

BIG-IP[®] Acceleration: Implementations

Version 13.0.0



Table of Contents

Using Acceleration Policies to Manage and Respond to HTTP Requests.....	9
Overview: Acceleration policies.....	9
Task summary for using acceleration policies to manage and respond to HTTP requests.....	9
Accessing the Policy Viewer screen.....	9
Copying an acceleration policy.....	10
Creating a user-defined acceleration policy from a predefined acceleration policy.....	10
Creating a new user-defined acceleration policy.....	10
Importing an acceleration policy.....	11
Publishing a user-defined acceleration policy.....	11
Modifying an acceleration policy's rules.....	11
Viewing rules for an acceleration policy.....	12
Saving an acceleration policy to an XML file.....	12
Deleting a user-defined acceleration policy.....	12
Overview: Policy Editor screen.....	13
Task summary for using the Policy Editor.....	13
Accessing the Policy Editor screen.....	13
Viewing the Policy Tree for an acceleration policy.....	14
Differentiating Requests and Responses with Variation Rules.....	14
Overview: Variation rules.....	14
Configuring variation rule settings.....	14
Configuring ambiguous query parameters as unnamed.....	15
Variation rule example.....	15
 Proxying Requests and Responses.....	 17
Overview: Proxying rules.....	17
Always proxying requests for a node.....	17
Configuring proxy rules for a node.....	18
Overriding proxy requests.....	19
About creating the example proxying rule.....	19
Configuring the example proxying rule.....	20
About configuring the proxying rule parameters example.....	20
About configuring the example proxy override rule.....	21
 Managing Requests and Responses with Lifetime Rules.....	 23
Overview: Lifetime rules.....	23
Configuring lifetime cache settings.....	23
An example lifetime rule.....	24
 Invalidating Cached Content.....	 27
Overview: Invalidating cached content for an application.....	27
Invalidating cached content for an application.....	27
Clearing cached content from the command line.....	27
Overview: Invalidating cached content for a node.....	27
Creating an Invalidations rule.....	28
Blog invalidations rule example.....	29
Creating leaf nodes for the blog example invalidations rule	30

Specifying matching rules for the invalidations example View node.....	30
Specifying matching rules for the invalidations example Post node.....	31
Task summary for specifying invalidations rules for the invalidations example Post node.....	32
Managing Object Types.....	35
Overview: Object classification.....	35
Task summary for managing object types.....	35
Creating a user-defined object type.....	35
Editing an object type.....	36
Deleting a user-defined object type.....	36
Caching Objects in a VIPRION Cluster.....	37
Overview: Acceleration in a cluster.....	37
Caching objects in a cluster or cluster member.....	37
Immediately Caching Dynamic Objects.....	39
Overview: Caching an object on first hit.....	39
Caching an object on first hit.....	39
Disabling the Caching an object on first hit setting.....	39
Accelerating Parallel HTTP Requests.....	41
Overview: HTTP request queuing.....	41
Enabling HTTP request queuing.....	42
Disabling HTTP request queuing.....	42
Managing HTTP Traffic with the HTTP/2 Profile.....	45
Overview: Managing HTTP traffic with the HTTP2 profile.....	45
About HTTP/2 profiles.....	46
HTTP/2 profile settings.....	46
Creating a pool to manage HTTPS traffic.....	47
Creating a virtual server to manage HTTP traffic.....	48
Creating an HTTP/2 profile.....	48
Creating a virtual server to manage HTTP/2 traffic.....	49
Accelerating Requests and Responses with Intelligent Browser Referencing.....	51
Overview: Reducing conditional GET requests with Intelligent Browser Referencing.....	51
Task summary for reducing conditional GET requests with Intelligent Browser Referencing.....	51
Configuring Intelligent Browser Referencing advanced settings.....	51
Enabling content assembly on proxies.....	52
Enabling Intelligent Browser Referencing.....	52
Adjusting the adaptive Intelligent Browser Referencing lifetime.....	53
Implementation result.....	54
Accelerating JavaScript and Cascading Style Sheet Files.....	55
Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files.....	55
Minifying cascading style sheet and JavaScript files.....	55
Task summary for prefetching cascading style sheet, JavaScript, and inline image files.....	55

Task summary for inlining cascading style sheet, JavaScript, and inline image files.....	57
Task summary for reordering cascading style sheet and JavaScript files.....	60
Task summary for concatenating cascading style sheet and JavaScript files.....	62
Implementation results.....	64
Establishing Additional TCP Connections with MultiConnect.....	65
Overview: Accelerating requests and responses with MultiConnect.....	65
Task summary for establishing additional TCP connections with MultiConnect.....	65
Enabling content assembly on proxies.....	65
Configuring DNS subdomains for use with MultiConnect.....	66
Enabling MultiConnect for HTTP traffic.....	67
Enabling MultiConnect for HTTPS traffic.....	67
Implementation result.....	68
Serving Specific Hyperlinked Content with Parameter Value Substitution.....	69
Overview: Serving specific hyperlinked content with parameter value substitution.....	69
Serving specific hyperlinked content with parameter value substitution.....	69
Accelerating Access to PDF Content.....	71
Overview: Accelerating access to PDF content with PDF linearization.....	71
Task summary for accelerating access to PDF content.....	71
Accelerating content with PDF linearization.....	71
Disabling PDF linearization for a specific node.....	72
Accelerating Images with Image Optimization.....	73
Overview: Accelerating images with image optimization.....	73
Task summary for optimizing images.....	73
Accelerating images by optimization.....	73
Disabling image optimization for a node.....	75
Accelerating Video Streams with Video Delivery Optimization.....	77
Overview: Optimizing video delivery.....	77
Task summary for optimizing video streams.....	77
Creating a video advertisement policy.....	77
Modifying a video advertisement policy.....	78
Deleting a video advertisement policy.....	78
Enabling video delivery optimization.....	78
Modifying video delivery optimization.....	79
Overview: Video Quality of Experience profile.....	80
Creating an iRule to collect video Quality of Experience scores.....	80
Creating an iRule to collect static information about video files.....	81
Creating a video Quality of Experience profile.....	82
Creating a pool.....	82
Creating a video Quality of Experience virtual server.....	82
Compressing Content from an Origin Web Server.....	85
Overview: Enabling content compression from an origin web server.....	85
Enabling content compression from an origin web server.....	85
Accelerating Responses with Metadata Cache Responses.....	87

Overview: Using Metadata cache responses to accelerate responses.....	87
Accelerating Metadata responses.....	87
Disabling Metadata responses	87
Accelerating Traffic with a Local Traffic Policy.....	89
About classifying types of HTTP traffic with a local traffic policy.....	89
Accelerating traffic for applications with a local traffic policy.....	89
Accelerating Traffic with Intelligent Client Cache.....	91
Overview: Accelerating traffic with Intelligent Client Cache.....	91
Accelerating traffic for HTML5-compliant browsers.....	91
Using the Request Logging Profile.....	93
Overview: Configuring a Request Logging profile.....	93
Creating a pool with request logging to manage HTTP traffic.....	93
Creating a request logging profile.....	93
Configuring a virtual server for request logging.....	95
Deleting a request logging profile.....	95
Request Logging profile settings.....	96
Request Logging parameters.....	97
Monitoring BIG-IP Acceleration Application Performance.....	101
Overview: Monitoring the performance of a BIG-IP acceleration application.....	101
Enabling performance monitoring for a BIG-IP application.....	101
Disabling performance monitoring for a BIG-IP application.....	101
Overview: Reporting ROI statistics.....	102
Creating an SMTP server configuration.....	102
Creating a pool.....	102
Overview: Configuring ROI reports.....	103
Modifying an application to export ROI statistics.....	109
Using Forward Error Correction to Mitigate Packet Loss.....	111
Overview: Using forward error correction (FEC) to mitigate packet loss.....	111
About forward error correction (FEC).....	111
Task summary.....	112
Customizing a FEC profile.....	112
Creating a FEC tunnel for receiving traffic.....	112
Creating a FEC tunnel for initiating traffic.....	113
Viewing FEC tunnel statistics.....	113
Managing Deduplication.....	115
Overview: Symmetric data deduplication.....	115
Task Summary.....	115
Enabling symmetric data deduplication.....	115
Disabling symmetric data deduplication.....	116
Clearing the deduplication cache.....	116
Managing the Settings for Subnet Discovery.....	119
Overview: Managing advertised routing.....	119
Task Summary.....	119
Adding a virtual server to advertised routes.....	119

Adding advertised routes manually.....	120
Modifying automatic discovery of advertised routes.....	120
Verifying subnet discovery.....	121
Managing Remote Endpoint Discovery.....	123
Overview: Managing dynamic discovery of remote endpoints.....	123
Task Summary.....	123
Verifying subnet discovery.....	123
Modifying dynamic discovery of remote endpoints.....	123
Setting Up an iSession Connection Using the Quick Start Screen.....	125
Overview: Setting up an iSession connection using the Quick Start screen.....	125
Setting up an iSession connection using the Quick Start screen.....	125
Forwarding Non-Optimized IP Traffic Through an IPsec Tunnel.....	127
Overview: Forwarding Non-Optimized IP traffic through an IPsec tunnel.....	127
Creating a virtual server for all IP iSession traffic.....	127
Adding compression to an IPsec policy.....	129
Securing an iSession Deployment.....	131
Overview: Securing an iSession deployment.....	131
Task summary.....	131
Generating and importing SSL certificates for a secure iSession connection.....	131
Customizing SSL profiles for a secure iSession connection.....	132
Configuring the remote endpoints for a secure iSession connection.....	134
Implementation result.....	137
Encrypting Application Traffic with iSession.....	139
Overview: Encrypting application traffic with iSession.....	139
Task summary for encrypting application traffic using IPsec.....	139
Encrypting application traffic using IPsec on the Quick Start screen.....	139
Creating a custom IPsec policy for iSession traffic.....	140
Legal Notices.....	143
Legal notices.....	143

Using Acceleration Policies to Manage and Respond to HTTP Requests

Overview: Acceleration policies

An *acceleration policy* is a collection of defined rule parameters that dictate how the BIG-IP® system handles HTTP requests and responses. The BIG-IP system uses two types of rules to manage content: matching rules and acceleration rules. *Matching rules* are used to classify requests by object type and match the request to a specific acceleration policy. Once matched to an acceleration policy, the BIG-IP system applies the associated *acceleration rules* to manage the requests and responses.

Depending on the application specific to your site, information in requests can sometimes imply one type of response (such as a file extension of `.jsp`), when the actual response is a bit different (like a simple document). For this reason, the BIG-IP system applies matching rules twice: once to the request, and a second time to the response. This means that a request and a response can match to different acceleration rules, but it ensures that the response is matched to the acceleration policy that is best suited to it.

Task summary for using acceleration policies to manage and respond to HTTP requests

Perform these tasks to use policies to manage and respond to HTTP requests.

Task summary

Accessing the Policy Viewer screen

Copying an acceleration policy

Creating a user-defined acceleration policy from a predefined acceleration policy

Creating a new user-defined acceleration policy

Importing an acceleration policy

Publishing a user-defined acceleration policy

Modifying an acceleration policy's rules

Viewing rules for an acceleration policy

Saving an acceleration policy to an XML file

Deleting a user-defined acceleration policy

Accessing the Policy Viewer screen

The Policy Viewer displays the matching rules and acceleration rules for predefined acceleration policies.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a predefined acceleration policy.

The Policy Viewer screen appears for the predefined acceleration policy.

Note: *You cannot edit the settings for a predefined policy.*

Copying an acceleration policy

Create a user-defined acceleration policy most efficiently by copying an existing acceleration policy and modifying its rules to meet your unique requirements.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. In the Tools column, click **Copy** for the acceleration policy you want to copy.
3. Name the policy.
4. Specify a folder, based on your configuration.
 - For a symmetric or farm configuration, from the **Sync Folder** list, select the name of a symmetric folder.
 - For an asymmetric configuration, from the **Sync Folder** list, select **No Selection**.
5. In the **Description** field, type a description.
6. Click **Copy**.

The acceleration policy appears in the Policy column.

Creating a user-defined acceleration policy from a predefined acceleration policy

You can copy a predefined acceleration policy, and modify applicable nodes, matching rules, and acceleration rules, to create a user-defined acceleration policy.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. In the Tools column, click **Copy** for the predefined acceleration policy you want to copy.
3. Name the policy.
4. Specify a folder, based on your configuration.
 - For a symmetric or farm configuration, from the **Sync Folder** list, select the name of a symmetric folder.
 - For an asymmetric configuration, from the **Sync Folder** list, select **No Selection**.
5. Click **Copy**.
6. Click the name of the new user-defined acceleration policy.
7. Create, delete, or modify nodes, matching rules, and acceleration rules, as necessary.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The user-defined acceleration policy appears in the Policy column.

Creating a new user-defined acceleration policy

You can create a new user-defined acceleration policy and define each matching rule and acceleration rule individually.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Create**.
3. Name the policy.

4. In the **Description** field, type a description.
5. Click **Create**.
6. Click the name of the new user-defined acceleration policy.
7. Create the Policy Tree by defining branch nodes for groups of content, and leaf nodes for specific content.
8. Specify the matching and acceleration rules for each node.
9. Click **Exit Policy Editor**.

The acceleration policy appears in the Policy column.

Importing an acceleration policy

An acceleration policy must have been saved.

You can import a saved acceleration policy file, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Import**.
3. Click **Browse**, and browse to the location of the XML file that you want to import.
4. Name the policy.
5. Specify a folder, based on your configuration.
 - For a symmetric or farm configuration, from the **Sync Folder** list, select the name of a symmetric folder.
 - For an asymmetric configuration, from the **Sync Folder** list, select **No Selection**.
6. Specify the overwrite method:
 - Select the **Overwrite existing policy of the same name** check box to replace an existing acceleration policy with the imported acceleration policy with the same name.
 - Clear the **Overwrite existing policy of the same name** check box to replace the existing acceleration policy.
7. Click **Import**.

The acceleration policy appears in the Policy column.

Publishing a user-defined acceleration policy

When you create a new acceleration policy, you must publish it before you can assign it to an application. The BIG-IP device uses a modified acceleration policy to manage traffic only after you publish it.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. In the Tools column, click **Publish** for the user-defined acceleration policy you want to publish.
3. Click **Publish Now**.

The user-defined acceleration policy is published.

Modifying an acceleration policy's rules

You can modify the acceleration rules for a user-defined acceleration policy to meet your unique requirements.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click an item on the **Acceleration Rules** menu to view the rules for that item.
6. Edit the settings, as necessary.
7. Click **Save**.

The modified rules appear for the acceleration rules menu item.

Viewing rules for an acceleration policy

You can view the acceleration rules for an acceleration policy and determine what rules apply to each node.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Specify the type of acceleration policy:
 - Click the name of a user-defined acceleration policy.
 - Click the name of a predefined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click an item on the **Acceleration Rules** menu to view the rules for that item.

The acceleration rules for the selected item appear for the node.

Saving an acceleration policy to an XML file

When you change an acceleration policy, you can export it to an XML file so that you always have a copy of the most recent acceleration policy.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Export**.
3. In the **Export** list, select one of the following:

Item	Description
Published Policy	Exports an acceleration policy that an application is currently using. If the acceleration policy has not been published, this option does not display.
Development Policy	Exports an unpublished acceleration policy.

4. Click **Export**.
5. Click **Save**.
6. Navigate to the location where you want to save the file.
7. Click **Save**.

The acceleration policy is exported as an XML file.

Deleting a user-defined acceleration policy

You can delete a user-defined policy, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.

2. Select the check box for the applicable acceleration policy.
3. Click **Delete**.

The policy is deleted.

Overview: Policy Editor screen

From the Policy Editor screen, you can view the matching rules and acceleration rules for user-defined and predefined acceleration policies, as well as create or modify user-defined acceleration policies.

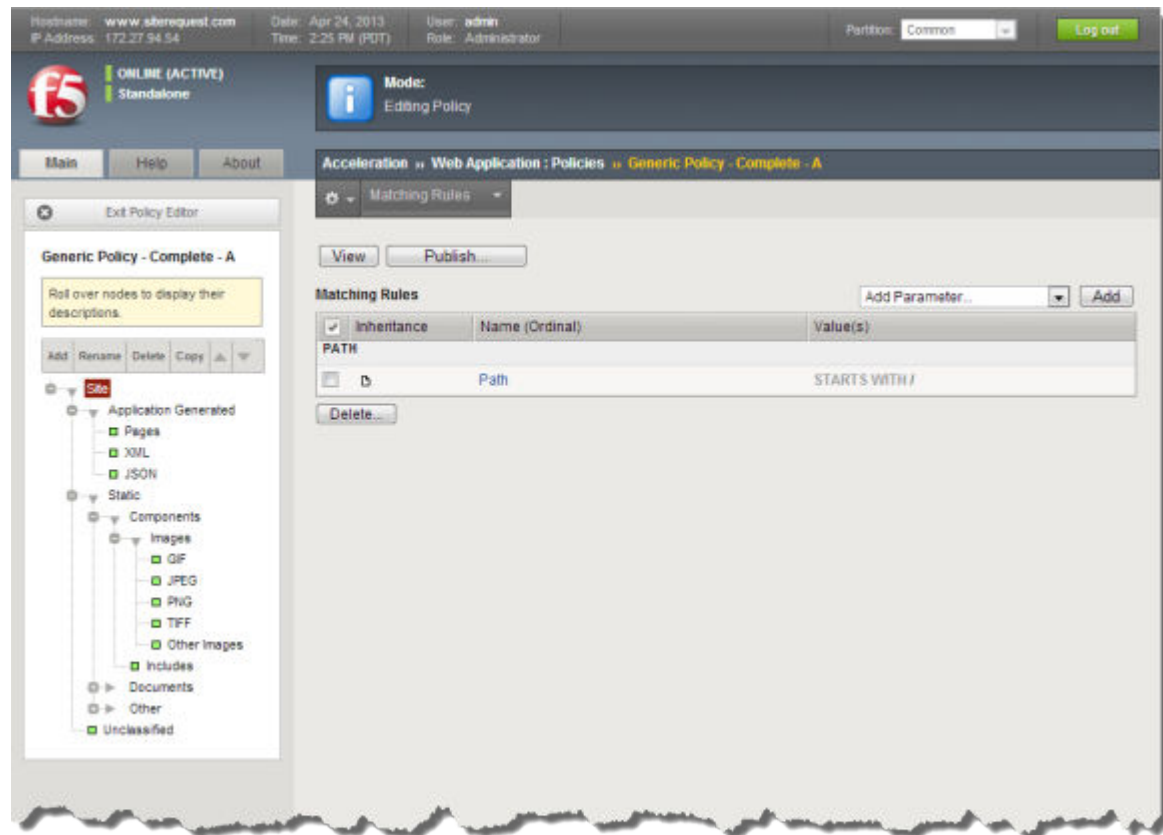


Figure 1: Policy Editor screen for an example acceleration policy

Task summary for using the Policy Editor

Perform these tasks to use the Policy Editor.

Task summary

Accessing the Policy Editor screen

Viewing the Policy Tree for an acceleration policy

Accessing the Policy Editor screen

The Policy Editor displays the matching rules and acceleration rules for user-defined acceleration policies, and enables you to modify user-defined acceleration policies.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.

The Policy Editor screen appears for the user-defined acceleration policy.

Viewing the Policy Tree for an acceleration policy

Matching rules and acceleration rules for acceleration policies are organized on the Policy Tree, which you access from the Policy Editor screen.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Specify the type of acceleration policy:

- Click the name of a user-defined acceleration policy.
- Click the name of a predefined acceleration policy.

The Policy Tree appears on the left side of the Policy Editor screen.

Differentiating Requests and Responses with Variation Rules

Overview: Variation rules

When the BIG-IP® system caches responses from the origin web server, it uses certain HTTP request parameters to create a Unique Content Identifier (UCI). The BIG-IP system stores the UCI in the form of a compiled response, and uses the UCI to easily match future requests to the correct content in the module's cache.

You can configure variation rules to add or modify the parameters on which the BIG-IP system bases its caching process. If the BIG-IP system receives two requests that are identical except for the value of a query parameter defined in the variation rule, it creates a different UCI for each, and caches each response under its unique UCI.

Consider a site that receives requests from customers and partners, and wants to serve different content to each. For this site, you could create a variation rule in which you specify that when a request contains a `version` cookie set to a value of 1, the BIG-IP system serves a page specifically for customers, and when the `version` cookie is set to a value of 2, it serves a page specifically for partners. For this rule, the BIG-IP system caches the following three compiled responses.

- For content produced for `Cookie: version=1`.
- For content produced for `Cookie: version=2`.
- For content produced when the `version` cookie does not appear in the request.

Note: When configuring this variation rule, you must specify a value for the `version` cookie parameter. If you do not, the BIG-IP system ignores the cookie's value and produces, at most, two compiled responses: one for requests that contain the cookie, and one for requests that do not contain the cookie. The BIG-IP system then serves the first response it caches to any subsequent requests that contain that cookie.

Configuring variation rule settings

You can configure variation rules to add or modify the parameters on which the BIG-IP system bases its caching process. If the BIG-IP system receives two requests that are identical except for the value of a

variation rule parameter, it creates a different UCI for each, and caches each response under its unique UCI.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Variation**.
6. Define the parameters that the BIG-IP system must match in a request.

The BIG-IP system creates a unique UCI for requests with different variation rule parameter values, and caches each response under its unique UCI.

Configuring ambiguous query parameters as unnamed

By default the BIG-IP system treats all ambiguous query parameters as a named query parameter without a value. You can, however, override this default behavior.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Variation**.
6. Click **All Query Parameters**.
7. Select the **Treat ambiguous parameters as unnamed** check box.
8. Click **Save**.

The BIG-IP system treats all ambiguous query parameters as unnamed query parameters.

Variation rule example

This example shows a variation rule for a user-defined acceleration policy. For this example site, you have three top-level nodes on the Policy Tree.

- **Home**. This branch node specifies the rules related to the home page.
- **Applications**. This branch node specifies the rules related to the applications for the site, with the following leaf nodes.
 - **Default**. This leaf node specifies the rules related to non-search related applications.
 - **Search**. This leaf node specifies the rules related to your site's search application.
- **Images**. This branch node specifies the rules related to graphics images.

This site also has the following two considerations.

- It needs to provide different branding information if the `REFERER` request header begins with, `http://www.siterequest.com`. Any other value for the `REFERER` request header does not affect the content served by your site.
- It uses a query parameter called `sessionID` to track site users. This query parameter does not affect page content, and is used for tracking purposes only.

For this example, you create the following two variation rules.

- **REFERRER rule**. You base this rule on the `REFERER` data type, and set it on the Applications node.
- **Query Parameter rule**. You base this rule on the `Query Parameter` data type, and set on the **Search** node.

Creating an example Referrer rule

This example provides steps to configure a Referrer rule.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click the **Applications** node.
4. From the **Matching Rules** list, select **Acceleration Rules**.
5. Click **Variation**.
6. Click **Referrer**.
7. In the Value Groups area, click **Add**.
8. Select the **Values matches** check box, and in the adjacent box, type the regular expression that matches the value you expect on the REFERER request header, as follows: `http://www\.\siterequest\.com.*`.
9. From the **Values Define** list, select **Different Content**, to prompt the BIG-IP system to provide a unique page to matched requests.
10. Click **Save**.
11. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The example Referrer variation rule is configured.

Creating an example Query Parameter rule

The example provides steps to configure a Query Parameter rule.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click the **Search** node.
4. From the Matching Rules list, select **Acceleration Rules**.
5. Click **Variation**.
6. From the **Add Parameter** list, select **Query Parameter** and click the **Add** button.
7. In the **Name** box, type `sessionID`.
8. In the Values Groups area, click the **Add** button.
9. Select the **Values matches** check box, and in the adjacent box, type the following regular expression, to indicate that the query parameter can have any value: `.*`.
10. Select the **Value is an empty string** check box.
11. From the **Values Define** list, select **Same Content**, to indicate that page content is not affected by this parameter.
12. Click **Save**.

The example Query Parameter variation rule is configured.

Proxying Requests and Responses

Overview: Proxying rules

In general, the BIG-IP system attempts to service all HTTP requests from the system's cache. However, if you have certain types of content that you do not want the BIG-IP system to service from the system's cache, you can configure proxying rules. Using proxying rules, you identify HTTP request elements that prompt the BIG-IP system to send a request to your origin web servers for content.

You configure proxying rules from the proxying screen.

Always proxying requests for a node

When the BIG-IP system matches a request to a node that has the **Always proxy requests for this node** setting enabled, it sends the request to the origin web server, and responds to the request without caching the content. This option overrides any configured proxying rule, and is useful for specific content that is private.

You can, if necessary, specify objects for the BIG-IP system to cache by defining **Proxy Override Rules**, thus eliminating a 304 (Not Modified) response from the origin web server for a conditional request. Any objects defined by **Proxy Override Rules**, however, should no longer be considered private.

***Note:** If you configure a node to cache conditional `POST` requests, using `If-None-Match`, `If-Match`, `If-Modified-Since`, or `If-Unmodified-Since` headers, you must configure the node to always proxy the requests.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Proxying**.
6. Select the **Always proxy requests for this node** option.

***Note:** Selecting **Always proxy requests for this node** only proxies requests. This setting does not proxy responses. To avoid caching a response that matches the node, on the menu bar click **Lifetime**, clear the **Honor Headers From Origin Web Server** check box in the WebAccelerator Cache Settings area, and type 0 in the **Maximum Age** field.*

7. From the **Caching Mode** list, select one of the following to cache proxy override objects.

Option	Description
Memory & Disk Cache	Caches objects for the selected node to memory and disk cache.
Memory-only Cache	Caches objects for the selected node only to memory. In the event of power loss, memory is cleared, providing greater security.

8. From the **Proxy Override Rules** list, select parameters and configure settings, as applicable.
Proxy override rules and associated conditions enable the BIG-IP to ignore proxying rules and, consequently, cache objects defined by this setting.

Important: To cache objects in accordance with **Caching Mode** settings, you must also define **Proxy Override Rules**. Unless you define **Proxy Override Rules**, all objects are proxied.

9. Click **Save**.
10. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP system always proxies a request that matches the associated node to the origin web server.

Configuring proxy rules for a node

When you configure proxy rules for a node, the BIG-IP system applies the configured proxying rules to requests that match the associated node.

