F5® BIG-IQ® Centralized Management: Security

Version 5.4

Table of Contents

Overview: BIG-IQ Centralized Management Security

Understanding Network Security and firewall management

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 BIG-IP Advanced Firewall Manager™ (AFM™) module.

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). 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 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 Network Security working configuration set. Unless local changes are reconciled, the deployment process overwrites any local changes.

In addition, 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

BIG-IQ® Centralized Management 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 Web Application Security and application management

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®Centralized Management 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 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 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.

About BIG-IQ Centralized Management configuration sets

The BIG-IQ® Centralized Management system 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 Centralized Management. The *current configuration* is updated during a re-discover and re-import, and before calculating differences during the deployment process.

Working configuration set

The configuration as maintained by BIG-IQ Centralized Management. The *working configuration* is the configuration that is edited on BIG-IQ Centralized Management and deployed back to BIG-IP devices.

The working configuration is created when the administrator first manages the BIG-IP device from the BIG-IQ Centralized Management system. The working configuration is updated when a device is re-imported or re-discovered.

If conflicts are observed during a re-discover and re-import, the object in conflict is only updated in the working configuration when the **Use BIG-IP** resolution conflict option is used.

About managing BIG-IP devices with BIG-IQ Centralized Management

Once you have placed a BIG-IP® device under management by the BIG-IQ® Centralized Management 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 Centralized Management system to avoid errors.

A BIG-IP software release may include features that the BIG-IQ Centralized Management system does not yet manage. If changes are made to the configuration of that feature directly on the BIG-IP device, the BIG-IQ Centralized Management system might remove those changes when a subsequent deployment is made to the BIG-IP device.

During the deployment process, the BIG-IQ Centralized Management 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 Centralized Management system to reconcile the configuration differences.

Filtering content in firewall policies

There are several filter fields you can use to select the data displayed for firewall objects. 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. Go to **Configuration** > **SECURITY** > **Network Security**.
2. Edit one of the firewall policy objects, such as the firewall policy.
3. In the appropriate filter text field, type the text you want to filter on, and press Enter.

Option	Description
Filter field at top right of screen	Use the filter field at the right top of the screen 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.

Option	Description
	• **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. This match is not case-sensitive. 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 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 Toolbox at bottom	Use the filter field in the upper right of the toolbox (displayed at the bottom of the page when active) to search the shared objects 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 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.*

Managing Object Pinning

What is object pinning?

You *pin* an object, such as a logging profile, to a pinning policy to have it included in a deployment. The pinning policy is associated with a BIG-IP® device and has the same name as the BIG-IP device. You do not create pinning policies. Pinning policies always exist to contain objects that get pinned to a policy.

You pin an object to a pinning policy for a BIG-IP device to mark the object as being used by the BIG-IP device configuration, and to have it deployed with that configuration and not deleted from the device. When an object is pinned for deployment to a BIG-IP device that is part of a cluster, the object is deployed to the other member of the cluster as well.

You use the Pinning Policies screen to pin policy objects so that they are deployed to a BIG-IP device, or to view the objects that are already pinned to be deployed to a BIG-IP device. The objects that can be selected for pinning differ depending on which service is being used. For example, only the Network Security service allows you to pin firewall policy objects, and only the Local Traffic service allows you to pin SMTP server objects. You can pin objects to, or unpin objects from, multiple BIG-IP device pinning policies at once.

***Note:** Both the system and users can pin an object. But users can unpin only objects that are labeled as user pinned. For easy identification, objects pinned by a user are listed with the User identifier in the Pin Source Tags column on the Pinning Policy Properties screen. Any user can unpin a user pinned object.*

Pin objects to a BIG-IP device pinning policy

You pin objects, such as logging profiles, to BIG-IP®device pinning policies to ensure that the objects are deployed to BIG-IP devices. The process for pinning to a single BIG-IP device pinning policy differs from the process for pinning to several BIG-IP device pinning policies.

1. Open the Pinning Policies screen. How you access the screen depends on the service you are using.
   - To pin Local Traffic service objects, click **Configuration** > **LOCAL TRAFFIC** > **Pinning Policies**.
   - To pin Network Security service objects, click **Configuration** > **SECURITY** > **Network Security** > **Pinning Policies**.
   - To pin Shared Security service objects, click **Configuration** > **SECURITY** > **Shared Security** > **Pinning Policies**.
   - To pin Access service objects, click **Configuration** > **ACCESS** > **Access Groups** > **Pinning Policies**. An Access group must exist to see this menu item.
2. Decide whether to pin to a single BIG-IP device pinning policy, or multiple BIG-IP device pinning policies.
   - Go to Step 3 to pin objects to a single BIG-IP device pinning policy.
   - Go to Step 4 to pin objects to multiple BIG-IP device pinning policies.
3. To pin objects to a pinning policy for a single BIG-IP device:
   a) Click the name of the BIG-IP device pinning policy to which you will pin objects. (It has the same name as the associated BIG-IP device.)
   The properties screen opens.
   b) At the top of the area near the bottom of the screen, select the type of object to be pinned.

The screen lists objects of the type you selected.

c) Select the check box to the left of the objects to be pinned, and click **Add Selected**.

4. To pin objects to multiple BIG-IP device pinning policies:

a) Select the check boxes for the BIG-IP device pinning policies to which to pin objects, and click **Pin to Multiple Policies**.
The properties screen opens and displays the selected BIG-IP device pinning policies.

b) In the area near the bottom of the screen, select the type of object to be pinned.
The screen lists objects of the type you selected.

c) Select the check box for objects to be pinned and click **Add Selected**.

5. Save your work.

A dialog box displays the success of the pinning operation. The object, or objects, are pinned to the pinning policy for the BIG-IP device, or devices, and will be deployed with them.

Unpin objects from a BIG-IP device pinning policy

You unpin objects, such as logging profiles, from a BIG-IP® device pinning policy when they no longer need to be deployed with the BIG-IP device. The process for unpinning from a single BIG-IP device pinning policy differs from the process for unpinning from several BIG-IP device pining policies.

***Note:** Both the system and users can pin an object. But users can unpin only objects that are labeled as user pinned. For easy identification, objects pinned by a user are listed with the User identifier in the Pin Source Tags column on the Pinning Policy Properties screen. Any user can unpin a user pinned object.*

1. Open the Pinning Policies screen. How you access the screen depends on the service you are using.

- To unpin Local Traffic service objects, click **Configuration** > **LOCAL TRAFFIC** > **Pinning Policies**.
- To unpin Network Security service objects, click **Configuration** > **SECURITY** > **Network Security** > **Pinning Policies**.
- To unpin Shared Security service objects, click **Configuration** > **SECURITY** > **Shared Security** > **Pinning Policies**.
- To pin Access service objects, click **Configuration** > **ACCESS** > **Access Groups** > **Pinning Policies**. An Access group must exist to see this menu item.

2. Decide whether to unpin from a single BIG-IP device pinning policy, or from multiple BIG-IP device pinning policies.

- Go to Step 3 to unpin objects from a single BIG-IP device pinning policy.
- Go to Step 4 to unpin objects from multiple BIG-IP device pinning policies.

3. To unpin objects from a single BIG-IP device pinning policy:

a) Click the name of the BIG-IP device pinning policy from which to unpin objects.
The properties screen opens.

b) In the Selected Resources area, expand the resource type of the object you want to unpin.
The screen lists objects of the type you selected.

c) Select the check box for the objects to be unpinned and click **Remove**.

Both the system and users can pin an object. But users can unpin only objects that are labeled as user pinned. For easy identification, objects pinned by a user are listed with the User identifier in the Pin Source Tags column on the Pinning Policy Properties screen. Any user can unpin a user pinned object.

4. To unpin objects from multiple BIG-IP device pinning policies:

   a) Select the check boxes for the BIG-IP device pinning policies from which to unpin objects, and click **Unpin from Multiple Policies**.
      The properties screen opens and displays the selected BIG-IP device pinning policies.
   b) In the lower area of the screen, select the type of object to be unpinned.
      The screen lists objects of the type you selected.
   c) Select the check box for the objects to unpin and click **Add Selected**.

      The Selected Resources area lists the objects to be unpinned. Both the system and users can pin an object. But users can unpin only objects that are labeled as user pinned. For easy identification, objects pinned by a user are listed with the User identifier in the Pin Source Tags column on the Pinning Policy Properties screen. Any user can unpin a user pinned object.

5. Save your work.

A dialog box displays the success of the unpinning operation. The object or objects are unpinned from the BIG-IP device pinning policy and will no longer be deployed to it.

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 Address Lists

About address lists

Address lists, also called network address lists, are collections of IPv4 or IPv6 addresses, address ranges, nested address lists, geolocations, and subnets. These can be used by other parts of the BIG-IQ® Centralized Management system, such as firewall rules or firewall policies.

You can manage address lists from the following locations:

- **Configuration** > **NETWORK** > **Address Lists**
- **Configuration** > **SECURITY** > **Network Security** > **Address Lists**

Be aware of the following considerations about address lists.

- 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.
- 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.
- To pin an address list to a deployment, you must do so from the Local Traffic pinning policy user interface: **Configuration** > **LOCAL TRAFFIC** > **Pinning Policies**.
- You can add geolocation awareness to address lists, which enables you to specify source or destination IP addresses by geographic location rather than by their IP addresses. The geolocation is validated when the address list is saved. If you use a geolocation specification that is valid on the BIG-IQCentralized Management 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 specification causes a discovery failure for that device.

Create address lists

You create address lists so that you can use them with other parts of the BIG-IQ® Centralized Management system, such as firewall rules. Address lists are a collection of addresses. You can access address lists from either the network or the network security configuration menu.

- To use the network configuration, click **Configuration** > **NETWORK** > **Address Lists**.
- To use the security configuration, click **Configuration** > **SECURITY** > **Network Security** > **Address Lists**.

1. Open the Address Lists screen.

   You can access the address list from either the network or network security configuration menu and it will behave in the same way.
2. Click **Create**.
   The New Address List screen opens.
3. On the left, click **Properties**.
4. Supply the properties for the address list.
   - In the **Name** setting, type a unique name for the address list.
   - In the **Description** setting, type an optional description for the address list.
   - In the **Partition** setting, type a partition if needed. The `Common` partition is the default.
5. On the left, click **Addresses**.
6. Supply the addresses for the address list.

The screen displays a template address for you to complete. An address list must contain at least one address.

7. In the **Type** column, select the address type, and then provide the address information in the **Addresses** column. You can also add a description for each address in the **Description** column.

   - To add a single address, select **Address** and type an IPV4 or IPV6 address.
   - To add an address list, select **Address List** and select the name of the address list.
   - To add a range of addresses, select **Address Range** and type the beginning and ending IPV4 or IPV6 addresses.
   - To add a location to the address list, select **Country/Region** and select the country and optionally, the region of the country. You can also select `Unknown` as the country or region option. Address locations can be used when defining rules based on where a system is located (the geolocation of the system), rather than on the IP address of the system.
   - To add a domain name, select **Domain Name** and type the domain name.
8. In the **Add/Remove** column, click **+** to add the address to the list.

   You can click **X** to delete an address from the list.
9. Continue to add or delete addresses to the address list until the address list is complete.
10. Save your work.

Edit address lists

You edit address lists to change the properties of the address list or to add, modify, or remove addresses from the address list, or both. You can access address lists from either the network or the network security configuration menu.

- To use the network configuration, click **Configuration** > **NETWORK** > **Address Lists**.
- To use the security configuration, click **Configuration** > **SECURITY** > **Network Security** > **Address Lists**.

1. Open the Address Lists screen.

   You can access an address list from either area and it will behave in the same way.
2. Click the name of the address list to edit it.
3. To modify the address list **Description**, click **Properties** and in the **Description** setting, type or revise an optional description for the address list.
4. On the left, click **Addresses**.
5. Add, modify, or delete addresses for the address list.

   - To modify that address, click the pencil icon to the left of the address.
   - To delete an address, click **X** in the **Add/Remove** column.
   - To add an address, click **+** in the **Add/Remove** column.

   An address list must contain at least one address.
6. If you are adding or modifying an address, supply or modify the settings.

   In the **Type** column, select the address type, and then provide the address information in the **Addresses** column. You can also add a description for each address in the **Description** column.

   - To add a single address, select **Address** and type an IPV4 or IPV6 address.
   - To add an address list, select **Address List** and select the name of the address list.
   - To add a range of addresses, select **Address Range** and type the beginning and ending IPV4 or IPV6 addresses.

- To add a location to the address list, select **Country/Region** and select the country and optionally, the region of the country. You can also select `Unknown` as the country or region option. Address locations can be used when defining rules based on where a system is located (the geolocation of the system), rather than on the IP address of the system.
- To add a domain name, select **Domain Name** and type the domain name.

7. In the **Add/Remove** column, click **+** to add the address to the list.

   You can click **X** to delete an address from the list.
8. Continue to add, modify, or delete addresses in the address list until the address list is complete.
9. Save your work.

Clone address lists

You can clone an address list to create a copy of it, which you can then edit to address any special considerations. You can access address lists from either the network or the network security configuration menu.

- To use the network configuration, click **Configuration** > **NETWORK** > **Address Lists**.
- To use the security configuration, click **Configuration** > **SECURITY** > **Network Security** > **Address Lists**.

1. Open the Address Lists screen.

   You can access an address list from either area and it will behave in the same way.
2. Select the check box next to the address list to clone.
3. Click **Clone**.
   The system makes a copy of that address list with the same name, but with `-CLONE` appended to the name and a blank **Description** field.
4. Change the address list properties and contained addresses as needed, such as providing a meaningful name or changing an address within the list.
5. Save your work.

The new address list is now defined and you can assigned it to an object.

Deploy address lists

If you want to do a quicker deployment by only deploying the address list portion of a configuration, you can do a partial deployment of the address list, instead of deploying the entire configuration. You can access address lists from either the network or the network security configuration menu.

- To use the network configuration, click **Configuration** > **NETWORK** > **Address Lists**.
- To use the security configuration, click **Configuration** > **SECURITY** > **Network Security** > **Address Lists**.

1. Open the Address Lists screen.

   You can access an address list from either area and it will behave in the same way.
2. Select the check box next to the address list to deploy.
3. Click **Deploy**.

The system displays the selected address list, with options for partial deployment selected. You can now continue the partial deployment process.

Delete address lists

You delete address lists you no longer use to avoid confusion in the user interface. You can access address lists from either the network or the network security configuration menu.

- To use the network configuration, click **Configuration** > **NETWORK** > **Address Lists**.
- To use the security configuration, click **Configuration** > **SECURITY** > **Network Security** > **Address Lists**.

1. Open the Address Lists screen.

   You can access an address list from either area and it will behave in the same way.
2. Click the check box next to the address list to delete.
3. Click **Delete**.
4. In the confirmation dialog box that opens, click **Delete** to confirm the removal.

   If the address list is pinned to a BIG-IP device pinning policy, the deletion will fail.

Managing Port Lists

About port lists

Port lists are collections of ports, port ranges, or port lists that can be assigned to firewall rules.

Firewall rules use 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 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.

***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.*

Create port lists

You create port lists so that you can use them when creating firewall rules.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Port Lists**.
2. Click **Create**.
   The New Port List screen opens.
3. On the left, click **Properties** and specify the settings for the port list.
   a) In the **Name** setting, type a unique name for the port list.
   b) In the **Description** setting, type an optional description for the port list.
   c) In the **Partition** setting, type a different partition if needed. The `Common` partition is the default.
4. On the left, click **Ports**.

   You must supply at least one entry in the port list, such as a port, port list, or port range.

   The screen displays a blank port entry template for you to complete.
5. In the Type column, supply the type of port entry to add to the port list.
   - To add a single port, select **Port**.
   - To add a port list, select **Port List**.
   - To add a range of ports, select **Port Range**.

   ***Note:** Before nesting a port list inside a port list, check that this option is supported on the specific version of your BIG-IP® device.*
6. In the Ports column, supply the port details for the port type you selected.
   - If you selected **Port** as the type, type a port number.
   - If you selected **Port List** as the type, select the name of the port list.
   - If you selected **Port Range** as the type, type the beginning and ending port numbers.
7. In the Description column, provide an optional meaningful description of the port entry.
8. To add more than one port entry to the port list, click the **+** in the **Add/Remove** column, and provide the details.

   Click **Update** to update and save the port entry you are currently editing.
9. Continue to add or delete ports until the port list is complete.
10. Click **Save** to save your work on the port list.

Edit port lists

You edit port lists to change the properties of the port list, or to add, modify, or remove port entries from the port list, or both.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Port Lists**.
2. Click the name of the port list to edit.
3. If you are modifying the port list description, click **Properties**, and type or modify the **Description** setting.

   Only the description property can be modified.
4. If you are modifying the port entries, on the left, click **Ports**.
5. Add, modify, or delete port entries in the port list.
   - To modify a port entry, click the pencil icon in that row.
   - To delete a port entry, click **X** in the **Add/Remove** column.
   - To add a port entry, click **+** in the **Add/Remove** column.

   A port list must contain at least one port entry.
6. If you are adding or modifying a port entry, supply or modify the settings.
7. In the Type column, supply the type of port entry to add to the port list.
   - To add a single port, select **Port**.
   - To add a port list, select **Port List**.
   - To add a range of ports, select **Port Range**.

   Note: *Before nesting a port list inside a port list, check that this option is supported on the specific version of your BIG-IP® device.*
8. In the Ports column, supply the port details for the port type you selected.
   - If you selected **Port** as the type, type a port number.
   - If you selected **Port List** as the type, select the name of the port list.
   - If you selected **Port Range** as the type, type the beginning and ending port numbers.
9. In the Description column, provide an optional meaningful description of the port entry.
10. Continue to add, modify, or delete port entries in the port list until the port list is complete.
11. Save your work.

Clone port lists

You can clone a port list to create a copy of it, which you can then edit.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Port Lists**.

   The Port Lists screen opens.
2. Click the check box next to the port list you want to clone.
3. Click **Clone**.
   A copy of that port list is created with the same name, but with `-CLONE` appended to the name and a blank **Description** field.
4. Change the port list as needed.
5. Save your work.

The new port list is defined and you can now assign it to a firewall rule.

Deploy port lists

If you want to do a quicker deployment by only deploying the port list portion of a configuration, you can do a partial deployment of the port list, instead of deploying the entire configuration.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Port Lists**.
   The Port Lists screen opens.
2. Click the check box next to the port list you want included in the partial deployment.
3. Click **Deploy**.

The system displays the selected port list, with options for partial deployment selected.

Continue the partial deployment process.

Delete port lists

You can delete port lists that you no longer use to avoid confusion in the user interface.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Port Lists**.
   The Port Lists screen opens.
2. Click the check box next to the port list to delete.
3. Click **Delete**.
4. In the confirmation dialog box that opens, click **Delete** to confirm the removal.

Managing Rule Schedules

About rule schedules

By default, all rules run continuously. You can apply a rule schedule to a rule to make that rule active only when scheduled.

Use the Rule Schedules screen to view the defined rule schedules. Rule schedules are *continuously active* if created without any scheduling specifics (such as the hour that the rule schedule starts).

Create rule schedules

You create rule schedules so that you can control when firewall rules are active.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Rule Schedules**.
2. Click **Create**.
   The New Rule Schedule screen opens.
3. In the **Name** setting, type a unique name for the rule schedule.
4. In the **Description** setting, type an optional description for the rule schedule.
5. In the **Partition** setting, type a partition if needed. The `Common` partition is the default.
6. In the **Date Range** setting, specify the date and time when the rule can be active.

   - **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 window.
   - **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 window.

   Using the calendar settings to specify the start and end dates and times is the preferred method. However, if you do specify dates by typing, use the format: MMM DD, YYYY for the date.
7. In the **Time Range** setting, specify 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 that the rule schedule starts at the specified time and runs until the specified end time. Click in the fields to choose the start and end times.
8. In the **Day** setting, specify the days the rule schedule is active. Select the check boxes for all days that apply.

   You must select at least one day per week.
9. Save your work.

The rule schedule is defined and you can assign it to a rule in a management IP firewall context, firewall policy, or rule list.

Clone rule schedules

You can clone a rule schedule to create a copy of it, which you can then edit to address any special considerations.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Rule Schedules**.

   The Rule Schedules screen opens.
2. Click the check box next to the rule schedule to clone and click**Clone**.
   The system makes a copy of that rule schedule with the same name, but with `-CLONE` appended to the name, and a blank **Description** field.
3. Change the rule schedule as needed.
4. Save your work.

The new rule schedule is now defined.

Delete rule schedules

You delete rule schedules that you no longer use to avoid confusion in the user interface.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Rule Schedules**.

   The Rule Schedules screen opens.
2. Click the check box next to the rule schedule that you want to delete.
3. Click **Delete**.
4. In the confirmation dialog box that opens, click **Delete** to confirm the removal.

The rule schedule is removed from the list.

Managing Rules and Rule Lists

About rules and rule lists

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.

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.

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.

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 **Firewall 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**.

Option	Description
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.

Reorder rules in rule lists

You can optimize your network security firewall policy by reordering rules in rule lists to change the order in which they are evaluated. Rules are evaluated from top to bottom in the list (lowest Id number first, highest Id number last).

1. Click **Configuration** > **SECURITY** > **Network Security** > **Rule Lists**.
2. Click the specific rule list you want to edit in the right pane.
3. On the left, click **Rules** to ensure that it is selected.
4. Drag and drop the rules until they are in the correct order.

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

   Alternatively, you can reorder rules using the **Cut Rule** option. Right-click 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.

   Note: *You can use* ***Copy Rule*** *and then paste rules between rule lists. However, if you use* ***Cut Rule*** *and then paste between rule lists, the cut rule will not be removed from the rule list.*
5. 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.

Option	Description
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**.

Option	Description
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 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, the drag-and-drop function does 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 paste, 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. On the left, click **Rules** to ensure that 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.

Option	Description
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 Viewer or Network Security Deployer cannot clone policies.*

1. Click **Configuration** > **SECURITY** > **Network Security** > **Rule Lists**..

   The Rule Lists screen opens.
2. Click the checkbox to the left of the rule list to clone, and click **Clone**.
3. 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.

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

   You can also click **Create Rule** to add a new rule.
5. 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 check box 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 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 describes the properties required when you are configuring network firewall rules.

Property	Description
ID	The evaluation order identifier of the rule within the policy. Rules are evaluated from the lowest number to the highest. If a rule is contained within a rule list, it will be numbered with the number of the rule list, with the contained rule numbered after the decimal point. For example, a policy with 3 rules, followed by a rule list containing 2 rules, followed by another rule outside of the rule list, would be numbered as: `1, 2, 3, 4, 4.1, 4.2, 5`. In the example, 4 represents the rule list, and 4.1 and 4.2 are the rules within that rule list.
Name	In a rule list, the unique, user-provided name for the rule. Alternatively, in a firewall context or firewall policy, a rule list name, preceded by: `Reference_To_` , such as `Reference_To_sys_self_allow_all`.
Address (Source or Destination)	An IPv4 or IPv6 source or destination IP address, address range, or address list, to which the firewall rule applies. • **Address** specifies an IP address. You type a single address in the **Addresses** field. • **Address Range** specifies a range of IP addresses. You specify the beginning and ending addresses of the range in the areas provided. • **Address List** specifies a list that contains IP addresses. You can select the address list from those listed. • **Domain Name** specifies a valid domain name. • **Country/Region** specifies a country and optionally a region. Once you select a country, the second list automatically updates with all available regions for that country. You can specify `Unknown` as the country if needed. Note that geolocation information, such as the country and region, is not supported on the management IP firewall context. ***Note:*** *You can specify subnets using forward slash (/) notation using either IPv4 or IPv6, such as `60.63.10.0/24` or `2001:db8:a::/64`. You can also append a route domain to an address using the format %RouteDomainID/Mask. For example, `12.2.0.0%44/16`.* You can add additional addresses, address ranges, address lists, or countries/regions (**Add**) and delete addresses, address ranges, address lists, or countries/regions (**X**). To recover an address that was marked for deletion using **X**, re-enter the address and click **Add**.

