BIG-IP® Access Policy Manager®: Implementations

Version 13.1

Table of Contents

Web Access Management

Overview: Configuring APM for web access management

Access Policy Manager® (APM®) web access management provides the ability to access web applications through a web browser without the use of tunnels or specific resources. With this type of access, APM communicates with backend web servers, forwarding requests from the client to web servers within a local traffic pool.

In a configuration that controls traffic and requests directed to your internal servers, using Access Policy Manager® (APM®) with Local Traffic Manager® provides additional security. APM communicates with backend web servers, forwarding requests from the client to web servers within a local traffic pool. APM allows access to the local traffic pool only after the user passes through an access policy that typically contains authentication actions, endpoint security checks, and ACLs.

Task summary

Creating a pool
Creating an access profile
Verifying log settings for the access profile
Creating an access policy for web access management
Creating a virtual server

About ways to time out a web access management session

The web access management access type does not have a logout mechanism; as a result configuring a timeout is important. Access Policy Manager® (APM®) provides these options.

The Windows Cache and Session Control access policy item
Terminates a user session when it detects that the browser screen has closed. You can also configure it to provide inactivity timeouts for the user session using the Terminate session on user inactivity setting.

Maximum Session Timeout access profile setting
Provides an absolute limit for the duration of the access policy connection, regardless of user activity. To ensure that a user session closes after a certain number of seconds, configure this setting.

Inactivity Timeout access profile setting
Terminates the session after there is no traffic flow for a specified number of seconds.

Note: *Depending on the application, you might not want to set this to a very short duration, because many applications cache user typing and generate no traffic for an extended period. In this scenario, a session can time out while the application is still in use, but the content of the user input is not relayed back to the server.*

.

Creating a pool

You can create a pool of servers for Access Policy Manager® (APM®) to perform access control for web application servers configured as local traffic pool members.

Important: *When you implement a service with multiple hosts, access through the virtual server for new requests causes the load balancing algorithm for the associated member pool to select a new server. This can cause problems if persistence to a particular host is required.*

Note: *When you add web servers as members of the pool, select the HTTPS service if the web server uses SSL, to maintain consistency between APM and the web servers.*

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. In the Resources area, for the **New Members** setting, add to the pool the application servers that host the web application:
   a) Type an IP address in the **Address** field.
   b) In the **Service Port** field, type a port number (for example, type `80` for the HTTP service), or select a service name from the list.
   c) Click **Add**.
5. Click **Finished**.

The new pool appears in the Pools list.

Creating an access profile

You create an access profile to provide the per-session policy configuration for a virtual server that establishes a secured session. In the access profile, you can also specify a timeout to use to terminate a web access management connection

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click **Create**.
   The New Profile screen opens.
3. In the **Name** field, type a name for the access profile.

   Note: *A access profile name must be unique among all access profile and any per-request policy names.*

4. From the **Profile Type** list, select **LTM-APM**.
   With this type selected, when you configure the access policy, only access policy items that are applicable for web access management are displayed.
5. In the **Inactivity Timeout** field, type the number of seconds that should pass before the access policy times out. Type `0` to set no timeout.

   The web access management connection type does not provide a logout mechanism. You should configure at least one timeout for the connection, either in this access profile, or by including the Windows Cache and Session Control item in the access policy and configuring a timeout in it.
6. In the **Maximum Session Timeout** field, type the maximum number of seconds the session can exist.

   Type `0` to set no timeout.
7. In the Language Settings area, add and remove accepted languages, and set the default language.

   A browser uses the highest priority accepted language. If no browser language matches the accepted languages list, the browser uses the default language.
8. Click **Finished**.

The access profile displays in the Access Profiles List. Default-log-setting is assigned to the access profile.

Verifying log settings for the access profile

Confirm that the correct log settings are selected for the access profile to ensure that events are logged as you intend.

Note: *Log settings are configured in the* ***Access*** > ***Overview*** > ***Event Log*** > ***Settings*** *area of the product. They enable and disable logging for access system and URL request filtering events. Log settings also specify log publishers that send log messages to specified destinations.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. On the menu bar, click **Logs**.
   The access profile log settings display.
4. Move log settings between the **Available** and **Selected** lists.

   You can assign up to three log settings that enable access system logging to an access profile. You can assign additional log settings to an access profile provided that they enable logging for URl request logging only.

   Note: *Logging is disabled when the* ***Selected*** *list is empty.*

   ---

5. Click **Update**.

An access profile is in effect when it is assigned to a virtual server.

Creating an access policy for web access management

You create an access policy to specify, at a minimum, logon and authentication. You can add other items to the policy to direct traffic and grant or deny access appropriately, increasing your security.

Note: *In an access policy for web access management, you do not need to assign resources, such as, webtops, portal access or network access resources, application access tunnels, or remote desktops.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.
3. On a policy branch, click the **(+)** icon to add an item to the policy.
   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
4. On the Logon tab, select **Logon Page** and click the **Add Item** button.
   The Logon Page Agent properties screen opens.
5. Make any changes that you require to the logon page properties and click **Save**.
   The properties screen closes and the policy displays.
6. On a policy branch, click the **(+)** icon to add an item to the policy.

   Repeat this action from the visual policy editor whenever you want to add an item to the policy.

   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
7. From the Authentication tab, select an authentication item.

8. Configure the properties for the authentication item and click **Save** when you are done.

   You can configure multiple authentication items in an access policy.

   You have now configured a basic access policy.
9. Add endpoint security checks or other items that you require to the access policy.

   Optionally, you can assign a pool of web servers in the access policy using the Pool Assign action; if you do, this pool takes precedence over the pool you assign to the virtual server configuration.

   Note: *You can add a* ***Windows Cache and Session Control*** *item to configure a way to terminate the session.*
10. To grant access at the end of any branch, change the ending from **Deny** to **Allow**:
    a) Click **Deny**.

       The default branch ending is **Deny**.

       A popup screen opens.
    b) Select **Allow** and click **Save**.
       The popup screen closes. The **Allow** ending displays on the branch.
11. Click the **Apply Access Policy** link to apply and activate the changes to the policy.

This creates an access policy that is appropriate for web access management connections.

To apply this access policy to network traffic, add the access profile to a virtual server.

Note: *To ensure that logging is configured to meet your requirements, verify the log settings for the access profile.*

Creating a virtual server

This task creates a standard, host type of virtual server for application traffic. A host type of virtual server listens for traffic destined for the specified destination IP address and service. Using this virtual server, Access Policy Manager® (APM®) can provide access control for web applications on web servers in a local traffic pool without using tunnels or specific resources.

Note: *By default, the health monitor is set to none and the load balancing method is set to Round Robin. You can add a health monitor or select an alternative load balancing method for this 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 for a host virtual server.

   This field accepts an address in CIDR format (IP address/prefix). However, when you type the complete IP address for a host, you do not need to type a prefix after the address.
5. In the **Service Port** field, type `80` (for HTTP) or `443` (for HTTPS), or select **HTTP** or **HTTPS** from the list.
6. For the **HTTP Profile** setting, verify that the default HTTP profile, **http**, is selected.
7. (Optional) For the **SSL Profile (Client)** setting, select a client SSL profile.

   If the web server uses SSL, the client should use SSL.
8. (Optional) For the **SSL Profile (Server)** setting, select an SSL server profile.

   If the web server uses SSL, the virtual server should use SSL.
9. In the Content Rewrite area, retain the default settings.

The web access management access type eliminates the need for content rewriting. The default values for the **Rewrite Profile** and the **HTML Profile** settings are **None**.

10. In the Access Policy area, from the **Access Profile** list, select the access profile you configured previously.

    Retain the default values for other settings in the Access Policy area.

11. (Optional) From the **HTTP Compression Profile** list, select **httpcompression**.

    You can use compression to provide a better end user experience, particularly where there is limited bandwidth or high latency between the virtual server and the client.

12. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
13. Click **Finished**.

You have a virtual server that supports web access management connections.

Protecting Internal Resources Per-Request

Overview: Protecting internal resources on a per-request basis

You can use a per-request policy to protect your internal resources and to be more selective about who accesses them and when. After a user starts a session, a per-request policy makes it possible to apply additional criteria for access any time the user makes a request. These steps are for use in a reverse proxy configuration; that is, with APM® and LTM® set up for web access management.

Task summary

Creating a per-request policy

You must create a per-request policy before you can configure it in the visual policy editor.

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. Click **Create**.
   The General Properties screen opens.
3. In the **Name** field, type a name for the policy and click **Finished**.

   A per-request policy name must be unique among all per-request policy and access profile names.

   The policy name appears on the Per-Request Policies screen.

Configuring policies to branch by local database user group

If you plan to look up local database groups from the per-request policy, you must configure local database-related items in the access policy and the per-request policy to use the same session variable.

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.
3. On a policy branch, click the **(+)** icon to add an item to the policy.
   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
4. In the search field, type `local`, select **Local Database**, and click **Add Item**.
   A popup properties screen opens.
5. Configure properties for the Local Database action:
   a) From the **LocalDB Instance** list, select a local user database.
   b) Click **Add new entry**

A new line is added to the list of entries with the Action set to **Read** and other default settings.

c) In the **Destination** column in the **Session Variable** field, type the name of the variable in which to store the user groups retrieved from the local database.

In the per-request policy, the default value that the LocalDB Group Lookup item uses is `session.localdb.groups`. If you enter a differentvalue, note it. You will need it to update the advanced expression in the LocalDB Group Lookup item in the per-request policy.

d) In the **Source** column from the **DB Property** list, select **groups**.

e) Click **Save**.
The properties screen closes. The policy displays.

This is not a complete access policy, but you can return to it and complete it later. You can close the visual policy editor or leave it open.

The access policy includes a Local Database action that can read groups into a session variable.

6. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
The Per-Request Policies screen opens.
7. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
The visual policy editor opens in another tab.
8. Click the **(+)** icon anywhere in the per-request policy to add a new item.
9. In the search field, type `local`, select **LocalDB Group Lookup**, and click **Add Item**.
A popup properties screen opens.
10. Click the Branch Rules tab.
11. Click the **change** link in the entry for the default expression.
A popup screen opens.
12. If the session variable you typed in the access policy Local Database action was `session.localdb.groups`, perform these substeps.
    a) In the **User is a member of** field, remove `MY_GROUP` and type the name of a group.
    b) Click **Finished**.
    The popup screen closes.
    c) Click **Save**.
    The properties screen closes and the policy displays.
13. If you typed a session variable other than `session.localdb.groups` in the access policy Local Database action, perform these substeps.
    a) Click the Advanced tab.
    In the field, this expression displays. `expression is expr { [mcget {session.localdb.groups}] contains "MY_GROUP" }`
    a) In the expression, replace `session.localdb.groups` with the name of the session variable you typed into the Local Database action.
    b) In the expression, replace `MY_GROUP` with the name of a group that should match a local database group.
    c) Click **Finished**.
    The popup screen closes.
    d) Click **Save**.
    The properties screen closes and the policy displays.

    This is not a complete per-request access policy, but you can return to it and complete it later.

The access and per-request policies are configured to use the same session variable. The access policy is configured to support the use of LocalDB Group Lookup in the per-request policy.

Complete the configuration of the access and per-request policies.

Categorizing URLs using custom categories in a per-request policy

Important: *These steps apply to a BIG-IP® system on which URL categories are available only by creating them in Access Policy Manager® (APM®).*

If you haven't configured URL categories and URL filters yet in APM, configure them before you start this task.

Look up the category for a URL request and use it in a policy branch rule, or to assign a URL filter, and so on.

Note: *These steps provide guidance for adding items to control traffic based on the URL category; they do not specify a complete per-request policy.*

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. Add a **Category Lookup** item and set its properties:

   Important: *A Category Lookup item triggers event logging for URL requests and provides categories for a URL Filter Assign item.*

   a) From the **Categorization Input** list, select an entry based on the type of traffic to be processed. .
      - For HTTP traffic, select **Use HTTP URI (cannot be used for SSL Bypass decisions)**.
      - For SSL-encrypted traffic, select **Use SNI in Client Hello (if SNI is not available, use Subject.CN)**.
      - **Use Subject.CN in Server Cert** is not supported for reverse proxy.

   b) For **Category Lookup Type**, you can only retain the default setting **Process custom categories only**.

   a) Click **Save**.
      The properties screen closes. The policy displays.
4. To add a **URL Filter Assign** item, do so anywhere on a branch after a Category Lookup item.

   A URL filter applies to the categories that a Category Lookup item returns. If the filter specifies the **Block** action for any URL category, URL Filter Assign blocks the request.

   Note: *If URL Filter Assign does not block the request and the filter specifies the confirm action for any URL category, URL Filter Assign takes the **Confirm** per-request policy branch and the policy exits on the ending for it.*

   a) From the **URL Filter** list, select a URL filter.
   b) To simplify the display in the visual policy editor if the URL filter does not specify confirm actions, select Branch Rules, and click **x** on the **Confirm** entry.
   c) Click **Save**.
      The properties screen closes and the policy displays.

Now the per-request policy includes an item that looks up the URL category. You can add other items to the policy to control access according to your requirements.

Note: *SSL bypass and SSL intercept are not supported when you are protecting internal resources from incoming requests. They are supported in a forward proxy configuration.*

A per-request policy goes into effect when you add it to a virtual server.

Configuring a per-request policy to control access to applications

Access Policy Manager® (APM®) supports a preset group of application families and applications. You can configure your own application filters or use one of the filters that APM provides: block-all, allow-all, and default.

Configure a per-request policy to specify the logic that determines whether to allow access to the applications or application families.

Note: *This task provides the steps for adding items to control requests based on the application name or application family or based on an application filter. It does not specify a complete per-request policy.*

---

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. Add an **Application Lookup** item to the policy.
   a) Click the **(+)** icon anywhere in the per-request policy to add a new item.
   A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.
   b) From the General Purpose tab, select **Application Lookup**, and click **Add Item**.
   A Properties popup screen opens.
   c) Click **Save**.
   The Properties screen closes. The visual policy editor displays. A single branch, fallback, follows the **Application Lookup** item.
4. To branch by application family or application name, add branch rules to the **Application Lookup** item.
   a) Click the name of the application lookup item.
   A Properties popup screen displays.
   b) Click the Branch Rules tab.
   c) Click **Add Branch Rule**.
   A new entry with **Name** and **Expression** settings displays.
   d) Click the **change** link in the new entry.
   A popup screen opens.
   e) Click the **Add Expression** button.
   Settings are displayed.
   f) For **Agent Sel**, select **Application Lookup**.
   g) For **Condition** select **Application Family** or **Application Name**.

   a) From the list, **Application Family is** or **Application Name is**, select a family or name.

   a) Click **Add Expression**.
   The expression displays.
   b) Continue adding branches and when you are done, click **Finished**.
   The popup screen closes. The Branch Rules popup screen displays.
   c) Click **Save**.
   The visual policy editor displays.

   Newly created branches follow the **Application Lookup** item.
5. To apply an application filter to the request, add an **Application Filter Assign** item on a branch somewhere after the Application Lookup item.
   A Properties popup screen displays.

6. From the **Application Filter** list, select an application filter and click **Save**.
   The popup screen closes.

To put the per-request policy into effect, add it to the virtual server.

Important: *To support application filtering, classification must be enabled on the virtual server.*

Configuring a per-request policy to branch by group or class

Add a group or class lookup to a per-request policy when you want to branch by user group or class.

Note: *The access policy must be configured to populate session variables for a group or class lookup to succeed. This task provides the steps for adding items to branch by group or class. It does not specify a complete per-request policy.*

---

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. On a policy branch, click the **(+)** icon to add an item to the policy.

   A small set of actions are provided for building a per-request policy.

   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
4. On the Authentication tab, select an option: **AD Group Lookup**, **LDAP Group Lookup**, or **RADIUS Class Lookup** to the per-request policy.
5. Click **Add Item**.
   A properties popup screen opens.
6. Click the Branch Rules tab.
7. To edit an expression, click the **change** link.
   An additional popup screen opens, displaying the Simple tab.
8. Edit the default simple expression to specify a group or class that is used in your environment.
   In an LDAP Group Lookup item, the default simple expression is **User is a member of** `CN=MY_GROUP, CN=USERS, CN=MY_DOMAIN`. You can use the simple expression editor to replace the default values.
9. Click **Finished**.
   The popup screen closes.
10. Click **Save**.
    The popup screen closes. The visual policy editor displays.

A per-request policy goes into effect when you add it to a virtual server.

Adding a per-request policy to the virtual server

To add per-request processing to a configuration, associate the per-request policy with the virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server.
3. In the Access Policy area, from the **Per-Request Policy** list, select the policy that you configured earlier.
4. Click **Update**.

The per-request policy is now associated with the virtual server.

Example policy: URL filter per user group

Each URL Filter Assign item in this per-request policy example should specify a filter that is applicable to the user group.

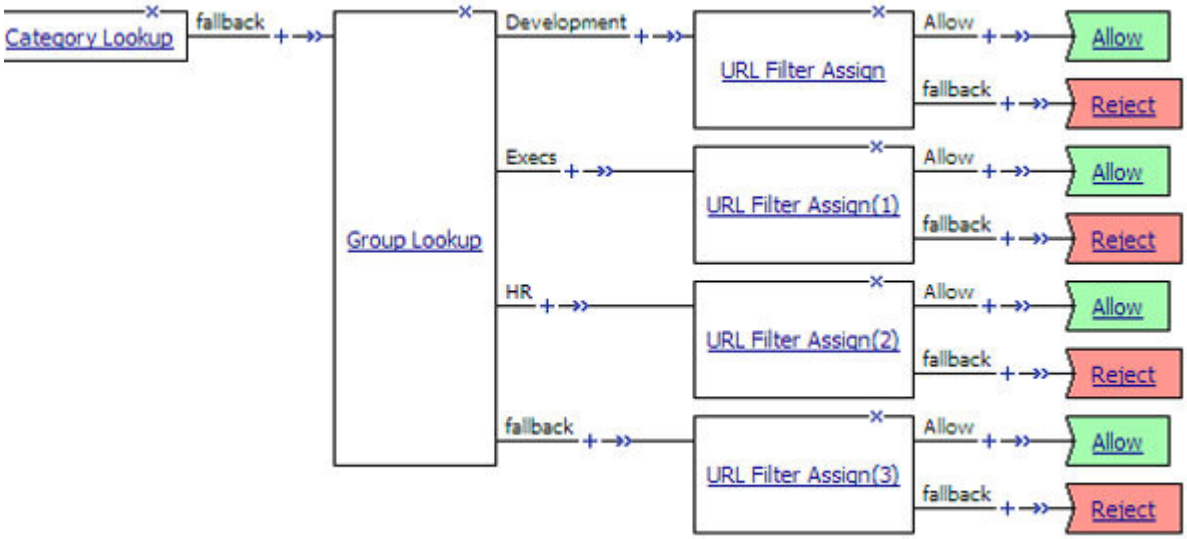

Figure 1: URL filter based on group membership

Example policy: Access control by date, time, and user group

This per-request policy example applies specific URL filters for weekends and weeknights, and restricts access during work hours based on user group.