Note: If you configure a node to cache conditional *POST* requests, using *If-None-Match*, *If-Match*, *If-Modified-Since*, or *If-Unmodified-Since* headers, you must configure the node to always proxy the requests.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Proxying**.
6. Select the **Configure and use Proxy Rules for this node** option.
7. From the **Caching Mode** list, select one of the following to cache proxy override objects.

Option	Description
Memory & Disk Cache	Caches objects for the selected node to memory and disk cache.
Memory-only Cache	Caches objects for the selected node only to memory. In the event of power loss, memory is cleared, providing greater security.

8. From the **Proxy Rules** list, select parameters and configure settings, as applicable.

Note: In general, proxy rules options are relevant to only requests that match the applicable node, rather than to matched responses. If you configure proxying rules based on HTTP response header parameters, you can use them only in terms of how the BIG-IP caches the responses, because the BIG-IP has already sent the request to the origin web servers when it reviews the response headers.

9. From the **Proxy Override Rules** list, select parameters and configure settings, as applicable.
Proxy override rules and associated conditions enable the BIG-IP to ignore proxying rules and, consequently, cache objects defined by this setting.

Important: To cache objects in accordance with **Caching Mode** settings, you must also define **Proxy Override Rules**. Unless you define **Proxy Override Rules**, all objects are proxied.

10. Click **Save**.
11. Publish the acceleration policy.
 - a) Click **Publish**.

- b) In the **Comment** field, type a description.
- c) Click **Publish Now**.

The BIG-IP system applies configured proxy rules to requests that match the associated node.

Overriding proxy requests

You can configure a proxy override rule to prevent an origin web server from managing excessive traffic. For example, one common use of proxy override rules is for sites that receive a high volume of traffic related to web crawlers and robots clients.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Proxying**.
6. Select the **Configure and use Proxy Rules for this node** option.
7. Select one of the following **Proxying Options**.
 - **Always proxy requests for this node**
 - **Configure and use Proxy Rules for this node**
8. From the **Proxy Override Rules** list, select parameters and configure settings, as applicable.
Proxy override rules and associated conditions enable the BIG-IP to ignore proxying rules and, consequently, cache objects defined by this setting.

Important: To cache objects in accordance with **Caching Mode** settings, you must also define **Proxy Override Rules**. Unless you define **Proxy Override Rules**, all objects are proxied.

9. Click **Save**.
10. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

A proxy override rule is configured.

About creating the example proxying rule

In the example for this site, you have three top-level nodes on the Policy Tree as follows:

- **Home**. This branch node specifies the rules related to the home page.
- **Applications**. This branch node specifies the rules related to the applications for the site, with the following leaf nodes:
 - **Default**. This leaf node specifies the rules related to non-search related applications.
 - **Search**. This leaf node specifies the rules related to your site's search application.
- **Images**. This branch node specifies the rules related to graphics images.

For this example, you use a segment parameter to contain identifying information for your shopping cart application. Requests for your applications are all in the following form:

```
http://www.example.com/apps/doSomething.jsp;AAy23BV39
```

If the session tracking string does not appear in the segment parameter at the end of the URI, you want the BIG-IP system to send the request to the origin web servers for special handling. Create a proxying rule for your **Applications** node with just one parameter based on the path segment data type. This rule should identify the subject as being path ordinal 1 in the full path, as counted from right to left.

You set this rule to go into effect if the parameter value is an empty string or if the parameter is absent from the request.

Configuring the example proxying rule

When defining parameters for proxy rules, you specify the parameter identity and its value that must appear in a request.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Proxying**.
6. Select the **Configure and use Proxy Rules for this node** option.
7. From the **Caching Mode** list, select **Memory & Disk Cache**.
8. From the **Proxy Rules** list, select **Cookie**, and click the **Add** button.
9. Name the cookie `version`.
10. Under **Value(s)**, select the check box beside the **Value matches:** list, and select **Value does not match** from the list.
11. From the **Value does not match** field, type `version`.
12. Select the **Value is an empty string** check box.
13. Click **Save**.
14. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The example proxying rule is configured.

About configuring the proxying rule parameters example

You define a parameter for proxy rules by specifying a parameter identity and its value, which must appear in a request for the BIG-IP system to send it to the origin web server for content.

For example, if you create a proxy rule based on a cookie named `version`, the BIG-IP system sends the request to the origin server if one of the following conditions is true:

- The `version` cookie does not appear on the request.
- The `version` cookie appears on the request, but has no value set for it (`version` is empty).

Creating the example proxying rule parameters

When defining parameters for proxy rules, you identify the parameter and its value that must appear in a request, for the BIG-IP system to send it to the origin web server for content.

For example, if you create a proxy rule based on a cookie named `version`, the BIG-IP system sends the request to the origin web server if one of the following conditions is true:

- The `version` cookie does not appear on the request.

- The `version` cookie appears on the request, but has no value set for it (version is empty).
1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
 2. Click the name of a user-defined acceleration policy.
 3. Expand the Policy Tree to a branch node or leaf node, and click the node.
 4. From the Matching Rules menu, choose Acceleration Rules.
 5. Click **Proxying**.
 6. Select the **Configure and use Proxy Rules for this node** option.
 7. From the **Caching Mode** list, select **Memory & Disk Cache**.
 8. From the **Proxy Rules** list, select **Cookie**, and click the **Add** button.
 9. Name the cookie `version`.
 10. Under **Value(s)**, select the check box beside the **Value matches:** list, and select **Value does not match** from the list.
 11. From the **Value does not match** field, type `version`.
 12. Select the **Value is an empty string** check box.
 13. Click **Save**.
 14. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The example proxying rule is configured.

About configuring the example proxy override rule

One common application for proxy override rules is for sites that receive a high volume of traffic related to web crawlers and robots clients. You can avoid having your origin web server manage this excessive traffic by configuring a proxy override rule. It is especially important to consider a proxy override rule for this application if your proxy rules are based on a set cookie, because web crawlers and robots rarely present cookies in their requests.

Before configuring the rule, you could examine the origin server's log files to see if the web crawler presents a specific string that you can use for a proxy override rule. For example, the web crawler might present a string for the `HTTP USER_AGENT` request header that looks very much like the value presented for MSIE 5.0 browsers. However, the web crawler adds `badcrawler` to the `HTTP USER_AGENT` string. For this example, you use a proxy override rule for the node's proxying rule, based on the user agent parameter to match the regular expression `.*badcrawler.*`.

When the BIG-IP system finds the user agent with this regular expression in an HTTP request, it services the request from cache, even if the matched proxying rule dictates that the BIG-IP system should send the request to the origin server.

Once you publish the acceleration policy, the BIG-IP system attempts to service, from cache, all requests that it receives with a `HTTP USER_AGENT` value that matches `.*badcrawler.*`, regardless of any proxy rules that the request matches.

Configuring the example proxy override rule

In this proxy override rule example, when the BIG-IP system finds the user agent with the regular expression `.*badcrawler.*` in an HTTP request, it services the request from cache, even if the matched proxying rule dictates that the BIG-IP system should send the request to the origin server.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Proxying**.
6. From the **Caching Mode** list, select **Memory & Disk Cache**.
7. From the **Proxy Override Rules** list, select **User Agent**.
8. Click the **Add** button.
9. Select the **Value matches** check box, and select **Value matches** from the list.
10. From the **Value matches** field, type `*badcrawler.*` (the regular expression that matches the user agent value).
11. Click **Save**.
12. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP system services HTTP requests with a user agent that contains `badcrawler` from cache.

Managing Requests and Responses with Lifetime Rules

Overview: Lifetime rules

The length of time that a client browser, upstream device, or BIG-IP system keeps compiled content in its cache before refreshing it is called *content lifetime*. Content lifetime is expressed in the form of a time to live (TTL) value, and can vary for each cached response.

When content is in cache longer than its TTL value, the BIG-IP system considers the content expired. When the BIG-IP system receives a request for expired content, it sends that request to the origin web servers for fresh content, replaces the expired cached content with the fresh response, and then responds to the request.

Configuring lifetime cache settings

The **WebAccelerator Cache Settings** and **Client Cache Settings** enable you to specify lifetime and privacy settings for origin web server and client objects that are cached by the BIG-IP and downstream devices.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Lifetime**.
6. Configure the system to honor headers in a request from a client.
 - a) Select the **Honor Headers In Request** check box.
 - b) For the **Request Headers** setting, select a header in the **Available** field, and move the header to the **Selected** field using the **Move** button.
7. Configure the system to honor headers from an origin web server.
 - a) Select the **Honor Headers From Origin Web Server** check box.
 - b) For the **Origin Web Server Headers** setting, select a header in the **Available** field, and move the header to the **Selected** field using the **Move** button.

***Note:** The BIG-IP inserts an *Expires* header into the response when *expires* is selected, if the origin web server does not include an *Expires* header. It does not insert an *Expires* header into the response when *all* is selected.*

8. Do one of the following to specify the maximum age.
 - Select the **Use HTTP Lifetime Heuristic** check box, and type a number in the **%Heuristic** field to specify a heuristic percentage based on the HTTP Last-Modified header.
 - Clear the **Use HTTP Lifetime Heuristic** check box, type a number for the duration in the **Maximum Age** field, and select a unit of time from the list.
9. In the **Stand-in Period** field, type a number for the duration, and select a unit of time from the list.
10. For the **Stand-in Codes** setting, select a code entry in the **Available** field, and move the code entry to the **Selected** field using the **Move** button.
11. Select one of the following **Client Cache Settings** options, and configure as applicable.
 - **Preserve Origin Web Server headers/directives to downstream devices.**

- a) In the **Maximum Age** field, type a number for the duration, and select a unit of time from the list.
- b) In the **S-Max Age** field, type a number for the duration, and select a unit of time from the list.
- c) In the **Custom Cache Extensions Add** field, type an extension, and click **Add**.
The extension appears in the **Custom Cache Extensions** list.

- **Custom Cache-Control Directives.**

- a) In the **Maximum Age** field, type a number for the duration, and select a unit of time from the list.
- b) In the **S-Max Age** field, type a number for the duration, and select a unit of time from the list.
- c) In the **Custom Cache Extensions Add** field, type an extension, and click **Add**.
The extension appears in the **Custom Cache Extensions** list.

- **Replace Origin Web Server Headers/Directives with no-cache.**

- a) In the **Custom Cache Extensions Add** field, type an extension, and click **Add**.
The extension appears in the **Custom Cache Extensions** list.

The lifetime and privacy settings are configured for client and origin web server objects that are cached by the BIG-IP and downstream devices.

An example lifetime rule

For this example site, you have three top-level nodes on the Policy Tree as follows:

- **Home.** This branch node specifies the rules related to the home page.
- **Applications.** This branch node specifies the rules related to the applications for the site, with the following leaf nodes:
 - **Default.** This leaf node specifies the rules related to non-search related applications.
 - **Search.** This leaf node specifies the rules related to your site's search application.
- **Images.** This branch node specifies the rules related to graphics images.

Example of lifetime rule for the Home node

You change your site's content approximately every 4 hours. You want the BIG-IP system to cache content for no longer than 24 hours. If the origin web servers are not responding for request for fresh content, you are willing to allow the BIG-IP system to serve content that is 8 hours old (or twice the age of the content).

To ensure that you can manage content invalidation, you do not want to rely solely on the browser's local cache settings for the home node, so you do not have a minimum time set for content residing in the browser cache before performing a check for content freshness.

1. On the **Lifetime** tab for the **Home** node, select the **Honor Headers From Origin Web Server** check box.
2. For the **Origin Web Server Headers** setting, select **expires** in the **Available** field, and move the header to the **Selected** field using the **Move** button.
3. Specify a **Maximum Age** of 24 **Hours** and a **Stand-in Period** of 8 **Hours**.
4. Leave all other options at the default settings.

An example **Home** branch node lifetime rule is configured.

An example lifetime rule for the Default leaf node

The content served for your general applications changes about once every 4 hours. You use an invalidations rule to force a refresh when content changes, but you do not want content to remain in the system's cache for more than 5 hours without a refresh.

If the origin web servers are not responding to the BIG-IP system's refresh requests, you are willing to allow the BIG-IP system to serve content that is 8 hours old (or twice the age of the content).

1. Create a lifetime rule for the **Default** leaf node.
2. Select the **Honor Headers From Origin Web Server** check box.
3. For the **Origin Web Server Headers** setting, select **expires** in the **Available** field, and move the header to the **Selected** field using the **Move** button.
4. Specify a **Maximum Age** of **5 Hours** and a **Stand-in Period** of **8 Hours**.
5. Leave all other options at the default settings.

An example **Default** leaf node lifetime rule is configured.

An example lifetime rule for the **Search** leaf node

Your search application returns data that has various expiration times; some content expires in as little as 10 minutes, and some content expires at 8 hours. You intend to use the HTTP Cache-Control Expire header max-age directive to identify the cache time for content served by the search application.

1. On the **Lifetime** tab for the **Search** leaf node, select the **Honor Headers From Origin Web Server** check box.
2. For the **Origin Web Server Headers** setting, select **expires** in the **Available** field, and move the header to the **Selected** field using the **Move** button.
3. Specify a **Maximum Age** of **8 Hours** and a **Stand-in Period** of **10 Hours**.
4. Leave all other options at the default settings.

An example **Search** leaf node lifetime rule is configured.

An example lifetime rule for the **Images** node

You change the images for your applications approximately every 4 hours. You want the BIG-IP system to cache images for no longer than 24 hours and, if the origin web servers are not responding to the BIG-IP system's refresh requests, you are willing to allow the BIG-IP system to serve images that are 8 hours old (or twice the age of the image).

1. On the **Lifetime** tab for the **Images** node, select the **Honor Headers From Origin Web Server** check box.
2. For the **Origin Web Server Headers** setting, select **expires** in the **Available** field, and move the header to the **Selected** field using the **Move** button.
3. Specify a **Maximum Age** of **24 Hours** and a **Stand-in Period** of **8 Hours**.
4. In the **Client Cache Settings** area, select the **Custom Cache-Control Directives** option, and specify a **Maximum Age** of **4 Days**.
5. Leave all other options at the default settings.

An example **Images** leaf node lifetime rule is configured.

Invalidating Cached Content

Overview: Invalidating cached content for an application

You typically use invalidations rules to expire specific content stored in cache. When you invalidate content, that content is immediately expired but not removed from the disk.

Invalidating cached content for an application

You can invalidate cached content for one or more applications, which expires, but does not clear, the cached content for the specified applications.

***Note:** Invalidating the cached content on the BIG-IP device temporarily increases traffic to the origin web servers until the BIG-IP device repopulates the cache.*

1. On the Main tab, click **Acceleration > Web Application > Invalidate Content**.
The Invalidate Content screen opens.
2. Select the check box for the application with cached content that you want to invalidate.
3. Click **Invalidate**.

The cached content for the specified applications is invalidated.

Clearing cached content from the command line

There might be occasions, such as when troubleshooting cache issues, when you want to clear cache and remove content that the BIG-IP device has stored.

***Note:** Clearing cached content temporarily increases traffic to the origin web servers until the BIG-IP device repopulates the cache.*

1. Log on to tmsh.
2. Type `tmsh delete ltm profile wa-cache wa_profile_name` at the command line, where `wa_profile_name` is the name of the Web Acceleration profile.

The content cached by the BIG-IP device is cleared.

Overview: Invalidating cached content for a node

Cache invalidation is a powerful tool that you can use to maintain tight coherence between the content on your origin web servers and the content that the BIG-IP system caches.

If you update content for your site at regular intervals, such as every day or every hour, you can use lifetime rules to ensure that the system's cache is refreshed with the same frequency. Invalidations rules, however, allow you to expire cached content before it has reached its time to live (TTL) value, and is a good tool to use when content updates are event-driven, such as when an item is added to a shopping cart, a request contains a new auction bid, or a poster has submitted content on a forum thread.

When you configure invalidations rules, you define elements in a request that prompt the BIG-IP system to invalidate and refresh specific cached content. When the BIG-IP system receives a request that matches the parameters that you specified for the invalidations rule, it performs the following steps.

- Invalidates the cached content that it would have served.
- Sends the request to the origin web server for fresh content.
- Replaces the specified content, which it previously had in cache, with the new content it receives from the origin web server.
- Responds to the request with the refreshed content.

You can create invalidations rules that are based on a specific parameter, for example, an invalidation rule based on a certain cookie.

Important: *Although there might be situations that require you to invalidate a significant portion of the cache, it is important to keep in mind that such a broad invalidation process can tax the origin web server as it attempts to respond to multiple requests for new content. For this reason, F5 Networks® recommends that you make the invalidations rule parameters as specific as possible, whenever possible.*

Creating an Invalidations rule

You can define a trigger in a policy that invalidates the obsolete cached post, and retrieves the updated post from the origin web server.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a leaf node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Invalidations**.
6. Click **Create**.
7. In the **Description** field, type a description.
8. Define the parameters that the BIG-IP system must match in a request.
These parameters trigger the invalidations rule.

Important: *All parameters are optional except for the **Path** parameter. If you do not specify the **Path** parameter for the **Request Header Matching Criteria** and the **Cached Content to Invalidate** settings, the invalidations rule does not trigger the BIG-IP system to invalidate the specified cache. If you do not want to define a specific path, you can use a single slash (/).*

- a) In the **Add Parameter** list, select **Path**.
 - b) Click **Add**.
 - c) In the **Value(s)** field, type the path for an application.
For example, to invalidate cached content for the application `/apps/doSomething.jsp`, type `/apps/doSomething.jsp`.
The BIG-IP system triggers the invalidations rule only for this application. All other requests that match to the Default node do not trigger this invalidations rule.
 - d) Click **Save**.
The path appears in the **Request Header Matching Criteria** table.
9. Define the content to invalidate and refresh.
This content is invalidated and refreshed when the **Request Header Matching Criteria** setting specifies the parameters in the HTTP request header.

- a) In the **Add Parameter** list, select **Path**.
- b) Click **Add**.
- c) In the **Type** list, select **Value Group**.
- d) In the **Value Group** field, type the path for the content.
For example, to invalidate cached content for `/srch/doSimpleSearch.jsp`, type `/srch/doSimpleSearch.jsp`.
- e) Click **Save**.
The path appears in the **Request Header Matching Criteria** table.

The trigger in the policy invalidates the obsolete cached post, and retrieves the updated post from the origin web server.

Blog invalidations rule example

In this example, a participant posts an update to a blog, which you want to make available through an invalidation rule, instead of sending the cached version without the updated post. You can define a trigger in a policy that invalidates the obsolete cached post, and retrieves the updated post from the origin web server.

First, you can create a **View** leaf node that includes **Matching** rules with these settings:

Parameter	Name	Value
Path	Path	STARTS WITH /
Query Parameter	view	MATCHES topic
Content Type	Content Type	MATCHES pages.html OR IS UNCLASSIFIED

Then, you can create a **Post** leaf node that includes **Matching** rules with these settings:

Parameter	Name	Value
Path	Path	STARTS WITH /
Query Parameter	view	MATCHES post
Query Parameter	forumid	MATCHES *
Content Type	Content Type	MATCHES pages.html OR IS UNCLASSIFIED

Finally, you can specify the **Invalidations** rules for the **Post** leaf node to include these settings:

Request Header Matching Criteria

Parameter	Name	Value
Path	Path	STARTS WITH /
Query Parameter	view	MATCHES post
Query Parameter	forumid	MATCHES *
Query Parameter	postid	MATCHES *

Cached Content to Invalidate

Parameter	Name	Value
Path	Path	STARTS WITH /
Query Parameter	view	MATCHES topic
Query Parameter	forumid	Source forumid [Query Parameter]
Query Parameter	postid	Source postid [Query Parameter]

Creating leaf nodes for the blog example invalidations rule

In this task, you create **View** and **Post** leaf nodes in the Policy Tree for the example blog invalidations rule.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click the **Blog** branch node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Click **Invalidations**.
6. Create a leaf node named **View**.
 - a) Click **Add** on the Policy Tree function bar.
 - b) Name the node `View`.
 - c) Type a description.
 - d) Click **Create**.
7. Create a leaf node named **Post**.
 - a) Click **Add** on the Policy Tree function bar.
 - b) Name the node `Post`.
 - c) Type a description.
 - d) Click **Create**.

The **View** and **Post** nodes appear in the Policy Tree.

Specifying matching rules for the invalidations example View node

A **View** and **Post** leaf node have been created.

This task configures matching rules for the invalidations example **View** node.

1. Click the **View** node.
2. For **Matching Rules**, specify the **Path** settings.
 - a) In the **Add Parameter** list, select **Path**.
 - b) Click **Add**.
 - c) In the **Value** field, type `/`.
 - d) Click **Save**.
3. For **Matching Rules**, specify the **view** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.

- c) In the **Name** field, type `view`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `topic`.
 - f) Click **Save**.
4. For **Matching Rules**, specify the **Content Type** settings.
 - a) In the **Add Parameter** list, select **Content Type**.
 - b) Click **Add**.
 - c) Select the check box beside the **Value matches** list.
 - d) In the **Enter a regular expression** field, type `pages.html`.
 - e) Select the **Match if not yet classified** check box.
 - f) Click **Save**.

The matching rules for the invalidations example **View** node are configured.

Specifying matching rules for the invalidations example **Post** node

A **View** and **Post** leaf node have been created.

This task configures matching rules for the invalidations example **Post** node.

1. Click the **Post** node.
2. For **Matching Rules**, specify the **Path** settings.
 - a) In the **Add Parameter** list, select **Path**.
 - b) Click **Add**.
 - c) In the **Value** field, type `/`.
 - d) Click **Save**.
3. For **Matching Rules**, specify the **forumid** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `forumid`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `*`.
 - f) Click **Save**.
4. For **Matching Rules**, specify the **view** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `view`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `post`.
 - f) Click **Save**.
5. For **Matching Rules**, specify the **Content Type** settings.
 - a) In the **Matching Rules Add Parameter** list, select **Content Type**.
 - b) Click **Add**.
 - c) Select the check box beside the **Value matches** list.
 - d) In the **Enter a regular expression** field, type `pages.html`.
 - e) Select the **Match if not yet classified** check box.
 - f) Click **Save**.

The matching rules for the invalidations example **Post** node are configured.

Task summary for specifying invalidations rules for the invalidations example Post node

Perform these tasks to configure the invalidations blog example **Post** node.

Task summary

Configuring blog example Path settings for the Post node invalidations rules

Configuring blog example view settings for the Post node invalidations rules

Configuring blog example forumid settings for the Post node invalidations rules

Configuring blog example postid settings for the Post node invalidations rules

Configuring blog example Path settings for the Post node invalidations rules

A **View** and **Post** leaf node have been created.

In this task for configuring blog invalidations rules, you specify **Path** settings for the **Post** node invalidations rules.

1. Click the **Post** node.
2. In the **Matching Rules** list, select **Acceleration Rules**.
3. Click **Invalidations**.
4. Click **Create**.
5. In the **Description** field, type a description.
6. For **Request Header Matching Criteria**, specify the **Path** settings.
 - a) In the **Add Parameter** list, select **Path**.
 - b) Click **Add**.
 - c) In the **Value(s)** field, type **/**.
 - d) Click **Save**.
7. For **Cached Content to Invalidate**, specify the **Path** settings.
 - a) In the **Add Parameter** list, select **Path**.
 - b) Click **Add**.
 - c) In the **Value Group** field, type **/**.
 - d) Click **Save**.

The **Path** settings for the **Post** node are configured for the example blog invalidations rules.

Configuring blog example view settings for the Post node invalidations rules

A **View** and **Post** leaf node have been created.

In this task for configuring blog invalidations rules, you specify **view** settings for the **Post** node invalidations rules.

1. Click the **Post** node.
2. In the **Matching Rules** list, select **Acceleration Rules**.
3. Click **Invalidations**.
4. Click **Create**.
5. In the **Description** field, type a description.
6. For **Request Header Matching Criteria**, specify the **view** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.

- c) In the **Name** field, type `view`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `post`.
 - f) Click **Save**.
7. For **Cached Content to Invalidate**, specify the **view** settings.
- a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `view`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `topic`.
 - f) Click **Save**.

The **view** settings for the **Post** node are configured for the example blog invalidations rules.

Configuring blog example forumid settings for the Post node invalidations rules

A **View** and **Post** leaf node have been created.

In this task for configuring blog invalidations rules, you specify **forumid** settings for the **Post** node invalidations rules.

1. Click the **Post** node.
2. In the **Matching Rules** list, select **Acceleration Rules**.
3. Click **Invalidations**.
4. Click **Create**.
5. In the **Description** field, type a description.
6. For **Request Header Matching Criteria**, specify the **forumid** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `forumid`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `*`.
 - f) Click **Save**.
7. For **Cached Content to Invalidate**, specify the **forumid** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `forumid`.
 - d) In the **Type** list, select **Query Parameter from Request**.
 - e) In the **Name** field, type `forumid`.
 - f) Click **Save**.

The **forumid** settings for the **Post** node invalidations rules are configured for the example blog invalidations rules.

Configuring blog example postid settings for the Post node invalidations rules

A **View** and **Post** leaf node have been created.

In this task for configuring blog invalidations rules, you specify **postid** settings for the **Post** node invalidations rules.

1. Click the **Post** node.
2. In the **Matching Rules** list, select **Acceleration Rules**.
3. Click **Invalidations**.

4. Click **Create**.
5. In the **Description** field, type a description.
6. For **Request Header Matching Criteria**, specify the **postid** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `postid`.
 - d) Select the check box beside the **Value matches** list.
 - e) In the **Enter a regular expression** field, type `*`.
 - f) Click **Save**.
7. For **Cached Content to Invalidate**, specify the **postid** settings.
 - a) In the **Add Parameter** list, select **Query Parameter**.
 - b) Click **Add**.
 - c) In the **Name** field, type `postid`.
 - d) In the **Type** list, select **Query Parameter from Request**.
 - e) In the **Name** field, type `postid`.
 - f) Click **Save**.

The **postid** settings for the **Post** node invalidations rules are configured for the example blog invalidations rules.

Managing Object Types

Overview: Object classification

Before sending a response to a client, the BIG-IP system enters an informational `X-WA-Info` response header into the response to describe how it handled the response. You cannot change these informational headers, and they do not affect processing, however, they can provide useful information for evaluating the efficiency of your acceleration policies.

Part of the information included in the `X-WA-Info` response header is the object type. The BIG-IP system classifies, by object type and group, every response it receives from the origin web servers. The object type and group classification determine how the BIG-IP system handles compression for the response.

Task summary for managing object types

Perform these tasks to manage predefined and user-defined object types.

