field.
6. Type an optional description for the report in the **Description** field.
7. Select a report type from those listed in the **Report Type** field.

   You can generate these types of reports::

   - Stale Rule Report
   - Overlap Status Stats Report
   - Compilation Status Stats Report

   If the **Stale Rule Report** report type is selected, the screen displays the Stale Rule Criteria property, otherwise that property is not displayed.
8. If you select **Stale Rule Report**, you can refine the report using the options listed in the **Stale Rule Criteria** setting:

   - To specify that the report should include only rules with a hit count less than the number specified, select **Rules with count less than** and specify a number in the provided field.
   - To specify that the report should include only rules that have not been hit since the date specified, select **Rules that haven't been hit since** and specify a date in the provided field.
9. From the **Available Devices** setting, select the BIG-IP devices or device group to use for the report:

   - Select **Group** and select a group of BIG-IP devices from the list.
   - Select **Device** and select individual BIG-IP devices by moving them from the **Available** list to the **Selected** list.
10. Save the report:

- Select **Save** to save the report. The system displays the Firewall Rule Reports page for that one report, and generates the report data.
- Select **Save & Close** to save the report. The system displays the Firewall Rule Reports page that lists all reports, and generates the report data.

11. Select the format for the report:

- Select **CSV Report** to have the report formatted as a CSV file.
- Select **HTML Report** to have the report formatted as an HTML file. The HTML file is displayed in the Web browser when complete.

You can save or print these reports.

Deleting firewall rule reports

You can delete firewall rule reports that are no longer needed.

1. Log in to the **BIG-IQ** system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **Firewall Rule Reports**.
4. Select one or more reports to delete, and click **Delete**.
   The reports are deleted from the list on the Firewall Rule Reports screen.

Managing Firewall Policies in BIG-IQ Network Security

About firewall policies in BIG-IQ Network Security

A *firewall policy* is a set of rules and/or rule lists. BIG-IP® network firewalls use policies to specify traffic-handling actions and to define the parameters for filtering network traffic. You can assign rule lists, or a policy to a firewall. Policies facilitate the assigning of a common collection of rules consistently across multiple firewalls.

The network software compares IP packets to the criteria specified in policies. If a packet matches the criteria, then the system takes the action specified by the policy. If a packet does not match any rule in the policy, the software accepts the packet or passes it to the next policy, rule, or rule list.

In BIG-IQ® Network Security, the Policies list displays the policies available for assignment to firewalls.

You can configure firewall policies as enforced or staged:

- An *enforced* policy refers to a policy whose actions are executed. Actions include: accept, accept decisively, drop, and reject.

  You are restricted to assigning a single, enforced policy on any specific firewall.
- A *staged* policy refers to a policy that is evaluated but policy actions are not enforced. All activity is logged.

  You are restricted to assigning a single, staged policy on any specific firewall. You can have rule lists assigned to a firewall (in the enforced area) and have a configured staged policy on that firewall. You cannot have rule lists in the staged area.

You can stage a firewall policy first and then examine logs to determine how the policy has affected traffic. Then you can determine the timing for turning the policy from staged to enforced.

Firewall policies can contain any combination of rules and rule lists. Policies cannot contain other policies. You can re-order rules within a policy.

Note: *The BIG-IQ® Network Security system is aware of functionality implemented in one BIG-IP version but not in another. In terms of firewall policies, this means that you are prohibited from dropping a policy onto a firewall on a BIG-IP device that does not have the software version required to support it.*

Filtering policies

To filter the system interface to display only those objects related to a selected policy, hover over the policy name, right-click and then click **Filter 'related to'**. The interface is filtered and a count appears to the right of each object type. The frame to the right provides its own filter field where you can enter text and click on the filter icon to constrain the display to those items that match the filter.

Creating firewall policies

To fine tune your network firewalls, you can configure policies and assign them to firewalls using the Firewall Policies screen Rules & Rule Lists settings.

1. Click **Policy Editor**.
2. On the left, click **Firewall Policies** and click **Create** to open the Firewall Policies - New Item screen.
3. Click **Properties** and complete the properties fields as required.

All boxes outlined in gold are required fields.

Option	Description
Name	User-provided name for the policy. This field is editable when creating or cloning a policy, and read-only when editing a policy.
Description	Optional description for the policy.
Partition	Although it is pre-populated with `Common` (default), you can set the partition when creating or cloning policies by typing a unique partition name. ***Note:*** *The partition with that name must already exist on the BIG-IP device.* No whitespace is allowed in the partition name. No editing of the partition is allowed.
Available Devices	Select the BIG-IP device to use, if your firewall policy is referenced by a self-IP context with a static (non-floating) IP address. If your firewall policy is not referenced by a static self-IP context, this property should not be set. You may be directed to set this property as a result of an evaluation critical error issued after performing a configuration evaluation prior to deployment. This property must be set for the peer BIG-IP device that is part of a DSC cluster managed by the BIG-IQ system. You select the BIG-IP device to use by moving it from the **Available** list to the **Selected** list using the arrow buttons. You can filter the list of available BIG-IP devices using the filter field at the top of the **Available** list. Moving a BIG-IP device that is part of a cluster to the **Selected** list will cause the other member of the cluster to move to that list as well.

4. Click **Rules & Rule Lists**, and then click either:
   - **Create Rule** to create rules.
   - **Add Rule List** to add rule lists.
5. Click **Save** to save the firewall policy, or click **Save & Close** to save the firewall policy and return to the Firewall Policies screen.

A new firewall policy is added.

Managing firewall policies

To fine tune your network firewalls, you can edit policies, create or edit rules, and add rule lists. You can also reorder rules in firewall policies. You cannot edit rule lists or reorder rules within rule lists.

1. Click **Policy Editor**.
2. On the left, click **Firewall Policies** to see the list of firewall policies.
3. Click the name of the firewall policy to edit.
4. Click **Properties** and review or change the properties fields as needed.

   All boxes outlined in gold are required fields.

Option	Description
Name	User-provided name for the policy. This field is editable when creating or cloning a policy, and read-only when editing a policy.
Description	Optional description for the policy.
Partition	Although it is pre-populated with `Common` (default), you can set the partition when creating or cloning policies by typing a unique partition name.

Option	Description
	Note: *The partition with that name must already exist on the BIG-IP device.* No whitespace is allowed in the partition name. No editing of the partition is allowed.
Pin Policy to Device(s)	Select the BIG-IP devices to be pinned to this policy, if needed. Pinning a BIG-IP device to a policy enables the policy to be deployed even if it is not associated with a firewall context for that device. If you have a self IP context with a static (non-floating) IP address, you may be required to assign the device depending on you cluster deployment settings. For example, this property must be set for a peer BIG-IP device that is part of a DSC cluster managed by the BIG-IQ system. You may be directed to set this property as a result of an evaluation critical error. You select the BIG-IP device to use by moving it from the **Available** list to the **Selected** list using the arrow buttons. You can filter the list of available BIG-IP devices using the filter field at the top of the **Available** list. Moving a BIG-IP device that is part of a cluster to the **Selected** list will cause the other member of the cluster to move to that list as well.

5. Click **Rules & Rule Lists**, and edit the existing rule list or click either:
   - **Create Rule** to add rules.
   - **Add Rule List** to add rule lists.
6. Click **Save** to save your changes.
7. When you are finished, click **Save & Close** to save your edits, and return to the Firewall Policies screen.

The edited firewall policy appears on the **Firewall Policies** screen.

Cloning firewall policies

Cloning creates an exact copy with a different name. It enables you to quickly and easily create firewall policies tailored to address any unique aspects of your network firewall environment. When you clone a firewall policy, you create an exact copy of the policy which you can then edit to address any special considerations.

Users with the roles of Network_Security_View or Network_Security_Deploy cannot clone policies.

1. Click **Policy Editor**.
2. On the left, click **Firewall Policies** to see the list of firewall policies.
3. Select a firewall policy in the list using the check box on the left and click **Clone** to copy and modify an existing firewall policy.
4. Click **Properties** and complete the properties fields as required.

   All boxes outlined in gold are required fields.

Option	Description
Name	User-provided name for the policy. This field is editable when creating or cloning a policy, and read-only when editing a policy.
Description	Optional description for the policy.
Partition	Although it is pre-populated with `Common` (default), you can set the partition when creating or cloning policies by typing a unique partition name.

Option	Description
	Note: *The partition with that name must already exist on the BIG-IP device.*
	No whitespace is allowed in the partition name. No editing of the partition is allowed.
Available Devices	Select the BIG-IP device to use, if your firewall policy is referenced by a self-IP context with a static (non-floating) IP address. If your firewall policy is not referenced by a static self-IP context, this property should not be set. You may be directed to set this property as a result of an evaluation critical error issued after performing a configuration evaluation prior to deployment. This property must be set for the peer BIG-IP device that is part of a DSC cluster managed by the BIG-IQ system. You select the BIG-IP device to use by moving it from the **Available** list to the **Selected** list using the arrow buttons. You can filter the list of available BIG-IP devices using the filter field at the top of the **Available** list. Moving a BIG-IP device that is part of a cluster to the **Selected** list will cause the other member of the cluster to move to that list as well.

5. Click **Rules & Rule Lists**, and then click either:
   - **Create Rule** to create rules.
   - **Add Rule List** to add rule lists.
6. Click **Save** to save the firewall policy, or click **Save & Close** to save the firewall policy and return to the Firewall Policies page.

The cloned policy appears in the Firewall Policies screen. In an HA configuration, the cloned policy appears on the standby BIG-IQ® system as soon as it is saved.

Reordering rules in firewall policies

Using the Firewall Policies screen, you can reorder rules in firewall policies to optimize your network firewall policies. You cannot edit rule lists or reorder rules inside rule lists.

1. Click **Policy Editor**.
2. On the left, click **Firewall Policies** to see the list of firewall policies.
3. Click the name of the firewall policy to edit.
4. Click **Rules & Rule Lists**.
5. To reorder rule lists or rules, simply drag-and-drop them until they are in the correct order.

   You can also right-click a rule name and select among the ordering options.
6. Click **Save** to save your changes.
7. When you are finished, click **Save & Close** to save your edits, and return to the Firewall Policies screen.

Deleting firewall policies

You can remove obsolete firewall policies to keep network firewalls up-to-date.

If a firewall policy is in use, you cannot remove it.

To see where a firewall policy is used, right click the firewall policy name and click **Filter 'related to'** . The BIG-IQ system displays a count of where the policy is used in the list to the left.

1. Click **Policy Editor**.

2. On the left, click **Firewall Policies** to see the list of firewall policies.
3. Select the firewall policy to be deleted using the check box to the left of the firewall policy.
4. Click **Delete** and then confirm the permanent removal in the popup dialog box.

The policy is deleted and no longer occurs in the list of firewall policies.

About managing firewall policies using snapshots

It is possible to introduce errors during the editing of the working-configuration set. In some cases, you might not detect these errors immediately. When you discover these errors, you might want to roll back to a previous state as quickly as possible to restore service. Then, you can triage to discover the root causes of any errors.

In one scenario, you might perform multiple emergency deployments in an attempt to fix a problem. If such attempts did not fix the issue, you might want to roll back to the most stable state prior to where you first saw the problem.

In another scenario, you might want to roll back after importing a device. For example, an administrator might import a device and as part of the import process, decide to overwrite the objects stored in the BIG-IQ® database. Subsequently, the administrator decides that the import was a mistake and wants to roll back to the state of the objects before the import.

You can address all of these scenarios by restoring from a snapshot.

The BIG-IQ system provides the ability to create snapshots in these ways:

- During discovery, the BIG-IQ system takes a snapshot of the working-configuration set on the device.
- During deployment, BIG-IQ Network Security takes a snapshot when you create an evaluation.
- At any time, you can create a user-defined snapshot using the Snapshot and Restore - Network Security screen in Change Management.

Managing Security Reports

About security reporting

Reporting for BIG-IQ® Network Security

You can use BIG-IQ® Network Security Reporting to view reports for managed BIG-IP® devices that are provisioned for Application Visibility and Reporting (AVR). Reports can be for a single BIG-IP device or can contain aggregated data for multiple BIG-IP devices (that are of the same BIG-IP device version).

Network Firewall, DoS and IP Intelligence reports can be created. Analytic reports provide detailed metrics about application performance such as transactions per second, server and client latency, request and response throughput, and sessions. Metrics are provided for applications, virtual servers, pool members, URLs, specific countries, and additional detailed statistics about application traffic running through one or more managed devices. You can view the analytics reports for a single device, view aggregated reports for a group of devices, and create custom lists to view analytics for only specified devices.

Reporting for BIG-IQ® Web Application Security

You can use BIG-IQ® Web Application Security Reporting to view reports for managed BIG-IP® devices that are provisioned for Application Visibility and Reporting (AVR). Similar to the availability of the AVR reporting on a single device, you have the ability to get visibility into application traffic passing through a single managed BIG-IP device or an aggregated system (aggregated data for multiple BIG-IP devices.

You can generate reports and charts in the following areas:

- Application. You can view information about requests based on applications (iApps), virtual servers, security policies, attack types, violations, URLs, client IP addresses, IP address intelligence (reputation), client countries, severities, response codes, request types, methods, protocols, viruses detected, usernames, and session identification numbers.
- Anomalies. You can view charts of statistical information in graphs about anomaly attacks, such as brute force attacks and web scraping attacks. You can use these charts to evaluate traffic to the web application, and to evaluate the vulnerabilities in the security policy.
- DoS. If you have configured DoS protection on the BIG-IP system, you can view charts and reports that show information about DoS attacks and mitigations in place on the system.

Managing Virtual Servers in Shared Security

About virtual servers

On BIG-IP® devices, virtual servers can have security objects, such as DoS profiles, SSH profiles, or log profiles, attached to them.

The BIG-IQ® Centralized Management system can successfully discover only BIG-IP devices that are using supported profile types. If the system attempts to discover a BIG-IP device that is using an unsupported type of profile:

- You might see an invalid profile error during discovery.
- The BIG-IQ Centralized Management system does not successfully discover the BIG-IP device.

You create virtual servers using the ADC service and then use them from within Shared Security. Select **ADC** from the BIG-IQ menu and click **Virtual Servers** to create a virtual server.

Edit virtual servers

The virtual server must exist to be edited.

You can modify virtual servers within Shared Security to manage the DoS profiles, SSH profiles, and log profiles that might be part of a virtual server.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Virtual Servers**.
4. On the Virtual Servers screen, click the name of the virtual server to edit.
5. On the Virtual Servers - *virtual server name* screen, select the profile to use with the virtual server. Note that all other properties cannot be changed.

Property	Description
DoS Profile	Select the DoS profile to use. You define DoS profiles using the Shared Security DoS Profiles screen.
SSH Profile	Select the SSH profile to use. You define SSH profiles using the Shared Security SSH Profiles screen.
Log Profiles	Select the one or more log profiles to use. You define log profiles using the Shared Security Logging Profiles screen.

6. Click **Save & Close** to save changes and return to the Virtual Servers screen.

The virtual server is updated with your changes. You use the ADC service to make other changes to the virtual server by selecting **ADC** from the BIG-IQ menu.

Managing DoS Profiles in Shared Security

About DoS profiles

A *denial-of-service attack (DoS attack)* makes a victim's resource unavailable to its intended users, or obstructs the communication media between the intended users and the victimized site so that they can no longer communicate adequately. Perpetrators of DoS attacks typically target sites or services, such as banks, credit card payment gateways, and e-commerce web sites.

Using BIG-IQ® Shared Security, you can configure profiles to help prevent network, SIP, and DNS DoS and DDoS attacks and to detect and protect against DoS (Denial of Service) attacks aimed at the resources that are used for serving the application (the web server, web framework, and the application logic).

HTTP-GET attacks and page flood attacks are typical examples of application DoS attacks. These attacks are initiated either from a single user (single IP address) or from thousands of computers (distributed DoS attack), which overwhelms the target system. In page flood attacks, the attacker downloads all the resources on the page (images, scripts, and so on) while an HTTP-GET flood repeatedly requests specific URLs regardless of their place in the application.

DoS attack detection and prevention:

- Detects and automatically drops packets that are malformed or contain errors.
- Logs unusual increases in packets of any type, including packets that are malformed, packets that contain errors, or packets of any other type that appear to rapidly increase.

Creating DoS profiles

You can create a DoS profile and configure the circumstances under which the system considers traffic to be a DoS attack and how the system handles a DoS attack.

1. Click **DoS Profiles**.
2. In the DoS Profiles screen, click **Create**.
3. In the DoS Profiles - New Item screen, add and set the properties as appropriate.

Property	Description
Name	Required. Specify a unique name for the DoS profile.
Description	Specify an optional description for the DoS profile.
Partition	Required. Specify the partition to which the DoS profile belongs. Although this field is pre-populated with **Common** (default), you can set the partition when creating DoS profiles by typing a unique name for the partition. ***Note:*** *The partition with that name must already exist on the BIG-IP® device. No whitespace is allowed in the partition name.*

4. Select the check box to the right of one or more protection types. When you select a protection type, the system adds a tab dynamically. For example, when you select **Application Security**, the system adds a tab with that label.

Option	Description
Application Security	When enabled, protects your web application against DoS attacks. Your virtual server must include an HTTP profile to use this feature. Click the **Application Security** tab to configure the protection type. Supply or modify any necessary property values. For configuration setting details, consult the relevant following section.
Protocol DNS	When enabled, protects your DNS server against DoS attacks. Note that your virtual server must include a DNS profile to work with this feature. Click the **Protocol DNS** tab to configure the protection type. Supply or modify any necessary property values. For configuration setting details, consult the relevant following section.
Protocol SIP	When enabled, protects against SIP DoS attacks. Note that your virtual server must include a SIP profile to work with this feature. Click the **Protocol SIP** tab to configure the protection type.For configuration setting details, consult the relevant following section.
Network	When enabled, protects your server against network DoS attacks. Click the **Network** tab to configure the protection type. Supply or modify any necessary property values. For configuration setting details, consult the relevant following section.

5. When finished, click **Save** to save the DoS profile, or click **Save & Close** to save the DoS profile and return to the DoS Profiles screen.

The new DoS profile is added to the list of profiles.

Configuring for Application Security

You can configure the conditions under which the system determines your application is under a DoS attack, and how the system reacts to a suspected attack. Your virtual server must include an HTTP profile to use this feature.

1. In Shared Security, click **DoS Profiles**.
2. In the DoS Profiles screen, click the profile name to configure.
3. On the Properties tab, select the check box to enable Application Security.
4. Select the **Application Security** tab.
5. Specify the settings as described for each tab. The system saves settings as you enter them.

Tab	Description of Settings
General Settings tab	• **Trigger iRule**. Enable this setting if you have an iRule that manages DoS events in a customized manner. When enabled, specifies that the system activates an Application DoS iRule event. Enable this setting if you write an iRule that tells the system how to manage after a DoS attack. The default is disabled. • **IP Whitelist**. Specifies IP addresses, including subnet masks, that the system considers legitimate and does not examine when performing DoS prevention. Note that after you add an IP address to this whitelist, the system automatically adds this IP address to all Anomaly Detection whitelists, and to the IP Address Exceptions list on the BIG-IP device. To add an IP address to the whitelist, type an IP address in the text box and click **Add**.

Tab	Description of Settings
	• **Geolocation Whitelist**. Overrides the DoS profile's Geolocation Detection Criteria threshold settings by selecting countries from which to allow traffic during a DoS attack. To add countries to the whitelist, select from the Country list and click **Add**. • **Geolocation Blacklist**. Overrides the DoS profile's Geolocation Detection Criteria threshold settings by selecting countries from which to block traffic during a DoS attack. To add countries to the blacklist, select from the Country list and click **Add**.
Proactive Bot Defense tab	• **Operation Mode**. Specifies the conditions under which the system detects and blocks bots. Select **Off**, **During Attacks** or **Always**. If **Off** is selected, no other settings are displayed on this tab. • **Block requests from suspicious browsers**. Strengthens the bot defense by blocking suspicious browsers. The system completely blocks highly suspicious browsers; it challenges with CAPTCHA moderately suspicious browsers. • **Grace Period**. Gives time for the system to validate that browsers are not bots. During this period, the system does not block requests that were not validated. • **Cross-Domain Requests**. You can add additional security by allowing only configured domains to reference resources of the site. From the list, select one of the options. Domains can also be configured after selecting one of the **Cross-Domain Requests** options. • **URL Whitelist**. Specifies excluded URLs. Proactive Bot Defense will not block requests to these URLs, although requests may still be blocked by the TPS-based / Stress-based attack mitigation. To add URLs to the whitelist, type a URL in the text box, and click **Add**.
Bot Signatures tab	• **Bot Signature Check**. Select **Enabled** or **Disabled**. You cannot disable the **Bot Signature Check** property while **Proactive Bot Defense** is enabled. To disable the **Bot Signature Check** property, you must disable **Proactive Bot Defense** . • **Bot Signature Categories**. There are two category lists that are handled similarly: **Malicious Categories** and **Benign Categories**. For either category, select **None**, **Report** or **Block**. The selected setting is then applied to all the listed categories. The categories can also be individually changed to another value. If you change them individually, the value for the **Malicious Categories** or **Benign Categories** changes to **Custom Configuration**. A user cannot set all categories to **None** and keep **Proactive Bot Defense** enabled. • **Bot Signatures List**. Specifies bot signatures that are available and disabled. Use the arrow buttons to move bot signatures between the Available Signatures and the Disabled Signatures lists.
TPS Based Detection tab	In TPS-based detection mode, if the ratio of the transaction rate detection interval to the transaction rate history interval is greater than the specific percentage you configure on this tab (the TPS increased by percentage), the system detects that the URL/site is under attack, or the IP address/geolocation is attacking. To stop the attack, the system blocks some, or all, requests from the detected IP address/ geolocation and/to the attacked URL/site, depending on the configuration of the DoS profile. • **Operation Mode**. Specifies how the system reacts when it detects an attack.

Tab	Description of Settings
	• **Source IP-based**. Specifies the criteria that determine when the system treats the IP address as an attacker. If the system reaches these thresholds, it prevents further attacks by limiting the number of requests per second to the history interval. The system does not return the blocking response page. • **Geolocation-based**. Specifies that if both criteria are met, the system treats the country as an attacker. If the system reaches these values, it prevents further attacks by limiting the number of requests per second to the history interval. The system does not return the blocking response page. The settings exclude blacklisted and whitelisted geolocations. • **URL-based**. Specifies the criteria that determine when the system determines that a URL is under attack. If requests for URLs meet either of the conditions in these settings, the system prevents further attacks by limiting the number of requests per second to the history interval. The system does not return the blocking response page. • **Site-wide**. Specifies the criteria that determine when the system determines an entire website is under attack. The system prevents further attacks by limiting the number of requests per second to the history interval. The system does not return the blocking response page. • **Prevention Duration**. Specifies the time spent in each mitigation step before moving to the next mitigation step.
Stress Based Detection tab	In this tab, configure the system to prevent DoS attacks based on the server's health condition. An attack is detected if the system finds the server to be under stress and either of the TPS thresholds are crossed, or the system found a behavioral anomaly. • **Operation Mode**. Specifies how the system reacts when it detects an attack. • **Source IP-based**. Specifies the criteria under which the system treats the IP address as suspicious (suspects the IP address is an attacker). If the system detects an attack according to the detection criteria, IP rate limiting will be done on the suspicious IP addresses. The system prevents the attack by limiting the number of requests per second. The system does not return the blocking response page. The system considers an IP as an attacking entity if either conditions occurs. • **Geolocation-based**. Specifies the conditions under which the system considers requests from a country as suspicious. The system performs mitigation methods on traffic from suspicious countries if at least one By Geolocation mitigation method is enabled and both conditions are met. The settings exclude blacklisted and whitelisted geolocations. • **URL-based**. Specifies the criteria that determine when the system suspects the URL to be attacked. If an attack is detected according to the detection criteria, URL rate limiting will be done on the suspicious URLs. The system prevents the attack by limiting the number of requests per second. The system does not return the blocking response page. The system considers a URL as an attacked entity if either condition occurs. • **Site-wide**. Specifies the conditions under which the entire site is considered suspicious and provides mitigation options. • **Prevention Duration**. Specifies the time spent in each mitigation step before moving to the next mitigation step.
Heavy URL Protection tab	• **Heavy URL Protection**. Select **Enabled** to protect heavy URLs during DoS attacks. • **Automatic Detection**. Select **Enabled** to automatically detect heavy URLs of the application, in addition to the URLs entered manually.

Tab	Description of Settings
	• **Latency Threshold**. If **Automatic Detection** is enabled, set the **Latency Threshold** field to be the number of milliseconds for the system to use as the threshold for automatically detecting heavy URLs. The default value is `1000` milliseconds. Click **Set default threshold** to reset the value to 1000. • **Heavy URLs**. Enables you to configure a list of heavy URLs to protect in addition to the automatically detected ones. Type a URL in the text box, and click **Add**. • **Ignored URLs**. Enables you to configure a list of URLs which are excluded from automatic detection as heavy URLs. The system supports wildcards. Type a URL in the text box, and click **Add**.
Record Traffic tab	This tab enables the recording of traffic (by performing a TCP dump) when a DoS attack is underway, to diagnose the attack vectors and attackers, observe whether and how it was mitigated, and draw conclusions for changing the DoS profile configuration. • **Record Traffic During Attacks**. The system records traffic during DoS attacks on the virtual server in which it detected the attack. You can collect the TCP dump files into the QuickView file so that F5 support can use it for solving customer cases. The files have a pcap extension and are located in the following path on the BIG-IP device: `/shared/dosl7/tcpdumps` The default is disabled. Note that the system records SSL traffic encrypted. Select **Enabled** to specify that the system record traffic when a DoS attack is underway. • **Maximum TCP Dump Duration**. Displays the maximum time, in seconds, for one dump cycle. Legal values are between 1 and 300. The default is 30 seconds. • **Maximum TCP Dump Size**. Displays the maximum size, in MB, for a dump cycle. Legal values are between 1 and 50. The default is 10 MB. • **TCP Dump Repetition**. Specifies whether the system performs one dump or multiple dumps for each DoS attack.

The settings are incorporated into the profile.

Configuring for Protocol DNS

You can use this tab to configure the conditions under which the system determines that your DNS server is under a DoS attack.

1. In Shared Security, click **DoS Profiles**.
2. In the DoS Profiles screen, click the **profile** you want to configure.
3. On the Properties tab, be sure the check box for Protocol DNS is selected.
4. Select the **Protocol DNS** tab.
5. To enable Protocol Errors Attack Detection, select **Enabled** from the list.
6. Specify the adjustable settings as necessary for your configuration.
   The system saves settings as you enter them.

Setting	Description
Rate increased by	Specifies that the system considers traffic to be an attack if the rate of requests increases greater than this number. By default, the system calculates this number every hour and updates it every minute. The default is 500 percent.
Rate threshold	Specifies the number of packets per second that must be exceeded in order to indicate to the system that there is an attack. The default is 250,000 packets per second.
Rate Limit	Specifies the limit in packets per second. The default is 2,500,000 packets per second.
DNS Query Attack Detection	The screen lists commonly known DNS query types that you want the system to detect in packets. To enable individually, select the **Enabled** check box to the right of the query type (and under Detection Status). Then, specify threshold, rate increase, and rate limit for the particular query type.

The settings are incorporated into the profile.

Configuring for Protocol SIP

You can use this tab to configure the conditions under which the system determines that your server, running the SIP protocol, is under a DoS attack.

1. In Shared Security, click **DoS Profiles**.
2. In the DoS Profiles screen, click the **profile** you want to configure.
3. On the Properties tab, be sure the check box for Protocol SIP is selected.
4. Select the **Protocol SIP** tab.
5. To enable Protocol Errors Attack Detection, select **Enabled** from the list.

   When enabled, the system detects SIP attacks based on a high volume of protocol errors, and displays both how many packets with errors per second are allowed before the system tracks SIP traffic anomalies, and in percentage, how much of an increase in SIP traffic with errors is legal before the system tracks SIP traffic anomalies.

   Specify the following settings as necessary for your configuration. The system saves settings as you enter them.

Setting	Description
Rate increased by	Specifies that the system considers traffic to be an attack if the rate of requests increases greater than this number. The system calculates this number, by default, every hour and updates it every minute. The default setting is 500 percent.
Rate threshold	Specifies the number of packets per second that must be exceeded in order to indicate to the system that there is an attack. The default setting is 250,000 packets per second.
Rate Limit	Specifies the limit in packets per second. The default setting is 2,500,000 packets per second.
SIP Method Attack Detection	The table identifies commonly-known SIP method types that you want the system to detect in packets. Enable a method type by selecting the **Enabled** check box under Detection Status. Then, specify threshold, rate increase, and rate limit for the particular method type.

The settings are incorporated into the profile.