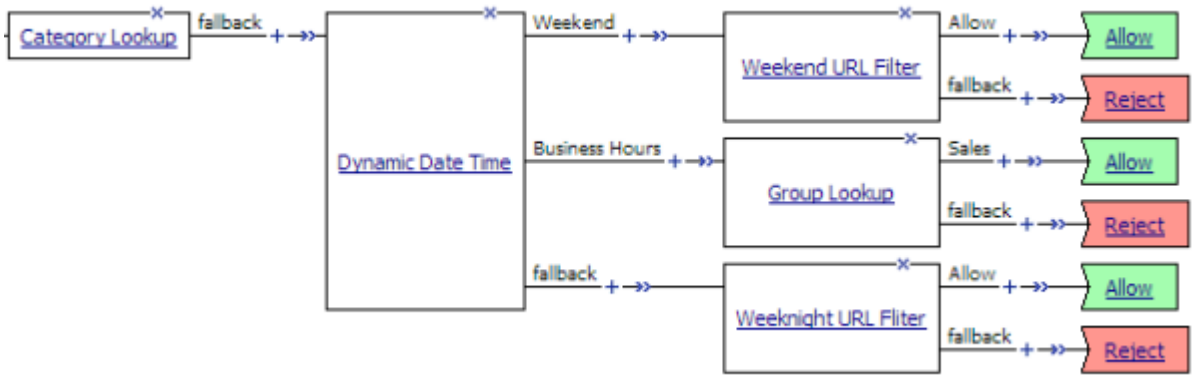

Figure 2: Deny or allow access based on date and time and group membership

Example policy: User-defined category-specific access control

In this per-request policy example, only recruiters are allowed to access URLs in the user-defined category Employment. The policy also restricts access to entertaining videos during business hours.

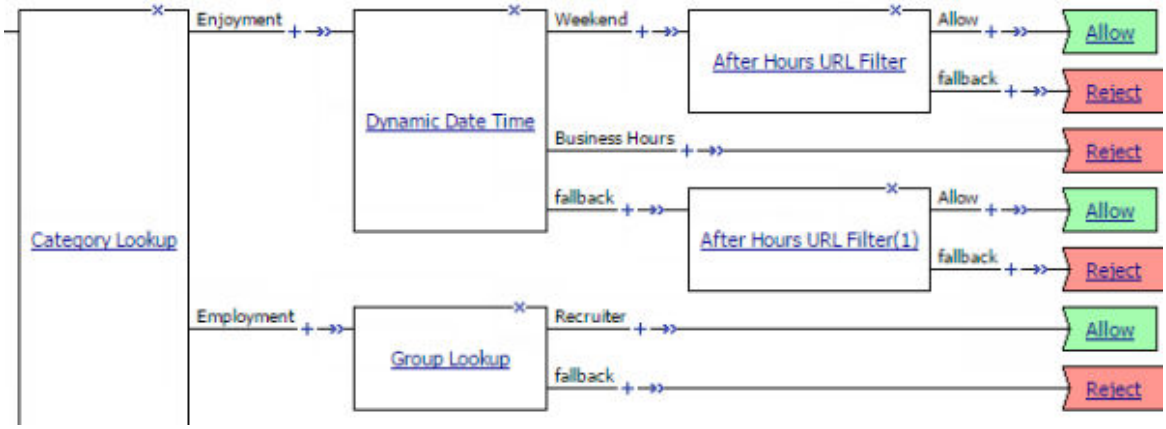

Figure 3: Category-specific access restrictions (using user-defined categories)

Example policy: Application lookup and filter

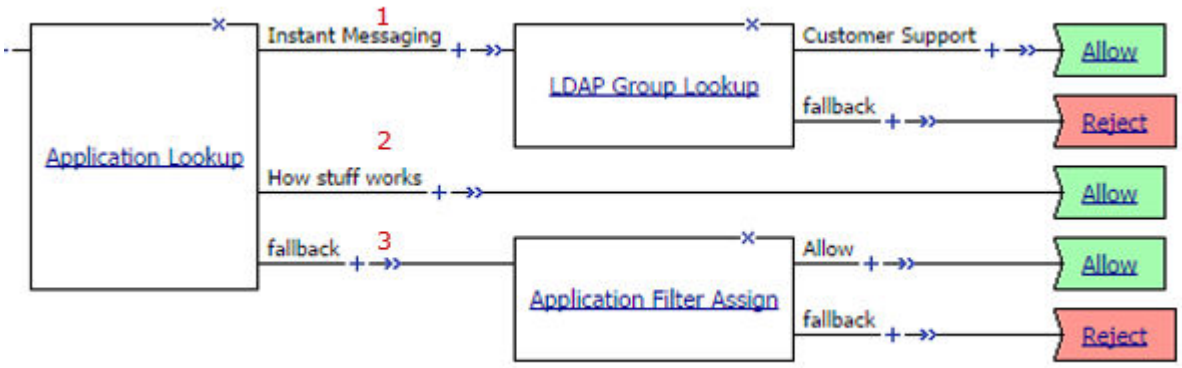

Figure 4: Application access control by application family, application name, and application filter

1	A user-defined branch for the instant messaging application family.
2	A user-defined branch for a specific application.
3	The default fallback branch, on which an application filter is applied. Application Filter Assign needs the information provided by Application Lookup.

LTM SSL Forward Proxy and Per-Request Policies

Overview: Adding a per-request policy to LTM SSL forward proxy

If you have an LTM® SSL forward proxy configuration, you can add a per-request policy to it. Every time a client makes a URL request, the per-request policy runs. The policy can contain any available per-request policy action item, including those for URL and application categorization and filtering.

Complete these tasks before you start:

- Configure any application filters that you want to use.
- Configure any URL filters (and user-defined URL categories) that you want to use.
- Configure a per-request policy.
- Have an LTM SSL forward proxy configuration set up.

Task summary

Creating an access profile for LTM-APM
Creating a per-request policy
Creating a DNS resolver
Updating the virtual server for SSL forward proxy

Creating an access profile for LTM-APM

You create an access profile to provide the access policy configuration for a virtual server that establishes a secured session.

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click **Create**.
   The New Profile screen opens.
3. In the **Name** field, type a name for the access profile.

   Note: *A access profile name must be unique among all access profile and any per-request policy names.*

4. From the **Profile Type** list, select **LTM-APM**.
   Additional settings display.
5. In the Language Settings area, add and remove accepted languages, and set the default language.

   A browser uses the highest priority accepted language. If no browser language matches the accepted languages list, the browser uses the default language.
6. Click **Finished**.
   This creates an access profile with a default access policy.

The access profile displays in the Access Profiles List. Default-log-setting is assigned to the access profile.

You can configure the access policy further but you are not required to do so.

Verifying log settings for the access profile

Confirm that the correct log settings are selected for the access profile to ensure that events are logged as you intend.

Note: *Log settings are configured in the* ***Access*** > ***Overview*** > ***Event Log*** > ***Settings*** *area of the product. They enable and disable logging for access system and URL request filtering events. Log settings also specify log publishers that send log messages to specified destinations.*

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. On the menu bar, click **Logs**.
   The access profile log settings display.
4. Move log settings between the **Available** and **Selected** lists.

   You can assign up to three log settings that enable access system logging to an access profile. You can assign additional log settings to an access profile provided that they enable logging for URl request logging only.

   Note: *Logging is disabled when the* ***Selected*** *list is empty.*
5. Click **Update**.

An access profile is in effect when it is assigned to a virtual server.

Creating a per-request policy

You must create a per-request policy before you can configure it in the visual policy editor.

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. Click **Create**.
   The General Properties screen opens.
3. In the **Name** field, type a name for the policy and click **Finished**.

   A per-request policy name must be unique among all per-request policy and access profile names.

   The policy name appears on the Per-Request Policies screen.

Processing SSL traffic in a per-request policy

To use SSL forward proxy bypass in a per-request policy, both the server and client SSL profile must enable SSL forward proxy and SSL forward proxy bypass; and, in the client SSL profile, the default bypass action must be set to **Intercept**.

Important: *Configure a per-request policy so that it completes processing of HTTPS requests before it starts the processing of HTTP requests.*

Note: *These steps describe how to add items for controlling SSL web traffic to a per-request policy; the steps do not specify a complete per-request policy.*

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. To process the HTTPS traffic first, configure a branch for it by adding a **Protocol Lookup** item at the start of the per-request policy.

a) Click the **(+)** icon anywhere in the per-request policy to add a new item.
A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.
b) In the Search field, type `prot`, select **Protocol Lookup**, and click **Add Item**.
A properties popup screen opens.
c) Click **Save**.
The properties screen closes. The policy displays.

The Protocol Lookup item provides two default branches: HTTPS for SSL traffic and fallback.

4. Before you add an SSL Bypass Set, or an SSL Intercept Set, item to the per-request policy, you can insert any of the following policy items to do logging or to base how you process the SSL traffic on group membership, class attribute, day of the week, time of day, or URL category:

- AD Group Lookup
- LDAP Group Lookup
- LocalDB Group Lookup
- RADIUS Class Lookup
- Dynamic Date Time
- Logging
- Category Lookup

Important: *Category Lookup is valid for processing SSL traffic only when configured for SNI or Subject.CN categorization input and only before any HTTP traffic is processed.*

If you insert other policy items that inspect the SSL payload (HTTP data) before an SSL Bypass Set item, the SSL bypass cannot work as expected.

5. At any point on the HTTPS branch where you decide to bypass SSL traffic, add an **SSL Bypass Set** item.

The per-request policy includes items that you can use to complete the processing of SSL traffic. Add other items to the policy to control access according to your requirements.

A per-request policy goes into effect when you add it to a virtual server. Depending on the forward proxy configuration, you might need to add the per-request policy to more than one virtual server.

Configuring policies to branch by local database user group

If you plan to look up local database groups from the per-request policy, you must configure local database-related items in the access policy and the per-request policy to use the same session variable.

1. On the Main tab, click **Access** > **Profiles / Policies**.
The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
The visual policy editor opens the access policy in a separate screen.
3. On a policy branch, click the **(+)** icon to add an item to the policy.
A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
4. In the search field, type `local`, select **Local Database**, and click **Add Item**.
A popup properties screen opens.
5. Configure properties for the Local Database action:
a) From the **LocalDB Instance** list, select a local user database.
b) Click **Add new entry**
A new line is added to the list of entries with the Action set to **Read** and other default settings.

c) In the **Destination** column in the **Session Variable** field, type the name of the variable in which to store the user groups retrieved from the local database.

In the per-request policy, the default value that the LocalDB Group Lookup item uses is `session.localdb.groups`. If you enter a differentvalue, note it. You will need it to update the advanced expression in the LocalDB Group Lookup item in the per-request policy.

d) In the **Source** column from the **DB Property** list, select **groups**.

e) Click **Save**.
The properties screen closes. The policy displays.

This is not a complete access policy, but you can return to it and complete it later. You can close the visual policy editor or leave it open.

The access policy includes a Local Database action that can read groups into a session variable.

6. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
The Per-Request Policies screen opens.
7. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
The visual policy editor opens in another tab.
8. Click the **(+)** icon anywhere in the per-request policy to add a new item.
9. In the search field, type `local`, select **LocalDB Group Lookup**, and click **Add Item**.
A popup properties screen opens.
10. Click the Branch Rules tab.
11. Click the **change** link in the entry for the default expression.
A popup screen opens.
12. If the session variable you typed in the access policy Local Database action was `session.localdb.groups`, perform these substeps.
    a) In the **User is a member of** field, remove `MY_GROUP` and type the name of a group.
    b) Click **Finished**.
    The popup screen closes.
    c) Click **Save**.
    The properties screen closes and the policy displays.
13. If you typed a session variable other than `session.localdb.groups` in the access policy Local Database action, perform these substeps.
    a) Click the Advanced tab.
    In the field, this expression displays. `expression is expr { [mcget {session.localdb.groups}] contains "MY_GROUP" }`
    a) In the expression, replace `session.localdb.groups` with the name of the session variable you typed into the Local Database action.
    b) In the expression, replace `MY_GROUP` with the name of a group that should match a local database group.
    c) Click **Finished**.
    The popup screen closes.
    d) Click **Save**.
    The properties screen closes and the policy displays.

    This is not a complete per-request access policy, but you can return to it and complete it later.

The access and per-request policies are configured to use the same session variable. The access policy is configured to support the use of LocalDB Group Lookup in the per-request policy.

Complete the configuration of the access and per-request policies.

Categorizing URLs using custom categories in a per-request policy

Important: *These steps apply to a BIG-IP® system on which URL categories are available only by creating them in Access Policy Manager® (APM®).*

If you haven't configured URL categories and URL filters yet in APM, configure them before you start this task.

Look up the category for a URL request and use it in a policy branch rule, or to assign a URL filter, and so on.

Note: *These steps provide guidance for adding items to control traffic based on the URL category; they do not specify a complete per-request policy.*

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. Add a **Category Lookup** item and set its properties:

   Important: *A Category Lookup item triggers event logging for URL requests and provides categories for a URL Filter Assign item.*

   a) From the **Categorization Input** list, select an entry based on the type of traffic to be processed. .
      - For HTTP traffic, select **Use HTTP URI (cannot be used for SSL Bypass decisions)**.
      - For SSL-encrypted traffic, select **Use SNI in Client Hello (if SNI is not available, use Subject.CN)**.
      - **Use Subject.CN in Server Cert** is not supported for reverse proxy.

   b) For **Category Lookup Type**, you can only retain the default setting **Process custom categories only**.

   a) Click **Save**.
      The properties screen closes. The policy displays.
4. To add a **URL Filter Assign** item, do so anywhere on a branch after a Category Lookup item.

   A URL filter applies to the categories that a Category Lookup item returns. If the filter specifies the **Block** action for any URL category, URL Filter Assign blocks the request.

   Note: *If URL Filter Assign does not block the request and the filter specifies the confirm action for any URL category, URL Filter Assign takes the **Confirm** per-request policy branch and the policy exits on the ending for it.*

   a) From the **URL Filter** list, select a URL filter.
   b) To simplify the display in the visual policy editor if the URL filter does not specify confirm actions, select Branch Rules, and click **x** on the **Confirm** entry.
   c) Click **Save**.
      The properties screen closes and the policy displays.

Now the per-request policy includes an item that looks up the URL category. You can add other items to the policy to control access according to your requirements.

Note: *SSL bypass and SSL intercept are not supported when you are protecting internal resources from incoming requests. They are supported in a forward proxy configuration.*

A per-request policy goes into effect when you add it to a virtual server.

Configuring a per-request policy to control access to applications

Access Policy Manager® (APM®) supports a preset group of application families and applications. You can configure your own application filters or use one of the filters that APM provides: block-all, allow-all, and default.

Configure a per-request policy to specify the logic that determines whether to allow access to the applications or application families.

Note: *This task provides the steps for adding items to control requests based on the application name or application family or based on an application filter. It does not specify a complete per-request policy.*

---

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. Add an **Application Lookup** item to the policy.
   a) Click the **(+)** icon anywhere in the per-request policy to add a new item.
   A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.
   b) From the General Purpose tab, select **Application Lookup**, and click **Add Item**.
   A Properties popup screen opens.
   c) Click **Save**.
   The Properties screen closes. The visual policy editor displays. A single branch, fallback, follows the **Application Lookup** item.
4. To branch by application family or application name, add branch rules to the **Application Lookup** item.
   a) Click the name of the application lookup item.
   A Properties popup screen displays.
   b) Click the Branch Rules tab.
   c) Click **Add Branch Rule**.
   A new entry with **Name** and **Expression** settings displays.
   d) Click the **change** link in the new entry.
   A popup screen opens.
   e) Click the **Add Expression** button.
   Settings are displayed.
   f) For **Agent Sel**, select **Application Lookup**.
   g) For **Condition** select **Application Family** or **Application Name**.

   a) From the list, **Application Family is** or **Application Name is**, select a family or name.

   a) Click **Add Expression**.
   The expression displays.
   b) Continue adding branches and when you are done, click **Finished**.
   The popup screen closes. The Branch Rules popup screen displays.
   c) Click **Save**.
   The visual policy editor displays.

   Newly created branches follow the **Application Lookup** item.
5. To apply an application filter to the request, add an **Application Filter Assign** item on a branch somewhere after the Application Lookup item.
   A Properties popup screen displays.
6. From the **Application Filter** list, select an application filter and click **Save**.

The popup screen closes.

To put the per-request policy into effect, add it to the virtual server.

Important: *To support application filtering, classification must be enabled on the virtual server.*

Configuring a per-request policy to branch by group or class

Add a group or class lookup to a per-request policy when you want to branch by user group or class.

Note: *The access policy must be configured to populate session variables for a group or class lookup to succeed. This task provides the steps for adding items to branch by group or class. It does not specify a complete per-request policy.*

1. On the Main tab, click **Access** > **Profiles / Policies** > **Per-Request Policies**.
   The Per-Request Policies screen opens.
2. In the **Name** field, locate the policy that you want to update, then in the **Per-Request Policy** field, click the **Edit** link.
   The visual policy editor opens in another tab.
3. On a policy branch, click the **(+)** icon to add an item to the policy.

   A small set of actions are provided for building a per-request policy.

   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
4. On the Authentication tab, select an option: **AD Group Lookup**, **LDAP Group Lookup**, or **RADIUS Class Lookup** to the per-request policy.
5. Click **Add Item**.
   A properties popup screen opens.
6. Click the Branch Rules tab.
7. To edit an expression, click the **change** link.
   An additional popup screen opens, displaying the Simple tab.
8. Edit the default simple expression to specify a group or class that is used in your environment.
   In an LDAP Group Lookup item, the default simple expression is **User is a member of** `CN=MY_GROUP, CN=USERS, CN=MY_DOMAIN`. You can use the simple expression editor to replace the default values.
9. Click **Finished**.
   The popup screen closes.
10. Click **Save**.
    The popup screen closes. The visual policy editor displays.

A per-request policy goes into effect when you add it to a virtual server.

Creating a DNS resolver

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

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

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

Adding forward zones to a DNS resolver

Before you begin, gather the IP addresses of the nameservers that you want to associate with a forward zone.

Add a forward zone to a DNS resolver when you want the BIG-IP® system to forward queries for particular zones to specific nameservers for resolution in case the resolver does not contain a response to the query.

Note: *Creating a forward zone is optional. Without one, a DNS resolver can still make recursive name queries to the root DNS servers; the virtual servers using the cache must have a route to the Internet.*

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

1. On the Main tab, click **Network** > **DNS Resolvers** > **DNS Resolver List**.
   The DNS Resolver List screen opens.
2. Click the name of the resolver you want to modify.
   The properties screen opens.
3. On the menu bar, click **Forward Zones**.
   The Forward Zones screen displays.
4. Click the **Add** button.