Property	Description
Port (Source or Destination)	Specifies source or destination port entries (ports, port ranges, or port lists) to which the firewall rule applies. • **Port** specifies a port number. • **Port Range** specifies a range of port numbers. You specify the beginning and ending port numbers in the range in the areas provided. • **Port List** specifies a list of port entries, such as ports or port ranges. You can select the port list from those listed. You can add additional ports, port ranges, or port lists (**Add**) and to delete ports, port ranges, or port lists (**X**). To recover a port that was marked for deletion using **X**, re-enter the port and click **Add**.
VLAN (Source)	Specifies a VLAN or tunnel from which the packet source originates, to which the rule applies. This VLAN is 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`.
Action	Specifies the action taken when the firewall rule is matched, such as whether it is accepted or rejected. • **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.
iRule	Specifies an iRule that is applied to the rule. Optionally, you can enter a number in the **Sampling Rate** field to indicate how often to take a sample. 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.
Protocol	Specifies the IP protocol to compare against the packet. If you select **ICMP**, or **IPv6-ICMP**, additional fields open where you can specify **Type** and **Code** combinations. If you select **Other**, only a **Type** field is displayed. 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	Specifies the activity state of the rule, such as whether it is enabled or disabled. • **disabled** specifies that the rule does not apply at all. • **enabled** specifies that the system applies the firewall rule to the given context and addresses. • **scheduled** specifies that the system applies the rule according to the specified schedule.
Send to Virtual	Specifies a virtual server to which packets matched by the firewall rule classifiers are routed. When a firewall rule is routed to a virtual server, the firewall rule action is not applied. This option is available only for rules on the global, route domain, or self IP context.
Service Policy	Specifies a service policy to associate with a rule. A service policy allows you to associate network idle timers or timer policies with firewall contexts and rules. You can add a service policy to a rule by dragging the service policy from the Shared Objects area onto the Service Policy column for the rule. This field is available with BIG-IP devices version 12.0 or higher.
Log	Specifies whether 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).

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 with 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, and then deploy it to a BIG-IP device version 12.0, or later.

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. Then you associate the service policy with the firewall context or rule.

- You use a *timer policy*, also known as a *firewall idle timer*, to configure timer 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, and then deploy it to a BIG-IP device version 12.0, or later.
- You use a *port misuse policy* 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, 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.

1. Navigate to the Timer Policies screen: Click **Configuration** > **SECURITY** > **Network Security** > **Timer Policies**.
2. Click **Create**.
   The New Timer Policy screen opens.
3. In the **Name** field, type a name for the timer policy.
4. In the **Description** field, type an optional description for the timer policy.
5. If needed, change the default `Common` partition in the **Partition** field.
6. To add timer rules, click the **Rules** on the left, and click **Create Rule**.
   A new rule is displayed with default name and values.
7. Click the edit icon to the left of the new rule to enable editing for the rule fields.
8. In the **Name** field, you may specify a more meaningful name than the default.
9. From the **Protocol** list, select the protocol to be used.

   If you select **all-other**, the rule will apply to all protocols not specified in another timer rule in the policy.
10. 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 **Add**. 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 **Add**. 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**.

11. 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.
12. Save your changes.

The timer policy is now configured.

You now need to add the timer policy to a service policy.

Clone a timer policy

Using the clone function, you can make a copy of a timer policy, and then modify it.

1. Navigate to the Timer Policies screen: click **Configuration** > **SECURITY** > **Network Security** > **Timer Policies**.
2. Select the check box to the left of any timer policy you want to clone.
3. Click **Clone**.

The system displays the New Timer Policy screen with the cloned policy displayed.

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. Navigate to the Timer Policies screen: click **Configuration** > **SECURITY** > **Network Security** > **Timer Policies**.
2. Select the check box to the left of any timer policy that you want to remove.
3. Click **Delete**.
4. Confirm that you want to remove the timer policy by clicking **Delete** in the confirmation dialog box.

The system removes the selected timer policies.

Create a port misuse policy

You create a port misuse policy containing port misuse rules to add to a service policy.

1. Go to the Port Misuse Policies screen: Click **Configuration** > **SECURITY** > **Network Security** > **Port Misuse Policies**.
2. Click **Create**.
   The New Port Misuse Policy screen opens.

3. Type a name and an optional description for the port misuse policy.
4. If needed, change the default `Common` partition in the **Partition** field.
5. 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.
6. To add port misuse rules, on the left, click **Rules**, and then click **Create Rule**.
   The screen displays a new port misuse rule with default name and values.
7. Click the edit icon to the left of the name of the new rule to enable editing for the rule fields.
8. In the **Name** field, you may specify a more meaningful name than the default.
9. 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.
10. In the **IP Protocol** field, select the IP protocol for the port matching rule.
11. 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.
12. 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.
13. 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.
14. Save your changes.

You have configured the port misuse policy.

You now can add the port misuse policy to a service policy.

Clone a port misuse policy

Using the clone option, you can make a copy of a port misuse policy that you can modify.

1. Navigate to the Port Misuse Policies screen: click **Configuration** > **SECURITY** > **Network Security** > **Port Misuse Policies**.
2. Select the check box to the left of any port misuse policy you want to clone.
3. Click **Clone**.

The system displays the New Port Misuse Policy screen with the cloned policy displayed.

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. Navigate to the Port Misuse Policies screen: click **Configuration** > **SECURITY** > **Network Security** > **Port Misuse Policies**.
2. Select the check box to the left of any port misuse policy that you want to remove.
3. Click **Delete**.
4. 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.

Create a service policy

You create a service policy to contain timer policies, port misuse policies, or both. Service policies can be applied to firewall contexts and added to a rule in a rule list or a rule on a security policy.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Service Policies**.
2. Click **Create**.
   The New Service Policy screen opens.
3. In the **Name** field type a name for the service policy.
4. If needed, change the default `Common` partition in the **Partition** field.
5. In the **Description** field, type an optional description for the service policy.
6. If needed, select a timer policy from those listed in the **Timer Policy** list.

   If no timer policy is listed, create one and then assign it to the service policy.
7. If needed, select a port misuse policy from those listed in the **Port Misuse Policy** list.

   If no port misuse policy is listed, create one and then assign it to the service policy.
8. Save your changes.

You have defined the service policy. You can now assign it to a firewall context. You can also add it to a rule in a rule list, or a rule on a security policy.

Clone a service policy

Using the clone option, you can make a copy of a service policy to modify..

1. Go to the Service Policies screen: Click **Configuration** > **SECURITY** > **Network Security** > **Service Policies**.
2. Select the check box to the left of any service policy you want to clone.
3. Click **Clone**.

The system displays the New Service Policy screen with the cloned policy displayed.

Deploy a service policy

You can do a partial deployment of only a service policy instead of an entire configuration.

1. Go to the Service Policies screen: Click **Configuration** > **SECURITY** > **Network Security** > **Service Policies**.
2. Select the check box to the left of any service policy you want to deploy.
3. Click **Deploy**.

The system displays the New Deployment - Network Security screen with the selected service policy on it. You can now continue the deployment process.

Delete a service policy

You can delete service policies that are no longer used, to simplify your view, using the Service Policies screen.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Service Policies**.
2. Select the check box to the left of any service policy you want to remove.
3. Click **Delete**.
4. In the confirmation dialog box, click **Delete** to confirm that you want to remove the service policy.

The system removes the selected service policies.

Apply a service policy to a firewall context

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

1. Navigate to the Contexts screen: Click **Configuration** > **SECURITY** > **Network Security** > **Contexts**.
2. Click the name of the context to open it for editing.
3. 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 Objects 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.
4. Save your changes.

The service policy is now associated with the context.

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. Display the list of rules from a rule list or from a firewall security policy.

Option	Description
If the rule is in a rule list:	Navigate to the Rule Lists screen: click **Configuration** > **SECURITY** > **Network Security** > **Rule Lists**. Click the name of the rule list containing the rule. The screen lists the rules.

Option	Description
If the rule is associated with a policy:	Navigate to the Firewall Policies screen: click **Configuration** > **SECURITY** > **Network Security** > **Firewall Policies**. Click the name of the policy containing the rule. The screen lists the rules.

2. To make it editable, click the edit icon to the left of the name of the rule to which you want to add the service policy.
3. 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 Objects area, select **Service Policies**, and then drag the service policy from that list and drop it into the Service Policy column for the rule.

4. Save your changes.

The service policy is added to the rule.

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 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® Centralized Management system, and then deploy it to a BIG-IP device version 12.1 or later.

Note: *When you view differences that include NAT policy changes to the global context, those changes appear under the global-device-context object rather than the global object.*

Create a NAT policy

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

1. Go to the NAT Policies screen: Click **Configuration** > **SECURITY** > **Network Security** > **Network Address Translation** > **NAT Policies**.
2. Click **Create**.
   The New NAT Policy screen opens with the Properties displayed.
3. Type a name for the NAT policy in the **Name** field.
4. Type an optional description for the NAT policy in the **Description** field.
5. If needed, change the default `Common` partition in the **Partition** field.
6. On the left, click **Rules** and then click **Create Rule**.
   A new row appears in the table of rules. The row contains a rule template, including defaults, for the new rule.
7. Click the edit icon to the left of the rule name to edit the default rule properties.
8. 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.
9. Save your changes.

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

NAT rule properties

This 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	Enter the name of a logging profile in the field. This logging profile must already be defined using Logging Profiles in Shared Security and should be pinned to the BIG-IP device using the Shared Security pinning policy.
State	Specify whether the rule is enabled or disabled. The field is updated.

Create NAT source translations

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

1. Click **Configuration** > **SECURITY** > **Network Security** > **Network Address Translation** > **NAT Source Translations**.
2. Click **Create**.
   The New NAT Source Translations screen opens.
3. Type a name for the NAT source translations in the **Name** field.
4. In the **Description** field, type an optional description for the NAT source translations.
5. If needed, change the default `Common` in the **Partition** field.
6. From the **Type** list, specify 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.
7. If you selected **Static NAT** for the **Type**, 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.
ICMP Echo	Specify whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes.

Property	Description
	• Select **disabled** to disable ICMP echoes.
Egress Interfaces	Specify 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.

8. If you selected **Static PAT** for the **Type**, 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.
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	Specify whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes. • Select **disabled** to disable ICMP echoes.
Egress Interfaces	Specify whether egress interfaces are available. • Select **Disabled on** to disable egress filtering interfaces. • Select **Enabled on** to disable egress filtering interfaces.

9. If you selected **Dynamic PAT** for the **Type**, 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	Specify whether ICMP echoes are available. • Select **enabled** to enable ICMP echoes. • Select **disabled** to disable ICMP echoes.

Property	Description
PAT Mode	Specify 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	Specify 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	Specify 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. • 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	Enables or disables hairpinning for incoming connections to active translation end-points (address/port combinations). Specify 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	Specify 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.

Property	Description
	• **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	Specify whether egress interfaces are available. • Select **Disabled on** to disable egress filtering interfaces. • Select **Enabled on** to disable egress filtering interfaces.
PCP	Specify the PCP profile to use. • In the **Profile** setting, select the PCP profile to use. • Specify either a self IP or a DS-Lite tunnel where PCP requests can be sent. • Select **Self IP**, and then select a self IP address. • Select **DSlite**, and then select a DS-Lite tunnel. ***Note:*** *DS-Lite tunnels cannot be created by BIG-IQ® Centralized Management. You must create them on the BIG-IP® device and then import them to BIG-IQ Centralized Management.*

10. Save your work.

The NAT source translations are now defined, and you can assign them to a rule used by a NAT policy.

Creating NAT destination translations

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

1. Click **Configuration** > **SECURITY** > **Network Security** > **Network Address Translation** > **NAT Destination Translations**.
2. Click **Create**.
   The NAT Destination Translations - New Item screen opens.
3. Type a name for the NAT destination translations in the **Name** field.
4. In the **Description** field, type an optional description for the NAT destination translations.
5. If needed, in the **Partition** field change the default `Common` partition.
6. 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.

7. 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.
8. 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.
9. 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 screen.

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

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 to add and schedule a new notification rule.

Creating notification rules

1. Click **Configuration** > **SECURITY** > **Network Security** > **Notification Rules**.
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.

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. 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. Save your work.

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 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. Click **Configuration** > **SECURITY** > **Network Security** > **Change Verifications**.

   The Change Verifications screen opens.
2. Click **Create**.
3. In the **Name** setting, type a name for the change verification.
4. In the **Description** setting, type a description of the change verification.
5. 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.

6. 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.
7. 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.

   Note: *When creating change verifications, you may encounter a critical error dialog box that indicates that an object, such as a logging profile, does not exist on a BIG-IP device. This critical error dialog box also contains a* ***Pin Object*** *button. Click* ***Pin Object*** *to correct the error by pinning the object to the BIG-IP device pinning policy.*

Viewing change verification properties

You view change verifications 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. Click **Configuration** > **SECURITY** > **Network Security** > **Change Verifications**.

   The Change Verifications screen opens.
2. Click the name of a change verification to view the properties, the device used, and the number of errors.

   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.
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.

Property	Description
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 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.

Managing Firewall Policies

About firewall policies

A *firewall policy* is a set of rules, or rule lists, or both. 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.

***Note:** When you are managing clustered BIG-IP devices in the BIG-IQ® Centralized Management system, avoid assigning a firewall policy to a cluster member that is a non-floating self IP. Doing so may cause unexpected results when performing partial deployments and other actions.*

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 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 Network Security system is aware of functionality implemented in one BIG-IP software 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 **Configuration** > **SECURITY** > **Network Security** > **Firewall Policies**.
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.

4. Click **Rules**, 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.

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 Viewer or Network Security Deployer cannot clone policies.

1. Click **Configuration** > **SECURITY** > **Network Security** > **Firewall Policies**.
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. ***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.

5. Click **Rules**, 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.

Reorder rules in firewall policies

Using the Firewall Policies screen, you can reorder rules in firewall policies to optimize your network firewall policies by reordering rules to change the order in which they are evaluated. Rules are evaluated from top to bottom in the list (lowest Id number first, highest Id number last).

1. Click **Configuration** > **SECURITY** > **Network Security** > **Firewall Policies**.
2. Click the name of the firewall policy to edit.
3. Click **Rules**.
4. To reorder rule lists or rules, drag and drop them until they are in the correct order.

   You can also right-click a rule row and select among the ordering options.

   Note: *You can use* ***Copy Rule*** *and then paste rules between policies. However, if you use* ***Cut Rule*** *and then paste between policies, the cut rule will not be removed from the policy.*

5. Click **Save** to save your changes.
6. 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 **Configuration** > **SECURITY** > **Network Security** > **Firewall Policies**.
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.

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. You create virtual servers using the Local Traffic service, and then use them from within Shared Security. You can modify only some of the virtual server properties within Shared Security. To modify other virtual server properties, use the Local Traffic service.

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.

Edit virtual servers

You can modify virtual servers within Shared Security to manage the security objects that might be part of a virtual server. To modify other virtual server properties, use the Local Traffic service.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Virtual Servers**.
2. On the Virtual Servers screen, click the name of the virtual server to edit.
3. On the virtual servers properties screen, modify the editable the properties as needed.

   You can change only some of the properties.
4. In the **DoS Profile** setting, select the DoS profile to use.

   You define DoS profiles using the Shared Security DoS Profiles screen.
5. In the **SSH Profile** setting, select the SSH profile to use.

   You define SSH profiles using the Shared Security SSH Profiles screen.
6. In the **IP Intelligence** setting, select the IP intelligence policy to use.

   You define IP intelligence policies using the Shared Security IP Intelligence Policies screen.
7. In the **Maximum Bandwidth** setting, type the maximum bandwidth allowed, in Mbps.

   Set this value to 0 to specify no bandwidth limit.
8. In the **Log Profiles** setting, select one or more log profiles to use.

   You define log profiles using the Shared Security Logging Profiles screen.
9. Save your work.

The virtual server is updated with your changes. You make other changes to the virtual server using this screen: **Configuration** > **LOCAL TRAFFIC** > **Virtual Servers**.

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 Shared Security, you can configure DoS 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).

DoS profile considerations when deploying to BIG-IP device clusters

In some cases, deploying a configuration containing a DoS profile from BIG-IQ® Centralized Management to a BIG-IP® device cluster can cause the cluster to become unsynchronized. If that occurs, manually synchronize the BIG-IP device cluster. Then, reimport the BIG-IP system configuration to BIG-IQ Centralized Management, and select **Use BIG-IP** system as the operation to resolve any differences.

DoS profile considerations when managing multiple BIG-IP device versions

You use BIG-IQ Centralized Management to manage multiple BIG-IP devices which can have multiple versions. In most cases, this is handled seamlessly. However, in certain cases, objects differ significantly between BIG-IP device versions, and these objects require special handling when shared between BIG-IP device versions.

- Address lists in DoS profiles

  DoS profiles that have address lists configured cannot be shared between BIG-IP devices that are version 12.1 or earlier and BIG-IP devices that are version 13.0 or later.
- Whitelists in DoS profiles

  DoS profiles that have whitelists configured cannot be shared between BIG-IP devices that are version 12.1 or earlier and BIG-IP devices that are version 13.0 or later. In the BIG-IQ Centralized Management DoS Profile, you configure whitelists differently, based on the BIG-IP device version you are managing.

  - To use a DoS profile to manage a BIG-IP device version 12.1 or earlier, select a whitelist value using the **IP Address Whitelist** setting on the DoS Profile Application Security Properties screen.
  - To use a DoS profile to manage a BIG-IP device version 13.0 or later, select a whitelist value using the **HTTP Whitelist** setting on the DoS Profile Properties screen.

  Do not select a value for both the **HTTP Whitelist** and the **IP Address Whitelist** settings in the same DoS profile.

Create 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 **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.

2. In the DoS Profiles screen, click **Create**.
3. In the New DoS Profile screen, add and set the properties as appropriate.
4. In the **Name** setting, specify a unique name for the DoS profile.
5. in the **Description** setting, specify an optional description for the DoS profile.
6. In the **Partition** setting, specify the partition to which the DoS profile belongs. You can replace the default `Common` partition when creating DoS profiles by typing a unique name for a new partition.

   The partition with that name must already exist on the BIG-IP® device. No whitespace is allowed in the partition name.
7. In the **Threshold Sensitivity** setting, specify the threshold sensitivity for the DoS profile. Thresholds for detecting attacks are higher when sensitivity is **Low** , and lower when sensitivity is **High**.

   This property is not used with the Application Security protection type.
8. In the **Source IP Address Whitelist** setting, specify the configuration of the Source IP address white list.

   This property is not used with the Application Security protection type.
9. In the **HTTP Whitelist** setting, specify the HTTP whitelist to use.

   This setting is applied only to BIG-IP devices version 13.0, or later.
10. Select a DoS protection type from the list on the left.

Option	Description
Application Security	Click **Application Security** > **Properties**, then select the **Application Security**check box, **Enabled**. When enabled, this protects your web application against DoS attacks. Your virtual server must include an HTTP profile to use this feature. Supply or modify any necessary property values.
Protocol DNS	Click **Protocol DNS**, then select the **Protocol DNS Protection** check box, **Enabled**. When enabled, this protects your DNS server against DoS attacks. Note that your virtual server must include a DNS profile to work with this feature. Supply or modify any necessary property values.
Protocol SIP	Click **Protocol SIP**, then select the **Protocol SIP Protection**check box, **Enabled**. When enabled, this protects against SIP DoS attacks. Note that your virtual server must include a SIP profile to work with this feature.
Network	Click **Network**, then select the **Network Protection** check box, **Enabled**. When enabled, this protects your server against network DoS attacks. Supply or modify any necessary property values.

11. When you are finished, save your work.

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

Configure for application security

Your virtual server must include an HTTP profile to use the application security feature.

You can configure the conditions under which the system determines that your application is under a DoS attack, and how the system reacts to a suspected attack.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.
2. In the DoS Profiles screen, click the profile name to configure.
3. On the left, click **Application Security** to expand the list.

4. Click **Properties** to display the General Settings screen and configure the application security general settings.
   a) In the **Application Security**setting, select **Enabled** to use application security protection and display additional properties.
   b) In the **IP Address Whitelist** setting, specify the IP addresses that the system considers legitimate and does not examine when performing DoS prevention.

      - To add an IP address to the whitelist, type it in the upper field, and click **Add**. The IP address is added to the whitelist in the lower field.
      - To delete an IP address from the whitelist, select the IP address from the whitelist in the lower field, and click **Remove**.

      This setting is applied only to BIG-IP devices earlier than version 13.0.
   c) In the **Geolocations** setting, specify that you want to override the DoS profile's Geolocation Detection Criteria threshold settings by selecting countries from which to allow or block traffic during a DoS attack.

      - To allow traffic from a country, select the country and move it to the Geolocation Whitelist.
      - To block traffic from a country, select the country and move it to the Geolocation Blacklist.
   d) In the **Trigger iRule** setting, enable this setting if you have an iRule that manages DoS events in a customized manner.
   e) In the **Single Page Application** setting, enable this setting if your website is a single page application.
   f) In the **URL Patterns** setting, Configure the URL patterns to be used. Each URL pattern defines a set of URLs which are logically the same URL with the varying part of the pattern acting as a parameter, such as `/product/*php`

      - To add the URL pattern to the list, type the URL pattern and click **Add**.
      - To remove the URL pattern from the list, select the pattern from the URL Patterns list, and click **Remove**.
5. To use the Proactive Bot Defense screen to configure those settings, click **Proactive Bot Defense**.

Property	Description
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. By default, the system completely blocks highly suspicious browsers and uses CAPTCHA challenges for moderately suspicious browsers. • Select the **Block Suspicious Browsers** check box to enable or disable blocking of suspicious browsers. • Select the **CAPTCHA Challenge** check box to enable or disable issuing a challenge. Click **CAPTCHA Response Settings** to select the responses to use.
Grace Period	Specifies time in seconds for the system to validate that browsers are not bots. During this period, the system does not block requests that were not validated. Modify the number or click **Reset to Default** to reset the value.
Cross-Domain Requests	You can add additional security by allowing only configured domains to reference resources of the site. From the list, select an option. You can also configure domains after selecting one of the **Cross-Domain Requests** options.
Related Site Domains	Specifies the domains that are part of the web site and protected by Proactive Bot Defense. Add domains by typing a domain in the text box and clicking **Add**. Remove a domain by selecting it and clicking **Remove**.

Property	Description
Related External Domains	Specifies the external domains (those not part of your web site) that are allowed to reference resources in your website. Add domains by typing a domain in the field and clicking **Add**. Remove a domain by selecting it in the text box and clicking **Remove**.
URL Whitelist	Specifies URLs that are not blocked by Proactive Bot Defense. Requests may still be blocked by the TPS-based / Stress-based attack mitigation. Add URLs to the whitelist by typing a URL in the text box and clicking **Add**. Remove a URL by selecting it and clicking **Remove**.

6. To use the Bot Signatures screen to configure those settings, click **Bot Signatures**.

Property	Description
Bot Signature Check	Select **Enabled** to display settings. You cannot disable the **Bot Signature Check** property while **Proactive Bot Detection**, **TPS-based Detection** with **By Device ID** selected, or **Stress-based Detection** with **By Device ID** selected, is enabled. To disable the **Bot Signature Check** property, you must first disable the previously listed properties. Alternatively, rather than disabling all bot signature checking by disabling **Bot Signature Check**, you can disable categories of bot signatures individually.
Malicious Categories and Benign Categories	These two category lists are handled similarly. For either category, select **None**, **Report**, or **Block**. That setting is then applied to all the listed items in the category. 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.
Disabled Bot Signatures	Specifies bot signatures that are available and disabled. Use the arrow buttons to move bot signatures between the **Available Signatures** list and the **Disabled Signatures** list.