Configuring for Network Security

In this tab, you can configure the conditions under which the system determines that your server is under a network DoS attack.

1. In Shared Security, click **DoS Profiles**.
2. In the DoS Profiles screen, click the **profile** you want to configure.
3. On the Properties tab, be sure the check box to enable Network is selected.
4. Select the **Network** tab.
5. Attack types appear along with adjustable settings for thresholds, rate increases, and rate limits for each attack type you enable. Specify the settings as necessary for your configuration.

Setting	Description
Attack Types	Enable each attack type by selecting the **Enabled** check box to the right of the attack type (and in the Detection Status column). Then, specify settings for thresholds, rate increases, and rate limits.
Threshold	Specifies the number of packets per second, averaged over the previous minute, that must be exceeded to indicate that there is an attack underway. The default setting is 1000 packets per second.
Rate Increase	If the rate of requests increases greater than the number specified here, the system considers the traffic to be an attack. By default, the system calculates this number every hour and updates it every minute. The default setting is 500 percent.
Rate Limit	Specifies the absolute limit of such packets allowed per second.

The settings are incorporated into the profile.

Editing DoS profiles

You can edit DoS profiles to fine tune the circumstances under which the system considers traffic to be a DoS attack and how the system handles a DoS attack.

1. Click **DoS Profiles**.
2. In the DoS Profiles screen, click the name of the profile to modify. The profile is locked for editing.

   For details, consult the sections in this guide:

   - Configuring for Application Security
   - Configuring for Protocol DNS
   - Configuring for Protocol SIP
   - Configuring for Network Security
3. Perform edits as needed for your configuration. The system saves edits as you make them.

Changes to the DoS profile are saved.

Managing Device DoS in Shared Security

About device DoS

You can use the BIG-IQ® Device DoS screen to manage the device DoS configuration on the BIG-IP®, including managing up to 8 network white lists.

You edit a device DoS configuration on the BIG-IQ system by clicking the name of the BIG-IP device to manage on the Device DoS screen.

To get help on any screen, click the **(?)** icon in the upper right corner.

Editing device DoS

Use the Device DoS screen to edit and view the device DoS properties.

1. Click the name of the BIG-IP® device configuration to edit on the Device DoS screen.
2. Modify the properties as needed. Note that not all properties can be modified. Some properties are read-only.
   - Click **Device Configuration** to view values within the configuration.
   - Click **Network White List** to add, delete or modify whitelist entries.
3. When finished, save your changes in one of two ways:
   - Click **Save** to save the device DoS properties.
   - Click **Save & Close** to save the device DoS properties and return to the Device DoS screen.

Editing device configuration entries

1. Click **Device Configuration** to view values within the configuration.
2. Locate the configuration category containing the entry to modify, and click the + to the left of it. The category expands.
3. Click the value to change and then edit it in the popup window displayed.
4. Click **OK** to save the value changes; to cancel the change, click **Cancel**.
5. When finished, save your changes in one of two ways:
   - Click **Save** to save the device DoS properties.
   - Click **Save & Close** to save the device DoS properties and return to the Device DoS screen.

Adding network whitelist entries

1. Click **Network White List** and then click **Add White List**. The Properties screen displays. Only 8 whitelist entries can exist at a time.
2. Type or modify the properties as needed. You can specify IPv4 or IPv6 addresses in CIDR notation as values to address fields. You can specify a source address or destination address but not both in the same whitelist entry.
3. Click **Finished** to complete the whitelist entry.
4. Click **Repeat** to create a copy of the current whitelist entry.
5. Click **Cancel** to not change the whitelist entry.
6. When finished, save your changes in one of two ways:

- Click **Save** to save the whitelist changes.
- Click **Save & Close** to save the whitelist changes and return to the Device DoS screen.

Editing network whitelist entries

1. Click **Network White List**.
2. Click **Edit** at the end of the row containing the whitelist to edit. The Properties screen displays.
3. Modify the properties as needed.
4. Click **Finished** to save the value change; to cancel the change, click **Cancel**.
5. When finished, save your changes in one of two ways:

   - Click **Save** to save the whitelist changes.
   - Click **Save & Close** to save the whitelist changes and return to the Device DoS screen.

Deleting network whitelist entries

1. Click **Network White List**.
2. Select the check box to the left of the whitelist to delete.
3. Click **Remove White List**.
4. When finished, save your changes in one of two ways:

   - Click **Save** to save the deletions.
   - Click **Save & Close** to save the deletions and return to the Device DoS screen.

Managing Logging Profiles in Shared Security

About logging profiles

The Logging Profiles screen in Shared Security lists logging profiles, scaled so that a subset of profiles is visible in the screen at any given time.

A *logging profile* records requests to the virtual server. A logging profile determines where events are logged, and which items (such as which parts of requests, or which type of errors) are logged. Events can be logged either locally by the system and viewed in the Event Logs screens, or remotely by the client's server. The system forwards the log messages to the client's server using the Syslog service.

The logging profile can be associated with multiple virtual servers from multiple devices. Multiple logging profiles can be associated with a virtual server, but the multiple logging profiles cannot have an overlap subset configured. For example, two logging profiles with application security configured and enabled cannot be associated with the same virtual server. The application security and protocol security cannot be configured on the same logging profile or associated with the same virtual server. BIG-IQ® Security supports importing logging profiles with spaces in the name. An imported logging profile with spaces in the name can be modified on the BIG-IQ Security system and deployed back to a BIG-IP device. However, the BIG-IQ system does not support creating logging profiles with spaces in the name.

The logging publisher cannot be created or modified by the BIG-IQ Security system. The logging publisher specified by the BIG-IQ logging profile should be the same as that configured on the BIG-IP device.

To get help on any screen, click the **(?)** icon in the upper right corner.

Create logging profiles

You create logging profiles to configure the kind of information that is logged for each object that supports logging.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click **Create** on the Logging Profiles screen.
   The Logging Profiles - New Item screen opens with the Properties tab displayed.
5. In the **Name** field, type a name for the logging profile.
6. In the **Description** field, type an optional description for the logging profile.
7. If needed, change the default `Common` partition in the **Partition** field.

   The partition with that name must already exist on the BIG-IP®device. No whitespace is allowed in the partition name. Only users with access to a partition can view the objects (such as the logging profile) that it contains. If the logging profile resides in the `Common` partition, all users can access it.
8. Enable each logging type that you want to use by selecting the **Enabled** check box in the appropriate row.
   - Enable **Application Security** to specify that the system logs traffic to the web application. You cannot enable both **Application Security** and **Protocol Security**.

- Enable **Protocol Security**to specify that the system logs any dropped, malformed, and/or rejected requests sent through the given protocol.
- Enable **Network Firewall** to specify that the system logs ACL rule matches, TCP events, and/or TCP/IP errors sent to the network firewall.
- Enable **Network Address Translation** to specify which Network Address Translation (NAT) events the system logs, and where those events are logged.
- Enable **DoS Protection** to specify that the system logs detected DoS attacks, and where DoS events are logged.

A configuration tab is added for each enabled logging type.

9. Click the configuration tab to configure the logging type.
10. When finished, save your changes.

The logging profile is saved.

You must configure each enabled logging type before you can use it.

Configure for Application Security logging

You need to configure application security logging profiles after you have enabled them. This configuration determines the kind of information that is logged.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click the name of the logging profile to configure on the Logging Profiles screen.
   The Logging Profiles - *logging-profile-name* screen opens with the Properties tab displayed.
5. Click the **Application Security** tab.
   The Application Security configuration tab is displayed.
6. Supply the Application Security Configuration settings.

Property	When enabled:
Local Storage	Specifies that the system stores all traffic in the system.
Guarantee Local Logging	Specifies that the system logs all requests, even though this might slow your web application. When cleared (disabled), specifies that the system logs the requests as long as it does not slow your web application. The default is disabled. In either case, the system does not drop requests.
Response Logging	Specifies whether the system logs HTTP responses. • **Off:** The system does not log responses. This is the default. • **For Illegal Requests Only:** The system logs responses to illegal requests. • **For All Requests:** The system logs all responses if the **Request Type** setting in the Storage Filter area is set to **All Requests**.
Remote Storage	Specifies that the system stores all traffic on a remote logging server. Provides additional remote storage options: • **Logging Format:** Specifies the logging format for the remote storage: • Select **Comma-Separated Values** to store traffic on a remote logging server like syslog. Messages are in syslog CSV format.

Property	When enabled:

- Select **Key-Value Pairs** to store traffic on a third party reporting server (for example, Splunk) using a pre-configured storage format. Key value pairs are used in the log messages.
- Select **Common Event Format (ArcSight)** if your network uses ArcSight servers. Log messages are in Common Event Format (CEF).
- Select **BIG-IQ** if you are using a BIG-IQ system as your logging server and you are using a BIG-IP device version 12.0 or later that has enabled the option to use a BIG-IQ system as a logging server.

The logging format you select determines what other options are displayed.

- **Protocol:** Specifies the protocol that the remote storage server uses.
- **Server Addresses:** Specifies one or more remote servers, reporting servers, ArcSight servers or BIG-IQ systems on which to log traffic. Type the values for the **IP Address** and **Port**, and click **Add** for each server.

Note: *The default value for* ***Port*** *is 514 for all types of remote storage other than* ***BIG-IQ****. If* ***BIG-IQ*** *is selected for the* ***Remote Storage Type****, the default port value is 8514.*

- **Facility:** Specifies the facility category of the logged traffic. The possible values are **LOG_LOCAL0** through **LOG_LOCAL7**.

Note: *If you have more than one security policy, you can use the same remote logging server for both applications, and use the facility filter to sort the data for each.*

- **Storage Format:** Specifies how the log displays information and which traffic items the server logs, and in what order it logs them:
  1. To determine how the log appears: select **Field-List** to display the items in the **Selected** list in CSV format with a delimiter you specify; select **User-Defined** to display the items in the **Selected** list in addition to any free text you type in the **Selected** list.
  2. To specify which items appear in the log and in what order, move items from the **Available** list into the **Selected** list.
- **Maximum Query String Size:** Specifies how much of a request the server logs.
  - Select **Any** to log the entire request.
  - Select **Length** and type the maximum number of bytes to log to limit the number of bytes that are logged per request. The value you specify for **Length** must be less than the value specified for **Maximum Entry Length**.
- **Maximum Entry Length:** Specifies how much of the entry length the server logs. The default length is 1K for remote servers that support UDP, and 2K for remote servers that support TCP and TCP-RFC3195. You can change the default maximum entry length for remote servers that support TCP.
- **Report Detected Anomalies:** Select **Enabled** if you want the system to send a report string to the remote system log when a brute force attack or web scraping attack starts and ends.

7. Supply the Application Security Storage Filter settings.

Property	When enabled:
Logic Operation	Specifies whether requests must meet one or all criteria in the Storage Filter area for the system, or server, to log the requests.

Property	When enabled:
	• **OR:** Specifies that requests must meet at least one of the criterion in the Storage Filter settings in order for the system, or server, to log the requests. This is the default. • **AND:** Specifies that requests must meet all of the criteria in the Storage Filter settings in order for the system, or server, to log the requests.
Request Type	Specifies which kind of requests the system, or server, logs. • **Illegal requests only:** Specifies that the system, or server, logs only illegal requests. This is the default. • **Illegal requests, and requests that include staged attack signatures:** Specifies that the system, or server, logs illegal requests, and logs requests that include attack signatures in staging (even though the system considers those requests legal). • **All requests:** Specifies that the system, or server, logs all requests.
Protocols	Specifies whether request logging occurs for all protocols or only for selected protocols. • **All:** Specifies that the system, or server, logs requests for all protocols. This is the default. • **Only:** Specifies that the system, or server, logs requests for only the specified protocol. **HTTP** and **HTTPS** are available for all supported BIG-IP device versions. **WS** and **WSS** are available only with BIG-IP devices version 12.1 or later. You can select more than one protocol for BIG-IP devices version 12.1 or later.
Response Status Codes	Specifies whether request logging occurs for all response status codes or only for selected response status codes. This setting applies only to requests that are not blocked by the system. • **All:** Specifies that the system, or server, logs all requests that generate all response status codes. This is the default. • **Only:** Specifies that the system, or server, logs only requests that generate specific response status codes. When selected, displays additional options where you specify the type of response status code to log. Unused status codes are in the **Available** list, selected status codes are in the **Selected** list.
HTTP Methods	Specifies whether request logging occurs for all HTTP methods or only for selected HTTP methods. • **All:** Specifies that the system, or server, logs requests for all HTTP methods. This is the default. • **Only:** Specifies that the system, or server, logs requests for the specified HTTP method. When selected, displays options where you specify the type of HTTP method to log.
Request Containing String	Specifies whether the request logging is dependent on a specific string. • **All:** Specifies that the system logs all requests, regardless of string. This is the default. • **Search In:** Specifies that the system logs only requests containing a specific string in a particular part of the request. • Select the part of the request to search from the list (**Request**, **URI**, **Query String**, **Post Data**, or **Headers**).

Property	When enabled:
	• Type the string to search for in the request in the field to the right. The search is case-sensitive.

8. When finished, save your changes.

The Application Security configuration settings are saved.

Configure for Protocol Security logging

You need to configure protocol security logging profiles after you have enabled them. This configuration determines the kind of information that is logged.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click the name of the logging profile to configure on the Logging Profiles screen.
   The Logging Profiles - *logging-profile-name* screen opens with the Properties tab displayed.
5. Click the **Protocol Security** tab.
   The Protocol Security configuration tab is displayed.
6. Select the log publisher to use for the HTTP, FTP and SMTP protocols in the **Publisher** property in the HTTP, FTP, and SMTP Security area, or accept the default of **None**.

   This value specifies where the system sends log messages.
7. Supply the Protocol Security DNS Security settings to configure where the system logs any dropped, malformed, rejected, and malicious DNS requests.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging DNS security events. Select a log publisher from the list, or accept the default of **None**.
Log Dropped Requests	Specifies that the system logs dropped DNS requests.
Log Filtered Dropped Requests	Specifies that the system logs filtered dropped DNS requests.
Log Malformed Requests	Specifies that the system logs malformed DNS requests.
Log Rejected Requests	Specifies that the system logs rejected DNS requests.
Log Malicious Requests	Specifies that the system logs malicious DNS requests.
Storage Format	Specifies the format type for log messages. You can set the following options: • **None** Specifies that the system uses the default format type to log the messages to a Remote Syslog server. This is the default setting. • **Field-List** Specifies that the system uses a set of fields, set in a specific order, to log messages. When this is selected, specify the field list as follows. • Specify the delimiter string in the **Delimiter** field. The default delimiter is the comma character (,).

Property	When enabled:
	Note: *You may not use the $ character because it reserved for internal usage.* • Select the fields to use. Unused fields are in the **Available** list, selected fields are in the **Selected** list. • **User-Defined** Specifies that the format the system uses to log messages is in the form of a user-defined string. Select the items for the server to log. Unused items are in the **Available** list, selected items are in the **Selected** list.

8. Supply the Protocol Security SIP Security settings to configure where the system logs any dropped and malformed malicious SIP requests, global and request failures, redirected responses, and server errors.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging SIP protocol security events. Select a log publisher configured in your system.
Log Dropped Requests	Specifies that the system logs dropped requests.
Log Global Failures	Specifies that the system logs global failures.
Log Malformed Requests	Specifies that the system logs malformed requests.
Log Redirection Responses Requests	Specifies that the system logs redirection responses.
Log Request Failures	Specifies that the system logs request failures.
Log Server Errors	Specifies that the system logs server errors.
Log Server Errors	Specifies that the system logs server errors.
Storage Format	Specifies the format type for log messages. You can configure the following options: • **None** Specifies that the system uses the default format type to log the messages to a Remote Syslog server. This is the default setting. • **Field-List** Specifies that the system uses a set of fields, set in a specific order, to log messages. When **Field-List** is selected, specify the field list as follows. • Specify the delimiter string in the **Delimiter** field. The default delimiter is the comma character (,). ***Note:*** *You may not use the $ character because it reserved for internal usage.* • Select the fields to use. Unused fields are in the **Available** list, selected fields are in the **Selected** list. • **User-Defined** Specifies that the format the system uses to log messages is in the form of a user-defined string. Select the items for the server to log. Unused items are in the **Available** list, selected items are in the **Selected** list.

9. When finished, save your changes.

The Protocol Security configuration settings are saved.

Configure for Network Firewall logging

You need to configure network firewall logging profiles after you have enabled them. This configuration determines the kind of information that is logged.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click the name of the logging profile to configure on the Logging Profiles screen.
   The Logging Profiles - *logging-profile-name* screen opens with the Properties tab displayed.
5. Click the **Network Firewall** tab.
   The Network Firewall configuration tab is displayed.
6. Supply the Network Firewall settings to configure which network firewall events the system logs, and where they are logged.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging Network events. Select a log publisher configured in your system.
Aggregate Rate Limit	Defines a rate limit for all combined network firewall log messages per second. Beyond this rate limit, log messages are not logged. You can select a **Rate Limit** value of **Indefinite**, which sets the rate limit to the maximum of 4294967295, or you can select **Specify** to specify a lower rate limit as an integer between 0 and 4294967295.
Log Rule Matches	Specifies that the system logs packets that match the ACL rules. • **Accept** Specifies that the system logs packets that match ACL rules configured with `action = Accept`. • **Drop** Specifies that the system logs packets that match ACL rules configured with `action = Drop`. • **Reject** Specifies, that the system logs packets that match ACL rules configured with `action = Reject`. When specifying the **Rate Limit** for all network firewall log messages of one of the match types: • A value of **Indefinite** sets the rate limit to the maximum of 4294967295, and a value of **Specify** allows you to specify a lower rate limit as an integer between 0 and 4294967295. • If the rate limit is exceeded, log messages of the matched action type are not logged until the threshold drops below the specified rate.
Log IP Errors	Specifies that the system logs IP error packets. When enabled, you can specify a rate limit for all network firewall log messages of this type. If this rate limit is exceeded, log messages of this type are not logged until the threshold drops below the specified rate. You can select a **Rate Limit** value of **Indefinite**, which means the rate limit is set to the maximum of 4294967295, or you can select **Specify** and specify an integer between 0 and 4294967295 that represents the number of messages per second.
Log TCP Errors	Specifies that the system logs TCP error packets. If this rate limit is exceeded, log messages of this type are not logged until the threshold drops below the specified rate. You can select a **Rate Limit** value of **Indefinite** which means the rate limit is

Property	When enabled:
	set to the maximum of 4294967295, or you can select **Specify** and specify an integer between 0 and 4294967295 that represents the number of messages per second.
Log TCP Events	Specifies that the system logs TCP events (open and close of TCP sessions). If this rate limit is exceeded, log messages of this type are not logged until the threshold drops below the specified rate. You can select a **Rate Limit** value of **Indefinite** which means the rate limit is set to the maximum of 4294967295, or you can select **Specify** and specify an integer between 0 and 4294967295 that represents the number of messages per second.
Log Translation Fields	Specifies that translation values are logged if and when a network firewall event is logged.
Always Log Region	Specifies that the geographic location should be logged when a geolocation event causes a network firewall event.
Storage Format	Specifies the format type for log messages. You can configure the following options: • **None** Specifies that the system uses the default format type to log the messages to a Remote Syslog server. This is the default setting. • **Field-List** Specifies that the system uses a set of fields, set in a specific order, to log messages. When **Field-List** is selected, specify the field list as follows. • Specify the delimiter string in the **Delimiter** field. The default delimiter is the comma character (,). ***Note:*** *You may not use the $ character because it reserved for internal usage.* • Select the fields to use. Unused fields are in the **Available** list, selected fields are in the **Selected** list. • **User-Defined** Specifies that the format the system uses to log messages is in the form of a user-defined string. Select the items for the server to log.

7. Supply the Network Firewall IP Intelligence settings to configure where IP intelligence events are logged.

   If the IP intelligence feature is enabled and licensed, you can configure the system to log source IP addresses that match an IP intelligence blacklist or whitelist category, as determined by the database of preconfigured categories, or as determined from an IP intelligence feed list.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging IP address intelligence events. Select a log publisher configured in your system.
Aggregate Rate Limit	Defines a rate limit for all combined IP intelligence log messages per second. Beyond this rate limit, log messages are not logged until the threshold drops below the specified rate. You can select a rate limit value of **Indefinite** which means the rate limit is set to the maximum of 4294967295, or you can select **Specify** and specify an integer between 0 and 4294967295 that represents the number of messages per second.
Log Translation Fields	Specifies that translation values are logged if and when a network firewall event is logged.

8. Supply the Network Firewall Traffic Statistics settings to configure logging of traffic statistics.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging traffic statistics. Select a log publisher configured in your system.
Log Timer Events	Specifies: • **Active Flows** - Logs the number of active flows each second. • **Reaped Flows** - Logs the number of reaped flows, or connections that are not established because of system resource usage levels. • **Missed Flows** - Logs 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. • **SYN Cookie (Per Session Challenge)** - Logs the number of SYN cookie challenges generated each second. • **SYN Cookie (White-listed Clients)** - Logs the number of whitelisted SYN cookie clients each second.

9. Supply the Network Firewall Port Misuse settings to configure logging of port misuse policies.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging port misuse policies. Select a log publisher configured in your system.
Aggregate Rate Limit	Defines a rate limit for all port misuse policy log messages per second. Beyond this rate limit, log messages are not logged until the threshold drops below the specified rate. You can select a rate limit value of **Indefinite** which means the rate limit is set to the maximum of 4294967295, or you can select **Specify** and specify an integer between 0 and 4294967295 that represents the number of messages per second.

10. Supply the Network Firewall SSH Proxy settings to configure logging of SSH proxy use.

Property	When enabled:
Publisher	Specifies the name of the log publisher used for logging SSH proxies. Select a log publisher configured in your system.
Allowed Channel Action	Logs allowed channel action events.
Disallowed Channel Action	Logs disallowed channel action events.
Non SSH Traffic	Logs non SSH traffic events.
SSH Timeout	Logs SSH timeout events.
Successful Client Side Auth	Logs successful client side authentication events.
Successful Server Side Auth	Logs successful server side authentication events.
Unsuccessful Client Side Auth	Logs unsuccessful client side authentication events.
Unsuccessful Server Side Auth	Logs unsuccessful server side authentication events.

11. When finished, save your changes.

The Network Firewall configuration settings are saved.

Configure for Network Address Translation logging

You need to configure network address translation (NAT) logging profiles after you have enabled them. This configuration determines the kind of information that is logged.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.

2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click the name of the logging profile to configure on the Logging Profiles screen.
   The Logging Profiles - *logging-profile-name* screen opens with the Properties tab displayed.
5. Click the **Network Address Translation** tab.
   The Network Address Translation configuration tab is displayed.
6. Supply the Network Address Translation settings to configure which NAT events the system logs, and where they are logged.

Property	When enabled:
LNS Legacy Mode	Specifies that events be logged in Carrier Grade Network Address Translation (CGNAT) LSN format for backward compatibility. If not enabled, the newer HSL logging format is used, which is the default.
Start Outbound Session	Specifies logging for the start of an outbound translation session, when the outbound flow is created. Select **Backup Allocation Only** to log the translation event if the translation occurred due to backup addresses being configured in a NAT Source Translations object. You create NAT Source Translations backup addresses by modifying NAT Source Translations properties within the Network Security policy editor.
End Outbound Session	Specifies logging for the end of an outbound translation session, when the outbound flow is deleted. Select **Backup Allocation Only** to log the translation event if the translation occurred due to backup addresses being configured in a NAT Source Translations object.
Start Inbound Session	Specifies logging for the start of an incoming connection to a translated address. Select **Backup Allocation Only** to log the translation event if the translation occurred due to backup addresses being configured in a NAT Source Translations object.
End Inbound Session	Specifies logging for the end of an incoming connection to a translated address. Select **Backup Allocation Only** to log the translation event if the translation occurred due to backup addresses being configured in a NAT Source Translations object.
Quota Exceeded	Specifies logging when a client exceeds the allocated resource limit.
Errors	Specifies logging when errors are encountered while attempting translation for clients.
Publisher	Specifies the name of the log publisher used for logging NAT events. Select a log publisher configured in your system.

7. When finished, save your changes.

The Network Address Translation configuration settings are saved.

Configure for DoS Protection logging

You need to configure DoS protection logging profiles after you have enabled them. This configuration determines the kind of information that is logged.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.

3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **Logging Profiles**.
4. Click the name of the logging profile to configure on the Logging Profiles screen.
   The Logging Profiles - *logging-profile-name* screen opens with the Properties tab displayed.
5. Click the **DoS Protection** tab.
   The DoS Protection configuration tab is displayed.
6. Supply the DoS Application Protection settings to configure where DoS application protection events are logged.
   - Enable **Local Publisher** to specify that the system logs DoS events to the local database.
   - Select a value for **Remote Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
7. Supply the DNS DoS Protection setting to configure where DNS DoS protection events are logged. Select a value for **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
8. Supply the SIP DoS Protection setting to configure where SIP DoS protection events are logged. Select a value for **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
9. Supply the Network DoS Protection setting to configure where Network DoS protection events are logged. Select a value for **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
10. When finished, save your changes.

The DoS Protection configuration settings are saved.

Editing logging profiles

Use the Logging Profiles screen to edit logging profiles.

1. Click the name of the logging profile on the Logging Profiles screen. The Logging Profiles - logging profile name screen displays, where logging profile name is the name of the logging profile you are editing.
2. In the Logging Profiles - logging profile name screen, review and add or modify the properties as appropriate. The logging profile properties are described in *Creating logging profiles* in this section.
3. When finished, save your changes in one of two ways:
   - Click **Save** to save the logging profile.
   - Click **Save & Close** to save the logging profile and return to the Logging Profiles screen.

Deleting logging profiles

Use the Logging Profiles screen to delete logging profiles.

1. Select the name of the logging profile on the Logging Profiles screen.
2. Click **Delete**.

The logging profile is removed from the list of defined logging profiles.

Managing SSH Profiles in Shared Security

About SSH profiles

You can configure SSH profiles to manage SSH connections. Once the SSH profile is created, you assign it to a virtual server. You enable logging for SSH proxies using logging profiles.

You use the BIG-IQ® Centralized Management system to manage SSH profiles for BIG-IP® devices running version 12.1.1 HF1, or later. For additional details about SSH proxy security, refer to the BIG-IP documentation.

Create SSH profiles

You create SSH proxy profiles to manage user access through SSH connections. This includes selecting what commands are available to users within an SSH connection.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **SSH Profiles**.
4. Click **Create**.
   The SSH Profiles - New Item screen opens with the Properties tab displayed.
5. In the **Name** field, type a name for the SSH profile.
6. In the **Description** field, type an optional description for the SSH profile.
7. If needed, change the default `Common` partition in the **Partition** field.

   The partition with that name must already exist on the BIG-IP device. No whitespace is allowed in the partition name.
8. In the **Timeout** field, if the default value of 0 is not appropriate, type how long, in seconds, before the connection times out.
9. Click **Save & Close** to save the SSH profile and return to the SSH Profiles screen.

The SSH profile has been created.

You add SSH proxy permissions and authentication keys to the SSH profile, as needed, to make it complete. Once complete, you can add the SSH profile to an appropriate virtual server.

Configure SSH proxy permissions

You must create an SSH profile before you can configure the permissions for that profile.

You use the SSH Proxy Permissions tab to configure rules for SSH proxy permissions for the SSH profile. These rules specify what channel actions are allowed for all users and for selected users. A single SSH connection may contain multiple channels and actions, such as Shell, SCP Up, and others.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **SSH Profiles**.
4. Click the name of the SSH profile on which you want to configure permissions.
5. Click the SSH Proxy Permissions tab, and click **Create Rule**.

Each SSH profile has the rule DEFAULT ACTIONS defined which initially allows all listed permissions for all users with no logging enabled. You can modify the permission and logging options for the DEFAULT ACTIONS rule. Review the DEFAULT ACTIONS rule before you create a new rule for specific users.

A new row appears in the table of rules. The row contains a rule template, including defaults, for the new rule.

6. Click the name of the rule to edit the default rule properties.
7. In the **Name** field, type a more meaningful name for the rule.
8. Create the list of SSH user accounts handled by the rule, by adding and removing those accounts from the **Users** column.
   - Add a new SSH user account to the list by typing the account name in the empty **Users** field, and then clicking + to the right of that field.
   - Delete an existing SSH user account from the list by clicking **X** to the right of the user account.
9. Review and, if needed, modify each SSH channel action. You can set each of the SSH channel actions listed in the table columns (such as **Shell**, or **Sub System**) to one of these options:
   - **Allow** permits the session to be set up for the SSH channel action. This is the default.
   - **Disallow** denies an SSH channel action, and sends a `command not accepted` message. Note that many SSH clients disconnect when this occurs.
   - **Terminate** ends an SSH connection by sending a reset message when a channel action is received.
10. To enable logging for any action, select the **Log** check box below the SSH channel action.
11. Review your settings, and click **Save**.