Task summary

Creating a user-defined object type

Editing an object type

Deleting a user-defined object type

Creating a user-defined object type

You can create an object type with custom parameters that determine how the BIG-IP system manages the specified object type.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Object Types**.
3. Click **Create**.
4. In the **Description** field, type a description.
5. In the **Object Type** field, type a name.
6. In the **Group** list, select a group.
7. Click **Add**, and, in the field, type an extension, omitting any leading period.
Click **Add** again to add another extension.
8. Click **Add**, and, in the field, type a MIME-type, for example, `application/rtf`.
Click **Add** again to add another MIME-type.
9. In the **Client Compression** list, select a compression setting.
10. In the **Symmetric Compression** list, specify whether to compress the object type in a symmetric configuration.
11. Click **Save**.

The new object type appears in the User-defined Object Types table and the BIG-IP system applies the new object type to all acceleration policies.

Editing an object type

You can modify an existing object type and enable the BIG-IP system to apply the object type changes globally to all acceleration policies.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Object Types**.
3. Click the name of an object type.
4. Do the following to edit extensions, as applicable:
 - a) Click **Add** to add an extension, and, in the field, type an extension, omitting any leading period.
 - b) Click **Delete** to delete an extension.
 - c) Edit an extension in the applicable field.
5. Do the following to edit a MIME-type, as applicable:
 - a) Click **Add**, and, in the field, type a MIME-type, for example, `application/rtf`.
 - b) Click **Delete** to delete a MIME-type.
 - c) Edit a MIME-type in the applicable field.
6. Click **Save**.

The BIG-IP system applies the object type changes to all acceleration policies.

Deleting a user-defined object type

You can permanently delete a user-defined object type.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click **Object Types**.
3. Select the check box next to the user-defined object type that you want to delete.
4. Click **Delete**.

The user-defined object type is permanently deleted.

Caching Objects in a VIPRION Cluster

Overview: Acceleration in a cluster

A VIPRION® system provides you with the ability to cache objects either for a policy node in a cluster or on a single cluster member. Typically, caching objects in a cluster achieves optimum acceleration for large, static objects. Comparatively, caching objects on a single cluster member achieves optimum acceleration for small, dynamic objects.

Caching objects in a cluster or cluster member

You can cache objects for a policy node in a cluster or on a single cluster member, optimizing acceleration for specific objects by policy node.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Responses Cached**.
6. (Optional) For the **Object Min Size** setting, clear the **Use Profile Setting** check box, and type the minimum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
7. (Optional) For the **Object Max Size** setting, clear the **Use Profile Setting** check box, and type the maximum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
8. From the **Content Coherency** list, select a cache coherency setting.

Option	Description
Cluster (default)	Provides optimum acceleration for large, static objects.
Blade	Provides optimum acceleration for small, dynamic objects.

9. Click **Save**.
10. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

Objects for a policy node are cached in a cluster or on a single cluster member.

Immediately Caching Dynamic Objects

Overview: Caching an object on first hit

The BIG-IP® system provides you with the ability to cache an object on the first cache hit, that is, the first time that an object is seen by the BIG-IP system. Typically, the BIG-IP system waits until the object is known to be popular, allowing a limited number of hits against the original content. Caching the object on first hit, however, causes the BIG-IP system to immediately cache the object, even if it is not popular. You only want to apply this setting for highly dynamic objects (for example, stock quotes provided by a stock ticker) that you want to cache for a limited time to offload the origin web server. If you use this setting on content that omits a `Content-Length` header, or if the object is an HTML, JavaScript (JS), or cascading style sheet (CSS) object, a significant degradation in performance can occur.

Caching an object on first hit

You can cache an object when the cache receives that object for the first time. Typically, you want to apply this setting only for highly dynamic objects, for example, stock quotes provided by a stock ticker.

Important: *A significant degradation in performance can occur when using this setting on an HTML, JavaScript (JS), or cascading style sheet (CSS) object, or content that omits a `Content-Length` header.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. Select the **Cache content on first hit** check box.

Important: *The **Queue Parallel Requests** check box must be clear to enable the **Cache content on first hit** setting. If the **Queue Parallel Requests** check box is selected, the **Cache content on first hit** setting is disabled.*

6. (Optional) For the **Object Min Size** setting, clear the **Use Profile Setting** check box, and type the minimum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
7. (Optional) For the **Object Max Size** setting, clear the **Use Profile Setting** check box, and type the maximum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
8. Click **Save**.

The BIG-IP system caches an object when the cache receives that object for the first time.

Disabling the Caching an object on first hit setting

You can disable the ability of the system to cache an object when the cache receives that object for the first time.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.

Immediately Caching Dynamic Objects

3. From the Matching Rules menu, choose Acceleration Rules.
4. On the menu bar, click **Responses Cached**.
5. Clear the **Cache content on first hit** check box.
6. Click **Save**.
7. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

Caching an object on first hit is disabled.

Accelerating Parallel HTTP Requests

Overview: HTTP request queuing

You can use the BIG-IP® system to accelerate responses and reduce the number of requests sent to origin web servers, freeing them to perform other tasks, by queuing HTTP requests. *HTTP request queuing* provides the ability to queue a large number of parallel HTTP requests for the same object, and provide a cached response to those requests, resulting in accelerated responses from the BIG-IP system and a reduction in requests sent to origin web servers.

The BIG-IP system manages the queued HTTP requests in accordance with the cache status of the requested object, that is, whether the requested object is uncached, cached and valid, cached and expired, uncacheable, or nonexistent.

Object cache status	Description
Initial requests for an uncached object	When the BIG-IP system receives a large number of parallel requests for an object that is not yet cached, it queues the requests, and then sends one of the requests to an origin web server. When the BIG-IP system receives the first response, it determines whether the response is cacheable while sending the response to the client. If the response is cacheable, the BIG-IP system sends a second request to the origin web server, caches the response with the object, and uses that cached response to service the queued requests.
Requests for a cached object	When the BIG-IP system receives a large number of parallel requests for a valid cached object, it services the requests with the cached response.
Requests for an expired cached object	If a cached object is expired, instead of sending all requests to the origin web server, the BIG-IP system queues the requests, and then sends one request to an origin web server for fresh content. When the BIG-IP system receives the fresh response, it caches the response with the fresh content, and uses the cached response to service the queued requests.
Requests for an invalidated cached object	If a cached object requires validation, the BIG-IP system can queue the requests, and then send one request to an origin web server for fresh content. When the response is received, the BIG-IP system caches the response with the fresh content, and uses the cached response to service the queued requests.
Requests for an uncacheable object	Sometimes, an object cannot be cached, for example, if the object exceeds the maximum object size or if the response includes a <code>no-store</code> response header. When the BIG-IP system first receives a large number of parallel requests for an object that cannot be cached, instead of sending each request to an origin web server, the BIG-IP system queues the requests, and then sends one request to an origin web server. When the BIG-IP system receives the response, it sends the queued requests to the origin web server. Subsequent requests for the uncacheable object bypass the BIG-IP system and are sent directly to the origin web server.
Requests for a nonexistent object	When the BIG-IP system receives a large number of parallel requests for an object that does not exist or no longer exists, the BIG-IP system can queue the requests, and then send one request to an

Object cache status	Description
	origin web server. When the BIG-IP system receives the response with a 404 (Not Found) status code, it services the queued requests with the 404 (Not Found) response. Note that the 404 (Not Found) response is not cached, and all subsequent requests for the nonexistent object are sent to the origin web server.

Enabling HTTP request queuing

You can queue parallel HTTP requests for the same object, and provide a cached response to those requests, resulting in accelerated responses from the BIG-IP and a reduction in requests sent to origin web servers.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Responses Cached**.
6. Select the **Queue Parallel Requests** check box.

Important: The *Cache content on first hit* check box must be clear to enable the **Queue Parallel Requests** setting. If the *Cache content on first hit* check box is selected, the **Queue Parallel Requests** setting is disabled.

7. (Optional) For the **Object Min Size** setting, clear the **Use Profile Setting** check box, and type the minimum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
8. (Optional) For the **Object Max Size** setting, clear the **Use Profile Setting** check box, and type the maximum object size (in bytes) to cache, overriding the Web Acceleration profile setting.
9. Click **Save**.
10. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP policy is enabled to queue concurrent HTTP requests for an object, cache a response, and provide the cached response to the queued requests.

Disabling HTTP request queuing

You can disable HTTP request queuing for an object, sending all requests for an object to the origin web servers or using the cached responses, as applicable.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Responses Cached**.

6. Clear the **Queue Parallel Requests** check box.
7. Click **Save**.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

HTTP request queuing for an object is disabled.

Managing HTTP Traffic with the HTTP/2 Profile

Overview: Managing HTTP traffic with the HTTP2 profile

You can configure a virtual server with the BIG-IP® Local Traffic Manager™ (LTM®) HTTP/2 profile to provide gateway functionality for HTTP 2.0 traffic, minimizing the latency of requests by multiplexing streams and compressing headers.

You can configure the BIG-IP® Acceleration HTTP/2 profile to provide full-proxy functionality for HTTP 2.0 traffic, minimizing the latency of requests by multiplexing streams and compressing headers.

Important: *The BIG-IP system supports HTTP/2 for client-side traffic only. More specifically, when client-side HTTP traffic reaches a virtual server with an HTTP/2 profile assigned to it, the BIG-IP system sends the server-side traffic to the pool members over HTTP/1.1.*

A client initiates an HTTP/2 request to the BIG-IP system, the HTTP/2 virtual server receives the request on port 443, and sends the request to the appropriate server. When the server provides a response, the BIG-IP system compresses and caches it, and sends the response to the client.

Note: *Source address persistence is not supported by the HTTP/2 profile.*

Summary of HTTP/2 profile functionality

By using the HTTP/2 profile, the BIG-IP system provides the following functionality for HTTP/2 requests.

Creating concurrent streams for each connection.

You can specify the maximum number of concurrent HTTP requests that are accepted on a HTTP/2 connection. If this maximum number is exceeded, the system closes the connection.

Limiting the duration of idle connections.

You can specify the maximum duration for an idle HTTP/2 connection. If this maximum duration is exceeded, the system closes the connection.

Enabling a virtual server to process HTTP/2 requests.

You can configure the HTTP/2 profile on the virtual server to receive HTTP and HTTP/2 traffic, or to receive only HTTP/2 traffic, based in the activation mode you select. (Note the HTTP/2 profile to receive only HTTP/2 traffic is primarily intended for troubleshooting.)

Inserting a header into the request.

You can insert a header with a specific name into the request. The default name for the header is x-HTTP/2.

Important: *The HTTP/2 protocol is incompatible with NTLM protocols. Do not use the HTTP/2 protocol with NTLM protocols.*

Task summary

Creating a pool to manage HTTPS traffic

Creating a virtual server to manage HTTP traffic

Creating an HTTP/2 profile

Creating a virtual server to manage HTTP/2 traffic

About HTTP/2 profiles

The BIG-IP® system includes an HTTP/2 profile type that you can use to manage HTTP/2 traffic, improving the efficiency of network resources while reducing the perceived latency of requests and responses. The LTM HTTP/2 profile enables you to achieve these advantages by multiplexing streams and compressing headers with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security.

The BIG-IP® system's Acceleration functionality includes an HTTP/2 profile type that you can use to manage HTTP/2 traffic, improving the efficiency of network resources while reducing the perceived latency of requests and responses. The Acceleration HTTP/2 profile enables you to achieve these advantages by multiplexing streams and compressing headers with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security.

Important: The BIG-IP system supports HTTP/2 for client-side traffic only. More specifically, when client-side HTTP traffic reaches a virtual server with an HTTP/2 profile assigned to it, the BIG-IP system sends the server-side traffic to the pool members over HTTP/1.1.

The HTTP/2 protocol uses a binary framing layer that defines a frame type and purpose in managing requests and responses. The binary framing layer determines how HTTP messages are encapsulated and transferred between the client and server, a significant benefit of HTTP 2.0 when compared to earlier versions.

All HTTP/2 communication occurs by means of a connection with bidirectional streams. Each stream includes messages, consisting of one or more frames, that can be interleaved and reassembled using the embedded stream identifier within each frame's header. The HTTP/2 profile enables you to specify a maximum frame size and write size, which controls the total size of combined data frames, to improve network utilization.

Multiplexing streams

You can use the HTTP/2 profile to multiplex streams (interleaving and reassembling the streams), by specifying a maximum number of concurrent streams permitted for a single connection.

Additionally, you can specify the way that the HTTP/2 profile controls the flow of streams. The **Receive Window** setting allows HTTP/2 to stall individual upload streams, as needed. For example, if the BIG-IP system is unable to process a slow stream on a connection, but is able to process other streams on the connection, it can use the **Receive Window** setting to specify a frame size for the slow stream, thus delaying that upload stream until the size is met and the receiver is able to process it, while concurrently proceeding to process frames for another stream.

Compressing headers

When you configure the HTTP/2 profile's **Header Table Size** setting, you can compress HTTP headers to conserve bandwidth. Compressing HTTP headers reduces the object size, which reduces required bandwidth. For example, you can specify a larger table value for better compression, but at the expense of using more memory.

HTTP/2 profile settings

This table provides descriptions of the HTTP/2 profile settings.

Setting	Default	Description
Name		Specifies the name of the HTTP/2 profile.

Setting	Default	Description
Parent Profile	http2	Specifies the profile that you want to use as the parent profile. Your new profile inherits all settings and values from the parent profile specified.
Concurrent Streams Per Connection	10	Specifies the number of concurrent requests allowed to be outstanding on a single HTTP/2 connection.
Connection Idle Timeout	300	Specifies the number of seconds an HTTP/2 connection is left open idly before it is closed.
Insert Header	Disabled	Specifies whether an HTTP header that indicates the use of HTTP/2 is inserted into the request sent to the origin web server.
Insert Header Name	X-HTTP/2	Specifies the name of the HTTP header controlled by the Insert Header Name setting.
Activation Modes	Select Modes	Specifies how a connection is established as a HTTP/2 connection.
Selected Modes	ALPN NPN	Used only with an Activation Modes selection of Select Modes , specifies the extension used in the HTTP/2 profile. The order of the extensions in the Selected Modes Enabled list ranges from most preferred (first) to least preferred (last). Clients typically use the first supported extension. At least one HTTP/2 mode must be included in the Enabled list. The values ALPN and NPN specify that the TLS Application Layer Protocol Negotiation (ALPN) and Next Protocol Negotiation (NPN) will be used. Clients that use TLS, but only support HTTP will work as if HTTP/2 is not present. The value Always specifies that all connections function as HTTP/2 connections. Selecting Always in the Activation Mode list is primarily intended for troubleshooting.
Receive Window	32	Specifies the <i>receive window</i> , which is HTTP/2 protocol functionality that controls flow, in KB. The receive window allows the HTTP/2 protocol to stall individual upload streams when needed.
Frame Size	2048	Specifies the size of the data frames, in bytes, that the HTTP/2 protocol sends to the client. Larger frame sizes improve network utilization, but can affect concurrency.
Write Size	16384	Specifies the total size of combined data frames, in bytes, that the HTTP/2 protocol sends in a single write function. This setting controls the size of the TLS records when the HTTP/2 protocol is used over Secure Sockets Layer (SSL). A large write size causes the HTTP/2 protocol to buffer more data and improves network utilization.
Header Table Size	4096	Specifies the size of the header table, in KB. The HTTP/2 protocol compresses HTTP headers to save bandwidth. A larger table size allows better compression, but requires more memory.

Creating a pool to manage HTTPS traffic

You can create a pool (a logical set of devices, such as web servers, that you group together to receive and process HTTPS traffic) to efficiently distribute the load on your server resources.

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

The Pool List screen opens.

2. Click **Create**.

The New Pool screen opens.

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

4. For the **Health Monitors** setting, assign **https** or **https_443** by moving it from the **Available** list to the **Active** list.

5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

The default is **Round Robin**.

6. For the **Priority Group Activation** setting, specify how to handle priority groups:

- Select **Disabled** to disable priority groups. This is the default option.
- Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.

7. Use the **New Members** setting to add each resource that you want to include in the pool:

- a) In the **Address** field, type an IP address.
- b) In the **Service Port** field type 443 , or select **HTTPS** from the list.
- c) (Optional) Type a priority number in the **Priority** field.
- d) Click **Add**.

8. Click **Finished**.

The HTTPS load balancing pool appears in the Pool List screen.

Creating a virtual server to manage HTTP traffic

You can create a virtual server to manage HTTP traffic redirected from an HTTP/2 virtual server.

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

The Virtual Server List screen opens.

2. Click the **Create** button.

The New Virtual Server screen opens.

3. In the **Name** field, type a unique name for the virtual server.

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

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

***Note:** The IP address for this field needs to be on the same subnet as the external self-IP.*

5. In the **Service Port** field, type 443 or select **HTTPS** from the list.

6. From the **HTTP Profile** list, select **http**.

7. For the **SSL Profile (Client)** setting, from the **Available** list, select **clientssl**, and using the Move button, move the name to the **Selected** list.

8. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.

9. Click **Finished**.

The HTTP virtual server is now available with the specified settings.

Creating an HTTP/2 profile

You can create an HTTP/2 profile for a virtual server, which responds to clients that send HTTP/2 requests.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **HTTP/2**.
The HTTP/2 profile list screen opens.
2. On the Main tab, click **Acceleration** > **Profiles** > **HTTP/2**.
The HTTP/2 profile list screen opens.
3. Click **Create**.
The New HTTP/2 Profile screen opens.
4. In the **Name** field, type a unique name for the profile.
5. From the **Configuration** list, select **Advanced**.
6. Select the **Custom** check box.
7. In the **Concurrent Streams Per Connection** field, type the number of concurrent connections to allow on a single HTTP/2 connection.
8. In the **Connection Idle Timeout** field, type the number of seconds that a HTTP/2 connection is left open idly before it is closed.
9. (Optional) From the **Insert Header** list, select **Enabled** to insert a header name into the request sent to the origin web server.
10. (Optional) In the **Insert Header Name** field, type a header name to insert into the request sent to the origin web server.
11. From the **Activation Modes** list, accept the default enabled modes.
12. In the **Selected Modes** setting, select the protocol modes that you want to enable.

Option	Description
All Modes Enabled	Enables all supported protocol versions: HTTP/2 and HTTP1.1.
Select Modes	Enables one or more specific protocol versions that you specify. For the Selected Modes setting, select a protocol entry in the Available field, and move the entry to the Selected field using the Move button.

13. From the **Priority Handling** list, select how the HTTP/2 profile handles priorities of concurrent streams within the same connection.

Option	Description
Strict	Processes higher priority streams to completion before processing lower priority streams.
Fair	Enables higher priority streams to use more bandwidth than lower priority streams, without completely blocking the lower priority streams.
14. In the **Receive Window** field, type the flow-control size for upload streams, in KB.
15. In the **Frame Size** field, type the size of the data frames, in bytes, that the HTTP/2 protocol sends to the client.
16. In the **Write Size** field, type the total size of combined data frames, in bytes, that the HTTP/2 protocol sends in a single write function.
17. In the **Header Table Size** field, type the size of the header table, in KB, for the HTTP headers that the HTTP/2 protocol compresses to save bandwidth.
18. Click **Finished**.

An HTTP/2 profile is now available with the specified settings.

Creating a virtual server to manage HTTP/2 traffic

You can create a virtual server to manage HTTP/2 traffic.

Important: Do not use the HTTP/2 protocol with NTLM protocols as they are incompatible.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.
The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is 10.0.0.1 or 10.0.0.0/24, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

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

5. In the **Service Port** field, type 443 or select **HTTPS** from the list.
6. From the **HTTP Profile** list, select **http**.
7. For the **SSL Profile (Client)** setting, from the **Available** list, select **clientssl**, and using the Move button, move the name to the **Selected** list.
8. From the **Acceleration** list, select **Advanced**.
9. From the **HTTP/2 Profile** list, select **http2**, or a user-defined HTTP/2 profile.
10. From the **Default Pool** list, select a pool that is configured for an HTTP/2 profile.
11. Click **Finished**.

The HTTP/2 virtual server is now ready to manage HTTP/2 traffic.

Accelerating Requests and Responses with Intelligent Browser Referencing

Overview: Reducing conditional GET requests with Intelligent Browser Referencing

You can increase the efficiency of the client's web browser's local cache and improve perceived access to your site by enabling the *Intelligent Browser Referencing* (IBR) feature, which reduces or eliminates requests to your site for relatively static content, such as images and cascading style sheet (CSS) files.

Task summary for reducing conditional GET requests with Intelligent Browser Referencing

Perform these tasks to reduce or eliminate requests to your site for relatively static content by enabling the web browser's cache to serve qualifying content.

Task summary

Configuring Intelligent Browser Referencing advanced settings

Enabling content assembly on proxies

Enabling Intelligent Browser Referencing

Adjusting the adaptive Intelligent Browser Referencing lifetime

Configuring Intelligent Browser Referencing advanced settings

You can customize the default prefix, duration of Intelligent Browser Referencing (IBR) lifetime, and duration of adaptive IBR lifetime by specifying advanced settings.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. Scroll down to the IBR Options area and modify the settings, as necessary.
 - a) In the **IBR Prefix** field, type a string to prepend to links or URLs that are embedded in your web pages.

***Note:** If you change the IBR prefix, be sure to test thoroughly to ensure that your application functions properly.*

The default string is ;wa.

- b) In the **IBR Default Lifetime** field, type a number and select a unit of time from the list to indicate the lifetime of links that match the node.
The initial default lifetime is 26 **Weeks**.
 - c) In the **IBR Adaptive Lifetime** field, type a number and select a unit of time from the list to indicate the duration of linked or imported URLs within externally linked CSS files.
The initial default lifetime is 10 **Days**.
5. Click **Save**.

This specifies the IBR default prefix and duration.

Enabling content assembly on proxies

When you enable content compression or Intelligent Browser Referencing (IBR), you should select the **Enable Content Assembly on Proxies** check box on the **Assembly** tab.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Select the **Enable Content Assembly on Proxies** check box.

***Note:** When selected (enabled), the BIG-IP system requests gzip-encoded or deflate-encoded content from the origin web server. The origin web server complies only if it supports the compression mode; otherwise, the origin web server provides uncompressed content.*

7. Click **Save**.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP system can compress content as required, and manage content by using IBR, even if the content is not served from the BIG-IP system's cache.

Enabling Intelligent Browser Referencing

The following prerequisites apply to enable Intelligent Browser Referencing (IBR) for a specific node.

- Specific variation rules must be configured with ordinals **Method**, **Cookie**, **User Agent**, **Referrer**, **Header**, and **Client IP** using a Values Define setting of **Same Content**. Variation rules cannot be defined for the node using a Values Define setting of **Different Content**. For example, if the link matches to a variation rule that identifies a cookie as being significant for content, the BIG-IP system cannot apply the IBR feature.
- The **Always proxy requests for this node** option is not selected.
- No **Proxy Override Rules** are defined for the node.

You can increase the efficiency of the client's web browser's local cache, and improve perceived access to your site by enabling IBR, which reduces or eliminates requests to your site for relatively static content, such as images and style sheet (CSS) files.

***Note:** When an object is matched to a rule in which the IBR feature is enabled, the BIG-IP system ignores the client cache minimum age settings for that object. However, the HTML page, into which those objects are loaded, honors all client cache minimum age settings.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.

The screen refreshes to show assembly options.

6. Select one or both of the following settings, as applicable.

Option	Description
Enable Intelligent Browser Referencing To	Specifies the system modifies external links that match the node with unique subdomains.
Enable Intelligent Browser Referencing Within	Specifies the system modifies URLs within a CSS file that is linked or imported into an HTML document.

***Important:** Setting any of the following **Variation** rule parameters to **Different Content** disables Intelligent Browser Referencing (IBR) and MultiConnect to the policy node. You must set these parameters to **Same Content** to enable IBR and MultiConnect to the policy node.*

- **Method**
 - **Cookie**
 - **User Agent**
 - **Referrer**
 - **Header**
 - **Client IP**
-

7. Click **Save**.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP system applies IBR to the policy node for the linked object in the HTML page or within the externally linked CSS file, typically, an image, style sheet, or JavaScript document.

Adjusting the adaptive Intelligent Browser Referencing lifetime

Before you can adjust the adaptive IBR lifetime, the **Enable Intelligent Browser Referencing Within** check box must be selected for a policy node used within an application, providing Intelligent Browser Referencing (IBR) for embedded image links within externally linked cascading style sheet (CSS) files.

You can adjust the **IBR Adaptive Lifetime** setting to the shortest lifetime of an image defined within externally linked CSS files, enabling assembly of linked image files before all of the image files are cached, and enabling the embedded image files to refresh before a client uses stale image files from a browser's cache.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
The IBR Options area appears.
4. In the **IBR Adaptive Lifetime** field, type a number and select a unit of time in the list to indicate the duration of linked or imported URLs within externally linked CSS files.
The initial default lifetime is 10 **Days**.
5. Click **Save**.

This specifies the shortest lifetime of an image defined within externally linked CSS files, enabling assembly of linked image files before all of the image files are cached, and enabling the embedded image files to refresh before a client uses stale image files from a browser's cache.

Implementation result

BIG-IP[®] acceleration reduces or eliminates requests to your site for relatively static content, such as images and style sheet (CSS) files, by enabling the web browser's cache to serve qualifying content.

Accelerating JavaScript and Cascading Style Sheet Files

Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files

You can improve acceleration by reducing the number and sizes of cascading style sheet (CSS), HTML, and JavaScript files transferred across a network, and by improving the ability for browsers to render content. The BIG-IP® system uses inlining and concatenation of CSS, HTML, and JavaScript files to reduce the number and sizes of files transferred across a network, thus improving the acceleration of traffic, and uses minification and reordering to improve the speed that browsers render content.

Task summary

Minifying cascading style sheet and JavaScript files

Task summary for prefetching cascading style sheet, JavaScript, and inline image files

Task summary for inlining cascading style sheet, JavaScript, and inline image files

Task summary for reordering cascading style sheet and JavaScript files

Task summary for concatenating cascading style sheet and JavaScript files

Minifying cascading style sheet and JavaScript files

You can use minification to remove whitespaces, comments, and unnecessary special characters from CSS and JavaScript files.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Select the **Enable Javascript and CSS Minification** check box.
7. Click **Save**.

Minification is enabled to remove whitespaces, comments, and unnecessary special characters from CSS and JavaScript files.