7. To configure settings for the detection of DoS attacks based on a high volume of incoming traffic, click **TPS-based Detection**.

Property	Description
Operation Mode	Specifies how the system reacts when it detects an attack, and can be **Off**, **Transparent**, or **Blocking**. If set to **Off**, no other properties are shown.
Thresholds Mode	Specifies how thresholds are configured. • To configure each mitigation behavior threshold manually, select **Manual**. • To use the system default mitigation threshold settings, select **Automatic**. Your **Thresholds Mode** selection affects which threshold options are available in the other sections on this screen.
By Source IP	Specifies the criteria that determine when the system treats the IP address as an attacker, and the mitigation method to be used for the attacking IP address.
By Device ID	Specifies the criteria that determine when the system treats the device ID as an attacker, and the mitigation method to be used for the attacking device.
By Geolocation	Specifies the criteria that determine when the system treats the geolocation as an attacker, and the mitigation method to be used for the attacking geolocation. The settings exclude blacklisted and whitelisted geolocations.
By URL	Specifies the criteria that determine when the system treats the URL as an attacker, and the mitigation method to be used for the attacking URL. Heavy

Property	Description
	URL Protection can also be enabled, but needs to be configured. Click the **Click to configure** link next to the option to do so.
Site Wide	Specifies the criteria that determine when the system determines an entire website is under attack, and the mitigation method to be used.
Prevention Duration	Specifies the time spent in each mitigation step before moving (escalating or de-escalating) to the next mitigation step.

8. To configure settings for the detection of DoS attacks based on server stress, click **Stress-based Detection**.

Property	Description
Operation Mode	Specifies how the system reacts when it detects a stress-based attack, and can be **Off**, **Transparent** or **Blocking**. If set to **Off**, no other properties are shown.
Thresholds Mode	Specifies how thresholds are configured. • To configure each mitigation behavior threshold manually, select **Manual**. • To use the system default mitigation threshold settings, select **Automatic**. Your **Thresholds Mode** selection affects which threshold options are available in the other sections on this screen.
By Source IP	Specifies the criteria that determine when the system treats the IP address as an attacker, and the mitigation method to be used for the attacking IP address.
By Device ID	Specifies the criteria that determine when the system treats the device ID as an attacker, and the mitigation method to be used for the attacking device.
By Geolocation	Specifies the criteria that determine when the system treats the geolocation as an attacker, and the mitigation method to be used for the attacking geolocation. The settings exclude blacklisted and whitelisted geolocations.
By URL	Specifies the criteria that determine when the system treats the URL as an attacker, and the mitigation method to be used for the attacking URL. Heavy URL Protection can also be enabled, but needs to be configured. Click the **Click to configure** link next to the option to do so.
Site Wide	Specifies the criteria that determine when the system determines an entire website is under attack, and the mitigation method to be used.
Behavioral Detection and Mitigation	Specifies the mitigation behavior, and when enabled, the selected level of mitigation to use. • For the **Bad actors behavior detection** setting, select **Enabled** to perform traffic behavior, server capacity learning, and anomaly detection. • For the **Request signatures detection** setting, select **Enabled** to perform signature detection. Select **Use approved signatures only** to use only approved signatures. • For the **Mitigation** setting, select the type of mitigation to be used. Review the description of each mitigation type to select the best one for your environment,
Prevention Duration	Specifies the time spent in each mitigation step before moving (escalating or de-escalating) to the next mitigation step.

9. To configure settings for protecting heavy URLs during DoS attacks, click **Heavy URL Protection**.

   Heavy URLs are those which have a potential to cause stress on the server, even with a low TPS count.

Property	Description
Automatic Detection	Select **Enabled** to automatically detect heavy URLs of the application, in addition to the URLs entered manually.
Heavy URLs	You can configure a list of heavy URLs to protect in addition to the automatically detected ones. Type a URL in the top field, and click **Add**. Optionally, for a BIG-IP device version 13.0 or later, enter a threshold value. To remove a URL from the list, select the URL from the text box, and click **Remove**
Ignored URLs	You can configure a list of URLs that are excluded from automatic detection as heavy URLs. The system supports wildcards. Type a URL in the top field, and click **Add**. To remove a URL from the list, select the URL from the text box, and click **Remove**
Latency Threshold	If **Automatic Detection** is enabled, set the **Latency Threshold** setting 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 **Reset to default** to reset the value to 1000.

10. To define the responses to use when issuing a challenge, click **CAPTCHA Response Settings**.

__Note:__ The exact format of a response body differs depending on the version of the BIG-IP device. Test and verify that any custom response you create works with your installed BIG-IP version.

a) For the **First Response Type**, select **Default** to use the default response, or select **Custom** to create your own first response body by entering it into the **First Response Body** area.

Here is an example first response body:

```
This question is for testing whether you are a human visitor and to prevent automated spam
submission.
<br>
%DOSL7.captcha.image% %DOSL7.captcha.change%
<br>
<b>What code is in the image?</b>
%DOSL7.captcha.solution%
<br>
%DOSL7.captcha.submit%
<br>
<br>
Your support ID is: %DOSL7.captcha.support_id%
```

b) For the **Failure Response Type**, select **Default** to use the default response or select **Custom** to create your own failure response body by entering it into the **Failure Response Body** area.

Here is an example failure response body:

```
You have entered an invalid answer for the question. Please, try again.
<br>
%DOSL7.captcha.image% %DOSL7.captcha.change%
<br>
<b>What code is in the image?</b>
%DOSL7.captcha.solution%
<br>
%DOSL7.captcha.submit%
<br>
<br>
Your support ID is: %DOSL7.captcha.support_id%
```

11. Click **Record Traffic** to configure settings for 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.

You can record traffic and 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 this path on the BIG-IP device: `/shared/dosl7/tcpdumps`.

Property	Description
Record Traffic During Attacks	Controls whether traffic recording is used. The default is disabled and causes other properties to be hidden. Note that the system records SSL traffic encrypted. Select **Enabled** to specify that the system record traffic when a DoS attack is underway, and display settings.
Maximum TCP Dump Duration	Specifies 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	Specifies 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.

12. Save your work.

The settings are incorporated into the DoS profile.

Configure for protocol DNS security

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

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.
2. In the DoS Profiles screen, click the profile name you want to configure.
3. On the left, click **Protocol DNS Security** to display the Properties screen.
4. On the Properties screen, select the **Enabled** check box for **Protocol DNS Protection**.
5. To enable **Protocol Errors Attack Detection**, select the **Enabled** check box.
6. Specify the adjustable settings as necessary for your configuration.

   The system saves settings as you enter them.

   a) In the **Rate increased by** setting, specify that the system considers traffic to be an attack if the rate of requests increases above this number.

      By default, the system calculates this number every hour, and updates it every minute. The default is 500 percent.

   b) In the **Rate threshold** setting, specify the number of packets per second that must be exceeded to indicate to the system that there is an attack.

      The default is 250,000 packets per second.

   c) In the **Rate limit** setting, specify the limit in packets per second.

      The default is 2,500,000 packets per second.
7. At the bottom of the screen, review the Known Attack Types list that shows commonly known DNS query types that you want the system to detect in packets.
8. Enable and customize attack types individually:

   a) Click the name of the attack type to open the properties screen for it.

   b) Enable the **Detection Status** and specify the properties for the attack type detection.

   Refer to the BIG-IP® system documentation, *BIG-IP® Systems: DoS Protection and Protocol Firewall Implementations*, for information on each attack type.
9. Save your work.

Configure for protocol SIP security

Your virtual server must include a SIP profile to configure protocol SIP security in the DoS profile.

You can configure the conditions under which the system determines that your server, running SIP (Session Initiation Protocol), is under a DoS attack.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.
2. In the DoS Profiles screen, click the profile name you want to configure.
3. On the left, click **Protocol SIP Security** to display the Protocol SIP Security Properties screen.
4. On the Properties screen, select the **Enabled** check box for **Protocol SIP Protection**.

   The screen displays additional properties.
5. To enable Protocol Errors Attack Detection, select the **Enabled** check box.
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. 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.

7. At the bottom of the screen, review the Known Attack Types list that shows commonly known SIP method types that you want the system to detect in packets.
8. Enable and customize attack types individually:
   a) Click the name of the attack type to open the properties screen for it.
   b) Enable the **Detection Status** and specify the properties for the attack type detection.

   Refer to the BIG-IP® system documentation, *BIG-IP Systems: DoS Protection and Protocol Firewall Implementations* for information on each attack type.

Configure for network security

You can configure the conditions under which the system determines that your server is under a network DoS attack.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.
2. In the DoS Profiles screen, click the profile name you want to configure.
3. On the left, click **Network Security** to display the Properties screen.
4. On the Properties screen, select the check box for **Network Protection**.
   The screen displays an area for configuring dynamic signatures, and a list of commonly-known network attack types that the system can detect.
5. In the **Enforcement** setting, select the enforcement state for dynamic signatures.

   This setting is available only for BIG-IP devices version 13.0 or later.

- To enable enforcement of dynamic DoS vectors, select **Enabled**. When enforcement is enabled, all thresholds and threshold actions are applied. Enabling enforcement causes additional options to be displayed.
- To apply no action or thresholds to dynamic vectors, select **Disabled**.
- To track dynamic vector statistics, without enforcing any thresholds or limits, select **Learn-Only**.

6. In the **Mitigation Sensitivity** setting, specify the mitigation sensitivity for dynamic signatures (**None**, **Low**, **Medium**, or **High**).
7. In the **Redirection/Scrubbing** setting, specify whether to enable redirection and scrubbing of IP addresses identified by dynamic vectors.

   This enables handling of the dynamic vector hits by an IP intelligence category. Enabling redirection and scrubbing causes additional options to be displayed.
8. In the **Scrubbing Category** setting, select the IP intelligence blacklist category to which scrubbed IP addresses are sent.
9. In the **Scrubbing Advertisement Time** setting, type the duration in seconds for which an IP address is added to the blacklist category.
10. In the Known Attack Types list, enable and customize attack types individually:
    a) Click the name of the attack type to open the properties screen for it.
    b) Enable the **Detection Status** and specify the properties for the attack type detection.

    Refer to the BIG-IP® system documentation, *BIG-IP Systems: DoS Protection and Protocol Firewall Implementations* for information on each attack type.

Edit DoS profiles

You can edit DoS profiles to fine tune what the system considers to be a DoS attack, and how the system handles a DoS attack.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **DoS Profiles**.
2. In the DoS Profiles screen, click the name of the profile to modify.

   This locks the profile for editing and opens the properties screen.

   For details, consult these topics:

   - *Configure for application security*
   - *Configure for protocol DNS security*
   - *Configure for protocol SIP security*
   - *Configure for network security*
3. Make edits as needed for your configuration.
   The system saves edits as you make them.

Managing Device DoS Configurations in Shared Security

About device DoS configurations

You use the Device DoS Configurations screen to manage the device DoS configuration on the BIG-IP® devices.

To review or edit a device DoS configuration, click the name of the BIG-IP device.

Edit device DoS configurations

You can view and edit device DoS configuration properties using the Device DoS Configuration Properties screen to better protect your systems against DoS attacks.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **Device DoS Configurations**.
2. In the Device DoS Configurations screen, click the name of the device configuration to view or edit.
3. From the **Log Publisher** list, specify whether to use a log publisher, and if so, which one.
4. Below the **Log Publisher** list, there might be a threshold field, depending on the version of BIG-IP® device you are managing.
   - If you are managing a BIG-IP device version earlier than version 12.1, there is no threshold field.
   - If you are managing a BIG-IP device version 12.1.x, you can use the **Auto Threshold Sensitivity** field to select a sensitivity value between 1 - 100.
   - If you are managing a BIG-IP device version 13.0.x or later, you can use the **Threshold Sensitivity** field to select the sensitivity.
5. In the **Enforcement** setting, specify the enforcement state for dynamic signatures.

   This setting is only available for BIG-IP devices version 13.0 or later.
   - To enable enforcement of dynamic DoS vectors, select **Enabled**. When enforcement is enabled, all thresholds and threshold actions are applied. Enabling enforcement causes additional options to be displayed.
   - To apply no action or thresholds to dynamic vectors, select **Disabled**.
   - To track dynamic vector statistics, without enforcing any thresholds or limits, select **Learn-Only**.
6. In the **Mitigation Sensitivity** setting, specify the mitigation sensitivity for dynamic signatures.
7. In the **Redirection/Scrubbing** setting, specify whether to enable redirection and scrubbing of IP addresses identified by dynamic vectors.

   This enables handling of the dynamic vector hits by an IP intelligence category. Enabling redirection and scrubbing causes additional options to be displayed.
8. In the **Scrubbing Category** setting, select the IP intelligence blacklist category to which scrubbed IP addresses are sent.
9. In the **Scrubbing Advertisement Time** setting, type the duration in seconds for which an IP address is added to the blacklist category.
10. In the Category area, click the triangle to the left of a category to expand the category, and view or modify attack types within the category.
11. In the Attack Type list, click the name of an attack type to modify its properties.

    Some properties are read-only.

12. When you are finished modifying an attack type, click **OK** to save your changes to that attack type.

13. When you are finished modifying all attack types for the BIG-IP device, save your changes.

Managing Bot Signatures and Bot Signature Categories

About bot signatures and bot signature categories

You use bot signatures to identify web robots by looking for specific patterns in the headers of incoming HTTP requests. You can create, modify, and delete only those bot signatures that are user-defined.

Bot signatures are organized by categories. You can assign a bot signature to an existing category, or create your own.

Create bot signatures

You use bot signatures to identify web robots by looking for specific patterns in the headers of incoming HTTP requests. Refer to the BIG-IP ® Application Security Manager™ (ASM) documentation on attack and bot signatures for more information.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Bot Signatures**.
2. Click **Create**.
3. In the **Name** field, type a name for the bot signature.
4. In the **Partition** setting, the `Common` partition is listed and cannot be changed.
5. In the **Domains** setting, you can add or delete domains.
   - To add a domain, in the **Domain Name** field, type the name and click **Add**.
   - To delete a domain, select a domain from the list and click **Remove**.
6. From the **Category** list, select the appropriate category for the bot signature.
7. In the **Rule** setting, create a rule for the bot signature using either simple or advanced editing.
   - Select **Simple Edit Mode** to create a rule by supplying what content the user agent and the URL should match.
     - From the **User-agent** list, select the type of match, and then type the string to be matched in the user agent.
     - From the **URL** list, select the type of match, and then type the string to be matched in the URL.
   - Select the **Advanced Edit Mode** to create more complex rules, such as those containing multiple search strings or a conditional text match. You type the rule expression using `Snort` control syntax. `Snort` control syntax is explained fully in the BIG-IP Application Security Manager documentation.
8. From the **Risk** list, select the risk level associated with the bot signature.
9. You see that **User-defined** is selected for any new or modified bot signature defined by the user: this cannot be changed.
10. You see that the **References** setting is read-only and set to `N/A`.
11. Save any changes.

Create bot signature categories

You use bot signature categories to label groups of bot signatures. You can create and modify only those bot signature categories that are user-defined.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Bot Signature Categories**.
2. Click **Create**.
   The New Bot Signature Category screen opens.
3. Type a **Category Name** for the bot signature category, and use the **Partition** setting default of `Common`.
4. From the **Category Type** list, select the appropriate type for the bot signature category, either **Malicious** or **Benign**.
5. Since you can only create user-defined bot signature categories, the **User-defined** setting is selected and cannot be changed.
6. Save any changes.

Managing Network Whitelists in Shared Security

About network whitelists

You use network whitelists to define network addresses that are allowed to bypass the checks in a DoS profile. The Network White Lists screen displays the managed BIG-IP® devices that might have network whitelists defined. Click the name of a BIG-IP device to display the network whitelists that are defined. A maximum of 8 network white lists are allowed for each BIG-IP device

Create network whitelist

You create network whitelists to bypass checks in a DoS profile.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **DoS Protection** > **Network White Lists**.
2. Click the name of the BIG-IP ®device on which to create the network white list.
3. In the **Source Address List** setting, select the IP address from which the packet is coming.
4. Click **Create** to add a network white list.
5. Type a **Name** for the network white list, and an optional **Description** that will be useful in your environment.
6. In the **Protocol** setting, leave the default value, **Any**, or select the appropriate network protocol.
7. For the **Address Type** setting, specify the type of addresses being handled: **Source** or **Destination**.

   The properties available change based on your choice.
8. In the **Address** setting, leave the default value, **Any**, or select **Specify** and provide the address in the provided field.

   You can specify IPv4 or IPv6 addresses in CIDR notation as the address. You can specify a source address or destination address, but not both in the same white list entry.
9. If you selected a source address type, in the **VLAN** setting, leave the default value, **Any**, select the appropriate VLAN, or select **Other** and provide a VLAN tag number.
10. If you selected a destination address type, in the **Port** setting, leave the default value, **Any**, or select the appropriate port.

    The system provides the default port number value for each port type when the **Protocol** is set to **TCP** or **UDP**.
11. When you are finished, click **OK**.
12. Save your changes.

Managing IP Intelligence Settings

Overview of IP intelligence settings

In a network firewall, you can configure IP intelligence policies to check traffic against an IP intelligence database. Such traffic can be handled automatically if it originates from known-bad or questionable IP addresses.

You can dynamically adjust the blacklists and whitelists used in the policy by creating feed lists. A *feed list* retrieves blacklists and whitelists from specified URLs. You can also set up blacklist matching criteria within the IP intelligence policy, and you may create additional blacklist categories to use in the matching criteria.

You can use global IP intelligence policies to select options that will be used for all your IP intelligence policies.

BIG-IQ® Centralized Management supports the IP Intelligence feature in BIG-IP® versions 12.0 or later.

Create blacklist categories

You create blacklist categories to use when matching blacklists in an IP intelligence policy when existing categories are insufficient. The blacklist category groups related untrustworthy IP addresses.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **IP Intelligence** > **Blacklist Categories**.
2. On the Blacklist Categories screen, click **Create**.
3. In the **Category Name** field, type the name of the category.

   You cannot change this when modifying a category.
4. In the **Description** field, type a description of the category.
5. In the **Match Type** setting, specify the criteria that defines a blacklist match.

   You can require a source match, a destination match, or both a source and destination match.

   - Select **Both Source and Destination** to require that both the source and the destination match the blacklist.
   - Select **Destination** to have the destination only match the blacklist.
   - Select **Source** to have the source only match the blacklist.
6. Save your work.

You can now use this blacklist category in an IP intelligence policy.

Create feed lists

You create feed lists containing URLs to dynamically adjust the blacklists and whitelists in an IP intelligence policy to allow more automatic handling of those lists.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **IP Intelligence** > **Feed Lists**.
2. On the Feed Lists screen, click **Create**.
3. In the **Name** field, type a unique name for the feed list.
4. In the **Description** field, type an optional description for the feed list.
5. In the **Partition** setting, the default is `Common`. Type a different partition if needed.

6. In the Feed URLs area, click **Create** to create a feed URL and add it to the feed list.

   The Feed URL properties screen opens. You may want to add multiple feed URLs to the feed list.
7. In the **Name** field, type a name for the feed URL.
8. In the **URL** field, type the URL for the feed.
9. For the **List Type** setting, select the list type to specify whether the list is by default a whitelist or blacklist. This applies only to items on the list that are not specified as blacklist or whitelist items.
10. For the **Blacklist Category** setting, select a default category for the list.
11. In the **Poll Interval** field, type a number that specifies how often the feed URL is polled for new feeds, in seconds.

    The default value is 300, which is the minimum.
12. In the **Username** field, type a user name used to access the feed list file, if required.
13. In the **Password** field, type a password used to access the feed list file, if required.

    Note: *In some cases, the value of the Password setting may be falsely displayed as changed when performing an evaluation prior to a deployment. This is due to encryption salt changes, and you can ignore it.*
14. If the **Password** setting is used, in the **Confirm Password** field, type the password again to confirm it.
15. Click **OK** to save the changes to the feed URL.
16. Continue to add or change the feed URLs in the feed list until it is complete.
17. Save your work.

You can now create and add more feed URLs to the feed list or add the feed list to an IP intelligence policy.

Create IP intelligence policies

You create an IP intelligence policy to check traffic against an IP intelligence database and determine whether to allow it.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **IP Intelligence** > **IP Intelligence Policies**.
2. In the IP Intelligence Policies screen, click **Create**.

   The IP Intelligence Policy Properties screen opens.
3. In the **Name** setting, type a unique name for the policy.
4. In the **Description** setting, type an optional description.
5. The **Partition** setting shows the default, `Common`, but you can type a different partition if needed.
6. In the **Feed Lists** setting, specify the feed lists to be used in the policy.
7. For the **Default Action** setting, specify the default action that the policy takes on identified blacklist items (for which no action is specified).
8. In the **Default Log Actions** setting, specify what actions to log by default.
   a) In the **Log Whitelist Overrides** setting, select whether to log whitelist overrides.
   b) In the **Log Blacklist Category Matches** setting, select whether to log blacklist category matches.
9. Click **Save** to save your work before creating a black list matching policy.
10. In the Blacklist Matching Policies area, click **Create** to create a new blacklist matching policy for the IP intelligence policy.
    The blacklist matching policy properties screen opens, which has the same name as the IP intelligence policy.

11. For the **Blacklist Categories** setting, select the category for which you are configuring settings in this policy.
12. For the **Action** setting, select the action for this policy.
    - Select **Use Policy Default** to use the default action for this policy.
    - Select **Drop** for the policy to use the drop action.
    - Select **Accept** for the policy to use the accept action.
13. For the **Log Blacklist Category Matches** setting, select the log action for this policy.
    - Select **Use Policy Default** to use the default log action for logging blacklist category matches.
    - Select **Yes** to override the default action and enable logging of blacklist category matches.
    - Select **No** to override the default log action, and disable logging of blacklist category matches.
    - Select **Limited** to override the default action and enable limited logging of blacklist category matches.
14. For the **Log Whitelist Overrides** setting, select **Use Policy Default** to use the default log action for whitelist overrides. Select **Yes** or **No** to override the default action.
    - Select **Use Policy Default** to use the default log action for logging whitelist overrides.
    - Select **Yes** to override the default action and enable logging of whitelist overrides.
    - Select **No** to override the default log action, and disable logging of whitelist overrides.
15. For the **Match Override** setting, specify the matching criteria that overrides a blacklist match.

    You can require a source match, a destination match, or both a source and destination match to override a blacklist match with a whitelist (**Match Source and Destination**, **Match Source**, or **Match Destination**).
16. Click **OK** to save your work on the blacklist matching policy

    The screen closes and the blacklist matching policy you created is listed on the IP intelligence policy screen.
17. Save your work on the IP intelligence policy.

Configure the global IP intelligence policy

You can configure an IP Intelligence policy to be used globally to apply blacklist and whitelist matching actions and logging to all traffic on the BIG-IP device.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **IP Intelligence** > **Global Policies**.
2. Click the name of the BIG-IP device on which to use the global IP intelligence policy.
3. In the **Description** field, type a description for the global IP intelligence policy.
4. In the **IP Intelligence Policy** setting, select the policy to use as the global IP intelligence policy.

   The default policy is `Common/ip-intelligence`.
5. Save your work.

Managing External Redirection Settings

Overview of external redirection settings

You use scrubber profiles, blacklist publishers, and blacklist publisher profiles to protect your network by detecting and redirecting DoS and DDoS attacks.

You use scrubber profiles to configure network traffic scrubbing and redirection for your environment, including enabling F5® Silverline® DDoS protection. You use blacklist publisher profiles and blacklist publishers to advertise blacklists to routers in your network.

Create blacklist publishers

You create blacklist publishers to advertise blacklists to routers in your network.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **External Redirection** > **Blacklist Publishers**.
2. On the Blacklist Publishers screen, click **Create**.
   The New Blacklist Publisher screen opens.
3. For the **Blacklist Category** setting, specify the blacklist category to use.
4. For the **Blacklist Publisher Profile** setting, select a black list publisher profile to use, if one is defined.