The SSH proxy permissions are defined for the SSH profile.

If not already defined, you can now configure the authentication keys to complete the SSH profile.

Configure SSH authentication keys

You must create an SSH profile before you can configure the authentication keys for that profile.

You use the Key Management tab to configure authentication key information for the SSH profile, such as proxy client authentication, proxy server authentication, and real server authentication.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **SSH Profiles**.
4. Click the name of the SSH profile on which you want to configure authentication keys.
5. Click the Key Management tab and click **Add**.
   A popup screen opens where you supply authentication key information.
6. In the **Name** field, type a name for the authentication information.
7. Supply the public, and if needed, private keys for the authentication types to be used in the fields provided.

   Proxy client authentication and Proxy server authentication require both a public and a private key. Real server authentication requires only a public key. Refer to the BIG-IP®AFM documentation on how to generate and use these keys.
8. Click **Add** to add the new authentication information and close the popup screen.
9. Review your settings, and click **Save**.

The authentication keys are defined for the SSH profile.

If not already defined, you can now configure the SSH proxy permissions to complete the SSH profile.

Delete SSH profiles

An SSH profile must be unused by any virtual server before you can delete it.

You can delete obsolete SSH profiles that are no longer used to avoid clutter in the user interface.

1. Log in to the BIG-IQ® Centralized Management system with your user name and password.
2. At the top left of the screen, select **Network Security** from the BIG-IQ menu.
3. Click **Shared Security** from the top menu bar, and then from the list on the left, click **SSH Profiles**.
   The SSH Profiles screen opens.
4. Select the check box to the left of the SSH profile to delete.
5. Click **Delete**.
   The delete confirmation dialog box opens.
6. Click **Delete** to confirm that you want to delete the SSH profile.
   If the SSH profile is in use by a virtual server, you cannot delete it.

If the SSH profile is not in use, it is deleted.

Managing Application Security Policies in BIG-IQ Web Application Security

About security policies in BIG-IQ Web Application Security

BIG-IQ® Web Application Security imports ASM™ application security policies from discovered BIG-IP® devices, and lists them on the Web Application Security policy editor Policies screen. Each security policy is assigned a unique identifier that it carries across the enterprise. This ensures that each policy is shown only once in the Policies screen, no matter how many devices it is attached to. In the BIG-IQ Web Application Security repository, policies are in XML format.

About subcollections in policies

Subcollections are groups of like objects related to the Web Application Security policy. Not all subcollections are visible in the Web Application Security policy editor. Other subcollections can be imported and deployed without being displayed. Generally, you can import subcollections from a BIG-IP® device and then deploy them without editing. Note that you cannot manage wildcard ordering for subcollections using the BIG-IQ user interface.

The following are the supported versions of the BIG-IQ system and the BIG-IP device for each subcollection. Refer to the release notes for BIG-IQ® Centralized Management for detailed information on BIG-IP device and BIG-IQ system support, such as the minimum TMOS version supported for this release.

Subcollection	Discovery and Deployment Support	Edit Support using BIG-IQ GUI	Minimum BIG-IP Device Version Support	Comments
Policy and properties	Yes	Yes	Any	
Character Sets	Yes	Yes	Any	The BIG-IQ user interface can be used to edit parameter names and parameter values.
Data Guard	Yes	Yes	Any	
File Types	Yes	Yes	Any	
IP Address Exceptions	Yes	Yes	Any	
Parameters	Yes	Yes	Any	
Extractions	Yes	Yes	11.6.0	
Response Pages	Yes	Yes	Any	
Signatures	Yes	Yes	Any	
Signature Sets and attack signature configuration	Yes	Yes	Any	No support for manual sets.

Subcollection	Discovery and Deployment Support	Edit Support using BIG-IQ GUI	Minimum BIG-IP Device Version Support	Comments
Blocking settings - violations	Yes	Yes	Any	No support for user defined violations.
Blocking settings - evasions	Yes	Yes	Any	
Blocking settings - http protocol compliance	Yes	Yes	Any	
Blocking settings - web services securities	Yes	Yes	Any	
Policy Builder	Yes	No	Any	
Allowed methods	Yes	Yes	Any	
Headers	Yes	Yes	Any	
Cookies	Yes	No	Any	
Host names	Yes	No	Any	
Geolocation enforcement	Yes	No	11.6.0	
IP Intelligence	Yes	No	11.6.0	
Redirection protection	Yes	No	11.6.0	
Sensitive parameters	Yes	No	Any	
Web scraping	Yes	No	12.0.0	
CSRF protection	Yes	No	11.6.0	
JSON Profiles	Yes	No	11.6.0	
XML Profiles	Yes	No	11.6.0	
GWT Profiles	Yes	No	11.6.0	
URLs	Yes	No	Any	
Login Pages	Yes	No	11.6.0	
Login Enforcement	Yes	No	11.6.0	
Brute Force Attack Preventions	Yes	No	11.6.0	
Session Tracking Configuration	Yes	No	11.6.0	Only configuration is supported, there is no support for online tracking data.

Editing security policies

Application security policies are often created on BIG-IP® devices and come into the BIG-IQ® Web Application Security configuration when you discover the devices. You can view and modify the properties of individual application security policies.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Policies screen: click **Web Application Security** > **Policy Editor**.
3. Click the name of a policy you want to edit.
   The policy is placed under administrative lock. Policy objects that you can view or edit are listed on the left.
4. Edit the properties of each policy object as appropriate. Click the object to edit in the Policy objects list, and click the **Edit** button on the right.

   For the Attack Signatures List object only, click the Attack Signatures List object, then click the signature name to edit in the Name column and click **Edit**.

   Consult the documentation for each policy object to edit it individually.
5. Click **Save** to save the modifications to each object and unlock the policy.

Changes to the policy object are saved in the working-config of the BIG-IQ system. Assuming the policy is assigned to a virtual server, the next deployment task sends the new configuration to one or more BIG-IP devices.

Editing properties settings

Application security policies are often created on BIG-IP® devices and come into the BIG-IQ® Web Application Security configuration when you discover the devices. You can view and modify the properties of individual application security policies.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Properties screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Properties**.
3. In the Properties screen, click the **Edit** button.
   The policy is placed under administrative lock and fields become editable.
4. Edit the properties as appropriate.
5. Click **Save** to save the modifications to each object and unlock the policy.

The system saves changes in the working configuration of the BIG-IQ system. Assuming the policy is assigned to a virtual server, the next deployment task sends the new configuration to one or more BIG-IP devices.

Properties settings

These properties are the general configuration options and settings that determine the overall behavior and functionality of the security policy.

Property	Description
Name	Unique name of the security policy. The **Name** field is editable only on policy creation.
Partition	Partition to which the security policy belongs. Only users with access to a partition can view the

Property	Description
	objects that it contains. If the policy resides in the `Common` partition, all users can access it.
Description	Optional description of the security policy. Type in any helpful details about the policy. ***Note:*** *This field is limited to 255 characters.*
Full Path	Full path to which the security policy belongs.
Application Language	A Language encoding for the web application, which determines how the security policy processes the character sets. The default language encoding determines the default character sets for URLs, parameter names, and parameter values.
Security is case sensitive	If enabled, the security policy treats file types, URLs, and parameters as case-sensitive. When this setting is disabled (not checked), the system stores these policy elements in lowercase in the policy configuration.
Learning Mode	Select one of the options: • **Automatic**: The system examines traffic, makes suggestions, and enforces most suggestions after sufficient traffic over a period of time from various users make it reasonable to add them. A few suggestions must be enforced manually. • **Manual**: The system examines traffic and makes suggestions on what to add to the security policy. You manually examine the changes and accept, delete, or ignore the suggestions. • **Disabled**: The system does not do any learning for the security policy, and makes no suggestions.
Enforcement Mode	Specifies how the system processes a request that triggers a security policy violation. If **Transparent**, specifies that when the system receives a request that violates a policy parameter, the system logs the violation event, but does not block the request. If **Blocking**, specifies that when the system receives a request that violates a policy parameter, the system logs the violation event, blocks the request, and responds to the request by sending the Blocking Response page and Support ID information to the client.
Enforcement Readiness Period	Use the control to indicate the number of days in the period. The default is 7 days. Both security policy entities and attack signatures remain in staging mode before the system suggests you enforce them. The system does not enforce

Property	Description
	policy entities and attack signatures in staging. Staging allows you to test the policy entities and the attack signatures for false positives without enforcing them.
Mask Credit Card Numbers in Request Log	If enabled, credit card numbers are masked in the request log. If disabled (cleared), credit card numbers are not masked.
Maximum HTTP Header Length	Maximum length of an HTTP header name and value that the system processes. The default setting is 8192 bytes. The system calculates and enforces the HTTP header length based on the sum of the length of the HTTP header name and value. To specify a value for length, type a different value in the field. To specify that any length is acceptable, clear the field. An empty field (a value of any) indicates that there are no restrictions on the HTTP header length up to 8192 bytes.
Maximum Cookie Header Length	Maximum length of a cookie header name and value that the system processes. The default setting is 8192 bytes. The system calculates and enforces a cookie header length based on the sum of the length of the cookie header name and value. To specify a value for length, type a different value in the field. To specify that any length is acceptable, clear the field. An empty field (a value of any) indicates that there are no restrictions on the cookie header length up to 8192 bytes.
Allowed Response Status Code	Specifies which requests the security policy permits, based on the HTTP response status codes they return. Click the gear icon to add or delete response codes.
Dynamic Session ID in URL	Specifies how the security policy processes URLs that use dynamic sessions. Click the gear icon to change the setting or create a custom pattern. • **Disabled**: The policy does not enforce dynamic sessions in URLs. • **Default pattern**: The policy uses the default regular expression for recognizing dynamic sessions in URLs. The default pattern is (\/sap\([^)]+\)). Note that you cannot edit the default regular expression. • **Custom pattern**: User-defined regular expression that the security policy uses to recognize dynamic sessions in URLs. Type an appropriate regular expression in the **Value** field and a description in the **Description** field.
Trigger ASM iRule Events	When enabled, specifies that Web Application Security activates ASM™ iRule events. Specifies, when disabled, that Web Application Security does not activate ASM iRule events. The default setting

Property	Description
	is disabled. Leave this option disabled if you either have not written any ASM iRules® or have written iRules that are not ASM iRules. iRule events that are not ASM are triggered by the Local Traffic Manager™. Enable this option if you have written iRules that process ASM iRule events, and assigned them to a specific virtual server.
Trust XFF Header	When set to **No** (the default), specifies that the system does not have confidence in an XFF (X-Forwarded-For) header in the request. Leave this option disabled if you think the HTTP header may be spoofed, or crafted, by a malicious client. With this setting disabled, if Web Application Security is deployed behind an internal proxy, the system uses the internal proxy's IP address instead of the client's IP address. If Web Application Security is deployed behind an internal or other trusted proxy, you can click the gear icon to change the setting and specify that the system has confidence in an XFF header in the request. Select the **Trust XFF Headers** check box and add a required custom header (use a-z, A-Z, no whitespace allowed). The system then uses the IP address that initiated the connection to the proxy instead of the internal proxy's IP address.
Handle Path Parameters	Specifies how the system handles path parameters that are attached to path segments in URIs. • **As parameter**: The system normalizes and enforces path parameters. For each path parameter, the system removes it from URLs as part of the normalization process, finds a corresponding parameter in the security policy (first at the matching URL level, and if not found, then at the global level), and enforces it according to its attributes like any other parameters. • **As URL**: The system does not normalize nor enforce path parameters. Path parameters are considered an integral part of the URL. • **Ignore**: The system removes path parameters from URLs as part of the normalization process, but does not enforce them. ***Note:*** *The maximum number of path parameters collected in one URI path is 10. All the rest of the parameters (from the eleventh on, counting from left to right) are ignored as parameters, but are still stripped off the URI as part of the normalization process.*

Property	Description
	Note: *Path parameters are extracted from requests, but not extracted in responses.*

Editing blocking settings

You can view and edit the application security policy blocking settings to specify how the system responds (learn, alarm, or block) to a request that contains each type of illegal request. You edit the blocking settings for each policy object individually.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Blocking Settings screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Blocking Settings** .
3. In the Blocking Settings screen, click the **Edit** button.
   The policy is placed under administrative lock and fields become editable.
4. Click the arrows to open or close each category and display specific violation types available to configure for that category.
5. Edit the settings to meet your requirements.
6. When you are finished, click **Save** to save the modifications and unlock the policy.

The blocking settings are updated to use the new settings and any changes made are put into effect in the working configuration of the BIG-IQ ®system.

Blocking settings

Blocking Setting	Description
Enforcement Mode	Specifies whether blocking is active or inactive for the security policy. • **Transparent**. Specifies that blocking is disabled for the security policy. This disables blocking for all options on the screen, and the **Block** check boxes are unavailable. • **Blocking**. Specifies that blocking is enabled for the security policy, and you can enable or disable blocking for individual violations.
General Features	• **Request length exceeds defined buffer size**: The incoming request is larger than the buffer for the Security Enforcer parser. • **Illegal session ID in URL**: The incoming request contains a session ID value that does not match the session ID value from a previous request from the same client. • **Illegal meta character in value**: The incoming request includes a parameter, or content, whose value contains an illegal meta character, according to the security policy. • **Illegal HTTP status in response**: The server response contains an HTTP status code that is not defined in the security policy.

Blocking Setting	Description
	• **Failed to convert character**: The incoming request contains a character that does not comply with the encoding of the web application (the character set of the security policy), and the Policy Enforcer cannot convert the character to the current encoding. • **Virus Detected**: The incoming request contains a virus detected by the system. • **Illegal Base64 value**: The incoming request contains data (parameters, cookies, or headers) that does not contain a valid Base64 string in its value.
HTTP protocol compliance failed	When the check box is cleared, the system does not enforce this sub-violation. • **Content length should be a positive number**: Specifies, when checked (enabled), that the system examines requests to see whether their content length value is greater than zero, which is required. The default setting is enabled. • **High ASCII characters in headers**: High ASCII characters in headers: Specifies, when checked (enabled), that the system inspects request headers for ASCII characters greater than 127, which are not permitted. The default setting is disabled. • **Bad HTTP version**: Specifies, when checked (enabled), that the system inspects requests to see whether they request information from a client using a legal HTTP protocol version number (0.9 or higher). The default setting is enabled • **Chunked request with Content-Length header**: Specifies, when checked (enabled), that the system examines chunked requests for a content-length header, which is not permitted. The default setting is enabled. • **POST request with Content-Length**: 0: Specifies, when checked (enabled), that the system examines POST method requests for no content-length header, and for a content length of 0. The default setting is disabled. • **Check maximum number of parameters**: Specifies, when checked (enabled), that the system compares the number of parameters in the request against the maximum number you specify here. A request that contains more parameters than allowed by the security policy should be considered a possible attack on the server. Type a number in the box to specify how many parameters are allowed. The default

Blocking Setting	Description
	value is enabled set to a maximum of 500 parameters. • **Host header contains IP address**: Specifies, when checked (enabled), that the system verifies that the request's host header value is not an IP address. The default setting is disabled. • **Header name with no header value**: Specifies, when checked (enabled), that the system checks requests for valueless header names, which are considered illegal. The default setting is enabled. • **Bad multipart/form-data request parsing**: Specifies, when checked (enabled), that the system examines requests to see whether the content-disposition header field contains the required parameters, name="param_key". The default setting is enabled. • **Multiple host headers**: Specifies, when checked (enabled), that the system examines requests to ensure that they contain only a single "Host" header. The default value is enabled. • **No Host header in HTTP/1.1 request**: Specifies, when checked (enabled), that the system examines requests sent by a client using HTTP version 1.1 to see whether they contain a host header, which is required. The default setting is enabled. • **CRLF characters before request start**: Specifies, when checked (enabled), that the system examines requests to see whether they begin with the characters CRLF, which is not permitted. The default setting is enabled. • **Check maximum number of headers**: Specifies, when checked (enabled), that the system compares the number of headers in the requests against the maximum number you specify here. Type a number in the box to specify how many headers are allowed. The default setting is enabled with a maximum of 20 headers unless the security policy is based on an Application-Ready security policy template. In this case, the default value depends on which template you are using. • **Several Content-Length headers**: Specifies, when checked (enabled), that the system examines each request to see whether it has more than one content-length header, which is considered illegal. The default setting is enabled.

Blocking Setting	Description
	• **Body in GET or HEAD requests**: Specifies, when checked (enabled), that the system examines requests that use the GET or HEAD methods to see whether the requests contain data in their bodies, which is considered illegal. The default setting is disabled. • **Bad host header value**: Specifies, when checked (enabled), that the system inspects requests to see whether they contain a non RFC compliant header value. The default value is enabled. • **Null in request**: Specifies, when checked (enabled), that the system inspects requests to see whether they contain a Null character, which is not allowed. The default setting is enabled. • **Bad multipart parameters parsing**: Specifies, when checked (enabled), that the system examines requests to see whether the content-disposition header field matches the format: name=“param_key”;\r\n. The default setting is enabled. • **Unparsable request content**: Specifies, when checked (enabled), that the system examines requests for content that the system cannot parse, which is not permitted. The default setting is enabled.
Attack Signatures	The system examines HTTP messages for known attacks by comparing them against known attack patterns. Click the **Edit Settings** link to edit the properties of that signature set.
Evasion technique detected	When the check box is cleared, the system does not enforce this sub-violation. • **Bad unescape**: Indicates that the system discovers illegal URL-encoding. For example, if the two bytes after % are not hexadecimal characters, or if the four bytes after %u are not a hexadecimal characters. This violation applies to URI and parameter input. However, for this violation, the system does not change the input. • **Bare byte decoding**: Indicates that the system discovers characters higher than ASCII-127. This violation applies to URI input. However, for this violation, the system does not change the input. • **IIS Unicode codepoints**: Indicates that, when XXXX is greater than 0x00FF, the system decodes %u according to an ANSI Latin 1 (Windows 1252) code page mapping. For example, the system turns a%u2044b to a/b.

Blocking Setting	Description
	The system performs this action on URI and parameter input. • **%u decoding**: Indicates that the system performs %u decoding (%UXXXX where X is a hexadecimal digit). For example, the system turns a%u002fb to a/b. The system performs this action on URI and parameter input. • **IIS backslashes**: Indicates that the system turns backslashes (\) into slashes (/). The system performs this action on URI input. • **Directory traversals**: Indicates that the system clears self references and performs directory traversals so that attackers cannot try to access restricted Web server files residing outside of the Web server's root directory. For example, the system turns a/b/../c to a/c and a/./b to a/b. The system performs this action on URI input. • **Apache whitespace**: Indicates that the system discovers the bytes 0x09, 0x0b, or 0x0c (a non-RFC standard of using tab for a space delimiter). The violation applies to URI input. However, for this violation, the system does not change the input. • **Multiple decoding (considered an evasion after a specified number of decoding passes)**: Indicates that the system performs multiple decoding. For example, the system can turn a %252fb to a/b (since %252f becomes %2f after one pass, and / after the second pass). The system performs this action on URI and parameter input. Select a number to specify how many decoding passes the system performs, and the level at which the system responds with the appropriate Alarm or Block action. For example, if you set this to 3, the system performs two decoding passes, and when it performs the third decoding pass, it takes the action specified by the Learn/Alarm/Block settings of the Evasion technique detected category. Use the **decoding passes** control to specify the number (up to 5) of decoding passes.
File Types	When enabled, the system checks that the requested file type is configured as a valid file type or not configured as an invalid file type. • **Illegal query string length**: The incoming request contains a query string whose length exceeds the acceptable length specified in the security policy. • **Illegal request length**: The incoming request length exceeds the acceptable length specified

Blocking Setting	Description
	in the security policy per the requested file type. • **Illegal URL length**: The incoming request includes a URL whose length exceeds the acceptable length specified in the security policy. • **Illegal file type**: The incoming request references file types not found in the security policy. • **Illegal POST data length**: The incoming request contains POST data whose length exceeds the acceptable length specified in the security policy.
URLs	• **Illegal flow to URL**: The incoming request references a flow that is not found in the security policy. • **Illegal number of mandatory parameters**: The incoming request contains either too few or too many mandatory parameters. • **Illegal meta character in URL**: The incoming request includes a URL that contains an illegal meta character, according to the security policy. Click the arrow icon to view and edit the URL character set. • **Illegal query string or POST data**: The incoming request contains a query string or POST data that is not found in the security policy. • **Illegal URL**: The incoming request references a URL that is not found in the security policy. • **Illegal request content type**: The incoming request for a URL contains header content that is not allowed according to the requested URL's Header-Based Content Profile Parsed As setting. • **Illegal entry point**: The incoming request references an entry point that is not found in the security policy.
Parameters	• **Illegal empty parameter value**: The incoming request contains a parameter whose value is empty when it should have contained a value. • **Illegal repeated parameter name**: The incoming request contains multiple parameters with the same name. • **Illegal parameter value length**: The incoming request contains a parameter whose value length does not match the value length that is defined in the security policy. This violation is relevant only for user input parameters.

Blocking Setting	Description
	• **Illegal value does not comply with regular expression**: The incoming request contains an alphanumeric parameter value that does not match the expected pattern specified by the regular-expression field for that parameter. • **Null in multi-part parameter value**: The incoming multi-part request has a parameter value that contains a NULL character (0x00). • **Illegal parameter date type**: The incoming request contains a parameter for which the data type (for example, alphanumeric) does not match the data type that is defined in the security policy. • **Illegal static parameter value**: The incoming request contains a static parameter whose value is not defined in the security policy. • **Illegal dynamic parameter value**: The incoming request contains a dynamic parameter value that does not comply with the security policy. • **Illegal parameter**: The incoming request contains a parameter that is not defined in the security policy. • **Disallowed file upload content detected**: load content detected: The incoming request contains a user-input File Upload parameter whose value contains binary executable content, and the security policy disallows this. • **Illegal parameter numeric value**: The incoming request contains a parameter whose numeric value is not within the range of decimal or integer values defined in the security policy. • **Illegal meta character in parameter name**: The incoming request includes a parameter name that contains an illegal meta character, according to the security policy. Click the arrow icon to view and edit the parameter name character set.
Sessions and Logins	• **Login URL bypassed**: The user accessed the target URL without having first visited the login URL configured in the security policy. Click the arrow icon to view and edit the security policy's login enforcement configuration. • **Brute Force**: Maximum login attempts are exceeded: The number of times a user tried to log on to a URL is more than what is allowed by the security policy. This indicates an attempt to access secured parts of the website by guessing user names and passwords. Click the

Blocking Setting	Description
	arrow icon to view and edit the security policy's brute force configuration. • **Access from disallowed User/Session/IP**: The system detected that the number of violations from the same User/Session/IP address within the specified time frame is above the configurable limit within session awareness. Click the arrow icon to view and edit the security policy's session tracking configuration. • **Login URL expired**: The login URL configured in the security policy timed-out. Click the arrow icon to view and edit the security policy's login enforcement configuration.
Cookies	• **Cookie not RFC-compliant**: The format of the Cookie header in the request does not comply with the standards as specified in the RFCs for HTTP. • **Illegal cookie length**: The incoming request includes a cookie that exceeds its length setting in the security policy. Click the arrow icon to open the Security Policy Properties screen. • **Modified domain cookie(s)**: The domain cookie(s) in the HTTP request does/do not match the original domain cookies. • **ASM cookie Hijacking**: The incoming request contains an Application Security Manager (ASM) cookie that was created in another session. • **Modified ASM cookie**: The incoming request contains an Application Security Manager (ASM) cookie that has been modified or tampered with. • **Expired timestamp**: The time stamp in the HTTP cookie is old, which indicates that a client session has expired.
Content Profiles	• **XML data does not comply with format settings**: The incoming request contains XML that does not match at least one of the XML defense configuration settings configured in the security policy. • **Malformed JSON data**: The incoming request contains JSON content that is not formed according to the general JSON formatting rules. Click the arrow icon to view JSON profiles that exist in the security policy. • **XML data does not comply with schema or WSDL document**: The incoming request contains an XML document that does not adhere to the rules of a schema or WSDL

Blocking Setting	Description
	document defined in an XML profile configured in the security policy. • **Malformed GWT data**: The incoming request contains data that is not formed according to the general GWT formatting rules. Click the arrow icon to view GWT profiles that exist in the security policy. • **Malformed XML data**: The incoming request contains an XML document whose structure is not formed according to the general XML formatting rules. Click the arrow icon to view XML profiles that exist in the security policy. • **JSON data does not comply with format settings**: The incoming request contains JSON content that does not comply with the request limits of a JSON profile configured in the security policy. • **SOAP method not allowed**: The incoming request invokes an illegal SOAP method according to an XML profile configured in the security policy. • **Illegal attachment in SOAP message**: The incoming request contains a SOAP message in which there is an attachment that is not permitted by the security policy. • **GWT data does not comply with format settings**: The incoming request contains data that does not comply with the various payload limits of a GWT profile configured in the security policy.

Web Services Security Failure	• **Certificate Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks responses when the client certificate extracted from the document is invalid. The default setting is enabled. Possible causes include the following instances:
	  • The client certificate structure is invalid, and cannot be parsed.
	  • The client certificate is not found in the keystore.
	• **Expired Timestamp**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the timestamp has expired. The default setting is enabled.
	• **Missing Timestamp**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the timestamp is missing from the document. The default setting is enabled.

Blocking Setting	Description

	• **Verification Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the underlying crypto library failed to perform digital signature verification, or there is information missing in the payload. The default setting is enabled.
	• **Internal Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests and/or responses when the system's web services security offload engine confronts an unexpected scenario, for example, if a resource fails to allocate. The default setting is enabled.
	• **Timestamp expiration is too far in the future**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the timestamp lifetime is greater than configured. The default setting is enabled.
	• **Encryption Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks responses when the system cannot encrypt a section requested by the user. For example, the message cannot be encrypted if no key information was detected in the request. The default setting is enabled.
	• **Signing Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks responses when the underlying crypto library failed to digitally sign the document, or the response contains an unknown or unsupported algorithm. The default setting is enabled.
	• **UnSigned Timestamp**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the timestamp is not digitally signed. The default setting is enabled.
	• **Invalid Timestamp**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when the timestamp is not formatted according to the specifications. The default setting is enabled.
	• **Decryption Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests when an encrypted section in the request could not be decrypted. The default setting is enabled. Possible causes include the following instances:
	  • The message could not be decrypted since no key information was found.
	  • The encryption algorithm is not supported.

Blocking Setting	Description

	• **Malformed Error**: Specifies, when checked (enabled), that the system learns, logs, or blocks requests and/or responses when the system's web services security offload engine confronts a malformed document, for example, if the document contains characters that are illegal according to the W3C XML 1.0 recommendation. The default setting is enabled.
	• **Certificate Expired**: Specifies, when checked (enabled), that the system learns, logs, or blocks responses when the client certificate extracted from the document has expired. The default setting is enabled. Possible causes include the following instances:
	  • The client certificate structure is invalid and cannot be parsed.
	  • The client certificate is not found in the key-store.
	  The system does not perform this check if the **Save Expired/Untrusted Certificate** option is enabled when you add the certificate to the system's certificate pool.

CSRF Protection	*Cross-site request forgery (CSRF)* is an attack that forces a user to execute unwanted actions on a web application in which the user is currently authenticated. When enabled, this setting specifies that you want the system to protect the web application against CSRF attacks. Click the arrow icon to view and edit the security policy's CSRF protection configuration. • **CSRF authentication expired**: The incoming request may be a Cross-Site Request Forgery (CSRF) attack. The request may come from a clicked link, embedded malicious HTML, or JavaScript in another application, and may involve transmission of unauthorized commands through an authenticated user. • **CSRF attack detected**: The incoming request includes an expired cross-site request forgery (CSRF) session cookie.
IP Addresses / Geolocations	• **Access from disallowed Geolocation**: The incoming request is sent from a country that is configured as disallowed in the security policy. Click the arrow icon to view and edit the security policy's geolocation enforcement configuration. • **Access from malicious IP address**: The incoming request is sent from an IP address

Blocking Setting	Description
	that is considered high risk according to an IP Address Intelligence database. Click the arrow icon to view and edit the security policy's IP address intelligence configuration.
Headers	• **Illegal meta character in header**: The incoming request includes a header whose value contains an illegal meta character, according to the security policy. Click the arrow icon to view and edit the security policy's headers character set. • **Illegal header length**: The incoming request includes a header that exceeds its length setting in the security policy. Click the arrow icon to navigate to the security policy properties screen where you can view and edit the maximum HTTP header length. • **Illegal method**: The incoming request references an HTTP request method that is not found in the security policy. Click the arrow icon to view and edit the security policy's allowed methods. • **Mandatory HTTP header is missing**: The request does not include an HTTP header that is configured in the security policy as being mandatory.
Redirection Protection	• **Illegal redirection attempt**: The server tries to redirect the user to a target domain that is not defined in the security policy.
Data Guard	Specifies which information the system considers sensitive, including credit card numbers, U.S. Social Security numbers, custom patterns, and file content. • **Data Guard: Information leakage detected**: The response from the web server includes data commonly considered sensitive.