   Note: *You add more than one zone to forward based on the needs of your organization.*

5. In the **Name** field, type the name of a subdomain or type the fully qualified domain name (FQDN) of a forward zone.

   Note: *To forward all requests (such as when creating an OAuth server), specify . (period) as the name.*

   For example, either `example` or `site.example.com` would be valid zone names.
6. Add one or more nameservers:
   a) In the **Address** field, type the IP address of a DNS nameserver that is considered authoritative for this zone.

      Based on your network configuration, add IPv4 or IPv6 addresses, or both.
   b) Click **Add**.
      The address is added to the list.

   Note: *The order of nameservers in the configuration does not impact which nameserver the system selects to forward a query to.*

7. Click **Finished**.

Adding a DNS resolver to the http-explicit profile

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

Note: *APM® provides a default* ***http-explicit*** *profile for Secure Web Gateway (SWG) explicit forward proxy. You must add a DNS resolver to the profile.*

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **HTTP**.

The HTTP profile list screen opens.

2. Click the **http-explicit** link.
   The Properties screen displays.
3. Scroll down to the Explicit Proxy area.
4. From the **DNS Resolver** list, select the DNS resolver you configured previously.
5. Ensure that you retain the default values for the **Tunnel Name** and **Default Connect Handling** fields.

   The default value for **Tunnel Name** is **http-tunnel**. The default value for **Default Connect Handling** is **Deny**.
6. Click **Finished**.

Updating the virtual server for SSL forward proxy

To add per-request processing to an LTM® SSL forward proxy configuration, associate the access profile, custom HTTP profile, and per-request policy with the virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of the virtual server that is configured for LTM SSL forward proxy.

   SSL client and server profiles that are configured specifically for SSL forward proxy are associated with this virtual server.
3. From the **HTTP Profile** list, select **http-explicit**.
4. From the **HTTP Profile** list, select the HTTP profile you configured earlier.
5. In the Access Policy area, from the **Access Profile** list, select the access profile that you configured earlier.
6. From the **Per-Request Policy** list, select the per-request policy that you configured earlier.
7. Click **Update**.

The access policy and per-request policy are now associated with the virtual server.

Example policy: SSL forward proxy bypass

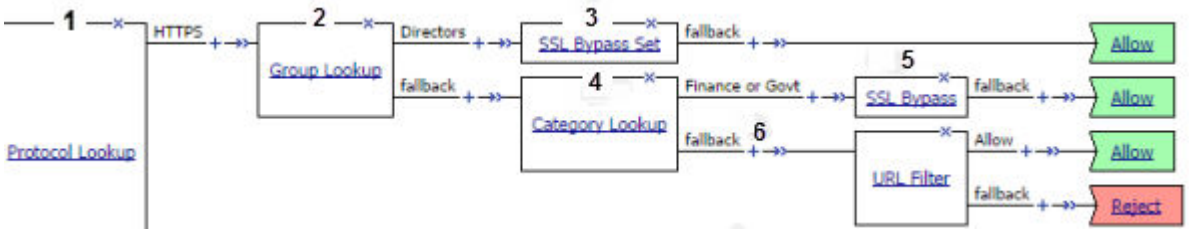

Figure 5: SSL bypass decision based on group membership and URL category

1	SSL traffic exits on the HTTPS branch of Protocol Lookup.
2	A lookup type item, such as LocalDB Group Lookup, identifies users in a group, Directors.
3	With SSL Bypass Set, any SSL request on the Directors branch is not intercepted or inspected.
4	Category Lookup processes HTTPS traffic when configured to use SNI or Subject.CN input. ***Note:*** *Finance or Govt is a standard URL category that SWG maintains on a system with an SWG subscription. User-defined URL categories can provide an alternative on systems without an SWG subscription.*

5	For users in a group other than Directors, bypass only requests that contain private information (determined through Category Lookup).
6	SSL traffic processing is complete. Now is the time to start processing HTTP data with actions that inspect the SSL payload. Using data provided by Category Lookup, URL Filter Assign item determines whether to allow or block traffic.

(For this example to be valid, both the server and client SSL profiles on the virtual server must enable SSL forward proxy and SSL forward proxy bypass; the client SSL profile must set the default bypass action to **Intercept**.)

Overview: SSL forward proxy client and server authentication

With the BIG-IP® system's *SSL forward proxy* functionality, you can encrypt all traffic between a client and the BIG-IP system, by using one certificate, and to encrypt all traffic between the BIG-IP system and the server, by using a different certificate.

A client establishes a three-way handshake and SSL connection with the wildcard IP address of the BIG-IP system virtual server. The BIG-IP system then establishes a three-way handshake and SSL connection with the server, and receives and validates a server certificate (while maintaining the separate connection with the client). The BIG-IP system uses the server certificate to create a second unique server certificate to send to the client. The client receives the second server certificate from the BIG-IP system, but recognizes the certificate as originating directly from the server.

Important: *To enable SSL forward proxy functionality, you can either:*

- Disassociate existing Client SSL and Server SSL profiles from a virtual server and configure the SSL Forward Proxy settings.
- Create new Client SSL and Server SSL profiles and configure the SSL Forward Proxy settings.

Then with either option, select the Client SSL and Server SSL profiles on a virtual server. You cannot modify existing Client SSL and Server SSL profiles while they are selected on a virtual server to enable SSL forward proxy functionality.

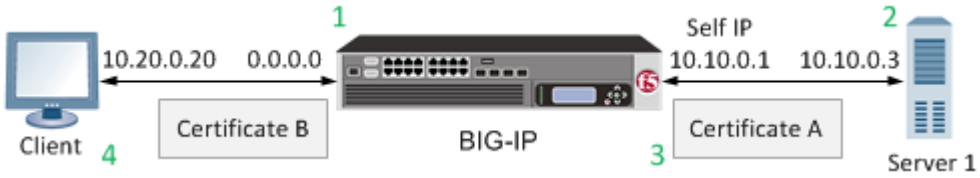

Figure 6: A virtual server configured with Client and Server SSL profiles for SSL forward proxy functionality

1. Client establishes three-way handshake and SSL connection with wildcard IP address.
2. BIG-IP system establishes three-way handshake and SSL connection with server.
3. BIG-IP system validates a server certificate (Certificate A), while maintaining the separate connection with the client.
4. BIG-IP system creates different server certificate (Certificate B) and sends it to client.

Task summary

To implement SSL forward proxy client-to-server authentication, as well as application data manipulation, you perform a few basic configuration tasks. Note that you must create both a Client SSL and a Server SSL profile, and enable the SSL Forward Proxy feature in both profiles.

Task list

Creating a custom Client SSL forward proxy profile

You perform this task to create a Client SSL forward proxy profile that makes it possible for client and server authentication while still allowing the BIG-IP® system to perform data optimization, such as decryption and encryption. This profile applies to client-side SSL forward proxy traffic only.

1. On the Main tab, click **Local Traffic** > **Profiles** > **SSL** > **Client**.
   The Client SSL profile list screen opens.
2. Click **Create**.
   The New Client SSL Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. From the **Parent Profile** list, select **clientssl**.
5. From the **SSL Forward Proxy** list, select **Advanced**.
6. Select the **Custom** check box for the SSL Forward Proxy area.
7. Modify the SSL Forward Proxy settings.
   a) From the **SSL Forward Proxy** list, select **Enabled**.
   b) From the **CA Certificate** list, select a certificate.

   Important: *If the BIG-IP system is part of a DSC Sync-Failover group, always select a non-default certificate name, and ensure that this same certificate name is specified in every instance of this SSL profile in the device group. Taking these actions helps to ensure that SSL handshakes are successful after a failover event.*

   c) From the **CA Key** list, select a key.

   Important: *If the BIG-IP system is part of a DSC Sync-Failover group, always select a non-default key name, and ensure that this same key name is specified in every instance of this SSL profile in the device group. Taking these actions helps to ensure that SSL handshakes are successful after a failover event.*

   d) In the **CA Passphrase** field, type a passphrase.
   e) In the **Confirm CA Passphrase** field, type the passphrase again.
   f) In the **Certificate Lifespan** field, type a lifespan for the SSL forward proxy certificate in days.
   g) (Optional) From the **Certificate Extensions** list, select **Extensions List**.
   h) (Optional) For the **Certificate Extensions List** setting, select the extensions that you want in the **Available extensions** field, and move them to the **Enabled Extensions** field using the **Enable** button.
   i) Select the **Cache Certificate by Addr-Port** check box if you want to cache certificates by IP address and port number.
   j) From the **SSL Forward Proxy Bypass** list, select **Enabled**.
   Additional settings display.
   k) From the **Bypass Default Action** list, select **Intercept** or **Bypass**.

   The default action applies to addresses and hostnames that do not match any entry specified in the lists that you specify. The system matches traffic first against destination IP address lists, then source IP address lists, and lastly, hostname lists. Within these, the default action also specifies whether to search the intercept list or the bypass list first.

Note: *If you select* ***Bypass*** *and do not specify any additional settings, you introduce a security risk to your system.*

8. Click **Finished**.

The custom Client SSL forward proxy profile now appears in the Client SSL profile list screen.

Creating a custom Server SSL forward proxy profile

You perform this task to create a Server SSL forward proxy profile that makes it possible for client and server authentication while still allowing the BIG-IP® system to perform data optimization, such as decryption and encryption. This profile applies to server-side SSL forward proxy traffic only.

1. On the Main tab, click **Local Traffic** > **Profiles** > **SSL** > **Server**.
   The Server SSL profile list screen opens.
2. Click **Create**.
   The New Server SSL Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. From the **Parent Profile** list select **serverssl**.
5. Select the **Custom** check box for the Configuration area.
6. From the **SSL Forward Proxy** list, select **Enabled**.
7. Click **Finished**.

The custom Server SSL forward proxy profile now appears in the Server SSL profile list screen.

Creating a load balancing pool

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

Note: *You must create the pool before you create the corresponding virtual server.*

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, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   Tip: *Hold the Shift or Ctrl key to select more than one monitor at a time.*
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. 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**.

8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server for client-side and server-side SSL traffic

You can specify a virtual server to be either a host virtual server or a network virtual server to manage application traffic.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For a network, in the **Destination Address** field, type an IPv4 or IPv6 address in CIDR format to allow all traffic to be translated.

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `0.0.0.0/0`, and an IPv6 address/prefix is `::/0`.
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. For the **SSL Profile (Client)** setting, from the **Available** list, select the name of the Client SSL forward proxy profile you previously created, and using the Move button, move the name to the **Selected** list.

   Important: *To enable SSL forward proxy functionality, you can either:*

   - Disassociate existing Client SSL and Server SSL profiles from a virtual server and configure the SSL Forward Proxy settings.
   - Create new Client SSL and Server SSL profiles and configure the SSL Forward Proxy settings.

   Then with either option, select the Client SSL and Server SSL profiles on a virtual server. You cannot modify existing Client SSL and Server SSL profiles while they are selected on a virtual server to enable SSL forward proxy functionality.

   ---
7. For the **SSL Profile (Server)** setting, from the **Available** list, select the name of the Server SSL forward proxy profile you previously created, and using the Move button, move the name to the **Selected** list.

   Important: *To enable SSL forward proxy functionality, you can either:*

   - Disassociate existing Client SSL and Server SSL profiles from a virtual server and configure the SSL Forward Proxy settings.
   - Create new Client SSL and Server SSL profiles and configure the SSL Forward Proxy settings.

   Then with either option, select the Client SSL and Server SSL profiles on a virtual server. You cannot modify existing Client SSL and Server SSL profiles while they are selected on a virtual server to enable SSL forward proxy functionality.

   ---
8. Assign other profiles to the virtual server if applicable.
9. In the Resources area, from the **Default Pool** list, select the name of the pool that you created previously.
10. Click **Finished**.

The virtual server now appears in the Virtual Server List screen.

Implementation result

After you complete the tasks in this implementation, the BIG-IP® system ensures that the client system and server system can authenticate each other independently. After client and server authentication, the BIG-IP system can intelligently decrypt and manipulate the application data according to the configuration settings in the profiles assigned to the virtual server.

Custom URL Categorization

How can I control traffic to URL categories?

A custom URL category enables you to group URLs to distinguish different types of web traffic and allow you to control it. Having custom URL categories available enables you to look up the category on a per-request basis. You can configure the per-request policy to specify whether anyone can access a URL category and when.

Example policy: User-defined category-specific access control

In this per-request policy example, only recruiters are allowed to access URLs in the user-defined category Employment. The policy also restricts access to entertaining videos during business hours.

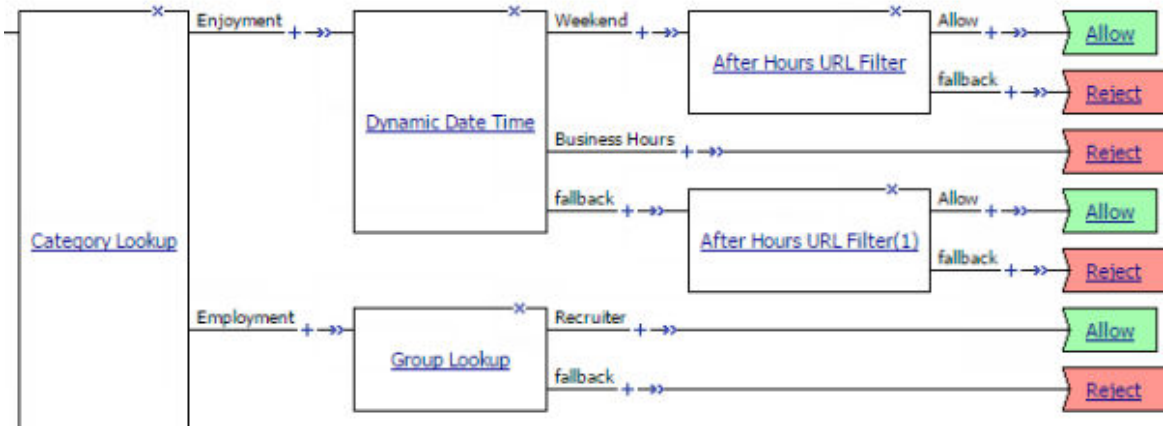

Figure 7: Category-specific access restrictions (using user-defined categories)

How can I block access to URLs?

If you have custom URL categories configured, you can also configure URL filters. A URL filter specifies an action (block, allow, or confirm) to take for each custom URL category. Having URL categories and filters available enables you to look up and filter URLs on a per-request basis.

Example policy: URL filter per user group

Each URL Filter Assign item in this per-request policy example should specify a filter that is applicable to the user group.

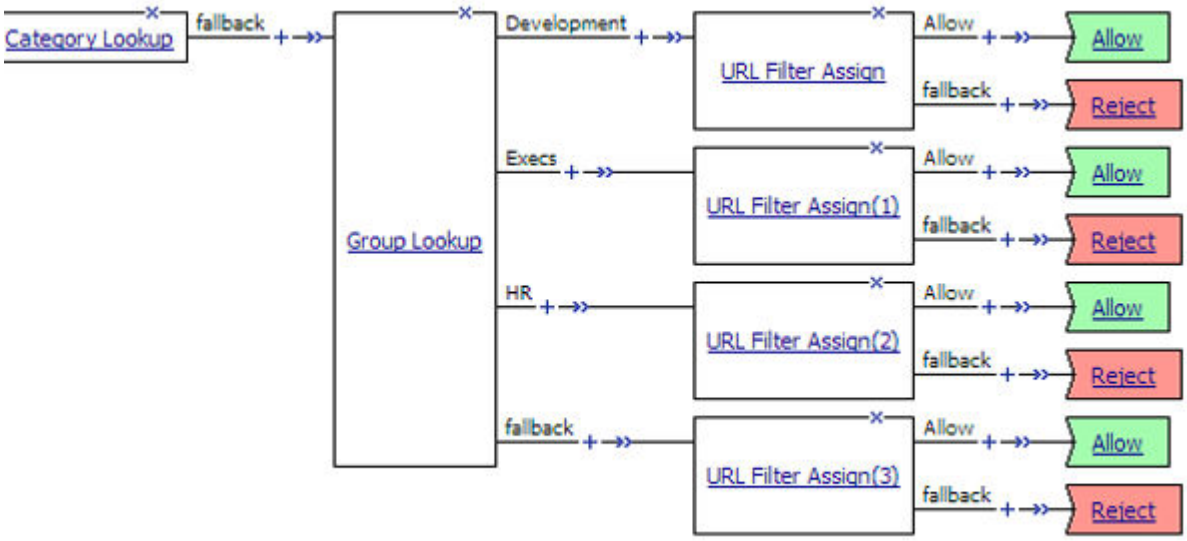

Figure 8: URL filter based on group membership

Overview: Configuring user-defined URL categories and filters

On a BIG-IP® system without a URL database, if you want to control traffic based on the type of URL being requested, and you have many URLs to consider, you should configure user-defined URL categories and user-defined URL filters. This approach provides good performance, ease-of-use, and the ability to use the **URL Category** and the **URL Filter Assign** agents in a per-request policy.

If you have only a few URLs that you want to treat differently, you can probably skip creating user-defined URL categories and filters and use a simple **URL Branching** agent in a per-request policy. In this case, you specify the URLs that you want to match directly in the **URL Branching** agent.

To configure user-defined URL categories and URL filters, complete these tasks.

Task summary

Configuring user-defined URL categories
Configuring URL filters

Configuring user-defined URL categories

Configure a user-defined URL category to specify a group of URLs over which you want to control access.

1. On the Main tab, click **Access Policy** > **Secure Web Gateway** > **URL Categories**.
   The URL Categories table displays. If you have not created any URL categories, the table is empty.
2. Click **Create**.
   The Category Properties screen displays.
3. In the **Name** field, type a unique name for the URL category.
4. From the **Default Action** list, retain the default value **Block**; or, select **Allow**.

Note: *A Confirm Box action in a per-request policy subroutine serves the purpose of enabling appropriate choices in a forward proxy (outbound) configuration. Currently, Access Policy Manager® does not support a similar action for reverse proxy.*

5. Add, edit, or delete the URLs that are associated with the category by updating the **Associated URLs** list.
6. To add URLs to the **Associated URLs** list:
   a) In the **URL** field, type a URL.
      You can type a well-formed URL that the system must match exactly or type a URL that includes globbing patterns (wildcards) for the system to match URLs.
   b) If you typed globbing patterns in the **URL** field, select the **Glob Pattern Match** check box .
   c) Click **Add**.
      The URL displays in the **Associated URLs** list.

   These are well-formed URLs:

   - `https://www.siterequest.com/`
   - `http://www.siterequest.com:8080/`
   - `http://www.sitequest.com/docs/siterequest.pdf/`
   - `http://www.sitequest.com/products/application-guides/`

   This URL `*siterequest.[!comru]` includes globbing patterns that match any URL that includes `siterequest`, except for `siterequest.com` or `siterequest.ru`.
   This URL `*://siterequest.com/education/*` includes globbing patterns that match any HTTP URL that includes `siterequest.com/education`, but that do not match any HTTPS URLs if Category Lookup specifies that the input is SNI or CN.Subject.