   Using the profile is optional. You can create blacklist publishers without using the profile.
5. Save your work.

Create blacklist publisher profiles

You create a blacklist publisher profile to use with your blacklist publisher to advertise blacklists to routers in your network.

***Note:** You cannot delete an unused blacklist publisher profile from a BIG-IP® device version 13.0 or earlier during deployment, even though the deployment difference shows it will be deleted. Deploying the configuration again causes the blacklist publisher profile to be deleted.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **External Redirection** > **Blacklist Publisher Profiles**.
2. On the Blacklist Publisher Profiles screen, click **Create**.
   The New Blacklist Publisher Profile screen opens.
3. In the **Name** field, type the name of the profile.
4. In the **Description** field, type a description for the profile.
5. For the **Route Domain** setting, specify the route domain on which blacklisted addresses are advertised.
6. In the **Advertisement Next-Hop** field, type an IP address for the next hop IP address of the BGP (Border Gateway Protocol) router to which you want to advertise blacklisted addresses.
7. For the **Traffic Group** setting, select the traffic group on which you want to advertise blacklisted addresses.
8. Save your work.

Edit the scrubber profile

You modify the scrubber profile to configure network traffic scrubbing, including enabling F5® Silverline® DDoS protection, if needed.

***Note:** Before deploying a change to the scrubber configuration, such as changing the route domain used by the scrubber, you should make sure the scrubber is inactive on the BIG-IP device. Deploying a changed configuration while the scrubber is active on the BIG-IP device can cause the following error:* `Deployment failed, with error: Cannot configure scrubber property when scrubber is active. Stop active scrubbering on scrubberName to make configuration changes.`

1. Click **Configuration** > **SECURITY** > **Shared Security** > **External Redirection** > **Scrubber Profiles**.
2. On the Scrubber Profiles screen, click the device name for the scrubber profile to modify.

   Each BIG-IP® device has only one scrubber profile.
3. On the left, click **Properties** and modify the settings as needed.
   a) For the **Advertisement TTL** setting, specify the amount of time, in seconds, that scrubbed IP addresses are advertised to the BGP router or to Silverline DDoS protection.
      - To allow an infinite amount of time, select **Infinite**.
      - To allow a specific amount of time, select the other option and type the number of seconds to advertise.

   b) For the **Silverline** setting, select **Enabled** to use Silverline DDoS protection to offload scrubbed IP addresses, and to display the Silverline configuration properties.
   c) In the **URL** field, type the URL of the Silverline DDoS account.
   d) In the **User** field, type the user name for the Silverline DDoS account.
   e) In the **Password** field, type the password for the Silverline DDoS account.

      ***Note:** In some cases, the value of the **Password** setting might be falsely displayed as changed when performing an evaluation prior to a deployment. This is due to encryption salt changes, and you can ignore it.*

   f) In the **Confirm Password** field, type the password for the Silverline DDoS account again to confirm it.
4. To create new or edit route domain scrubber definitions, click **Route Domains**.
   - To create a new route domain scrubber definition, click **Create**. Then edit the definition to add details, such as the route domain.
   - To edit a route domain scrubber definition, click the pencil icon in the definition row.
   - To delete a route domain scrubber definition, right click in the definition row and select **Delete Row**.
5. When creating or editing a route domain scrubber definition, specify the route domain scrubber definition settings.
   a) In the Name column, type the optional name of the route domain definition.
   b) In the Route Domain column, select the route domain to use.
   c) In the Scrubbing Threshold column, in the top field, select the type of value: **Absolute** or **Percentage**.
   d) In the Scrubbing Threshold column, in the bottom field, specify that the value is **Infinite**, or select **Specify** and type a numeric value in Mbps in the provided field.

e) In the Advertisement Method column, specify the method for this route domain: **BGP**, **Silverline**, or **None**.

f) In the Scrubber Details column, use the **Type** setting to specify how to advertise. Your selection determines what other settings are available.

- To advertise all scrubbed IP addresses to a BGP router, select **Advertise All**. The **IPv4** and **IPv6** settings are displayed. Type the IP address of the BGP router in the appropriate field for the IP address.
- To advertise specific prefixes to a BGP router or to Silverline, select **Prefix Specific Advertisement**. The **IP Address** and **BGP Scrubber Destination** settings are displayed.

  1. In the **IP Address** field, type the IP address and prefix to be scrubbed, in CIDR notation.
  2. In the **BGP Scrubber Destination** field, type the IP address of the scrubber. This field is only used when the Advertisement Method is set to **BGP**.
  3. Click **Add** to add the entry to the list.

***Note:** Scrubber profiles imported from a BIG-IP device might contain the following as IP address values: `any`, `any6`, `0.0.0.0`, or `::` in the route domain scrubber details when **Prefix Specific Advertisement** is selected. These values are not supported on the BIG-IQ® Centralized Management system and will cause differences when importing or deploying configurations. You can remove these differences by changing these values to values that BIG-IQ Centralized Management supports. For example, you can replace `any` and `any6` on the BIG-IP device with a blank value on the BIG-IQ Centralized Management system, since all indicate that any IP address is valid for that field.*

6. To create or edit virtual server scrubber definitions, click **Virtual Servers**.

   - To create a new virtual server scrubber definition, click **Create**. Then edit the definition to add details, such as the virtual server.
   - To edit a virtual server scrubber definition, click the pencil icon in the definition row.
   - To delete a virtual server scrubber definition, right click in the definition row and select **Delete Row**.

7. Specify the virtual server scrubber definition settings.

   a) In the Name column, type the optional name of the virtual server definition.
   b) In the Virtual Server column, select the virtual server to use.
   c) In the Scrubbing Threshold column, in the top list, select the type of value: **Absolute** or **Percentage**.
   d) In the Scrubbing Threshold column, in the bottom field, specify that the value is **Infinite**, or select **Specify** and type a numeric value in Mbps in the provided field.
   e) In the Advertisement Method column, select the method for this virtual server.
   f) In the Scrubber Details column, type the IP address of the scrubber. This value is only used when the Advertisement Method is set to **BGP**.

8. To create or edit blacklist category scrubber definitions, click **Categories**.

   - To create a new blacklist category scrubber definition, click **Create**. Then edit the definition to add details, such as the advertisement method.
   - To edit a blacklist category scrubber definition, click the pencil icon in the definition row.
   - To delete a blacklist category scrubber definition, right click in the definition row and select **Delete Row**.

9. When creating or editing a blacklist category scrubber definition, specify the blacklist category scrubber definition settings.

   a) In the Name column, type the optional name of the blacklist category scrubber definition.
   b) In the Blacklist Category column, select the category to use. In most cases, you will want to select `attacked_ips`. This is a category created for IP addresses that are under attack.

  c) In the Route Domain column, select the route domain to use.
  d) In the Advertisement Method column, select the method for this blacklist category scrubber definition.
  e) In the Scrubber Details column, if you selected BGP as the advertisement method, type the destination IP address in the **IPv4** or **IPv6** setting, whichever is appropriate. If you selected another advertisement method, you do not supply any scrubber details.

10. Save your work.

Managing Logging Profiles in Shared Security

About logging profiles

The Logging Profiles screen lists both default logging profiles that cannot be modified, and other logging profiles that can be modified. The default logging profiles are: `Log all requests`, `Log illegal requests`, `global-network`, and `local-dos`. Default logging profiles are imported from the BIG-IP® device, and only top-level information about them can be viewed on the Logging Profiles screen, such as the logging profile name, description, partition, and devices.

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® Centralized Management 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 system and deployed back to a BIG-IP device. However, BIG-IQ does not support creating logging profiles with spaces in the name.

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

Create logging profiles

You create logging profiles to configure the kind of information to log for objects that support logging.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click **Create**.
   The New Logging Profile screen opens with the Properties displayed.
3. In the **Name** field, type a name for the logging profile.
4. In the **Description** field, type an optional description for the logging profile.
5. 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.
6. On the left, click the logging type that you want to use, and then select the **Enabled** check box to display the related settings.

   - Enable **Application Security** to specify that the system logs traffic to the web application. You cannot enable both **Application Security** and **Protocol Security**. Refer to the *Configure for Application Security logging* section of *BIG-IQ Centralized Management: Security* on `support.f5.com` for configuration information.

- Enable **Protocol Security** to specify that the system logs any dropped, malformed, and/or rejected requests sent through the given protocol. Refer to the *Configure for Protocol Security logging* section of *BIG-IQ Centralized Management: Security* on `support.f5.com` for configuration information.
- Enable **Network Firewall** to specify that the system logs ACL rule matches, TCP events, and/or TCP/IP errors sent to the network firewall. Refer to the *Configure for Network Firewall logging* section of *BIG-IQ Centralized Management: Security* on `support.f5.com` for configuration information.
- Enable **Network Address Translation** to specify which Network Address Translation (NAT) events the system logs, and where those events are logged. Refer to the *Configure for Network Address Translation logging* section of *BIG-IQ Centralized Management: Security* on `support.f5.com` for configuration information.
- Enable **DoS Protection** to specify that the system logs detected DoS attacks, and where DoS events are logged.
- Enable **Bot Defense** to specify that the system logs bot defence events. Refer to the *Configure for Bot Defense logging* section of *BIG-IQ Centralized Management: Security* on `support.f5.com` for configuration information.

You must configure each enabled logging type before you can use it. You can do that now, or save the profile and configure the logging types later.

7. Specify the settings needed for each logging type you use.

   You can configure multiple logging types while editing the logging profile.
8. When finished, save your changes.

Configure for Application Security logging

You need to configure application security logging profile settings after you have enabled them to specify what information is logged.

***Note:** You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. Click the name of the logging profile to configure on the Logging Profiles screen.
   The `logging-profile-name` screen opens with the Properties displayed.
3. On the left, click **Application Security**.
   The Application Security configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the Application Security configuration settings.
5. Supply the Application Security Configuration settings.

Property	When enabled:
Local Storage	When enabled, specifies that the system stores all traffic in the system. This setting can only be disabled when **Remote Storage** is enabled.
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. This setting is displayed only when **Local Storage** is enabled.
Response Logging	Specifies whether, and how, the system logs HTTP responses. • **Off:** The system does not log responses. This is the default.

Property	When enabled:
	• **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**.
Guarantee Local Response Logging	Specifies that the system logs all responses, even though this may slow your web application. When cleared (disabled), specifies that the system logs responses as long as it does not slow your web application. The default is disabled. In either case, the system does not drop responses. This setting is displayed only when **Guarantee Local Logging** is enabled, and **Response Logging** is set to **For Illegal Requests Only** or **For All Requests**.
Remote Storage	When enabled, specifies that the system stores all traffic on a remote logging server. This setting can only be disabled when **Local Storage** is enabled. Also 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. • 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 Centralized Management 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 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.

Property	When enabled:
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. Select an appropriate value. The value you can select is determined by what protocol is selected. When logging Web Application Security traffic, the **Maximum Entry Length** setting should be set to **64K**.
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.

6. Supply the Application Security settings for the Storage Filter area.

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. • **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

Property	When enabled:
	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**). • Type the string to search for in the request in the field to the right. The search is case-sensitive.
Login Result	Specifies whether request logging occurs for all login results or only for selected login results. • **All:** Specifies that the system, or server, logs all login results. This is the default. • **Only:** Specifies that the system, or server, logs login results of the specified type. When selected, displays options where you specify the login results to log. This option is only valid with BIG-IP devices version 13.0 or later.

7. When you are 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.

***Note:** You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

***Note:** You cannot enable **Protocol Security** if you have already enabled **Application Security**.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click the name of the logging profile to configure.
   The `logging-profile-name` screen opens with the Properties displayed.
3. On the left, click **Protocol Security**.
   The Protocol Security configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the Protocol Security configuration settings.
5. In the HTTP, FTP, and SMTP Security area, in the **Publisher** setting, select the log publisher to use for the HTTP, FTP and SMTP protocols , or accept the default of **None**.

This value specifies where the system sends log messages.

6. In the DNS Security area, 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 (,). *Note: Do not use the $ character: it is 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.

7. In the SIP Security area, 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	Specifies that the system logs redirection responses.

Property	When enabled:
Log Request Failures	Specifies that the system logs request failures.
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:*** *Do not use the $ character; it is 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. In the SSH Proxy area, supply the Protocol Security SSH Proxy settings to configure logging of SSH proxy use. Select **Enabled** to make the other settings available.

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.
Log Client Auth Partial Event	Logs client side partial authentication events.
Log Server Auth Partial Event	Logs server side partial authentication events.

9. When you are 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.

***Note:** You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click the name of the logging profile to configure.
   The `logging-profile-name` screen opens with the Properties displayed.
3. On the left, click **Network Firewall**.
   The Network Firewall configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the Network Firewall properties.
5. In the Properties area, 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 **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: • **Indefinite** sets the rate limit to the maximum of 4294967295, and **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** 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** 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 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** of **Indefinite**, which

Property	When enabled:
	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:*** *Do not use the $ character; it is 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.

6. In the IP Intelligence area, 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 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.
Log Shun Events	Specifies that IP Intelligence shun list events are logged.
Log RTBH Events	Specifies that remotely triggered black holing (RTBH) events are logged.
Log Scrubber Events	Specifies that IP Intelligence scrubber events are logged.

7. In the Traffic Statistics area, 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.

8. In the Port Misuse area, 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 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.

9. When you are 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.

Note: *You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click the name of the logging profile to configure.
   The `logging-profile-name` screen opens with the Properties displayed.
3. On the left, click **Network Address Translation**.
   The Network Address Translation configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the Network Address Translation properties.
5. Supply the Network Address Translation settings to configure which NAT events the system logs, and where they are logged.

Property	When enabled:
LSN Legacy Mode	When enabled, 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.

Property — **When enabled:**

Aggregate Rate Limit

Specifies, when enabled, a rate limit for all combined NAT firewall log messages per second. Above this rate limit, log messages are not logged.

- To enable a limit, select **Specify** and provide a numeric value for the number of messages per second.
- To have no limit, select **Indefinite**.

Start Outbound Session

Specifies logging options for the start of an outbound translation session, when the outbound flow is created.

Select one of the following from the list.

- Select **Enabled** to log Start Outbound Session events.
- Select **Disabled** to not log Start Outbound Session events. This is the default.
- 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.
- Select **Include Destination Address/Port** to include the destination address/port.
- In the **Rate Limit** setting, specify a rate limit for these events.
    - To enable a limit, select **Specify** and provide a numeric value for the number of messages per second.
    - To have no limit, select **Indefinite**.

End Outbound Session

Specifies logging options for the end of an outbound translation session, when the outbound flow is deleted.

Select one of the options from the list.

- Select **Enabled** to log End Outbound Session events.
- Select **Disabled** to not log End Outbound Session events. This is the default.
- 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.
- Select **Include Destination Address/Port** to include the destination address/port.
- In the **Rate Limit** setting, specify a rate limit for these events.
    - To enable a limit, select **Specify** and provide a numeric value for the number of messages per second.
    - To have no limit, select **Indefinite**.

Start Inbound Session

Specifies logging options for the start of an incoming connection to a translated address.

Select one of the options from the list.

- Select **Enabled** to log Start Inbound Session events.
- Select **Disabled** to not log Start Inbound Session events. This is the default.
- 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.
- In the **Rate Limit** setting, specify a rate limit for these events.
    - To enable a limit, select **Specify** and provide a numeric value for the number of messages per second.
    - To have no limit, select **Indefinite**.

Property	When enabled:
End Inbound Session	Specifies logging options for the end of an incoming connection to a translated address. Select one of the options from the list. • Select **Enabled** to log End Inbound Session events. • Select **Disabled** to not log End Inbound Session events. This is the default. • 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. • In the **Rate Limit** setting, specify a rate limit for these events. • To enable a limit, select **Specify** and provide a numeric value for the number of messages per second. • To have no limit, select **Indefinite**.
Quota Exceeded	When enabled, specifies logging when a client exceeds the allocated resource limit. In the **Rate Limit** setting, specify a rate limit for these events. • To enable a limit, select **Specify** and provide a numeric value for the number of messages per second. • To have no limit, select **Indefinite**.
Errors	When enabled, specifies logging when errors are encountered while attempting translation for clients. In the **Rate Limit** setting, specify a rate limit for these events. • To enable a limit, select **Specify** and provide a numeric value for the number of messages per second. • To have no limit, select **Indefinite**.
Publisher	Specifies the name of the log publisher used for logging NAT events. Select a log publisher configured in your system.

6. When you are 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.

***Note:** You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

---

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click the name of the logging profile to configure.
   The `logging-profile-name` screen opens with the Properties displayed.
3. On the left, click **DoS Protection**.
   The DoS Protection configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the DoS Protection properties.
5. 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 **Remote Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.

6. In the DNS DoS Protection area, configure where DNS DoS protection events are logged: Select a **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
7. For the SIP DoS Protection area, configure where SIP DoS protection events are logged: Select a **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
8. For the Network DoS Protection area, configure where Network DoS protection events are logged: Select a **Publisher** to specify the name of the log publisher used for logging events. Select a log publisher configured in your system.
9. When you are finished, save your changes.

The DoS Protection configuration settings are saved.

Configure for Bot Defense logging

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

***Note:** You can do this configuration when initially creating a logging profile (in which case go directly to step 5), or perform the configuration later in a logging profile that already exists.*

1. Click **Configuration** > **SECURITY** > **Shared Security** > **Logging Profiles**.
2. On the Logging Profiles screen, click the name of the logging profile to configure.
   The *`logging-profile-name`* screen opens with the Properties displayed.
3. On the left, click **Bot Defense**.
   The Bot Defense configuration screen opens.
4. For **Status**, select the **Enabled** check box.
   The screen displays the Bot Defense request logging properties.
5. In the Request Log area, select the request logging options to use.
   - To use a local publisher, in the **Local Publisher** setting, select **Enabled**.
   - To use a remote publisher, in the **Remote Publisher** setting, select the remote publisher to use. Select **None** to not use a remote publisher.
   - To log illegal requests, in the **Log Illegal Requests** setting, select **Enabled**.
   - To log Captcha challenged requests, in the **Log Captcha Challenged Requests** setting, select **Enabled**.
   - To log challenged requests, in the **Log Challenged Requests** setting, select **Enabled**.
   - To log bot signature matched requests, in the **Log Bot Signature Matched Requests** setting, select **Enabled**.
   - To log legal requests, in the **Log Legal Requests** setting, select **Enabled**.
6. When you are finished, save your changes.

The Bot Defense 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 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 *channel action* is an action on a channel, A single SSH connection may contain multiple channels and actions, such as `Shell`, `SCP Up`, and others. The channel actions you can use in rules are shown in columns in the user interface.

1. Click **Configuration** > **SECURITY** > **Shared Security** > **SSH Profiles**.
2. Click the name of the SSH profile for which you want to configure permissions.
3. On the left, click **SSH Proxy Permissions**, and then click the **Create Rule** button.

   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.

4. Click the pencil icon next to the name of the rule to edit the default rule properties.
5. In the **Name** field, type a more meaningful name for the rule.
6. 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 **Add** 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.
7. 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.
   - **Unspecified** indicates that the DEFAULT ACTIONS rule value be used for the rule. The DEFAULT ACTIONS rule is shown at the bottom of the rule list.
8. To enable logging for any action, select the **Log** check box below the SSH channel action.
9. Review your settings, and click **Save**.

The SSH proxy permissions are defined for the SSH profile.

If they are 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 Web Application Security

About application security policies in Web Application Security

Web Application Security imports BIG-IP® Application Security Manager™ (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 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® Centralized Management user interface.

Following is a list of the supported versions of the BIG-IQ Centralized Management 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 Centralized Management system support, such as the minimum F5® 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 Centralized Management 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	Learning using the Central Policy Builder is not applicable for use with response pages.
Signatures	Yes	Yes	Any	

Subcollection	Discovery and Deployment Support	Edit Support using BIG-IQ GUI	Minimum BIG-IP Device Version Support	Comments
Signature Sets and attack signature configuration	Yes	Yes	Any	Filter-based sets are supported, manual sets are not.
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	Yes	Any	Learning using the Central Policy Builder is not applicable for use with the local Policy Builder.
Central Policy Builder	Yes	Yes	13.1.0	Learning using the Central Policy Builder is not supported with GWT content profiles, and is non-applicable for response pages, local policy builder, web scraping, and brute force attack prevention.
Allowed methods	Yes	Yes	Any	
Headers	Yes	Yes	Any	
Cookies	Yes	Yes	Any	
Host names	Yes	Yes	Any	
Geolocation enforcement	Yes	Yes	11.6.0	
IP Intelligence	Yes	Yes	11.6.0	
Redirection protection	Yes	Yes	11.6.0	
Sensitive parameters	Yes	Yes	Any	
Web scraping	Yes	No	12.0.0	Learning using the Central Policy Builder is not applicable for use with web scraping.
CSRF protection	Yes	Yes	11.6.0	When editing policies deployed to BIG-IP device versions earlier than 13.1, URLs added

Subcollection	Discovery and Deployment Support	Edit Support using BIG-IQ GUI	Minimum BIG-IP Device Version Support	Comments
				to the CSRF URLs list must have the following settings: • **Method** setting of **Any** • **Enforcement Action** setting of **Verify CSRF Token** • **Required Parameters** setting of **At Least One**
JSON Content Profiles	Yes	Yes	11.6.0	
XML Content Profiles	Yes	Yes	11.6.0	Schemas and WSS are not supported.
GWT Content Profiles	Yes	No	11.6.0	Learning using the Central Policy Builder is not supported with GWT content profiles.
Plain Text Content Profiles	Yes	Yes	12.1.0	
URLs	Yes	Yes	Any	Flow configuration for URLs is not supported, such as referrer, check flows, check pd/qa, allow pd/qs, or isEntryPoint.
Websocket URLs	Yes	Yes	12.1.0	
Login Pages	Yes	Yes	11.6.0	
Login Enforcement	Yes	Yes	11.6.0	
Brute Force Attack Preventions	Yes	Yes	11.6.0	Learning using the Central Policy Builder is not applicable for use with brute force attack prevention.
Session Tracking Configuration	Yes	Yes	11.6.0	Only configuration is supported, there is no support for online tracking data.
Layered Policy	Yes	Yes	13.0.0	
Inheritance Settings	Yes	Yes	13.0.0	
Enforcement Readiness	Yes	Yes	Any	
Server Technologies	Yes	Yes	13.1.0	

Overview of layered policies and inheritance

You can use Web Application Security to create and manage two layers of security policies: parent policies and child policies. This feature is new with BIG-IP® version 13.0 and BIG-IQ® Centralized Management version 5.2. Parent policies include mandatory policy elements, and child policies inherit those attributes from the parent. When the parent policy is updated, the associated child policies are automatically updated.

With parent policies you can:

- Create and maintain common elements and settings.
- Impose mandatory elements on child policies.
- Push a change to multiple child policies.

You can specify which parts of the security policy must be inherited, which are optional, and which are not inherited. This allows you to keep child policies synchronized with the changes in the global mandatory policies and still allow the child policies to address their own unique requirements.

You establish the parent and child policy relationship as follows:

1. Identify the current policy as a parent policy.

   On the General Properties screen for the policy, set the **Policy Type** to **Parent Policy**. Navigate to **Configuration** > **SECURITY** > **Web Application Security** > **Policies**, then click the policy to edit, and click **POLICY PROPERTIES** > **General Properties**
2. Set a policy to be the child policy of the parent policy.

   On the Inheritance Settings screen for the policy, select the parent policy for a child policy by selecting the parent policy name in the **Parent Policy** setting. Navigate to **Configuration** > **SECURITY** > **Web Application Security** > **Policies**, then click the policy to become a child policy and click **POLICY PROPERTIES** > **Inheritance Settings**.
3. Click **Save** to save this policy as a child policy and display the inheritance properties.
4. Continue to use the Inheritance Settings screen to accept or decline what is to be inherited from the parent policy.