Editing response page settings

You can view and edit the default response page, the login page response, the XML response page, and the AJAX response page.

Response page settings specify the content of the response that the system sends to the user when the security policy blocks a client request. You can also configure a redirect URL so that the system redirects the client to another location instead of displaying a response page. Edit the response pages for each policy object individually.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Response Page screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Response Page**.

3. In the Response Page screen, click the **Edit** button to revise the settings as needed.
   The policy is placed under administrative lock and fields become editable.
4. Edit the default response page to specify the content of the response that the system sends to the user when the security policy blocks a client request: Click the **Default Response Page** tab and from the **Response Type** list, select one of the options.
   - **Default Response**: Specifies the system-supplied response text written in HTML. You cannot edit this text.
   - **Custom Response**: Specifies a modified response text. The system provides additional options where you can type the response header and response body you prefer. Click **Upload** to browse to and select a file. When you click **Open**, the contents of the file are loaded into the **Response Body** field.
   - **Redirect URL**: Type a URL that the system is to use to redirect the user to a specific web page instead of viewing a response page.
   - **SOAP Fault**: Displays the system-supplied response written in SOAP fault message structure. Use this type when a SOAP request is blocked due to an XML related violation. You cannot edit this text.
5. Edit the login page response to specify the response that the system sends after the user violates one of the preconditions when requesting the target URL of a configured login page: Click the **Login Page** tab and select one of the options.
   - **Default Response**: Specifies the system-supplied response text written in HTML. You cannot edit this text.
   - **Custom Response**: Specifies a modified response text. The system provides additional options where you can type the response header and response body you prefer. Click **Upload** to browse to and select a file. When you click **Open**, the contents of the file are loaded into the Response Body text box.
   - **Redirect URL**: Type a URL that the system is to use to redirect the user to a specific web page instead of viewing a response page.
   - **SOAP Fault**:Displays the system-supplied response written in SOAP fault message structure. Use this type when a SOAP request is blocked due to an XML related violation. You cannot edit this text.
6. Click the **XML Response Page** tab and then select an option from the **Response Type** list to display and edit settings.
   - **SOAP Fault**: Displays the system-supplied response written in SOAP fault message structure. Use this type when a SOAP request is blocked due to an XML-related violation. You cannot edit this text.
   - **Custom Response**: Specifies a modified response text. The system provides additional options where you can type the response header and response body you prefer.

   The system sends the XML response page when the security policy blocks a client request that contains XML content that does not comply with the settings of an XML profile configured in the security policy.
7. Click the **AJAX Response Page** tab and then select the **Ajax Blocking Enabled** check box to display and edit settings.

   When enabled, the system injects JavaScript code into responses. You must select this check box to configure an Application Security Manager AJAX response page which is returned when the system detects an AJAX request that does not comply with the security policy. The default setting is disabled.

   a) **Default Response Page Action**: Specifies which content, or URL, the system sends to the client as a response to an AJAX request that does not comply with the security policy. Select from the following actions:

- • **Custom Response**: Click the **Upload** link and select the file containing the text that is to serve as the custom response. The text box is populated with the text.
  - **Popup Message**: Accept the default text or type a message.
  - **Redirect URL**: Type a URL (required).

b) **Login Page Response Action**: Specifies which content, or URL, the system sends to the client after the user sends an AJAX request that attempts to directly access a URL that is allowed to be accessed only after visiting a login page. Select from the following actions:

- **Custom Response**: Click the **Upload** link and select the file containing the text that is to serve as the custom response. The text box is populated with the text.
- **Popup Message**: Accept the default text or type a message.
- **Redirect URL**: Type a URL (required).

8. When you are finished, click **Save** to save the modifications and unlock the policy.

The response page settings are updated to use the new settings and any changes made are put into effect in the working configuration of the BIG-IQ system.

Editing Data Guard settings

You can view and edit Data Guard settings to specify which information the system considers sensitive, including credit card numbers, U.S. Social Security numbers, custom patterns, and file content.

1. Log in to BIG-IQ Security with Administrator, Security Manager or Web App Security Manager credentials.
2. Navigate to the Policy objects screen: click **Web Application Security** > **Policy Editor**, select a policy name, and click **Data Guard**.
3. At the right of the screen, click **Edit**.
   The policy is placed under administrative lock and fields become editable.
4. For the **Data Guard** setting, select how the system treats sensitive data:

Option	Description
Protect credit card numbers	Specifies that the system considers credit card numbers as sensitive data. The system returns asterisks to the client instead of the sensitive data.
Protect U.S. security card numbers	Specifies the system considers U.S. security card numbers as sensitive data.
Mask sensitive data	When checked, the system masks sensitive data returned by the web server by returning asterisk (*) characters to the client instead of the sensitive data. When disabled, the system sends the response, including the sensitive information, to the user. If the policy's enforcement mode is Transparent and the **Mask Data** check box is selected, the system encodes the sensitive data by returning asterisks to the client instead of the sensitive data. (The system also returns asterisks if the enforcement mode is Blocking, the **Data Guard: Information leakage detected** violation **Block** check box is cleared, and the **Alarm** check box is checked.) If the policy's enforcement mode is Blocking, and the **Block** check box for the **Data Guard: Information leakage detected** violation is checked, the system blocks the response.

5. To have the system recognize customized patterns as sensitive data, select the **Custom Patterns** check box and then type in the field a pattern you want the system to consider as sensitive data, and click **Add**.
   Use PCRE regular expression syntax. For example, `999-[/d][/d]-[/d][/d][/d][/d].`
   To delete a selected pattern, click the **x**.

6. To have the system recognize exception patterns as not being sensitive data, select the **Exception Patterns** check box and then type in the text box a pattern you want the system to consider as an exception to sensitive data, and click **Add**.

   Use PCRE regular expression syntax. For example, `999-[/d][/d]-[/d][/d][/d][/d].`

   To delete a selected pattern, click the **x**.
7. To specify whether or not the system checks responses for file content, and if so, which types of file content are considered as sensitive data (and not returned to users), select the **File Content Detection** check box. The default setting is cleared (disabled). When checked (enabled), the system displays possible types of content the system could consider as sensitive data. You must select at least one check box.
8. To specify the URLs for which the system enforces Data Guard protection and the URLs it ignores, select one of the check boxes shown when **File Content Detection** is checked.

Option	Description
Enforce URLs in List	The system enforces data guard protection only for those URLs in the Enforcement Mode list. It enforces Data Guard protection for these URLs even if they are not in the policy. Type a URL in the text box and click **Add**. ***Note:*** *When adding URLs, you can type either explicit (`/index.html`) or wildcard (`*xyz.html`) URLs.*
Ignore URLs in List	The system enforces data guard protection for all URLs except for those URLs in the Enforcement Mode list (default). Add a URL to the list by typing in the field and clicking **Add**. ***Note:*** *When adding URLs, you can type either explicit (`/index.html`) or wildcard (`*xyz.html`) URLs.*

9. When finished, click **Save** to save the modifications and unlock the policy.

The Data Guard settings are updated to use the new settings and any changes made are put into effect in the working configuration of the BIG-IQ® system.

Editing Headings and Methods settings

In the application security policy, you can view and edit (modify, add, and delete):

- Allowed methods. You can specify methods other web applications may use when requesting a URL from another domain.

  All security policies accept standard HTTP methods by default. If your web application uses HTTP methods other than the default allowed methods (GET, HEAD, and POST), you can add them to the security policy.
- HTTP headers. You can specify a list of request headers for other web applications hosted in different domains to use when requesting this URL.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Headers and Methods screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Headers And Methods**.
3. In the Headers and Methods screen, click the **Edit** button.
   The **Add** and **Delete** buttons become visible. Settings become editable and the policy is placed under administrative lock.
4. In the Allowed Methods area, click **Add** to add a method and specify how it will function.

a) From the **Method** list, select a method.
b) Specify whether the method should act as the GET method or as the POST method.
   - **GET**. Specifies that the request does not contain HTTP data following the header part of the request.
   - **POST**. Specifies that the request contains HTTP data following the header part of the request.
c) When you are finished, click **Save**.
   The new method is added to the Allowed Methods table. The method appears in blue meaning that you can edit it and you can also delete it.

5. In the HTTP Headers area, click **Add**.

The system can perform normalization and attack signature checks on HTTP headers. In this context, *headers* refers to HTTP request headers, not response headers.

a) Select or modify the settings as appropriate.
   - **Name**. From the list, select the name of the HTTP header.
   - **Type**. Specifies `Explicit` or `Wildcard`. The only wildcard header in the system is the default pure wildcard header (*****).
   - **Mandatory**. If enabled, requires this header to appear in requests.
   - **Check Signatures**. If enabled, the system performs attack signature checks on this header.
   - **Base64 Decoding**. When enabled, specifies that the security policy checks the parameter's value for Base64 encoding, and decodes the value. The default is disabled.
   - **Normalization**.

     Header normalization is a process whereby the Application Security Manager™ buffers the contents of request headers to change them into a standard format that can be more easily checked for discrepancies.

     Normalizing deals with special characters (such as percent encoding), non-ASCII text, URL paths and parameters, Base64 encoded binary content, non-printable characters, HTML codes, and many other formats that may be used in headers that could potentially hide malicious code.

     You do not need to normalize all headers. You should normalize referer headers, and custom headers containing binary data, URLs, or other encoded information. There is a performance trade-off when using normalization, so implement it only when needed.

     Select the type of normalization to be performed.
   - **Evasion Techniques Violations**. Coding methods that attackers use to avoid detection by attack signatures and intrusion prevention systems.
b) When you are finished, click **Save**.

The system updates the Headers and Methods settings to use the new settings, and puts any changes made into effect in the working configuration of the BIG-IQ® system.

Editing IP address settings

You can view and edit configured IP address exceptions, including particular characteristics, such as:

- Never blocking nor logging traffic sent from configured IP address exceptions.
- Not generating Learning Suggestions for traffic sent from configured IP address exceptions.
- Always allowing traffic sent from configured IP address exceptions.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.

2. Navigate to the IP Address screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **IP Address**.

3. In the IP Address screen, click the **Edit** button.
The policy is placed under administrative lock and fields become editable.

4. Click **Add new IP**.
   a) Type an IP address.

      Enter the IP address that you want the system to trust. To add a route domain, type `%n` after the IP address where n is the route domain identification number.
   b) Type a netmask.

      If you omit the netmask value, the system uses a default value of `255.255.255.255`.
   c) Select **Settings** as appropriate.

      - **Policy Builder Trusted IP**. Select to specify that the Policy Builder considers traffic from this IP address to be legitimate. The Policy Builder automatically adds to the security policy data logged from traffic sent from this IP address.
      - **Ignore in Anomaly Detection**. Select to specify that the system considers traffic from this IP address to be safe. The security policy does not take this IP address into account when performing brute force prevention and web scraping detection.
      - **Ignore in Learning Suggestion**. Select to specify that the system not generate learning suggestions from traffic sent from this IP address.
      - **Never block traffic from this IP Address**. Select to specify that the system not block requests sent from this IP address, even if the security policy is configured to block all traffic.
      - **Never log traffic from this IP Address**. Select to specify that the system not log requests or responses sent from this IP address, even if the security policy is configured to log all traffic.
      - **Ignore in IP Address Intelligence**. Select to specify that the system considers traffic from this IP address to be safe even if it matches an IP address in the IP Address Intelligence database.
   d) Enter a brief description for the IP address.
5. When you are finished, click **Save** to save the modifications and unlock the policy.

The IP Address settings are updated to use the new configured IP address exceptions, and any changes made are put into effect in the working configuration of the BIG-IQ® system.

Adding file types settings

You can add and configure settings for file types that are allowed (or disallowed) in traffic to the web application being protected. These settings determine how the security policy reacts to requests referring to files with these extensions.

1. Log in to BIG-IQ Security with Administrator, Security Manager or Web App Security Manager credentials.
2. Navigate to the File Types screen: (click **Web Application Security** > **Policy Editor**, select a policy name) and click **File Types**.
3. Click **Edit**, then click **Add** and specify either:

   - **Allowed file type**: To add file types in the web application that the security policy considers legal, and to view information about each file type.
   - **Disallowed file type**: To add file types in the web application that the security policy considers illegal, and to exclude file types that are included in allowed wildcard file types.

   The screen displays fields applicable to your selection.
4. If you selected **Allowed file type**, fill in the settings.
   a) For **File type**, select either option and type a name.

      Explicit: Specifies that the system displays only explicit file types.

      Wildcard: Specifies that the system displays only wildcard file types.
   b) For **Perform Staging**, select the **Enabled** check box to have the system perform staging.
   c) For **URL Length**, type the maximum acceptable length, in bytes, of a URL containing this file type.

d) For **Request Length**, type the maximum acceptable length, in bytes, of the request containing this file type.
e) For **Query String Length**, type the maximum acceptable length, in bytes, for the query string portion of a URL that contains this file type.
f) For **POST Data**, type the maximum acceptable length, in bytes, for the POST data of an HTTP request that contains the file type.
g) For **Apply Response Signature Staging**, select the check box to apply response signature staging.

5. If you selected **Disallowed file type**, fill in the name and click **Save**.
6. When you are finished, click **Save** to save the modifications and unlock the policy.

The file types settings are updated to use the new settings, and any changes you made are put into effect in the working configuration of the BIG-IQ® system.

Editing parameters settings

You can view and edit settings for parameters, such as the parameter type and whether the parameter is allowed to contain an empty value. In addition, depending on the parameter type, you can view the parameter's values, extraction methods, enabled attack signatures, and allowed regular expressions.

The GUI lists the configured parameters in a table of 5 columns: check box, Name, Value Type, Level, Staging. The name of the parameter is a link for entering the detail and/or edit mode for the parameter. A default parameter exists for all policies: ***** (asterisk). It is a pure wildcard with a value type of **user-input**, a level of **global**, and staging **enabled**.

Use the filter to search the column values of the parameter (Name, Value Type, Level, Staging).

Using BIG-IQ® Web Application Security's support for extractions, you can:

- Define a parameter's characteristics and add them to the policy by clicking **Edit**. The policy is placed under administrative lock and the fields become editable. To unlock, click **Unlock**.
- Create a new parameter by first clicking **Edit** and then **Add**.
- Delete a parameter by clicking **Edit** to enter edit mode, selecting the check box of the parameter, clicking **Delete**, and clicking **Yes** to confirm.

  From the parameters list, you can select directly by clicking the parameter or you can select multiple parameters using the selection check box. Once selected, you can delete by clicking clicks **Delete** and confirming.

1. Log in to BIG-IQ with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Parameters screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Parameters**.
3. On the Parameters screen, be sure the **Parameters**tab is selected.
4. On the Parameters screen, click the **Edit** button.
   The policy is placed under administrative lock, and the fields become editable.
5. To add a new parameter, click **Add**.
6. In the Parameter details area, specify parameter details.

Option	Description
Name	Note that you cannot change the Name after creation, but you can add a parameter with the same name but a different level. • **explicit**. Specifies that this a regular named parameter. • **wildcard**. Specifies that any parameter name that matches the wildcard expression is permitted by the security policy. (For example, typing the

Option	Description
	wildcard * specifies that the security policy allows every parameter.) The syntax for wildcard entities is based on shell-style wildcard characters. • **unnamed**. Specifies that this parameter does not have a name. The system automatically names the parameter **UNNAMED**. Behaves the same as an explicit parameter. The system configures one unnamed parameter per policy. It is an explicit parameter with an empty name.
Level	Specifies whether or not the screen displays the parameters filtered by level. From the list, select **global** (parameters not associated with flows or URLs) or **URL** (parameters associated with specific URLs). If the security policy is configured to differentiate between HTTP and HTTPS URLs, then you can additionally filter URL parameters by the HTTP and HTTPS protocols.
Perform Staging	Displays the staging status on this parameter (**Enabled** if checked or **Disabled**). Entities in staging do not cause violations, which allows you to fine-tune their settings without causing false positives.
Allow Empty Value	Select to enable (allow empty value).
Allow Repeated Occurrences	Select to enable (allow repeated occurrences).
Sensitive Parameter	In a validated request, protects sensitive user input, such as a password or a credit card number. The contents of sensitive parameters are not visible in logs nor in the user interface. The status is **Enabled** if checked.
Value Type	Note that you cannot change the Value Type after creation. From the list, select among supported types. • **user-input**. Specifies parameters whose data are provided by user-input. • **ignore**. Specifies parameters whose values the system does not check. • **XML**. Specifies parameters fetched from server and not editable. • **JSON**. Specifies parameters fetched from server and not editable. • **dynamic-content**.
Data Type	Specifies whether the screen displays parameters, whose data is provided by user-input, based on their data type. From the list, select among supported types. • **email**. Specifies parameters whose data must be text in email format only. • **alpha-numeric**. Specifies parameters whose data can be any text consisting of letters, digits, and the underscore character. The system provides additional options for this parameter's values to contain meta characters, or to match regular expressions. • **integer**. Specifies parameters whose data must be whole numbers only (no decimals). • **decimal**. Specifies parameters whose data is numbers only and can include decimals. • **phone**. Specifies parameters whose data can be text in telephone number format only.

7. In the Data type attributes area, specify whether or not the screen displays parameters, whose data is provided by user input, based on their data type.

Option	Description
Maximum Length	Specifies that this parameter's value has a restricted maximum length. The system considers as illegal any request that passes this parameter's value with a larger length. To set the maximum length, type in the box the maximum length (number of bytes) this parameter's value may contain. Leave empty to specify Any (no restrictions to the maximum length of this parameter's value).
Regular Exp.	Specifies, when checked (enabled), that an alpha-numeric parameter value includes a parameter pattern. To define the expression type it in the box. This is a positive regular expression that defines what is legal. The character length limit for this setting is 254. Specifies, when cleared (disabled), that this parameter does not include a pattern. The default is disabled.
Base64 encoding	You can enable the security policy to check whether user input parameter values contain a Base64 encoded string. If the value is indeed Base64 encoded, the system decodes this value and continues with its security checks. This option is available for user input Alpha-Numeric parameters and File Upload parameters. Specifies when enabled that the security policy checks the parameter's value for Base64 encoding, and decodes the value. The default setting is disabled.

8. In the Value Meta Character area, configure the security policy to allow or prohibit characters that may appear in the value of a specific parameter. You can override the global parameter value settings for the parameter.
   When you enable this setting, the system displays a list, from which you can select and name meta characters.
9. In the Attack Signatures area, enable **Attack Signature Checks** (default) to override the security policy settings of an attack signature for the parameter.

Option	Description
Value Meta Character	Specifies when checked (enabled) that you want to override the global parameter value settings for the parameter. After you enable this setting, the system displays a list of characters. The default is enabled.
Select Meta Character	From the list of characters, select your override.

10. In the Attack Signatures area, change the security policy settings for a specific parameter for this attack signature.

Option	Description
Attack Signature	Specifies when checked (enabled) that you want to override the security policy settings of an attack signature for the parameter. When this setting is enabled, the system displays a list of attack signatures. The default is enabled.
Attack Signature	From the list of characters, select your attack signature.

11. When you are finished, click **Save** to save the modifications and unlock the policy.

The parameter settings are updated to use the new settings, and any changes made are put into effect in the working configuration of the BIG-IQ® system.

Editing extractions settings

You can create and manage defined extractions and extraction properties. An *extraction* is a subcollection that isolates a parameter from an object. Other subcollections (such as parameters) reference extractions by name (not by URL). Use extractions when creating parameters of type dynamic content value. Using this screen, you can manage how the system extracts dynamic values for dynamic parameters from the responses returned by the web application server.

As with the other policy subcollections, the GUI for extractions displays the **Edit** button (unless the policy is already locked for edit). Press this button to lock the policy and the system displays the **Add**, **Delete** and **Unlock** buttons. Note that **Delete** button is disabled until you select at least one extraction.

Use the filter to search the column values of the extraction.

To create an extraction configuration, click **Edit**. The policy is placed under administrative lock and the fields become editable. For details, consult the following steps. Then, click **Add**. The system displays the Extraction Configuration screen in with the default values filled in.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Parameters screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Parameters**.
3. On the Parameters screen, select the **Extractions**tab.
4. Click the **Edit** button.
   The policy is placed under administrative lock, and the fields become editable.
5. To add a new extraction, click **Add**.
6. In the Create new extraction area, specify the **Name** attribute.

Option	Description
Name	Displays the name of the dynamic parameter for which the system extracts values from responses. If you are defining new extraction properties for a dynamic parameter, select one of the following from the list: • **New**. Type a name of the dynamic parameter in the adjacent box. This is the default. • **no name**. Specifies that the extraction is for the parameter **UNNAMED**.

7. In the Extracted Items Configuration area, specify the items from which the system should extract the values for dynamic parameters.

Option	Description
Extract From	• **File Types**. Specifies, when checked (enabled), that the system extracts the values of dynamic parameters from responses to requests for file types that exist in the security policy. To add a file type to be extracted, select an file type from the list, and click the **Add** button. That file type is added. Specifies when cleared (disabled), that the system does not extract values of dynamic parameters from file types. The default is disabled. • **URLs**. Specifies, when checked (enabled), that the system extracts the values of dynamic parameters from responses to requests for the listed URLs. To specify the URLs from which the system extracts dynamic parameter values, select either HTTP or HTTPS from the list, type the URL in the adjacent box, and click the **Add** button. If you enter a URL that does not yet exist in the security policy, the URL is added to the security policy. Specifies when cleared (disabled), that the system does not extract values of dynamic parameters from URLs. The default is disabled. • **RegEx**. Specifies, when checked (enabled), that the system extracts the values of dynamic parameters from responses to requests that match the listed pattern (regular expression). Type the regular expression in the box. Specifies when cleared (disabled), that the system does not extract values of dynamic parameters from regular expressions. The default is disabled.
Extract From All Items	Specifies when selected (enabled), that the system extracts the values of the dynamic parameters from all URLs found in the web application. Specifies when cleared (disabled), that the system extracts the values of the dynamic parameters from limited items found in the web application. The default is disabled.

8. In the Extracted Items Configuration area, configure the methods from which the system extracts the values for dynamic parameters.

Option	Description
Search in Links	Specifies, when checked (enabled), that the system searches for dynamic parameter values within links that appear in the response body. Specifies when cleared (disabled) that the system does not search for dynamic parameter values within links that appear in the response body. The default is checked (enabled).
Search Entire Form	Specifies, when checked (enabled), that the system searches for dynamic parameter values in the entire form found on a web page. Specifies, when cleared (disabled), that the system does not search for dynamic parameter values in web page forms. The default is checked (enabled).
Search Within Form	Specifies, when checked (enabled), that the system searches for dynamic parameter values in a specific location within forms found on a web page that contains the dynamic parameter. You must provide all of this information: • **Form Index**. Type the HTML index of the form that contains the dynamic parameter. • **Parameter Index**. Type the HTML index of the input parameter within the form that contains it.
Search Within XML	Specifies, when checked (enabled), that the system searches for dynamic parameter values within the URL's XML. Type the XPath specification in the XPath box. Specifies, when cleared (disabled), that the system ignores the URL's XML. The default is cleared (disabled).

Search Response Body	Specifies, when checked (enabled), that the system searches for dynamic parameter values in the body of the response. Specifies, when cleared (disabled), that the system does not search in the body of the response. Use the additional options to further refine the system's search. You can specify one or more of the following options, but you must specify the RegExp value if you enable this setting.
	• **Number of Occurrences**.
	  • **All**. Specifies a search for all incidences of the parameter values in the body of the request.
	  • **Number**. Specifies that the search is restricted to the number you type in the box.
	• **Prefix**. Specifies that the system extracts values only if they are preceded by the HTML segment you type in the box.
	• **Match Regular Expression Value**. Specifies that the system extract must match the parameter pattern (regular expression) you type in the box. The default is **.+?**.
	• **Suffix**. Specifies that the system extracts values only if they are followed by the HTML segment that you type in the box.

The extraction settings are updated to use the new settings, and any changes made are put into effect in the working configuration of the BIG-IQ® system.

Editing character sets settings

You can view and edit the characters (letters, digits, and symbols) available, and how the security policy responds to each when a request includes that character in parameter values. You configure the security policy to allow or disallow certain characters if they appear in parameter values and/or parameter names.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.

2. Navigate to the Character Sets screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Character Sets**.
3. Click **Edit**.
   The policy is placed under administrative lock and fields become editable.
4. Select the **Parameter Name** tab and edit the settings.

   Use these check boxes to configure the security policy to allow or disallow certain characters if they appear in wildcard parameter names. You can restore the system defaults even after you have made and saved changes on this tab.

Column label	Description
Hex	Shows the hexadecimal value of each character.
Character	Displays the character itself where applicable, otherwise displays a symbolic representation of the meta character, like TAB.
Allowed	This check box specifies which action the policy takes when it discovers one of these characters in parameter values. Select one of the following actions: • If checked (Allowed), the policy permits this character in all incoming requests. • If cleared (Disallowed), the policy does not permit this character in incoming requests.

5. Select the **Parameter Value** tab and edit the settings.

   Use this tab to configure the security policy to allow or disallow certain characters if they appear in parameter values. You can restore the system defaults even after you have made and saved changes on this tab.

Column label	Description
Hex	Shows the hexadecimal value of each character.
Character	Displays the character itself where applicable, otherwise displays a symbolic representation of the meta character, like TAB.
Allowed	This check box specifies which action the policy takes when it discovers one of these characters in parameter values. Select one of the following actions: • If checked (Allowed), the policy permits this character in all incoming requests. • If cleared (Disallowed), the policy does not permit this character in incoming requests.

6. When you are finished, click **Save** to save the modifications and unlock the policy.

The Character Sets settings are updated to use the new settings and any changes made are put into effect in the working configuration of the BIG-IQ® system.

Editing attack signatures settings

Attack signatures are rules or patterns that identify attacks or classes of attacks on a web application and its components. You can configure aspects of attack signatures to specify whether the signatures should be put into staging before being enforced, and whether or not to apply signatures to responses.

1. Log in with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Attack Signatures screen.
3. Click **Edit**.
   The policy is placed under administrative lock and the fields become editable.
4. Revise the settings as needed.

a) For **Signature Staging** select the **Enabled** check box to enable signature staging on the security policy.

If cleared (**Disabled**), then signature staging is disabled on the security policy.

b) Select the **Place updated signatures in staging** check box to have the system place new or updated signatures in staging for the number of days specified in the enforcement readiness period (specified in Policy properties).

The system does not enforce signatures that are in staging, even if it detects a violation. Instead, the system records the request information.

c) For **Attack Signature Set Assignment**, from the list, select one or more signature sets to specify the attack signatures the screen displays.

Select or clear the **Learn**, **Alarm**, and **Block** options to customize the settings.

You can click the **X** to the left of any signature set to remove it from the list you are using.

d) For the **Apply Response Signatures** setting, select a file type, and click the + icon.

5. When you are finished, click **Save** to save the modifications and unlock the policy.

The Attack Signatures settings are updated to use the new settings, and any changes you made are put into effect in the working configuration of the BIG-IQ® system.

Viewing attack signatures lists

Attack signatures are rules or patterns that identify attacks or classes of attacks on a web application and its components. You can view the list of security policy attack signatures that belong to signature sets assigned to the security policy, and you can edit individual attack signatures, including enabling or disabling attack signatures.

1. Log in with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Attack Signatures List screen: click **Web Application Security** > **Policy Editor**, in the Policy list, click the policy you want to edit, and from the Policy objects list, select **Attack Signatures List**.

   The screen displays a list of attack signatures with these properties:

   - **Name**. Signature name.
   - **ID**. Signature ID number. The system automatically provides the signature ID which cannot be changed.
   - **In Staging**. `Yes/No`. Specifies whether or not the signatures displayed are in staging, based on each signature's Perform Staging setting. This option is available only if signature staging is enabled in the policy.
   - **Enabled**. `Yes/No`. Specifies whether the screen displays signatures based on their Signature State setting. Specifies only whether the signatures displayed are enabled or disabled for the entire security policy, without regard to parameter-specific settings.
3. To edit an individual signature, click the signature name and then click **Edit**.
   The attack signature is placed under administrative lock. You can now edit the **Signature State** and **Perform Staging** settings.
4. Select the **Signature State** check box to specify only whether the signature displayed is enabled or disabled for the entire security policy, without regard to parameter-specific settings.
5. Select the **Perform Staging** check box to specify whether or not the signatures displayed are in staging.

   This option is available only if signature staging is enabled in the policy.
6. Hover over the Information icon to view additional information.
7. When finished, click **Save** to save the modifications and unlock the policy.