Important: *For SNI or CN.Subject input, Category Lookup uses `scheme://host` for matching, instead of matching the whole URL.*

7. Click **Finished**.
   The URL Categories screen displays.
8. To view the newly created URL category, expand **Custom Categories**.
   The custom URL category displays in the Sub-Category column.

Add or edit a URL filter to specify an action (allow, block, or confirm) for the custom category.

Configuring URL filters

You configure a URL filter to specify whether to allow or block requests for URLs in URL categories. You can configure multiple URL filters.

1. On the Main tab, click **Access Policy** > **Secure Web Gateway** > **URL Filters**.
   You can click the name of any filter to view its settings.

Note: *On a BIG-IP® system with an SWG subscription, default URL filters, such as **block-all** and **basic-security**, are available. You cannot delete default URL filters.*

   The URL Filters screen displays.
2. To configure a new URL filter, click one of these options.
   - **Create** button: Click to start with a URL filter that allows all categories.
   - **Copy** link: Click for an existing URL filter in the table to start with its settings.
3. In the **Name** field, type a unique name for the URL filter.
4. Click **Finished**.

The screen redisplays. An Associated Categories table displays. It includes each URL category and the filtering action that is currently assigned to it. The table includes a Sub-Category column. Any URL categories that were added by administrators are subcategories within **Custom Categories**

5. Select the actions to take:
   a) To block access to particular categories or subcategories, select them and click **Block**.

   Important: *When you select a category, you also select the related subcategories. You can expand the category and clear any subcategory selections.*

   b) To allow access to particular categories or subcategories, select them and click **Allow**.

   The confirm action is not fully supported in a reverse proxy configuration.

   Note: *A Confirm Box action in a per-request policy subroutine serves the purpose of enabling appropriate choices in a forward proxy (outbound) configuration. Currently, Access Policy Manager® does not support a similar action for reverse proxy.*

To put a URL filter into effect, you must assign it in a per-request policy. A per-request policy runs each time a URL request is made.

Application Filter Configuration

About application families

Access Policy Manager® (APM®) supports a predefined set of application families and applications. An *application family* name characterizes the type of applications associated with it. Users cannot add applications or application families to APM.

About application filters

An *application filter* specifies the applications (and application families) that Access Policy Manager® (APM®) supports and a filtering action (allow or block) for each application. An application filter can be used in a per-request policy in a supported APM configuration to control access to supported applications.

APM provides predefined application filters: block-all, allow-all, and default. The default application filter allows access to some application families and blocks access to others. Users can define their own application filters and use those that APM provides.

Overview: Configuring filters for application access

Access Policy Manager® (APM®) provides a few default application filters and you can configure additional filters. Application filtering is effected in a per-request policy.

Task summary

Specifying the default filter action for an application
Configuring application filters

Specifying the default filter action for an application

You can change the default filter action (block or allow) for any application. When you create a new application filter, the applications in it specify the default filter action.

Note: *A change to the default filter action for an application has no effect on existing application filters.*

1. On the Main tab, click **Access Policy** > **Secure Web Gateway** > **Applications**.
   The Applications screen displays.
2. To view applications, expand an application family.
3. To modify the default filter action for an application:
   a) Click the application name.
      An Application Properties screen displays.
   b) From the **Default Filter Action** list, retain the displayed setting or select another.
      The options are **Block** and **Allow**.
   c) Click **Update**.
      The Applications screen displays.

The default filtering action for the application is updated and is used when a new application filter is created.

Configuring application filters

Configure an application filter to specify how to process requests for access to applications or application families. You can configure multiple application filters.

1. On the Main tab, click **Access Policy** > **Secure Web Gateway** > **Application Filters**.
   Click the name of any filter to view its settings.

   Note: *Default application filters, such as block-all, allow-all and default, are available. You cannot delete default application filters.*

   The Application Filters screen displays.
2. To configure a new application filter, click one of these:
   - **Create** button - Click to start with an application filter with the default filter action specified for each application.
   - **Copy** link - Click this link for an existing application filter in the table to start with its settings.

   Another screen opens.
3. In the **Name** field, type a unique name for the application filter.
4. In the **Description** field, type any descriptive text.
5. Click **Finished**.
   The properties screen displays with an Associated Applications table.
6. To block access to particular applications or entire application families, select them and click **Block**.

   Important: *When you select an application family, you also select the related applications. You can expand the application family and clear any applications that are selected.*

   Important: *To block any applications that Secure Web Gateway cannot categorize, select the application family* ***Unknown****.*
7. To allow access to particular applications or entire application families, select them and click **Allow**.

To use an application filter, you must assign it in a per-request policy. A per-request policy runs each time a request is made.

Configuring Dynamic ACLs

Overview: Applying ACLs from external servers

You can apply ACLs from Active Directory, RADIUS, or LDAP servers using the Dynamic ACL action from an Access Policy Manager® access policy.

Task summary

After you configure ACLs in a supported format on an Active Directory, LDAP, or RADIUS server, you can configure a dynamic ACL action to extract and use the ACLs.

Task list

Configuring a dynamic ACL container
Adding a dynamic ACL to an access policy

About Dynamic ACL

Use this action to add dynamic ACLs to a policy branch.

Note: *Access Policy Manager® supports dynamic ACLs in an F5 ACL format, and in a subset of the Cisco ACL format.*

A Dynamic ACL action provides these configuration elements and options:

Source
Specifies a type of session variable (**Custom** or **CiscoAV-PairVSA**) and the source session variable from which the dynamic ACL is derived. For **Custom** dynamic ACL entries, this is any session variable that is populated with an f5 format ACL. For **CiscoAV-PairVSA** dynamic ACL entries, this is predefined as `session.radius.last.attr.vendor-specific.1.9.1`.

ACL
Specifies the dynamic ACL container configured on the BIG-IP® system.

Format
Specifies the format (F5 or Cisco) in which the ACL is specified.

Note: *To succeed, a Dynamic ACL action must follow actions that populate the session variables with ACLs.*

Configuring a dynamic ACL container

A dynamic ACL container provides an unconfigured ACL that you select when you configure a dynamic ACL action in an access policy.

1. On the Main tab, click **Access** > **Access Control List**.
   The ACLs screen opens.
2. Click **Create**.
   The New ACL screen opens.
3. In the **Name** field, type a name for the access control list.

4. From the **Type** list, select **Dynamic**.
5. (Optional) In the **Description** field, add a description of the access control list.
6. (Optional) From the **ACL Order** list, specify the order in which to add the new ACL relative to other ACLs:
   - Select **After** to add the ACL after a specific ACL and select the ACL from the list.
   - Select **Specify** to type the specific number of the ACL in the field.
   - Select **Last** to add the ACL at the last position in the list.
7. From the **Match Case for Paths** list, select **Yes** to match case for paths, or **No** to ignore path case.

   This setting specifies whether alphabetic case is considered when matching paths in an access control entry.
8. Click the **Create** button.
   The ACL Properties screen opens; it displays the newly configured dynamic ACL container.

Adding a dynamic ACL to an access policy

Before you start this task, configure an access profile and a dynamic ACL container. Add an authentication action to the access policy before the dynamic ACL action so that Access Policy Manager® can first capture authentication variables that contain the dynamic ACL specification.

Configure a dynamic ACL action to extract and apply an ACL from an AAA server (Active Directory, LDAP, or RADIUS).

Note: *Because a dynamic ACL is associated with a user directory, you can use one to assign ACLs specifically per the user session.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.
3. Click the **(+)** icon anywhere in the access policy to add a new item.

   Note: *Only an applicable subset of access policy items is available for selection in the visual policy editor for any access profile type.*

   A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.
4. From the Assignment tab, select **Dynamic ACL**, and click **Add Item**.
   A properties screen opens.
5. To add an ACL, click the **Add new entry** button.
   A new row opens in the table.
6. Select one of these from the list:
   - **Custom** Select to use an F5 ACL from an AD, RADIUS, or LDAP directory.
   - **Cisco AV-Pair VSA** Select to use a Cisco AV-Pair ACL from a RADIUS directory.
7. In the **Source** field, type the attribute from which the Dynamic ACL action extracts ACLs.

   If you are using Cisco AV-Pair VSA from a RADIUS server, the field is prepopulated with `session.radius.last.attr.vendor-specific.1.9.1`.
8. From the **ACL** list, select the dynamic ACL container that you configured previously.
9. From the **Format** list, select the format in which the ACL is specified.
10. (Optional) To configure another ACL, click the **Add new entry** button and repeat the configuration steps.

11. Select **Save** to save any changes and return to the policy.
12. Complete the policy:
    a) Add any additional policy items you require.
    b) Change the ending from **Deny** to **Allow** on any access policy branch on which you want to grant access.
13. Click the **Apply Access Policy** link to apply and activate the changes to the policy.

The access policy is configured to extract an ACL from an AAA server and apply it when processing occurs on the access policy branch.

To apply this access policy to network traffic, add the access profile to a virtual server.

Note: *To ensure that logging is configured to meet your requirements, verify the log settings for the access profile.*

F5 ACL format

Specifies F5® ACL syntax and provides examples. This syntax applies to both static and dynamic ACLs.

Specify an F5 ACL using this syntax.

```
comment { action [logging_options] context } comment { action
[logging_options] context }...
```

The syntax allows multiple ACLs in a single string along with comments.

comment

Any characters before an open curly brace ({) or after a closed curly brace (}) are treated as comments. Comments are optional. They have no effect on the ACLs. These examples show identical ACLs with different comments; APM® interprets them as being the same.

String comments

```
This is my HTTP server ACL { allow tcp any 1.2.3.4:80 } This is my default
ACL { reject ip any any }
```

A space as a comment

```
{ allow tcp any 1.2.3.4:80 } { reject ip any any }
```

Newline comments

```
{ allow tcp any 1.2.3.4:80 }\n
{ reject ip any any }\n
```

Vertical bar comments

```
| { allow tcp any 1.2.3.4:80 } | { reject ip any any } |
```

action

This is an action that the ACL takes on traffic that matches the ACL context.

`allow` Allows the specified traffic.

`reject` Rejects the specified traffic and sends a TCP RST code to the initiator.

`discard` Silently drops the packets.

`continue` Skips checking against the remaining access control entries in this ACL, and continues evaluation at the next ACL.

logging_options

Specifying a logging option is optional.

`log` Enables default logging for the ACL

`log-packet` Writes packet-level logs to the packet filter log file

`log-verbose` Writes verbose logs

`log-summary` Writes summary logs

`log-config` Writes configuration logs to the configuration log file

context

Context specifies a protocol followed by addresses, networks, and ports for the ACL action.

`http` HTTP protocol traffic. Requires that you specify an HTTP or HTTPS URL in the ACL definition

`udp` UDP traffic only

`tcp` TCP traffic only

`ip` IP protocol traffic

Note: *F5 ACL format treats IP protocol number zero (0) as a wildcard, meaning that it applies to all IPv4 and IPv6 traffic.*

For example, `{ reject ip 0 any any }` is the equivalent of `{ reject ip any any }`.

Address, network, and port specification

Specify addresses in a pair separated by a space. The first address in the pair should match the host, and the second address in the pair should match the destination. This syntax:

```
any[/mask][:port]
```

matches any host or IP address with an optional subnet mask or a port. For example,

```
{ allow tcp any 1.2.3.4 }
```

allows TCP traffic between any host and the destination IP address 1.2.3.4.

```
{ allow tcp any/8 1.2.3.4 }
```

allows TCP traffic between any host within the subnet 255.0.0.0 and the destination IP address 1.2.3.4.

```
{ allow tcp any/8:8000 1.2.3.4 }
```

allows TCP traffic between any host within the subnet 255.0.0.0 on port 8000 and the destination IP address 1.2.3.4.

This syntax:

```
IP address[/mask][:port]
```

matches a specific IP address with an optional subnet mask or a port. For example,

```
{ allow tcp 1.1.1.1 1.2.3.4 }
```

allows TCP traffic between the host IP address 1.1.1.1 and the destination IP address 1.2.3.4.

```
{ allow tcp 1.1.1.1:22 1.2.3.4 }
```

allows TCP traffic between the host IP address 1.1.1.1 on port 22 and the destination IP address 1.2.3.4.

F5 ACL with the IP protocol

This example shows how to specify an IP protocol address in F5 ACL format. An IP protocol number, 51, and an address pair specification follow the context word `ip`.

```
{ allow ip 51 any 1.2.3.4 }
```

F5 ACL with the TCP or UDP protocol

This example shows how to specify a TCP or UDP protocol address in F5 ACL format. An address pair specification follows the context word (`tcp` or `udp`).

```
{ allow tcp any 1.2.3.4 }
{ allow udp any 1.2.3.4 }
```

F5 ACL with the HTTP protocol

These examples show how to specify an HTTP protocol address in F5 ACL format. A host address, destination address, and URL follow the context word `http`. The URL specification supports wildcards with glob matching.

```
{ allow http any 1.2.3.4 https://www.siterequest.com/* }
{ allow http any 1.2.3.0/24 http://*.siterequest.com/* }
{ allow http any 1.2.3.0/24 http://*.siterequest.???/* }
```

Cisco ACL format

Specifies the subset of Cisco ACL syntax that Access Policy Manager® supports and provides examples.

Usage

On a RADIUS server, Access Policy Manager supports dynamic ACLs that use the subset of the Cisco ACL format described here.

Prefix

You must specify a prefix. For IPv4, use `ip:inacl#`*X*= where *X* is an integer used as a rule identifier. For IPv6, use `ipv6:inacl#`*X*=.

Keywords

These keywords are mapped with the F5 log-packet format: `log` and `log-input`.

These keywords are not supported: `tos`, `established`, `time-range`, `dynamic`, and `precedence`.

Supported specification for Cisco ACL for IP protocol

```
{ip|ipv6}:inacl#X={deny|permit}
      {ip|ipv6} source source-wildcard destination destination-wildcard
      [log|log-input]
```

For example:

```
ipv6:inacl#10=permit ipv6 any any log
```

Supported specification for Cisco ACL for TCP protocol

```
{ip|ipv6}:inacl#X={deny|permit}
      tcp source source-wildcard [operator [port]] destination destination-wildcard
     [operator [port]] [log|log-input]
```

For example:

```
ip:inacl#10=permit tcp any host 10.168.12.100 log
```

Supported specification for Cisco ACL for UDP protocol

```
{ip|ipv6}:inacl#X={deny|permit}
      udp source source-wildcard [operator [port]] destination destination-wildcard
     [operator [port]] [log|log-input]
```

For example:

```
ip:inacl#2=deny udp any any log
```

Configuring Routing for Access Policies

Overview: Selecting a route domain for a session (example)

A *route domain* is a BIG-IP® system object that represents a particular network configuration. Route domains provide the capability to segment network traffic, and define separate routing paths for different network objects and applications. You can create an access policy that assigns users to different route domains using the Route Domain and SNAT Selection action based on whatever criteria you determine appropriate.

You might use policy routing in a situation such as this: your company has switched from RADIUS authentication to Active Directory authentication, but has not yet completed the full transition. Because of the state of the authentication changeover, you would like your legacy RADIUS users to pass through to a portal access connection on a separate router, instead of allowing full access to your network.

This implementation provides configuration steps for this example.

Task summary

Creating a route domain on the BIG-IP system
Creating an access profile
Verifying log settings for the access profile
Configuring policy routing

Creating a route domain on the BIG-IP system

Before you create a route domain:

- Ensure that an external and an internal VLAN exist on the BIG-IP® system.
- Verify that you have set the current partition on the system to the partition in which you want the route domain to reside.

You can create a route domain on BIG-IP system to segment (isolate) traffic on your network. Route domains are useful for multi-tenant configurations.

1. On the Main tab, click **Network** > **Route Domains**.
   The Route Domain List screen opens.
2. Click **Create**.
   The New Route Domain screen opens.
3. In the **Name** field, type a name for the route domain.
   This name must be unique within the administrative partition in which the route domain resides.
4. In the **ID** field, type an ID number for the route domain.
   This ID must be unique on the BIG-IP system; that is, no other route domain on the system can have this ID.
   An example of a route domain ID is `1`.
5. For the **Parent Name** setting, retain the default value.
6. For the **VLANs** setting, from the **Available** list, select a VLAN name and move it to the **Members** list.
   Select the VLAN that processes the application traffic relevant to this route domain.

Configuring this setting ensures that the BIG-IP system immediately associates any self IP addresses pertaining to the selected VLANs with this route domain.

7. Click **Finished**.
   The system displays a list of route domains on the BIG-IP system.

You now have another route domain on the BIG-IP system.

Creating an access profile

You create an access profile to provide the access policy configuration for a virtual server that establishes a secured session.

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click **Create**.
   The New Profile screen opens.
3. In the **Name** field, type a name for the access profile.

   Note: *A access profile name must be unique among all access profile and any per-request policy names.*

4. From the **Profile Type** list, select one these options:
   - **LTM-APM**: Select for a web access management configuration.
   - **SSL-VPN**: Select to configure network access, portal access, or application access. (Most access policy items are available for this type.)
   - **ALL**: Select to support LTM-APM and SSL-VPN access types.
   - **SSO**: Select to configure matching virtual servers for Single Sign-On (SSO).

     Note: *No access policy is associated with this type of access profile*

   - **RDG-RAP**: Select to validate connections to hosts behind APM when APM acts as a gateway for RDP clients.
   - **SWG - Explicit**: Select to configure access using Secure Web Gateway explicit forward proxy.
   - **SWG - Transparent**: Select to configure access using Secure Web Gateway transparent forward proxy.
   - **System Authentication**: Select to configure administrator access to the BIG-IP® system (when using APM as a pluggable authentication module).
   - **Identity Service**: Used internally to provide identity service for a supported integration. Only APM creates this type of profile.

     Note: *You can edit Identity Service profile properties.*

   Note: *Depending on licensing, you might not see all of these profile types.*

   Additional settings display.
5. In the Language Settings area, add and remove accepted languages, and set the default language.
   A browser uses the highest priority accepted language. If no browser language matches the accepted languages list, the browser uses the default language.
6. Click **Finished**.

The access profile displays in the Access Profiles List. Default-log-setting is assigned to the access profile.

Verifying log settings for the access profile

Confirm that the correct log settings are selected for the access profile to ensure that events are logged as you intend.

Note: *Log settings are configured in the* ***Access*** > ***Overview*** > ***Event Log*** > ***Settings*** *area of the product. They enable and disable logging for access system and URL request filtering events. Log settings also specify log publishers that send log messages to specified destinations.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. On the menu bar, click **Logs**.
   The access profile log settings display.
4. Move log settings between the **Available** and **Selected** lists.

   You can assign up to three log settings that enable access system logging to an access profile. You can assign additional log settings to an access profile provided that they enable logging for URl request logging only.