Task summary for prefetching cascading style sheet, JavaScript, and inline image files

Perform these tasks to prefetch cascading style sheet (CSS) files, JavaScript files, and embedded images in externally linked CSS files.

Task summary

Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files

Specifying cascading style sheet, JavaScript, and image URL resources

Creating Domain Lists for DNS Prefetching

Configuring DNS prefetching

Specifying cascading style sheet, JavaScript, and image URL resources

You can specify the URL resources for cascading style sheet (CSS), JavaScript, and inline image files to use in inlining.

1. On the Main tab, click **Acceleration > Web Application > Policies > URL Resources**.
The URL Resources screen displays lists of URLs available to reorder.
2. In the **CSS URLs** setting, add the CSS URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a CSS file, and click **Add**.
For example, `http://www.siterequest.com/css_file.css`.
The URL appears in the **CSS URLs** list.
3. In the **JavaScript URLs** setting, add the JavaScript URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a JavaScript file, and click **Add**.
For example, `http://www.siterequest.com/javascript_file.js`.
The URL appears in the **JavaScript URLs** list.
4. In the **Image URLs** setting, add the image URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for an image file, and click **Add**.
For example, `http://www.siterequest.com/image_file.png`.
The URL appears in the **Image URLs** list.
5. Click **Save**.

Each CSS and JavaScript URL that you specified appears in the **CSS URLs** list and the **JavaScript URLs** list, respectively.

Creating Domain Lists for DNS Prefetching

DNS prefetching improves page load time by resolving domain names to an IP address before the user tries to follow a link. When DNS prefetching headers are inserted by the end client, HTML5 compliant browsers can do DNS resolution of dynamic links in the background while other items are being downloaded.

1. On the Main tab, click **Acceleration > Web Application > Policies > Domain Lists**.
The Policies screen displays a list of existing user-defined domain lists.
2. In the Name field, type a descriptive and useful name for the domain list.
3. Select a Sync Folder, if applicable.
4. Type optional information that might help you later identify the associated video advertisement policy.
5. Type domain names that you would like to add to your domain list.
6. Click **Save**.

A domain list has been created to add to a DNS Prefetching acceleration profile.

Configuring DNS prefetching

Before you can enable DNS prefetching, you must create a Domain List. (You can do this by going to **Acceleration > Web Application > Policies > Domain Lists > New Domain List**.)

You can enable *DNS prefetching* to improve page load time on HTML5 compliant browsers by resolving domain names to an IP address prior to a browser requesting content from third parties. When DNS prefetching headers are inserted by the BIG-IP system, HTML5 compliant browsers can do DNS resolution of dynamic links in the background while other items are being downloaded.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. In the Advanced Assembly area, configure DNS Prefetch settings, as applicable.
 - a) Select the **Enable DNS Prefetch** check box.
 - b) For the **Domain Lists** setting, select a domain in the **Available** field, and move the entry to the **Selected** field using the Move button.
 - c) Select the **Force Injection on HTTPS** check box.
 - d) Select the **HTTPS Automatic Page Prefetch** check box.

DNS prefetching is typically disabled on pages served over HTTPS to avoid leaking information about which particular document is served. Turning on Force Injection on HTTPS will enable DNS prefetching specifically for the domains listed in a domain list. Turning on HTTPS Automatic Page Prefetch will turn on DNS prefetching for the entire document served. Force Injection on HTTPS must be enabled in order to enable HTTPS Automatic Page Prefetch.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

DNS Prefetching is enabled to resolve domain names to an IP address before the user follows a link.

Task summary for inlining cascading style sheet, JavaScript, and inline image files

Perform these tasks to inline cascading style sheet (CSS) files, JavaScript files, and embedded images in externally linked CSS files.

Task summary

Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files

Specifying cascading style sheet, JavaScript, and image URL resources

Inlining cascading style sheet files

Inlining JavaScript files

Inlining image files

Specifying cascading style sheet, JavaScript, and image URL resources

You can specify the URL resources for cascading style sheet (CSS), JavaScript, and inline image files to use in inlining.

1. On the Main tab, click **Acceleration > Web Application > Policies > URL Resources**.
The URL Resources screen displays lists of URLs available to reorder.
2. In the **CSS URLs** setting, add the CSS URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a CSS file, and click **Add**.
For example, `http://www.siterequest.com/css_file.css`.
The URL appears in the **CSS URLs** list.

3. In the **JavaScript URLs** setting, add the JavaScript URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a JavaScript file, and click **Add**.
For example, `http://www.siterequest.com/javascript_file.js`.
The URL appears in the **JavaScript URLs** list.
4. In the **Image URLs** setting, add the image URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for an image file, and click **Add**.
For example, `http://www.siterequest.com/image_file.png`.
The URL appears in the **Image URLs** list.
5. Click **Save**.

Each CSS and JavaScript URL that you specified appears in the **CSS URLs** list and the **JavaScript URLs** list, respectively.

Inlining cascading style sheet files

Before you can apply inlining to cascading style sheet (CSS) files, you need to specify the CSS files in the URL Resources list.

You can use inlining to replace specified URLs to CSS files with an inline copy of the document.

***Note:** In order for content to be inlined, the inlined content must expire later than the parent content.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure CSS inlining settings, as applicable.
 - a) Select the **Enable CSS Inlining** check box.
 - b) For the **CSS Inlining URLs** setting, select a CSS URL entry in the **Available** field, and move the entry to the **Selected** field using the Move button.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

CSS inlining is enabled to replace the specified URLs to CSS files with an inline copy of the document.

Inlining JavaScript files

Before you can apply inlining to JavaScript files, you need to have specified the JavaScript files in the URL Resources list.

You can use inlining to replace specified URLs to JavaScript files with an inline copy of the document.

***Note:** In order for content to be inlined, the inlined content must expire later than the parent content.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure the JavaScript inlining settings, as applicable.
 - a) Select the **Enable JavaScript Inlining** check box.
 - b) For the **JavaScript Inlining URLs** setting, select a JavaScript URL entry in the **Available** field, and move the entry to the **Selected** field using the Move button.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

JavaScript inlining is enabled to replace the specified URLs to JavaScript files with an inline copy of the document.

Inlining image files

Before you can apply inlining to image files, you need to specify the image files in the URL Resources list.

You can use image inlining to replace specified URLs to external images with image data.

Note: In order for content to be inlined, the inlined content must expire later than the parent content.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure the image inlining settings, as applicable.
 - a) Select the **Enable Image Inlining** check box.
 - b) In the **Image Inlining Max Size** field, type a preferred maximum size for an image.
You can adjust the maximum size according to the maximum image size that you want to relocate. The default is 2 KB.
 - c) For the **Image Inlining URLs** setting, select an image inlining URL entry in the **Available** field, and move the entry to the **Selected** field using the Move button.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

Image inlining is enabled to replace specified URLs to external images with image data.

Task summary for reordering cascading style sheet and JavaScript files

Perform these tasks to reorder cascading style sheet (CSS) files and JavaScript files in externally linked CSS files.

Task summary

Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files

Specifying cascading style sheet and JavaScript resources

Reordering URLs to cascading style sheet files

Reordering URLs to JavaScript files

Specifying cascading style sheet and JavaScript resources

You can specify the URL resources for cascading style sheet (CSS) and JavaScript files to use in reordering.

1. On the Main tab, click **Acceleration > Web Application > Policies > URL Resources**.
The URL Resources screen displays lists of URLs available to reorder.
2. In the **CSS URLs** setting, add the CSS URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a CSS file, and click **Add**.
For example, `http://www.siterequest.com/css_file.css`.
The URL appears in the **CSS URLs** list.
3. In the **JavaScript URLs** setting, add the JavaScript URLs that you want to use.
 - a) In the **Name** field, type a name.
 - b) In the **URL** field, type a URL for a JavaScript file, and click **Add**.
For example, `http://www.siterequest.com/javascript_file.js`.
The URL appears in the **JavaScript URLs** list.
4. Click **Save**.

Each CSS and JavaScript URL that you specified appears in the **CSS URLs** list and the **JavaScript URLs** list, respectively.

Reordering URLs to cascading style sheet files

Before you can reorder URLs with this procedure, you must first specify each cascading style sheet (CSS) URL that you want to reorder in the HTML page in the URL Resources list.

You can enable progressive rendering of content as an HTML page loads, by configuring a user-defined policy to reorder links to CSS files.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure CSS reordering settings, as applicable.
 - a) Select the **Enable CSS Reordering** check box.

- b) In the **CSS Reorder Cache Size** field, type a preferred cache size.
You can adjust the cache size according to the number of CSS links that you want to relocate. The default is **8 KB**.
- c) For the **CSS Reorder URLs** setting, select a CSS URL entry in the **Available** field, and move the entry to the **Selected** field using the Move button.

Important: *If you configure CSS reordering in a policy with **Enable Content Assembly on Proxies** enabled, you cannot also select the **Always proxy requests for the node** option. If you select the **Always proxy requests for the node** option and enable **Enable Content Assembly on Proxies**, then CSS reordering becomes disabled.*

- 8. Click **Save**.
- 9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

CSS information is moved to the start of the HTML page (preceding the `</head>` element) to enable progressive rendering as the page loads.

Reordering URLs to JavaScript files

Before you can use this procedure, you must first specify each JavaScript URL that you want to reorder in the HTML page in the URL Resources list.

You can give browsers the ability to download objects in parallel, by configuring a user-defined policy that reorders links to JavaScript files.

- 1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
- 2. Click the name of a user-defined acceleration policy.
- 3. Click a node in the Policy Tree.
- 4. From the Matching Rules menu, choose Acceleration Rules.
- 5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
- 6. From the **Content Assembly Options** list, select **Advanced**.
- 7. Configure JavaScript reordering settings, as applicable.
 - a) Select the **Enable JavaScript Reordering** check box.
 - b) In the **JavaScript Reorder Cache Size** field, type a preferred cache size.
You can adjust the cache size according to the number of JavaScript links that you want to relocate. The default is **8 KB**.
 - c) For the **JavaScript Reorder URLs** setting, select a JavaScript URL entry in the **Available** field, and move the entry to the **Selected** field using the Move button.
- 8. Click **Save**.
- 9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

JavaScript information is moved to the end of the HTML page (preceding the `</body>` element), giving browsers the ability to download objects in parallel. Exceptions to reordering JavaScript information include JavaScript URLs and scripts that use `document.write` to insert content for the page.

Task summary for concatenating cascading style sheet and JavaScript files

Perform these tasks to concatenate Cascading Style Sheet (CSS) files and JavaScript (JS) files in externally linked CSS files.

Task summary

Overview: Accelerating JavaScript, HTML, cascading style sheet, and inline image files

About concatenation of JavaScript and cascading style sheet files

Specifying Cascading Style Sheet and JavaScript resources

Creating a Cascading Style Sheet concatenation set

Concatenating Cascading Style Sheet files

Creating a JavaScript concatenation set

Concatenating JavaScript files

About concatenation of JavaScript and cascading style sheet files

Concatenation combines a specified list of JavaScript (JS) or Cascading Style Sheet (CSS) files into a single concatenated file, which reduces the number of requests and responses, and the time required to transfer serialized files.

For each user-defined policy, you can specify lists of JS URLs and lists of CSS URLs for concatenation. Listings in each URL list appear in the specified order.

During the process of concatenation, the first JS and CSS URL within the HTML file that is specified in a JS or CSS URL list is replaced with the optimized URL, and each subsequent specified JS and CSS URL is removed. The TTL of a concatenated response is determined by the earliest expiration of the concatenated objects.

Task summary for concatenating cascading style sheet and JavaScript files

Specifying Cascading Style Sheet and JavaScript resources

You can specify the URL resources for Cascading Style Sheet (CSS) and JavaScript (JS) files to use in concatenating.

1. On the Main tab, click **Acceleration > Web Application > Policies > URL Resources**.

The URL Resources screen displays lists of URLs available to reorder.

2. In the **CSS URLs** setting, add the CSS URLs that you want to use.

a) In the **Name** field, type a name.

b) In the **URL** field, type a URL for a CSS file, and click **Add**.

For example, `http://www.siterequest.com/css_file.css`.

The URL appears in the **CSS URLs** list.

3. In the **JavaScript URLs** setting, add the JavaScript URLs that you want to use.

a) In the **Name** field, type a name.

b) In the **URL** field, type a URL for a JavaScript file, and click **Add**.

For example, `http://www.siterequest.com/javascript_file.js`.

The URL appears in the **JavaScript URLs** list.

4. Click **Save**.

Each CSS and JavaScript URL that you specified appears in the **CSS URLs** list and the **JavaScript URLs** list, respectively.

Creating a Cascading Style Sheet concatenation set

You can create a concatenation set for Cascading Style Sheet (CSS) URLs to use in optimization.

1. On the Main tab, click **Acceleration > Web Application > Policies > Concatenation Sets**.
The Concatenation Sets screen displays lists of URLs available to reorder.
2. In the **CSS Set Name** field, type the name of the concatenation set.
3. For the **Concatenation Set Type** setting, select **CSS**.
4. For the **URLs** setting, select a CSS URL entry in the **Available CSS URLs** field, and move the entry to the **Selected** field using the Move button.
5. Click **Save**.

Each CSS URL that you specified for optimization appears in the concatenation set.

Concatenating Cascading Style Sheet files

Before you can apply concatenation to Cascading Style Sheet (CSS) files, you need to specify each CSS URL in a Concatenation Set.

Concatenation combines a specified list of Cascading Style Sheet (CSS) files into a single concatenated file, which reduces the number of requests and responses and the time required to transfer serialized files.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure the CSS Concatenation Sets, as applicable.
 - a) Select the **URL Concatenation** check box.
 - b) For the **CSS Concatenation Sets** setting, select a URL Concatenation entry in the **Available** field, and move the entry to the **Selected** field using the Move button.

The selected CSS Concatenation Sets combine the specified list of cascading style sheet (CSS) files into a single concatenated file.

Creating a JavaScript concatenation set

You can create a concatenation set for JavaScript (JS) URLs to use in optimization.

1. On the Main tab, click **Acceleration > Web Application > Policies > Concatenation Sets**.
The Concatenation Sets screen displays lists of URLs available to reorder.
2. In the **JavaScript Set Name** field, type the name of the concatenation set.
3. For the **Concatenation Set Type** setting, select **JavaScript**.
4. For the **URLs** setting, select a JavaScript URL entry in the **Available JavaScript URLs** field, and move the entry to the **Selected** field using the Move button.
5. Click **Save**.

Each JavaScript URL that you specified for optimization appears in the concatenation set.

Concatenating JavaScript files

Before you can use this procedure, you need to specify each JavaScript URL that you want to optimize in a Concatenation Set.

Concatenation combines a specified list of JavaScript (JS) files into a single concatenated file, which reduces the number of requests and responses, and the time required to transfer serialized files.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Configure the JavaScript Concatenation Sets, as applicable.
 - a) Select the **URL Concatenation** check box.
 - b) For the **JavaScript Concatenation Sets** setting, select a URL Concatenation entry in the **Available** field, and move the entry to the **Selected** field using the Move button.

The selected JavaScript Concatenation Sets combine the specified list of JavaScript files into a single concatenated file

Implementation results

The system is configured to accelerate cascading style sheet, JavaScript, and inline image files.

Establishing Additional TCP Connections with MultiConnect

Overview: Accelerating requests and responses with MultiConnect

Most web browsers create a limited number of persistent TCP connections when requesting data, which restricts the amount of content a client can receive at one time. You can provide faster data downloads to your clients using the BIG-IP® device's MultiConnect feature.

The *MultiConnect* feature enables you to specify unique subdomains that prompt the browser to open more persistent TCP connections (up to five per HTTP subdomain and five per HTTPS subdomain generated by the BIG-IP device). The origin web servers never get a request from these additional subdomains; they are used exclusively on externally linked URLs or links that request images or scripts and are only for requests or responses between the client and the BIG-IP device. If the BIG-IP device needs to send a request to the origin server, it removes the subdomain prefixes before sending the request.

The BIG-IP device uses the MultiConnect feature only on the following types of links:

- Image tags: ``
- Script tags: `<script src="...">`
- Forms whose input type is an image: `<form><input type="image" src="..."></form>`

Task summary for establishing additional TCP connections with MultiConnect

Perform these tasks to establish additional TCP connections with MultiConnect.

Task summary

Enabling content assembly on proxies

Configuring DNS subdomains for use with MultiConnect

Enabling MultiConnect for HTTP traffic

Enabling MultiConnect for HTTPS traffic

Enabling content assembly on proxies

When you enable content compression or Intelligent Browser Referencing (IBR), you should select the **Enable Content Assembly on Proxies** check box on the **Assembly** tab.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Select the **Enable Content Assembly on Proxies** check box.

***Note:** When selected (enabled), the BIG-IP system requests gzip-encoded or deflate-encoded content from the origin web server. The origin web server complies only if it supports the compression mode; otherwise, the origin web server provides uncompressed content.*

7. Click **Save**.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP system can compress content as required, and manage content by using IBR, even if the content is not served from the BIG-IP system's cache.

Configuring DNS subdomains for use with MultiConnect

Before you can configure DNS subdomains for use with MultiConnect, you must ensure that you have completed these tasks:

- Configure DNS with entries for the additional subdomains.
- Map the additional DNS entries to the same IP address as the base origin web server (for example, `www.siterequest.com`).
- Assign specific prefixes to the additional subdomains. For example, if the requested host for the mapping is `www.siterequest.com` and you request additional subdomains for the HTTP protocol, you assign a subdomain prefix of `wa`.

The BIG-IP system changes the domain on qualifying embedded URLs and links so that they use the domains that you specify.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. Click **Options** for a requested host.
4. From the **HTTP Subdomains** list, select the number of HTTP subdomains that you want the BIG-IP system to generate for each protocol.
The default is 0.
5. From the **HTTPS Subdomains** list, select the number of HTTPS subdomains that you want the BIG-IP system to generate for each protocol.
The default is 0.

***Important:** Some client browsers close HTTPS connections to one domain before opening HTTPS connections to a new domain. This type of browser behavior can decrease the speed of access to applications for which the MultiConnect feature is enabled; therefore, do not enable the MultiConnect feature for HTTPS connections.*

6. In the **Subdomain Prefix** field, type a prefix.
The default prefix is `wa`.
7. Click **Save**.

The BIG-IP system now changes the domain on qualifying embedded URLs and links so that they use the domains you specified. For example:

- `wa1.www.siterequest.com`
- `wa2.www.siterequest.com`

Enabling MultiConnect for HTTP traffic

Before configuring MultiConnect settings, you must configure DNS subdomains for use with the MultiConnect feature.

The MultiConnect feature opens additional persistent TCP connections and is optimum for sites that have a high number of first-time visitors who are downloading a large number of images or scripts.

Note: Use the MultiConnect feature only if you have low-latency, high-bandwidth links, because the additional TCP connections also increase the amount of traffic to your site.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Select one or both of the following options, as applicable.

Option	Description
Enable MultiConnect To	Modifies externally linked URLs to images or scripts that match a node with unique subdomains.
Enable MultiConnect Within	Modifies URLs to images or scripts within a CSS file that is linked or imported into an HTML document.

Important: Setting any of the following Variation rule parameters to the **Different Content** option disables Intelligent Browser Referencing (IBR) and MultiConnect to the policy node. You must set these parameters to the **Same Content** option to enable IBR and MultiConnect to the policy node.

- **Method**
- **Cookie**
- **User Agent**
- **Referrer**
- **Header**
- **Client IP**

-
7. Click **Save**.
 8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

This applies the MultiConnect feature to the node, and modifies externally linked URLs with unique subdomains, prompting the browser to open more persistent TCP connections (up to five per subdomain generated by the BIG-IP device).

Enabling MultiConnect for HTTPS traffic

Before configuring MultiConnect settings, you must configure DNS subdomains for use with the MultiConnect feature.

The MultiConnect feature opens additional persistent TCP connections and is optimum for sites that have a high number of first-time visitors who are downloading a large number of images or scripts.

Note: Use the MultiConnect feature only if you have low-latency, high-bandwidth links, because the additional TCP connections also increase the amount of traffic to your site.

Important: F5 Networks® recommends that you do not enable the MultiConnect feature for HTTPS connections. Some client browsers close HTTPS connections to one domain before opening HTTPS connections to a new domain, which can decrease the speed of access to applications for which the MultiConnect feature is enabled.

1. On the Main tab, click **Acceleration > Web Application > Policies**.

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Select one or both of the following options, as applicable.

Option	Description
Enable MultiConnect To	Modifies externally linked URLs to images or scripts that match a node with unique subdomains.
Enable MultiConnect Within	Modifies URLs to images or scripts within a CSS file that is linked or imported into an HTML document.

Important: Setting any of the following Variation rule parameters to the **Different Content** option disables Intelligent Browser Referencing (IBR) and MultiConnect to the policy node. You must set these parameters to the **Same Content** option to enable IBR and MultiConnect to the policy node.

- **Method**
 - **Cookie**
 - **User Agent**
 - **Referrer**
 - **Header**
 - **Client IP**
-

7. Click **Save**.
8. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

This applies the MultiConnect feature to the node, and modifies externally linked URLs with unique subdomains, prompting the browser to open more persistent TCP connections (up to five per subdomain generated by the BIG-IP device).

Implementation result

This implementation configures BIG-IP® acceleration to modify externally linked URLs with unique subdomains, prompting the browser to open more persistent TCP connections.

Serving Specific Hyperlinked Content with Parameter Value Substitution

Overview: Serving specific hyperlinked content with parameter value substitution

You can use parameter value substitution functionality to change a targeted parameter's value on a specific page serviced from cache, so that the client-specified parameter appears on the URL embedded in the page. This substitution is especially beneficial when a query parameter contains identification information for a site's visitors.

Some requested pages include hyperlinks that vary according to the request to provide dynamic information. For example, you can configure parameter value substitution so that a request with a query parameter called `shopper` produces HTML output with its embedded hyperlinks varying the value for `shopper`. Thus, when a query parameter contains identification information for a site's visitors, it prompts the BIG-IP® device to serve different content for the request, based on the specific visitor.

Conversely, if parameter value substitution is not configured, the BIG-IP device uses the value that it cached for the original request, for all subsequent requests after the first, even if the subsequent requests have different values that the origin web server used in the response.

Serving specific hyperlinked content with parameter value substitution

When you configure parameter value substitution, you specify a source definition in an HTTP request, specifically a value that you want the BIG-IP system to embed in the URL in place of the cached (target) value, and a target definition, specifically a parameter by data type and name or location in the request. You also have the option to provide a URL prefix for the target, to limit the URLs to which the BIG-IP system performs the substitution.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. Click **Create**.
7. For the **Source Definition** setting, select a source type from the **Type** list.

***Note:** The configuration steps change depending on the type you select.*

Source type	Description	Configuration steps
Number Randomizer	Generates a random number and places it in a targeted location in the embedded URL.	None.
Request URL	Limits the BIG-IP system to target-specific URLs embedded in a page, as defined in the prefix that an embedded	<ul style="list-style-type: none">• Select the Relative URL or Absolute URL option.

Source type	Description	Configuration steps
Query Parameter	URL must match before the BIG-IP system performs substitution. The BIG-IP system substitutes the URL parameter with the parameter value you specify.	In the Name field, type the query parameter value.
Unnamed Query Parameter	The BIG-IP system substitutes the URL parameter, as specified by the Ordinal setting.	<ul style="list-style-type: none"> • Alias: Type an alias. • Ordinal: Type a number that represents the ordinal location of the unnamed query parameter. Numbering for the ordinal location starts at one (1).
Path Segment	The BIG-IP system substitutes the URL parameter, as specified by the Segment Ordinal setting.	<ul style="list-style-type: none"> • Alias: Type an alias. • Segment Ordinal: Type a number that represents the path segment location in the URL, and select a numbering scheme from the list.

8. For the **Target Definition** setting, select a target type from the **Type** list.

Note: The configuration steps change depending on the type you select.

Target type	Description	Configuration steps
Query Parameter	The BIG-IP system substitutes the URL parameter, as specified by the Name setting.	<ul style="list-style-type: none"> • Name: Type the query parameter value. • All URLs: Select this option if you want the substitution to apply to all URLs. • Selected URLs: Select this option if you want the substitution to apply to only those URLs that you specify.
Unnamed Query Parameter	The BIG-IP system substitutes the URL parameter, as specified by the Ordinal setting.	<ul style="list-style-type: none"> • Alias: Type an alias. • Ordinal: Type a number that represents the ordinal location of the unnamed query parameter. Numbering for the ordinal location starts at one (1). • All URLs: Select this option if you want the substitution to apply to all URLs. • Selected URLs: Select this option if you want the substitution to apply to only those URLs that you specify.
Path Segment	The BIG-IP system substitutes the URL parameter, as specified by the Segment Ordinal setting.	<ul style="list-style-type: none"> • Alias: Type an alias. • Segment Ordinal: Type a number that represents the path segment location in the URL, and select a numbering scheme from the list.

9. Click **Save**.

A source definition and target definition are configured for parameter value substitution, so that when a query parameter contains identification information for a site's visitors, it prompts the BIG-IP system to serve different content for the request, based on the specific visitor.

Accelerating Access to PDF Content

Overview: Accelerating access to PDF content with PDF linearization

Large PDF files can provide a slow response in displaying content when the entire file must download before a requested page can be accessed. The BIG-IP® device provides the ability to display a requested page more quickly by using PDF linearization (optimization). PDF linearization prepares the PDF file for byte serving, which enables the BIG-IP device to provide individual pages to a client when it receives byte-range requests.

All PDF files are constructed in one of two formats:

- Nonlinear. A *nonlinear* (not optimized) PDF file typically provides slower access to specific pages than a linear PDF file because a page-offset index for the document's pages is omitted. For example, PDF files that are created for high quality print output are often nonlinear.
- Linear. A *linear* (optimized) PDF file, in comparison, provides faster access to specific pages because a page-offset index for the document's pages is written at the beginning, enabling a web browser to send byte-range requests to access and display initial or specific pages before the entire file is downloaded.

When you enable PDF linearization, the BIG-IP device provides a linear PDF file, thus allowing expedient access to a requested page.

Task summary for accelerating access to PDF content

Perform these tasks to accelerate access to PDF content by using PDF linearization.

Task summary

Accelerating content with PDF linearization

Disabling PDF linearization for a specific node

Accelerating content with PDF linearization

You can enable PDF linearization, or optimization, for a specific node to accelerate access to content within a PDF file without requiring a download of the entire PDF file.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Select the **Enable PDF Linearization** check box.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.