Any modified attack signatures settings are updated and put into effect in the working configuration of the BIG-IQ system.

Customizing attack signatures lists

You can use the filter on the Attack Signatures Lists screen to customize which attack signatures the system displays, so you can view only those that you are interested in.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Attack Signatures screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the Policy objects list, select **Attack Signatures Lists**.
3. Type search text in the filter text field and press Enter.
4. For finer control over the filtering options, click **Advanced Filter** and specify the appropriate settings.

Option	Description
Containing String	Specify whether the screen displays signatures based on a string found in the signature name. The default empty field indicates that the screen displays all signatures. Type part of the signature name in the field to view signature names that contain a specific string. This search is not case-sensitive.
Signature Type	Specify what type of signatures the system displays: • **All**. Specifies all signatures, which is the default. • **Request**. Specifies signatures that are configured to inspect the client request. • **Response**. Specifies signatures that are configured to inspect the server response.
Risk	Specify the risk level: • **Low**. Indicates the attack may assist the user in gathering knowledge to perpetrate further attacks, but does not cause direct damage or reveal highly sensitive data. • **Medium**. Indicates the attack may reveal sensitive data, or cause moderate damage. • **High**. Indicates the attack may cause a full system compromise, denial of service, and the like.
Enabled	Specify whether the screen displays signatures based on their **Enabled** setting. This filter specifies only whether the signatures displayed are enabled or disabled for the entire security policy, without regard to parameter-specific settings. • **All**. Specifies all signatures, which is the default. • **No**. Specifies signatures that are disabled. • **Yes**. Specifies signatures that are enabled.
Signature ID	Specify whether the screen displays signatures based on their ID number. (The system automatically provides the signature ID which cannot be changed.) The default empty field indicates that the screen displays all signatures. Type the signature ID in the field to view signatures with a specific signature ID.
User Defined	Specify whether the screen displays signatures based on who created them. • **All**. Specifies all signatures, which is the default. • **No**. Specifies system-defined signatures. • **Yes**. Specifies user-defined signatures.
Accuracy	Specify the accuracy levels:

Option	Description
	• **Low**. Indicates a high likelihood of false positives. • **Medium**. Indicates some likelihood of false positives. • **High**. Indicates a low likelihood of false positives.
In Staging	Specify whether the screen displays signatures based on each signature's **Perform Staging** setting. This filter specifies whether the signatures displayed are in staging or not. This option is available only if signature staging is enabled in the security policy. • **All**. Specifies all signatures, which is the default. • **No**. Specifies signatures currently not in staging. • **Yes**. Specifies signatures currently in staging.

5. When you have finished making your selections, click **Apply Filter**.
6. When you are finished, click **Save** to save the modifications and unlock the policy.

Any modified Attack Signatures settings are updated and put into effect in the working configuration of the BIG-IQ® system.

Adding security policies

You can use BIG-IQ® Web Application Security to add new application security policies for possible later deployment.

1. Navigate to **Web Application Security** > **Policy Editor**.

   Policies are listed on the Policies screen.
2. In the Policies screen, click **Add** to display a screen for creating a new policy.

   The newly-created policy contains only the editable configuration (the configuration deployed to the BIG-IP® device). It acquires the configuration default values from it.
3. Specify the following information about the new Web Application Security policy:
   a) Type the **Name** (required) of the security policy.
   b) Specify the **Partition** (required) to which the security policy belongs.

      Only users with access to a partition can view the objects that it contains. If the security policy resides in the `Common` partition, all users can access it.
   c) For **Application Language**, select the language encoding (required) for the web application, which determines how the security policy processes the character sets.

      The default language encoding determines the default character sets for URLs, parameter names, and parameter values.
   d) For **Enforcement Mode**, specify whether blocking is active or inactive for the security policy.

      You can enable or disable blocking for individual violations in the subsequent tables of settings and properties. If `transparent` appears, blocking is disabled for the security policy. This disables blocking for all options, and the check boxes to enable blocking are unavailable.
4. When you are finished editing the properties, click **Save**.
   This makes the remaining policy objects available for editing.
5. In the Policy objects list on the left, click the next object to edit, and then click the **Edit** button.

   For the **Attack Signatures List** object only, click the **Attack Signatures List** object, then in the Name column, click the signature name you want to edit, then click **Edit**.
6. Click **Save** to save the modifications to each policy object before moving to another one.
7. Click **Save** when you are finished editing.

The newly-created policy is added to the list of application security policies, and the new policy object exists in the working configuration of the BIG-IQ system. At this point, you can add it to any virtual server object in Web Application Security.

Importing security policies

Before you import a security policy from another system, make sure that the attack signatures and user-defined signatures are the same on both systems. You also need access to the exported policy file.

You can use BIG-IQ® Web Application Security to import security policies that were previously exported from another Application Security Manager™ system.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Policies screen: click **Web Application Security** > **Policy Editor**.
3. In the Policies screen, click the **Import** button.
4. In the Import Policy dialog box, select the security policy file by clicking **Choose File** or **Browse**, and navigating to the file location, or you can drag and drop a file to the **Drop Policy File Here** field.

   You can drag and drop a policy file onto the Source File area to view the content of the XML file.
5. Click **Import**.

After import, the policy is listed in the Policies screen. The uploaded policy will have the same name as the XML file.

If you replaced an existing policy, the imported security policy completely overwrites the existing security policy. Also, the imported policy is then associated with the virtual server and local traffic policy that was previously associated with the policy you replaced. The replaced policy is automatically archived with the inactive security policies.

Exporting security policies

You can use BIG-IQ® Web Application Security to export security policies. The exported security policy can be used as backup, or you can import it onto another system.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Policies screen: click **Web Application Security** > **Policy Editor**.
3. Select the check box to the left of the security policy you want to export.
   The **Export** button becomes active.
4. Click the **Export** button.
5. In the Export Policy dialog box, from the **Compatibility Version** list, select a version and click **Export**.

You can use the exported security policy as backup, or you can import it onto another system. Note that the exported security policy includes any user-defined signature sets that are in the policy, but not the user-defined signatures themselves.

Removing security policies

BIG-IQ® Web Application Security provides a way to remove ASM™ application security policies from the BIG-IQ database.

1. Log in to BIG-IQ Security with Administrator, Security Manager, or Web Application Security Manager credentials.
2. Navigate to the Policies screen: click **Web Application Security** > **Policy Editor**.
3. Select the check box to the left of the security policy you want to remove.
   The **Remove** button becomes active.
4. Click the **Remove** button.
5. In the Remove Policies dialog box, confirm the removal by clicking **Remove**.

The application security policy is removed from the BIG-IQ system, and can be managed locally.

Managing Signature Files

About signature files in BIG-IQ Web Application Security

Through BIG-IQ® Web Application Security, you can view and manage signature files and signature file updates centrally for multiple BIG-IP® devices. For each signature file, the system displays the file name, the file version, the version of BIG-IP with which it is compatible, and its source.

You can also update certain signature file settings.

Note: *You can lock for editing only the settings of signature files. You cannot edit signature files; therefore, there is no need to lock them.*

By managing signature files from the BIG-IQ platform, the administrator can spend less time on signature updates and can view the signatures update information in a single central location.

The BIG-IP system includes an attack signature pool and a bot signature pool. These pools include the system-supplied attack signatures and bot signatures, which are shipped with the BIG-IP Application Security Manager, and any user-defined signatures.

BIG-IQ Web Application Security fetches all new and relevant signature files automatically from an external server proxy configured from the system interface. It can then push the signature files to the relevant BIG-IP device or to multiple BIG-IP devices. It displays the signature version for each device.

Note: *You can lock (for edit or update) only the settings of signature files. You cannot edit signature files; therefore, there is no need to lock them.*

Viewing signature file properties

An application security policy exported from BIG-IP® Application Security Manager™ includes any attack signature sets that are in use by the policy, but not the actual signatures. Therefore, it is good practice to make sure that the attack signatures (both system-supplied and user-defined) are the same on the two systems. Use the BIG-IQ® Web Application Security Signature Files screen to view signature file properties.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.
2. Navigate to the Signature Files screen: click **Web Application Security** > **Signature Files**.
3. In the Signatures file screen, click a specific signature file to view properties.
4. When you are finished, click **Cancel**.

Signature file properties

Signature file properties are read-only and displayed for informational purposes only.

Property	Description
Name	Name of the signature file. Example: ASM-SignatureFile_20150917_152714.im123456789
Version	Version of the signature file. Example: 1445276000000

Property	Description
Compatibility	Version running on the BIG-IP® device with which the signature file is compatible. Example: BIG-IP 11.5.3
Source	F5 Networks. Example: F5

Updating and pushing signature files

You can use the BIG-IQ® Web Application Security Signature files screen to update the signature files and push them to BIG-IP® devices.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.
2. Navigate to the Signature Files screen: click **Web Application Security** > **Signature Files** > **Update Process/Push Status**.
3. In Update Process/Push Status, edit the settings as needed.
   a) From the **Interval** list, select how often the update should run.
   b) This field is pre-populated with the current date and time. To change, type a starting date and time in the format: `dd/mm/yyyy, hh:mm:ss AM (or PM)`. Example: `2/11/2016, 9:00:00 AM`. Or, click in the field to bring up a calendar, select a date, and use the **Hr** and **Min** controls to select the hour and minute. You can also select **Now**, **Clear**, and **Done**.
   c) Select the **Run Manual Sync** check box to have the system synchronize the configuration with the standby device when a signature file is pushed to the primary BIG-IP device.

   Note that some fields are display only:

   - **Last update**: Specifies the last time the file was updated, and whether the update was done manually or automatically; for instance, `Tue Oct 27 2015 10:40:27. (Triggered by scheduler)`.
   - **Next update**: Specifies the time of the next scheduled file update; for instance, `Tue Oct 27 2015 10:40:27`.
   - **Last run status**: Specifies the status of the last file update. Possible statuses include: `Passed`, `Failed`.
4. When you are finished, click **Save**.

   You can click **Cancel** to close the screen without saving your changes.

   Signatures are updated.
5. In the Current running task area, edit the settings as needed.
   a) **Run now**: To update the signature files and push them to the server, click the **Update & push** button.

      When the task has run to completion, the status displays as `Completed`. Ensure that the **Auto update enabled** check box on the Devices properties screen is checked, or updated files will not be pushed.
   b) Select the **Run Manual Sync** check box, and when a signature file is pushed to the primary BIG-IP device,the system synchronizes the configuration with the standby device.

   Note that the **Current status** setting specifies the status of the current file update. Possible statuses include: `Passed`, `Failed`.
6. When you are finished, click **Save**.

   You can click **Cancel** to close the screen without saving your changes.

Signatures are updated.

If a signature file is pushed to a clustered system, the configuration of the nodes is synchronized. The ASM™ configuration is deployed to the active device and then synchronized with the standby device.

Managing Custom Attack Signatures and Signature Sets

About custom attack signatures

Attack signatures are rules or patterns that identify attacks on a web application. When Application Security Manager® (ASM) receives a client request (or a server response), the system compares the request or response against the attack signatures associated with the security policy. If a matching pattern is detected, ASM™ triggers an attack-signature-detected violation, and either alarms or blocks the request, based on the enforcement mode of the security policy.

An ideal security policy includes only the attack signatures needed to defend the application. If too many are included, you waste resources on keeping up with signatures that you do not need. On the other hand, if you do not include enough, you might let an attack compromise your application without knowing it. If you are in doubt about a certain signature set, it is a good idea to include it in the policy rather than to omit it.

There are system-supplied signatures and custom (user-defined) signatures.

- *System-supplied signatures* enforce policies for best-known attacks. F5 Networks provides:
  - Over 2,500 signatures to guard against many different types of attacks and protect networking elements such as operating systems, web servers, databases, frameworks, and applications.
  - Signatures that include rules of attack that are F5 intellectual property.
  - Signatures that you can view but not edit or remove. Also, you cannot view the rules governing these signatures.
  - Periodic updates.

  To learn more about system-supplied attack signatures, consult the BIG-IP® system documentation.
- *Custom (user-defined) signatures* are created by your organization for specific purposes in your environment. These signatures:
  - Are added to the attack signatures pool where F5 Networks stores them along with the system-supplied signatures.
  - Must adhere to a specific rule syntax (like system-supplied signatures).
  - Can be combined with system-supplied signatures or system-supplied sets to create custom signature sets.
  - Are never updated by F5 Networks, but are carried forward as-is when the system is updated to a new software version.

In BIG-IQ® Web Application Security, you can obtain system-supplied or custom attack signatures through the device discovery process. These signatures are automatically deployed to all policies when the system performs a deployment.

Creating custom attack signatures

Custom (user-defined) attack signatures can handle security policy enforcement unique to your networking environment, emergency situations, or analysis of specific activity on the network. If your organization needs a custom attack signature, you can use the BIG-IQ® Web Application Security Policy Editor to create one. You can then assign the new signature to system-supplied or custom attack signature sets.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.

2. Navigate to the Policy Editor screen: click **Web Application Security** > **Policy Editor**.
3. On the left, click **Attack Signatures**.
   The Attack Signatures screen opens and lists all signatures available to the BIG-IQ system. The system lists the system-supplied (factory) signatures in static black text, and lists any custom signatures in blue text. Blue indicates a hyperlink. System-supplied signatures are locked as indicated by a green padlock icon.

   Note that you can click anywhere in a row to display the Signature Properties tab and the Documentation tab for the signature.
4. At the right of the screen, click **Add** and use the Attack Signatures - New Item screen to supply the required information.
   The screen displays a blank template for signature properties.
5. On the Signature properties tab, fill in fields and select options to define the new custom signature:
   a) In the **Name** field, type a unique name.

      If you attempt to create a custom signature with the same name as a system-supplied signature, you will receive an error message and the system will not create the signature.
   b) In the **Description** field, type an (optional) description.
   c) From the **Signature Type** list, select what the signature should examine:

      - **Request**. Use this signature to examine requests only.
      - **Response**. Use this signature to examine responses only.

   d) For **Attack Type**, select the threat classification.
   e) Select the **Systems** that you want protected by the signature: use the Move button to shift your choices from the **Available** list to the **Enabled** list.
   f) For the **Rule** setting, type a rule, according to the syntax guidelines, to specify the content of the signature.

      The rule is the heart of the attack signature. All attack signatures must adhere to the F5 attack signature syntax. Refer to the BIG-IP® system documentation on signature options and signature syntax for details.
   g) For **Accuracy**, select the level that you want for the signature.

      The accuracy level indicates the ability of the attack signature to identify the attack, including susceptibility to false-positive alarms. Higher accuracy results in fewer false positives.
   h) For **Risk**, select the level of potential damage this attack might cause, if it were successful.

      - **Low** indicates the attack may assist the user in gathering knowledge to perpetrate further attacks, but does not cause direct damage or reveal highly sensitive data.
      - **Medium** indicates the attack may reveal sensitive data, or cause moderate damage.
      - **High** indicates the attack may cause a full system compromise, denial of service, and the like.

   i) The **User-defined** field specifies whether the screen displays signatures based on who created them. Currently, it defaults to **Yes**, indicating that the signature was created by a user. You cannot change the setting.
6. When you are finished, click **Save** to save the new custom attack signature.

   Clicking **Save and Close** prompts the system to return to the Attack Signatures screen.

   Custom signatures appear in blue and are hyperlinks to an edit screen. Click anywhere on the row except the link to display Signature Properties at the bottom of the screen.

The system places the new custom attack signature into the attack signature pool, and adds it to the signature sets for the systems you specified. The custom signature is put in staging for all policies that have this signature in their assigned signature sets. It is a good idea to make sure that the system added the new signature to the appropriate security policies.

About signature staging

When you first activate a security policy, the system places the attack signatures into staging (if staging is enabled for the policy). *Staging* means that the system applies the attack signatures to the web application traffic, but does not apply the blocking policy action to requests that trigger those attack signatures. The default staging period is seven days.

Whenever you add or change signatures in assigned sets, those signatures are also placed in staging. You also have the option of placing updated signatures in staging.

Placing new and updated attack signatures in staging helps to reduce the number of violations triggered by false-positive matches. When signatures match attack patterns during the staging period, the system generates learning suggestions. If you see that an attack signature violation has occurred, you can view and evaluate these attack signatures. After evaluation, if the signature is a false-positive, you can disable the signature, and the system no longer applies that signature to traffic for the corresponding web application. Alternately, if the detected signature match is legitimate, you can enable the corresponding attack signature.

Note: *Enabling the signature removes it from staging, and puts the blocking policy into effect.*

About custom attack signature sets

An *Attack signature set* is a group of attack signatures. Rather than applying individual attack signatures to a security policy, you can apply one or more attack signature sets. The Application Security Manager™ ships with several system-supplied signature sets.

Each security policy has its own attack signature set assignments. By default, a generic signature set is assigned to new security policies. You can assign additional signature sets to the security policy. Sets are named logically so you can tell which ones to choose. Additionally, you can combine custom attack signatures with system-supplied signatures or system-supplied sets to create custom signature sets.

An ideal security policy includes only the attack signature sets needed to defend the application. If too many are included, you waste resources on keeping up with signatures that you do not need. On the other hand, if you do not include enough, you might let an attack compromise your application without knowing it. If you are in doubt about a certain signature set, it is a good idea to include it in the policy rather than to omit it.

In Web Application Security, you can obtain system-supplied or custom attack signature sets through the device discovery process. You can assign these sets to security policies. Then, you can deploy those policies to BIG-IP® devices.

Add custom attack signature sets

You can use the Web Application Security policy editor to add custom (user-defined) attack signature sets. Like system-supplied signature sets, *custom signature sets* contain signatures from the signature pool. Once you create a custom signature set, you can apply it to the security policy to protect web applications against known attacks.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.
2. At the top left of the screen, select **Web Application Security** from the BIG-IQ menu.

   The Web Application Security Policy Editor screen opens.

3. On the left, click **SIGNATURE SETS**.
   The default, system-supplied signature sets are displayed on the Signature Sets screen, along with any user-defined sets. By default, the system lists signature sets in alphabetical order by name.
4. Click **Add** and use the Signature Sets - New Item screen to supply the required information.
5. On the Properties tab, type a unique name for the signature set.
6. From the **Type** list, select how to create the signature set.
   - Select **Filter-based** to create a signature set by using a filter only.
   - Select **Manual** to manually assign signatures to a signature set.

   Selecting **Manual** causes the Signatures Filter tab to be hidden, since it will not be used, and changes the fields displayed on the Signatures tab.

   You can create or edit a signature set by configuring a filter to select from the signature pool signatures that meet specific criteria. Using a filter enables you to focus on the criteria that define the signatures you are interested in. When you update the signatures database, the system also updates any signature sets affected by the update.
7. For **Default Blocking Actions**, select the blocking actions you want the system to enforce for the set when you associate it with a new security policy.

   The **Learn**, **Alarm**, and **Block** actions take effect only when you assign this set to a new security policy. If this set is already assigned to an existing security policy, these settings have no effect.
8. If you want the system to automatically include this set in any newly-created security policies, enable the **Assign to Policy by Default** setting.
9. Click the Signatures Filter tab, and select the filter options to narrow the scope of the signatures to include in the new signature set. This tab is only displayed when the signature set type is set to **Filter-based**.
   a) Select a **Signature Type** to include the type of signatures the system displays.
      - **All** traffic is the default.
      - **Request** only. Signatures that are configured to inspect the client request.
      - **Response** only. Signatures that are configured to inspect the server response.

   b) From the **Attack Type** list, specify the threat classifications for which to include signatures in the set.
      - Select **All** for signatures with all Attack Type values, which is the default.
      - Select an attack type for signatures configured to protect against that specific attack type.

   c) From the **Systems** lists, specify the systems (for example web applications, web server databases, and application frameworks) that you want protected by the set.
   d) From the **Accuracy** list, select the accuracy association.
      - **All** specifies signatures that match all accuracy levels, which is the default.
      - **Equals** specifies signatures whose accuracy levels exactly match the accuracy level you set.
      - **Greater Than/Equal To** specifies signatures whose accuracy levels are more precise than, or the same as, the accuracy level you set.
      - **Less Than/Equal To** specifies signatures whose accuracy levels are less precise than, or the same as, the accuracy level you set.

   e) From the resulting list, select the accuracy level.
      - **Low** indicates a high likelihood of false positives.
      - **Medium** indicates some likelihood of false positives.
      - **High** indicates a low likelihood of false positives.

   f) From the **Risk** list, select the risk association.
      - **All** specifies signatures that protect against attacks of all risk levels, which is the default.

- **Equals** specifies signatures whose risk levels exactly match the risk level you set.
- **Greater Than/Equal To** specifies signatures whose risk levels are higher than, or the same as, the risk level you set.
- **Less Than/Equal To** specifies signatures whose risk levels are lower than, or the same as, the risk level you set.

g) From the resulting list, select the risk level; the level of potential damage for attacks protected by the signatures in the set.

- **Low** indicates the attack may assist the user in gathering knowledge to perpetrate further attacks, but does not cause direct damage or reveal highly sensitive data.
- **Medium** indicates the attack may reveal sensitive data, or cause moderate damage.
- **High** indicates the attack may cause a full system compromise, denial of service, and the like.

h) For **User-defined**, specify whether to include signatures based on who created them: the user (**Yes**), the system (**No**), or both (**All**).

i) For **Update Date**, specify whether to include all signatures in the set based on the date the signature was changed (**All**), only signatures added before the date the signature was changed (**Before**), or only signatures added after the signature was changed (**After**).

If specifying **Before** or **After**, use the calendar icon to specify a date.

10. Click the Signatures tab.

The Signatures tab appears differently depending on whether the signature set is user-defined (also called custom) or system-supplied (also called a factory signature set), and if user-defined, then whether **Type** on the Properties tab is set to **Filter-based** or **Manual**.

- If the signature set is system-supplied, the Signatures tab lists the signatures selected for the signature set.
- If the signature set is user-defined and **Type** is set to **Filter-based**, the Signatures tab lists the signatures selected using the criteria set by the Signature Filters tab. The list content changes dynamically based on changes to the Signature Filters tab.
- If the signature set is user-defined and **Type** is set to **Manual**, the Signatures tab lists a selectable list of signatures. If you want to view only a subset of the signatures, click **Signatures Advanced Filter** at the top of the Signatures tab to filter the signatures shown.

11. In the Included Policies tab, view the policies (if any) that enforce this signature set.

Each security policy enforces one or more signature sets. The decision about which signature sets to include occurs when creating a security policy. You can assign additional signature sets to the security policy.

12. When you are finished, click **Save** to save the new custom attack signature set.

Clicking **Save and Close** prompts the system to return to the Signature Sets screen and display the new set.

Sets are listed in alphabetical order; custom sets appear in blue.

The new signature set is added to the list of signature sets that are available on the system, and is available to be applied when creating new security policies. If, in the future, you no longer need a custom signature set, you can delete it. Note that when you delete a custom signature set, you are deleting the set; you are not deleting the signatures that made up the set.

Edit custom attack signature sets

You can use the Web Application Security policy editor to edit custom attack signature sets. Once you edit a custom signature set, you can apply it to the security policy to protect your web applications in ways that are unique to your needs.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.

2. At the top left of the screen, select **Web Application Security** from the BIG-IQ menu.

   The Web Application Security Policy Editor screen opens.
3. On the left, click **Signature Sets.**
   The system displays the default, system-supplied signature sets, along with any user-defined sets. By default, the system lists signature sets in alphabetical order by name.
4. Click the name of the signature set that you want to change and use the Signature Sets screen to modify the settings.
5. On the Properties tab, revise the settings for this custom attack signature set, as needed.

   Note that **Name** and **Category** are not editable fields.
6. From the **Type** list, you can modify how to create the signature set.

   - Select **Filter-based** to create a signature set by using a filter only.
   - Select **Manual** to manually assign signatures to a signature set.

   Selecting **Manual** causes the Signatures Filter tab to be hidden since it will not be used, and changes the fields displayed on the Signatures tab.

   You can create or edit a signature set by configuring a filter to select from the signature pool signatures that meet specific criteria. Using a filter enables you to focus on the criteria that define the signatures you are interested in. When you update the signatures database, the system also updates any signature sets affected by the update.
7. For **Default Blocking Actions**, select the blocking actions you want the system to enforce for the set when you associate it with a new security policy.

   The **Learn**, **Alarm**, and **Block** actions take effect only when you assign this set to a new security policy. If this set is already assigned to an existing security policy, these settings have no effect.
8. If you want the system to automatically include this set in any newly-created security policies, enable the **Assign to Policy by Default** setting.
9. Click the Signatures Filter tab, and select the filter options to narrow the scope of the signatures to include in the new signature set.

   This tab is only displayed when the signature set type is set to **Filter-based**.

   a) Select a **Signature Type** to include the type of signatures the system displays.

      - **All** traffic is the default.
      - **Requests** only. Include signatures that are configured to inspect the client request.
      - **Responses** only. Include signatures that are configured to inspect the server response.

   b) From the **Attack Type** list, specify the threat classifications for which to include signatures in the set.

      - Select **All** for signatures with all Attack Type values, which is the default.
      - Select an attack type for signatures configured to protect against that specific attack type.

   c) From the **Systems** lists, specify the systems (for example web applications, web server databases, and application frameworks) that you want protected by the set.

   d) From the **Accuracy** list, select the accuracy association.

      - **All** specifies signatures that match all accuracy levels, which is the default.
      - **Equals** specifies signatures whose accuracy levels exactly match the accuracy level you set.
      - **Greater Than/Equal To** specifies signatures whose accuracy levels are more precise than, or the same as, the accuracy level you set.
      - **Less Than/Equal To** specifies signatures whose accuracy levels are less precise than, or the same as, the accuracy level you set.

   e) From the resulting list, select the accuracy level.

      - **Low** indicates a high likelihood of false positives.

- **Medium** indicates some likelihood of false positives.
- **High** indicates a low likelihood of false positives.

f) From the **Risk** list, select the risk association.

- **All** specifies signatures that protect against attacks of all risk levels, which is the default.
- **Equals** specifies signatures whose risk levels exactly match the risk level you set.
- **Greater Than/Equal To** specifies signatures whose risk levels are higher than, or the same as, the risk level you set.
- **Less Than/Equal To** specifies signatures whose risk levels are lower than, or the same as, the risk level you set.

g) From the resulting list, select the risk level; the level of potential damage for attacks protected by the signatures in the set.

- **Low** indicates the attack may assist the user in gathering knowledge to perpetrate further attacks, but does not cause direct damage or reveal highly sensitive data.
- **Medium** indicates the attack may reveal sensitive data, or cause moderate damage.
- **High** indicates the attack may cause a full system compromise, denial of service, and the like.

h) For **User-defined**, specify whether to include signatures based on who created them: the user (**Yes**), the system (**No**), or both (**All**).

i) For **Update Date**, specify whether to include all signatures in the set based on the date the signature was changed (**All**), only signatures added before the date the signature was changed (**Before**), or only signatures added after the signature was changed (**After**).

If specifying **Before** or **After**, use the calendar icon to specify a date.

10. Click the Signatures tab.

The Signatures tab appears differently depending on whether the signature set is user-defined (also called custom) or system-supplied (also called a factory signature set), and if user-defined, then whether **Type** on the Properties tab is set to **Filter-based** or **Manual**.

- If the signature set is system-supplied, the Signatures tab lists the signatures selected for the signature set.
- If the signature set is user-defined and **Type** is set to **Filter-based**, the Signatures tab lists the signatures selected using the criteria set by the Signature Filters tab. The list content changes dynamically based on changes to the Signature Filters tab.
- If the signature set is user-defined and **Type** is set to **Manual**, the Signatures tab lists a selectable list of signatures. If you want to view only a subset of the signatures, click **Signatures Advanced Filter** at the top of the Signatures tab to filter the signatures shown.

11. Click the Included Policies tab, and view the policies (if any) that enforce this signature set.

Each security policy enforces one or more signature sets. The decision about which signature sets to include occurs when creating a security policy. You can assign additional signature sets to the security policy.

12. When you are finished, click **Save** to save the new custom attack signature set.

Clicking **Save and Close** prompts the system to return to the Signature Sets screen and display the new set.

The system lists sets in alphabetical order, custom sets appear in blue

The edited signature set is available for application when creating new security policies. If, in the future, you no longer need a custom signature set, you can delete it. Note that when you delete a custom signature set, you are deleting the set; you are not deleting the signatures that made up the set.

Signatures advanced filter properties

The **Signatures Advanced Filter** option and properties are only available on the Signatures tab when the signature set type is manual.

Signatures Advanced Filter Property	Description
Signature Type	Specifies what type of signatures to include in the signature set. • Select **All** to include both requests and responses. • Select **Request** to include only requests. • Select **Response** to include only responses.