   Note: *Logging is disabled when the* ***Selected*** *list is empty.*

   ---
5. Click **Update**.

An access profile is in effect when it is assigned to a virtual server.

Configuring policy routing

To follow the steps in this example, you must have Access Policy Manager® AAA server objects created for Active Directory and RADIUS as well.

You configure an access policy similar to this one to route users depending on whether they pass Active Directory authentication or RADIUS authentication. This example illustrates one way to handle a company-wide transition between one type of authentication and another, and to ensure that users get access to the correct resources, however they authenticate.

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile for which you want to edit the access policy.
   The properties screen opens for the profile you want to edit.
3. On the menu bar, click **Access Policy**.
4. In the General Properties area, click the **Edit Access Policy for Profile** ***profile_name*** link.
   The visual policy editor opens the access policy in a separate screen.
5. On a policy branch, click the **(+)** icon to add an item to the policy.
   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
6. On the Logon tab, select **Logon Page** and click the **Add Item** button.
   The Logon Page Agent properties screen opens.
7. Make any changes that you require to the logon page properties and click **Save**.
   The properties screen closes and the policy displays.
8. On the fallback branch after the previous action, click the **(+)** icon to add an item to the policy.
   A popup screen opens.

9. On the Authentication tab, select **AD Auth**.
   A properties screen displays.
10. From the **Server** list, select a server.
11. Click **Save**.
    The properties screen closes and the policy displays.
12. On the Successful branch after the previous action, click the **(+)** icon.
    A popup screen opens.
13. Assign resources to the users that successfully authenticated with Active Directory.
    a) On the Assignment tab, select the **Advanced Resource Assign** agent, and click **Add Item**.
       The Resource Assignment window opens.
    b) Click **Add new entry**.
       An **Empty** entry displays.
    c) Click the **Add/Delete** link below the entry.
       The screen changes to display resources on multiple tabs.
    d) On the Network Access tab, select a network access resource.
    e) (Optional) Optionally, on the Webtop tab, select a network access webtop.
    f) Click **Update**.
       The popup screen closes.
    g) Click **Save**.
       The properties screen closes and the policy displays.
    h) Click the ending that follows the Advanced Resource Assign action and change it to an allow ending, by selecting **Allow** and clicking **Save**.
14. On the fallback branch after the Active Directory action, click the **(+)** icon to add an item to the access policy.

    In this case, fallback indicates failure. For users that did not pass Active Directory authentication, you can configure RADIUS authentication and select a route domain for them so that they go to a different gateway.

    A popup screen opens.
15. Type `radi` in the search field, select **RADIUS Auth** from the results, and click **Add Item**.
    A popup screen opens.
16. From the **AAA Server** list, select a RADIUS server and click **Save**.
    The popup screen closes and the visual policy editor displays.
17. On the Successful branch after the previous action, click the **(+)** icon.
    A popup screen opens.
18. On the Assignment tab, select **Route Domain and SNAT Selection** and click the **Add Item** button.
    This opens the popup screen for the action.
19. From the Route Domain list, select a route domain and click **Save**.
    The popup screen closes and the visual policy editor displays.
20. On the successful branch after the route domain selection action, click the **(+)** icon.
    A popup screen opens.
21. Assign resources to the users that successfully authenticated with RADIUS.
    a) On the Assignment tab, select the **Advanced Resource Assign** agent, and click **Add Item**.
       The Resource Assignment window opens.
    b) Click **Add new entry**.
       An **Empty** entry displays.
    c) Click the **Add/Delete** link below the entry.
       The screen changes to display resources on multiple tabs.
    d) On the Network Access tab, select a network access resource.

Note that you can assign the same network access resource to clients whether they authenticate with Active Directory or RADIUS. You assigned a different route domain to the clients that successfully authenticated with RADIUS. As a result, both types of clients will reach separate routers.

e) (Optional) Optionally, on the Webtop tab, select a network access webtop.
f) Click **Update**.
The popup screen closes.
g) Click **Save**.
The properties screen closes and the policy displays.
h) Click the ending that follows the Advanced Resource Assign action and change it to an allow ending, by selecting **Allow** and clicking **Save**.

22. Click the **Apply Access Policy** link to apply and activate the changes to the policy.

To apply this access policy to network traffic, add the access profile to a virtual server.

Note: *To ensure that logging is configured to meet your requirements, verify the log settings for the access profile.*

Synchronizing Access Policies

Overview: Syncing access policies with a Sync-Only device group

Syncing access policies from one BIG-IP® Access Policy Manager® (APM®) device to another Access Policy Manager device, or to multiple devices in a device group, allows you to maintain up-to-date access policies on multiple APM devices, while adjusting appropriate settings for objects that are specific to device locations.

To synchronize access policies between multiple devices, you configure a Sync-Only device group that includes the devices between which you want to synchronize access policies.

Note: *To add devices to a device group, the devices must all belong to the same trust domain.*

For policy sync to work seamlessly, the Sync-Only device group configuration must specify an automatic type of sync. However, the process to perform a policy sync remains manual. The process involves selecting an access policy, running a sync, and resolving conflicts as needed.

Task summary

Establishing device trust
Configuring a Sync-Only device group for access policy sync
Synchronizing a policy across devices initially
Configuring static resources with policy sync
Configuring dynamic resources with policy sync
Resolving policy sync conflicts

Understanding policy sync device group setup for Active-Standby pairs

To add devices to a device group, all devices must belong to the same local trust domain. If you want to sync access policies with a device that does not belong to the local trust domain, but also belongs to a Sync-Failover group, you must reset the trust between the devices and remove them from the Sync-Failover device group. (For more information, see *BIG-IP® Device Service Clustering: Administration* on the AskF5™ web site located at `http://support.f5.com/`.)

After you establish device trust between your BIG-IP system and the devices, you can add them to a Sync-Failover group again.

Understanding policy sync for Active-Standby pairs

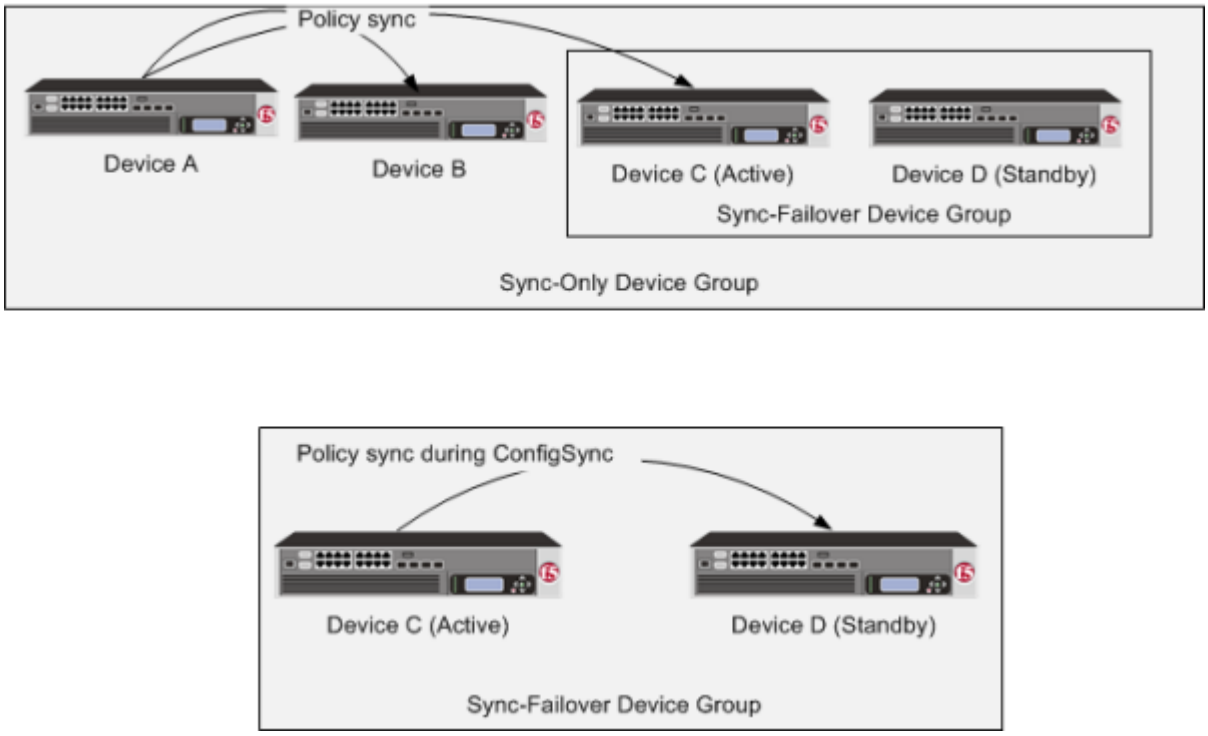

Figure 9: Access policy synchronization in Sync-Only and Sync-Failover device groups

Before you configure device trust

Before you configure device trust, you should consider the following:

- Only version 11.x or later systems can join the local trust domain.
- You can manage device trust when logged in to a certificate signing authority only. You cannot manage device trust when logged in to a subordinate non-authority device.
- If you reset trust authority on a certificate signing authority by retaining the authority of the device, you must subsequently recreate the local trust domain and the device group.
- As a best practice, you should configure the ConfigSync and mirroring addresses on a device before you add that device to the trust domain.
- You must configure DNS on all systems.
- You must configure NTP on all systems, preferably to the same NTP server.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3`

each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management** > **Device Trust** > **Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
   - If the BIG-IP device is an appliance, type the management IP address for the device.
   - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
   - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
   - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Configuring a Sync-Only device group for access policy sync

You configure a device group with specific settings for use in synchronizing access policies across devices. You can perform this task on any BIG-IP® device within the local trust domain.

Important: *When you initiate the sync of an access policy to a device group, the only device groups that you can select are those configured with the settings specified in this task.*

---

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. In the **Name** field, type a name for the device group.
4. From the **Group Type** list, select **Sync-Only**.
5. For the **Members** setting, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group. Use the Move button to move the host name to the **Includes** list.

   The list shows any devices that are members of the device's local trust domain.
6. From the **Sync Type** list, select one of these options:
   - **Automatic with Incremental Sync** - F5 Networks recommends that you select this option for optimal performance.
   - **Automatic with Full Sync**
7. Click **Finished**.

You now have a device group that you can select when you initiate policy sync for an access policy.

Synchronizing a policy across devices initially

After you set up a sync-only device group for your Access Policy Manager® devices, you can sync a policy from one device to other devices in the group. You can perform a policy sync from any device in the group.

1. On the Main tab, click **Access** > **Profiles / Policies** > **Policy Sync**.
   A list of policies and related sync status information opens. The sync status is either:

   Policies with no sync pending
   No synchronization is currently in progress for policies on this list.

   Policies with sync pending
   A synchronization is in progress for these policies. Select a policy from this list to view the Sync Details or Resolve Conflicts panel for it.
2. Select a policy and click the **Sync Policy** button.
   The **Policy Sync** screen opens.
3. From the **Device Group** list, select the device group to which to sync the policy.
   This list displays only Sync-Only device groups with automatic sync and full sync enabled.
4. In the **Description** field, type a description of the reason for the policy sync operation.
5. From the **Ignore errors due to Variable Assign Agent during sync** list, select whether to ignore errors caused by syncing the variable assign agent.

   Note: *If the policy includes a Variable Assign action, errors occur when resources are missing from the target device. If you select* ***Yes****, you might need to manually configure the resources on the target device.*
6. Click **Sync**.
   The sync process begins.

The policy is synced between devices in the device group.

Important: *A policy sync operation takes 25-30 seconds, depending on the number of devices.*

Configuring static resources with policy sync

A BIG-IP® Access Policy Manager® might exist in a different physical location from another BIG-IP in the same device group, and might use different resources that are specific to that location or local network. For example, different authentication servers might exist in each location. Configure static resources to set these static resources for devices in different locations.

1. On the Main tab, click **Access** > **Profiles / Policies** > **Policy Sync**.
   If policies are present and configured for sync, a list of policies and related sync status information opens.
2. Select a policy and click the **Sync Policy** button.
   The **Policy Sync** screen opens.
3. Click the **Advanced Settings** button, then click **Static Resources**.
   The list displays a name, type, and **Location Specific** check box for each resource. You might need to configure a location-specific resource differently on a remote system. With the Location Specific check box selected, the first time a resource is synced as part of a policy, you must resolve its configuration on the remote system. Subsequent policy sync operations do not modify a previously synced location-specific resource.

Important: *Many resource types are marked as location-specific by default. If a resource is not location-specific in this configuration, clear the* ***Location Specific*** *check box.*

4. Click the **OK** button.
   The APM Policy Sync screen is displayed.
5. Click the **Sync** button.

The policy is synced between devices in the device group.

If this is the first time you sync a policy with location-specific resources, or you have added location-specific resources to the policy sync operation, you must resolve the location-specific issues on each affected target system.

Configuring dynamic resources with policy sync

When policies are configured with the Variable Assign action, some dynamically assigned resources might not be available on sync target machines. You can specify that such resources are included in a policy sync operation and will be created on the target devices.

1. On the Main tab, click **Access** > **Profiles / Policies** > **Policy Sync**.
   A list of policies and related sync status information opens.
2. Select a policy and click the **Sync Policy** button.
   The **Policy Sync** screen opens.
3. Click the **Advanced Settings** button, then click **Dynamic Resources**.
   The list displays a name, type, **Dynamic Resource**, and **Location Specific** check box for each resource.
4. Select the dynamic resources by clicking the check boxes.
5. Click the **OK** button.
   The APM Policy Sync screen is displayed.
6. Click the **Sync** button.

The policy is synced between devices in the device group.

Resolve the location-specific issues on each affected target system.

Resolving policy sync conflicts

After you sync a policy, you might need to resolve conflicts on the target devices. Conflicts occur when a policy contains new location-specific resources.

1. On a target system that requires conflicts to be resolved, on the Main tab, click **Access** > **Profiles / Policies** > **Policy Sync**.
   A list of policies and related sync status information opens.
2. From the **Policies with Sync Pending** list, select a policy for which you want to resolve conflicts.
   If conflicts exist, the Resolve Conflicts panel displays one entry and an Unresolved link for each location-specific or dynamic resource that is in conflict.
3. Click an **Unresolved** link.
   A popup window opens displaying two panes.
   - A navigation pane with one or more groups of settings. In the navigation pane, an icon indicates that data is required.
   - A data entry pane in which you can type or select values. The data entry pane displays the values from the source device, with labels for required fields asterisked (*) and filled with yellow.
4. Select a group of settings from the left pane, and type or select the required information in the right pane until you have added the required information.

You can fill in the required information only, or any other information and settings you wish to configure.

In the navigation pane, an icon indicates that required information for a group of settings is complete.

5. Click the **OK** button.
   The popup window closes. If no more **Unresolved** links remain, the **Finish** button is active.
6. After you resolve all conflicts, click the **Finish** button.

Access Policy Manager® creates the resolved policy on the device. After sync is completed on all target devices, sync status on the source device will be updated to **Sync completed**.

About ignoring errors due to the Variable Assign agent

The **Ignore errors due to Variable Assign Agent during sync** setting affects system behavior only when a Variable Assign agent is included in an access policy, and the Variable Assign agent uses resources.

Important: *The user name and password fields are not considered to be resources.*

If you set **Ignore errors due to Variable Assign Agent during sync** to **Yes**:

- If you do not select any dynamic resources, after the policy sync completes you must create all needed resources on each target system.
- If you select the appropriate dynamic resources, after the policy sync completes, you must resolve any conflicts that exist on the target systems. If you do not select all the dynamic resources that are required, you must create them on each target system.

If you set **Ignore errors due to Variable Assign Agent during sync** to **No**:

- If you do not select any dynamic resources, an error is displayed and the policy sync does not start.
- If you select the appropriate dynamic resources, after the policy sync completes, you must resolve any conflicts that exist on the target systems.

Implementation result

To summarize, you now have synchronized access policies between devices in a sync-only device group.

Understanding sync details

On the **Sync Details** tab, you can see sync status for an access policy.

Column	Description
Device	The specific device to which the access policy was synced.
Sync Status	One of the following: • `Sync initiated` - This status indicates that the sync is in progress, initiated from this device. • `Sync Completed` - This status indicates that the sync completed successfully to the specified device. • `Not available` - This status indicates that the device to which the sync was initiated was not available, or not available yet.

Column	Description
	• `Sync cancelled` - This status indicates that the sync was cancelled before it could complete to the specified device. • `User Changes Failed` - This status indicates that policy creation failed after the administrator resolved the conflicts. Sync success is set to Standby. • `Pending location specific updates` - This status indicates that the access policy on the specified device requires updates because of conflicts due to location-specific information. Resolve the conflicts to complete the sync successfully.
Status End Time	The time at which the last status entry completed on the specific device.
Sync Status Details	More information about the Sync Status for a specific device.

Understanding sync history

On the **Sync History** tab, you can see the sync history for an access policy.

Column	Description
Last sync	The last time a sync was initiated for this access policy.
Last Sync Status	The outcome of the last sync for this access policy.
Device Group	The device group to which the access policy was synced.
Description	A clickable icon that presents information about the sync operation for the device group.
Non Location Specific Objects	An access policy was created with certain resources which the sync process indicates are not location-specific, but that might in fact be location-specific on the target device. This column lists such objects, which you can then verify by checking the objects on the remote systems, and modifying if necessary.

Load balancing Access Policy Manager

Overview: Load balancing BIG-IP APM with BIG-IP DNS

After you integrate BIG-IP® DNS into a network with BIG-IP Local Traffic Manager™ (LTM®), or vice versa, the BIG-IP systems can communicate with each other. If Access Policy Manager® (APM®) is also installed on one of the BIG-IP systems with LTM, APM calculates virtual server scores and provides them to BIG-IP DNS.

The calculation is based on the number of active access sessions. APM calculates two usage scores and assigns the higher of the two to the virtual server:

- One usage score is based on the BIG-IP system licensed maximum access concurrent sessions and the sum of the current active sessions on all the access profiles configured on the system.
- The other usage score is based on the maximum concurrent user sessions configured on the access profile attached to the virtual server and the current active sessions count on the access profile.

A value of 0 indicates no capacity and a value of 100 means full capacity available on the device.

Note: *The calculations do not include connectivity session usage.*

Use a BIG-IP DNS global load-balancing pool for BIG-IP DNS to load balance APM users based on the virtual server score. BIG-IP DNS uses virtual server score in the VS Score and Quality of Service load balancing methods for global load-balancing pools.

Task summary

These tasks must already be complete before you begin.

- BIG-IP DNS and APM must be installed and configured.
- Either BIG-IP DNS must be integrated with other BIG-IP systems on a network or BIG-IP LTM® must be integrated into a network with BIG-IP DNS.
- The health monitors defined for the BIG-IP DNS and LTM servers must include bigip; otherwise, APM does not calculate virtual server scores and send them to BIG-IP DNS.