- b) In the **Comment** field, type a description.
- c) Click **Publish Now**.

The BIG-IP® device applies PDF linearization to the node, enabling a client to access a requested page without requiring a download of the entire PDF file.

Disabling PDF linearization for a specific node

PDF linearization, or optimization, provides access to content within the PDF file without requiring a download of the entire PDF file. You can, if necessary, disable PDF linearization for a specific node, requiring the entire file to download before being able to access the content.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Clear the **Enable PDF Linearization** check box.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The BIG-IP device disables PDF linearization for the associated node.

Accelerating Images with Image Optimization

Overview: Accelerating images with image optimization

You can configure *image optimization* in a BIG-IP® policy to reduce the size of image files, for example, by removing unnecessary metadata, by changing the format, or by increasing compression, and, consequently, accelerate the transfer of image objects across a network.

When an image object is matched to a policy node, it is modified in accordance with the acceleration rules of the policy. Configurable acceleration rules for an image object include several parameters.

***Note:** Image optimization only benefits raster images. Vector images, such as SVG files, benefit little from image optimization, but can benefit from file compression. You can use file compression to improve the performance of vector images.*

Task summary for optimizing images

Perform these tasks to reduce the file size of images by using image optimization.

Task summary

Accelerating images by optimization

Disabling image optimization for a node

Accelerating images by optimization

You can reduce the size of image files by removing unnecessary metadata, changing the file format, or using compression, and, consequently, accelerate the transfer of image objects across a network.

***Note:** F5® Networks recommends examination of converted file sizes for different file formats to optimize performance with a reduced file size.*

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Scroll to the Image Optimization Settings area of the screen, select an image optimization option, and configure the displayed settings, as applicable.

***Note:** The configuration steps change depending on the setting you select.*

Image optimization option	Description	Configuration steps
No Optimization to JPEG	<p>Default for all policies except Generic Complete. No image optimization is performed.</p> <p>Converts the image associated with the policy node into a JPEG, in accordance with specified settings.</p>	<p>None.</p> <ul style="list-style-type: none"> For the JPEG Quality setting, select one of the options. <ul style="list-style-type: none"> Absolute. Specifies a quality level for compression Relative. Specifies a percentage of compression relative to the original JPEG file. In the JPEG Quality Factor field, type a quality factor number for the selected JPEG Quality option. Select the Strip EXIF keeps copyright check box to preserve copyright metadata when the BIG-IP[®] system strips other metadata from the EXIF header. For the Strip/Safe-strip JPEG EXIF Header setting, select one of the options. <ul style="list-style-type: none"> Don't Strip EXIF. The EXIF header is not changed. Always Strip EXIF. The EXIF header is always removed from the JPEG file. Strip EXIF if safe. The EXIF header is always removed, unless the header includes a color profile. Apply color profile, then strip EXIF. After the color profile is applied, the EXIF header is always removed. This option converts the image to the default color profile, so that the EXIF header can be safely removed. For the JPEG Sampling Factor setting, select one of the options. <ul style="list-style-type: none"> Preserve. The sampling factor matches the brightness and color values of the original file. 1x1. Provides the same sampling factor for the brightness and color values. 2x1. Averages color values for horizontal pixels. 1x2. Averages color values for vertical pixels. 2x2. Averages color values for both vertical and horizontal pixels. Select the Use Progressive Encoding check box to enable browsers to quickly render a low-quality version of the entire image. Select the Optimize For Client check box to convert images associated with the policy node into a WebP or a JPEG-XR image format. In the WebP Quality Factor field, type a quality factor number for the selected Optimize for Client option. In the JPEG-XR Factor field, type a quality factor number for the selected Optimize for Client option.
to GIF	Converts the image associated with the policy node into a GIF.	<ul style="list-style-type: none"> Select the Optimize For Client check box to convert images associated with the policy node into a WebP or a JPEG-XR image format.

Image optimization option	Description	Configuration steps
to PNG	Converts the image associated with the policy node into a PNG, in accordance with specified settings.	<ul style="list-style-type: none"> • In the WebP Quality Factor field, type a quality factor number for the selected Optimize for Client option. • In the JPEG-XR Factor field, type a quality factor number for the selected Optimize for Client option. • Select the Reduce to 256 Colors check box to reduce the file size with minimal degradation in the quality of the image. • Select the Optimize For Client check box to convert images associated with the policy node into a WebP or JPEG-XR image format. • In the WebP Quality Factor field, type a quality factor number for the selected Optimize for Client option. • In the JPEG-XR Factor field, type a quality factor number for the selected Optimize for Client option.
to TIFF	Converts the image associated with the policy node into a TIFF.	<ul style="list-style-type: none"> • Select the Optimize For Client check box to convert images associated with the policy node into a WebP or JPEG-XR image format. • In the WebP Quality Factor field, type a quality factor number for the selected Optimize for Client option. • In the JPEG-XR Factor field, type a quality factor number for the selected Optimize for Client option.

8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

The size of the image files is reduced to accelerate the transfer of image objects across a network.

Disabling image optimization for a node

You must have already enabled image optimization for image objects on a node.

You can, if necessary, disable image optimization for a specific node, thus requiring the unoptimized, original image files to download.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Scroll to the Image Optimization Settings area of the screen, and for the **Optimize image** setting, select the **No Optimization** option.
8. Click **Save**.

9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

This disables image optimization for the node.

Accelerating Video Streams with Video Delivery Optimization

Overview: Optimizing video delivery

BIG-IP® *video delivery optimization* provides you with the ability to retrieve and accelerate an on-demand video stream from an origin web server. The BIG-IP system sends client requests for the video stream to an origin web server, caches the response video segments, and sequentially sends optimized video responses to all authorized users.

Additionally, video delivery optimization enables you to associate video advertisements with a video stream, providing the ability to preroll advertisements or to insert advertisements as specified by a video advertisement policy.

Task summary for optimizing video streams

Perform these tasks to accelerate video segments by using video delivery optimization.

Task summary

Creating a video advertisement policy

Modifying a video advertisement policy

Deleting a video advertisement policy

Enabling video delivery optimization

Modifying video delivery optimization

Creating a video advertisement policy

You can create a video advertisement policy to manage the video advertisements that you want to associate with a video stream.

1. On the Main tab, click **Acceleration > Web Application > Policies > Video Ad Policies**.
The Video Ad Policies screen displays a list of existing video advertisement policies.
2. Click **Create**.
3. In the **Policy Name** field, type a name for the video advertisement policy.
4. Specify a folder, based on your configuration.
 - For a symmetric or farm configuration, from the **Sync Folder** list, select the name of a symmetric folder.
 - For an asymmetric configuration, from the **Sync Folder** list, select **No Selection**.
5. (Optional) In the **Description** field, type a description.
6. For the **Mode** setting, select the applicable option:
 - **Sequential** displays video advertisements sequentially.
 - **Random** displays video advertisement randomly.
7. In the **Name** field, type a name.
8. In the **URL** field, type the URL for the video advertisement.
9. Select the **Preroll** check box to enable prerolling of the video advertisements.
10. Click **Add**.

The advertisement resource appears in the **Ad URL's** list.

11. Click **Save.**

The video advertisement policy is configured, as specified, to manage the advertisements that you want to associate with a video stream.

Modifying a video advertisement policy

You can modify a user-defined video advertisement policy, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Policies > Video Ad Policies.**

The Video Ad Policies screen displays a list of existing video advertisement policies.

2. Click the name of a user-defined video advertisement policy.

3. (Optional) In the **Description field, type a description.**

4. For the **Mode setting, select the applicable option:**

- **Sequential** displays video advertisements sequentially.
- **Random** displays video advertisement randomly.

5. In the **Name field, type a name.**

6. In the **URL field, type the URL for the video advertisement.**

7. Select or clear the **Preroll check box to enable or disable prerolling of the video advertisements.**

8. Click **Add.**

The advertisement resource appears in the **Ad URL's** list.

9. From the **Ad URL's list, click **Delete** for the advertisement resource that you want to modify, and then complete the following steps in the **Ad Resources** area.**

- In the **Name** field, type a name.
- In the **URL** field, type the URL for the video advertisement.
- Select or clear the **Preroll** check box to enable or disable insertion of the specified advertisement into the video stream.
- Click **Add**.

10. Click **Save.**

The video advertisement policy is modified, as specified, to manage the advertisements that you want to associate with a video stream.

Deleting a video advertisement policy

You can delete a video advertisement policy, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Policies > Video Ad Policies.**

The Video Ad Policies screen displays a list of existing video advertisement policies.

2. Select the check box for each video ad policy that you want to delete.

3. Click **Delete.**

4. Click **Delete.**

The specified video advertisement policies are deleted.

Enabling video delivery optimization

In an acceleration policy, you can use video delivery optimization to retrieve and accelerate an on-demand video stream from an origin web server.

1. On the Main tab, click **Acceleration > Web Application > Policies.**

The Policies screen displays a list of existing acceleration policies.

2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Video**.
The screen refreshes to show video options.
6. (Optional) In the Video Optimization Options area, configure the options, as necessary.
 - a) Select the **Enable Fast Start** check box to enable caching in accordance with cache priority settings.
 - b) In the **Maximum Bitrate** field, type a maximum bit rate (in kbps) for the video stream.
7. (Optional) In the Video Advertisement Options area, configure the options, as necessary.
 - a) Select the **Enable Ad insertion** check box to insert the specified advertisement into the video stream.
 - b) In the **Ad Insertion Period** field, type the period (in seconds) to display the advertisement.
 - c) Select the **Enable Preroll Ads** check box to preroll the specified advertisements.
 - d) From the **Ad Policy** list, select a video advertisement policy.
8. Click **Save**.
9. On the menu bar, click **Responses Cached**.
10. From the **Cache Priority** list, select a priority to determine the caching priority for the video segments associated with the node.

Important: If you have selected the **Cache content on first hit** check box, the **Cache Priority** setting is overridden and not used. You must clear the **Cache content on first hit** check box to enable the **Cache Priority** functionality.

11. Click **Save**.

The video advertisement policy is configured according to the specified settings.

Modifying video delivery optimization

You can modify the settings for video delivery optimization, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a branch node or leaf node, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Video**.
The screen refreshes to show video options.
6. (Optional) In the Video Optimization Options area, configure the options, as necessary.
 - a) Select the **Enable Fast Start** check box to enable caching in accordance with cache priority settings.
 - b) In the **Maximum Bitrate** field, type a maximum bit rate (in kbps) for the video stream.
7. (Optional) In the Video Advertisement Options area, configure the options, as necessary.
 - a) Select the **Enable Ad insertion** check box to insert the specified advertisement into the video stream.
 - b) In the **Ad Insertion Period** field, type the period (in seconds) to display the advertisement.
 - c) Select the **Enable Preroll Ads** check box to preroll the specified advertisements.
 - d) From the **Ad Policy** list, select a video advertisement policy.
8. Click **Save**.

9. On the menu bar, click **Responses Cached**.
10. From the **Cache Priority** list, select a priority to determine the caching priority for the video segments associated with the node.

Important: If you have selected the **Cache content on first hit** check box, the **Cache Priority** setting is overridden and not used. You must clear the **Cache content on first hit** check box to enable the **Cache Priority** functionality.

11. Click **Save**.

Video optimization for the node is modified, as specified.

Overview: Video Quality of Experience profile

The BIG-IP® system's video Quality of Experience (QoE) profile enables you to assess an audience's video session or overall video experience, providing an indication of customer satisfaction. The QoE profile uses static information, such as bitrate and duration of a video, and video metadata, such as URL and content type, in monitoring video streaming. Additionally, the QoE profile monitors dynamic information, which reflects the real-time network condition.

By considering both the static video parameters and the dynamic network information, the user experience can be assessed and defined in terms of a single mean opinion score (MOS) of the video session, and a level of customer satisfaction can be derived. QoE scores are logged in the `ltm` log file, located in `/var/log`, which you can evaluate as necessary.

Task summary

Creating an iRule to collect video Quality of Experience scores

Creating an iRule to collect static information about video files

Creating a video Quality of Experience profile

Creating a pool

Creating a video Quality of Experience virtual server

Creating an iRule to collect video Quality of Experience scores

You can create an iRule to use with a video Quality of Experience (QoE) profile that defines the QoE scores to collect.

1. On the Main tab, click **Local Traffic > iRules**.
The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a name, such as `my_irule`.
The full path name of the iRule cannot exceed 255 characters.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.
For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site (<http://devcentral.f5.com>).
For example, the following iRule saves `Content-Type` to session DB with a 600-second lifetime.

```
...
when HTTP_REQUEST {
    set LogString "Client [IP::client_addr]:[TCP::client_port] ->
    [HTTP::host] [HTTP::uri]"
}
```



```

    set x_playback_session_id [HTTP::header "X-Playback-Session-Id"]
}

when HTTP_RESPONSE {
    set content_type [HTTP::header "Content-Type"]
}

when CLIENT_CLOSED {
    catch {
        if { ($content_type contains "video") &&
            ([QOE::video available] == 1) } {
            set qoe_params [list available width height duration nominal_bitrate
            average_bitrate freeze_period freeze_frequency mos]
            foreach param $qoe_params {
                set value [QOE::video $param]
                append params "$param=$value "
            }
            if {[string length $x_playback_session_id]}{
                log local0. "$LogString X-Playback-Session-Id:
                $x_playback_session_id QOE::video $params"
            } else {
                log local0. "$LogString QOE::video $params"
            }
        }
    }
}

```

5. Click Finished.

The new iRule appears in the list of iRules on the system.

There is now an available iRule to use with a QoE profile that collects specified QoE scores.

Creating an iRule to collect static information about video files

You can create an iRule to collect static information specific to video files, primarily for use with Policy Enforcement Manager™ (PEM).

1. On the Main tab, click Local Traffic > iRules.

The iRule List screen opens, displaying any existing iRules.

2. Click Create.

The New iRule screen opens.

3. In the Name field, type a name, such as my_irule.

The full path name of the iRule cannot exceed 255 characters.

4. In the Definition field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.

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

For example, the following iRule collects static information specific to video files.

```

when QOE_PARSE_DONE {
    set w [QOE::video width]
    set h [QOE::video height]
    set d [QOE::video duration]
    set b [QOE::video nominal_bitrate]
    log local0. "QOE_PARSE_DONE_ENABLED: width=$w height=$h
    bitrate=$b duration=$d"
}

```

5. Click Finished.

The new iRule appears in the list of iRules on the system.

There is now an iRule available to collect static information specific to video files.

Creating a video Quality of Experience profile

You can use the Traffic Management shell (`tmsh`) to create a video Quality of Experience (QoE) profile to use with Policy Enforcement Manager™ (PEM™) or Application Acceleration Manager™ (AAM™) and determine a customer's video Quality of Experience.

1. Log in to the command-line interface of the system using the root account.
2. Open the Traffic Management Shell (`tmsh`).

```
tmsh
```

3. Create a video QoE profile.

```
create ltm profile qoe qoe_profile_name video true
```

This creates the video QoE profile.

Creating a pool

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

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
5. Click **Finished**.
6. Repeat these steps for each pool you want to create.

The new pool appears in the Pools list.

Creating a video Quality of Experience virtual server

Before creating a video Quality of Experience (QoE) virtual server, you need to have created and configured a video QoE profile.

You can assign video QoE profile to a virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. From the **HTTP Profile** list, select **http**.
4. In the Resources area, for the **iRules** setting, from the **Available** list, select the name of the iRule that you want to assign, and move the name into the **Enabled** list.
5. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
6. Click **Finished**.
7. Log in to the command-line interface of the system using the root account.
8. Open the Traffic Management Shell (`tmsh`).

tmsh

9. Assign the video QoE profile to the virtual server.

```
modify virtual_server_name profile add qoe_profile_name
```

This assigns the video QoE profile and iRules to the virtual server.

Compressing Content from an Origin Web Server

Overview: Enabling content compression from an origin web server

The BIG-IP® device can request gzip-encoded or deflate-encoded content from the origin web server to accelerate responses. When the **Enable Assembly Compression OWS** check box is selected (enabled), the BIG-IP® device sends an `Accept-Encoding: gzip, deflate` header to the origin web server. The origin web server complies only if it supports the compression mode; otherwise, the origin web server provides uncompressed content.

This functionality occurs independently of selecting (enabling) the **Enable Content Compression** check box, which sets the compression for the response that the BIG-IP device sends back to the client.

Enabling content compression from an origin web server

The BIG-IP system can request compressed content from the origin web server to accelerate responses.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Click a node in the Policy Tree.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. Select the **Enable Assembly Compression OWS** check box.
8. Click **Save**.
9. Publish the acceleration policy.
 - a) Click **Publish**.
 - b) In the **Comment** field, type a description.
 - c) Click **Publish Now**.

If configured for compression, the origin web server provides compressed content in the response.

Accelerating Responses with Metadata Cache Responses

Overview: Using Metadata cache responses to accelerate responses

Responses from origin web servers include *entity tags* (ETags), which are arbitrary strings attached to a document that specify some characteristic of the document, such as a version, serial number, or checksum of content. A changed document includes a different ETag, enabling a client's GET request to use an `If-None-Match` conditional header to acquire a new copy of the document. Because not all web applications generate ETags consistently, the BIG-IP device creates its own ETag for each cached document that is based on a signature, or checksum, of the document's content. The BIG-IP device stores content signatures in the Metadata cache for other optimizations, including Intelligent Browser Referencing.

BIG-IP applications provide options to always or never send metadata. All BIG-IP applications share the same Metadata cache.

BIG-IP policies cache ETag headers, which include the following:

- Request URL
- Content signature of the response body
- Application name for the matching request
- Metadata, including the expiration time, read time, and update time for content

Accelerating Metadata responses

The BIG-IP system creates its own Metadata response for each cached document that is based on a signature, or checksum, of the document's content. The BIG-IP system stores content signatures in the Metadata cache for other optimizations, including Intelligent Browser Referencing.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. From the **Send Metadata** list, select **Always**.
5. In the **Metadata Cache Max Size** field, type a size in megabytes (MB) for the maximum cache size.
6. Click **Save**.

The BIG-IP system stores content signatures in the Metadata cache.

Disabling Metadata responses

If necessary, you can disable the BIG-IP system from sending Metadata responses.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. From the **Send Metadata** list, select **Never**.
5. Click **Save**.

The BIG-IP system no longer sends Metadata responses.

Accelerating Traffic with a Local Traffic Policy

About classifying types of HTTP traffic with a local traffic policy

An application that runs on a virtual server accelerates all HTTP traffic. You can, however, use a local traffic policy to classify types of HTTP traffic for the BIG-IP® system to accelerate, by specifying hosts, paths, headers, and cookies.

Important: Although you can use a local traffic policy to classify the types of HTTP traffic to accelerate, the local traffic policy overrides the **Web Acceleration** profile on the virtual server. Acceleration of HTTP traffic with the BIG-IP system should primarily be configured through a **Web Acceleration** profile, instead of a local traffic policy.

Accelerating traffic for applications with a local traffic policy

Ensure that the configuration includes a **Web Acceleration** profile configured with an enabled BIG-IP® acceleration application.

A local traffic policy uses the HTTP header, cookie, host, and path to classify and accelerate traffic for applications that are running on a virtual server. You can assign multiple local traffic policies to a virtual server, as needed.

Important: If you configure a local traffic policy to accelerate traffic, the policy overrides settings configured on the virtual server for an acceleration application in the **Web Acceleration** profile.

1. On the Main tab, click **Local Traffic > Policies > Policy List**.
The Policy List Page screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a unique name for the policy.
4. From the **Strategy** list, select a matching strategy.
5. For the **Requires** setting, select **http** from the **Available** list, and move the entry to the **Selected** list using the Move button.
6. For the **Controls** setting, select **acceleration** from the **Available** list, and move the entry to the **Selected** list using the Move button.
7. Click **Add**.
The New Rule screen opens.
8. In the **Rule** field, type a unique name for the rule.
9. From the **Operand** list, select the **http-host** operand for the rule, configure the applicable settings, and click **Add**.
10. From the **Operand** list, select the **http-uri** operand for the rule, configure the applicable settings, and click **Add**.
11. From the **Operand** list, select the **http-header** operand for the rule, configure the applicable settings, and click **Add**.
12. From the **Operand** list, select the **http-cookie** operand, configure the applicable settings, and click **Add**.
13. Using the **Actions** setting, configure the applicable options.

- a) From the **Target** list, select a target.
- b) From the **Event** list, select an event.
- c) From the **Action** list, select an action.
- d) From the **Parameters** list, select a type of parameter to apply.
- e) In the **Parameters** field, type the text that applies to the type of parameter and click **Add**.

The configured parameter appears in the **Parameters** list box.

- f) At the lower left, click **Add**.

The configured settings for the action appear in the **Actions** list.

14. Click **Finished**.

A traffic policy classifies and accelerates traffic for applications that are running on a virtual server.

Accelerating Traffic with Intelligent Client Cache

Overview: Accelerating traffic with Intelligent Client Cache

Intelligent Client Cache (ICC) is a web acceleration technique for mobile and desktop browsers that support HTML5. ICC uses HTML5 local storage to build a cache of documents and resources. Client-side javascript code tracks the resources cached and interacts with the serverside code to ensure that only changed resources are downloaded on subsequent requests.

Accelerating traffic for HTML5-compliant browsers

Intelligent Client Cache (ICC) is a web acceleration technique for mobile and desktop clients who support HTML5.

1. On the Main tab, click **Acceleration > Web Application > Policies**.
The Policies screen displays a list of existing acceleration policies.
2. Click the name of a user-defined acceleration policy.
3. Expand the Policy Tree to a node that matches html content, and click the node.
4. From the Matching Rules menu, choose Acceleration Rules.
5. On the menu bar, click **Assembly**.
The screen refreshes to show assembly options.
6. From the **Content Assembly Options** list, select **Advanced**.
7. In the Intelligent Client Cache Settings, select the check box to enable Intelligent Client Cache.
8. Configure the Intelligent Client Cache settings.
 - a) Set the max size of images to ICC.
The max size can be limited by the specific browser used. The default setting is 32KB.
 - b) Set the max size of CSS to ICC.
The default setting is 50KB.
 - c) Set the max size of JS to ICC.
The default setting is 50KB.
 - d) Set the number of links to ICC.
The maximum number of links that can be cached with ICC is 100. The default is 10.
 - e) Set the client expiry of resource.
This specifies the minimum expiry of linked content that will be inlined for ICC. The object must have a greater expiry than this minimum in order to be inlined. The default setting is 2 days.
9. Click **Save**.

You have now configured Intelligent Client Cache (ICC).

Using the Request Logging Profile

Overview: Configuring a Request Logging profile

The Request Logging profile gives you the ability to configure data within a log file for HTTP requests and responses, in accordance with specified parameters.

Task summary

Perform these tasks to log HTTP request and response data.

Creating a pool with request logging to manage HTTP traffic

Creating a request logging profile

Configuring a virtual server for request logging

Deleting a request logging profile

Creating a pool with request logging to manage HTTP traffic

For a basic configuration, you need to create a pool to manage HTTP connections.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, from the **Available** list, select the **http** monitor and move the monitor to the **Active** list.
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.
The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
 - Select **Disabled** to disable priority groups. This is the default option.
 - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Add the IP address for each logging server that you want to include in the pool, using the **New Members** setting:
 - a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
 - b) Type the port number for the logging server in the **Service Port** field.
 - c) (Optional) Type a priority number in the **Priority** field.
 - d) Click **Add**.
8. Click **Finished**.

The new pool appears in the Pools list.

Creating a request logging profile

You must have already created a pool that includes logging servers as pool members before you can create a request logging profile.

With a request logging profile, you can log specified data for HTTP requests and responses, and then use that information for analysis and troubleshooting.

1. On the Main tab, click **Local Traffic > Profiles > Other > Request Logging**.
The Request Logging profile list screen opens.
2. Click **Create**.
The New Request Logging Profile screen opens.
3. From the **Parent Profile** list, select a profile from which the new profile inherits properties.
4. Select the **Custom** check box for the Request Settings area.
5. Configure the request settings, as necessary.
6. Select the **Custom** check box for the Response Settings area.
7. Configure the response settings, as necessary.
8. Click **Finished**.

This makes a request logging profile available to log specified data for HTTP requests and responses.

You must configure a virtual server for request logging.

Configuring a request logging profile for requests

Ensure that the configuration includes a pool that includes logging servers as pool members.

You can use a request logging profile to log specified data for HTTP requests, and then use that information for analysis and troubleshooting.

1. On the Main tab, click **Local Traffic > Profiles > Other > Request Logging**.
The Request Logging profile list screen opens.
2. Click **Create**.
The New Request Logging Profile screen opens.
3. From the **Parent Profile** list, select a profile from which the new profile inherits properties.
4. Select the **Custom** check box for the Request Settings area.
5. From the **Request Logging** list, select **Enabled**.
6. In the **Template** field, type the request logging parameters for the entries that you want to include in the log file.
7. From the **HSL Protocol** list, select a high-speed logging protocol.
8. From the **Pool Name** list, select the pool that includes the log server as a pool member.
9. (Optional) You can also configure the error response settings.
 - a) From the **Respond On Error** list, select **Enabled**.
 - b) In the **Error Response** field, type the error response strings that you want to include in the log file.
These strings must be well-formed for the protocol serving the strings.
 - c) Select the **Close On Error** check box to drop the request and close the connection if logging fails.
10. (Optional) You can also configure the logging request errors settings.
 - a) From the **Log Logging Errors** list, select **Enabled**.
 - b) In the **Error Template** field, type the request logging parameters for the entries that you want to include in the log file.
 - c) From the **HSL Error Protocol** list, select a high-speed logging error protocol.
 - d) From the **Error Pool Name** list, select a pool that includes the node for the error logging server as a pool member.
11. Click **Update**.

This configures a request logging profile to log specified data for HTTP requests.

Configuring a request logging profile for responses

You must have already created a pool that includes logging servers as pool members before you can configure a request logging profile for responses.

With a request logging profile, you can log specified data for HTTP requests and responses, and then use that information for analysis and troubleshooting.