Signature Scope	Specifies whether the system displays all signatures, or only those that do, or do not, apply to parameters, cookies, XML documents, JSON data, GWT data, headers, URI content, and request or response content.
	• Select **All** to include all signatures, which is the default.
	• Select **Parameter** to specify whether the system displays signatures that apply to alpha-numeric user-input parameters. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to parameters.
	  • **Yes** specifies only signatures that apply to parameters.
	• Select **Cookie** to specify whether the system displays signatures that apply to allowed cookies. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to allowed cookies.
	  • **Yes** specifies only signatures that apply to allowed cookies.
	• Select **XML** to specify whether the system displays signatures that apply to XML documents. XML documents may appear as the values of XML parameters defined in the security policy, or as the body of requests to URLs to be parsed as XML, as defined in the security policy. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to XML documents.
	  • **Yes** specifies only signatures that apply to XML documents.
	• Select **JSON** to specify whether the system displays signatures that apply to JSON data. JSON data may appear as the values of JSON parameters defined in the security policy, or as the body of requests to URLs to be parsed as JSON, as defined in the security policy. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to JSON data.
	  • **Yes** specifies only signatures that apply to JSON data.
	• Select **GWT** to specify whether the system displays signatures that apply to GWT data. GWT data may appear as the body of requests to URLs to be parsed as GWT, as defined in the security policy. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to GWT data.
	  • **Yes** specifies only signatures that apply to GWT data.
	• Select **Header** to specify whether the system displays signatures that apply to headers. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to headers.
	  • **Yes** specifies only signatures that apply to headers.
	• Select **URI** to specify whether the system displays signatures that apply to URI content. Then select **No** or **Yes**.
	  • **No** specifies only signatures that do not apply to URI content.
	  • **Yes** specifies only signatures that apply to URI content.
	• Select **Request Content** to specify whether the system displays signatures that apply to the entire request content. Then select **No** or **Yes**.

Signatures Advanced Filter Property	Description
	• **No** specifies only signatures that do not apply to the entire request content. • **Yes** specifies only signatures that apply to the entire request content. • Select **Response Content** to specify whether the system displays signatures that apply to the entire response content. Then select **No** or **Yes**. • **No** specifies only signatures that do not apply to the entire response content. • **Yes** specifies only signatures that apply to the entire response content.
Attack Type	Specifies which attack type should be included in the set. Select **All** to include all attack types.
Systems	Specifies the systems (for example web applications, web server databases, and application frameworks) that you want protected by the set.
Accuracy	Specifies the accuracy level of the signature. Higher accuracy results in fewer false positives. • Select **All** to specify that all signatures should be included, regardless of accuracy level. • Select **Equals** to specify signatures with a single accuracy level, then select the accuracy level to be **Low**, **Medium**, or **High**. • Select **Greater Than/Equal To** to specify that the accuracy level of the signatures should be greater than or equal to the specified accuracy level, then select the level to be **Low**, **Medium**, or **High**. • Select **Less Than/Equal To** to specify that the accuracy level of the signatures should be less than or equal to the specified accuracy level, then select the level to be **Low**, **Medium**, or **High**.
Risk	Specifies the level of potential damage that the signature protects against. • Select **All** to specify that all signatures should be included, regardless of risk level. • Select **Equals** to specify a single risk level, then select the risk level to be **Low**, **Medium**, or **High**. • Select **Greater Than/Equal To** to specify that the risk level should be greater than or equal to the specified risk level, then select the level to be **Low**, **Medium**, or **High**. • Select **Less Than/Equal To** to specify that the risk level should be less than or equal to the specified risk level, then select the level to be **Low**, **Medium**, or **High**.
User-defined	Specifies whether to include attack signatures based on who created them. • Select **All** to specify that all signatures should be included, including those defined by the system and by users. • Select **Yes** to specify that only user-defined signatures should be included. • Select **No** to specify that only system-defined signatures should be included.
Update Date	Specifies whether to include signatures in the set based on when the signature was last updated or added. • Select **All** to include all signatures, regardless of when they were last updated.

Signatures Advanced Filter Property	Description
	• Select **Before** to include all signatures updated before a specified date, and then select the date using the displayed **Select Date** button. • Select **After** to include all signatures updated after a specified date, and then select the date using the displayed Select Date button.
Signatures	Specifies the signatures that should be included in the signature set. The available signatures list displayed changes based on the **Signatures Advanced Filter** settings. You can use the **Filter** field above the Available list to search for particular signatures. Add signatures to the signature list by moving them from the Available list to the Selected list.

Assign custom attack signature sets

You use the Web Application Security policy editor to assign a custom attack signature set to a policy.

Each security policy enforces one or more attack signature sets. You can assign additional attack signature sets to the security policy.

1. Log in with Administrator, Security Manager, or Web App Security Manager credentials.
2. Navigate to the Policy Editor screen: click **Web Application Security** > **Policy Editor**, select a policy name, and from the **Policy objects** list, select **Attack Signatures Configuration**.
3. Click **Edit**.
   The policy is placed under administrative lock and fields become editable.
4. From the **Attack Signature Set Assignment** list, select attack signature sets to assign to the policy. Any newly-created custom signature sets appear in the list.
5. When you are finished, click **Save** to save the new assignment and unlock the policy.

The system assigns the signature sets to the security policy, and the blocking policy applies to all of the signatures in the signature set. Any changes made subsequently are put into effect in the working configuration of the BIG-IQ Centralized Management system.

Managing Virtual Servers in BIG-IQ Web Application Security

About virtual servers in the policy editor

BIG-IQ® Web Application Security displays virtual servers in the policy editor for each discovered BIG-IP® device, and enables you to view the properties for these virtual servers and manage the policies attached to them.

For each device discovered, the BIG-IQ system creates an extra virtual server to hold all security policies not related to any virtual server in the discovered device.

Displaying virtual server properties and managing policies

With BIG-IQ® Web Application Security, you can view virtual server properties and manage policies attached to the virtual server.

1. Log in to the BIG-IQ® system with your user name and password.
2. At the top left of the screen, select **Web Application Security** from the BIG-IQ menu.
3. Click **Policy Editor**, and then from the list on the left, click **Virtual Servers**.

   A list of the virtual servers displays.
4. Click the name of the virtual server to view properties or to manage the policy attached to it.

Property	Description
Name	Name of the virtual server.
Full Path	Full path, including partition, to the virtual server on the BIG-IP® device.
IP Address	Self IP address of the BIG-IP device.
Device	Fully qualified domain name of the BIG-IP device.
Attached Policy	Lists any attached policy. To attach a policy, click **Add Policy**. Alternatively, drag-and-drop a policy from those listed in the Policy Editor toolbox at the bottom of the page to the Attached Policy row. To delete a policy from the virtual server, click the **X** to the right of the policy.

5. Click **Save** to save the virtual server changes, or click **Save & Close** to save the virtual server changes and return to the Virtual Servers screen.

Managing Event Logs for Web Application Security

How do I manage events with a Logging Node?

Viewing the events as implemented on BIG-IQ® eases processing of Web Application Security events, and provides a way to obtain useful insights regarding the activity on client applications. The BIG-IQ platform enables a single view of all filters and log entries (and details for each entry) from multiple BIG-IP® devices.

It also provides a more intuitive navigation path through the log items.

To properly configure event log viewing, you should:

- Discover and activate a BIG-IQ Logging Node.
- License and provision a BIG-IQ Logging Node.
- Define an external machine to which periodic data snapshots are sent.
- Configure a BIG-IP system to collect events and send them to the BIG-IQ Logging Node. Part of this configuration includes a virtual server configured with an FPS profile.
- Configure the BIG-IQ for HA, if needed.
- A minimum of 1 logging node is required, but up to 5 logging nodes are supported.

An FPS profile created on the BIG-IP is used to configure when and how events are triggered on the client. The profile then directs security events to a BIG-IQ Logging Node, and the BIG-IQ system retrieves them from that node.

Logging Node uses a search engine that requires separate services for management and traffic. Keeping those services on separate networks reduces unnecessary congestion. The network designs described here are not required, but considered best practice.

BIG-IQ Networks

- A cluster management network to perform Elasticsearch configuration and status operations
- A cluster traffic network for inter-node communication

Logging Node Networks

- A cluster management network to perform Elasticsearch configuration and status operations
- A cluster traffic network for inter-node communication
- A listener network to handle inbound data traffic

This figure illustrates the network topology required to deploy a logging node for your events.

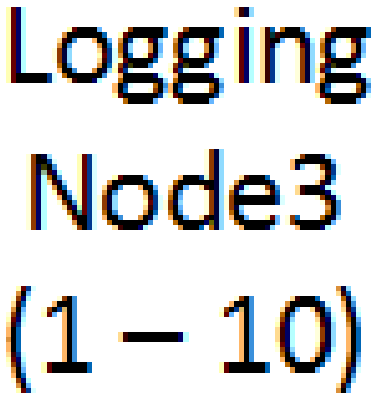

Figure 1: Logging Node network topology

Important: *F5 Networks strongly recommends that the Listener Network and Management Networks be separate. This separation, can help with data protection and management network availability in case the Listener Network is flooded with data.*

What is a BIG-IQ Logging Node?

A *BIG-IQ Logging Node* is a specially-provisioned BIG-IQ® system, which runs the same software version as the BIG-IQ device that you use to manage your security and the rules that determine your alert responses. After you provision the BIG-IQ Logging Node, you discover it from BIG-IQ and then add the service. After you configure the service, the logging node stores events from one or more BIG-IP® systems. The BIG-IQ system can then retrieve and manage those events.

Note: *The software version on the Logging Node must be the same as the version on its partner BIG-IQ system. If you need to upgrade the Logging Node, follow the instructions in Upgrading BIG-IQ Systems.*

Discover and activate a logging node

Using BIG-IQ® System Management, you can discover a Logging Node and add it to the Logging Group. The BIG-IQ can then access all event on the discovered Logging Node. You can then receive these events from multiple BIG-IP® systems. This unified view makes browsing easier, and provides a complete view of application event activity.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING**.
4. Under **BIG-IQ LOGGING**, select **Logging Nodes**.
5. Click **Add Node**.
6. On the New Logging Device screen, fill in as appropriate:
   a) In **IP Address**, type the management IP address.
   b) In **User Name**, type the user name for an administrator on the Logging Node (for example, `admin`.

c) In **Password**, type the password for an administrator on the Logging Node (for example, `admin`.
d) In **Transport Address**, type the IP address of the logging node internal self IP address.
e) For **Transport Port**, the default value is 9300. The BIG-IQ uses this port for internal polling and communication with the logging nodes.

7. Click the **Add** button at the bottom of the screen to add the Logging Node to the system. Or, click **Discard** to cancel the operation.

Note: *This operation might take a minute or two.*

---

8. Repeat these 7 steps for each Logging Node you want to configure.
9. To activate this logging node for the service you want to monitor, in the Services column, click **Add Services**.
   The Logging Node Services screen opens.
10. For the service you want to add, confirm that the **Listener Address** correctly specifies the external self IP address of the Logging Node, and click **Activate**.
    When the service is successfully added, the **Service Status** changes to `Active`.
11. Click **Close**.

Once discovered and activated, this logging node collects the events generated by the configured BIG-IP systems. Thus, BIG-IQ provides a single view of all event entries.

Important: *The **Total Document Count** is not a report of the number of alerts sent to the Logging Node. Instead, it is a sum of various document types sent to the Logging Node. Alerts are included in this list, but this total includes other document types as well.*

Modifying event log indices

Event log indices determine the physical characteristics of what is sent to the Logging Node.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, select **Logging Configuration**.
   The Logging Configuration screen opens to display the current state of the logging node cluster defined for this device.
4. In the Web Application Security row in the bottom half of the screen, click the **Configure** button.
   The Web Application Security Indices screen opens.
5. For the **Rotation Type**, keep the default setting: **Size Based**.
6. For the **Max Index Size**, type the maximum size of the indices you want to send to the logging node.

   For example, if you type `1000`, when the event log data reaches a size of 1 Gig, it is sent to the logging node.
7. For the **Retained Index Count**, type the total number of indexes you want to store on the logging node.

   The maximum amount of data stored on the Logging Node is the product of the **Max Index Size** and the **Retained Index Count**. When the amount of data reaches this size, the oldest event data is truncated or discarded.
8. Click **Save** to save the indices configuration settings.

Define event snapshot storage locations

Before you can configure the external snapshot storage location, you need the following information on the machine you will use to store the event snapshots:

- storage-machine-IP-address

- storage-file-path
- Read/Write permissions for the storage file path

You need snapshots of your alert data to perform software upgrades, hotfix upgrades, and to restore your .

When event snapshots are created, they need to be stored on a machine other than the Logging Node that stores the events. You define the location for the snapshot by editing the `fstab` file on your Logging Node machines and on the BIG-IQ® and HA peer devices.

Important: *You must perform this task on each Logging Node device, on the BIG-IQ device, and on the BIG-IQ HA peer.*

---

1. On the first device, in the folder `/var/config/rest/elasticsearch/data/`, create a new folder named `essnapshot`.
   `mkdir /var/config/rest/elasticsearch/data/essnapshot`
2. Edit the `/etc/fstab` file to add `/var/config/rest/elasticsearch/data/essnapshot`.
   For example, `//<storage machine ip-address>/<storage-file-path> /var/config/rest/elasticsearch/data/essnapshot cifs iocharset=utf8,rw,noauto,uid=elasticsearch,gid=elasticsearch, 0 0`
3. Run the mount command to mount the snapshot storage location to the new folder.
   For example, from `/var/config/rest/elasticsearch/data` type: `mount essnapshot`.
4. Confirm that the `essnapshot` folder has full read, write, and execute permissions, (specifically `Chmod 777 essnapshot`), and that the owner and group are `elasticsearch` for this folder.
   For example, `ls-l` would yield: `drwxrwxrwx 3 elasticsearch elasticsearch 0 Apr 25 11:27 essnapshot`.
5. Create a test file to confirm that the storage file-path has been successfully mounted.
   For example: `touch testfile`.
   The test file should be created on the storage machine at the location storage file path.
6. Repeat these five steps for each Logging Node, the BIG-IQ, and the BIG-IQ HA peer.

The storage location should now be accessible to the BIG-IQ devices and to the logging node machines.

Define Web Application Security snapshot schedules

Before you define snapshot schedules, you must have defined the snapshot storage locations.

Snapshots of the events sent to your Logging Nodes are an essential safeguard for your data. If the machine that stores the events fails, the data can be restored using these snapshots. These snapshots are created based on the snapshot schedules you define. F5 recommends that you schedule snapshots at least every 6 hours and retain at least 4 snapshots.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING**.
4. Under **BIG-IQ LOGGING**, select **Logging Configuration**.
5. For the **Snapshot Schedules** setting, click **Create**.
   The New Logging Snapshot screen opens.
6. For the **Snapshot Name Prefix**, type the string that you want to use to identify the snapshots created by this schedule.

   For example `snapshot_`.
7. In **Snapshots to Keep**, specify the number of snapshots that you want to accumulate before they are deleted for space constraints.

   For example, if you specify 25, then the system will retain a maximum of 25 snapshots before it starts to delete older snapshots as new snapshots are created. You can save up to 100.

8. Define how you want the snapshots scheduled.

Option	Description
Schedule the interval at which you want to create snapshots:	You schedule the system to take snapshots indefinitely. Snapshots are created at the frequency you specify. 1. Select **Repeat Interval**. 2. Specify the **Snapshot Frequency**. 3. Select a time increment. For example, if you set the frequency to **6** and **Hours**, the first log event data snapshot is taken immediately (on **Save**). Subsequent snapshots are taken every 6 hours.
Schedule specific days on which you want to create snapshots:	You schedule the system to take snapshots on specific days. 1. Select **Days of the Week**. 2. For the **Days of the Week** setting, select the days on which you want backups to occur. 3. For the **Start Date**, select the time (date, hour, minute, and AM or PM) on which you want backups to start.

9. Click **Save** to save the new schedule.

How do I license and do the basic setup to start using a Logging Node?

The BIG-IQ® Logging Node runs as a virtual machine in supported hypervisors, or on the BIG-IQ 7000 series platform. You license the Logging Node using the base registration key you purchased. The *base registration key* is a character string that the F5 license server uses to provide access to Logging Node features.

You license Logging Node in one of the following ways:

- If the system has access to the internet, you can have the Logging Node contact the F5 license server and automatically activate the license.
- If the system is not connected to the internet, you can manually retrieve the activation key from a system that is connected to the internet, and transfer it to the Logging Node.
- If your Logging Node is in a closed-circuit network (CCN) that does not allow you to export any encrypted information, you must open a case with F5 support.

When you license the Logging Node, you:

- Specify a host name for the system.
- Assign a management port IP address.
- Specify the IP address of your DNS server and the name of the DNS search domain.
- Specify the IP address of your Network Time Protocol (NTP) servers and select a time zone.
- Change the administrator's default admin and root passwords.

Automatically license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`http://www.f5.com`).

If the BIG-IQ® system is connected to the public internet, you can follow these steps to automatically perform the initial license activation and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://<management_IP_address>`, where `<management_IP_address>` is the address you specified for device management.
2. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

3. Click **Activate**.
   The Base Registration Key field is added to the screen.
4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

   Important: *The registration key you use must support a Logging Node capable license.*

5. In the **Add-On Keys** field, paste any additional license key you have.
6. To add another additional add-on key, click the + sign and paste the additional key in the new **Add-On Keys** field.
7. For the **Activation Method** setting, select **Automatic**, and click the **Activate License** button.
   The End User Software License Agreement (EULA) displays.
8. To accept the license agreement, click the **Agree** button.
9. Click the **Next** button at the right of the screen.
   If the license you purchased supports both Logging Node and BIG-IQ Central Management Console, the License Feature Selection popup screen opens. Otherwise the Management Address screen opens.
10. If you are prompted with the License Feature Selection, select **BIG-IQ Logging Node**, and then click **OK**. If you are not prompted, proceed to the next step.

    Important: *This choice cannot be undone. Once you license a device as a Logging Node, you cannot change your mind and license it as a BIG-IQ Management Console.*

    The Management Address screen opens.
11. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.
    You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).
12. In the **Management Port IP Address** field, type the IP address for the management port IP address.

    Note: *The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

13. In the **Management Port Route** field that the system creates, type the IP address for the management port route.
14. Specify what you want the BIG-IQ to use for the **Discovery Address**.
    - To use the management port, select **Use Management Address**.
    - To use the internal self IP address, select **Self IP Address**, and type the IP address.

      Important: *If you are configuring a Logging Node device, you must use the internal self IP address.*

      Note: *The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

15. Click the **Next** button at the right of the screen.
16. In the **DNS Lookup Servers** field, type the IP address of your DNS server.
    You can click the **Test Connection** button to verify that the IP address is reachable.
17. In the **DNS Search Domains** field, type the name of your search domain.
    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
18. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.
    You can click the **Test Connection** button to verify that the IP address is reachable.
19. From the **Time Zone** list, select your local time zone.

20. Click the **Next** button at the right of the screen.
21. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.
22. Click the **Next** button at the right of the screen.

Manually license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`http://www.f5.com`).

If the BIG-IQ® system is not connected to the public internet, use this procedure to manually activate the license and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://<management_IP_address>`, where `<management_IP_address>` is the address you specified for device management.
2. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
3. Click **Activate**.
   The Base Registration Key field is added to the screen.
4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

   Important: *The registration key you use must support a Logging Node capable license.*

5. In the **Add-On Keys** field, paste any additional license key you have.
6. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button.
   The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
7. Select and copy the text displayed in the **Device Dossier** field.
8. Click the **Access F5 manual activation web portal** link.
   The Activate F5 Product site opens.
9. Into the **Enter your dossier** field, paste the dossier.

   Alternatively, if you saved the file, click the **Choose File** button and navigate to it.

   After a pause, the license key text displays.
10. Click the **Next** button.

    The Accept User Legal Agreement screen opens.
11. To accept the license agreement, select the **I have read and agree to the terms of this license**, and click **Next**. button.
    The licensing server creates the license key text.
12. Copy the license key.
13. In the **License Text** field on BIG-IQ, paste the license text.
14. Click the **Activate License** button.
15. Click the **Next** button at the right of the screen.
    If the license you purchased supports both Logging Node and BIG-IQ Central Management Console, the License Feature Selection popup screen opens. Otherwise the Management Address screen opens.
16. If you are prompted with the License Feature Selection, select **BIG-IQ Logging Node**, and then click **OK**. If you are not prompted, proceed to the next step.

    Important: *This choice cannot be undone. Once you license a device as a Logging Node, you cannot change your mind and license it as a BIG-IQ Management Console.*

    The Management Address screen opens.
17. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.

    You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).

18. In the **Management Port IP Address** field, type the IP address for the management port IP address.

    Note: *The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

19. In the **Management Port Route** field that the system creates, type the IP address for the management port route.
20. Specify what you want the BIG-IQ to use for the **Discovery Address**.
    - To use the management port, select **Use Management Address**.
    - To use the internal self IP address, select **Self IP Address**, and type the IP address.

      Important: *If you are configuring a Logging Node device, you must use the internal self IP address.*

      Note: *The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

21. Click the **Next** button to save your configuration.
22. In the **DNS Lookup Servers** field, type the IP address of your DNS server.

    You can click the **Test Connection** button to verify that the IP address is reachable.
23. In the **DNS Search Domains** field, type the name of your search domain.

    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
24. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.

    You can click the **Test Connection** button to verify that the IP address is reachable.
25. From the **Time Zone** list, select your local time zone.
26. Click the **Next** button at the right of the screen.
27. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.
28. Click the **Next** button at the right of the screen.

Configuring the BIG-IP logging profile

Each properly-configured BIG-IP® system sends its event log to a BIG-IQ® Logging Node. You configure the BIG-IP system to do so by creating a logging profile and assigning the logging profile to a virtual server, and then deploying it to the BIG-IP system. The *logging profile* defines the content of the events, and identifies the Logging Node to which the events are sent.

1. Log in to BIG-IQ system with your administrator user name and password.
2. Access the BIG-IQ component you are setting up, using the BIG-IQ menu and options near the top of the screen.
   - If you are setting up Web Application Security, select **Web Application Security** and then **Shared Security**.
   - If you are setting up Fraud Protection Service, select **Network Security** and then **Shared Security**.
3. On the left, expand **SECURITY PROFILES** and click **Logging Profiles**.
4. On the Logging Profiles screen, click **Create**.
   The Logging Profiles properties screen opens, showing the Properties tab.
5. On the Properties tab, edit as appropriate:
   a) In the **Name** field, type a unique name for this new profile. This field is required.

b) For **Application Security**, select **Enabled**. The Application Security tab appears.

When you select this option, the **Protocol Security** option is not available, but you can select any of the other options.

c) To use **Network Firewall**, select the **Enabled** check box; the Network Firewall tab appears.

d) To use **DoS Protection**, select the **Enabled** check box; the DoS Protection tab appears.

6. On the Application Security tab, select **Remote Storage**.
Several new fields appear, including the **Protocol** list.
7. Specify the appropriate **Logging Format**.
   - If the BIG-IP device runs version 12.0 or later, select **BIG-IQ**.
   - If the BIG-IP device runs a version prior to 12.0, select **Comma-Separated Values**. Several new fields appear.
     - For Storage Format, select **User Defined**.
     - In the **Selected Items** field, paste the following text:

```
unit_hostname="%unit_hostname%",management_ip_address="%management_ip_address%",
http_class_name="%http_class_name%",web_application_name="%http_class_name
%",policy_name="%policy_name%",
policy_apply_date="%policy_apply_date%",violations="%violations%",support_id="%support_id%",
request_status="%request_status%",response_code="%response_code%",ip_client="%ip_client%",
route_domain="%route_domain%",method="%method%",protocol="%protocol
%",query_string="%query_string%",
x_forwarded_for_header_value="%x_forwarded_for_header_value%",sig_ids="%sig_ids
%",sig_names="%sig_names%",
date_time="%date_time%",severity="%severity%",attack_type="%attack_type
%",geo_location="%geo_location%",
ip_address_intelligence="%ip_address_intelligence%",username="%username
%",session_id="%session_id%",
src_port="%src_port%",dest_port="%dest_port%",dest_ip="%dest_ip
%",sub_violations="%sub_violations%",
virus_name="%virus_name%",uri="%uri%",request="%request
%",violation_details="%violation_details%",
header="%headers%",response="%response%
```

> ***Note:*** *The line breaks in the example above were necessary due to screen width; remove all of them after you paste this data. It must be a single string with no white space.*

8. For **Protocol**, select **TCP**.
9. For the **Server Addresses** settings, specify the address you want to use:
   a) In the **IP Address** field, type the Logging Node's management IP address.
   b) Specify the port to use for your data.
      - If you are setting up a logging profile for Web Application Security, type `8514` in the **Port** field.
      - If you are setting up a logging profile for Fraud Protection Service, type `8008` in the **Port** field.
   c) Click the **Add** button to add the address and port to the list of servers.
10. To specify the **Storage Format** for this data, select the preferred formats in the **Available Items** list, and click the right arrow to add them to the **Selected Items** list.
11. For the **Maximum Entry Length**, select **64k**.
12. In the Storage Filter area, from the **Request Type** list, select **All requests**.
13. Click **Save** to save the new profile.

Virtual servers that remote logging uses to route event logs

You can either create a new virtual server on the BIG-IP® device that creates the event, or you can use a virtual server that already exists on that device.

Creating a virtual server for remote logging

If the device for which you are configuring remote logging does not have a virtual server, you need to create one.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays a list of virtual servers defined on this device.
5. Click **Create**.
   The Virtual Servers - New Item screen opens.
6. In the **Name** field, type in a name for the virtual server you are creating.
7. From the **Device** list, select the device on which to create the virtual server.
8. In the **Description** field, type in a brief description for the virtual server you are creating.
9. For the **Destination Address**, type the IP address of the destination you want to add to the Destination list.

   The format for an IPv4 address is `I<a>.I<b>.I<c>.I<d>`. For example, `172.16.254.1`.

   The format for an IPv6 address is `I<a>:I<b>:I<c>:I<d>:I<e>:I<f>:I<g>:I<h>.`.

   For example, `2001:db8:85a3:8d3:1319:8a2e:370:7348`.
10. In the **Service Port** field, type a service port number, or select a type from the list.

    When you select a type from the list, the value in the **Service Port** field changes to reflect the associated default, which you can change.
11. Click **Save**.
    The system creates the new virtual server with the settings you specified.
12. Click **Save** to save the assignment. Or, click **Save & Close** to save the assignment and return to the Virtual Servers screen.

A virtual server that can be used to routeevent data to the logging node is created for the BIG-IP® device.

Before the BIG-IP device can actually use this new virtual server, you must deploy it to the device.

Assigning the logging profile to a virtual server

After configuring a logging profile on the BIG-IQ® system, you must assign it to a virtual server and deploy it to the BIG-IP® system from which you want to collect event logs.

1. Log in to BIG-IQ system with your administrator user name and password.
2. Access the BIG-IQ component you are setting up, using the BIG-IQ menu and options near the top of the screen.
   - If you are setting up Web Application Security, select **Web Application Security** and then **Shared Security**.
   - If you are setting up Fraud Protection Service, select **Network Security** and then **Shared Security**.
3. If necessary, expand **ADC**, then select **Virtual Servers**.

The screen displays a list of virtual servers that are configured with devices that have been provisioned and discovered.

4. On the Virtual Servers screen, click the name of the virtual server you want to use.
   The Virtual Servers - Properties screen opens.
5. From the **Log Profiles** list, under **Available**, click a logging profile and move it to the **Selected** list.
6. Click **Save** to save the assignment. Or, click **Save & Close** to save the assignment and return to the Virtual Servers screen.

You are now ready to deploy your virtual server (with logging profile) to the BIG-IP system(s) from which you want to collect event logs.

Logging Node management

There are a number of useful concepts to consider when you manage Logging Nodes for off-box log storage. This reference material might prove helpful in setting up and maintaining your Logging Node configuration.

Logging Node sizing guide

Logging Nodes are specialized BIG-IQ® devices designed to provide sufficient CPU, memory, and disk capacity to store and search logging data from BIG-IP® devices. The underlying technology to provide these services is Elasticsearch. (Information about general Elasticsearch comments can be found on their website: `https://www.elastic.co/guide/en/elasticsearch/reference/current/_basic_concepts.html`)

Logging Nodes managed by BIG-IQ provide an Elasticsearch (ES) cluster that can scale horizontally (more nodes = more capacity), but each node in that cluster has limits on disk space. To mitigate that, there are a number of configuration elements that control how much disk is used by the system.

Logging Node Minimum Recommended Configuration

CPU	8 Cores
Memory	32 GB
Disk	10 GB (`/var` file system)

The `/var` file system on the Logging Node (which includes ES data) is only 10GB in size. To store more data on the file system, you need to extend the size. Refer to *Index rotation policy* for details on managing the data requirements. Extending the file system to 500GB is straightforward, assuming overall disk allocation on the BIG-IQ virtual machine is adequate. Log in as `root` to the Logging Node, and run the following commands.