By default, the **Parent Policy** field is set to **None**, and there is no layered policy use (no child or parent policies).

Refer to the *BIG-IP Application Security Manager: Getting Started* guide for additional information on using parent and child layered policies.

Overview of central policy building

You can use the Central Policy Builder feature to receive policy building learning suggestions from multiple BIG-IP® devices, rather than have each BIG-IP device perform policy learning in isolation. Using the Central Policy Builder, the learning suggestions from all BIG-IP devices are combined to improve the policy.

Using the Central Policy Builder:

1. Each BIG-IP device sends learning suggestions for a particular policy to the centralized policy building device or devices. These devices are BIG-IQ data collection devices (DCD) that have the Web Application Security service activated on them.
2. The policy learning suggestions from all BIG-IP devices are combined so that you can view and manage them on the BIG-IQ Centralized Management system.

3. As suggestions are accepted and the policy is tuned, you can choose to deploy the policy to one or more of the BIG-IP devices that use it.

Configuring central policy building

To configure and use central policy building, you need:

- A BIG-IP® version 13.1 or later device with ASM® provisioned and licensed.
- A BIG-IQ® Centralized Management version 5.4 or later system, with the Web Application Security service installed and the BIG-IP device discovered and imported. Data collections devices also need to be configured.

You use central policy building to combine policy learning suggestions from multiple BIG-IP devices to have more sources for policy suggestions. You configure central policy building by configuring both the data collection devices and the policies that will use central policy building.

1. Configure the data collection devices to be used to collect suggestions from BIG-IP devices.

   You must configure the data collection devices before setting up central policy building.

   a) Click **System** > **BIG-IQ DATA COLLECTION** > **BIG-IQ Data Collection Devices**.
   b) Click the name of the data collection device you want to use for central policy building.

      The data collection device properties screen opens.
   c) On the left, click **Services**, and then on the right, in the Web Application Security area, for the **Activate Service** setting, click **Activate**.
   d) Save your work.
   e) Repeat this process for each data collection device you want to use for central policy building.
2. Verify that learning mode is enabled in the Web Application Security policy.

   Learning mode must be enabled when using either local or centralized policy building.

   a) Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
   b) Click the name of the policy you want to use with central policy building, and then on the left click **POLICY PROPERTIES** > **General Properties**.

      The general properties screen opens.
   c) For the **Learning Mode** setting, select either **Automatic** or **Manual** if one is not already selected.
   d) If you made any changes, save your work.
3. Enable centralized policy building for the Web Application Security policy.

   a) Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
   b) Click the name of the policy you want to use with central policy building, and then on the left click **POLICY BUILDING** > **Settings**.

      The Settings screen opens.
   c) For the **Policy Building Mode** setting, select **Central**.

      When local policy building is configured, this value is set to **Local**.
   d) Save your work.
   e) Repeat this process for each policy that you want to use with central policy building.

Central policy building is now enabled for this policy.

How do I determine what permissions to apply to roles accessing child and parent policies?

When adding or modifying the role type permissions associated with a Web Application Security policy, you need to be aware of whether the policy is a standalone policy without inheritance, a parent policy, or a child policy. You define access to policies using the New Role Type properties screen.

1. Click **System** > **ROLE MANAGEMENT** > **Custom Role Types**.
2. Click **Add**. The New Role Type properties screen opens.
3. Select Web Application Security (ASM) as the service. Those object types are displayed.
4. Select **Policies: Web Application Security** as the object type, and click **Add Selected**.

- To define access to standalone policies that do not use inheritance, select from the permissions without the Child or Parent prefix: Read, Add, Edit, or Delete.
- To define access to only child policies, select permissions with the Child prefix: Child Create, Child Delete, or Child Edit.
- To define access to only parent policies, select permissions with the Parent prefix: Parent Create, Parent Delete, or Parent Edit.

***Note:** If you assign general permissions (Read, Add, Edit, or Delete) to a child or parent policy, you are assigning access to both parent and child policies. For example, assigning the Delete permission to a role allows that role to delete standalone policies, parent policies, and child policies. But, assigning the Child Delete permission to a role allows that role to delete only child policies, and not parent or standalone policies.*

Regardless of the type of policy, you should always allow users Read access to the policy.

Edit application security policies

You modify application security policies to customize how they protect your web application server. Application security policies can be created in Web Application Security. But more often, they are created on BIG-IP® devices and come into the Web Application Security configuration when you discover the devices.

1. Navigate to the Policies screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. 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.
3. Edit the properties of each policy object as needed.

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

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

Edit general property 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. Go to the General Properties screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy to modify, and then on the left click **General Properties**.
3. Edit the properties as appropriate.
4. Save your changes to the general properties of the policy.

The system saves changes in the working configuration of the BIG-IQ Centralized Management system.

General property 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. You can set the **Name** only when you create the policy.
Partition	Partition to which the security policy belongs. Only users with access to a partition can view the 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 the security policy.
Policy Type	Indicates the type of policy. • **Security Policy** specifies a policy that does not use inheritance, or that uses inheritance and is a child policy. • **Parent Policy** specifies a policy that uses inheritance, and is a parent policy.
Parent Policy	Specifies the parent policy associated with this policy, if any. • Select **None** to indicate that there is no parent policy. • Select the appropriate parent policy from the list if there is a parent policy.
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 Policy 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.
Event Correlation Reporting	If enabled, events are reported in groups (correlated), rather than as individual transactions.

Property	Description
	You can only disable this setting for BIG-IP devices version 13.1 or later.
Learning Mode	Select one of the options to indicate how the policy learns: • **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. • **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. • **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	Indicates 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 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	When enabled, they system masks credit card numbers in the request log. If disabled (cleared), credit card numbers are not masked.
Maximum HTTP Header Length	Specifies the 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

Property	Description
	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	Specifies the 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**: Specifies a 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 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

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

Edit inheritance settings

You use the Inheritance Settings screen to change the properties that are part of a policy by editing the inheritance settings of a child or parent policy. .

1. Navigate to the Inheritance Settings screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the appropriate policy name to display the policy properties screen.

3. Click **Inheritance Settings**.
4. Review or modify the inheritance settings.

   The contents of this screen differ depending on whether the policy is a parent policy, a child policy, or neither.
5. If the current policy is neither a parent policy nor a child policy, the **Parent Policy** list is set to **None**, and no other properties are shown on the screen.
6. If the current policy is a child policy or will be a child policy, do the following.
   a) From the **Parent Policy** list, review or select a parent policy. By default, the setting is **None**.
   b) Review the list of properties that are displayed, and where needed, select **Accept** or **Decline**.
   c) Optionally, you can add comments about the inheritance settings by clicking the comment icon in the Comments column and then typing text in the space provided.
7. If the current policy is a parent policy, do the following.
   - In the Inheritance column, review or change the inheritance settings for each property in each property row.
     - If the property must be inherited by a child policy, click **Mandatory**.
     - If the property is optional for a child policy, click **Optional**.
     - If the property is not available to the child policy, click **None**.
   - The Accepted, Declined, Unread, and Comments columns show the number of child policies for each category for that property. Optionally, you can click the number to display additional information on the Child Policy Overview screen.
8. Click **Save** to save your changes.

The inheritance settings for the policy are updated.

Edit child policy overview settings

You can edit the inheritance settings for child policies associated with a parent policy. A parent policy can be associated with multiple child policies.

1. Navigate to the Child Policy Overview screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to review, and click **Child Policy Overview**.
3. Review the inheritance settings for child policies associated with the parent policy.
   - Click **All** to view all properties that could be inherited by a child policy.
   - Click **Declined** to view only the properties that a child policy declined to inherit.
4. Expand each policy section in the list to review the inheritance status (declined or accepted) for each child policy.
5. Indicate whether you have reviewed declined inheritance properties. In the Policy Section row for a child policy property:
   - Click **Mark as Read** to indicate that you have reviewed a declined property for a child policy.
   - Click **Mark as Unread** to indicate that you have not reviewed a declined property for a child policy.
   - Click **Mark All as Read** to indicate that you have reviewed all declined properties within that heading.
   - To enter a comment, click the comment icon in the row. To remove all comments in a section, click **Clear All** in the heading row for a policy section.
6. Click **Save** to save your changes.

The child policy overview is updated.

Response page editing

You can review and change the settings on various types of response pages. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

Edit Ajax response page settings

You use the Ajax Response Page screen to view and edit the settings for the Ajax response page, which is one of several response pages. Response page settings specify the content of the response that the system sends to the user when the security policy blocks a client request.

1. Go to the Ajax Response Page screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and then click **RESPONSE PAGES** > **Ajax**.
3. In the **AJAX Blocking** setting, click the **Enabled** check box to view and edit settings.

   When this is checked (enabled), the system injects JavaScript code into responses.

   Important: *You must enable this check box to configure an ASM Ajax response page which is returned when the system detects an Ajax request that does not comply with the security policy.*
4. From the **Default Response Page Action** list, select an action. Your selection determines the settings.

Option	Description
Popup Message	The screen displays a sample pop up message which you can edit. Click **Preview On** to preview the response.
Custom Response	The screen displays the default response page which you can edit to create a custom response. Alternatively, you can upload the response. • You click **Choose File** to select the file containing the response, and then click **Upload** to insert it. • Click **Preview On** to preview the response. • If you want to return to the original default response text, click **Paste Default Response Body**.
Redirect URL	The system redirects the user to a specific web page instead of displaying a response page. You must enter a URL for the redirect.

5. In the **Login Page Response Action** list, select an action.

   Your selection determines the settings. The actions are the same as those for the **Default Response Page Action** list.
6. In the **Failed Login Honeypot Page Response Action** list, select an action.

   Your selection determines the settings. The actions are the same as those for the **Default Response Page Action** list.
7. When you are finished, save your changes.

The response page settings are updated.

Edit CAPTCHA response page settings

You use the CAPTCHA Response Page screen to view and edit the settings for CAPTCHA responses. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.

2. Click a policy name, on the next screen, on the left click **Response Pages** and then for the Response Pages type, click **CAPTCHA Fail**.
3. For the **Response Type** setting, specify whether to use the default or a custom response.

   - To use the displayed response header and response body, select **Default Response**.
   - To use a modified response header or response body, select **Custom Response**.

   Selecting **Custom Response** makes editing options available.
4. In the **Response Header** setting, review or change the response header.

   - If you selected the default response type, you can review but not modify the response header.
   - If you selected the custom response type, you can modify the response header by editing the header text.
   - To replace your modifications with the default response header, click **Paste Default Response Header**.
5. For the **Response Body** setting, review or change the response body.

   - If you selected the default response type, you can review but not modify the response body.
   - If you selected the custom response type, you can modify the response body by editing the body text directly or by importing a file with that text. To import the response body:

     1. Click **Choose File**.
     2. In the displayed Open dialog box, select the file to import and click **Open**. The Open dialog box closes.
     3. Click **Upload**. The contents of the file are now in the response body text box.
   - To replace your modifications with the default response body, click **Paste Default Response Body**.
6. For the **Preview** setting, specify whether to see a preview of the response body.

   - To see a preview of how the response is displayed, click **Preview On**.
   - To skip the preview, click **Preview Off**.
7. When you are finished, save your changes.

Edit CAPTCHA fail response page settings

You use the CAPTCHA Fail Response Page screen to view and edit the settings for CAPTCHA Fail responses. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, on the next screen, on the left click **Response Pages** and then for the Response Pages type, click **CAPTCHA Fail**.
3. For the **Response Type** setting, specify whether to use the default or a custom response.

   - To use the displayed response header and response body, select **Default Response**.
   - To use a modified response header or response body, select **Custom Response**.

   Selecting **Custom Response** makes editing options available.
4. For the **Response Header** setting, review or change the response header.

   - If you selected the default response type, you can review but not modify the response header.
   - If you selected the custom response type, you can modify the response header by editing the header text.
   - To replace your modifications with the default response header, click **Paste Default Response Header**.
5. For the **Response Body** setting, review or change the response body.

- If you selected the default response type, you can review, but not modify, the response body .
- If you selected the custom response type, you can modify the response body by editing the body text directly or by importing a file with that text. To import the response body:
    1. Click **Choose File**.
    2. In the displayed Open dialog box, select the file to import and click **Open**. The Open dialog box closes.
    3. Click **Upload**. The contents of the file are now in the response body text box.
- To replace your modifications with the default response body, click **Paste Default Response Body**.

6. In the **Preview** setting, select whether to see a preview of the response body.
    - To see a preview of how the response is displayed, click **Preview On**.
    - To skip the preview, click **Preview Off**.
7. When you are finished, save your changes.

Edit default response page settings

You use the Default Response Pages screen to view and edit the settings for the default response page, which is one of several response pages. Response page settings specify the content of the response that the system sends to the user when the security policy blocks a client request.

1. Go to the Default Response Page screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and then click **RESPONSE PAGES** > **Default**.
3. Select a **Response Type** from the list. Your selection determines the additional settings.

Option	Description
Default Response	The screen displays the default response header and response body. The system sends the response body to the client as shown. You cannot edit these fields. Click **Preview On** to preview the response.
Custom Response	The screen displays the default response header and response body which you can edit to create a custom response. Alternatively, for the response body, you can upload a response. • Click **Choose File** to select the file containing the response body, and then click **Upload** to insert it. • Click **Preview On** to preview the response. • If you want to return to the original default response text for the header or the body, click **Paste Default Response Header** or **Paste Default Response Body**.
Redirect URL	The system redirects the user to a specific web page instead of displaying a response page. You must enter a URL in the **Redirect URL** field.
Soap Fault	The system blocks a SOAP request due to an XML-related violation. The system displays the system-supplied response written in the SOAP fault message structure. You cannot edit this text. Click **Preview On** to preview the response.
Erase Cookies	The system deletes all client side domain cookies. This is done in order to block web application users once, and not from the entire web application. The system displays this text in the response page. You cannot edit this text. The response header and response body are shown. Click **Preview On** to preview the response.

4. When you are finished, save your changes.

The response page settings are updated.

Edit failed login honeypot response page settings

You use the Failed Login Honeypot screen to view and edit the settings for the Failed Login Honeypot response page. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, on the left of the next screen, click **Response Pages**then for the Response Pages type, click **Failed Login Honeypot**.
3. For the **Response Type** setting, specify whether to use the default or a custom response.

   - To use the displayed response header and response body, select **Default Response**.
   - To use a modified response header or response body, select **Custom Response**.

   Selecting **Custom Response** makes editing options available.
4. For the **Response Header** setting, review or change the response header.

   - If you selected the default response type, you can review, but not modify the response header.
   - If you selected the custom response type, you can modify the response header by editing the header text.
   - To replace your modifications with the default response header, click **Paste Default Response Header**.
5. For the **Response Body** setting, review or change the response body.

   - If you selected the default response type, you can review, but not modify the response body.
   - If you selected the custom response type, you can modify the response body by editing the body text directly or by importing a file with that text. To import the response body:

     1. Click **Choose File**.
     2. In the displayed Open dialog box, select the file to import and click **Open**. The Open dialog box closes.
     3. Click **Upload**. The contents of the file are now in the response body text box.
   - To replace your modifications with the default response body, click **Paste Default Response Body**.
6. For the **Preview** setting, specify whether to see a preview of the response body.

   - To see a preview how the response is displayed, click **Preview On**.
   - To skip a preview, click **Preview Off**.
7. When you are finished, save your changes.

Edit cookie hijacking response page settings

You use the Cookie Hijacking Response Page screen to view and edit the settings for the Cookie Hijacking response page. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, on the left of the next screen, click **Response Pages**, and for the Response Pages type, click **Cookie Hijacking**.
3. For the **Response Type** setting, specify the type of response to use.

   - To use the default response header and body, select **Default Response**.
   - To use a modified response header or body, select **Custom Response**.
   - To use the SOAP fault response header and body, select **SOAP Fault**.
   - To use the erase cookies response header and body, select **Erase Cookies**.

The response header and body change based on the response type you select. Selecting **Custom Response** makes editing options available.

4. For the **Response Header** setting, review or change the response header.
   - If you did not select **Custom Response** as the response type, you can review but not modify the response header.
   - If you selected **Custom Response** as the response type, you can modify the response header by editing the header text.
   - To replace your modifications with the default response header, click **Paste Default Response Header**.
5. For the **Response Body** setting, review or change the response body.
   - If you did not select **Custom Response** as the response type, you can review but not modify the response body.
   - If you selected the custom response type, you can modify the response body by editing the body text directly or by importing a file with that text. To import the response body:
     1. Click **Choose File**.
     2. In the displayed Open dialog box, select the file to import and click **Open**. The Open dialog box closes.
     3. Click **Upload**. The contents of the file are now in the response body text box.
   - To replace your modifications with the default response body, click **Paste Default Response Body**.
6. For the **Preview** setting, specify whether to see a preview of the response body.
   - To see a preview how the response is displayed, click **Preview On**.
   - To skip a preview, click **Preview Off**.
7. When you are finished, save your changes.

Edit mobile application response page settings

You use the Mobile Application Response Page screen to view and edit the settings for the mobile application response page. Response page settings specify the response content that the system sends to the user when the security policy blocks a client request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, on the left of the next screen click **Response Pages** and for the Response Pages type, click **Mobile Application**.
3. for the **Response Type** setting, specify whether to use the default or a custom response.
   - To use the displayed response header and response body, select **Default Response**.
   - To use a modified response header or response body, select **Custom Response**.

   Selecting **Custom Response** makes editing options available.
4. For the **Response Header** setting, review or change the response header.
   - If you selected the default response type, you can review but not modify the response header.
   - If you selected the custom response type, you can modify the response header by editing the header text.
   - To replace your modifications with the default response header, click **Paste Default Response Header**.
5. For the **Response Body** setting, review or change the response body.
   - If you selected the default response type, you can review but not modify the response body.
   - If you selected the custom response type, you can modify the response body by editing the body text directly or by importing a file with that text. To import the response body:

1. Click **Choose File**.
2. In the displayed Open dialog box, select the file to import and click **Open**. The Open dialog box closes.
3. Click **Upload**. The contents of the file are now in the response body text box.

- To replace your modifications with the default response body, click **Paste Default Response Body**.

6. For the **Preview** setting, specify whether to see a preview of the response body.
   - To see a preview how the response is displayed, click **Preview On**.
   - To skip a preview, click **Preview Off**.
7. When you are finished, save your changes.

Edit login response page settings

You use the Login Pages Response Page screen to view and edit the settings for the login page response page, which is one of several response pages. Response page settings specify the content of the response that the system sends to the user when the security policy blocks a client request.

1. Go to the Login Pages Response Page screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and then click **RESPONSE PAGES** > **Login Page**.
3. Select a **Response Type** from the list. Your selection determines the additional settings.

Option	Description
Default Response	The screen displays the default response header and response body. The system sends the response body to the client as shown. You cannot edit these fields. Click **Preview On** to preview the response.
Custom Response	The screen displays the default response header and response body which you can edit to create a custom response. Alternatively, for the response body, you can upload a response. • Click **Choose File** to select the file containing the response body, and then click **Upload** to insert it. • Click **Preview On** to preview the response. • If you want to return to the original default response text for the header or the body, click **Paste Default Response Header** or **Paste Default Response Body**.
Redirect URL	The system redirects the user to a specific web page instead of displaying a response page. You must enter a URL in the **Redirect URL** field.
Soap Fault	The system blocks a SOAP request due to an XML-related violation. The system displays the system-supplied response written in the SOAP fault message structure. You cannot edit this text. Click **Preview On** to preview the response.
Erase Cookies	The system deletes all client side domain cookies. This is done in order to block web application users once, and not from the entire web application. The system displays this text in the response page. You cannot edit this text. The response header and response body are shown. Click **Preview On** to preview the response.

4. When you are finished, save your changes.

The response page settings are updated.

Edit XML response page settings

You use the XML Response Page screen to view and edit the settings for the XML response page, which is one of several response pages. Response page settings specify the content of the response that the system sends to the user when the security policy blocks a client request.

1. Go to the XML Response Page screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and then click **RESPONSE PAGES** > **XML**.
3. Select a **Response Type** from the list. Your selection determines the additional settings.

Option	Description
Custom Response	The screen displays the default response header and response body which you can edit to create a custom response. Alternatively, for the response body, you can upload a response. • Click **Choose File** to select the file containing the response body, and then click **Upload** to insert it. • Click **Preview On** to preview the response. • If you want to return to the original default response text for the header or the body, click **Paste Default Response Header** or **Paste Default Response Body**.
Soap Fault	The system blocks a SOAP request due to an XML-related violation. The system displays the system-supplied response written in the SOAP fault message structure. You cannot edit this text. Click **Preview On** to preview the response.

4. When you are finished, save your changes.

The response page settings are updated.

Edit policy building overview settings

You can review or change various overall aspects of policy building using the Policy Building Overview screen.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to edit, and on the left of the new screen, click **POLICY BUILDING** > **Overview**.
3. To review the status of the devices being used for central policy building, expand Devices of Central Policy Building.

   The screen lists the devices and their status. You can click **Refresh** to update the device information.
4. To review or change enforcement readiness for various policy entities, expand Enforcement Readiness Summary.

   The screen shows a summary list of entities in the security policy that can be enforced, along with their status.
5. For additional information, you can click the links shown in the summary table.
   - In the Entity Type column, click the name of the policy entity type to review a list of suggestions for it, if any exist.
   - In the Learn New Entities column, you can see the learning status for the policy entity.
   - In the Total, Not Enforced, or Not Enforced And Have Suggestions columns, click the number of entities link to be taken to the appropriate policy screen. Entities that are not enforced are in

staging, or have wildcard entities configured so that the security policy learns all explicit entities that match them.
- In the Ready to be Enforced column, review the numbers. To enforce these entities, select the check box to the left in the row and click **Enforce Selected Entities**.
- To update the data shown, click **Refresh**.

6. To review the suggestions used to reduce false positive alerts, expand Suggestions To Reduce Potential False-Positive Alerts.

   In this area, you see three lists: Top Violations, Top Violating Meta Characters, and Top Matched Signatures. You may need to scroll down to see all three. Each list contains suggestions for the entities listed, if there are any.

   - To see the suggestions associated with a listed entity, such as one of the top violations, click the name.
   - To update the data shown, click **Refresh**.
7. To review suggestions to add to new entries, expand Suggestions To Add New Entries.
   - You can review the new suggestions. These are listed by entity type.
   - To update the data shown, click **Refresh**.
8. Save your work.

Your changes are applied to the security policy.

Edit policy building suggestion settings

You can view policy building suggestions and decide how to respond to each suggestion, such as accepting, ignoring, or deleting the suggestion.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to edit, and on the left, click **POLICY BUILDING** > **Suggestions**.
3. To accept a suggestion, select it, click **Actions**, and select one of the accept options.
   - To accept a suggestion and have it added to the policy entity, select the **Accept** action.
   - To accept this suggestion and have it added to the policy entity in staged mode, select the **Accept and Stage** action.
   - To accept this suggestion and have it added globally at the policy level, select the **Accept Globally** action.

   Not all options are available for all suggestions. Unsupported options for a suggestion are not selectable. For example, the **Accept and Stage** option is only available for policy entities that support staging, such as signatures and URLs.
4. To delete a suggestion, select the check box to the left of the suggestion and click **Delete**.

   The policy builder can suggest a deleted suggestion again.
5. To ignore a suggestion, select the check box to the left of the suggestion and click **Ignore**.

   Once a suggestion is ignored, the policy builder will not suggest it again.
6. To see additional details about the suggestion, click the name of the suggestion.

   The additional details for the suggestion vary, but may include other related suggestions and the list of samples.
7. To add a comment to a suggestion, click the icon in the Comment column for that suggestion, and type your comment in the text box that opens.
8. To list either all suggestions or only a subset of the suggestions, select one of the options in the filtering area in the upper left of the screen, such as **Pending Suggestions** or **Ignored Suggestions**.

9. To perform a simple search of the suggestions, type the text to search for in the search area in the upper right of the screen.