1. On the Main tab, click **Local Traffic > Profiles > Other > Request Logging**.
The Request Logging profile list screen opens.
2. Click **Create**.
The New Request Logging Profile screen opens.
3. From the **Parent Profile** list, select a profile from which the new profile inherits properties.
4. Select the **Custom** check box for the Response Settings area.
5. In the Response Settings area, from the **Response Logging** list, select **Enabled**.
6. (Optional) Select the **Log By Default** check box.
The **Log By Default** check box is selected by default.
7. In the **Template** field, type the response logging parameters for the entries that you want to include in the log file.
8. From the **HSL Protocol** list, select a high-speed logging protocol.
9. From the **Pool Name** list, select the pool that includes the node log server as a pool member.
10. (Optional) Configure the logging request error settings.
 - a) From the **Log Logging Errors** list, select **Enabled**.
 - b) In the **Error Template** field, type the response logging parameters for the entries that you want to include in the log file.
 - c) From the **HSL Error Protocol** list, select a high-speed logging error protocol.
 - d) From the **Error Pool Name** list, select a pool that includes the node for the error log server as a pool member.
11. Click **Update** to save the changes.

This configures a request logging profile to log specified data for HTTP responses.

Configuring a virtual server for request logging

You can configure a virtual server to pass traffic to logging servers.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Resources**.
4. From the **Default Pool** list, select a pool name that is configured with pool members for request logging.
5. Click the **Properties** tab.
6. From the **Configuration** list, select **Advanced**.
7. From the **Request Logging Profile** list, select the profile you want to assign to the virtual server.
8. Click **Update**.

This virtual server can now pass traffic to the configured logging servers.

Deleting a request logging profile

You can delete a user-defined request logging profile that is obsolete or no longer needed.

1. On the Main tab, click **Local Traffic > Profiles > Other > Request Logging**.
The Request Logging profile list screen opens.
2. Select the check box for the applicable profile.
3. Click **Delete**.
4. Click **Delete**.

The profile is deleted.

Request Logging profile settings

With the Request Logging profile, you can specify the data and the format for HTTP requests and responses that you want to include in a log file.

General Properties

Setting	Value	Description
Name	No default	Specifies the name of the profile.
Parent Profile	Selected predefined or user-defined profile	Specifies the selected predefined or user-defined profile.

Request Settings

Setting	Value	Description
Request Logging	Disabled	Enables logging for requests.
Template		Specifies the directives and entries to be logged.
HSL Protocol	UDP	Specifies the protocol to be used for high-speed logging of requests.
Pool Name	None	Defines the pool associated with the virtual server that is logged.
Respond On Error	Disabled	Enables the ability to respond when an error occurs.
Error Response	None	<p>Specifies the response text to be used when an error occurs.</p> <p>For example, the following response text provides content for a 503 error.</p> <pre><html> <head> <title>ERROR</title> </head> <body> <p>503 ERROR-Service Unavailable</p> </body> </html></pre>
Close On Error	Disabled	When enabled, and logging fails, drops the request and closes the connection.
Log Logging Errors	Disabled	Enables the ability to log any errors when logging requests.
Error Template	None	Defines the format for requests in an error log.
HSL Error Protocol	UDP	Defines the protocol to be used for high-speed logging of request errors.
Error Pool Name	None	Specifies the name of the error logging pool for requests.

Response Settings

Setting	Value	Description
Response Logging	Disabled	Enables logging for responses.
Log By Default	Enabled	Defines whether to log the specified settings for responses by default.
Template	None	Specifies the directives and entries to be logged.
HSL Protocol	UDP	Specifies the protocol to be used for high-speed logging of responses.
Pool Name	None	Defines the pool name associated with the virtual server that is logged.
Log Logging Errors	Disabled	Enables the ability to log any errors when logging responses.
Error Template	None	Defines the format for responses in an error log.
HSL Error Protocol	UDP	Defines the protocol to be used for high-speed logging of response errors.
Error Pool Name	None	Specifies the name of the error logging pool for responses.

Request Logging parameters

This table lists all available parameters from which you can create a custom HTTP Request Logging profile. These are used to specify entries for the **Template** and **Error Template** settings. For each parameter, the system writes to the log the information described in the right column.

Table 1: Request logging parameters

Parameter	Communication method	Log file entry description
BIGIP_BLADE_ID	Request and Response	An entry for the slot number of the blade that handled the request.
BIGIP_CACHED	Response	An entry of <code>Cached status: true</code> , if the response came from BIG-IP® cache, or <code>Cached status: false</code> , if the response came from the server.
BIGIP_HOSTNAME	Request and Response	An entry for the configured host name of the unit or chassis.
CLIENT_IP	Request and Response	An entry for the IP address of a client, for example, 192.168.74.164.
CLIENT_PORT	Request and Response	An entry for the port of a client, for example, 80.
DATE_D	Request and Response	A two-character entry for the day of the month, ranging from 1 (note the leading space) through 31.
DATE_DAY	Request and Response	An entry that spells out the name of the day.
DATE_DD	Request and Response	A two-digit entry for the day of the month, ranging from 01 through 31.
DATE_DY	Request and Response	A three-letter entry for the day, for example, Mon.
DATE_HTTP	Request and Response	A date and time entry in an HTTP format, for example, Tue, 5 Apr 2011 02:15:31 GMT.

Parameter	Communication method	Log file entry description
DATE_MM	Request and Response	A two-digit month entry, ranging from 01 through 12.
DATE_MON	Request and Response	A three-letter abbreviation for a month entry, for example, APR.
DATE_MONTH	Request and Response	An entry that spells out the name of the month.
DATE_NCSA	Request and Response	A date and time entry in an NCSA format, for example, dd/mm/yy:hh:mm:ss ZNE.
DATE_YY	Request and Response	A two-digit year entry, ranging from 00 through 99.
DATE_YYYY	Request and Response	A four-digit year entry.
HTTP_CLASS	Neither Request nor Response: Deprecated	The name of the <code>httpclass</code> profile that matched the request, or an empty entry if a profile name is not associated with the request.
HTTP_KEEPALIVE	Request and Response	A flag summarizing the HTTP1.1 keep-alive status for the request:: aY if the HTTP1.1 keep-alive header was sent, or an empty entry if not.
HTTP_METHOD	Request and Response	An entry that defines the HTTP method, for example, GET, PUT, HEAD, POST, DELETE, TRACE, or CONNECT.
HTTP_PATH	Request and Response	An entry that defines the HTTP path.
HTTP_QUERY	Request and Response	The text following the first ? in the URI.
HTTP_REQUEST	Request and Response	The complete text of the request, for example, \$METHOD \$URI \$VERSION.
HTTP_STATCODE	Response	The numerical response status code, that is, the status response code excluding subsequent text.
HTTP_STATUS	Response	The complete status response, that is, the number appended with any subsequent text.
HTTP_URI	Request and Response	An entry for the URI of the request.
HTTP_VERSION	Request and Response	An entry that defines the HTTP version.
NCSA_COMBINED	Response	An NCSA Combined formatted log string, for example, \$NCSA_COMMON \$Referer \${User-agent} \$Cookie.
NCSA_COMMON	Response	An NCSA Common formatted log string, for example, \$CLIENT_IP - - \$DATE_NCSA \$HTTP_REQUEST \$HTTP_STATCODE \$RESPONSE_SIZE.
RESPONSE_MSECS	Response	The elapsed time in milliseconds (ms) between receiving the request and sending the response.
RESPONSE_SIZE	Response	An entry for the size of response in bytes.
RESPONSE_USECS	Response	The elapsed time in microseconds (μs) between receiving the request and sending the response.

Parameter	Communication method	Log file entry description
SERVER_IP	Response	The IP address of the pool member to which the HTTP request was sent, for example, 10.10.0.1.
SERVER_PORT	Response	The port of the pool member to which the HTTP request was sent, for example, 80.
SNAT_IP	Response	An entry for the self IP address of the BIG-IP-originated connection to the server when SNAT is enabled, or an entry for the client IP address when SNAT is not enabled.
SNAT_PORT	Response	An entry for the port of the BIG-IP-originated connection to the server when SNAT is enabled, or an entry for the client port when SNAT is not enabled.
TIME_AMPM	Request and Response	A twelve-hour request-time qualifier, for example, AM or PM.
TIME_H12	Request and Response	A compact twelve-hour time entry for request-time hours, ranging from 1 through 12.
TIME_HRS	Request and Response	A twelve-hour time entry for hours, for example, 12 AM.
TIME_HH12	Request and Response	A twelve hour entry for request-time hours, ranging from 01 through 12.
TIME_HMS	Request and Response	An entry for a compact request time of H:M:S, for example, 12:10:49.
TIME_HH24	Request and Response	A twenty-four hour entry for request-time hours, ranging from 00 through 23.
TIME_MM	Request and Response	A two-digit entry for minutes, ranging from 00 through 59.
TIME_MSECS	Request and Response	An entry for the request-time fraction in milliseconds (ms).
TIME_OFFSET	Request and Response	An entry for the time zone, offset in hours from GMT, for example, -11.
TIME_SS	Request and Response	A two-digit entry for seconds, ranging from 00 through 59.
TIME_UNIX	Request and Response	A UNIX time entry for the number of seconds since the UNIX epoch, for example, 00:00:00 UTC, January 1st, 1970.
TIME_USECS	Request and Response	An entry for the request-time fraction in microseconds (μs).
TIME_ZONE	Request and Response	An entry for the current Olson database or tz database three-character time zone, for example, PDT.
VIRTUAL_IP	Request and Response	An entry for the IP address of a virtual server, for example, 192.168.10.1.

Parameter	Communication method	Log file entry description
VIRTUAL_NAME	Request and Response	An entry for the name of a virtual server.
VIRTUAL_POOL_NAME	Request and Response	An entry for the name of the pool containing the responding server.
VIRTUAL_PORT	Request and Response	An entry for the port of a virtual server, for example, 80.
VIRTUAL_SNATPOOL_NAME	Request and Response	The name of the Secure Network Address Translation pool associated with the virtual server.
WAM_APPLICATION_NAME	Response	An entry that defines the name of the BIG-IP® acceleration application that processed the request.
WAM_X_WA_INFO	Response	An entry that specifies a diagnostic string (X-WA-Info header) used by BIG-IP acceleration to process the request.
NULL		Undelimited strings return the value of the respective header.

Monitoring BIG-IP Acceleration Application Performance

Overview: Monitoring the performance of a BIG-IP acceleration application

The BIG-IP's performance reports provide information about page requests, the frequency of those requests, and how well the BIG-IP system serviced those requests from cache. Additionally, performance reports provide information about the acceleration application, policy, policy node, HTTP response status, S-code, size range of the response, response object type, and ID of the BIG-IP system or browser making the request.

The BIG-IP system provides three types of performance reports.

- **Traffic Reports.** These reports display the number of requests (hits) received, and responses served, by the BIG-IP system.
- **Byte Reports.** These reports display the bytes of content that the BIG-IP system has sent in response to requests.
- **Response Reports.** These reports display the average amount of time it takes the BIG-IP system to respond to a request from the client.

You can use these performance reports to evaluate your acceleration policies, adjusting them as required to maximize client access to your applications. The individual performance reports display content according to the persistent parameters that you select for the filter. You can also save performance reports to a specified file type so that you can import them into specific applications.

***Note:** Enabling performance monitoring for a BIG-IP acceleration application can degrade overall performance and should only be used temporarily.*

Enabling performance monitoring for a BIG-IP application

You can enable performance monitoring for a BIG-IP application, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. From the **Performance Monitor** list, select **Enable**.
5. In the **Data Retention Period** field, type the number of days to retain data.
6. Click **Save**.

Performance monitoring is enabled for the application.

Disabling performance monitoring for a BIG-IP application

You can disable performance monitoring for a BIG-IP application, as necessary.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. From the **Performance Monitor** list, select **Disable**.
5. Click **Save**.

Performance monitoring is disabled for the application.

Overview: Reporting ROI statistics

ROI statistics provide you with acceleration data for an application that you can use to determine the benefits of using acceleration, or use to refine acceleration performance. You can email the resultant charts and statistics, either weekly or monthly, as necessary.

Task summary

Creating an SMTP server configuration

Creating a pool

Overview: Configuring ROI reports

Modifying an application to export ROI statistics

Creating an SMTP server configuration

You specify the SMTP server configuration so that you can send emails through an SMTP server.

1. On the Main tab, click **System > Configuration > Device > SMTP**.
2. Click the **Create** button.
The New SMTP Configuration screen opens.
3. In the **Name** field, type a name for the SMTP server that you are creating.
4. In the **SMTP Server Host Name** field, type the fully qualified domain name for the SMTP server host.
5. In the **SMTP Server Port Number** field, type a port number.
For no encryption or TLS encryption, the default is 25. For SSL encryption, the default is 465.
6. In the **Local Host Name** field, type the host name used in the SMTP headers in the form of a fully qualified domain name.
This host name is not the same as the BIG-IP® system's host name.
7. In the **From Address** field, type the email address that you want displayed as the reply-to address for the email.
8. From the **Encrypted Connection** list, select the encryption level required for the SMTP server.
9. To require that the SMTP server validates users before allowing them to send email, select the **Use Authentication** check box, and type the user name and password required to validate the user.
10. Click the **Finish** button.

You can now configure the system to use this SMTP server to send emails. For the SMTP mailer to work, you must make sure the SMTP server is on the DNS lookup server list, and configure the DNS server on the BIG-IP® system.

Creating a pool

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

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:

- a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
- b) In the **Address** field, type an IP address.
- c) In the **Service Port** field, type a port number, or select a service name from the list.
- d) (Optional) In the **Priority** field, type a priority number.
- e) Click **Add**.

5. Click **Finished**.

6. Repeat these steps for each pool you want to create.

The new pool appears in the Pools list.

Overview: Configuring ROI reports

You can evaluate the benefits and performance improvements of acceleration functionality for an application by examining ROI Statistics. Acceleration ROI Statistics provide data on compression, caching, minification, inlining, image optimization, and Intelligent Browser Referencing (IBR), which you can assess to determine the current acceleration performance, and, based on those results, refine the acceleration performance.

Task summary

Overview: Reporting ROI statistics

Configuring a Caching Bytes Saved report

Configuring a Compression Bytes Saved report

Configuring an Image Optimization Bytes Saved report

Configuring a Minification Bytes Saved report

Configuring a Caching Requests Saved report

Configuring Client IBR'd Links data

Configuring Client IBR'd Links Received data

Configuring Inlined Links data

Configuring ICC Inlined Links data

Configuring ICC Referenced Links data

Configuring a Caching Bytes Saved report

You can configure the Cache Bytes Saved report to provide statistics about the number of kilobytes served from cache, describing the reduction in bandwidth between the origin web server and the BIG-IP[®] system, and optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Byte Savings > Caching Bytes Saved**.
The Caching Bytes Saved screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for caching bytes saved.
 - Time: Lists the day for the statistical results.
 - Caching Bytes Saved (KB): Lists the number of kilobytes cached, and saved from being sent to origin web servers, for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.

- **CSV.** Exports data for the configured ROI report as a column separated value file.
- **Excel.** Exports data for the configured ROI report as a Microsoft® Excel® file.
- **XML.** Exports data for the configured ROI report as an XML file.

The ROI Report for Caching Bytes Saved data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring a Compression Bytes Saved report

When BIG-IP® system, instead of the origin web server, performs compression, you can configure the Compression Bytes Saved report to provide statistics about the reduction in kilobytes sent to the client, and optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Byte Savings > Compression Bytes Saved.**
The Compression Bytes Saved screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Compression Bytes Saved.
 - **Time:** Lists the day for the statistical results.
 - **Compressed Bytes (KB):** Lists the compressed size of objects in kilobytes, reducing the sizes of objects sent to origin web servers, for the respective day.
 - **Pre-Compressed Bytes (KB):** Lists the precompressed size of objects in kilobytes for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV.** Exports data for the configured ROI report as a column separated value file.
 - **Excel.** Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML.** Exports data for the configured ROI report as an XML file.

The ROI Report for Compression Bytes Saved data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring an Image Optimization Bytes Saved report

You can configure the Image Optimization Bytes Saved settings to provide statistics about optimized images in kilobytes, which reduces the sizes of images sent to the client, and you can optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Byte Savings > Image Optimization Bytes Saved.**
The Image Optimization Bytes Saved screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Image Optimization Bytes Saved.

- **Time:** Lists the day for the statistical results.
 - **Optimized Bytes (KB):** Lists the size of optimized images in kilobytes, reducing the sizes of images sent to origin web servers, for the respective day.
 - **Pre-Image Optimization Bytes (KB):** Lists the preoptimized size of images in kilobytes for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV.** Exports data for the configured ROI report as a column separated value file.
 - **Excel.** Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML.** Exports data for the configured ROI report as an XML file.

The ROI Report for Image Optimization Bytes Saved data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring a Minification Bytes Saved report

You can configure the Minification Bytes Saved settings to provide statistics about minified JavaScript and cascading style sheet (CSS) objects in kilobytes, reducing the sizes of objects sent to the client, and to optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Byte Savings > Minification Bytes Saved**.
The Minification Bytes Saved screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics. Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Minification Bytes Saved.
 - **Time:** Lists the day for the statistical results.
 - **Minification Bytes (KB):** Lists the size of minification objects in kilobytes, reducing the sizes of objects sent to origin web servers, for the respective day.
 - **Pre-JavaScript/CSS Minification Bytes (KB):** Lists the pre-minification size of JavaScript and CSS objects in kilobytes for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV.** Exports data for the configured ROI report as a column separated value file.
 - **Excel.** Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML.** Exports data for the configured ROI report as an XML file.

The ROI Report for Minification Bytes Saved data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring a Caching Requests Saved report

You can configure the Caching Requests Saved settings to provide statistics about the number of requests that are served from cache, which reduces the load for the origin web server, and you can optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Caching Requests Saved**.
The Caching Requests Saved screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.

Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.

3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Caching Requests Saved.
 - **Time**: Lists the day for the statistical results.
 - **Caching Requests Saved**: Lists the number of requests not made to an origin web server when using caching for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for Caching Requests Saved data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring Client IBR'd Links data

You can configure the Client IBR'd Links settings to provide statistics about the number of Intelligent Browser Referencing (IBR) links that have been sent to a client, and to optionally export the resultant data. You will want to use the statistics for this setting in combination with the Client IBR'd Links Received statistics, to determine the number of links that eliminated unnecessary conditional GET requests by increasing the content expire time.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > IBR Savings > Client IBR'd Links**.
The Client IBR'd Links screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Client IBR'd Links.
 - **Time**: Lists the day for the statistical results.
 - **Client IBR'd Links**: Lists the number of Intelligent Browser Referencing (IBR) links that have been sent to a client for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for Client IBR'd Links data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring Client IBR'd Links Received data

You can configure the Client IBR'd Links Received settings to provide statistics about the number of Intelligent Browser Referencing (IBR) links that have been received from a client, and optionally to export the resultant data. You will want to use the statistics for this setting in combination with the Client IBR'd Links statistics, to determine the number of links that eliminated unnecessary conditional GET requests by increasing the content expire time.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > IBR Savings > Client IBR'd Links Received**.
The Client IBR'd Links Received screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Client IBR'd Links Received.
 - **Time**: Lists the day for the statistical results.
 - **Client IBR'd Links Received**: Lists the number of Intelligent Browser Referencing (IBR) links that have been received from a client for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for Client IBR'd Links Received data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring Inlined Links data

You can configure the Inlined Links settings to provide statistics about the number of links that become inlined when sending a response to a client, reducing the number of GET requests from the client, and you can optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > Inlined Links**.
The Inlined Links screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for Inlined Links.
 - **Time**: Lists the day for the statistical results.
 - **Inlined Links**: Lists the number of links that become inlined when sending a response to a client for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for Inlined Links data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring ICC Inlined Links data

You can configure the ICC Inlined Links settings to provide statistics about the number of links that become inlined for a client by Intelligent Client Cache (ICC) functionality, reducing the number of GET requests from the client, and you can optionally export the resultant data.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > ICC Savings > ICC Inlined Links**.
The ICC Inlined Links screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for ICC Inlined Links.
 - **Time**: Lists the day for the statistical results.
 - **ICC Inlined Links**: Lists the number of links that become inlined for a client by Intelligent Client Cache (ICC) functionality for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for ICC Inlined Links data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Configuring ICC Referenced Links data

You can configure the ICC Ref'd Links settings to provide statistics about the number of links referencing content that might exist in the client's local storage, added by Intelligent Client Cache (ICC) functionality, and to optionally export the resultant data. Content residing in a client's local storage reduces the number of GET requests from the client, and reduces the amount of bytes sent to the client, by instead referencing the content in the client's local storage.

1. On the Main tab, click **Acceleration > Web Application > ROI Reports > ICC Savings > ICC Ref'd Links**.
The ICC Ref'd Links screen opens.
2. From the **Byte Rate: Daily Time Period** list, select a range for the period displaying daily statistics.
Settings for a daily time period range from **Last Hour** through **Last 4 Weeks**, providing statistical results in daily increments.
3. To refresh statistical results for a selected daily time period, click the **Refresh** button, or select a value for the **Auto Refresh** setting.
4. From the **Device** list, select a device group.
5. Observe the statistical results for ICC Reference Links.
 - **Time**: Lists the day for the statistical results.
 - **ICC Referenced Links**: Lists the number of links referencing content that exists in the client's local storage, added by Intelligent Client Cache (ICC) functionality, for the respective day.
6. (Optional) From the **Export options** controls, click an option to specify how to export data.
 - **CSV**. Exports data for the configured ROI report as a column separated value file.
 - **Excel**. Exports data for the configured ROI report as a Microsoft® Excel® file.
 - **XML**. Exports data for the configured ROI report as an XML file.

The ROI Report for ICC Referenced Links data is available for evaluation, advanced configuration within an application, and optional exporting to a specified file format.

Modifying an application to export ROI statistics

Exporting ROI statistics for an application requires the following items.

- Availability of an SMTP server.
- Preconfiguration of ROI reports.

You can modify an application to use advanced ROI statistics, which include configured ROI reports, and send the results by email for evaluation.

1. On the Main tab, click **Acceleration > Web Application > Applications**.
The Applications List screen opens.
2. Click the name of an application.
3. From the **General Options** list, select **Advanced**.
4. From the **Collect ROI Statistics** list, select **Enabled**.
5. In the **Schedule Title** field, type a name for the ROI statistics report.
6. From the **Schedule Frequency** list, select a report frequency.
 - **Monthly**. Provides an ROI statistics report monthly.
 - **Weekly**. Provides an ROI statistics report weekly.
7. Verify the time and date specified by the **Next Send Time** text.

***Tip:** To send an ROI statistics report immediately, click **Now**, and then click **Save**.*

8. From the **Smtp Server** list, select an SMTP server.

***Note:** You need to configure an SMTP server in order to send ROI reports. If you only need to provide data available through the Configuration utility, then an SMTP server is not required.*

9. In the **Send To** field, type an email address to which you want to send the ROI Statistics, and click **Add**.

***Note:** You can add multiple email addresses, as necessary.*

10. For the **Charts To Send** setting, from the **Available** list, select the charts that you want and transfer them to the **Selected** list using the Move button.
11. Click **Save**.

An application is available that sends advanced ROI statistics, with configured ROI reports, by email for evaluation.

Using Forward Error Correction to Mitigate Packet Loss

Overview: Using forward error correction (FEC) to mitigate packet loss

The BIG-IP® system performs forward error correction (FEC) by adding redundancy to the transmitted information. FEC provides a loss correction facility for all IP-based protocols optimized by Application Acceleration Manager™. All iSession™ traffic can benefit from FEC loss mitigation, which is preferred over aggressive TCP retransmission in shared network environments.

To implement forward error correction, the BIG-IP system aggregates packets for a specified amount of time, divides the load into the specified number of equal packets (source packets), and adds the specified number of redundant (repair) packets. With adaptive FEC, the system adjusts these numbers as it measures the link error rate.

If you are configuring FEC on a central BIG-IP device for a server that does not initiate traffic, you can configure a FEC tunnel with an undefined remote address. You then configure a separate FEC tunnel from each remote BIG-IP device that handles client-initiated traffic to the central BIG-IP device. You can also configure a FEC tunnel between the local BIG-IP device and any other BIG-IP device that has a FEC tunnel with an undefined remote address.




Figure 2: FEC configuration between BIG-IP devices

In addition to configuring FEC between two BIG-IP systems, you can configure FEC between an edge client and a BIG-IP system that has Access Policy Manager® licensed. Consult the Access Policy Manager (APM®) documentation for information about configuring the client access deployment.

Note: Before you can configure forward error correction (FEC), you must have licensed and provisioned Application Acceleration Manager (AAM™).

About forward error correction (FEC)

Forward error correction (FEC) is an acceleration technique for all kinds of traffic, including TCP and UDP traffic on lossy networks. FEC controls data transmission errors over unreliable or noisy communication channels. With FEC, the sender encodes messages with an extra error-correcting code (ECC). The redundancy allows the receiver to detect a limited number of errors that might occur anywhere in the message, and often to correct these errors without retransmission.

Packet loss occurs when one or more packets traveling across a network fail to reach their destination. Packet loss can be caused by a number of factors that inevitably result in highly noticeable performance issues, particularly with realtime protocols, streaming technologies, voice-over-IP, online gaming, and video conferencing. Some network transport protocols, such as TCP, provide for reliable delivery of packets. In the event of packet loss, the receiver might ask for retransmission, or the sender automatically

resends any segments that have not been acknowledged. Although TCP can recover from packet loss, retransmitting missing packets causes the overall throughput of the connection to decrease. Error correction occurs without the need for a reverse channel to request retransmission of data, but at the cost of a fixed, higher forward channel bandwidth. Therefore, FEC is most useful in situations where retransmissions are costly or impossible.

Task summary

The BIG-IP® system handles forward error correction according to the parameters in the FEC profile you select when you create a FEC tunnel. If the system-supplied FEC profile does not meet your network needs, you can customize the profile. For example, if you know that the bulk of the traffic is not compressible, you might want to disable LZO compression. The system-supplied FEC profile has both adaptive settings enabled, which means that it adjusts the number of source and repair packets according to network traffic conditions. This feature is particularly useful for unstable conditions. If your network conditions are stable, you might want to adjust the FEC profile accordingly.

Note: If you are using iSession™ with FEC, disable compression on either the iSession connection or the FEC profile you select for the FEC tunnel.

Task list

Customizing a FEC profile

Creating a FEC tunnel for receiving traffic

Creating a FEC tunnel for initiating traffic

Viewing FEC tunnel statistics

Customizing a FEC profile

You can customize the parameters for FEC packet loss mitigation to adjust to your network conditions.