Task list

Creating a load balancing pool
Creating a wide IP for BIG-IP DNS

Creating a load balancing pool

Ensure that at least one virtual server exists in the configuration before you start to create a load balancing pool.

Create a pool of systems with Access Policy Manager® to which the system can load balance global traffic.

1. On the Main tab, click **DNS** > **GSLB** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
3. In the General Properties area, in the **Name** field, type a name for the pool.

Names must begin with a letter, and can contain only letters, numbers, and the underscore (_) character.

Important: *The pool name is limited to 63 characters.*

4. From the **Type** list, depending on the type of the system (IPv4 or IPv6), select either an **A** or **AAAA** pool type.
5. In the Configuration area, for the **Health Monitors** setting, in the **Available** list, select a monitor type, and move the monitor to the **Selected** list.

Tip: *Hold the Shift or Ctrl key to select more than one monitor at a time.*

6. In the Members area, for the **Load Balancing Method** settings, select a method that uses virtual server score:
   - VS Score - If you select this method, load balancing decisions are based on the virtual server score only.
   - Quality of Service - If you select this method, you must configure weights for up to nine measures of service, including **VS Score**. Virtual server score then factors into the load balancing decision at the weight you specify.
7. For the **Member List** setting, add virtual servers as members of this load balancing pool.

   The system evaluates the virtual servers (pool members) in the order in which they are listed. A virtual server can belong to more than one pool.
   a) Select a virtual server from the **Virtual Server** list.
   b) Click **Add**.
8. Click **Finished**.

Creating a wide IP for BIG-IP DNS

Ensure that at least one load balancing pool exists in the configuration before you start creating a wide IP.

Create a wide IP to map an FQDN to one or more pools of virtual servers that host the content of the domain.

1. On the Main tab, click **DNS** > **GSLB** > **Wide IPs**.
   The Wide IP List screen opens.
2. Click **Create**.
   The New Wide IP List screen opens.
3. In the General Properties area, in the **Name** field, type a name for the wide IP.

Tip: *You can use two different wildcard characters in the wide IP name: asterisk (*) to represent several characters and question mark (?) to represent a single character. This reduces the number of aliases you have to add to the configuration.*

4. From the **Type** list, select a record type for the wide IP.
5. In the Pools area, for the **Pool List** setting, select the pools that this wide IP uses for load balancing.

   The system evaluates the pools based on the wide IP load balancing method configured.
   a) From the **Pool** list, select a pool.

      A pool can belong to more than one wide IP.
   b) Click **Add**.
6. Click **Finished**.

Using APM as a Gateway for RDP Clients

Overview: Configuring APM as a gateway for Microsoft RDP clients

Access Policy Manager® (APM®) can act as a gateway for Microsoft RDP clients, authorizing them on initial access and authorizing access to resources that they request after that. The APM configuration includes these elements.

APM as gateway

From a configuration point of view, this is a virtual server that accepts SSL traffic from Microsoft RDP clients and is associated with an access policy that authorizes the client.

Client authorization access policy

This access policy runs when the RDP client initiates a session with the gateway (APM). Only NTLM authentication is supported. This access policy should verify that NTLM authentication is successful and must assign an additional access policy to use for resource authorization throughout the session.

Resource authorization access policy

This access policy runs when the authorized RDP client requests access to a resource. The access policy must contain logic to determine whether to allow or deny access to the target server and port.

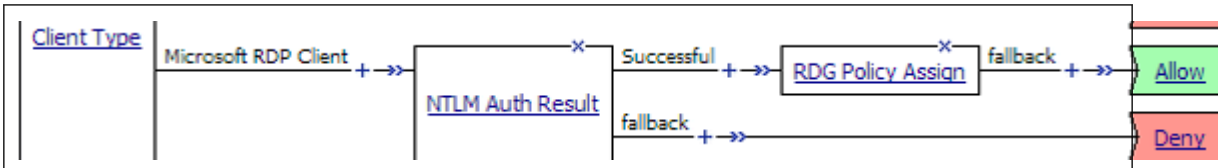

Figure 10: Sample client authorization policy

Notice the RDG Policy Assign item; it is used to specify the resource authorization policy.

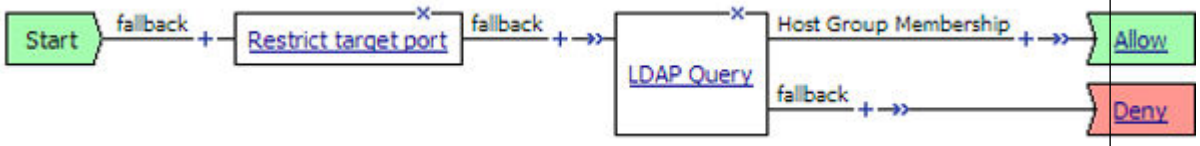

Figure 11: Sample resource authorization policy

Task summary

If you already have configured them, you can use existing configuration objects: a machine account, an NTLM authentication configuration, a VDI profile, a connectivity profile, and a client SSL profile.

Task list

Configuring an access profile for resource authorization
Verifying log settings for the access profile
Configuring an access policy for resource authorization
Creating an access profile for RDP client authorization

About supported Microsoft RDP clients

Supported Microsoft RDP clients can use APM® as a gateway. The configuration supports Microsoft RDP clients on Windows, Mac, iOS, and Android.

Refer to *BIG-IP® APM® Client Compatibility Matrix* on the AskF5™ web site at `http://support.f5.com/kb/en-us.html` for the supported platforms and operating system versions for Microsoft RDP clients.

About Microsoft RDP client login to APM

On a Microsoft RDP client, a user types in settings for a gateway and a connection. The names for the settings vary depending on the Microsoft RDP client.

RDP client gateway settings

1. Hostname setting: The hostname or IP address of the virtual server must be specified.
2. Port setting: If requested, `443` must be specified.
3. Credentials: Selection of specific logon method and entry of a user name and password should be avoided. In this implementation, APM® supports only NTLM authentication.

RDP client connection settings

Gateway setting: On some clients, you must configure a name and address for the gateway and at login type the gateway name. If requested, the gateway name must be specified as configured on the client.

1. Hostname setting: Hostname of the target server.
2. Port setting: Port on the target server.

Configuring an access profile for resource authorization

Configure an RDG-RAP type of access profile for Access Policy Manager® (APM®) before you create an access policy to authorize resource requests from Microsoft RDP clients.

Note: *After APM authorizes a Microsoft RDP client, subsequent resource requests are sent to APM.*

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click **Create**.
   The New Profile screen opens.
3. In the **Name** field, type a name for the access profile.

   Note: *A access profile name must be unique among all access profile and any per-request policy names.*

4. From the **Profile Type** list, select **RDG-RAP**.
5. Click **Finished**.
   The new access profile displays on the list.

The access profile displays in the Access Profiles List. Default-log-setting is assigned to the access profile.

You must configure an access policy that determines whether to deny or allow access to a resource.

Verifying log settings for the access profile

Confirm that the correct log settings are selected for the access profile to ensure that events are logged as you intend.

Note: *Log settings are configured in the* ***Access*** > ***Overview*** > ***Event Log*** > ***Settings*** *area of the product. They enable and disable logging for access system and URL request filtering events. Log settings also specify log publishers that send log messages to specified destinations.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. On the menu bar, click **Logs**.
   The access profile log settings display.
4. Move log settings between the **Available** and **Selected** lists.

   You can assign up to three log settings that enable access system logging to an access profile. You can assign additional log settings to an access profile provided that they enable logging for URl request logging only.

   Note: *Logging is disabled when the* ***Selected*** *list is empty.*

   ---

5. Click **Update**.

An access profile is in effect when it is assigned to a virtual server.

Configuring an access policy for resource authorization

Configure this access policy to perform resource authorization every time an RDP client requests access to a new resource.

Note: *The requested resource is specified in these session variables:* `session.rdg.target.host` *and* `session.rdg.target.port`.

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Access Policy column, click the **Edit** link for the RDG-RAP type access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.
3. Click the **(+)** icon anywhere in the access policy to add a new item.

   Note: *Only an applicable subset of access policy items is available for selection in the visual policy editor for any access profile type.*

   A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.

4. To restrict the target port to the RDP service only, perform these substeps:

   Note: *F5® strongly recommends this action.*

   a) In the search field, type **emp**, select **Empty** from the result list, and then click **Add Item**.
   A popup Properties screen opens.
   b) Click the Branch Rules tab.
   c) Click **Add Branch Rule**.
   A new entry with **Name** and **Expression** settings displays.
   d) In the **Name** field, replace the default name by typing a new name.
   The name appears on the branch in the policy.
   e) Click the **change** link in the new entry.
   A popup screen opens.
   f) Click the Advanced tab.
   g) In the field, type this expression: `expr { [mcget {session.rdg.target.port}] == 3389 }`
   h) Click **Finished**.
   The popup screen closes.
   i) Click **Save**.
   The properties screen closes and the policy displays.

5. To verify group membership for the requested host, add an **LDAP Query** to the access policy and configure properties for it:

   Adding an LDAP Query is one option. The visual policy editor provides additional items that you can use to determine whether to allow the client to access the resource.

   a) From the **Server** list, select an AAA LDAP server.
   An LDAP Query uses SSL connections when you select an LDAP AAA server that is configured for LDAPS.
   b) Type queries in the **SearchFilter** field.
   This query matches hosts with the fully qualified domain name (FQDN) of the host. `(DNSHostName=%{session.rdg.target.host})` When clients request a connection, they must specify the FQDN.
   This query matches hosts with the host name or with the FQDN of the host. `(|(name=%{session.rdg.target.host})(DNSHostName=%{session.rdg.target.host}))` When clients request a connection, they can specify a host name or an FQDN.
   c) Click **Save**.
   The properties screen closes and the policy displays.

6. To verify that the target host is a member of an Active Directory group, add a branch rule to the LDAP query item:

   a) In the visual policy editor, click the **LDAP Query** item that you want to update.
   A popup Properties screen displays.
   b) Click the Branch Rules tab, click **Add Branch Rule**, and type a descriptive name for the branch in the **Name** field.
   c) Click the **change** link in the new entry.
   A popup screen displays.
   d) Click the Advanced tab.
   e) Type an expression in the field.
   This expression matches the last LDAP memberOf attribute with an Active Directory group, *RDTestGroup*. `expr { [mcget {session.ldap.last.attr.memberOf}] contains "CN=`*RDTestGroup*`" }` The hypothetical members of the group in this example are the hosts to which access is allowed.
   f) Click **Finished**.

The popup screen closes.

g) Click **Save**.
The properties screen closes and the policy displays.

7. Click **Save**.
The properties screen closes and the policy displays.
8. Add any other items to the access policy and change any appropriate branch ending to **Allow**.
9. Click **Apply Access Policy** to save your configuration.

Important: *Do not specify this access policy in a virtual server definition. Select it from an RDG Policy Assign item in an access policy that authorizes Microsoft RDP clients.*

Creating an access profile for RDP client authorization

You create an access profile to provide the access policy configuration for a virtual server that establishes a secured session.

1. On the Main tab, click **Access** > **Profiles / Policies**.
The Access Profiles (Per-Session Policies) screen opens.
2. Click **Create**.
The New Profile screen opens.
3. In the **Name** field, type a name for the access profile.

Note: *A access profile name must be unique among all access profile and any per-request policy names.*

4. From the **Profile Type** list, select one of these options.
   - **LTM-APM**: Select for a web access management configuration.
   - **SSL-VPN**: Select to configure network access, portal access, or application access. (Most access policy items are available for this type.)
   - **ALL**: Select to support LTM-APM and SSL-VPN access types.

   Additional settings display.
5. Select the **Custom** check box.
6. In the **Access Policy Timeout** field, type the number of seconds that should pass before the access profile times out because of inactivity.

   The timeout needs to be at least 15 minutes long because an RDP client sends a keepalive to the gateway every 15 minutes.

Important: *To prevent a timeout, type* `0` *to set no timeout or type* `900` *or greater. 900 indicates a 15-minute timeout, which is enough time for the keepalive to prevent the timeout.*

7. Click **Finished**.

Verifying log settings for the access profile

Confirm that the correct log settings are selected for the access profile to ensure that events are logged as you intend.

Note: *Log settings are configured in the* ***Access*** > ***Overview*** > ***Event Log*** > ***Settings*** *area of the product. They enable and disable logging for access system and URL request filtering events. Log settings also specify log publishers that send log messages to specified destinations.*

1. On the Main tab, click **Access** > **Profiles / Policies**.

The Access Profiles (Per-Session Policies) screen opens.

2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. On the menu bar, click **Logs**.
   The access profile log settings display.
4. Move log settings between the **Available** and **Selected** lists.

   You can assign up to three log settings that enable access system logging to an access profile. You can assign additional log settings to an access profile provided that they enable logging for URl request logging only.

   Note: *Logging is disabled when the* ***Selected*** *list is empty.*
5. Click **Update**.

An access profile is in effect when it is assigned to a virtual server.

Configuring an access policy for an RDP client

Configure an access policy to authorize Microsoft RDP clients and to specify the access policy that APM® should use to authorize access to resources as the client requests them.

Note: *NTLM authentication occurs before an access policy runs. If NTLM authentication fails, an error displays and the access policy does not run.*

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.
3. Click the **(+)** icon anywhere in the access policy to add a new item.

   Note: *Only an applicable subset of access policy items is available for selection in the visual policy editor for any access profile type.*

   A popup screen opens, listing predefined actions on tabs such as General Purpose, Authentication, and so on.
4. (Optional) On the Endpoint Security (Server-Side) tab, select **Client Type**, and click **Add Item**.

   The Client Type action identifies clients and enables branching based on the client type.

   A properties screen opens.
5. Click **Save**.
   The properties screen closes; the **Client Type** item displays in the visual policy editor with a **Microsoft Client RDP** branch and branches for other client types.
6. On a policy branch, click the **(+)** icon to add an item to the policy.
7. To verify the result of client authentication:
   a) Type `NTLM` in the search field.
   b) Select **NTLM Auth Result**.
   c) Click **Add Item**.

   A properties screen opens.
8. Click **Save**.
   The properties screen closes and the policy displays.
9. Select the RDG-RAP access policy you configured earlier:
   a) Click the **[+]** sign on the successful branch after the authentication action.

b) Type `RDG` in the search field.
c) Select **RDG Policy Assign** and click **Add Item**.
d) To display available policies, click the **Add/Delete** link.
e) Select a policy and click **Save**.

Without an RDG policy, APM denies access to each resource request.

10. Click the **Apply Access Policy** link to apply and activate the changes to the policy.

To apply this access policy to network traffic, add the access profile to a virtual server.

Note: *To ensure that logging is configured to meet your requirements, verify the log settings for the access profile.*

Configuring a machine account

You configure a machine account so that Access Policy Manager® (APM®) can establish a secure channel to a domain controller.

1. On the Main tab, click **Access** > **Authentication** > **NTLM** > **Machine Account**.
   A new Machine Account screen opens.
2. In the Configuration area, in the **Machine Account Name** field, type a name.
3. In the **Domain FQDN** field, type the fully qualified domain name (FQDN) for the domain that you want the machine account to join.
4. (Optional) In the **Domain Controller FQDN** field, type the FQDN for a domain controller.
5. In the **Admin User** field, type the name of a user who has administrator privilege.
6. In the **Admin Password** field, type the password for the admin user.

   APM uses these credentials to create the machine account on the domain controller. However, APM does not store the credentials and you do not need them to update an existing machine account configuration later.
7. Click **Join**.

This creates a machine account and joins it to the specified domain. This also creates a non-editable **NetBIOS Domain Name** field that is automatically populated.

Note: *If the **NetBIOS Domain Name** field on the machine account is empty, delete the configuration and recreate it. The field populates.*

Creating an NTLM Auth configuration

Create an NTLM Auth configuration to specify the domain controllers that a machine account can use to log in.

1. On the Main tab, click **Access** > **Authentication** > **NTLM** > **NTLM Auth Configuration**.
   A new NTLM Auth Configuration screen opens.
2. In the **Name** field, type a name.
3. From the **Machine Account Name** list, select the machine account configuration to which this NTLM Auth configuration applies.

   You can assign the same machine account to multiple NTLM authentication configurations.
4. For each domain controller, type a fully qualified domain name (FQDN) and click **Add**.

   Note: *You should add only domain controllers that belong to one domain.*

By specifying more than one domain controller, you enable high availability. If the first domain controller on the list is not available, Access Policy Manager® tries the next domain controller on the list, successively.

5. Click **Finished**.

This specifies the domain controllers that a machine account can use to log in.

Maintaining a machine account

In some networks, administrators run scripts to find and delete outdated machine accounts on the domain controllers. To keep the machine account up-to-date, you can renew the password periodically.

1. On the Main tab, click **Access** > **Authentication** > **NTLM** > **Machine Account**.
   The Machine Account screen opens.
2. Click the name of a machine account.
   The properties screen opens and displays the date and time of the last update to the machine account password.
3. Click the **Renew Machine Password** button.
   The screen refreshes and displays the updated date and time.

This changes the machine account last modified time.

Configuring a VDI profile

Configure a VDI profile to specify NTLM authentication for Microsoft RDP clients that use APM® as a gateway.

1. On the Main tab, click **Access** > **Connectivity / VPN** > **VDI / RDP** > **VDI Profiles**.
   The VDI Profiles list opens.
2. Click **Create**.
   A popup screen opens with **General Information** selected in the left pane and settings displayed in the right pane.
3. In the **Profile Name** field, type a name.
4. From the **Parent Profile** field, select an existing VDI profile.

   A VDI profile inherits properties from the parent profile. You can override them in this profile.
5. In the left pane, click **MSRDP Settings**.
   Settings in the right pane change.
6. From the **MSRDP NTLM Configuration** list, select an NTLM authentication configuration.
7. In the left pane, click **VMware View Settings**.
   Settings in the right pane change.
8. From the **Transport Protocol (UDP-only)** list, select a protocol.
9. Click **OK**.
   The popup screen closes.

The VDI profile displays on the screen.

To apply the VDI profile, you must specify it in a virtual server.

Creating a connectivity profile

You create a connectivity profile to configure client connections.

1. On the Main tab, click **Access** > **Connectivity / VPN** > **Connectivity** > **Profiles**.
   A list of connectivity profiles displays.

2. Click **Add**.
   The Create New Connectivity Profile popup screen opens and displays General Settings.
3. Type a **Profile Name** for the connectivity profile.
4. Select a **Parent Profile** from the list.
   APM® provides a default profile, **connectivity**.
5. Click **OK**.
   The popup screen closes, and the Connectivity Profile List displays.

The connectivity profile displays in the list.

To provide functionality with a connectivity profile, you must add the connectivity profile and an access profile to a virtual server.

Creating a custom Client SSL profile

You create a custom Client SSL profile when you want the BIG-IP® system to terminate client-side SSL traffic for the purpose of:

- Authenticating and decrypting ingress client-side SSL traffic
- Re-encrypting egress client-side traffic

By terminating client-side SSL traffic, the BIG-IP system offloads these authentication and decryption/encryption functions from the destination server.