1. `tmsh show sys disk directory`

   The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            10485760        -
/var/log                        7168000         -
```

2. `tmsh modify sys disk directory /var new-size 500000000`

   `tmsh show sys disk directory`

The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            10485760        500000000
/var/log                        7168000         -
```

3. Reboot the system and then confirm the size disk size.

```
tmsh show sys disk directory
```

The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            500003840       -
/var/log                        7168000         -
```

Logging Node Capacity

The following table is a very rough guide to how much data can be stored on a given Logging Node. The estimate assumes that the Logging Node has been configured to the recommended /var filesystem size. This size is outlined in the *Index rotation policy*. Because all indexes share the same filesystem, the approximate maximum documents per node estimate assumes no other indexes exist on that node.

Module	**Index name**	Average document size (bytes)	Approximate maximum documents per node
Access	access-event-logs	730	500GB / 730 = 700 million
Access	access-stats	730	500GB / 730 = 700 million
ASM	asmindex	1400	500GB / 1400 = 350 million
FPS	websafe	1400	10GB / 1400 = 70 million

Index rotation policy

The optimum settings used to configure your logging node indices depend on a number of factors. Some of the key factors are discussed here.

The system provides the ability to dynamically create new indices based either on a specified interval or on a specified size. The primary goal to consider when you make these decisions is how to maintain a maximum disk allocation for the Logging Node data while maintaining capacity for new data that flows in.

Secondary considerations include search optimization, and the ability to optimize old indices to reduce their size.

Generally, the best policy is one that does not create unnecessary indices. The more indices, the lower the overall performance, because your searches have to deal with more shards. For example, if a module knows that it has a low indexing volume (thousands/day) then it makes the most sense to have a large

aggregation per rotation (5 days or 30 days). For components like Web Application Security that probably have high indexing volumes, it makes more sense to rotate every 8 hours (which reduces the number of retained indices).

Index rotation also allows changing sharding and replica counts by changing the template on a given index type. New indices created from that template will contain the new shard and replica count properties.

This table shows the default configuration values for each index running on the BIG-IQ®. These values are based on anticipated data ingestion rates and typical usage patterns.

Component	Index Name	Minimum Number of Logging Nodes	Rotation Policy	Retained Index Count	Approximate time window	Size of /var file system
Access	access-event-logs	2	Time/5 days	19	95 days	500 GB
Access	access-stats	2	Time/5 days	19	95 days	500 GB
Web Application Security	asmindex	2	Size/100000 MB	5	N/A	500 GB
FPS	websafe	2	Time/30 days	100	8 years	10 GB

If multiple modules are running on a given Logging Node or if you have higher inbound data rates, you might have to adjust these values to keep the `/var` file system from filling up (there is a default alert to warn of this when the file system becomes 80% full).

The simplest resolution is to revise the retained index count; lowering this value will reduce the disk space requirements but it will also reduce the amount of data available for queries. For details on changing this setting, refer to *Modifying event indices*.

How do I use the event log interface?

The event log interface consists of two filter fields and three main screens:

- Filter fields:
  - Selected devices filter. This filter appears below the Event Logs header. You can use it to select one or more devices for event viewing.
  - Filter field. Appears to the right of the selected devices field. You can use it to type text to rapidly narrow the search scope. You can also save filters that you use often.
- Screens:
  - Devices. At the far left, use this to select a group of requests, policies, saved filters, or pre-configured tags. The object you select determines the set of items that appears in the next screen.
  - Log items. Use this to browse log items, or select one and view log item details.
  - Details. Displays details of the item selected in the Log items screen.

Viewing event log details

You can view request and response details for a single log item.

1. Log in to the BIG-IQ system with your administrator user name and password.
2. From the BIG-IQ list, select **WebApplication Security**.

3. Click **Event Logs**.
4. On the Log Items screen (list of events), click a single event log.
   The Details screen displays a variety of information about the event.
5. On the Details screen, click **Request** to view request details.

   Details include:

   - Raw HTTP[S] request
   - General request details
   - Geolocation
   - Policy details
   - List of related tags
6. Click **Response** to view response details.

Using common filters

You can update common filters for requests and security policies.

1. Log in to the BIG-IQ system with your administrator user name and password.
2. From the BIG-IQ list, select **Web Application Security**.
3. Click **Event Logs**.
4. To update log items according to a selected filter (such as **Requests** or **Policies**), click any item under **Requests** or **Policies**.

The system updates log items according to the selected filter, and results appear in the Log Items screen.

Filtering the event logs (basic)

You can use the filter to refine your searches through the event logs, including searches through logs from multiple BIG-IP® devices.

1. Log in to BIG-IQ system with your administrator user name and password.
2. From the BIG-IQ list, select **Web Application Security**.
3. Click **Event Logs**.
4. In the Filter field, click the triangle to the right of the field.
   The Search filter popup screen opens to the basic view, which is the default.
5. Complete the fields.

Setting	Description
Request type	Type a request type or select from the list **All requests** or **Illegal requests** (log responses for illegal requests only).
Support ID	Type the complete support ID (unique ID given for a transaction), or select the **Last 4 digits** check box and type the last 4 digits of the support ID.
Violation	Use this list to select the policy violation that detects attacks, such as Attack Signature Detection or Illegal Cookie Length. You can select a violation type from the list or you can select none of the violations (indicating that any violation type matches).
Attack type	Use this list to select the type of service attacks (such as Denial of Service or HTTP Parser Attack) that you want to see. Select nothing (indicating that any attack type matches), or select a specific attack type.
Time Period	In the **From** and **To** fields, type a date and time in the format: `2015-12-01T15:15:29-05:00`. Or click the calendar icon and select dates.

Setting	Description
Policies	In the field, type a policy name, or click in the field and, from the list, select a policy.

6. Click the Search bar.
   The results of the filtering process appear in the Log Items list.
7. When you have configured a search that you will use repeatedly or frequently, click **Save the current filter**, type a filter name, and click **Save**.
   The saved filter appears in the left panel under **Saved filter**.

Filtering (advanced)

You can use the filter's advanced setting to refine your searches.

You can type a query in the filter box in the format `method:'value' protocol:'value' severity:'value'`. For example: `method:'GET' protocol:'HTTPS' severity:'error'`.

Or, you can open the filter and use the method described in the following section.

1. Log in to BIG-IQ system with your administrator user name and password.
2. From the BIG-IQ list, select **Web Application Security**.
3. Click **Event Logs**.
4. Open the Filter field.
   The Search filter popup screen opens to the basic view, which is the default.
5. Click **Advanced**.
6. Complete the fields.

Setting	Description
Method	From the list, select a method.
Protocols	From the list, select **HTTP** or **HTTPS**, depending on the security requirements.
Severity	From the list, select **Informational**, **Critical**, or **Error**.

7. Click the search bar.
   The results of the filtering process appear in the Log Items list.
8. When you have configured a search that you will use repeatedly or frequently, click **Save the current filter**, type a filter name, and click **Save**.
   The saved filter appears in the left panel under **Saved filter**.

Filtering by entering query parameters

You can use the BIG-IQ® Filter field to enter a query in ODATA format:

```
key1:'value' key2:'value' (key3:'value' OR key4:'value').
```

For example:

```
policy_name:'/Common/policy1'
```

Note: *The BIG-IQ system supports AND/OR constructs.*

- OR. Use this operator to log the data that meets one or more of the criteria.
- AND. Use this operator to log the data that meets all of the criteria.

Keys, values, and operators are listed and described in the following text.

1. Log in to the BIG-IQ system with your administrator user name and password.
2. From the BIG-IQ list, select **Web Application Security**.

3. Click **Event Logs**.
4. In the Filter field, type a query in ODATA format.
5. Type a key from the following list:

Key	Description
attack_type	Name of identified attack (string). For example: `Non-browser client`.
date_time	Current date and time. For example: `2016-09-19 13:52:29`
dest_ip	Requested service IP address, generally, the virtual server IP address. For example: `192.168.5.11`.
dest_port	Destination port of this transaction (non-negative integer). For example: `80`.
geo_location	Country/city location information, based on the source IP address. For example: `USA/NY`.
headers	List of request headers found in request logs. For example: `Host: myhost.com; Connection: close`.
http_class_name	Alias of policy name. For example: `/Common/topaz4-web4`.
ip_address_intelligence	List of IP intelligence categories found for an IP category such as proxy, phishing and so on. For example: `Scanners`.
ip_client	Client source (attacker) IP address. For example: `192.168.5.10`.
management_ip_address	BIG-IP® management IP address.
method	HTTP method requested by the client. For example: `GET`.
policy_apply_date	Last apply policy operation date and time.
policy_name	Name of the active security policy. For example: `ACME security policy`.
protocol	Transport protocol (string). For example: `HTTP`.
query_string	URI query string. For example: `/`.
request	Request string sent by the client. For example: `GET / HTTP/1.0\r\nUser-Agent: Wget/1.12 (linux-gnu)\r\nAccept: */*\r\nHost: 10.4.1.200\r\nConnection: Keep-Alive\r\n\r\n`.
request_status	Action applied to the client request. For example: `Blocked`.
response_code	The HTTP response code returned by the back-end server (application). This information is relevant only for requests that are not blocked. For example: `200`.
route_domain	Route domain number (non-negative integer). For example: `0`.
session_id	ID number (hexadeicmal number) assigned to the request to allow the system administrator to track requests by session. For example: `a9141b68ac7b4958`.
severity	Severity category to which the event belongs. For example: `Error`.
sig_ids	Signature ID number (positive non-zero integer). For example: `200021069`.

Key	Description
sig_names	Signature name(s). For example: `Automated client access %22wget%22`.
src_port	Client protocol source port of this transaction (non-negative integer). For example: `52974`.
sub_violations	Comma-separated list of sub-violation strings. for example: `Bad HTTP version, Null in request`.
support_id	Internally-generated integer to assist with client access support. For example: `18205860747014045721`.
unit_hostname	BIG-IP system FQDN (unit host name).
uri	URI requested by the client (string). For example: `/`.
username	User name for the client session. For example: `admin`.
violations	Comma-separated list of the violations that occurred during enforcement of the request or response. For example: `Attack signature detected`.
virus_name	Virus name (string). For example: `Melissa`.
x_forwarded_for_header_value	Value of the XFF HTTP header (string). For example: `192.168.5.10`

6. Type an operator from the following list:

Operator	Description
eq	Equal
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to

7. Type a value in any of the following formats:
   - 'value'. For example: `policy_name:'/Common/policy1'`
   - '*alue'. For example: `policy_name:'*Common/policy1'`
   - 'alu*'. For example: `policy_name:'Common/policy*'`
   - '*ue*'. For example: `policy_name:'policy*'`
8. Press **Enter** or click the search icon to start the search.

The system updates log items according to the typed query, and results appear in the Log Items list.

Restore event log snapshots

To submit the REST API calls required by this task, you must provide the administrator user name and password.

The BIG-IQ® user interface does not currently support restoring the event snapshots. However, if a logging node fails, you can manually restore the data up to the last snapshot.

Please note the following:

- The restore operation requires a down time during which no BIG-IQ or Logging Node work is performed.
- During the restore operation, no event data sent to the Logging Node is retained.
- The restore operation restores only the data from the time before the chosen snapshot was created. Data from the time that the chosen snapshot was created to the current time is not restored.
- Before initiating a snapshot restore, make sure that sufficient disk space is allocated to the /var folder on the device to which you are restoring the snapshot.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING** and select **Logging Configuration**.
   The Logging Configuration screen opens.
4. Click the **View History** button.
   The BIG-IQ Logging Snapshots screen opens.
5. Browse through the list to get an idea of which snapshot you want to restore.
6. On the Logging Configuration screen, next to **Last Snapshot/Time**, click **Restore**.

Managing Configuration Snapshots

What is snapshot management?

You can manage configuration snapshots for the configurations you have created on the BIG-IQ® Centralized Management system. A *snapshot* is a backup copy of a configuration. Configuration snapshots are created manually. This type of snapshot does not include events.

Creating a snapshot

You create a configuration snapshot to compare it to another configuration snapshot, or so you can save the current working configuration and then restore from that snapshot if needed.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service configuration from which to create the snapshot: **Network Security** or **Web Application Security**.
   The screen displays a list of snapshots that have been created for that service on this device.
4. At the top of the screen, click **Create**.
   The Create Snapshot screen opens.
5. Supply the values on the Create Snapshot screen, and click **Create**.

The system creates the snapshot and adds it to the list of snapshots on the Snapshot and Restore - screen, including information related to the snapshot, including the date it was created, what account created it, and any description.

Comparing snapshots

You can compare two snapshots to view their differences.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service containing the snapshots to compare: **Network Security** or **Web Application Security**.

The screen displays a list of snapshots that have been created on this device.

4. Select the check box to the left of each of the two snapshots to be compared.
5. Click **Compare**.
   The Compare Snapshots screen displays.
6. For the **Target**, select the snapshot to which you want to compare the snapshot listed as the Source.
7. Compare the snapshots selected:
   - To compare firewall object differences, click **Compare** in the Compare Firewall row. This option is only available with the AFM™ service.
   - To compare ASM differences, click **Compare** in the Compare ASM row. This option is only available with the ASM™ service.
   - To compare shared security object differences, click **Compare** in the Compare Shared Security row.
8. Analyze the configuration differences between the two snapshots, When you are finished, click **Cancel** to close the Differences screen, then click **Close**.
   The screen closes and you return to the Snapshot and Restore - screen.

Restoring a snapshot

You can restore a snapshot to change the working configuration to that of the snapshot. Restoring the snapshot merges objects from the snapshot into the BIG-IQ® configuration and removes all active locks. No objects in the BIG-IQ configuration are removed. Once the restore process starts, you cannot modify the BIG-IQ configuration until the process is completed or canceled. If the process is canceled, all configuration settings are rolled back.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service containing the snapshot to restore: **Network Security** or **Web Application Security**.
   The screen displays a list of snapshots that have been created on this device.
4. Select the check box to the left of the snapshot to use to restore the current working configuration to the configuration of the snapshot.
5. Click **Restore**.
   The Restore snapshot to Working Configuration screen opens.
6. Compare the snapshot to restore to the working configuration:
   - To compare firewall object differences, click **Compare** in the Compare Firewall row. This option is only available with the AFM™ service.
   - To compare ASM™ differences, click **Compare** in the Compare ASM row. This option is only available with the ASM service.
   - To compare shared security object differences, click **Compare** in the Compare Shared Security row.

   The differences screen for the comparison is displayed for you to review. Click **Cancel** when done.
7. Click **Restore** to restore the configuration in the snapshot and have it replace the working configuration.
8. Click **Restore** in the popup screen to confirm that you want to restore the configuration, or click **Cancel** in the popup screen to stop the restore process for this the snapshot.

   You can also click **Cancel** after starting the restore process to roll back the restore.

Deploying Changes

How do I evaluate changes made to managed objects?

To change the object settings on a managed device, there are four tasks to perform.

This figure illustrates the workflow you perform to manage the objects on BIG-IP® devices. Evaluating the changes you have made is the third step in this process.

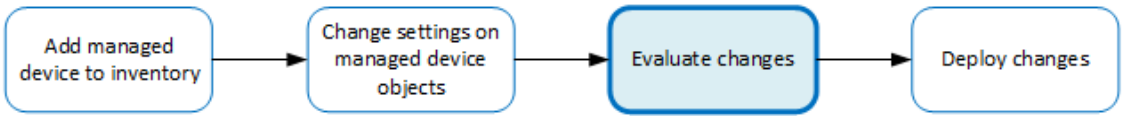

Figure 2: Overview of evaluating changes made to managed objects

Note: *If you need to make an urgent change, you can skip the evaluation step. However, we highly recommend evaluation in all but emergency situations. See Making an urgent deployment for details.*

Evaluating configuration changes

Evaluating your changes gives you a chance to spot critical errors and review your revisions one more time before deploying them.

Note: *Critical errors are issues with a configuration change that cannot be deployed successfully. Verification warnings are less serious in that they may not cause the deployment to fail, but should be reviewed nonetheless.*

Note: *If you have Local Traffic & Network (LTM) changes to deploy, deploy the LTM changes before deploying changes to other components, or those deployments may fail.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under EVALUATE & DEPLOY, select the type of evaluation and deployments to view. Select **Network Security** or **Web Application Security**, whichever is appropriate.
   The list of evaluations and deployments that have been created on this device opens.
5. Under Evaluations, click **Create**.
   The Create Evaluation screen opens.
6. In the **Name** field, type in a name for the evaluation task you are creating.
7. In the **Description** field, type in a brief description for the evaluation task you are creating.
8. For the **Source**, select what you want to evaluate.
   - To compare the object settings currently on the managed device with the object settings in the pending version, select **Current Changes**.
   - To compare the object settings currently on the managed device with the object settings in a stored snapshot, select **Existing Snapshot**, then choose the snapshot you want to use.
9. For the **Target** setting, identify the devices for which you want to evaluate changes.
   a) If the devices are in a device group, select **Group**, and select the group from the list.

b) If the devices are not in a device group, select **Device**.
c) Select the devices from the **Available** list, and use the arrow button to move the devices to the **Selected** list.

Important: *If you deploy changes to a device that is in a DSC® cluster, you must include both devices before you can create the evaluation.*

Important: *If the device in the **Selected** list has a filled circle in front of it, a deployment is needed for the BIG-IP device configuration to match the BIG-IQ working configuration for that BIG-IP device. This notification occurs only when creating Web Application Security evaluations.*

10. Click the **Create** button at the bottom of the screen.
The system adds the new evaluation to the list, and analyzes the changes for errors. When the configuration evaluation completes, you will see how many changes or errors the evaluation found.

11. Review the evaluation to determine whether you are going to deploy it or not.
a) If there are critical errors, you cannot deploy these changes. Click each error to see what it is, and then go back to where you made the change to fix it.
After resolving any critical errors, you can come back and repeat the evaluation.
b) If there are verification warnings, you can still deploy your changes, but you will probably want to resolve the warnings first. Click each warning to see what it is, and then go back to where you made the change to fix it.
After resolving any verification warnings, you can come back and repeat the evaluation.
c) If there are no critical errors or verification warnings, review the changes by clicking the **Difference** link.
Each change is listed. You can review each one by clicking the name.

To apply the object changes to the managed device, you must deploy them.

How do I deploy changes made to managed objects?

Deploying changes applies the revisions that you have made on the BIG-IQ® to the managed BIG-IP® devices.

Note: *Before the BIG-IQ deploys configuration changes, it first reimports the configuration from the managed device to ensure there are no unexpected differences. If there are issues, the default behavior is to discard any changes made on the managed device and then deploy the configuration changes.*

- To accept the default, proceed with the deployment. The settings from the managing BIG-IQ overwrite the settings on the managed BIG-IP device.
- To override the default, rediscover the device and reimport the service. Any changes that have been made using the BIG-IQ are overwritten with the settings from the managed BIG-IP device.

This figure illustrates the workflow you perform to manage the objects on BIG-IP devices. Deploying the settings is the last step in this process.

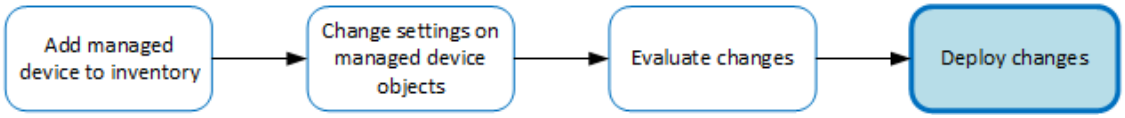

Figure 3: Change managed object workflow

Deploying configuration changes

To apply the changes you made on the BIG-IQ to your managed device, you must deploy those changes to the managed device.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under EVALUATE & DEPLOY, select **Network Security** or **Web Application Security**, whichever is appropriate.
   The list of evaluations and deployments defined on this device opens.
5. Click the name of the evaluation that you want to deploy.
   The View Evaluation screen opens.
6. Specify whether you want to deploy the changes immediately or schedule deployment for later.
   - To deploy this change immediately:
     1. Select **Deploy Now**.
     2. Click **Deploy** to confirm.
   - To deploy this change later:
     1. Select **Schedule for later**.
     2. Select the date and time.
     3. Click **Schedule Deployment**.
     4. Click **Schedule Deployment** again to confirm.

   The process of deploying changes can take some time, especially if there are a large number of changes. During this time, you can click **Cancel** to stop the deployment process.

Important: *If you cancel a deployment, some of the changes may have already deployed.* ***Cancel*** *does not roll back these changes.*

The evaluation you chose is added to the list of deployments on the bottom half of the screen.

- If you chose to deploy immediately, the changes begin to deploy and the Status column updates as it proceeds.
- If you choose to delay deployment, the Status column displays the scheduled date and time.

Making an urgent deployment

You can skip the evaluation task and deploy changes right now if you need to make urgent changes.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under EVALUATE & DEPLOY, select **Network Security** or **Web Application Security**, whichever is appropriate.
   The list of evaluations and deployments defined on this device opens.
5. Under Deployments, click **Create**.
   The Create Deployment screen opens.
6. In the **Name** field, type in a name for the deployment task you are creating.
7. In the **Description** field, type in a brief description for the deployment task you are creating.
8. For the **Source** setting, select what you want to deploy.

- To deploy your changes to the managed device, select **Current Changes**.
- To deploy the object settings from a stored snapshot, select **Existing Snapshot**, then choose the snapshot you want to use.

9. Using the **Target** settings, identify the devices for which you want to deploy changes.
   a) If the devices are in a device group, select **Group**, and select the group.
   b) If the devices are not in a device group, select **Device**.
   c) Select the devices from the **Available** list and use the arrow button to move the devices to the **Enabled** list.
10. Consider one more time how you want to deploy these changes.
    - To make the changes right now, click **Deploy immediately**.
    - If you want to review the changes, click **Create evaluation**.
11. Click **Create**.
    - If you selected **Deploy immediately**, the changes begin to deploy and the Status column updates as it proceeds.
    - If you selected **Create evaluation**, the new evaluation is added to the list and the changes are analyzed for errors. When the evaluation completes, you will see how many changes or errors the evaluation found.

Verifying firewall rules have compiled on all BIG-IP devices

Once a firewall deployment has completed successfully, **Check Rule Compilation** is enabled on the View Deployment screen.

Use **Check Rule Compilation** to verify that your firewall rules are active on the BIG-IP devices to which you deployed those rules.

1. On the Deployments screen, click the name of the deployment that contains the firewall rules you want to verify.

   The View Deployment screen for that deployment displays.
2. On the View Deployment screen, click **Check Rule Compilation** to determine if rules have been compiled on all the BIG-IP devices in the firewall deployment.
   The rule compilation status and last activation time for each BIG-IP device included in the deployment are listed in a popup.
3. Verify that the last activation time for each BIG-IP device is after the end time of the BIG-IQ deployment task to ensure that firewall rules have been compiled on each BIG-IP devices. You can repeat this step multiple times.

   Review the following considerations when using **Check Rule Compilation**:

   - Be aware of any time differences, due to time zones and so on, between the BIG-IQ system and the BIG-IP device.
   - BIG-IP device versions earlier than 11.5.1 HF4 do not support the compilation statistics used by this feature and will display the message, `Compilation stats not provided for this version of BIG-IP`.
   - If the Check Rule Compilation feature is used with an older deployment, where the state of the BIG-IP device has changed since the deployment, the status returned will include all active firewall rule changes on the BIG-IP device since the deployment.
   - If the Check Rule Compilation feature returns the message `Local Last Activation Time` or the message `No stats found on device`, then the state of the BIG-IP device has changed since the deployment, and compilation statistics have been reset. This can be caused by a reboot of the BIG-IP device.

Reviewing deployment process states to diagnose problems

When a firewall security policy or a web application security policy is deployed, that policy goes through several deployment states. Reviewing these states may be useful in understanding what occurred during deployment in order to diagnose a problem. Note that not all states may appear in the log, since what states are displayed depends on how the deployment was processed.

Review the `restjavad.n.log` file to view deployment states for either a firewall security policy or a web application security policy.

Device deployment states

This table displays states that can occur during the deployment process, and a brief description of each state.

Table 3: Deployment States

State	Description
CHECK_LICENSE	Licenses for BIG-IQ systems are checked to be valid.
CHECK_OTHER_RUNNING_TASKS	Verifies that no tasks are running that could cause errors during deployment. Tasks that could cause errors include: • Other BIG-IQ Security deployment tasks running at the same time as this deployment, even if they are from different modules. • Tasks to declare management authority over a BIG-IP device. • Tasks that rescind management authority of a BIG-IP device.
GET_DEVICES	Finds all devices managed by the BIG-IQ Security system.
CHECK_DEVICE_AVAILABILITY	Determines whether the devices to be deployed are available.
LOOKUP_CLUSTERS	Determines if any devices included in the deployment are part of a cluster, and if so, verifies that both devices in the cluster are configured with the same sync mode and sync failover group on the BIG-IP device.
REFRESH_CURRENT_CONFIG_SOAP	Using the SOAP API, refreshes the current configuration for all devices included in the deployment. This process adds any new configuration items from the BIG-IP device to the current configuration.
REFRESH_CURRENT_CONFIG_REST	Using the REST API, refreshes the current configuration for all devices included in deployment. This process adds any new configuration items from the BIG-IP device to the current configuration.
CREATE_SNAPSHOT	Creates a snapshot of the working configuration.
CREATE_DIFFERENCE	Generates the differences between the snapshot taken and the current configuration.
VERIFY_CONFIG	Verifies that devices to be deployed do not have configuration problems that could lead to deployment errors.

State	Description
GET_CHILD_DEPLOY_DEVICES	Finds all devices managed by Shared Security objects. These devices are considered to be child deployments of a parent firewall security or web application security deployment.
START_CHILD_DEPLOY	Starts the deployment of devices managed by Shared Security objects.
WAIT_FOR_CHILD_DEPLOY	Waits for deployment of devices managed by Shared Security objects to complete.
CLEANUP_PREVIOUS_EVALUATE	Cleans up processing artifacts from the previous evaluation.
DISTRIBUTE_DSC_CLUSTERS	Distributes changes to devices identified as being in a cluster by the LOOKUP_CLUSTERS process and that are configured to use the BIG-IP Device Service Clustering (DSC) to keep the BIG-IP devices synchronized.
DISTRIBUTE_CONFIG	Distributes configuration changes to the specified devices.
DISTRIBUTE_CONFIG_SOAP	Using the SOAP API, distributes configuration changes to the specified devices.
DISTRIBUTE_CONFIG_REST	Using the REST API, distributes configuration changes to the specified devices.
FOLDBACK_DEPLOYED_ADDITIONS	Inserts any newly-added objects directly into the current configuration to that the BIG-IQ system will already know about those objects on the next refresh of the current configuration.
DONE	Indicates the deployment process has completed.

Managing Configuration Snapshots

What is snapshot management?

You can manage configuration snapshots for the configurations you have created on the BIG-IQ® Centralized Management system. A *snapshot* is a backup copy of a configuration. Configuration snapshots are created manually. This type of snapshot does not include events.

Creating a snapshot

You create a configuration snapshot to compare it to another configuration snapshot, or so you can save the current working configuration and then restore from that snapshot if needed.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service configuration from which to create the snapshot: **Network Security** or **Web Application Security**.
   The screen displays a list of snapshots that have been created for that service on this device.
4. At the top of the screen, click **Create**.
   The Create Snapshot screen opens.
5. Supply the values on the Create Snapshot screen, and click **Create**.

The system creates the snapshot and adds it to the list of snapshots on the Snapshot and Restore - screen, including information related to the snapshot, including the date it was created, what account created it, and any description.

Comparing snapshots

You can compare two snapshots to view their differences.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service containing the snapshots to compare: **Network Security** or **Web Application Security**.
   The screen displays a list of snapshots that have been created on this device.
4. Select the check box to the left of each of the two snapshots to be compared.
5. Click **Compare**.
   The Compare Snapshots screen displays.
6. For the **Target**, select the snapshot to which you want to compare the snapshot listed as the Source.
7. Compare the snapshots selected:
   - To compare firewall object differences, click **Compare** in the Compare Firewall row. This option is only available with the AFM™ service.
   - To compare ASM differences, click **Compare** in the Compare ASM row. This option is only available with the ASM™ service.
   - To compare shared security object differences, click **Compare** in the Compare Shared Security row.

8. Analyze the configuration differences between the two snapshots, When you are finished, click **Cancel** to close the Differences screen, then click **Close**.
   The screen closes and you return to the Snapshot and Restore - screen.

Restoring a snapshot

You can restore a snapshot to change the working configuration to that of the snapshot. Restoring the snapshot merges objects from the snapshot into the BIG-IQ® configuration and removes all active locks. No objects in the BIG-IQ configuration are removed. Once the restore process starts, you cannot modify the BIG-IQ configuration until the process is completed or canceled. If the process is canceled, all configuration settings are rolled back.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select the service containing the snapshot to restore: **Network Security** or **Web Application Security**.
   The screen displays a list of snapshots that have been created on this device.
4. Select the check box to the left of the snapshot to use to restore the current working configuration to the configuration of the snapshot.
5. Click **Restore**.
   The Restore snapshot to Working Configuration screen opens.
6. Compare the snapshot to restore to the working configuration:
   - To compare firewall object differences, click **Compare** in the Compare Firewall row. This option is only available with the AFM™ service.
   - To compare ASM™ differences, click **Compare** in the Compare ASM row. This option is only available with the ASM service.
   - To compare shared security object differences, click **Compare** in the Compare Shared Security row.

   The differences screen for the comparison is displayed for you to review. Click **Cancel** when done.
7. Click **Restore** to restore the configuration in the snapshot and have it replace the working configuration.
8. Click **Restore** in the popup screen to confirm that you want to restore the configuration, or click **Cancel** in the popup screen to stop the restore process for this the snapshot.

   You can also click **Cancel** after starting the restore process to roll back the restore.

Managing Audit Logs

About audit logs

You use audit logs to review changes in the BIG-IQ® system. All BIG-IQ system roles have read-only access to the audit log, and can view and filter entries. Any user with the appropriate privileges can initiate an action.

You view audit logs by selecting **Audit Logging** from the BIG-IQ menu, and then selecting the appropriate service option from the list on the left.

To view or change the audit log archive settings, click the **Archive Settings** button on the Audit Logging screen. Archived audit log files are stored in the `archive-audit.n.txt` file in the appropriate subdirectory of the `/var/config/rest/auditArchive` directory on the BIG-IQ system:

All API traffic on the BIG-IQ system, and every REST service command for all licensed modules, is logged in a separate, central audit log (`restjavad-audit.n.log`) which is located in `/var/log` on the BIG-IQ system.

Considerations when using the audit log

When using the audit log, consider the following:

- The audit log does not record an entry for every generation of a task. It only records an entry when the task status changes.
- When an object is deleted and then recreated with the same name, partition, and other information, the difference between those objects may show the deleted object as being the previous generation of the new object.
- By default, not all columns are displayed by the audit log to conserve space. To review what columns are displayed, click the gear icon in the upper right of the Audit Logging screen.

Actions and objects that generate audit log entries

BIG-IQ® records in the audit log all user-initiated changes that occur on the BIG-IQ system. A change is defined as when certain objects are modified, when certain tasks change state, or when certain user actions are performed. For example, when the admin account is used to log in to the BIG-IQ system, the audit log records the time, the user (admin), the action (New) and the object type (Login). The log does not include changes that occurred on BIG-IP® devices that were imported.

Changes to working-configuration objects generate audit log entries. In addition, these actions generate log entries:

- Creating or deleting a user account.
- Users logging in and logging out, including when the user is logged out due to inactivity.
- Creating or cancelling a device discovery or a device reimport.
- Creating a Network Security Change Verifications action to verify the changes to a specific BIG-IP device or group.
- Deleting a previously discovered device.
- Creating or deleting a deployment task.
- Creating a difference task.
- Creating, restoring, or deleting a snapshot.

- Editing some system information (such as editing a host name, a root password, a DNS entry, or an SNMP entry).

Audit log entry properties

The audit log displays the following properties for each log entry.

Property	Description
Source	IP address of the client machine that made the change. This property is blank for actions that were initiated by an internal process. For example, when a user invokes a deployment action, the deployment action then invokes a difference task to find the differences between the current configuration and the one to be deployed. The difference task has no Source IP address.
Service	Indicates whether the change was made by the internal object synchronization service. This service synchronizes shared objects, such as virtual servers, from the Local Traffic & Network service to the Network Security or Web Application Security services. • If a check mark is displayed, the change was made by the internal object synchronization service, and no IP address is shown in the Source column. The check mark is only displayed in the Network Firewall Audit Log or the Web Application Security Audit Log screens. • If a check mark is not displayed, the change was not made by the internal object synchronization service.
Time	Time that the event occurred. The time is the BIG-IQ system local time and is expressed in the format: mmm dd, yyyy hh:mm:ss (time zone); for example: `Apr 19, 2016 13:09:03(EDT)`.
Node	Fully qualified domain name for the BIG-IQ system that recorded the event. This appears as the **Hostname** at the top of the BIG-IQ user interface.
User	Name of the account that initiated the action, such as an account named `Admin` for an administrative account.
Action	Type of modification. For operation changes, the action types include New, Delete, and Modify. For task changes, the action types include Start, Finish, Failed, and Cancelled.
Object Name	Object identified by a user-friendly name; for example: `newRule1`, `deploy-test`, or `Common/global`. When the name `RootNode` is listed, that indicates that the object is associated with a BIG-IP device. `RootNode` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Changes	Indicates whether there was a change in the object. If **View** occurs in this column, there is a change to the object. To view the detailed differences of the change, click **View**.
Object Type	Classification for this action. When the type `Root Node` is listed, that indicates that the object is associated with a BIG-IP device. `Root Node` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Parent	The administrative partition and name of the parent object. This property is displayed for firewall rules, logging profiles, and DoS profiles. For firewall rules, the parent shows the rule list, firewall, or policy that contains the rule. A change in a firewall rule often also affects the rule's parent object.

Property	Description
Parent Type	Class or group of the parent object.
Version	Version of the configuration object. Typically, when a configuration object changes, the version is increased by 1. However, other audit entries, such as those for finishing snapshot creation or finishing deployment, may increase the version by more than 1.

Viewing audit entry differences

In the audit log, when potential changes to an object are logged, the **View** link is shown in the Changes column for that entry. You can click **View** to examine the differences between generations of that object.

1. Log in to BIG-IQ® system with Administrator or Security Manager credentials.
2. Select **Audit Logging** from the BIG-IQ menu, and then from the list on the left, click the option from which to view audit entries.
3. To display differences for an object, click **View** in the Changes column.

   A popup screen opens, showing two columns that compare the differences between the two generations of the object in JSON. In these columns, additions to an object generation are highlighted in green, and differences are highlighted in gold.

   If the system cannot retrieve a generation of an object, the column displays either `Generation Not Available` or `Generation No previous generation`. Object information may not be available if it has been automatically purged from the system to conserve disk space, or if it has been deleted.

   The JSON difference displayed for a delete entry in the audit log shows the JSON difference from the previous operation because the generation identifier is not incremented when an object is deleted.
4. When you are finished, click **Close** on the popup screen to return to the Audit Logging screen.

Filtering entries in the audit log

You can use the Filter field at the top right of the Audit Logging screen to rapidly narrow the scope displayed, and to more easily locate an entry in the audit log.

- Filtering is text-based.
- Filtering is not case-sensitive.
- You can use wild cards, or partial text.
- All BIG-IQ® roles can filter entries.
- To clear the filter, click the **X** to the right of the search string in the **Filtered by** field on the left.

1. Log in to BIG-IQ system with Administrator or Security Manager credentials.
2. Select **Audit Logging** from the BIG-IQ menu, and then, on the left, click the area from which to view audit entries.
3. Use the Filter field to narrow your search:
   a) Use the arrow key to the left of the field to select the appropriate filter options.
   b) Type the information specific to the object you want to filter on.
   c) Press Enter.

Option	Description
All	Specifies that all objects should be filtered using the filter text. When this option is used, both the user-visible and the underlying data are searched for a match, so you may see matches to your filter text which do not appear to match it.
Source	Type the source IP address in the filter. When this option is used, both the user-visible and the underlying data are searched for a match, so you may see matches to your filter text which do not appear to match it. Using this option is equivalent to using the **All** option.
Time	Type both a date and a time. Displayed times are given in the local time of the BIG-IQ system. Supported time formats are highly Web browser-dependent. Time formats other than those listed might appear to filter successfully but are not supported. Entering a single date and time results in a filter displaying all entries from the specified date and time to the current date and time. For time formats that use letters and numbers, enter the date time in one of the following formats: • mmm dd yyyy hh:mm:ss. Example: `Jan 7 2014 8:30:00` • mmm dd, yyyy hh:mm:ss (time zone). Example: `Apr 28, 2016 13:09:03(EDT)` • mmm dd, yyyy. Example: `Apr 28, 2016` • mmm dd, yyyy hh:mm:ss. Example: `Apr 28, 2016 16:09:06` • ddd mmm dd yyyy hh:mm:ss. Example: `Thu Jan 16 2014 11:13:50` For time formats that use only numbers, enter the date time in one of the following formats: • mm/dd/yy hh:mm:ss. Example: `01/01/16 12:14:15` • m/d/yy hh:mm:ss. Example: `1/1/14 12:14:15` • mm/dd/yyyy hh:mm:ss. Example: `1/1/2014 12:14:15`
Node	Type the node name in the filter.
User	Type the user account name in the filter.
Action: Operation	Type the operation action name in the filter. Operation actions include: New, Delete, and Modify.
Action: Task Status	Type the task status action name in the filter. Task status actions include: Start, Finish, Cancelled, and Failed.
Object Name	Type the full or partial name of the object in the filter. If a partition name is displayed, do not include it in the filter. For example, Common/AddressList_4 would be entered as `AddressList_4`. Because the device-specific object name includes the BIG-IP® host name, you can enter a full or partial device name to get all objects for a specific BIG-IP device.
Object Type	Type the object type in the filter.
Parent	Type the parent name in the filter. Only appears for rules to show the rule list, firewall, or policy that contains the rule.
Parent Type	Type the Parent Type name in the filter. Only appears when the Parent field contains a value.
Contains	Specifies that the filter text is contained within the object specified. When you select **Contains**:

Option	Description
	• If the filter text is a string, the filter text matches an entire string or only a part of a string. • If the filter text is an IP address, the filter text matches an IPV4 or IPV6 address that is the same as the filter text, or matches an IPV4 address range or subnet that includes the filter text. IPV6 addresses can not be found within a range or subnet. • If the filter text is a port number, the filter text matches a port number that is the same as the filter text, or matches a port number range that includes the filter text.
Exact	Specifies that the filter text is exactly contained within the object specified. When **Exact** is selected: • If the filter text is a string, the filter text matches only the entire string. • If the filter text is an IP address, the filter text matches only an IPV4 or IPV6 address that is the same as the filter text. • If the filter text is a port number, the filter text matches only a port number that is the same as the filter text.

The result of a search filter operation is a set of entries that match the filter criteria, sorted by time.

Customizing the audit log display

You can customize the audit log display to assist you in locating information faster.

- To customize the order of columns displayed, click any column header and drag the column to the location you want.
- To sort by column, click the name of the column you want to sort. Not all columns can be sorted. When sorting items in the Object Name column, partition names are ignored. For example, the object name `Common/rule1` would be sorted without the common partition name, as if it were named `rule1`.
- To resize columns, click the column side and drag it to the preferred location.
- To select what columns are displayed, click the gear icon in the upper right of the Audit Logging screen. In the popup screen, select columns you want to display and clear columns you do not want to display. Move your cursor away from the screen to dismiss it.

Managing audit log archive settings

You can view or change the audit archive settings. The archived audit log files are stored in the `/var/config/rest/auditArchive/` directory on the BIG-IQ® system.

1. Log in to BIG-IQ system with Administrator or Security Manager credentials.
2. Select **Audit Logging** from the BIG-IQ menu.
3. Click the **Archive Settings** button in the upper left of the Audit Logging screen to display the audit log settings.
4. Complete or review the properties and status settings, and click **Save**.

Property	Description
Retain Entries	Specifies the number of days to keep audit log entries. The field must contain an integer between 1 and 366. The default is 30.

Property	Description
Weekly Update	Specifies which days of the week to update the audit log. Select the check box to the left of each day that you want the audit log to be updated. The default is every day.
Start Time	Specifies when the audit archiving should begin. The default is 12:00 am.
Items Expired	Displays the read-only number of entries that have expired.
Last Error	If an error has occurred, displays the read-only error text for any errors found.
Last Error Time	If an error has occurred, displays a read-only value that contains the time the last error was found. The time in the field is the BIG-IQ system local time and is expressed in the format: ddd mmm dd yyyy hh:mm:ss, for example, `Fri Jan 17 2014 23:50:00`.

About archived audit logs

You can view or change how audit logs are archived by clicking the **Archive Settings** button on the Audit Logging screen.

Archived audit log files are stored in the `archive-audit.n.txt` file in the appropriate subdirectory of the `/var/config/rest/auditArchive` directory on the BIG-IQ® Centralized Management system:

- Network Security audit log: `/var/config/rest/auditArchive/networkSecurity/`
- Web Application Security audit log: `/var/config/rest/auditArchive/webAppSecurity/`
- Fraud Protection Service audit log: `/var/config/rest/auditArchive/websafe/`
- Local Traffic and Network audit log: `/var/config/rest/auditArchive/adc/`
- Device Management audit log: `/var/config/rest/auditArchive/device/`
- Access audit log:`/var/config/rest/auditArchive/access/`

Audit entries are appended to the `archive-audit.0.txt` file. When the `archive-audit.0.txt` file reaches approximately 800 MB, the contents are copied to `archive-audit.1.txt`, compressed into the `archive-audit.1.txt.gz` file, and a new empty `archive-audit.0.txt` file is created, which then has new audit entries appended to it.

Up to five compressed archived audit files can be created before those files begin to be overwritten to conserve space. The compressed audit log archive is named `archive-audit.n.txt.gz`, where n is a number from 1 to 5. As the audit log archives are created and updated, the content of the archives is rotated so that the newest archive is always `archive-audit.1.txt.gz` and the oldest is always the highest numbered archive, typically, `archive-audit.5.txt.gz`.

The file content rotation occurs whenever `archive-audit.0.txt` is full. At that time, the content of each `rchive-audit.n.txt.gz` file is copied into the file with the next higher number, and the content of `archive-audit.0.txt` is copied into `archive-audit.1.txt` and then compressed to create `archive-audit.1.txt.gz`. If all five `archive-audit.n.txt.gz`files exist, during the rotation the contents of `archive-audit.5.txt.gz` are overwritten, and are no longer available.

About audit logs in high-availability configurations

In high-availability (HA) configurations, there is a primary and secondary BIG-IQ® system. During failover, the audit log entries and the audit archive settings are copied from the primary to the secondary system before the secondary system becomes the new primary system.

However, archived audit logs are not copied from the primary to the secondary BIG-IQ system.

About the REST API audit log

The REST API audit log records all API traffic on the BIG-IQ® system. It logs every REST service command for all licensed modules in a central audit log (`restjavad-audit.n.log`) located on the system.

Note: *The current iteration of the log is named `restjavad-audit.0.log`. When the log reaches a certain user-configured size, a new log is created and the number is incremented. You can configure and edit settings in `/etc/restjavad.log.conf`.*

Any user who can access the BIG-IQ system console (shell) has access to this file.

Managing the REST API audit log

The REST API audit log contains an entry for every REST API command processed by the BIG-IQ® system, and is an essential source of information about the modules licensed under the BIG-IQ system. It can provide assistance in compliance, troubleshooting, and record-keeping. With it, you can review log contents periodically, and save contents locally for off-device processing and archiving.

1. Using SSH, log in to the BIG-IQ Network Security system with administrator credentials.
2. Navigate to the `restjavad` log location: `/var/log`.
3. Examine files with the naming convention: `restjavad-audit.n.log`.

   The letter *n* represents the log number.
4. Once you have located it, you can view or save the log locally through a method of your choice.

Legal Notices

Legal notices

Publication Date

This document was published on January 26, 2017.

Publication Number

MAN-0520-05

Copyright

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.

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

address lists
about *24*
adding *20*
adding addresses *25*
adding to firewalls and rules *20*
and properties *25*
editing *20*
removing entries *25*
address types
adding to address lists *25*
addresses
adding to address lists *25*
and address list properties *25*
removing from address lists *25*
advanced filtering
for event log *179*
alert logs
about *165*
API (REST) audit log
about *199*
application security
configuring for *92*
application security policies
adding with BIG-IQ ASM *148*
adding with BIG-IQ Web Application Security *148*
exporting with BIG-IQ Web Application Security *149*
importing with BIG-IQ Web Application Security *149*
removing *149*
application security policies properties
description *119*
application security policy properties
displaying *119*
editing *119*
archived audit logs
about *198*
attack signature sets
about custom *155*
assigning custom *162*
assigning to security policies *162*
attack signature sets (custom)
adding *155*
editing *157*
attack signatures (custom)
creating *153*
attack signatures lists
customizing *147*
filtering display *147*
viewing *146*
attack signatures settings
editing *145*
audit log
about REST API *199*
customizing display of *197*
filtering entries *195*
audit log archive settings
managing *197*
audit log display
customizing *197*
audit log entries
filtering *195*
generation of *193*
properties of *194*
audit log filtering
of entries *195*
audit logs
about *193*
in high-availability configurations *198*
viewing differences *195*
authentication keys
configuring for SSH profiles *114*
authentication token
deleting *76*
requesting *75*

B

base registration key
about *171*
basic filtering
for event log *178*
behavior analysis
enabling *97*
BIG-IP device management
about *17*
BIG-IP logging profile
configuring *167*, *172*
BIG-IQ inventory
adding devices to *15*
BIG-IQ Logging Node
about discovering *166*
about management *175*
and best practice *175*
defined *166*
BIG-IQ Network Security
about *9*
BIG-IQ system
about licensing and initial setup for Logging Node *169*
BIG-IQ Web Application Security
about *10*
blocking settings
description *123*
editing *123*
for application security policies *123*
Bot Signatures tab
configuring for application security *92*
browser resolution
about *10*

C

capacity planning
for Logging Node *175*
centralized management
of BIG-IP devices *15*

change verifications
about *67*
adding *67*
properties *68*
viewing properties *68*
changes
about evaluating before deploying *185*
character sets
editing *144*
check rule compilation *188*
cloning process
for objects *21*
common filters
using for event logs *178*
configuration
and initial setup *169*, *171*
configuration changes
about deploying *186*
deploying to a device *187*
evaluating *185*
urgent *187*
configuration deployment
about *186*
configuration objects
about locking *71*
clearing locks *13*
viewing properties *16*
configuration snapshots
about managing *182*, *191*
contexts
firewall *31*
for firewalls, about *31*
current configuration
about *13*
custom attack signature sets
about *155*
adding *155*
assigning to security policies *162*
editing *157*
custom attack signatures
about *153*
creating *153*
custom signatures
about staging *155*
define *153*

D

Data Guard settings
editing *136*
deleting application security policies *149*
deployment
of configuration changes *186*, *187*
state *189*
states during *189*
device DoS
editing *99*
overview *99*
device inventory
about *15*
device management
about *15*
devices
about discovering *15*
adding to BIG-IQ inventory *15*
discovering *15*
Devices panel
using with event logs *178*
differences
in audit logs *195*
viewing in audit logs *195*
discovery
defined *15*
discovery address
defined *169*
DoS attack
defined *91*
DoS profiles
adding *91*
editing *97*
enabling protection types *91*
overview *91*
dossier
providing *169*, *171*

E

Elasticsearch
and Logging Node *175*
emergency deployment
of configuration changes *187*
enforced firewall policies *81*
enforced firewall policy
adding *35*
evaluation
of configuration changes *185*
evaluation of changes
before deploying *185*
event *181*
event log details screen
described *177*
event log filtering
using advanced functions *179*
using basic functions *178*
event log indices
defined *167*
event log snapshots
restoring *181*
event logs
about management interface *177*
configuring snapshot schedules *168*
configuring the logging profile *167*, *172*
using common filters *178*
viewing details *177*
event logs filtering
and query parameters *179*
using ODATA query parameters *179*
exporting application security policies
with BIG-IQ Web Application Security *149*
external logging devices
about *73*
accessing *76*
adding *73*
deleting authentication token *76*

external logging devices *(continued)*
launching user interface *76*
modifying *74*
removing *75*
requesting authentication token *75*
extraction settings
editing *142*

F

features
BIG-IQ Web Application Security *10*
for BIG-IQ Network Security *9*
file types settings
editing *139*
filter
bottom frame of Policy Editor *24*
Policy Editor *22*
toolbox *24*
using *22*
Filter field
and advanced options *179*
and basic options *178*
filter related to *24*
filtering
about *10*
firewall
contexts *31*
firewall contexts
about *31*
customizing the display of *11*
Firewall Contexts
and properties *33*
firewall idle timer
definition *47*
firewall policies
about *81*
adding rules *38*
cloning *83*
creating *81*
creating by cloning *83*
deleting *84*
editing *82*
enforced *81*
managing *82*
managing with snapshots *85*
removing rule lists *42*
removing rules *39*
reordering rules *84*
staged *81*
firewall policy
types of *33*
firewall policy (BIG-IQ Network Security)
adding enforced *35*
adding staged *35*
firewall properties
listed *33*
firewall rule
about reports *79*
applying a service policy *50*
creating reports *79*
deleting reports *80*
firewall rule hit count *77*
firewall rule reports
about *79*
creating *79*
deleting *80*
firewall rules
monitoring *77*, *78*
firewall security locked objects
deleting locks *71*
viewing *71*
firewall security locks
deleting *71*
viewing *71*
firewall security policy editing
by multiple users *12*
firewall types
customizing the display of *11*
firewalls
adding rules *38*
removing rule lists *42*
removing rules *39*
FQDN resolver
about *61*
configuring *61*
fully qualified domain name
definition *61*

G

General Settings tab
configuring for application security *92*
geolocation
adding to address lists *25*
global context
applying a service policy *50*
global firewalls
about *32*
global user settings
customizing *11*

H

headers (allowed) settings
editing *137*
Heavy URL Protection tab
configuring for application security *92*

I

importing application security policies
with BIG-IQ Web Application Security *149*
inactive virtual server
BIG-IQ Web Application Security *10*
index rotation policy
about *176*
default configuration values *176*
initial configuration
for BIG-IQ system *169*, *171*
performing automatically for BIG-IQ system *169*
performing manually for BIG-IQ system *171*
IP address settings
editing *138*

L

license
  activating automatically *169*
  activating manually *171*
license activation
  for BIG-IQ system *169*, *171*
locked objects
  about *71*
  viewing *12*
Locked Objects screen
  about *71*
locks
  clearing *13*
  for configuration objects *71*
log items list
  described *177*
Logging Node
  about capacity planning *175*
  about discovering for BIG-IQ *166*
  about management *175*
  activating *166*
  adding to a Logging Group *166*
  defined *175*
  discovering *166*
  extending file size *175*
  recommended configuration *175*
logging profile
  assigning to a virtual server *174*
  configuring *167*, *172*
  defined *165*
  sending events to Logging Node *167*, *172*
logging profiles
  configuring for application security *102*
  configuring for DoS protection *110*
  configuring for network address translation *109*
  configuring for network firewall *107*
  configuring for protocol security *105*
  creating *101*
  editing *111*
  overview *101*
logs
  restoring snapshots *181*

M

managed devices
  about discovering *15*
  revising *16*
  viewing properties *16*
managed objects
  about evaluating changes before deploying *185*
management IP firewalls
  about *33*
methods (HTTP) settings
  editing *137*
monitoring
  firewall rules *77*, *78*
multi-user editing
  about *12*

N

NAT destination translation
  about *53*
NAT destination translations
  cloning *60*
  creating *59*
  deleting *60*
NAT policy
  cloning *55*
  creating *53*
  definition *53*
  deleting *55*
NAT rules properties
  listed *53*
NAT source translation
  about *53*
NAT source translations
  cloning *58*
  creating *56*
  deleting *59*
nested address lists
  about *24*
network security
  configuring for *97*
Network Security objects
  about *19*
Network Security Reporting
  about *87*
notification email *63*
notification rules
  adding *63*
  deleting *65*
  scheduling *63*
notifications rules
  editing *65*

O

object locking
  for multi-user editing *12*
objects
  about *19*
  adding *20*
  adding to firewalls and rules *20*
  cloning *21*
  duplicating *21*
  editing *20*
  renaming *21*
ODATA
  filtering event logs *179*

P

parameters settings
  editing *140*
policies, *See* firewall policies
policy editor
  about *19*
Policy Editor
  filter *22*
Policy Editor objects

Policy Editor objects *(continued)*
removing *21*
policy, firewall
types of *33*
port list properties *27*
port lists
about *26*
adding *20*
adding ports *27*
adding to firewalls and rules *20*
editing *20*
removing entries *27*
port misuse policy
about *47*
creating *48*
definition *47*
deleting *52*
ports
adding to port lists *27*
preferences
setting *11*
privileges
of user roles *13*
Proactive Bot Defense tab
configuring for application security *92*
profiles
about SSH *113*
creating for SSH *113*
deleting for SSH *115*
properties
for application security policies *119*
for NAT rules *53*
for rule lists *42*
for rules *42*
for schedules *28*
for signature files *151*
of address lists *24*, *25*
of firewall policies *81*
of port lists *27*
of rule lists *37*
viewing for signature files *151*
protocol DNS
configuring for *95*
protocol SIP
configuring for *96*
proxy permissions
configuring for SSH profiles *113*

R

related items
showing *24*
remote logging
virtual servers *174*
removing application security policies *149*
reordering rules
in firewall policies *84*
Reporting screen
about *87*
reports
creating for firewall rules *79*
deleting for firewall rules *80*
reports *(continued)*
for firewall rules *79*
request details
viewing for events *177*
resolution
for browser *10*
response details
viewing for events *177*
response page settings
editing *134*
REST API audit log
about *199*
saving locally *199*
restjavad-audit.n.log
about *199*
roles
about *13*
roll back, *See* snapshots
route domain context
applying a service policy *51*
route domain firewalls
about *32*
rule lists
about *37*
adding *39*
and properties *37*
and properties for *42*
cloning *41*
editing *40*
editing rules *40*
removing *42*
removing rules *39*
reordering rules *38*
rule schedule properties *28*
rule schedules
about *28*
rules
about *37*
adding rule lists *39*
adding to rule lists, firewalls, firewall policies *38*
and cloning rule lists *41*
clearing fields *41*
creating *38*
deleting *39*
deleting fields *41*
editing in rule lists *40*
removing *39*
removing fields *41*
reordering *38*
rules properties
listed *42*

S

schedules
adding *20*
editing *20*
screens
customizing the display of *11*
Search filter
for event logs *179*
for Event Logs *178*

security policy properties
displaying 119
editing 119
security reporting
about 87
security user settings
customizing 11
self IP firewalls
about 33
service policies
about 47
service policy
applying to firewall rule 50
applying to global context 50
applying to route domain 51
applying to self IP address 51
creating 49
definition 47
deleting 52
sets
about custom attack signature 155
sets (custom attack signature)
adding 155
editing 157
setup
for BIG-IQ Logging Node 169
for BIG-IQ system 169, 171
SevOne PLA logging device
about 73
accessing 76
adding 73
deleting authentication token 76
launching user interface 76
modifying 74
removing 75
requesting authentication token 75
Shared Security
about 9
signature file properties
listed 151
viewing 151
signature files
about 151
in BIG-IQ Web Application Security 151
synchronizing 152
updating and pushing 152
signature sets
about custom attack 155
signature staging
about 155
signatures
about staging 155
signatures advanced filter properties
described 159
snapshot
creating 182, 191
restoring 183, 192
snapshot locations
defining 167
snapshot management
about 182, 191
snapshot schedules
snapshot schedules *(continued)*
defining 168
snapshot storage
defining locations 167
snapshots
comparing 182, 191
defining schedules 168
managing BIG-IQ Network Security policies 85
restoring 181
SSH profiles
configuring authentication keys 114
configuring proxy permissions 113
creating 113
deleting 115
overview 113
staged firewall policies 81
staged firewall policy
adding 35
state
deployment 189
Stress-based Detection tab
configuring for application security 92
system interface
about 10
and filtering 22
filtering 10
system license
about 169
system snapshot
restoring 183, 192
system snapshots
comparing 182, 191
system-supplied signatures
define 153

T

timer policies
about 47
timer policy
creating 47
definition 47
deleting 52
TPS-based Detection tab
configuring for application security 92

U

update and push process
for signature files 152
user preferences
setting 11
user roles
about 13
user-defined attack signature sets
adding 155
editing 157
user-defined attack signatures
creating 153

V

VIP firewalls, *See* virtual server firewalls
virtual server
    adding for logging profile assignment *174*
    creating for logging profile *174*
virtual server firewalls
    about *32*
virtual server properties
    displaying *163*
virtual servers
    about *163*
    and Web Application Security *163*
    displaying properties *163*
    editing *89*
    for remote logging *174*
    overview *89*

W

Web Application Security policies
    about *117*
web application security policy editing
    by multiple users *12*
Web Application Security Reporting
    about *87*
working configuration
    about *13*
    defined *9*