   You cannot use the simple search when looking for a violation or a refinement. You must use an advanced search filter instead and select the violation or refinement.

10. To perform an advanced search using a filter, click the icon to the left of the search area in the upper right of the screen. The filter dialog box opens.

    - To use an existing filter, click the filter name. The filter is applied.
    - To create a new filter, click **Create**. The New Filter dialog box opens.

      1. In the **Filter Name** setting, type a name for the filter.
      2. In the Query Parameter area, specify values for the parameters you want to use to create the search filter. As you select parameters, the system creates the query in the Query Expression area.
      3. When you are done, click **Save & Apply** to save your changes and apply the filter.

Edit policy building settings

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

1. Go to the Blocking Settings screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of a policy, and from the list on the left select **POLICY BUILDING** > **Settings**.
3. Click the arrows to open or close each category and display specific violation types available to configure for that category.
4. Edit the settings to meet your requirements.
5. When you are finished, save your work.

This updates the blocking settings in the application security policy.

Policy building settings properties

You configure the settings of the security policy to specify how the system responds to a request that contains each type of illegal request.

Blocking Setting	Description
Enforcement Mode	Specify 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.
Learning Mode	Specify how learning is, or is not, performed. • **Automatic** has the system examine traffic, make suggestions, and enforce most suggestions after sufficient traffic over a period of time from various users make it reasonable

Blocking Setting	Description
	to add them. A few suggestions must be enforced manually. • **Manual** has the system examine traffic and make suggestions on what to add to the security policy. • **Disabled** has the system do no learning for the security policy, and make no suggestions.
Policy Building Mode	Specify how policy building is performed. The option you select changes the other settings that are available. • To have policy building occur on the local BIG-IP® device, select **Local**. • To have policy building occur on a central policy builder device that can take information from multiple BIG-IP devices, select **Central**.
Policy Building Device	Specify which central policy building device to use. This option is available only when central policy building mode is selected. The policy building device is also a data collection device. • To have the central policy building device chosen automatically, select **Auto Select**. • To manually choose the device to use for central policy building, select the device.
Auto-Apply Policy	Specify when learning is automatically applied to the policy. • To have learning applied to the policy as it occurs, select **Real Time**. • To have learning applied to a policy at a scheduled time, select **Scheduled**, and then specify that time and day. • To have learning applied at any time during the day, select **All Day**. • To have learning applied at a certain time during the day, select **Specific Time** and provide that time. • To have learning applied every day of the week, select **All Week**. • To have learning applied on selected days of the week, select **Specific Days**and specify the days. • To have the system not apply learning to the security policy, select **Disabled**.
Learning Speed	Select the speed the Policy Builder uses for learning. • To have the Policy Builder learn traffic from a small number of requests, sessions, and IP

Blocking Setting	Description
	addresses, select **Fast**. Using this setting, there is a high chance that the Policy Builder will add false entities to the security policy. • To have the Policy Builder learn traffic from a medium amount of requests, sessions, and IP addresses, select **Medium**. Using this setting, there is a medium chance that the Policy Builder will add false entities to the security policy. • To have the Policy Builder learn traffic from a large number of requests, sessions, and IP addresses, select **Slow**. Using this setting, there is a low chance that the Policy Builder will add false entities to the security policy. • To have the Policy Builder use custom settings in the policy, select **Custom**. The **Custom** option is selected automatically when you customize settings in the policy.
All Violations	Select the **Learn**, **Alarm** or **Block** check boxes in this row to have those selections apply to all the violations in this group. You can select or clear these check boxes in the violation rows to change the behavior for individual violations, or groups of violations. • **Learn** specifies that when a request triggers this violation, the system learns the request. • **Alarm** specifies that if a request triggers this violation, the system records the request. • **Block** specifies that if this violation occurs, the system takes these actions: • Records the request. • Blocks the request that triggered the violation. • Returns the blocking response page to the client who sent the request.
Policy General Features	Expand this setting to see the contained violations. Click the information icon next to each violation for more information about it. Select **Learn**, **Alarm**, or **Block** for each, as appropriate for your policy.
HTTP protocol compliance failed	Expand this setting to see the sub-violations, and click the information icons for more information. Either select the **Enable** or **Learn** check box at the top of the section to select all HTTP protocol compliance failed sub-violations at once, or select the **Enable** or **Learn** check box to the left of each sub-violation to specify that the system enforces the sub-violation. When the check box is cleared, the system does not enforce this sub-violation.

Blocking Setting	Description
	This category contains the following sub-violations. • **Bad HTTP version**. When checked (enabled), 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 • **Bad host header value**. When checked (enabled), the system inspects requests to see whether they contain a non RFC compliant header value. The default value is enabled. • **Bad multipart parameters parsing**. When checked (enabled), 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. • **Bad multipart/form-data request parsing**. When checked (enabled), the system examines requests to see whether the content-disposition header field contains the required parameters, `name=‟param_key”`. The default setting is enabled. • **Body in GET or HEAD requests**. When checked (enabled), 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. • **CRLF characters before request start**. When checked (enabled), 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**. When checked (enabled), the system compares the number of headers in the requests against the maximum number you specify here. Type a number in the field to specify how many headers are allowed. The default setting is enabled with a maximum of 20 headers unless the policy is based on an Application-Ready security policy template. In this case, the default value depends on which template you are using. • **Check maximum number of parameters**. When checked (enabled), 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 policy should be considered a possible attack on the server. Type a number in

Blocking Setting	Description
	the field to specify how many parameters are allowed. The default value is enabled set to a maximum of 500 parameters. • **Chunked request with Content-Length header**. When checked (enabled), the system examines chunked requests for a content-length header, which is not permitted. The default setting is enabled. • **Content length should be a positive number**. When checked (enabled), the system examines requests to see whether their content length value is greater than zero, which is required. The default setting is enabled. • **Header name with no header value**. When checked (enabled), the system checks requests for valueless header names, which are considered illegal. The default setting is enabled. • **High ASCII characters in headers**. When checked (enabled), the system inspects request headers for ASCII characters greater than 127, which are not permitted. The default setting is disabled. • **Host header contains IP address**. When checked (enabled), the system verifies that the request's host header value is not an IP address. The default setting is disabled. • **Multiple host headers**. When checked (enabled), 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**. When checked (enabled), 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. • **Null in request**. When checked (enabled), the system inspects requests to see whether they contain a Null character, which is not allowed. The default setting is enabled. • **POST request with Content-Length 0**. When checked (enabled), the system examines POST method requests for no content-length header, and for a content length of 0. The default setting is disabled. • **Several Content-Length headers**. When checked (enabled), 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
	• **Unparseable request content**. When checked (enabled), 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	Expand this setting to see the evasion technique sub-violations and click information icons for more information. Either select the **Enable** or **Learn** check box at the top of the section to select all sub-violations at once, or select the **Enable** or **Learn** check box to the left of each sub-violation to specify that the system enforces the sub-violation. When the check box is cleared, the system does not enforce this sub-violation. This category contains the following sub-violations. • **%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. • **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. • **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. • **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.

Blocking Setting	Description
	• **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. 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. • **Multiple decoding** . Indicates that the system performs multiple decoding. Use the decoding passes drop-down control to specify the number (up to 5) of decoding passes. 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.
File Types	Expand this setting to see the file type sub-violations, and click information icons for more information. 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. This category contains the following sub-violations. • **Illegal query string length**. The incoming request contains a query string whose length exceeds the acceptable length specified in the policy. • **Illegal POST data length**. The incoming request contains POST data whose length exceeds the acceptable length specified in the policy. • **Illegal request length**. The incoming request length exceeds the acceptable length specified in the policy per the requested file type. • **Illegal file type**. The incoming request references file types not found in the policy. • **Illegal URL length**. The incoming request includes a URL whose length exceeds the acceptable length specified in the policy.

Blocking Setting	Description
	In the **Learn New File Types** setting, select under which circumstances the Policy Builder adds, or suggests you add, explicit file types to the security policy. As you select the setting, additional information about the setting is displayed below it. In the **Maximum Learned File Types** setting, type the maximum number. The default value changes based on the value of the **Learn New File Types** setting.
URLs	Expand this area to see the URL sub-violations and click the information icons for more information on each. • In the **Learn New HTTP URLs** setting, select under which circumstances the Policy Builder adds, or suggests you add, HTTP URLs to the security policy. As you select the setting, additional information about the setting is displayed below it. • In the **Maximum Learned HTTP URLs** setting, type the maximum number. The default value changes based on the value of the **Learn New HTTP URLs** setting. • In the **Learn New WebSocket URLs** setting, select under which circumstances the Policy Builder adds, or suggests you add, WebSocket URLs to the security policy. As you select the setting, additional information about the setting is displayed below it. • In the **Maximum Learned WebSocket URLs**setting, type the maximum number, • In the **Classify Request Content of Learned HTTP URLs** setting, use the **Enabled** check box to specify whether it should be enabled. When enabled, if the Policy Builder detects legitimate XML or JSON data to URLs configured in the security policy, the Policy Builder adds XML or JSON profiles to the security policy and configures their attributes according to the data it detects. • In the **Classify Client Message Payload Format of Learned WebSocket URLs** setting, use the **Enabled** check box to specify whether it should be enabled. When enabled, if the Policy Builder detects legitimate plain text or JSON data to WebSocket URLs configured in the security policy, the Policy Builder adds Plain Text or JSON profiles to the security policy and configures their attributes according to the data it detects. • In the **Learn Allowed Methods on HTTP URLs** setting, use the **Enabled** check box to

Blocking Setting	Description
	specify whether it should be enabled. When enabled, if the Policy Builder detects a method used in a request that is not in the security policy's list of generic methods, the Policy Builder adds the new method to the security policy and associates it to the specific requested URL. • In the **Collapse many common HTTP URLs into one wildcard HTTP URL** setting, use the **Enabled** check box to specify whether it should be enabled. When enabled, the system collapses many common explicit URLs into one wildcard URL with a common prefix and suffix. The Policy Builder collapse only URLs in the same directory (with the same prefix path), and if they have the same file extension. You can also type the number of occurrences and the depth. • In the **File types for which wildcard HTTP URLs will be configured** setting, select which file types to add or delete.
Parameters	Expand this area to see the parameter sub-violations and click the information icons for more information on each. • In the **Learn New Parameters** setting, select under which circumstances the Policy Builder adds, or suggests you add, explicit parameters to the security policy. As you select the setting, additional information about the setting is displayed below it. • In the **Maximum Learned Parameters** setting, type the maximum number of parameters that the security policy allows. • In the **Parameter Level** setting, select how the Policy Builder determines on which level to add, or suggest you add, parameters to the security policy. Select **Global** or **URL**, and review the information displayed about each. • In the **Collapse many common Parameters into one wildcard Parameter** setting, select the check box to specify that the system collapses many common parameters into one wildcard parameter. In the field, type how many explicit parameters the Policy Builder must detect (the number of occurrences) before collapsing them to one wildcard parameter. • In the **Classify Value Content of Learned Parameters** setting, when enabled, specifies that if the Policy Builder detects legitimate XML or JSON data to parameters configured in the security policy, the Policy Builder adds

Blocking Setting	Description
	XML or JSON profiles to the security policy and configures their attributes according to the data it detects. • In the **Learn Integer Parameters** setting, specifies, when enabled, that the Policy Builder learns integer parameters. • In the **Learn Dynamic Parameters** setting, you specify the conditions under which the Policy Builder adds dynamic parameters to the security policy.
Sessions and Logins	Expand this area to see the session and login sub-violations, and click the information icons for more information on each violation. In the **Detect login pages** setting, select the **Enabled** check box to have the Policy Builder detect login pages by examining traffic to the web application.
Cookies	Expand this area to see the cookie sub-violations and click the information icons for more information on each violation. • In the **Learn New Cookies** setting, select the circumstances the Policy Builder adds, or suggests you add, explicit cookies to the security policy. As you select the setting, additional information about the setting is displayed below it. • In the **Maximum Learned Cookies** setting, type the largest number of cookies that the policy allows. • In the **Learn and enforce new unmodified cookies** setting, specify, when the **Enabled** check box is selected, that the system enforces new cookies as long as they were not modified by the client. This option only appears if the **Learn New Cookies** option is set to **Selective** and the * wildcard cookie is of type **Allowed**. • In the **Collapse many common Cookies into one wildcard Cookie** setting, you specify, when the **Enabled** check box is selected, that the system collapses many common cookies into one wildcard cookie. Type in the box how many explicit cookies the Policy Builder must detect (the number of occurrences) before collapsing them to one wildcard cookie.
Content Profiles	Expand this area to see the content profile sub-violations, and click the information icons for more information on each violation. In the **Collapse many common Content Profiles into one wildcard Content profile** setting, you

Blocking Setting	Description
	specify, when the **Enabled** check box is selected, that the system collapses many common content profiles into one wildcard content profile. Type in the field how many explicit content profiles the Policy Builder must detect (the number of occurrences) before collapsing them to one wildcard content profile.
Web Services Security Failure	Expand this area to see the web services security failure sub-violations. At the top of the list of sub-violations, select either the **Enable** or **Learn** check box to select all sub-violations at once, or select the **Enable** or **Learn** check box to the left of each sub-violation to specify that individual sub-violation. • **Certificate Error**. When checked (enabled), 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. • **Certificate Expired**. When checked (enabled), 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. ***Note:*** *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.* • **Decryption Error**. When checked (enabled), 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
	• **Encryption Error**. When checked (enabled), 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. • **Expired Timestamp**. When checked (enabled), the system learns, logs, or blocks requests when the timestamp has expired. The default setting is enabled. • **Internal Error**. When checked (enabled), 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. • **Invalid Timestamp**. When checked (enabled), the system learns, logs, or blocks requests when the timestamp is not formatted according to the specifications. The default setting is enabled. • **Malformed Error**. When checked (enabled), 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. • **Missing Timestamp**. When checked (enabled), the system learns, logs, or blocks requests when the timestamp is missing from the document. The default setting is enabled. • **Signing Error**. When checked (enabled), 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. • **Timestamp expiration is too far in the future**. When checked (enabled), the system learns, logs, or blocks requests when the timestamp lifetime is greater than configured. The default setting is enabled. • **UnSigned Timestamp**. When checked (enabled), the system learns, logs, or blocks requests when the timestamp is not digitally signed. The default setting is enabled. • **Verification Error**. When checked (enabled), the system learns, logs, or blocks requests

Blocking Setting	Description
	when the underlying crypto library failed to perform digital signature verification, or there is information missing in the payload. The default setting is enabled.
CSRF Protection	Expand this area to see the cross-site request forgery (CSRF) protection sub-violations. *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 this setting is enabled, the system protects the web application against CSRF attacks. This category contains the following violations. • **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	Expand this area to see the IP address and Geolocation sub-violations, and click the information icons for more information on each violation.
Headers	Expand this area to see the header sub-violations and click the information icons for more information on each violation. In the **Learn Host Names** setting, Select the **Enabled** check box to specify that the Policy Builder suggests you add host names that have not yet been added to the policy.
Redirection Protection	Expand this area to see the redirection protection sub-violations, and click the information icons for more information on each violation. In the **Learn New Redirection Domains** setting, select under which circumstances the Policy Builder adds, or suggests you add, explicit redirection domains to the policy. As you select the setting, additional information about the setting is displayed below it. In the **Maximum Learned Redirection Domains** setting, type the largest number of redirection domains that the policy allows.
Bot Detection	Expand this area to see the WebSocket sub-violations. The Bot Detection category contains the **Web scraping detected** violation, which

Blocking Setting	Description
	detects when the web client, or user agent, does not demonstrate human behavior.
Data Guard	Expand this area to see the Data Guard sub-violations. The Data Guard category specifies which information the system considers sensitive, including credit card numbers, U.S. Social Security numbers, custom patterns, and file content. This category contains the **Data Guard. Information leakage detected** violation.
WebSocket protocol compliance	Expand this area to see the WebSocket protocol compliance sub-violations, and click the information icons for additional information about each violation.
Antivirus Protection	Expand this area to see the antivirus protection sub-violations, and click the information icons for additional information about each violation.

Edit policy building process settings

You can view and edit the building process settings for the application security policy to specify how Web Application Security builds policies.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of a policy, and on the left, click **POLICY BUILDING** > **Settings**.
3. Scroll to near the bottom of the settings properties screen, expand the Policy Building Process area, and specify settings as needed.
4. In the **Trust IP Addresses** setting, you specify IP addresses that the Policy Builder considers safe.
   - Select **All IP Addresses** to indicate that all IP addresses are safe.
   - Select **Address List** to indicate that all IP addresses in the displayed list are safe.
   - Click **Edit IP Addresses** to add IP addresses to the list.
5. For the **Loosen Policy** setting, you specify the number of sources that the system must detect during a specified time period, in order for the Policy Builder to accept and learn a policy change from traffic.

   For example, when the Policy Builder detects the same file type, URL, parameter, or cookie from enough different user sessions and IP addresses over time, it adds the entity to the security policy. You can configure values for both untrusted traffic and for trusted traffic.
6. For the **Tighten Policy (stabilize)** setting, you specify the number of requests and the amount of time that must pass for the Policy Builder to stabilize the policy element.

   Stabilizing the policy element may mean tightening it by deleting wildcard entities, removing entities from staging mode, and enforcing violations that did not occur, depending on the element.
7. For the **Minimize false positives (Track Site Changes)** setting, you specify whether, after stabilizing the policy, the Policy Builder remains enabled, and if enabled, how it handles trusted and untrusted traffic to minimize false positives.
   - Select **Enable** to specify that after the Policy Builder stabilizes the policy, the Policy Builder remains enabled, and may still make changes to the policy by loosening it, usually as a result of changes to the web application. Specifies, when cleared (disabled), that after the Policy Builder stabilizes the policy, it disables itself and makes no more changes to the policy, even if it detects that changes were made to the web application.

- Select **From Trusted and Untrusted Traffic** to specify that the Policy Builder loosens the policy based on traffic from trusted and not trusted sources. This setting is available only if **Enable** is selected.
- Select **Only from Trusted Traffic** to specify that the Policy Builder loosens the policy based on traffic from trusted sources. Click **IP Address Exceptions** to define a trusted IP addresses. This setting is available only if **Enable** is selected.

You configure values for trusted and untrusted traffic separately.

8. For the **Options** setting, you can establish options that determine what type of entities the Policy Builder adds to the policy.

   - When enabled, **Learn from responses** specifies that the Policy Builder adds elements found in responses to the policy, when either of the following circumstances is true:

     - The Policy Builder trusts the request IP address (because the request IP address appears in the Trusted IP Addresses list).
     - The Policy Builder does not trust the request IP address (because the request IP address does not appear in the Trusted IP Addresses list), but the request is legal and fully enforced. *Legal* means that the request does not trigger any violations, suggestions to learn explicit entities, and staging suggestions. *Fully enforced* means that the system is not currently determining whether URLs and parameters found in the request are to be parsed as JSON or XML and should be assigned to content profiles.

     If the Policy Builder learns from responses, it uses the trusted traffic thresholds configured on this screen. This does not include learning from violations in responses. In that case, the thresholds are determined by whether the client IP address is trusted or untrusted.

     Specifies, when **Learn from responses** is cleared (disabled), that the Policy Builder never adds elements found in responses to the policy. Violations occurring in responses are learned according to the learn flag of each violation and do not depend on this setting.

     Note: *This setting applies only to what can be learned from the response content such as occurrences of URLs and parameters. It does not apply to learning from violations that occur in responses, such as Data Guard leakage. Learning from these violations is determined by the Learn flag of the respective violation.*

   - When enabled, **Suggest to delete policy entity if it was not observed in traffic for more than** specifies that a suggestion to delete a policy entity should be made if that entity hasn't been observed in traffic for the specified number of days.
   - Select **Full Policy Inspection** to specify that the Policy Builder learns all policy elements. Specifies, when cleared (disabled), that you are limiting the number of entities the Policy Builder learns.

     Note: *Do not disable this check box unless F5 Support advises it.*

   - In the **HTTP Response Status Codes used to learn traffic** setting, you can specify that the Policy Builder extracts information from traffic based on transactions that return specific HTTP response status codes. In the field type the response code that must be returned in order for the Policy Builder to extract information from that traffic.

     Click **Add** to add the response status code to the response status codes list. You may enter either a specific response code number from 0 to 599, or a generic code, for example, `4xx`. The response status codes list displays the response codes allowed by the Policy Builder and in the policy.

     Click the **X** to the left of the response code to remove the selected response status code from the response status codes list and to make it disallowed by the Policy Builder.

9. When you are finished, save the modifications.

The policy building process settings are updated in the application security policy.

Edit server technologies settings

You can add server technologies to your security policy so that your policy can be automatically associated with the correct attack signature sets for the technology. Server technologies can be server-side applications, frameworks, programs, web servers, operating systems, and so on.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to edit, then on the next screen, on the left, click **Server Technologies**.
3. Select the server technology from the list.
   A confirmation dialog box opens listing the changes that will be made to the policy.
4. Confirm that you want to add the server technology by clicking **OK** in the dialog box.
   The technologies are added to the list of selected server technologies.
5. To remove a server technology entry, click the **X** to the left of that entry.
6. Save your work.

Edit 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. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and on the left click **Data Guard**.
3. In the Data Guard setting, select **Enabled** so that you can modify the other settings.

   When this setting is disabled, the system sends the response, including the sensitive information, to the user.
4. Modify the settings as needed to specify 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	When selected, the system considers U.S. security card numbers as sensitive data.
Mask sensitive data	When selected, the system masks sensitive data returned by the web server by returning asterisk (*) characters to the client instead of the sensitive data.
Custom Patterns	When selected, specifies that the system recognizes customized patterns as sensitive data. In the field, type a pattern that you want the system to consider as sensitive data, and click **Add**. Use PCRE regular expression syntax for the pattern, for example, `999-[/d][/d]-[/d][/d][/d][/d]`. To delete a selected pattern, click **X**.
Exception Patterns	When selected, the system recognize exception patterns as not being sensitive data. In the field, type a pattern that you want the system to consider as an exception to sensitive data, and click **Add**. Use PCRE regular expression syntax for the pattern, for example, `999-[/d][/d]-[/d][/d][/d][/d]`. To delete a selected pattern, click **X**.

Option	Description
File Content Detection	When this is selected, you specify the possible types of content the system could consider as sensitive data. The system checks responses for the selected file content, and if it finds it, that content is not returned,
Enforcement Mode	Specify whether the listed URLs should be enforced or ignored by Data Guard. • Select **Enforce URLs in List** to have Data Guard protect these URLs even if they are not in the policy. • Select **Ignore URLs in List** to have Data Guard protect all URLS except those in this list. Add a URL to the list by typing it in the field and clicking **Add**. ***Note:*** *When adding URLs, you can type either explicit (`/index.html`) or wildcard (`*xyz.html`) URLs.*

5. When you are finished, save your work.

The policy is updated to use the new Data Guard settings.

Edit CSRF protection settings

You can enable and modify CSRF protection properties in your security policy to better protect your applications from a CSRF attack. *Cross-site request forgery* (CSRF) is an attack method that exploits a pre-existing relationship of trust and forces a user to run unwanted actions on a web application in which the user is currently authenticated.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of a policy, and from the list on the left, select **CSRF Protection**.
3. For the **CSRF Protection** setting, select the **Enabled** check box.
   The screen displays the other property settings.
4. For the **SSL Only** setting, select the **Enabled** check box.
5. For the **Expiration Time** setting, select the **Enabled** check box, then provide the expiration time in seconds in the area provided.
6. Use the CSRF URLs area to add, remove, or modify CSRF URLs to be protected.

   Existing URLs are listed in their evaluation order at the bottom of the screen.