1. On the Main tab, click **Local Traffic** > **Profiles** > **SSL** > **Client**.
   The Client SSL profile list screen opens.
2. Click **Create**.
   The New Client SSL Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select **clientssl** in the **Parent Profile** list.
5. From the **Configuration** list, select **Advanced**.
6. Select the **Custom** check box.
   The settings become available for change.
7. Next to Client Authentication, select the **Custom** check box.
   The settings become available.
8. From the **Configuration** list, select **Advanced**.
9. Modify the settings, as required.
10. Click **Finished**.

Creating a virtual server for SSL traffic

Define a virtual server to process SSL traffic from Microsoft RDP clients that use APM® as a gateway.

Note: *Users must specify the IP address of this virtual server as the gateway or RDG gateway from the RDP client that they use.*

---

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 for a host virtual server.

This field accepts an address in CIDR format (IP address/prefix). However, when you type the complete IP address for a host, you do not need to type a prefix after the address.

5. For the **Service Port**, do one of the following:
   - Type `443` in the field.
   - Select **HTTPS** from the list.
6. In the **SSL Profile (Client)** list, select an SSL profile.
7. In the Access Policy area, from the **Access Profile** list, select the access profile for RDP client authorization that you configured earlier.
8. From the **Connectivity Profile** list, select a profile.
9. From the **VDI Profile** list, select the VDI profile you configured earlier.
10. Click **Finished**.

Implementation result

Supported Microsoft RDP clients can specify a virtual server on the BIG-IP® system to use as a remote desktop gateway. Access Policy Manager® (APM®) can authorize the clients and authorize access to target servers as the clients request them.

Overview: Processing RDP traffic on a device configured for explicit forward proxy

If you configure Access Policy Manager® APM® as a gateway for RDP clients and configure APM to act as an explicit forward proxy on the same BIG-IP® system, you need to complete an additional configuration step to ensure that APM can process the RDP client traffic. The configuration F5 recommends for explicit forward proxy includes a catch-all virtual server, which listens on all IP addresses and all ports, on an HTTP tunnel interface.

When a programmatic API queries listeners for a specific IP and port, the query covers all interfaces and tunnels. As a result, the catch-all virtual server will always match. Sending traffic using this tunnel results in all packets being dropped because this virtual server is configured as a reject type of virtual server.

To prevent RDP client traffic from being dropped, add an additional wildcard port-specific virtual server on the HTTP tunnel interface.

Note: *Removing the catch-all virtual server from the HTTP tunnel interface is not recommended because doing so is counterproductive for security.*

Creating a virtual server for RDP client traffic

You specify a port-specific wildcard virtual server to match RDP client traffic on the HTTP tunnel interface for the Secure Web Gateway (SWG) explicit forward proxy configuration.

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 `0.0.0.0` to accept any IPv4 traffic.
5. In the **Service Port** field, type `3389`.

6. From the **Configuration** list, select **Advanced**.
7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**.
8. For the **VLANs and Tunnels** setting, move the HTTP tunnel interface used in the SWG explicit forward proxy configuration to the **Selected** list.

   The default tunnel is **http-tunnel**.

   This must be the same tunnel specified in the HTTP profile for the virtual server for forward proxy.
9. For the **Address Translation** setting, clear the **Enabled** check box.
10. Click **Finished**.

The virtual server now appears in the Virtual Server List screen.

About wildcard virtual servers on the HTTP tunnel interface

In the recommended Secure Web Gateway explicit forward proxy configuration, client browsers point to a forward proxy server that establishes a tunnel for SSL traffic. Additional wildcard virtual servers listen on the HTTP tunnel interface. The listener that best matches the web traffic directed to the forward proxy server handles the traffic.

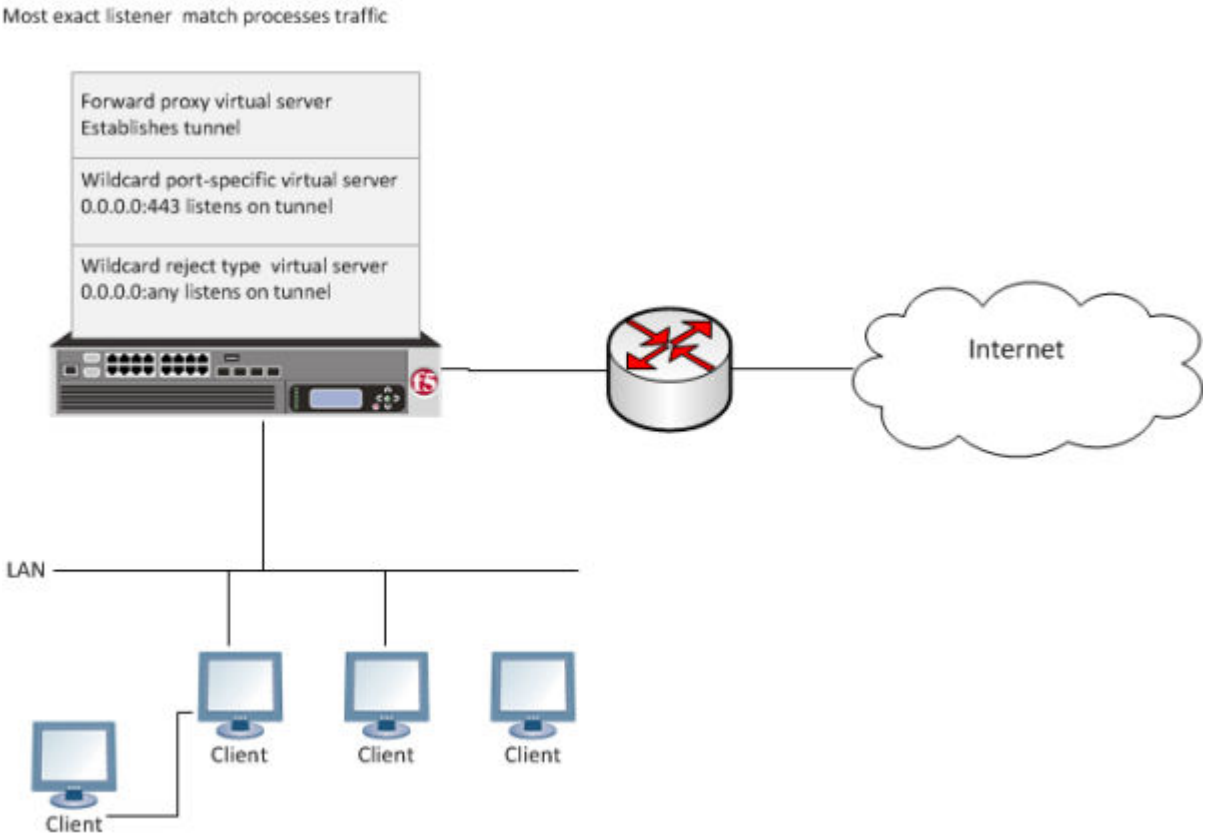

Figure 12: Explicit forward proxy configuration

Maintaining OPSWAT Libraries with a Sync-Failover Device Group

Overview: Updating antivirus and firewall libraries with a Sync-Failover device group

This implementation describes how to upload antivirus and firewall libraries from OPSWAT to one BIG-IP® Access Policy Manager® device, and to install an antivirus and firewall library to that device, or to multiple devices in a device group.

To download OPSWAT OESIS library updates, you must have an account with OPSWAT, and be able to download software updates.

To synchronize installation between multiple devices, you configure a Sync-Failover device group, which includes the devices between which you want to synchronize installation of updates. Device group setup requires establishing trust relationships between devices, creating a device group, and synchronization of settings.

About device groups and synchronization

When you have more than one BIG-IP® device in a local trust domain, you can synchronize BIG-IP configuration data among those devices by creating a device group. A *device group* is a collection of BIG-IP devices that trust each other and synchronize their BIG-IP configuration data. If you want to exclude certain devices from ConfigSync, you can simply exclude them from membership in that particular device group.

You can synchronize some types of data on a global level across all BIG-IP devices, while synchronizing other data in a more granular way, on an individual application level to a subset of devices.

Important: *To configure redundancy on a device, you do not need to explicitly specify that you want the BIG-IP device to be part of a redundant configuration. Instead, this occurs automatically when you add the device to an existing device group.*

Before you configure device trust

Before you configure device trust, you should consider the following:

- Only version 11.x or later systems can join the local trust domain.
- You can manage device trust when logged in to a certificate signing authority only. You cannot manage device trust when logged in to a subordinate non-authority device.
- If you reset trust authority on a certificate signing authority by retaining the authority of the device, you must subsequently recreate the local trust domain and the device group.
- As a best practice, you should configure the ConfigSync and mirroring addresses on a device before you add that device to the trust domain.

Task summary

The configuration process for a BIG-IP® system entails adding the OPSWAT library update to one system, then installing it to that same system, or to a device group. You must pre-configure a device group to install the update to multiple systems.

Task list

Establishing device trust
Adding a device to the local trust domain
Creating a Sync-Failover device group
Manually synchronizing the BIG-IP configuration
Uploading an OPSWAT update to Access Policy Manager
Installing an OPSWAT update on one or more Access Policy Manager devices
Viewing supported products in the installed OPSWAT EPSEC version

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management** > **Device Trust** > **Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
   - If the BIG-IP device is an appliance, type the management IP address for the device.
   - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
   - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
   - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.

8. Click **Add Device**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Adding a device to the local trust domain

Verify that each BIG-IP® device that is to be part of a local trust domain has a device certificate installed on it.

Follow these steps to log in to any BIG-IP® device on the network and add one or more devices to the local system's local trust domain.

Note: *Any BIG-IP devices that you intend to add to a device group at a later point must be members of the same local trust domain.*

---

1. On the Main tab, click **Device Management** > **Device Trust** > **Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
   - If the BIG-IP device is an appliance, type the management IP address for the device.
   - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
   - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
   - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
5. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
6. In the **Name** field, verify that the name of the remote device is correct.
7. Click **Add Device**.

After you perform this task, the local device and the device that you specified in this procedure have a trust relationship and, therefore, are qualified to join a device group.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP® devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management** > **Device Groups**.
2. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.
3. In the **Name** field, type a name for the device group.
4. From the **Group Type** list, select **Sync-Failover**.

5. In the **Description** field, type a description of the device group.

   This setting is optional.
6. From the **Configuration** list, select **Advanced**.
7. For the **Members** setting, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.

   The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only. Also, for vCMP-provisioned systems on platforms that contain a hardware security module (HSM) supporting FIPS multi-tenancy, the FIPS partitions on the guests in the device group must be identical with respect to the number of SSL cores allocated to the guest's FIPS partition and the maximum number of private SSL keys that the guest can store on the HSM.
8. From the **Sync Type** list:

   - Select **Automatic with Incremental Sync** when you want the BIG-IP system to automatically sync the most recent BIG-IP configuration changes from a device to the other members of the device group. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
   - Select **Manual with Incremental Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the latest BIG-IP configuration changes from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
   - Select **Manual with Full Sync** when you want to manually initiate a config sync operation. In this case, the BIG-IP system syncs the full set of BIG-IP configuration data from the device you choose to the other members of the device group. We strongly recommend that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
9. In the **Maximum Incremental Sync Size (KB)** field, retain the default value of `1024`, or type a different value.

   This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.
10. For the **Network Failover** setting, select or clear the check box:

    - Select the check box if you want device group members to handle failover communications by way of network connectivity. This is the default value and is required for active-active configurations.
    - Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

    For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.
11. In the **Link Down Time on Failover** field, use the default value of `0.0`, or specify a new value.

    This setting specifies the amount of time, in seconds, that interfaces for any external VLANs are down when a traffic group fails over and goes to the standby state. Specifying a value other than `0.0` for this setting causes other vendor switches to use the specified time to learn the MAC address of the newly-active device.

    Important: *This setting is a system-wide setting, and does not apply to this device group only. Specifying a value in this field causes the BIG-IP system to assign this value to the global bigdb variable* `failover.standby.linkdowntime`.

    ---
12. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Manually synchronizing the BIG-IP configuration

Before you perform this task, verify that device trust has been established and that all devices that you want to synchronize are members of a device group.

You perform this task when the automatic sync feature is disabled and you want to manually synchronize BIG-IP® configuration data among the devices in the device group. This synchronization ensures that any device in the device group can process application traffic successfully. You can determine the need to perform this task by viewing sync status in the upper left corner of any BIG-IP Configuration utility screen. A status of `Changes Pending` indicates that you need to perform a config sync within the device group.

Important: *You can log into any device in the device group to perform this task.*

---

1. On the Main tab, click **Device Management** > **Overview**.
2. In the Device Groups area of the screen, click the arrow next to the name of the relevant device group.
   The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, choose a device.
4. In the Sync Options area of the screen, choose an option:

Option	Description
Push the selected device configuration to the group	Select this option when you want to synchronize the configuration of the selected device to the other device group members.
Pull the most recent configuration to the selected device	Select this option when you want to synchronize the most recent configurations of one or more device group members to the selected device.

5. Click **Sync**.

After you initiate a manual config sync, the BIG-IP system compares the configuration data on the local device with the data on each device in the device group, and synchronizes the most recently-changed configuration data from one or more source devices to one or more target devices. Note that the system does not synchronize non-floating self IP addresses.

Uploading an OPSWAT update to Access Policy Manager

When new updates to OPSWAT antivirus and firewall libraries are made available, you can add these updates to the BIG-IP® system. To upload an update to the BIG-IP system, you must first download an update, using a registered account, from the OPSWAT web site.

1. On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Antivirus Check Updates screen displays a list of OPSWAT packages available on the device.
2. Click the **Upload** button to add an OPSWAT update.
   The Upload Package screen appears.
3. Click **Browse** and select an OPSWAT package ZIP file to upload.
4. Select an install option from the list.
   - Select **Do Not Install** to upload the package to the local device, but without installing the OPSWAT package on the system.

- Select **Install on this device** to upload the package to the local device, and then install the OPSWAT package to this device.
- Select **Install on device group** to upload the package to the local device, and then install the OPSWAT package on the device group. A list of available device groups appears, and you can select the device group on which to install.

5. Click **OK**.

The OPSWAT package file is added to the list on the Antivirus Check Updates screen. You can install or delete OPSWAT packages from this page.

Installing an OPSWAT update on one or more Access Policy Manager devices

After you have uploaded an OPSWAT antivirus and firewall library update to the BIG-IP® system, you can install the update to one or more BIG-IP systems in a device group.

1. On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Antivirus Check Updates screen displays a list of OPSWAT packages available on the device.
2. Double-click an OPSWAT package to view details about the update and included firewall or antivirus libraries.
3. Select an OPSWAT package and click **Install**.
   The Install Package screen opens.
4. Select **Install on device group** to upload the package to the local device, and then install the OPSWAT package on the device group. A list of available device groups appears, and you can select the device group on which to install.
5. Click **Ok**.

The OPSWAT update is installed on the selected systems. You can view the installed and available OPSWAT versions on the **Software Management** > **Antivirus Check Updates** screen.

Viewing supported products in the installed OPSWAT EPSEC version

You can always view details about any installed OPSWAT version, including supported antivirus, firewall, anti-spyware, hard disk encryption, peer-to-peer software, patch management software, and Windows Health Agent features for supported platforms.

1. To view the details for the current device group:
   a) Click the F5® logo to go to the start (Welcome) page.
   b) In the Support area, click the **OSWAT application integration support charts** link.
   The OPSWAT Integration web page opens in a new browser tab or window. By default, this page shows Antivirus Integration for Windows.
   c) From the lists at the top of the screen, select the page to view. You can select the supported EPSEC feature, and you can select to view supported products for **Windows**, **Mac**, or **Linux**.
   d) Click the **Show** button to view the list of supported products for the type and platform you selected.
2. To view the details for another device group or another OESIS version:
   a) On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Package Status screen displays a list of OPSWAT packages available on the device.
   b) Click the **Device EPSEC Status** button.
   The **Device EPSEC Status** screen appears and shows the installed OPSWAT version.
   c) To select a different device group on which to view the installed OPSWAT version, select the device group from the **Local Device/Device Group** list.
   d) Under **Installed OESIS version**, click the version number for which you want to view the OPSWAT features chart.

The OPSWAT Integration web page opens in a new browser tab or window. By default, this page shows Antivirus Integration for Windows.

e) From the lists at the top of the screen, select the page to view. You can select the supported EPSEC feature, and you can select to view supported products for **Windows**, **Mac**, or **Linux**.

f) Click the **Show** button to view the list of supported products for the type and platform you selected.

Implementation result

To summarize, you now have uploaded an OPSWAT update to one BIG-IP® system, and installed it to one system, or to multiple systems in a device group.

You can view the installed and available OPSWAT versions on the **Software Management** > **Antivirus Check Updates** screen.

Maintaining OPSWAT Libraries with a Sync-Only Device Group

Overview: Updating antivirus and firewall libraries with a Sync-Only device group

This implementation describes how to upload antivirus and firewall libraries from OPSWAT to one BIG-IP® Access Policy Manager® device, and to install an antivirus and firewall library to that device, or to multiple devices in a device group.

To download OPSWAT OESIS library updates, you must have an account with OPSWAT, and be able to download software updates.

To synchronize installation between multiple devices, you configure a Sync-Only device group, which includes the devices between which you want to synchronize installation of updates. Device group setup requires establishing trust relationships between devices, creating a device group, and synchronization of settings.

About device groups and synchronization

When you have more than one BIG-IP® device in a local trust domain, you can synchronize BIG-IP configuration data among those devices by creating a device group. A *device group* is a collection of BIG-IP devices that trust each other and synchronize their BIG-IP configuration data. If you want to exclude certain devices from ConfigSync, you can simply exclude them from membership in that particular device group.

You can synchronize some types of data on a global level across all BIG-IP devices, while synchronizing other data in a more granular way, on an individual application level to a subset of devices.

Important: *To configure redundancy on a device, you do not need to explicitly specify that you want the BIG-IP device to be part of a redundant configuration. Instead, this occurs automatically when you add the device to an existing device group.*

Before you configure device trust

Before you configure device trust, you should consider the following:

- Only version 11.x or later systems can join the local trust domain.
- You can manage device trust when logged in to a certificate signing authority only. You cannot manage device trust when logged in to a subordinate non-authority device.
- If you reset trust authority on a certificate signing authority by retaining the authority of the device, you must subsequently recreate the local trust domain and the device group.
- As a best practice, you should configure the ConfigSync and mirroring addresses on a device before you add that device to the trust domain.

Task summary

The configuration process for a BIG-IP® system entails adding the OPSWAT library update to one system, then installing it to that same system, or to a device group. You must pre-configure a device group to install the update to multiple systems.