1. On the Main tab, click **Network > Tunnels > Profiles > FEC > Create**.
The New FEC Profile screen opens.
2. In the **Name** field, type a unique name for the profile.
3. From the **Parent Profile** list, select a profile.
A default profile, `fec`, is available.
4. Select the **Custom** check box.
5. Modify the settings, as required.
6. Click **Finished**.

This FEC profile is now available for applying to a FEC tunnel.

To apply this FEC profile to traffic between BIG-IP® systems, you must select it from the **Encapsulation Type** list on the Acceleration Quick Start screen, the Symmetric Optimization Local Endpoint screen, or the New Tunnel screen.

Creating a FEC tunnel for receiving traffic

You can configure a FEC tunnel on a BIG-IP® device to receive requests for a FEC connection from a remote BIG-IP device.

1. On the Main tab, click **Network > Tunnels > Tunnel List > Create or Carrier Grade NAT > Tunnels > Create**.
The New Tunnel screen opens.

2. In the **Name** field, type a unique name for the tunnel.
3. From the **Profile** list, select **fec**.
This setting tells the system which tunnel profile to use. The system-supplied `fec` profile is configured for adaptive behavior for the number of source and repair packets. If you create a new FEC profile with custom settings, the profile then appears in this list, where you can select it.
4. In the **Local Address** field, type the IP address of the local endpoint.
If you are using an iSession connection, use the same IP address you used for the iSession local endpoint. Otherwise, use any self IP address on the BIG-IP system.
5. For the **Remote Address** list, retain the default selection, **Any**.
6. Click **Finished**.

You now have a tunnel that is configured for receiving FEC traffic from any BIG-IP system that has a FEC tunnel configured with the IP address of the local system specified as the **Remote Address**.

If you also want to initiate traffic through a FEC tunnel from the local BIG-IP system, you must create a FEC tunnel with the specific IP address of a remote BIG-IP system that is configured for receiving FEC traffic.

Creating a FEC tunnel for initiating traffic

You can configure a FEC tunnel between BIG-IP® devices to use forward error correction to mitigate data loss during transmission.

1. On the Main tab, click **Network > Tunnels > Tunnel List > Create** or **Carrier Grade NAT > Tunnels > Create**.
The New Tunnel screen opens.
2. In the **Name** field, type a unique name for the tunnel.
3. From the **Profile** list, select **fec**.
This setting tells the system which tunnel profile to use. The system-supplied `fec` profile is configured for adaptive behavior for the number of source and repair packets. If you create a new FEC profile with custom settings, the profile then appears in this list, where you can select it.
4. In the **Local Address** field, type the IP address of the local endpoint.
If you are using an iSession connection, use the same IP address you used for the iSession local endpoint. Otherwise, use any self IP address on the BIG-IP system.
5. From the **Remote Address** list, select **Specify**, and type the IP address of the BIG-IP device at the other end of the tunnel.
6. Click **Finished**.

You now have a tunnel that can transmit FEC traffic to the BIG-IP system specified by the remote IP address, provided the other BIG-IP system has a FEC tunnel that is open to receiving FEC transmissions.

If you also want to receive traffic through a FEC tunnel from the another BIG-IP system, you must create a FEC tunnel with an undefined IP address.

Viewing FEC tunnel statistics

You can view packet-level statistics for FEC tunnels that you have created.

1. Access the `tmsh` command-line utility.
2. At the prompt, type `tmsh show /net tunnels fec-stat all-properties`.
The following listing is an example of results for this command.

Using Forward Error Correction to Mitigate Packet Loss

Net::FEC Tunnel

Name	Profile	Out pkts Raw	Out bits Raw	Out pkts Rdnt	Out bits Rdnt
10.10.10.2	fec_1	51.5K	30.1M	19.3K	28.2M
	In pkts Raw	In bits Raw	In pkts Rdnt	In bits Rdnt	In pkts Rdnt Lost
	97.4K	1.1G	152.4K	1.7G	864
	In pkts Raw Lost	Rmt In Rdnt Pkts	Rmt In Raw Pkts	Rmt In Rdnt Lost	Rmt In Raw Lost
	613	18.2K	48.6K	28	63

Managing Deduplication

Overview: Symmetric data deduplication

Symmetric data deduplication reduces the amount of bandwidth consumed across a WAN link for repeated data transfers. This feature is available only with an Application Acceleration Manager™ (AAM™) license.

With data deduplication, the system performs pattern matching on the transmitted WAN data, rather than caching. If any part of the transmitted data has already been sent, the BIG-IP® system replaces the previously transmitted data with references. As data flows through the pair, each device records the byte patterns and builds a synchronized dictionary. If an identical pattern of bytes traverses the WAN more than once, the BIG-IP closest to the sender replaces the byte pattern with a reference to it, compressing the data. When the reference reaches the other side of the WAN, the remote BIG-IP device replaces the reference with the data, restoring the data to its original format.

Task Summary

Perform these tasks to manage symmetric data deduplication.

Task list

Enabling symmetric data deduplication

Disabling symmetric data deduplication

Clearing the deduplication cache

Enabling symmetric data deduplication

Ensure that you have licensed and provisioned Application Acceleration Manager™ (AAM™) on the BIG-IP® system.

Symmetric data deduplication (SDD) reduces the amount of bandwidth consumed across a WAN link. You can enable symmetric data deduplication on the iSession™ connection between the local endpoint and any remote endpoints. SDD is enabled by default when you provision AAM.

1. On the Main tab, click **Acceleration > Symmetric Optimization > Symmetric Deduplication**.
2. In the **Maximum Number of Remote Endpoints** field, type the number of BIG-IP systems that you expect to connect to this one.

This number specifies the maximum number of remote endpoints that can have symmetric data deduplication enabled, and thus, share the available cache. Any added iSession remote endpoint that exceeds this number receives no cache for deduplication. If you select **SSD v2** in the **Codec** field, the maximum supported is 8. If you select **SSD v3**, the set value is 128.
3. For the **Enable Symmetric Deduplication** setting, select **Yes**.
4. For the **Mode** setting, select the method of storage for symmetric data deduplication.

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

Disk	Specifies that iSession uses the disk, in addition to memory, for storing information used for optimization.
-------------	--

Option	Description
	<i>Note: If you enable data storage on the disk, you must restart the datastor service from the command line using the command sequence <code>bigstart restart datastor</code> for the change to take effect.</i>
Memory	Specifies that iSession uses only memory for storing information used for optimization.
	<i>Note: This setting can provide benefits for higher speed links.</i>

5. For the **Codec** setting, select the SDD version.

Option	Description
SDD v3	Supports a high spoke count, such as for connecting remote sites and for mesh topologies.
SSD v2	Supports a topology with fewer than eight spokes, such as replicating data between data centers.

For SDD to occur between iSession endpoints, you must select the same codec on both the local and remote BIG-IP systems.

6. Click **Update** to save changes.

***Important:** Updating any of these settings causes the deduplication cache to clear.*

Symmetric data deduplication starts after an iSession connection is established with a remote endpoint that also has symmetric data deduplication enabled, provided that the number of remote endpoints does not exceed the value in the **Maximum Number of Remote Endpoints** field.

If you changed the **Codec** setting, the system applies the new setting to any new data flows. However, if you enabled or disabled SDD, you must then restart the BIG-IP from the command line using the command sequence `bigstart restart`.

Disabling symmetric data deduplication

You can disable symmetric data deduplication on the iSession™ connections between the local endpoint and any remote endpoint.

1. On the Main tab, click **Acceleration > Symmetric Optimization > Symmetric Deduplication**.
2. For the **Enable Symmetric Deduplication** setting, select **No**.
3. Click **Update** to save the change.
4. Restart the BIG-IP from the command line by typing `bigstart restart`.

Symmetric data deduplication stops on all iSession connections between the local endpoint and any remote endpoints, and the deduplication cache clears.

Clearing the deduplication cache

Under some circumstances, you might want to clear the deduplication cache. For example, if you are testing the performance of symmetric data deduplication, you might want to clear the cache before you start and reset statistics, to ensure accurate performance data. The cache does not accumulate stale content, as a web cache does. It is more like a dictionary used in compression. To optimize compression ratios, the system manages storage automatically .

1. On the Main tab, click **Acceleration > Symmetric Optimization > Remote Endpoints**.

2. In the Remote Endpoints List screen, select the check box next to the remote endpoint for which you want to clear the cache, and then click **Clear Dedup Cache**.

***Note:** Make sure you select a check box. If you do not select a remote endpoint, the system does not clear any deduplication cache.*

Managing the Settings for Subnet Discovery

Overview: Managing advertised routing

An *advertised route* is a subnet that can be reached through a iSession connection. After the iSession connection is configured between two BIG-IPs, they automatically exchange advertised route specifications between the endpoints. The local endpoint needs to advertise the subnets to which it is connected so that the remote endpoint can determine the destination addresses for which traffic can be optimized. Advertised routes configured on the local endpoint become remote advertised routes on the remote endpoint; that is, the BIG-IP® on the other side of the WAN.

When a BIG-IP device is deployed in a large scale network with large number of servers, and many of them belong to different subnets, manually configuring local optimization subnets can be very time consuming. Subnet Discovery is designed to ease such configuration challenges. With local subnet discovery, instead of requiring manual configuration of local subnets for traffic optimization, the BIG-IP system automatically discovers the local optimization subnet when traffic flows from the client side BIG-IP device to a server-side BIG-IP device.

Task Summary

Perform these tasks to manage advertised routing

Task list

Adding a virtual server to advertised routes

Adding advertised routes manually

Modifying automatic discovery of advertised routes

Verifying subnet discovery

Adding a virtual server to advertised routes

You can add the IP address of a virtual server you created to intercept application traffic to the list of advertised iSession™ routes on the central BIG-IP® system. This configuration tells the BIG-IP system in the remote location that the iSession-terminating endpoint on the central BIG-IP system can route traffic to the application server.

1. On the Main tab, click **Acceleration > Symmetric Optimization > Advertised Routes**.
2. Click **Create**.
The New Advertised Routes screen opens.
3. In the **Name** field, type a name for the advertised route (subnet).
4. In the **Address** field, type the IP address of the virtual server you created for accelerating application traffic.
5. In the **Netmask** field, type the appropriate subnet mask, based on the destination.
6. Click **Finished**.

The remote BIG-IP system now knows that the iSession-terminating endpoint on the central BIG-IP system can route traffic to the application server.

Verify that the iSession profile on the iSession-terminating (endpoint) virtual server is configured to target this virtual server. The default profile `isession`, for which the default **Target Virtual** setting is

match all is appropriate, as long as the **Address** setting for this virtual server is not a wildcard (0.0.0.0).

Adding advertised routes manually

An *advertised route* is a subnet that can be reached through the local endpoint. You can add advertised routes manually, for example, if you disabled the **Discovery** setting on the Quick Start screen.

1. On the Main tab, expand **WAN Optimization** and click **Advertised Routes**.
2. Click **Create**.
The New Advertised Routes screen opens.
3. In the **Name** field, type a name for the subnet.
4. In the **Address** field, type the IP address of the subnet.
5. In the **Netmask** field, type the subnet mask.
6. In the **Label** field, type a descriptive label to identify the subnet.
7. For the **Mode** setting, specify whether traffic on the subnet is included in optimization.
If you select **Excluded**, the local and remote endpoints exchange subnet configuration information, but traffic on this subnet is excluded from optimization.

***Note:** You can define a subset of IP addresses to exclude from optimization within a larger included subnet. An excluded endpoint advertised route must be a valid address range subset of an included endpoint advertised route.*

8. Depending on how many advertised routes you want to add, click the appropriate button.

Option	Description
Repeat	Save this route and add more advertised routes.
Finished	You have finished adding advertised routes.

Modifying automatic discovery of advertised routes

You can modify the settings that pertain to the discovery of subnets that can be reached through the local endpoint. These settings determine how BIG-IP® learns about discovered subnets, and when to display the subnets. Using these settings, you can control the number and reach of the discovered subnets that are included.

1. On the Main tab, click **WAN Optimization > Advertised Routes > Discovery**.
2. From the **Configuration** list, select **Advanced** to view all the settings.
3. Ensure that the **Discover Routes** check box is selected.

***Note:** For server discovery to take place, the setting **Discover Other Endpoints** on the Remote Endpoints Dynamic Discovery screen, at the other end of the connection, must not be set to **Disabled**.*

4. In the **Stop discovery after** field, type the maximum number of servers or subnets (advertised routes) you want the system to discover before it stops looking.
5. In the **Do not add servers with RTT greater than** field, type the maximum round-trip time in milliseconds. The system does not add discovered servers that have an RTT over this value.
6. In the **Minimum prefix length for IPv4 address** field, type the minimum prefix length for route aggregation in IPv4 networks.

If you use the default value of 32/128, the BIG-IP adds the host address as the advertised route. If you change this value to 24, the system adds the /24 network in which the server resides as the advertised route.

7. In the **Minimum prefix length for IPv6 address** field, type the minimum prefix length for route aggregation in IPv6 networks.
8. In the **Allow idle time for routes** field, specify the minimum and maximum lengths of time a discovered route can be idle (no optimized traffic coming through) without being removed.
You can specify these limits in days, hours, or minutes, and the unit of measure must be the same for both limits. This setting does not affect manually configured routes.
9. In the **Do not add routes with ip ttl less than** field, leave the default value of 5, or type a number between 0 and 255.
The BIG-IP system matches the value you set with the IP TTL value of the discovery packets from the server. If the packet has an IP TTL value less than the configured value, it means the server is farther away than you want, so the system does not add the advertised route (server).
10. To save the discovered subnets in the configuration, ensure that the **Automatically save discovered routes** check box is selected.
11. In the **Filter Mode** field, you can exclude from discovery a subset you specify in the **Subnet Filter** field.
You can also narrow the scope of the subnet discovery by selecting **Include** and specifying only the subnets to include in discovery.

Important: If you select **Include** without entering an IP address in the **Subnet Filter** field, the system does not discover any subnets.

12. Click **Update** to save changes.

After the BIG-IP system discovers a subnet and adds the route to the list, the system automatically optimizes traffic to any hosts in that subnet without rediscovery.

Verifying subnet discovery

After sending a client request from the local BIG-IP® to a server behind a remote BIG-IP device, you can perform this procedure to verify that the destination subnet is discovered.

1. Using the browser interface on the client-side BIG-IP system, on the Main tab, click **WAN Optimization > Remote Endpoints**.
The Remote Endpoints List screen opens.
2. Verify that the status indicator is green, and the IP address is correct for the remote endpoint you are checking.
3. On the menu bar, click **Routes**, and verify that the list includes the IP address of the destination subnet.

This subnet is also displayed on the Advertised Routes List screen of the browser interface on the server-side BIG-IP system.

Managing Remote Endpoint Discovery

Overview: Managing dynamic discovery of remote endpoints

Dynamic discovery is a process through which the BIG-IP® identifies and adds remote endpoints automatically. The process occurs when the BIG-IP receives traffic that is matched by a virtual server with an iSession™ profile, but does not recognize the remote destination. When a BIG-IP receives a request destined for a location on the network behind the BIG-IP on the other side of the WAN, the first BIG-IP sends out TCP options or ICMP probes to discover, authenticate, and initiate communication with the new remote endpoint.

***Note:** A TCP request from the client to the server is the action that triggers discovery, not a ping between two endpoints.*

Task Summary

Perform these tasks to manage dynamic discovery of remote endpoints.

Task list

Verifying subnet discovery

Modifying dynamic discovery of remote endpoints

Verifying subnet discovery

After sending traffic between two configured BIG-IP® devices, you can perform this task to verify that BIG-IP has discovered the remote endpoint.

1. On the Main tab, click **Symmetric Optimization > Remote Endpoints**.
2. Verify that the status indicator is green, and the IP address is correct for the remote endpoint you are checking.

Modifying dynamic discovery of remote endpoints

You can modify the dynamic discovery settings, such as specifying the number and types of probe messages, or disabling dynamic discovery.

1. On the Main tab, click **Symmetric Optimization > Remote Endpoints > Discovery**.
2. From the **Dynamic Discovery** list, select **Advanced** to view all the settings.
3. Modify the settings, as required.
4. Click **Update** to save changes.

Setting Up an iSession Connection Using the Quick Start Screen

Overview: Setting up an iSession connection using the Quick Start screen

The Quick Start screen for WAN acceleration provides the settings you need to configure an iSession™ connection on one side of the WAN. To complete the iSession connection, you must use the Quick Start screen on the BIG-IP system on the other side of the WAN.

The Quick Start screen is for the initial BIG-IP symmetric acceleration setup. To change the settings for any iSession acceleration objects after you have completed the initial configuration on the Quick Start screen, use the screen that pertains to that object. For example, to change the settings for the local endpoint, use the Local Endpoint screen.

Setting up an iSession connection using the Quick Start screen

You cannot view the Quick Start screen until you have defined at least one VLAN and at least one self IP on a configured BIG-IP® system that is licensed and provisioned for acceleration.

Use the Quick Start screen to quickly set up the iSession™ endpoints on a BIG-IP system. To optimize WAN traffic, you must configure the iSession endpoints on the BIG-IP systems on both sides of the WAN.

1. Log in to the BIG-IP system that you want to configure.

The default login value for both user name and password is `admin`.

2. On the Main tab, click **Acceleration > Quick Start > Symmetric Properties**.

3. In the **WAN Self IP Address** field, type the local endpoint IP address, if it is not already displayed.

This IP address must be in the same subnet as a self IP address on the BIG-IP system, and to make sure that dynamic discovery properly detects this endpoint, the IP address must be the same as a self IP address on the BIG-IP system.

4. Verify that the **Discovery** setting is set to **Enabled**.

If you disable the **Discovery** setting, or discovery fails, you must manually configure any remote endpoints and advertised routes.

5. Specify the VLANs on which the virtual servers on this system receive incoming traffic.

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

LAN VLANs	Select the VLANs that receive incoming LAN traffic destined for the WAN.
------------------	--

WAN VLANs	Select the VLANs that receive traffic from the WAN through an iSession™ connection.
------------------	---

6. In the Authentication area, for the **Outbound iSession to WAN** setting, select the SSL profile to use for all encrypted outbound iSession connections.

To get WAN optimization up and running, you can use the default selection **serversssl**, but you need to customize this profile for your production environment.

7. For the **Inbound iSession from WAN** setting, select the SSL profile to use on the incoming iSession connection.

To get WAN optimization up and running, you can use the default selection **wom-default-clientssl**.

***Note:** If you configure the iSession connection to not always encrypt the traffic between the endpoints, this profile must be a client SSL profile for which the **Non-SSL Connections** setting is enabled, such as **wom-default-clientssl**.*

8. In the IP Encapsulation area, from the **IP Encapsulation Type** list, select the encapsulation type, if any, for outbound iSession traffic.
 - a) If you select **FEC**, select a FEC profile from the **FEC Profile** list that appears, or retain the default, **default-ipsec-policy-isession**.
 - b) If you select **IPsec**, select an IPsec policy from the **IPSEC Policy** list that appears, or retain the default, **default-ipsec-policy-isession**.
 - c) If you select **IPIP**, the system uses the IP over IP tunneling protocol, and no additional encapsulation setting is necessary.
 - d) If you select **GRE**, select a GRE profile from the **GRE Profile** list that appears, or retain the default, **gre**.
9. Click **Apply**.

To complete the setup, repeat this task on the BIG-IP system on the other side of the WAN.

Forwarding Non-Optimized IP Traffic Through an IPsec Tunnel

Overview: Forwarding Non-Optimized IP traffic through an IPsec tunnel

When you configure an iSession™ connection using the Quick Start screen, you can specify IPsec encapsulation for outbound iSession traffic. If you select IPsec, the BIG-IP® system also encrypts the TCP traffic for the applications you select when you create iApps® templates for optimizing applications.

If you also want to send secured and encrypted non-TCP traffic, you can create a forwarding virtual server that uses the iSession routing to send all IP traffic not matched by other virtual servers through the IPsec tunnel. To accelerate the traffic, you can add IP Payload Compression Protocol (IPComp) to the IPsec tunnel. You would choose IPComp when you expect a great deal of compressible non-TCP traffic.

***Note:** NAT traversal is not supported with iSession routing. For NAT traversal, you must configure a separate IPsec tunnel, and then route the IP traffic through the tunnel.*

Creating a virtual server for all IP iSession traffic

Before you create the virtual server, ensure that you have selected **IPsec** for the **IP Encapsulation Type** setting on the Quick Start screen or the Symmetric Optimization Local Endpoint screen, and chosen an IPsec policy. You can use the pre-defined default policy `default-ipsec-policy-isession`, or create a custom policy, for example, to compress all IP traffic that does not match another virtual server.

If you are using IPsec to encrypt iSession™ traffic, you can create a forwarding virtual server to send all IP traffic through the IPsec tunnel. Creating the virtual server avoids the need for any special routing for non-TCP traffic, such as UDP and ICMP.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
2. Click the **Create** button.
3. Type a unique name for the virtual server, such as `non_tcp_traffic`.
4. For the **Type** setting, select **Forwarding (IP)** from the list.
5. In the **Destination Address** field, type an IP address in CIDR format.

The supported format is address/prefix, where the prefix length is in bits. For example, to select all IP addresses, an IPv4 address/prefix is `0.0.0.0/0`, and an IPv6 address/prefix is `::/0`. To specify a network, an IPv4 address/prefix is `10.07.0.0` or `10.07.0.0/24`, and an IPv6 address/prefix is `ffe1::/64` or `2001:ed8:77b5::/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

***Note:** For best results, F5® recommends that you enter the subnet that matches your destination server network.*

6. In the **Service Port** field, type `*` or select `* All Ports` from the list.
7. In the Configuration area of the screen, from the **Protocol** list, select `*All Protocols`.
8. In the Acceleration area of the screen, from the **iSession Profile** list, select an iSession profile.

***Note:** This setting is available only if you have licensed and provisioned the Application Acceleration Manager™ (AAM™) product.*

9. Click **Finished**.

The completed screen looks similar to the following example.

Common » Local Traffic » Virtual Servers : Virtual Server List » New Virtual Server...

General Properties

Name	non_tcp_issession
Description	forward non-TCP iSession traffic through IPsec tunnel
Type	Forwarding (IP)
Source	
Destination	Type: <input type="radio"/> Host <input checked="" type="radio"/> Network Address: 10.107.0.0 Mask: 255.255.0.0
Service Port	* * All Ports
State	Enabled

Configuration: Basic

Protocol	* All Protocols
VLAN and Tunnel Traffic	All VLANs and Tunnels
Source Address Translation	None

Content Rewrite

Rewrite Profile	None
HTML Profile	None

Acceleration

iSession Profile	isession
Rate Class	None
SRDF Profile	<input type="checkbox"/>
SPDY Profile	None

Figure 3: Example of a completed virtual server screen for non-TCP iSession traffic, with destination subnet specified

Adding compression to an IPsec policy

You can create an IPsec policy that uses iSession™ routing to compress IP traffic through an IPsec tunnel.

1. On the Main tab, click **Network > IPsec > IPsec Policies**.
2. Click the **Create** button.
The New Policy screen opens.
3. In the **Name** field, type a unique name for the policy.
4. For the **IPsec Protocol** setting, retain the default selection, **ESP**.
5. From the **Mode** list, select **iSession Using Tunnel**.
6. For the **Authentication Algorithm** setting, retain the default value, or select the algorithm appropriate for your deployment.
7. For the **Encryption Algorithm** setting, retain the default value, or select the algorithm appropriate for your deployment.
8. For the **Perfect Forward Secrecy** setting, select the option appropriate for your deployment.
9. Only if you want to use IPComp to compress the traffic in the IPsec tunnel, from the **IPComp** list, select **DEFLATE**.
10. For the **Lifetime** setting, retain the default value, **1440**.
This is the length of time (in minutes) before the current security association expires.
11. Click **Finished**.
The screen refreshes and displays the new IPsec policy in the list.

For this IPsec policy to take effect, you must associate it with the iSession routing information, using the IP Encapsulation settings on either the Quick Start screen or the Symmetric Optimization Local Endpoint screen.

Securing an iSession Deployment

Overview: Securing an iSession deployment

For a secure iSession™ deployment, you must use SSL encryption to secure the endpoints of the iSession™ connection. The default SSL profile settings on BIG-IP® acceleration Quick Start screen are sufficient to get symmetric optimization up and running in a demo environment or for testing. F5® recommends that, to secure the endpoints, you specify SSL profiles that use a symmetric optimization-specific root certificate (cert) from a trusted certificate authority (CA).

This illustration shows the network setup for a secure iSession deployment. The example in this implementation uses the specified IP addresses.

- The local endpoint IP address on the BIG-IP SiteA system is 1.1.1.1.
- The local endpoint IP address on the BIG-IP SiteB system is 2.2.2.2.



Figure 4: Network topology for a secure iSession connection

Task summary

The process of securing an iSession™ deployment using SSL includes creating a cert for each iSession endpoint, and then specifying this cert (along with its associated key) in acceleration-related profiles and settings on the system. Before you start this procedure, ensure that you have configured the BIG-IP system on both sides of the WAN. This implementation is based on the default acceleration settings, except where noted.

Task list

Generating and importing SSL certificates for a secure iSession connection

Customizing SSL profiles for a secure iSession connection

Configuring the remote endpoints for a secure iSession connection

Generating and importing SSL certificates for a secure iSession connection

You need to generate and import SSL certificates for a secure iSession™ connection.

1. Generate a root certificate using external Certificate Authority (CA) software, such as the freeware program SimpleCA.
2. Import the generated root certificate into both BIG-IP® systems (for example, BIG-IP SiteA and BIG-IP SiteB).
3. On one of the BIG-IP systems, complete the following steps.
 - a) On the Main tab, click **System** > **File Management** > **SSL Certificate List** > **Import**.
 - b) From the **Import Type** list, select **Certificate**.
 - c) For the **Certificate Name** setting, click **Create New**, and type `wom-root-ca`.