   For BIG-IP® device versions earlier than 13.1, URLs added to the CSRF URLs list must have the following settings:

   - Specify the **Method** setting as **Any**.
   - Specify the **Enforcement Action** setting as **Verify CSRF Token**.
   - Specify the **Required Parameters** setting as **At Least One**.
7. To add a new URL to the list:
   a) In the **Method** setting, select the method type.
   b) In the **URL** setting, type the URL to be protected.
   c) In the **Enforcement Action** setting, select the type of enforcement.
   d) In the **Required Parameters** setting, specify whether there are any required parameters, and if needed, provide them.
   e) Click **Add**.
8. To edit existing CSRF URLs:

   - To modify a URL, change the required parameters or enforcement action in the list at the bottom of the screen.

- To change the evaluation order of a URL, drag the URL row to a new position in the list.
- To delete a URL from the list, click the **X** to the left of the URL.

9. Save your work.

Add or edit brute force attack prevention settings

You can protect login URLs against brute force attacks. A *brute force* attack is an outside attempt by hackers to access post login pages of a website by guessing user names and passwords. Brute force attacks are performed when a hacker tries to log in to a URL numerous times, running many combinations of user names and passwords, until he successfully logs in. The `Default` login URL is used for all defined login URLs that do not have their own brute force configuration.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of a policy, and on the left click **ANOMALY DETECTION** > **Brute Force Attack Prevention**.
3. Specify the action to take for brute force attack prevention settings:
   - To add a login URL to the security policy, click **Add**.
   - To modify the brute force prevention properties for a login URL, click the name of the login URL.

   The brute force prevention properties display.
4. Supply the general properties for brute force attack prevention for the login URL.
   a) In the **Login Page** setting, select a login page, or create a login page by clicking **Create login page**.
   b) In the **Configuration Support** setting, specify whether to use current or legacy settings. The other available properties differ based on this setting.
      - Select **Current** when managing a BIG-IP® device version later than 13.0.
      - Select **13.0 And Prior** when managing a BIG-IP device version 13.0 or earlier.

   c) In the **IP Address Whitelist** setting, review the settings or add new settings. To add an IP address, click the **IP Address Whitelist** setting link.
5. In the Source-based Brute Force Protection area, supply the source-based protection settings.

   This area is available only when **Configuration Support** is set to **Current**.

   a) In the **Detection Period** setting, type the number of minutes the detection period should last.
   b) In the **Maximum Prevention Duration** setting, type the number of minutes the prevention period should last.
   c) For each of the other settings in this section, set the trigger and the action:
      - In the **Trigger** setting, specify when the trigger for the action occurs by selecting either **Never** or **After** a specified value is reached.
      - For the **Action** setting, select the action that occurs when the trigger is reached.
6. In the Distributed Brute Force Protection area, supply the distributed protection settings.

   This area is available only when **Configuration Support** is set to **Current**.

   a) In the **Detection Period** setting, type the number of minutes for detection.
   b) In the **Maximum Prevention Duration** setting, type the number of minutes for maximum prevention duration.
   c) In the **Detect Distributed Attack** setting, select when the distributed attack detection occurs.
      - Select **Never** to have no distributed brute force attack protection.
      - Select **After x failed login attempts** to have distributed brute force attacks detected if x failed logins are detected within the **Detection Period** configured previously.

   d) In the **Detect Credential Stuffing** setting, select when the detection should occur.

- Select **Never** to have no credential stuffing detection.
- Select **After x login attempts that match stole credentials dictionary** to have it reported when the configured conditions are met.

e) In the **Mitigation** setting, select the distributed brute force protection mitigation option to use.

7. In the Session-based Brute Force Protection area, supply the session-based protection settings.

   This area is available only when the **Configuration Support** setting is set to **13.0 And Prior**.

   - In the **Login Attempts from the Same Client** setting, type the number of attempts to allow.
   - In the **Re-enable Login After** setting, type the number of seconds.
   - In the **Use Device ID** setting, specify whether it is enabled.
8. In the Dynamic Brute Force Protection area, supply the dynamic protection settings.

   This area is available only when the **Configuration Support** setting is set to **13.0 And Prior**.

   - For the **Operation Mode** setting, select one of the modes: **Off**, **Alarm**, or **Alarm and Block**.
   - In the **Measurement Period** field, type the number of seconds.
   - In the **Detection Criteria** field, type the values that define when a problem is detected.
   - For the **Prevention Policy** setting, select one or of the options to use for the policy. When **Source IP-Based Client Side Integrity Defense** is selected, the **Suspicious Criteria (per IP address)** setting is displayed and can be modified.
   - In the **Suspicious Criteria (per IP address)** setting, type the values that define when failed login attempts become suspicious.
   - In the **Prevention Duration** setting, select the duration. This setting is displayed only when **Source IP-Based Client Side Integrity Defense** is selected in the **Prevention Policy** setting.
     - To have no limit on the duration, select **Unlimited**.
     - To have a maximum duration, select **Maximum** and type a value for the number of seconds.
9. Save your work.

Add methods

In the application security policy, you can specify methods that 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.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the policy name, and then click **HEADERS** > **Methods**.
3. Click **Add** to add a method.
4. From the **Method** list, select a method.
5. When you are finished, click **Save**.
   The new method is added to the list on the Methods screen. The method appears in blue, meaning that you can edit it. The check box to the left indicates that you can also delete it.

The system updates the policy to use the new methods.

Add or edit HTTP header settings

In the application security policy, you can specify a list of HTTP request headers that other web applications hosted in different domains can use when requesting this URL.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the policy name, and then on the left, click **HEADERS** > **HTTP Headers**.
3. Select whether to add a new HTTP header or view or modify an existing HTTP header.

- Click **Add** to add a new header.
- Click the name of the header to view or modify the properties.

Only HTTP headers that are displayed in blue can be modified or viewed.

4. Provide or modify the header settings as appropriate.

   - **Name**. When adding a new header, select the name of the HTTP header from the list. When modifying a header, the name cannot be changed.
   - **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 this is 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**. Specifies whether the system normalizes headers. Select the options for which type of normalization the system should perform on headers. There is a performance trade-off when using normalization, so use it only when needed.
     - **Percent Decoding**: Specifies, when enabled, that the system performs the following actions on header content: `%XX` and `%uXX`, bad unescaping, Apache whitespace, IIS Unicode codepoints, and plus to space.
     - **URL Normalization and Percent Decoding:** Specifies, when enabled, that the system performs the these actions on header content: multiple slashes, directory traversal, backslash replacement, and path parameter removal, and all **Percent Decoding** checks.
     - **HTML Normalization:** Specifies, when enabled, that the system performs the following actions on header content: removes all non-printables, whitespaces and the "+" character, skips comments, decodes HTML entities, performs hex decoding, decimal decoding, 0xXX decoding, style sheet escaping, and removes backslashes.
   - **Evasion Techniques Violations**. Specifies, when enabled, that the system logs and/or suggests learning suggestions for evasion violations detected during the normalization process if there are problems during the normalization of the specific header. The default is disabled.
   - **Overridden Security Policy Settings**. If used, select the signature override from the list and then enable or disable it by clicking **Enabled** or **Disabled**.
5. When you are finished, click **Save**.

The system updates the policy to use the new settings.

Edit host name settings

You can review, add, and delete host names from the policy using the Host Names screen. This list of host names is used by several features of the application security policy.

1. Navigate to the Host Names screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Then click the name of the appropriate policy, and on the left click **HEADERS** > **Host Names**.
3. Review the list of host names.

   If no host names are listed, you can add them by clicking the **Add**.
4. To modify a host name, click the name of the host name.

   The Host Name properties screen opens.
5. Review the Host Name.
6. To allow users to be redirected to a sub-domain of this host name, select the **Include Sub-domains** check box.
7. To set the policy to transparent mode and forward all responses, select the check box for **Policy is always transparent for this host**.

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

The host name settings for the policy are updated.

Add or edit cookie settings

You can review, add, and remove cookies from a policy, and re-order cookie wildcards using the Cookies screen. You use the same process to modify or add a cookie. The only difference is that when you modify a cookie, the **Cookie Name** properties already exist and you cannot change them.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left click **HEADERS** > **Cookies**.

   The screen displays a list of cookies.
3. To add a new cookie, click **Add**, or click a cookie name to modify an existing cookie.

   You use the same process to modify or add a cookie. However, you can specify some properties only when adding a cookie, and not modifying an existing cookie.
4. Type or review the **Cookie Name**, and specify whether it is **Explicit** or is a **Wildcard** expression.

   You can specify a cookie name only when adding a new cookie.
5. Specify the **Cookie Type**:

   - Select **Allowed** to indicate the client may change the cookie.
   - Select **Enforced** to indicate that the cookie cannot be changed by the client.

   Allowed provides additional options.
6. Select the settings for the cookie.

   - For **Perform Staging**, select the **Enabled** check box to indicate that the cookie is placed in staging.
   - For **Insert HTTPOnly attribute**, select the check box to insert the attribute in the domain cookie response header.
   - For **Insert SameSite attribute**, specify whether the attribute should be set to **None**, **Strict**, or **Lax**. Only **None** can be selected for BIG-IP® devices earlier than version 13.1.
   - For **Insert Secure attribute**, select the check box to insert the attribute into the domain cookie response header.
   - For **Base64 Decoding**, select the check box to enable decoding of Base64 strings. (This setting is displayed only if the **Cookie Type** is set to **Allowed**.)
   - For **Attack Signatures Check**, select the check box to verify attack signatures and display attack signature override settings. (This setting is displayed only if the **Cookie Type** is set to **Allowed**.)
   - For **Attack Signature Overrides**, select a signature from the list, and then click **Enabled** or **Disabled** to indicate whether each signature should be overridden.
7. To remove a cookie from staging, select the check box for the cookie and click **Enforce Selected**.
8. To filter the list of cookies by their enforcement readiness, select an option from the **Enforcement Readiness**setting.

   Enforcement readiness is the state of enforcement for each cookie, such as not enforced,, has a suggestion, or is ready to be enforced.

   - To list all cookies, select **All**.
   - To list cookies that have one or more suggestions, select **Has suggestion**.
   - To list cookies that are not being enforced, select **Not enforced**.
   - To list cookies that are ready to be enforced, select **Ready to be enforced**.
9. Click **Save** to save your changes.

The cookie settings for the policy are updated.

Edit redirection protection settings

You can enable redirection protection and list those domains that are allowed by your security policy, using the Redirection Protection screen. By enabling redirection protection, you can help prevent users from being redirected to questionable, phishing, or malware websites.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left click **HEADERS** > **Redirection Protection**.
3. For the **Redirection Protection** setting, select the **Enabled** check box.
   The screen displays other property settings.
4. For **Domain Name**, type the domain name that is allowed by the security policy.
5. To have the security policy also allow sub-domains of the domain, select the **Include Sub-Domains** check box.
6. To add the domain to the **Allowed Redirection Domains** list, click **Add**.
7. To delete a domain from the **Allowed Redirection Domains**, click the **X** to the left of that domain name.
   The domain is removed without confirmation.
8. Save your work.

Edit header character set settings

You can configure the security policy to allow or disallow certain characters in the value field of an HTTP header and in uncommon header names.

1. Navigate to the Character Set screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, expand **HEADERS** and click **Character Set**.
3. Review the list of characters, and for each, determine whether it should be allowed.

   You can use the View options to select which group of characters are displayed.

   - To allow characters in a header, select the check box in the **Allowed** column of the table row .
   - For characters that should not be allowed in a header, clear the check box in the **Allowed** column of the table row.
4. Click **Save** to save your changes.

Edit IP addresses list settings

You can view and edit configured IP address exceptions and characteristics.

1. Navigate to the IP Address screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Select a policy name, expand **IP ADDRESSES**, and select **IP Addresses List**.
3. Click **Add**.
4. Type an **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.
5. Type a **Netmask**.

   If you omit the netmask value, the system uses a default value of `255.255.255.255`.
6. Select whichever of the following options should be enabled.

- Select the **Policy Builder Trusted IP** check box 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.
- Select the **Ignore in Anomaly Detection and do not Collect Device ID** check box 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.
- Select the **Ignore in Learning Suggestion** check box to specify that the system not generate learning suggestions from traffic sent from this IP address.
- Select the **Never log traffic from this IP Address** check box 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.
- Select the **Ignore IP Address Intelligence** check box 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.

7. In the **Block this IP Address** setting, select one of the blocking options.

   - Select **Policy Default** to use the policy blocking settings.
   - Select **Never Block This IP** to not block this IP address.
   - Select **Always Block This IP** to block this IP address.

   If **Always Block This IP** is selected, many of the options become invalid and are removed from the screen.
8. Type a brief description for the IP address.
9. 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® Centralized Management system.

Edit IP address intelligence settings

You can review and modify IP address intelligence settings. An *IP intelligence database* is a list of IP addresses with questionable reputations. Refer to the ASM documentation or online help for more information on IP address intelligence.

1. Navigate to the IP Address screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Select a policy name, expand **IP ADDRESSES**, and select **IP Address Intelligence**.
3. Select the **IP Address Intelligence Enabled** check box.
   Other properties display; you can review the descriptions of the properties for additional information.
4. For the **IP Address White List** setting, type the IP address and subnet mask for each IP address that should be whitelisted, and click **Add** after each addition.
5. In the **IP Address Intelligence Categories** area, specify the categories that you want to alarm or block.

   - Select the **Alarm** check box to specify that whenever a request is sent from a source IP address that matches the category, the system logs the IP Intelligence data.
   - Select the **Block** check box to specify that the system stops requests sent from a source IP address that matches the category.

   Note: *In order for the system to block requests, the security policy must be in Blocking mode.*
6. Click **Save** when you are done.

The IP address intelligence settings are updated.

Add or edit HTTP URL settings

You can view, add, modify, and remove HTTP URLs that are either allowed or disallowed in an application security policy. *Allowed URLs* are URLs that the security policy accepts in traffic to the web application being protected. *Disallowed URLs* are URLs that the security policy denies.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, then on the left expand **URLs**, and click **HTTP**.

   The screen displays the list of HTTP URLs. You can add, delete, or reorder the HTTP URLs that are allowed or disallowed.
3. To add an allowed or disallowed HTTP URL to a policy, click **Add** for the allowed or disallowed list.

   Allowed HTTP URLs are listed at the top of the screen and disallowed HTTP URLs are listed at the bottom. The Add URL screen displays the properties, which differ between allowed URLs and disallowed URLs.

   - For disallowed HTTP URLs, specify whether the protocol is HTTP or HTTPS, and type the URL name.
   - For allowed HTTP URLs, specify whether the URL is explicit or a wildcard, whether the protocol is HTTP or HTTPS, and type the URL name or wildcard. Specify or modify additional properties for the allowed HTTP URL as needed.
4. Save your work.
5. To review or edit the properties of a URL, click the URL to open the Properties screen.

   Allowed URLs are listed in the Allowed URL column in the upper table of URLs. Disallowed URLs are listed in the Disallowed URL column in the bottom table of URLs.
6. To change the processing order of allowed URLS with the wildcard type, click **Wildcards Order**. The Wildcard Order screen opens, where you can move the wildcard entries in the list to change their sequence, and save your work.
7. To remove an HTTP URL from staging, select the check box for the HTTP URL and click **Enforce Selected**.
8. To filter the list of HTTP URLs by their enforcement readiness, select a value from the **Enforcement Readiness** setting.

   - To list all HTTP URLs, select **All**.
   - To list HTTP URLs that have one or more suggestions, select **Has suggestion**.
   - To list HTTP URLs that are not being enforced, select **Not enforced**.
   - To list HTTP URLs that are ready to be enforced, select **Ready to be enforced**.
9. To delete an allowed or disallowed HTTP URL from the policy, select the check box in the row for that HTTP URL and click **Delete** in the upper or lower portion of the screen, whichever is appropriate.

Add or edit WebSocket URL settings

You can view, add, modify, and remove WebSocket URLs that are either allowed or disallowed in an application security policy. *Allowed URLs* are URLs that the security policy accepts in traffic to the web application being protected. *Disallowed URLs* are URLs that the security policy denies.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, then on the left expand **URLs** and click **WebSocket**. The WebSocket URLs screen opens where you can add, or edit, WebSocket URLs.
3. To remove the WebSocket URL from staging, select the check box for the WebSocket URL and click **Enforce Selected**.

4. To edit the properties of a WebSocket URL, click the URL in either the **Allowed WebSocket URLs** or **Disallowed WebSocket URLs** column.
   The WebSocket URL properties screen opens, and you can change the properties (as described in the details for adding a URL of that type).
5. To add a WebSocket URL to a policy, determine whether it is an allowed or disallowed WebSocket URL.
   - To add an allowed WebSocket URL, click **Add** in the upper portion of the screen. This opens the Add Allowed WebSocket URL screen, where you can supply the needed properties.
   - To add a disallowed WebSocket URL, click **Add** in the lower portion of the screen. This opens the Add Disallowed WebSocket URL screen, where you can supply the needed properties.
6. For disallowed WebSocket URLs:
   a) Specify whether the protocol is **WS** or **WSS**.
   b) Type the URL name.
7. For allowed WebSocket URLs, supply the needed properties.
   a) In the Properties area, supply or modify the overall properties for the WebSocket URL.
   b) In the Message Handling area, supply or modify the message handling properties for the WebSocket URL.
   c) For wildcard URLs, expand the Meta Characters area to specify how meta characters are handled.
      - For **Check Signatures on this URL**, select the **Enabled** check box.
      - For **Check characters on this URL**, select the meta characters from the list and then click **Allow** or **Disallow** as needed.

   d) In the HTML5 Cross-Domain Request Enforcement area, supply or modify the HTML5 cross-domain request enforcement properties for the WebSocket URL.
8. To filter the list of WebSocket URLs by their enforcement readiness, select an option from the **Enforcement Readiness** list.
   - To list all WebSocket URLs, select **All**.
   - To list WebSocket URLs that have one or more suggestions, select **Has suggestion**.
   - To list WebSocket URLs that are not being enforced, select **Not enforced**.
   - To list WebSocket URLs that are ready to be enforced, select **Ready to be enforced**.
9. Save your work.

Edit URL character set settings

You can view and edit how the security policy responds to each character contained in a URL.

1. Navigate to the Character Sets URL screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, on the left expand **URLs**, and click **Character Sets**.
3. Review the list of characters, and for each, determine whether it should be allowed.

   You can use the View options to select which group of characters are displayed.
   - To allow characters in a URL, select the check box in the **Allowed** column of the table row.
   - For characters that should not be allowed in a URL, clear the check box in the **Allowed** column of the table row.
4. Click **Save** to save your changes.

Add or edit 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. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy to change, and on the left, click **File Types**.

   The screen displays a list of file types.
3. To remove the file type from staging, select the check box for the file type and click **Enforce Selected**.
4. To add a file type to the policy, click **Add** in either the Allowed File Types area at the top of the screen, or in the Disallowed File Types area at the bottom of the screen.
   - Use the Allowed File Types area to add file types that the security policy considers legal, and to view information about each file type.
   - Use the Disallowed File Types area to add file types 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.
5. If you chose to add Disallowed File Types, fill in the name.
6. If you chose to add Allowed File Types, fill in these settings.
   a) For **File type**, select whether the file type is a wildcard or is explicit, and type a wildcard name or an explicit name.
   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 Length**, 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.
7. To filter the list of file types by their enforcement readiness, select an option from the **Enforcement Readiness** setting.
   - To list all file types, select **All**.
   - To list file types that have one or more suggestions, select **Has suggestion**.
   - To list file types that are not being enforced, select **Not enforced**.
   - To list file types that are ready to be enforced, select **Ready to be enforced**.
8. When you are finished, save your work.

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® Centralized Management system.

Edit or add JSON content profile settings

You use JSON content profile properties to define what the application security policy enforces and considers legal when it detects traffic that contains JSON data.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to modify, then on the left expand **CONTENT PROFILES**, and click **JSON Profiles**.
3. Click the name of the JSON profile to modify, or click **Add** to create a new one.
4. Review the existing name, or type a **Profile Name** for the new profile.
5. Revise or type an optional **Description** for the profile.

6. In the **Maximum Total Length Of JSON Data** field, type or revise the longest length, in bytes, allowed by the security policy of the request payload, or parameter value, where the JSON data was found.

   To have no length restriction, you can leave this field blank.
7. In the **Maximum Value Length** field, type or revise the maximum acceptable length, in bytes, of the longest JSON element value in the document allowed by the security policy.

   To have no length restriction, you can leave this field blank.
8. For **Maximum Structure Depth**, type or revise the greatest nesting depth found in the JSON structure allowed by the security policy.

   To have no depth restriction, you can leave this field blank.
9. In the **Maximum Array Length** field, type or revise the largest number of elements allowed for arrays.

   To have no array length restriction, you can leave this field blank.
10. For **Tolerate JSON Parsing Warnings**, specify whether to enable response signature staging.
    - Select the **Enabled** check box to specify that the system does not report when the security enforcer encounters warnings while parsing JSON content.
    - Clear the check box to specify that the security policy reports when the security enforcer encounters warnings while parsing JSON content.
11. For **Parse Parameters**, specify whether to enable parameter parsing.
    - To enable parsing, select the **Enabled** check box.
    - When this setting is disabled, the system displays more main areas (such as Attack Signature Overrides, Meta Characters, and Sensitive Data Configuration) with additional properties for review and modification.
12. Expand the Attack Signatures Overrides area to select any signature overrides. (This area is displayed only when **Parse Parameters** is disabled.)
    - For the **Attack Signatures Check** setting, select the **Enabled** check box.
    - For the **Attack Signatures Overrides** setting, select the signature from the list and then click **Enabled** or **Disabled** as needed for that signature.
13. Expand the Meta Characters area to select how meta characters are handled. (This area is displayed only when **Parse Parameters** is disabled.)
    - For the **Check Characters** setting, select the **Enabled** check box.
    - For the **Overrides** setting, select the meta characters from the list and then click **Allowed** or **Disallowed** as needed.
14. Expand the Sensitive Data Configuration area to select how sensitive data is handled. (This area is displayed only when **Parse Parameters** is disabled.)

    a) In the **Sensitive Data** setting, type an element name within the JSON data whose values the system should consider sensitive.

    b) Click **Add** to add the element name to the sensitive data list.
15. Click **Save** to save your changes.

Edit or add XML content profile settings

You use XML content profile properties to define what the application security policy enforces and considers legal when it detects traffic that contains XML data.

1. Navigate to the XML Profiles screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to work with, then, on the left, expand **CONTENT PROFILES**, and click **XML Profiles**.

3. Click the name of the XML profile to modify, or click **Add** to create a new one.
4. Review the existing name or type a **Profile Name** for the new profile.
5. Review, revise, or type an optional **Description** for the profile.
6. For the **Use XML Blocking Response Page** property, select the type of response page to send when the security policy blocks a client request that contains URL XML content that does not comply with the settings of this XML profile.
   - To have the system send an XML response page, select the **Enabled** check box.
   - To have the system send the default response page, do not select the **Enabled** check box.
7. To configure the validation and defense settings of an XML profile, expand the XML Firewall Configuration area and modify those settings as needed.
8. To configure the system to perform attack signature checks on the XML profile, expand the Attack Signatures area and modify those settings as needed.
9. To change the security policy settings for specific meta characters in XML values on the XML profile, expand the Meta Characters area and modify those settings as needed.
10. Expand the Sensitive Data Configuration area to program the system to mask sensitive data that appears in an XML document, as shown in the BIG-IP device configuration interface and internal Application Security logs.
11. Click **Save** to save your changes.

Edit or add plain text content profile settings

You use plain text content profile properties to define what the application security policy enforces and considers legal when it detects traffic that contains plain text data.

1. Navigate to the Plain Text Profiles screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to modify, at the left, expand **CONTENT PROFILES**, and click **Plain Text Profiles**.
3. Click the name of the plain text profile to modify, or click **Add** to create a new one.
4. Review the existing, or type a **Profile Name** for the new profile.
5. Review, revise, or type an optional **Description** for the profile.
6. In the **Maximum Total Length** field, type the longest length, in bytes, allowed by the security policy.

   You can leave this field blank to have no length restriction.
7. In the **Maximum Line Length** field, type the longest line length, in bytes, allowed by the security policy.

   You can leave this field blank to have no length restriction.
8. If you want the system to perform percent decoding, select the **Perform Percent Decoding Enabled** check box.
9. To configure attack signature overrides, expand Attack Signatures Overrides and supply the needed values.
   a) In the **Attack Signatures Check** setting, select the **Enabled** check box.
   b) In the **Attack Signatures Overrides** setting, select one or more attack signatures to override.
   c) For each attack signature, select whether the override is enabled or disabled.
10. To change the security policy settings for specific meta characters in values on the plain text profile, expand Meta Characters and supply the needed values.
    a) In the **Check Characters** setting, select the **Enabled** check box.
    b) In the **Overrides** setting, select one or more meta characters to override.
    c) For each meta character, select whether the override is allowed or disallowed.
11. Click **Save** to save your changes.

Edit character set JSON settings

You can configure the security policy to allow or disallow certain characters if they appear in JSON values.

1. Navigate to the JSON screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy you want to work with, and on the left expand **CONTENT PROFILES** and **CHARACTER SETS**, then click **JSON**.
3. Review the list of characters, and for each, determine whether it should be allowed or not.

   - To allow characters, select the check box in the Allowed column of the table row.
   - For characters that should not be allowed, clear the check box in the Allowed column of the table row.

   Use the View options to select which characters are displayed.

   - Click **All Characters** to display all characters.
   - Click **Allowed** to display only characters that are marked as allowed.
   - Click **Disallowed** to display only characters that are not allowed.
4. Click **Save** to save your changes.

Edit character set plain text settings

You can configure the security policy to allow or disallow certain characters if they appear in plain text values.

1. Navigate to the Plain Text screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left expand **CONTENT PROFILES** and **CHARACTER SETS**, then click **Plain Text**.
3. Review the list of characters, and for each, determine whether it should be allowed or not.

   - To allow characters, select the check box in the Allowed column of the table row.
   - For characters that should not be allowed, clear the check box in the Allowed column of the table row.

   Use the View options to select which characters are displayed.

   - Click **All Characters** to display all characters.
   - Click **Allowed** to display only characters that are marked as allowed.
   - Click **Disallowed** to display only characters that are not allowed.
4. Click **Save** to save your changes.

Edit character set XML settings

You can configure the security policy to allow or disallow certain characters if they appear in XML values.

1. Navigate to the XML screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left expand **CONTENT PROFILES** and **CHARACTER SETS**, then click **XML**.
3. Review the list of characters, and for each, determine whether it should be allowed or not.

- To allow characters, select the check box in the Allowed column of the table row .
- For characters that should not be allowed, clear the check box in the Allowed column of the table row.

Use the View options to select which characters are displayed.

- Click **All Characters** to display all characters.
- Click **Allowed** to display only characters that are marked as allowed.
- Click **Disallowed** to display only characters that are not allowed.

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

Add or edit parameter settings

You can add or edit settings for parameters that the security policy permits in requests, such as the parameter type and whether the parameter is allowed to contain an empty value. The default parameter is displayed for all policies, and can be edited. It is indicated by ***** (asterisk).

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click a policy name, and on the left, click **PARAMETERS** > **Parameters**.
3. You can add a new, or edit an existing, parameter.
   - To add a new parameter, click **Add**.
   - To edit an existing parameter, click the parameter name.

   The properties screen opens for the new or existing parameter.
4. To remove the parameter from staging, select the check box for the parameter and click **Enforce Selected**.
5. For a new parameter, for the **Name** setting, select the type, and then type a name for the new parameter.
   - Select **explicit** if this is a regular named parameter.
   - Select **wildcard** if any parameter name that matches the wildcard expression is permitted by the security policy. (For example, typing the wildcard ***** specifies that the security policy allows every parameter.) The syntax for wildcard entities is based on shell-style wildcard characters.
   - Select **no name** if this parameter does not have a name. The system automatically names the parameter `no name` and it behaves the same as an explicit parameter.

   The name setting cannot be changed once the parameter is created.
6. For **Level**, select the level of parameters to be displayed.
   - Select **global** to display global parameters not associated with flows or URLs.
   - Select **URL** to display parameters associated with flows or URLs, select **HTTP** or **HTTPS** as the protocol, and then select the URL.

   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.
7. To enable or allow any of these settings, click the **Enabled** check box for the setting:
   - Select **Perform Staging** to display the staging status on this parameter.
   - Select **Allow Empty Value** to allow empty values.
   - Select **Allow Repeated Occurrences** to allow repeated occurrences.
   - Select **Sensitive Parameter** to, in a validated request, protect sensitive user input, such as a password or a credit card number. The contents of sensitive parameters are not visible in logs or in the user interface.
8. Specify the **Value type** for the parameter.

The value type you specify might display additional fields. You cannot change the value type after it is created.

- Select **dynamic-content** for parameters whose data is dynamic.
- Select **ignore** for parameters whose values the system does not check.
- Select **json** for JSON parameters fetched from the server that are not editable.
- Select **static-content** for parameters whose data is static. In the Parameter Static values area displayed at the bottom of the screen, supply a value in the **Add New Value** setting, and click **Add**. Add or subtract values as needed.
- Select **user-input** for parameters whose data is provided by user-input. Use the **Data type** setting to provided additional information about the user input.
- Select **xml** for XML parameters fetched from the server that are not editable. In the XML Profile area displayed at the bottom of the page, select an XML profile.

9. For the **Data type** setting, select the data type to use for the user input.

   - Select **email** to specify that the data must be text in email format only. In the Data type attributes area, specify a value for the **Maximum Length** setting in bytes.
   - Select **alpha-numeric** to specify that the data can be any text consisting of letters, digits, and the underscore character.

     - In the Data type attributes area, specify a value for the **Maximum Length** setting in bytes, and select whether to enable regular expressions or Base64 encoding. When the **Regular Exp** setting is enabled, it specifies that the parameter value includes the specified parameter pattern. This is a positive regular expression that defines what is legal.
     - In the Value Meta Character area, select the **Enabled** check box and then select which meta character to allow or disallow as a value.
     - In the Attack Signatures area, select the **Enabled** check box and then select which attack signature overrides to enable or disable.
   - Select **integer** to specify that the data must be whole numbers only (no decimals). In the Data type attribute area, specify values for the **Minimum Value** , **Maximum Value**, and **Maximum Length** settings.
   - Select **decimal** to specify that the data is numbers only and can include decimals. In the Data type attributes area, specify values for the **Minimum Value** , **Maximum Value**, and **Maximum Length** settings.
   - Select **phone** to specify that the data can be text in telephone number format only. In the Data type attributes area, specify a value for the **Maximum Length** setting.
   - Select **file upload** to specify there is no text limit for the data (length checks only). In the Data type attributes area, specify a value for the **Maximum Length** setting, and specify whether to disallow file uploading or enable Base64 encoding.
10. To filter the list of parameters by their enforcement readiness, select an option from the **Enforcement Readiness** setting.

    - To list all parameters, select **All**.
    - To list parameters that have one or more suggestions, select **Has suggestion**.
    - To list parameters that are not being enforced, select **Not enforced**.
    - To list parameters that are ready to be enforced, select **Ready to be enforced**.
11. When you are finished, save your work.

The application security policy is updated to use the new settings.

Add or edit extraction settings

You use extraction settings to manage how the system extracts dynamic values for dynamic parameters from the responses returned by the web application server. An *extraction* is a subcollection that isolates a parameter from an object. Other subcollections (such as parameters) reference extractions by name (not by URL).

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies.**
2. Click the name of the policy and then on the left, click **PARAMETERS** > **Extractions**.
3. You can add a new or edit an existing extraction.
   - To add a new extraction, click **Add**.
   - To edit an existing extraction, click the extraction name.

   The properties screen opens for the new or existing extraction.
4. For a new extraction, specify the **Name** of the dynamic parameter for which the system extracts values from responses.
   - For a named parameter, select **New** and type the name in the field.
   - For the `UNNAMED` parameter, select **no name**.

   The name setting cannot be changed once the extraction is created.
5. 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 **Add**. • **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 field, and click **Add**. If you enter a URL that does not yet exist in the security policy, the URL is added to the security policy. • **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 field.
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.

6. In the Extracted Method Configuration area, specify the methods by 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.
Search Entire Form	Specifies, when checked (enabled), that the system searches for dynamic parameter values in the entire form found on a web page.
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.

Option	Description
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** field.
Search Response Body	Specifies, when checked (enabled), that the system searches for dynamic parameter values in the body of the response. Use the additional options to further refine the search. You can specify one or more of the following options, but you must specify the RegEx 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.

7. When you are finished, save your work.

The application security policy is updated to use the new settings.

Edit character set parameter name settings

You use character set parameter name settings in the security policy to allow or disallow certain characters in parameter names.

1. Go to the Policies screen: Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Continue to the parameter name screen: Click the name of the policy and then, on the left, click **PARAMETERS** > **CHARACTER SETS** > **Parameter Name**.
3. Review the list of characters, and for each, determine whether it should be allowed or not.
   - Select the **Allowed** check box for characters that should be allowed.
   - Clear the **Allowed** check box for characters that should not be allowed.

   Use the View options to select which characters are displayed.

   - Click **All Characters** to display all characters.
   - Click **Allowed** to display only characters that are marked as allowed.
   - Click **Disallowed** to display only characters that are not allowed.
4. Click **Save** to save your changes.

The system updates the security policy to use the new character set parameter name settings.

Edit character set parameter value settings

You use character set parameter value settings in the security policy to determine whether the security policy allows those values in a request.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.

2. Click the name of the policy and then, on the left, click **PARAMETERS** > **CHARACTER SETS** > **Parameter Value**.
3. Review the list of characters, and for each, determine whether it should be allowed or not.

   - Select the **Allowed** check box for characters that should be allowed.
   - Clear the **Allowed** check box for characters that should not be allowed.

   Use the View options to select which characters are displayed.

   - Click **All Characters** to display all characters.
   - Click **Allowed** to display only characters that are marked as allowed.
   - Click **Disallowed** to display only characters that are not allowed.
4. Click **Save** to save your changes.

The system updates the security policy to use the new character set parameter value settings.

Add sensitive parameters settings

You can add and delete sensitive parameters used by your security policy. Some requests include sensitive data, such as account numbers, in parameters. If you create sensitive parameters, the data in those parameters is replaced with asterisks (***) in the stored request and in logs.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left click **PARAMETERS** > **Sensitive Parameters**.
3. Click **Add** to add a sensitive parameter.

   The Sensitive Parameter properties screen opens.
4. In the **Name** setting, type the name of the sensitive parameter.
5. Save your work.

Configure attack signatures

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. Go to the Policies screen: Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Continue to the Attack Signatures Configuration screen: Click the name of a policy, and on the left click **Attack Signatures Configuration** .
3. Revise the settings as needed.

   - To enable staging of signatures, select the **Signature Staging Enabled** check box.
   - To place updated signatures in staging, select the **Place updated signatures in staging Enabled** check box. New signatures are always placed in staging, regardless of this setting.
   - For **Attack Signature Set Assignment**, select one or more signature sets from the list to be assigned to the policy, and then select the appropriate options for that signature set.

     - Select or clear the **Learn**, **Alarm**, and **Block** options for each signature set.

       - Select **Learn** to have the security policy learn all requests that match enabled signatures in the signature set.
       - Select **Alarm** to have the security policy logs the request data if a request matches a signature in the signature set.

        - Select **Block**, to have the security policy block all requests that match a signature included in the signature set.
    - From the **Actions** list, select, if needed, whether to enable or enforce signatures in the signature set.
  - For **Apply Response Signatures**, select a file type, if needed. The default wildcard character indicates all file types.
4. When you are finished, save your work.

The system updates the application security policy attack signatures settings.

View and modify attack signatures

You can view the list of attack signatures that belong to signature sets assigned to the policy, and specify whether they are enabled or in staging.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of a policy, and on the left click **Attack Signatures**.
3. To restrict the number of signatures displayed, use the filter field at the upper right of the screen.

   You can select both basic and advanced filter options by clicking the arrow to the left of the field.
4. To specify whether or not the attack signature is enabled, select the check box in the Enabled column of the table for that row.
5. To have an attack signature placed in staging, select the check box in the In Staging column of the table for that row.
6. When you are finished, save your work.

The system updates any modified attack signature settings.

Edit geolocation enforcement settings

You use geolocation enforcement to select which geolocations the policy does not allow.

1. Navigate to the Geolocation Enforcement screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left, click **Geolocation Enforcement**.
3. Select a geolocation that is not allowed by the policy from the **Disallowed Geolocations** list.
   Once you have selected the geolocation, it is listed below the drop-down list.
4. You remove a selected geolocation from the list by clicking the **X** to the left of the geolocation name.
5. Click **Save** to save your changes.

The system updates the list of geolocations that the policy does not allow.

Add or edit login page settings

You can view and manage login page settings for the security policy to better protect the login page URLs used by your web applications.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy to manage, and on the left click **SESSIONS AND LOGINS** > **Login Pages**.
3. You can add new, or edit existing login page settings.

   - Click **Add** to add a login page and settings.
   - Click the name of the login page to edit the settings.

   The Login Page Properties screen opens.

4. In the **Login URL** setting, select the appropriate options for the URL.
   a) Specify whether the URL uses wildcards or is explicitly named. Select **Wildcard** or **Explicit**.
   b) Specify the URL protocol. Select **HTTP** or **HTTPS**.
   c) Select the URL to use, or select **Custom URL** and specify the URL.
5. In the **Authentication Type** setting, select the type of authentication to use.
6. In the Access Validation area, specify how the login page should be validated by typing one or more setting values.

   You define validation criteria on the response of the login URL. You must configure at least one of the validation criteria. If you configure more than one validation criteria, then all the criteria must be fulfilled in order to access the authenticated URL.
7. Save your work.

Add or edit logout page settings

You can view and manage logout page settings for the security policy to better protect the logout page URLs used by your web applications.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left click **SESSIONS AND LOGINS** > **Logout Pages**.
3. Specify whether you are adding or editing logout page settings.
   - Click **Add** to add a logout page and settings.
   - Click the name of the logout page to edit the settings.

   The Logout Page Properties screen opens.
4. In the **Logout URL (explicit only)** setting, select the appropriate options for the URL.
   a) Specify the URL protocol. Select **HTTP** or **HTTPS**.
   b) Select the URL to use, or select **Custom URL** and specify the URL.
5. In the **A string that should appear in the response** setting, type a string that should appear in the request (either the query string or in its payload) to indicate that the request is a logout request.
6. In the **A string that should NOT appear in the response** setting, type a string that should not appear in the request (either the query string or in its payload) to indicate that the request is a logout request.
7. Save your work.

Add or edit login enforcement settings

You can add and modify login enforcement properties. Login enforcement specifies the authenticated login URLs and logout URLs for the web application.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the appropriate policy, and on the left click **SESSIONS AND LOGINS** > **Login Enforcements**.
3. For the **Expiration Time** setting, specify whether you want the login session to expire.
   - If you do not want the login session to expire, click **Disabled**.
   - If you want the login URL to be valid for a limited time, click the button to the left of the **Seconds** field, and type a value, in seconds (1-99999), that indicates how long the session will last. The login session ends after the number of seconds has passed.
4. For the **Authenticated URLs** setting, specify the target URLs that users can access only by using the login URL.
   a) In the provided field, type the target URL name in the format `/private.php`.

Wildcards are allowed.

b) Click **Add** to add the URL to the list of authenticated URLs.

c) Repeat to add as many authenticated URLs as needed.

You can remove a URL from the list of authenticated URLs by clicking **X**.

5. Save your work.

Edit session tracking settings

You can enable session hijacking and session tracking to track, enforce, and report on user sessions and IP addresses.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Click the name of the policy to work on, and on the left click **SESSIONS AND LOGINS** > **Session Tracking**.
3. To enable session hijack detection, for the **Detect Session Hijacking by Device ID Tracking** setting, select the **Enabled** check box.

   Review the notes displayed.
4. To configure session tracking, supply values for the following settings.

   a) Select the **Session Awareness Enabled** check box.

   b) For the **Application Username** setting, select the form of the username.

   - To use no application username, select **None**.
   - To use APM usernames and session IDs, select **Use APM Usernames and Session ID**.
   - To use individual login pages, select **Use Individual Login Pages** and then select the login page in the area provided.
   - To use all login pages, select **Use All Login Pages**.
5. To configure violation detection actions, specify additional settings.

   a) For **Track Violations and Perform Actions**, select the **Enabled** check box.

   b) For **Violation Detection Period**, type the number of seconds for the detection period.
6. In the Block All area, specify how the system performs when the Block All action is triggered.
7. In the Log All Requests area, specify how the system performs when the Log All Requests action is triggered.
8. In the Delay Blocking area, specify how the system performs when the Delay Blocking action is triggered.
9. Save your work.

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.

Import application 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 Web Application Security to import security policies that were previously exported from another BIG-IP® Application Security Manager™ (ASM) system.

1. Navigate to the Policies screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. On the Policies screen, click the **Import** button.
3. On the Import Policy screen, select the security policy file by clicking **Choose File** , and navigating to the file location.

   You can also drag and drop a file to the **Drop Policy File Here** field.
4. Select a policy name for the imported policy if you like.
5. Click **Import**.

After you have imported the policy, the system lists it in the Policies screen. The uploaded policy will have the same name as the XML file, unless you provided a policy name.

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.

Export application security policies

You can use Web Application Security to export security policies. You can use the exported security policy as backup, or you can import it onto another system.

1. Navigate to the Policies screen: click **Configuration** > **SECURITY** > **Web Application Security** > **Policies**.
2. Select the check box to the left of the security policy you want to export.
   The **Export** button becomes active.
3. Click the **Export** button to show a list, and select the BIG-IP version to use when exporting this security policy.

The policy is exported.

You can use the exported security policy as a 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 Web Application Security

Through 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. By managing signature files from the BIG-IQ® Centralized Management 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.

Web Application Security fetches all new and relevant signature files from an external server, which may use a proxy. You can configure a proxy from the BIG-IQ Centralized Management system (**System** > **PROXIES**). The BIG-IQ Centralized Management system can then push the signature files to the relevant BIG-IP device or devices. It displays the signature version for each device.

Web Application Security signature file processing, such as importing, downloading, installing (pushing to devices), and deleting signature files, requires the following built-in roles, or the equivalent permissions on a custom role: Administrator, Security Manager, or Web App Security Manager.

View and install individual signature files

Before you start this task, make sure that your current BIG-IQ® Centralized Management account has Administrator, Security Manager, or Web App Security Manager credentials, or a custom role with equivalent permissions. These permissions are required for importing, downloading, and installing signature files.

You can edit and install individual signature files with the Signature Files screen.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Signature Files** > **Signature Files List**.
2. To view and install a signature file, click the file name.

   The signature file properties screen opens.
3. Review the information about the signature file in the read only fields.
   - The **Name** setting displays the name of the signature file.
   - The **Version** setting displays the version of the signature file.
   - The **Compatibility** setting displays the BIG-IP device version that should be used with this signature file.
   - The **Source** setting displays the source of the signature file.
4. In the **Install to Devices** setting, specify which BIG-IP devices should receive the signature file by moving them from the **Available Devices** list to the **Selected Devices** list.
5. In the **Install To** setting, specify which grouping of BIG-IP devices should receive the signature file.
   - Select **All Devices** to install the signature file to all listed BIG-IP devices.
   - Select **Active Devices Only** to install the signature file to all listed BIG-IP devices, except for those devices that are the inactive members of a cluster.

Once a signature file is deployed to an active clustered BIG-IP device, a synchronization task will run on the BIG-IP device cluster.

6. In the Related Devices area, review the BIG-IP devices listed.
7. Expand the Readme area to view details about the changes to the signature file.
8. Click **Install** to have the signature file installed on the selected BIG-IP devices.

   Or you can click **Cancel** to remove any changes and not install the signature file.

Update and install all signature files

Before you start this task, make sure that your current BIG-IQ® Centralized Management account has Administrator, Security Manager, or Web App Security Manager credentials, or a custom role with equivalent permissions. These permissions are required for importing, downloading, and, installing signature files.

You can schedule signature file updates and installations for all signature files, using the Settings screen.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Signature Files** > **Signature Files List**.
2. Click **Settings**.

   The Settings screen opens.
3. For the **Remote Updates** setting, select **Enabled** to allow remote signature file updates.

   If this setting is disabled, the other settings are not displayed.
4. In the **Interval** setting, select how often the scheduled update should run.
5. For the **Starting at** setting, specify when the scheduled update and installation should begin.

   You must specify a day after the current day.
6. Review when the **Last Update** occurred.
7. In the **Proxy** setting, select the proxy to use when retrieving signature files, or select **None**.

   You can configure proxies from the BIG-IQ Centralized Management system (**System** > **PROXIES**).
8. In the **Next Update** setting, review when the next update is schedules to occur.
9. In the **Last Run Status** setting, review the status of the last file update.

   Possible statuses include: `Passed`, or `Failed`.
10. For the **Install To** setting, specify which grouping of BIG-IP devices should receive the signature file.

    - Select **All Devices** to install the signature file to all listed BIG-IP devices.
    - Select **Active Devices Only** to install the signature file to all listed BIG-IP devices, except for those devices that are the inactive members of a cluster.

    Once a signature file is deployed to an active clustered BIG-IP device, a synchronization task will run on the BIG-IP device cluster.
11. Save your work.

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 Web Application Security

About virtual servers in Web Application Security

Web Application Security displays virtual servers 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, Web Application Security creates an extra virtual server to hold all policies not related to any virtual server in the discovered device. You will see the IP address of this server expressed as dashes (----) in the virtual server list. All security policies that are inactive in the BIG-IP device, or are not attached to any virtual servers in the BIG-IP device, are assigned to this virtual server and will be deployed to the associated BIG-IP device without attaching it to any other virtual server. This is the only virtual server that is allowed to have multiple, editable policies assigned to it. BIG-IQ Centralized Management cannot assign policies to both active and inactive virtual servers on the same BIG-IP device.

Attaching Web Application Security policies to virtual servers

You can view virtual server properties and attach Web Application Security policies to virtual servers.

1. Click **Configuration** > **SECURITY** > **Web Application Security** > **Virtual Servers**.

   The screen shows a list of the virtual servers that can be used with Web Application Security policies.
2. Click the name of the virtual server to view properties or to manage the policy attached to the virtual server.

   The following properties cannot be changed from this screen.

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.

3. To change the policy attached to the virtual server, use the **Attached Policy** setting.
   - To attach a policy to the virtual server, select the policy from the list.
   - To remove the attached policy from the virtual server, click the **X** to the left of the policy name.
4. If you changed which policy is attached to the virtual server, save your work.

Security Deployment Best Practices

Understanding roles required for deploying security policies

When you want to deploy a Web Application Security configuration, or a Network Security configuration, you need to use one of the following built-in roles.

- To deploy Network Security or Web Application Security configurations, use an account with the Security Manager role.
- To deploy only Network Security configurations, use an account with the Security Manager, Network Security Manager, or Network Security Deployer roles.
- To deploy only Web Application Security configurations, use an account with the Security Manager, Web App Security Manager, or Web App Security Deployer roles.

For more information on roles, refer to the role descriptions on the Roles screen (**System** > **ROLE MANAGEMENT** > **Roles**) or refer to *F5® BIG-IQ® Centralized Management: Authentication, Roles, and User Management* on `support.f5.com`.

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.

Legal Notices

Legal notices

Publication Date

This document was published on December 29, 2017.

Publication Number

MAN-0520-08

Copyright

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

Trademarks

For a current list of F5 trademarks and service marks, see *http://www.f5.com/about/guidelines-policies/trademarks*.

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

Patents

This product may be protected by one or more patents indicated at: *https://f5.com/about-us/policies/patents*.

Link Controller Availability

This product is not currently available in the U.S.

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