- d) For the **Certificate Source** setting, either click **Upload File** and provide a file name by typing or browsing to the file, or click **Paste Text**, and paste the text copied from another source into the field.
- e) Click **Import**.
- f) Repeat these steps on the other BIG-IP system.
4. Create a certificate and key on one of the BIG-IP systems (for example, BIG-IP SiteA).
 - a) On the Main tab, click **System > File Management > SSL Certificate List**.
 - b) Click the **Create** button.
 - c) In the **Name** field, type `wom-endpoint`.
 - d) From the **Issuer** list, select **Certificate Authority**.
 - e) In the **Common Name** field, type the IP address of the local endpoint for the BIG-IP, for example, `1.1.1.1`.
 - f) Provide any additional information required by your organization.
 - g) Click **Finished**.
5. On the Certificate Signing Request screen, copy or download the certificate signing request for the certificate created in the previous step, and use it to generate a signed certificate using your external CA and the CA certificate that you generated in step 1.
6. Import the generated certificate into the BIG-IP system (for example, BIG-IP SiteA).
 - a) On the Main tab, click **System > File Management > SSL Certificate List**.
 - b) Click `wom-endpoint` (the certificate you created in step 4).
 - c) Select the file `wom-endpoint.crt`.
 - d) Click **Import**.
7. Repeat steps 4-6 on the other BIG-IP system (for example, BIG-IP SiteB), but type `2.2.2.2` in the **Common Name** field on the New SSL Certificate screen.

Customizing SSL profiles for a secure iSession connection

To create custom SSL profiles to use for securing an iSession™ connection, follow these steps.

1. On one of the BIG-IP® systems (for example, BIG-IP SiteA), create a new SSL client profile based on the parent profile `clientssl`.
 - a) On the Main tab, click **Local Traffic > Profiles > SSL > Client**.
 - b) Click the **Create** button.
 - c) In the **Name** field, type `wom-clientssl`.
 - d) From the **Configuration** list, select **Advanced** to display more options.
 - e) For the **Certificate** setting, select the associated Custom check box (to override the default setting), and select **wom-endpoint** from the list.
 - f) For the **Key** setting, select the associated Custom check box, and select **wom-endpoint** from the list.

Local Traffic >> Profiles : SSL : Client >> New Client SSL Profile...

General Properties

Name: my-wom-clientssl

Parent Profile: clientssl

Configuration: Advanced Custom ☐

Certificate: wom-endpoint ☒

Key: wom-endpoint ☒

- g) In the Client Authentication area (near bottom of screen), for the **Client Certificate** setting, select the associated Custom check box, and select **require** from the list.
- h) For the **Frequency** setting, select the associated Custom check box, and select **always** from the list.
- i) For the **Trusted Certificates Authorities** setting, select the associated Custom check box, and select **wom-root-ca** from the list.
- j) For the **Advertised Certificates Authorities** setting, select the associated Custom check box, and select **wom-root-ca** from the list.

Client Authentication Custom ☐

Client Certificate: require ☒

Frequency: always ☒

Retain Certificate: ☒ Enabled

Certificate Chain Traversal Depth: 9

Trusted Certificate Authorities: wom_root_ca_crt ☒

Advertised Certificate Authorities: wom_root_ca_crt ☒

Certificate Revocation List (CRL): None

Cancel Repeat Finished

- k) Click **Finished**.
- 2. Update the configuration on the BIG-IP system (BIG-IP SiteA in our example) to refer to the new client SSL profile.
 - a) On the Main tab, click **Acceleration > Quick Start**.
 - b) From the **Inbound iSession from WAN** list, select wom-clientssl.

Authentication

Outbound iSession to WAN + serverssl

Inbound iSession from WAN + my-wom-clientssl

IP Encapsulation

IP Encapsulation Type None

Apply

- c) Click **Apply**.

Alternatively, you can use the iSession Listener screen settings to create an iSession listener that refers to wom-clientssl.

3. Repeat steps 1-2 on the other BIG-IP system (BIG-IP SiteB in our example).

Configuring the remote endpoints for a secure iSession connection

To configure the remote endpoints using SSL profiles to secure the iSession connection, follow these steps.

1. On the first BIG-IP® system (for example, BIG-IP SiteA) create a new SSL server profile based on the parent profile serverssl.
 - a) On the Main tab, click **Local Traffic > Profiles > SSL > Server**.
 - b) Click the **Create** button.
 - c) In the **Name** box, type wom-serverssl-2.2.2.2.
 - d) From the **Parent Profile** list, select serverssl.
 - e) From the **Configuration** list, select **Advanced** to display more options.
 - f) For the **Certificate** setting, select the associated Custom check box (to override the default setting), and select **wom-endpoint** from the list.
 - g) For the **Key** setting, select the associated Custom check box, and select **wom-endpoint** from the list.

Local Traffic >> Profiles : SSL : Server >> New Server SSL Profile...

General Properties

Name wom-serverssl-2.2.2.2

Parent Profile serverssl

Configuration: Advanced Custom ☐

Certificate wom-endpoint ☒

Key wom-endpoint ☒

- h) In the Server Authentication area, for the **Server Certificate** setting, select the associated Custom check box, and select **require** from the list.
- i) For the **Frequency** setting, select the associated Custom check box, and select **always** from the list.
- j) For the **Authenticate Name** setting, select the associated Custom check box, and type 2.2.2.2.

- k) For the **Trusted Certificates Authorities** setting, select the associated Custom check box, and select **wom-root-ca** from the list.

Server Authentication		Custom <input type="checkbox"/>
Server Certificate	require	<input checked="" type="checkbox"/>
Expire Certificate Response Control	drop	<input type="checkbox"/>
Untrusted Certificate Response Control	drop	<input type="checkbox"/>
Frequency	always	<input checked="" type="checkbox"/>
Retain Certificate	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
Certificate Chain Traversal Depth	9	<input type="checkbox"/>
Authenticate Name	2.2.2.2	<input checked="" type="checkbox"/>
Trusted Certificate Authorities	wom_root_ca_crt	<input checked="" type="checkbox"/>
Certificate Revocation List (CRL)	None	<input type="checkbox"/>

Cancel Repeat Finished

- l) Click **Finished**.
2. On the first BIG-IP system (BIG-IP SiteA in our example), edit the remote endpoint settings.
 - a) On the Main tab, click **Acceleration > Symmetric Optimization > Remote Endpoints**.
 - b) In the IP Address column, click 2.2.2.2 to open the properties screen for that remote endpoint.
 - c) For the **Authentication and Encryption** setting, select wom-serverssl-2.2.2.2.

Acceleration » Symmetric Optimization : Remote Endpoints » **Properties**

⚙️ Properties

Remote Endpoint

Name	2.2.2.2
Partition / Path	Common
IP Address	2.2.2.2
State	<input checked="" type="checkbox"/> Enabled

Outbound iSession to WAN

Outbound Connections	<input checked="" type="checkbox"/> Enabled
Authentication and Encryption	wom-serverssl-2.2.2.2 ▼
Tunnel Port	443
IP Encapsulation Type	Default ▼

Inbound iSession from WAN

SNAT	Default ▼
------	-----------

Cancel Update

- d) Click **Update**.
3. On the second BIG-IP system (BIG-IP SiteB in our example), create a new SSL server profile based on the parent profile `serverssl`.
 - a) On the Main tab, click **Local Traffic** > **Profiles** > **SSL** > **Server**.
 - b) Click the **Create** button.
 - c) In the **Name** box, type `wom-serverssl-1.1.1.1`.
 - d) From the **Parent Profile** list, select `serverssl`.
 - e) From the **Configuration** list, select **Advanced** to display more options.
 - f) For the **Certificate** setting, select the associated Custom check box (to override the default setting), and select **wom-endpoint** from the list.
 - g) For the **Key** setting, select the associated Custom check box, and select **wom-endpoint** from the list.
 - h) In the Client Authentication area, for the **Server Certificate** setting, select the associated Custom check box, and select **require** from the list.
 - i) For the **Frequency** setting, select the associated Custom check box, and select **always** from the list.
 - j) For the **Authenticate Name** setting, select the associated Custom check box, and type `1.1.1.1`.
 - k) For the **Trusted Certificates Authorities** setting, select the associated Custom check box, and select **wom-root-ca** from the list.
 - l) Click **Finished**.
4. On the second BIG-IP system (BIG-IP SiteB in our example), edit the remote endpoint settings.
 - a) On the Main tab, click **Acceleration** > **Symmetric Optimization** > **Remote Endpoints**.

- b) In the IP Address column, click 1.1.1.1 to open the properties screen for that remote endpoint.
- c) For the **Authentication and Encryption** setting, select wom-serverssl-1.1.1.1.
- d) Click **Update**.

Implementation result

After you complete the tasks in this implementation, you have secured the iSession endpoints of your symmetric deployment. The iSession traffic is now secure. Next, you can encrypt data traffic with iSession, using either IPsec for all applications, or SSL on a per-application basis.

Encrypting Application Traffic with iSession

Overview: Encrypting application traffic with iSession

You can use either SSL or IPsec to encrypt application data traffic through a secured iSession™ connection, depending on how you configure symmetric optimization.

- If you are using IPsec, you specify IPsec encapsulation of the data traffic. After the trust relationship is established between the iSession endpoints, the data traffic is encapsulated, regardless of the application.
- If you are using SSL, you specify WAN encryption on a per-application basis when you create an iApps® template for that application. If you manually create an optimized application virtual server for outbound iSession traffic, ensure that you associate an iSession profile that has encryption enabled.

Note: Selecting IPsec encapsulation supersedes any per-application SSL data encryption settings.

Task summary for encrypting application traffic using IPsec

Before you begin encrypting application traffic, you must secure the iSession™ endpoints using SSL.

After the iSession connection is secure, the easiest and quickest method of configuring application data encryption using IPsec is on the Quick Start screen.

Note: For this implementation, creating a custom policy is an optional task.

Task list

Encrypting application traffic using IPsec on the Quick Start screen

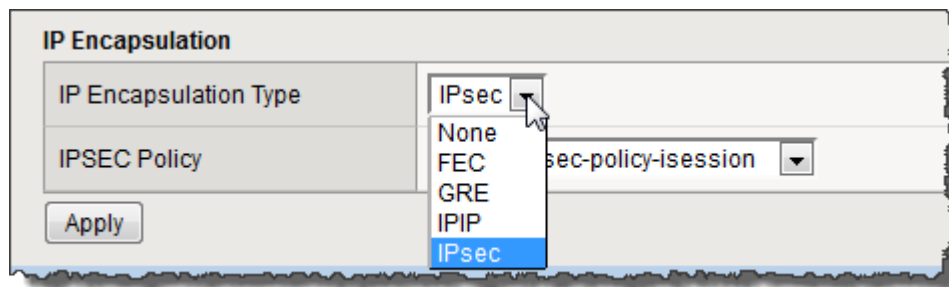
Creating a custom IPsec policy for iSession traffic

Encrypting application traffic using IPsec on the Quick Start screen

You cannot view the Quick Start screen until you have defined at least one VLAN and at least one self IP on a configured BIG-IP® system that is provisioned for acceleration.

You complete this task to encrypt application traffic over an iSession connection using IPsec.

1. On the Main tab, click **Acceleration > Quick Start > Symmetric Properties**.
2. In the IP Encapsulation area, select **IPsec** from the **IP Encapsulation Type** list.



The screen refreshes and displays the **IPSEC Policy** field.

3. From the **IPSEC Policy** list select an IPsec policy.

You can use the pre-defined default policy `default-ipsec-policy-isession`, or create a custom policy, which the system adds to the list.

4. Click **Apply**.

Application traffic is now encrypted over the iSession connection using IPsec, according to the settings in the selected IPsec policy.

Creating a custom IPsec policy for iSession traffic

You can create a custom IPsec policy for iSession traffic if you want settings that are different from the default values. For example, you might want to specify a different authentication algorithm or Diffie-Hellman group for IKE phase 2 negotiations.

1. On the Main tab, click **Network** > **IPsec** > **IPsec Policies**.
2. Click the **Create** button.
The New Policy screen opens.
3. In the **Name** field, type a unique name for the policy.
4. From the **Mode** list, select **iSession Using Tunnel**.
5. From the **Authentication Algorithm** list, select an algorithm.

These are the possible values:

- SHA-1
- AES-GMC128
- AES-GMC192
- AES-GMC256
- AES-GMAC128
- AES-GMAC192
- AES-GMAC256

6. From the **Perfect Forward Secrecy** list, select a Diffie-Hellman group.

These are the possible values:

- MODP768
- MODP1024
- MODP1536
- MODP2048
- MODP3072
- MODP4096
- MODP6144
- MODP8192

7. For the **IPComp** setting, specify whether to use IPComp encapsulation, which performs packet-level compression before encryption:

- Retain the default value **None**, if you do not want to enable packet-level compression before encryption.
- Select **DEFLATE** to enable packet-level compression before encryption.

8. Click **Finished.**

The screen refreshes and displays the new IPsec policy in the list.

For a custom IPsec policy to take effect, you must apply it to the iSession endpoints. You can select it on the Quick Start screen or the Local Endpoint screen. The selected policy settings must be the same on both endpoints of an iSession connection.

Legal Notices

Legal notices

Publication Date

This document was published on August 10, 2018.

Publication Number

MAN-0468-04

Copyright

Copyright © 2017, F5 Networks, Inc. All rights reserved.

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

Trademarks

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

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

Patents

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

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

A

- accelerating cascading style sheet files
 - task summary [55](#)
- accelerating CSS files
 - task summary [55](#)
- accelerating inline image files
 - task summary [55](#)
- accelerating JavaScript files
 - task summary [55](#)
- accelerating JS files
 - task summary [55](#)
- acceleration policies
 - about [9](#)
 - importing [11](#)
 - saving XML file [12](#)
- acceleration policy rules
 - modifying [11](#)
- adaptive IBR
 - adjusting lifetime [53](#)
- adaptive Intelligent Browser Referencing
 - adjusting lifetime [53](#)
- advanced settings
 - for IBR [51](#)
- advertised routes
 - adding manually [119](#), [120](#)
 - description [119](#)
 - modifying automatic discovery of [120](#)
 - verifying discovery [121](#)
- Always proxy requests for this node
 - enabling [17](#)
- application
 - invalidating content [27](#)
- application traffic
 - encrypting with IPsec on Quick Start screen [139](#)
 - encrypting with iSession [139](#)

B

- BIG-IP Cache Settings
 - configuring [23](#)

C

- cache
 - clearing content [27](#)
 - invalidating content [27](#)
- Cache on first hit setting
 - disabling [39](#)
 - enabling [39](#)
 - using [39](#)
- Caching Bytes Saved reports
 - configuring [103](#)
- Caching Requests Saved report
 - configuring [105](#)
- cascading style sheet files
 - about inlining [55](#)
 - about minification [55](#)

- cascading style sheet files (*continued*)
 - about optimization [55](#)
 - about reordering [55](#)
 - Inlining [58](#)
 - minifying [55](#)
 - reordering links [60](#)
 - specifying concatenating resources [62](#)
 - specifying inlining resources [56](#), [57](#)
 - specifying reordering resources [60](#)
- Cascading Style Sheet files
 - about concatenation [62](#)
 - creating concatenation set [62](#)
 - optimizing [63](#)
- Cascading Style Sheet files concatenation
 - task summary [62](#)
- Client Cache Settings
 - configuring [23](#)
- Client IBR'd Links Received report
 - configuring [106](#)
- Client IBR'd Links report
 - configuring [106](#)
- compression
 - adding to IPsec policies [129](#)
 - enabling content assembly on proxies [52](#), [65](#)
 - enabling from origin web server [85](#)
 - overview [77](#)
 - results [64](#)
- Compression Bytes Saved report
 - configuring [104](#)
- concatenation
 - and Cascading Style Sheet files [62](#)
 - and JavaScript files [62](#)
- Configure and use Proxying Rules for this node
 - enabling [18](#)
- content assembly
 - enabling on proxies [52](#), [65](#)
- CSS concatenation set
 - creating [62](#)
- CSS files
 - about inlining [55](#)
 - about minification [55](#)
 - about optimization [55](#)
 - about reordering [55](#)
 - Inlining [58](#)
 - minifying [55](#)
 - See also cascading style sheet files

D

- deduplication
 - described [115](#)
 - disabling [116](#)
 - enabling [115](#)
- deduplication cache
 - clearing [116](#)
- discovery
 - and advertised routes [119](#)
 - enabling for advertised routes [120](#)

- discovery (*continued*)
 - modifying for remote endpoints [123](#)
 - of local subnets [119](#)
 - of remote endpoints [123](#)
- DNS prefetching
 - configuring [56](#)
 - domain lists [56](#)
- DNS subdomains
 - configuring for MultiConnect [66](#)
- domain lists [56](#)

E

- emails
 - sending through SMTP server [102](#)
- ETag
 - including in metadata [87](#)

F

- FEC, *See* forward error correction (FEC)
- FEC profiles
 - customizing [112](#)
- forward error correction (FEC)
 - about [111](#)
 - configuring [112](#), [113](#)
 - overview [111](#)
 - viewing statistics for [113](#)

H

- HTML
 - about inlining [55](#)
 - about minification [55](#)
 - about optimization [55](#)
 - about reordering [55](#)
- HTTP request logging
 - and code elements [97](#)
 - and profile settings [96](#)
- HTTP request logging profile, overview [93](#)
- HTTP request queuing
 - disabling [42](#)
 - enabling [42](#)
 - overview [41](#)
- HTTP traffic
 - enabling MultiConnect [67](#)
 - overview of managing with HTTP2 profile [45](#)
- HTTP/2 profile
 - about [46](#)
 - creating [48](#)
 - overview [46](#)
- HTTP/2 profile settings
 - defined [46](#)
 - listed [46](#)
- HTTP/2 traffic
 - creating virtual servers for [48](#), [49](#)
- HTTP2 profile
 - overview [45](#)
- HTTPS traffic
 - creating a pool to manage [47](#)
 - enabling MultiConnect [67](#)

I

- ICC
 - enabling [91](#)
- ICC Inlined Links report
 - configuring [107](#)
- ICC Referenced Links report
 - configuring [108](#)
- image files
 - inlining [59](#)
 - specifying inlining resources [56](#), [57](#)
- image optimization
 - overview [73](#)
 - task summary [73](#), [77](#)
- Image Optimization Bytes Saved report
 - configuring [104](#)
- images
 - disabling optimization [75](#)
 - optimizing [73](#)
- Inlined Links report
 - configuring [107](#)
- inlining
 - and cascading style sheet files [55](#)
 - and CSS files [55](#)
 - and JavaScript files [55](#)
 - cascading style sheet files [58](#)
 - CSS files [58](#)
 - image files [59](#)
 - JavaScript files [58](#)
- inlining cascading style sheet files
 - task summary [57](#)
- inlining CSS files
 - task summary [57](#)
- inlining inline image files
 - task summary [57](#)
- inlining JavaScript files
 - task summary [57](#)
- inlining JS files
 - task summary [57](#)
- Intelligent Browser Referencing
 - advanced settings [51](#)
 - enabling [52](#)
 - enabling content assembly on proxies [52](#), [65](#)
 - implementation results [54](#)
 - overview [51](#)
 - task summary [35](#), [51](#)
- intelligent client cache
 - enabling [91](#)
 - overview [91](#)
- invalidation
 - overview [27](#)
- invalidations rule
 - blog example [29](#)
 - creating [28](#)
- invalidations rules
 - configuring blog example forumid settings [33](#)
 - configuring blog example Path settings [32](#)
 - configuring blog example postid settings [33](#)
 - configuring blog example view settings [32](#)
 - creating example Post node matching rules [31](#)
 - creating example View node matching rules [30](#)
 - creating leaf nodes for blog example [30](#)

- invalidations rules (*continued*)
 - overview [27](#)
 - specifying for invalidations example Post [32](#)
- IPComp
 - adding to IPsec policies [129](#)
 - compressing non-TCP traffic [127](#)
- IPsec
 - and application traffic [139](#)
 - encrypting iSession application traffic [139](#)
- IPsec policies
 - creating for compression [129](#)
 - customizing for iSession [140](#)
- IPsec tunnels
 - forwarding non-TCP traffic through [127](#)
- iSession
 - encrypting application traffic [139](#)
 - using IPsec to encrypt application traffic [139](#)
- iSession connection
 - customizing SSL profile for [132](#)
 - generating SSL certificates for [131](#)
 - securing with SSL [131](#)
- iSession endpoint security
 - about [131](#)
 - implementation result [137](#)

J

- JavaScript concatenation set
 - creating [63](#)
- JavaScript files
 - about concatenation [62](#)
 - about inlining [55](#)
 - about minification [55](#)
 - about optimization [55](#)
 - about reordering [55](#)
 - concatenating [63](#)
 - creating concatenation set [63](#)
 - inlining [58](#)
 - minifying [55](#)
 - reordering links [61](#)
 - specifying concatenating resources [62](#)
 - specifying inlining resources [56, 57](#)
 - specifying reordering resources [60](#)
- JavaScript files concatenation
 - task summary [62](#)
- JS files, *See* JavaScript files

L

- lifetime cache settings
 - configuring [23](#)
- lifetime rules
 - about [23](#)
- Lifetime rules
 - Default leaf node example [24](#)
 - example [24](#)
 - Home node example [24](#)
 - Image node example [25](#)
 - Search node example [25](#)
- local traffic policy
 - accelerating BIG-IP applications [89](#)
 - accelerating traffic [89](#)

- local traffic policy (*continued*)
 - using to classify types of HTTP traffic [89](#)

M

- Metadata responses
 - about using [87](#)
 - configuring settings [87](#)
 - disabling [87](#)
- minification
 - and HTML [55](#)
 - of cascading style sheet files [55](#)
 - of CSS files [55](#)
 - of JavaScript files [55](#)
- Minification Bytes Saved report
 - configuring [105](#)
- monitoring performance
 - about [101](#)
- MultiConnect
 - configuring subdomains [66](#)
 - enabling for HTTP traffic [67](#)
 - enabling for HTTPS traffic [67](#)
 - overview [65](#)
 - results [68](#)
 - task summary [9, 13, 65](#)

N

- non-TCP iSession traffic
 - forwarding with IPsec encapsulation [127](#)

O

- object
 - classification [35](#)
- Object type
 - creating user-defined [35](#)
 - deleting user-defined [36](#)
 - editing [36](#)
- optimization
 - of cascading style sheet files [55](#)
 - of CSS files [55](#)
 - of JavaScript files [55](#)
- optimized images
 - accelerating [73](#)
 - disabling for a node [75](#)

P

- packet loss
 - mitigating with FEC [111](#)
- parameter value substitution
 - configuring [69](#)
 - serving specific content [69](#)
- parameters
 - for HTTP request logging [97](#)
 - for request logging [97](#)
- PDF linearization
 - disabling [72](#)
 - enabling [71](#)
 - overview [71](#)

- PDF linearization (*continued*)
 - task summary 71
- performance monitoring
 - about 101
 - disabling 101
 - enabling 101
- Policies
 - copying 10
 - creating user-defined 10
 - creating user-defined from predefined 10
 - deleting user-defined 12
- Policies screen
 - viewing rules 12
- Policy Editor screen
 - accessing 13
 - overview 13
- Policy Tree
 - viewing 14
- Policy Viewer screen
 - accessing 9
- pools
 - creating 82, 102
 - creating for HTTPS traffic 47
 - creating with request logging 93
- predefined policy
 - copying 10
- profile settings
 - for HTTP/2 46
- profile types
 - for HTTP/2 46
- profiles
 - customizing for FEC tunnel 112
- proxy override rule
 - configuring example 21
 - example of configuring 21
- proxy requests
 - overriding 19
- proxying rule
 - about creating an example 19
 - configuring example 20
 - creating example 20
- proxying rule parameters
 - example of configuring 20
- proxying rules
 - overview 17

Q

- Quick Start screen
 - about 125
 - configuring iSession endpoints 125
 - encrypting application traffic with IPsec 139

R

- remote endpoints
 - about discovery of 123
 - configuring SSL on 134
 - modifying discovery of 123
 - verifying discovery 121, 123
- reordering
 - of cascading style sheet files 55

- reordering (*continued*)
 - of CSS files 55
 - of JavaScript files 55
- reordering cascading style sheet files
 - task summary 60
- reordering CSS files
 - task summary 60
- reordering JavaScript files
 - task summary 60
- reordering JS files
 - task summary 60
- request logging profile
 - creating 93
 - deleting 95
 - enabling for requests 94
 - enabling for responses 95
 - overview 93
 - settings 96
- request logging, and code elements 97
- ROI reports
 - configuring caching bytes saved data 103
 - configuring caching requests saved data 105
 - configuring client IBR'd links data 106
 - configuring client IBR'd links received data 106
 - configuring compression bytes saved data 104
 - configuring ICC inlined links data 107
 - configuring ICC referenced links data 108
 - configuring image optimization bytes saved data 104
 - configuring inlined links data 107
 - configuring minification bytes saved data 105
- ROI statistics
 - about 102
 - exporting to server 109
 - overview 103
 - task summary 102

S

- SDD, See symmetric data deduplication
- SMTP server
 - configuring 102
- SSL
 - and application traffic 139
 - and iSession endpoint security 131
 - configuring for iSession connection 131
 - customizing profile for iSession 132
 - generating certificates for iSession connection 131
- subnets
 - about discovery of 119
 - verifying discovery 121
- symmetric data deduplication
 - described 115
 - disabling 116
 - enabling 115

T

- tunnels
 - configuring for FEC 113
 - configuring for FEC, receiving 112
 - customizing FEC profile for 112

U

- URL resources
 - specifying for concatenating [62](#)
 - specifying for inlining [56](#), [57](#)
 - specifying for reordering [60](#)
- user-defined policy
 - copying [10](#)
 - creating [10](#)
 - creating from predefined [10](#)
 - deleting [12](#)
 - modifying [11](#)
 - publishing [11](#)

V

- variation rules
 - configuring [14](#)
 - configuring ambiguous query parameters [15](#)
 - example [15](#)
 - overview [14](#)
 - Query Parameter rule example [16](#)
 - Referrer rule example [16](#)
- video advertisement policy
 - creating [77](#)
 - deleting [78](#)
 - modifying [78](#)
- video delivery optimization
 - enabling [78](#)
 - modifying [79](#)
- video Quality of Experience
 - creating iRule to collect scores [80](#)
 - creating iRule to collect static information [81](#)
 - creating profile [82](#)
 - creating virtual server [82](#)
- Video Quality of Experience
 - overview [80](#)
- VIPRION
 - about acceleration in a cluster [37](#)
 - using clustered acceleration [37](#)
- virtual servers
 - adding to advertised routes [119](#)
 - assigning a Request Logging profile [95](#)
 - compressing non-TCP iSession traffic [127](#)
 - creating for HTTP/2 traffic [48](#), [49](#)
 - creating for IP iSession traffic [127](#)
 - creating for video Quality of Experience [82](#)
 - forwarding non-TCP iSession traffic [127](#)
 - using iSession routing [127](#)