Task list

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices `Bigip_1`, `Bigip_2`, and `Bigip_3` each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device `Bigip_1` and add devices `Bigip_2` and `Bigip_3` to the local trust domain; there is no need to repeat this process on devices `Bigip_2` and `Bigip_3`.

1. On the Main tab, click **Device Management** > **Device Trust** > **Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
   - If the BIG-IP device is an appliance, type the management IP address for the device.
   - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
   - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
   - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
5. Click **Retrieve Device Information**.
6. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
7. In the **Name** field, verify that the name of the remote device is correct.
8. Click **Add Device**.

After you perform this task, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat this task to specify each device that you want to add to the local trust domain.

Adding a device to the local trust domain

Verify that each BIG-IP® device that is to be part of a local trust domain has a device certificate installed on it.

Follow these steps to log in to any BIG-IP® device on the network and add one or more devices to the local system's local trust domain.

Note: *Any BIG-IP devices that you intend to add to a device group at a later point must be members of the same local trust domain.*

---

1. On the Main tab, click **Device Management** > **Device Trust** > **Device Trust Members**.
2. Click **Add**.
3. From the **Device Type** list, select **Peer** or **Subordinate**.
4. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
   - If the BIG-IP device is an appliance, type the management IP address for the device.
   - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
   - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
   - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
5. Verify that the certificate of the remote device is correct, and then click **Device Certificate Matches**.
6. In the **Name** field, verify that the name of the remote device is correct.
7. Click **Add Device**.

After you perform this task, the local device and the device that you specified in this procedure have a trust relationship and, therefore, are qualified to join a device group.

Creating a Sync-Only device group

You perform this task to create a Sync-Only type of device group. When you create a Sync-Only device group, the BIG-IP® system can then automatically synchronize configuration data in folders attached to the device group (such as security policies and acceleration applications) with the other devices in the group, even when some of those devices reside in another network.

Note: *You perform this task on any one BIG-IP device within the local trust domain; there is no need to repeat this process on the other devices in the device group.*

---

1. On the Main tab, click **Device Management** > **Device Groups**.
2. Find the **Partition** list in the upper right corner of the BIG-IP Configuration utility screen, to the left of the **Log out** button.
3. From the **Partition** list, pick partition `Common`.
4. On the Device Groups list screen, click **Create**.
   The New Device Group screen opens.

5. From the **Configuration** list, select **Advanced**.
6. For the **Members** setting, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group. Use the Move button to move the host name to the **Includes** list.

   The list shows any devices that are members of the device's local trust domain.
7. For the **Full Sync** setting, specify whether the system synchronizes the entire configuration during synchronization operations:

   - Select the check box when you want all sync operations to be full syncs. In this case, every time a config sync operation occurs, the BIG-IP system synchronizes all configuration data associated with the device group. This setting has a performance impact and is not recommended for most customers.
   - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

   If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.
8. In the **Maximum Incremental Sync Size (KB)** field, retain the default value of `1024`, or type a different value.

   This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.
9. Click **Finished**.

You now have a Sync-Only type of device group containing BIG-IP devices as members.

Uploading an OPSWAT update to Access Policy Manager

When new updates to OPSWAT antivirus and firewall libraries are made available, you can add these updates to the BIG-IP® system. To upload an update to the BIG-IP system, you must first download an update, using a registered account, from the OPSWAT web site.

1. On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Antivirus Check Updates screen displays a list of OPSWAT packages available on the device.
2. Click the **Upload** button to add an OPSWAT update.
   The Upload Package screen appears.
3. Click **Browse** and select an OPSWAT package ZIP file to upload.
4. Select an install option from the list.

   - Select **Do Not Install** to upload the package to the local device, but without installing the OPSWAT package on the system.
   - Select **Install on this device** to upload the package to the local device, and then install the OPSWAT package to this device.
   - Select **Install on device group** to upload the package to the local device, and then install the OPSWAT package on the device group. A list of available device groups appears, and you can select the device group on which to install.
5. Click **OK**.

The OPSWAT package file is added to the list on the Antivirus Check Updates screen. You can install or delete OPSWAT packages from this page.

Installing an OPSWAT update on one or more Access Policy Manager devices

After you have uploaded an OPSWAT antivirus and firewall library update to the BIG-IP® system, you can install the update to one or more BIG-IP systems in a device group.

1. On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Antivirus Check Updates screen displays a list of OPSWAT packages available on the device.
2. Double-click an OPSWAT package to view details about the update and included firewall or antivirus libraries.
3. Select an OPSWAT package and click **Install**.
   The Install Package screen opens.
4. Select **Install on device group** to upload the package to the local device, and then install the OPSWAT package on the device group. A list of available device groups appears, and you can select the device group on which to install.
5. Click **Ok**.

The OPSWAT update is installed on the selected systems. You can view the installed and available OPSWAT versions on the **Software Management** > **Antivirus Check Updates** screen.

Viewing supported products in the installed OPSWAT EPSEC version

You can always view details about any installed OPSWAT version, including supported antivirus, firewall, anti-spyware, hard disk encryption, peer-to-peer software, patch management software, and Windows Health Agent features for supported platforms.

1. To view the details for the current device group:
   a) Click the F5® logo to go to the start (Welcome) page.
   b) In the Support area, click the **OSWAT application integration support charts** link.
   The OPSWAT Integration web page opens in a new browser tab or window. By default, this page shows Antivirus Integration for Windows.
   c) From the lists at the top of the screen, select the page to view. You can select the supported EPSEC feature, and you can select to view supported products for **Windows**, **Mac**, or **Linux**.
   d) Click the **Show** button to view the list of supported products for the type and platform you selected.
2. To view the details for another device group or another OESIS version:
   a) On the Main tab, click **System** > **Software Management** > **Antivirus Check Updates**.
   The Package Status screen displays a list of OPSWAT packages available on the device.
   b) Click the **Device EPSEC Status** button.
   The **Device EPSEC Status** screen appears and shows the installed OPSWAT version.
   c) To select a different device group on which to view the installed OPSWAT version, select the device group from the **Local Device/Device Group** list.
   d) Under **Installed OESIS version**, click the version number for which you want to view the OPSWAT features chart.
   The OPSWAT Integration web page opens in a new browser tab or window. By default, this page shows Antivirus Integration for Windows.
   e) From the lists at the top of the screen, select the page to view. You can select the supported EPSEC feature, and you can select to view supported products for **Windows**, **Mac**, or **Linux**.
   f) Click the **Show** button to view the list of supported products for the type and platform you selected.

Implementation result

To summarize, you now have uploaded an OPSWAT update to one BIG-IP® system, and installed it to one system, or to multiple systems in a device group.

You can view the installed and available OPSWAT versions on the **Software Management** > **Antivirus Check Updates** screen.

Adding Hosted Content to Access Policy Manager

About uploading custom files to Access Policy Manager

You can upload custom files to BIG-IP® Access Policy Manager®(APM®) to provide resources directly to users.

For example, you can upload BIG-IP Edge Client® installers, antivirus or firewall update packages, or Citrix receiver files for your users to download. You can upload custom images, web pages, Java archives, JavaScript files, CSS files, archive files, and many other types of files as well.

Optionally, you can compress and upload multiple files as a single ZIP archive file. When you upload an archive file, you can choose to either upload the compressed file, or upload and extract the compressed file.

Upload Only
Select this option to upload an archived file that must remain in archive format. For example, you can upload a ZIP file for a user to download, containing a package of documents, or an application and related files. Some applications also use archived files; for example, you will upload a JAR file without extracting it.

Upload and Extract
Select this option to upload an archived file and extract it to the specified location. The folder hierarchy of the extracted file is preserved when you use this action. Select this option when you are uploading a collection of files that must be separated on the server for use by the end user; for example, to upload a web application that includes top-level HTML files, and subdirectories containing scripts, images, CSS, and other files.

Understanding hosted content

Hosted content is any type of file you would like to serve from Access Policy Manager® (APM®) to access policy users. Hosted content can include executable files, scripts, text, HTML, CSS files, and image files. You can serve hosted content from a webtop link, or from a portal access link.

About accessing hosted content

To access hosted content, a user must belong to an access profile that is associated with the hosted content. After content is uploaded to Access Policy Manager® (APM®), the entire hosted content library must be associated with one or more access profiles. These access profiles alone can view the content.

In addition, each file uploaded to the hosted content repository is assigned a permission level that determines the users who can access that content.

Permissions for hosted content

A permission level is assigned to each file in the hosted content repository, as described here.

Permission level	Description
policy	The file is available only to users who have successfully completed an access policy, with an **Allow** ending result, and an access profile

Permission level	Description
	associated with the hosted content repository. You can assign this to display an HTML file that only a verified user can see.
public	The file is available to anyone with an access profile associated with the hosted content repository. You can assign this to allow access to an installation package that a user needs to start an access session.
session	The file is available only to users with an active access policy session and an access profile associated with the hosted content repository. You can assign this to allow a user with an active session access to a required logon component.

Task summary

To add hosted content to Access Policy Manager® (APM®), complete these tasks.

Task list

Uploading files to Access Policy Manager

Associating hosted content with access profiles

Uploading files to Access Policy Manager

Before you upload multiple files to Access Policy Manager®, you can compress and combine the files into a ZIP archive file. Then, you can upload and extract the files in one step.

You can upload files to Access Policy Manager to provide content for public viewing, to provide pages and content to Portal Access connections, or to provide customized webtop links.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. Click the **Upload** button.
   The Create New File popup screen opens.
3. For the **Select File** setting, click the **Browse** button and select the file to upload.
   - To upload each file separately, select the first file, then repeat this step for all remaining files.
   - To upload all files at once from a compressed file, select the compressed file.

   The **Select File** and **File Name** fields are populated with the file name.
4. If you are uploading a compressed file that you want to extract, from the **File Action** list, select **Upload and Extract**.
5. Click **OK**.
   The file appears in the hosted content list.

You must associate any access profiles that will access hosted content with the hosted content repository.

Associating hosted content with access profiles

A user can access hosted content that is associated with that user's access profile. Each access profile that requires hosted content access must be associated with the entire hosted content repository.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. On the **Upload** button, click the right-side arrow to select **Manage Access** from the list.
   The Access Settings popup screen opens.
3. Select the access profiles to associate with hosted content, then click **OK**.

   A user must belong to an associated access profile to access hosted content.

View the hosted content list, and verify that the access policy association was successful.

Implementation result

As a result of these implementation tasks, you have edited files and deleted hosted files on Access Policy Manager® as necessary.

Editing Hosted Content with Access Policy Manager

About editing hosted files on Access Policy Manager

You can upload custom files to BIG-IP® Access Policy Manager® to provide resources directly to users.

You might need to edit files after you upload them to Access Policy Manager, such as to rename a file or change the file MIME type. You can make these changes using the hosted content settings.

Task summary

To edit hosted content on Access Policy Manager®, complete these tasks.

Task list

Renaming or moving hosted content files

You can rename or move a hosted content file on Access Policy Manager®.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. At bottom left of the screen, click the right-side arrow on the **Edit** button to select **Rename/Move File** from the list.
   The **Rename/Move File Properties** popup screen opens.
3. In the **New File Name** field, type a new name for the file.
4. In the **New File Destination Folder**, specify a new destination folder for the file.
5. Click **OK**.
   The file changes are saved, and the screen returns to the hosted content list.

Editing hosted content file properties

You can edit the permissions and MIME type for hosted content files on Access Policy Manager®.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. At bottom left of the screen, click the right-side arrow on the **Edit** button to select **Edit File Properties** from the list.
   The **Edit File Properties** popup screen opens.
3. If the MIME type for the file is incorrect or must be changed, from the **Mime Type** list, select the MIME type for the file.
4. From the **Secure Level** menu, select the access level for the file.

Option	Description
policy	The file is available only to users who have successfully completed an access policy, with an **Allow** ending result. You might use this to display an HTML file that only a verified user can see.
public	The file is available to anyone. You might use this to allow access to an installation package that a user needs to start an access session.
session	The file is available only to users with an active access policy session. You might use this to allow a user with an active session access to a required logon component.

5. Click **OK**.
   The file changes are saved, and the screen returns to the hosted content list.

The settings for the file are displayed in the Hosted Content list.

Replacing a hosted file

You can upload a new version of a file to hosted content, to replace the current file on Access Policy Manager®.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. At bottom left of the screen, click the right-side arrow on the **Edit** button to select **Upload New Version** from the list.
   The **Upload New File Version** popup screen opens.
3. For the **Select File** setting, click the **Browse** button and select the file to upload.
   The **Select File** and **File Name** fields are populated with the file name.
4. If the MIME type for the file is incorrect or must be changed, from the **Mime Type** list, select the MIME type for the file.
5. From the **Secure Level** menu, select the access level for the file.

Option	Description
policy	The file is available only to users who have successfully completed an access policy, with an **Allow** ending result. You might use this to display an HTML file that only a verified user can see.
public	The file is available to anyone. You might use this to allow access to an installation package that a user needs to start an access session.
session	The file is available only to users with an active access policy session. You might use this to allow a user with an active session access to a required logon component.

6. Click **OK**.
   The file changes are saved, and the screen returns to the hosted content list.

View the hosted content list to verify your changes to the file.

Deleting a hosted file

You can delete one or more files from the hosted content on Access Policy Manager®.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. Select one or more files to delete. To select all files, select the check box at the top of the list, next to the Name column.
3. Click **Delete**, and in the **Delete File** popup screen that opens, click **Yes**.

The files are removed from the list.

Implementation result

As a result of these implementation tasks, you have edited files and deleted hosted files on Access Policy Manager® as necessary.

Hosting a BIG-IP Edge Client Download with Access Policy Manager

About hosting a BIG-IP Edge Client file on Access Policy Manager

You can host files on BIG-IP® Access Policy Manager® (APM®) so clients can download them.

When you host a file on Access Policy Manager, you can provide the link to the file in a number of ways. In this example, the BIG-IP Edge Client® for Mac link is provided as a link on the user's webtop. The user connects through the web client, then clicks a link on the webtop to download the client file. To provide the BIG-IP Edge Client for Mac, first you must create a connectivity profile. Then, you can download the Mac client file as a ZIP file.

Task summary

To add the BIG-IP® Edge Client® for Mac file to the hosted content repository on Access Policy Manager®, so clients can download it, complete these tasks.

Task list

Configuring a connectivity profile for Edge Client for Mac
Downloading the ZIP file for Edge Client for Mac
Uploading BIG-IP Edge Client to hosted content on Access Policy Manager
Associating hosted content with access profiles
Creating a webtop link for the client installer
Adding a webtop, links, and sections to an access policy

Configuring a connectivity profile for Edge Client for Mac

Update the connectivity profile in your Network Access configuration to configure security settings, servers, and location-awareness for BIG-IP® Edge Client® for Mac.

1. On the Main tab, click **Access** > **Connectivity / VPN** > **Connectivity** > **Profiles**.
   A list of connectivity profiles displays.
2. Select the connectivity profile that you want to update and click **Edit Profile**.
   The Edit Connectivity Profile popup screen opens and displays General Settings.
3. From the left pane of the popup screen, select **Win/Mac Edge Client**.
   Edge Client settings for Mac and Windows-based systems display in the right pane.
4. Retain the default (selected) or clear the **Save Servers Upon Exit** check box.

   Specifies whether Edge Client maintains a list of recently used user-entered APM® servers. Edge Client always lists the servers that are defined in the connectivity profile, and sorts them by most recent access, whether this option is selected or not.
5. To support automatic reconnection without the need to provide credentials again, allow password caching.
   a) Select the **Allow Password Caching** check box.

      This check box is cleared by default.

The remaining settings on the screen become available.

b) To require device authentication to unlock the saved password, select **Require Device Authentication.**

This option links the option to use a saved password to a device authentication method. Supported device authentication methods include PIN, passphrase, and biometric (fingerprint) authentication on iOS and Android. Android devices also support pattern unlocking.

c) From the **Save Password Method** list, select **disk** or **memory**.

If you select **disk**, Edge Client caches the user's password (in encrypted form) securely on the disk where it is persisted even after the system is restarted or Edge Client is restarted.

If you select **memory**, Edge Client caches the user's password within the BIG-IP Edge Client application for automatic reconnection purposes.

If you select **memory**, the **Password Cache Expiration (minutes)** field displays with a default value of 240.

d) If the **Password Cache Expiration (minutes)** field displays, retain the default value or type the number of minutes to save the password in memory.

6. To enable automatic download and update of client packages, from the **Component Update** list, select **yes** (default).

If you select **yes**, APM® updates Edge Client software automatically on the client system when newer versions are available.

7. Specify the list of APM servers to provide when the client connects.

The servers you add here display as connection options in the BIG-IP Edge Client.

Note: *Users can select from these servers or they can type a hostname.*

---

a) From the left pane of the popup screen, select **Server List**.
A table displays in the right pane.

b) Click **Add**.
A table row becomes available for update.

c) You must type a host name in the **Host Name** field.

Typing an alias in the **Alias** field is optional.

d) Click **Update**.
The new row is added at the top of the table.

e) Continue to add servers, and when you are done, click **OK**.

8. Specify DNS suffixes that are considered to be in the local network.

Providing a list of DNS suffixes for the download package enables Edge Client to support the autoconnect option. With **Auto-Connect** selected, Edge Client uses the DNS suffixes to automatically connect when a client is not on the local network (not on the list) and automatically disconnect when the client is on the local network.

a) From the left pane of the popup screen, select **Location DNS List**.
Location DNS list information is displayed in the right pane.

b) Click **Add**.
An update row becomes available.

c) Type a name and click **Update**.

Type a DNS suffix that conforms to the rules specified for the local network.

The new row displays at the top of the table.

d) Continue to add DNS names and when you are done, click **OK**.

9. Click **OK**.
The popup screen closes, and the Connectivity Profile List displays.

Downloading the ZIP file for Edge Client for Mac

You can download a Mac Client package and distribute it to clients.

1. On the Main tab, click **Access** > **Connectivity / VPN** > **Connectivity** > **Profiles**.
   A list of connectivity profiles displays.
2. Select a connectivity profile.
3. Click the arrow on the **Customize Package** button and select **Mac**.
   The Customize Mac Client Package screen displays.
4. Click **Download**.

   The screen closes and the package, `BIGIPMacEdgeClient.zip`, downloads.

The ZIP file includes a Mac installer package (PKG) file and configuration settings.

Distribute the entire ZIP file to your users.

Uploading BIG-IP Edge Client to hosted content on Access Policy Manager

Upload the client file to the Access Policy Manager® hosted content repository so you can provide it to clients through a download link.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. Click the **Upload** button.
   The Create New File popup screen opens.
3. For the **Select File** setting, click the **Browse** button. Browse and select the `BIGIPMacEdgeClient.zip` file that you previously downloaded.
   The **Select File** and **File Name** fields are populated with the file name.
4. From the **File Action** list, select **Upload Only**.
5. In the **File Destination Folder** field, specify the folder path in which to place the file. For purposes of this example, the folder `/client` is specified.
6. Click **OK**.
   The file appears in the hosted content list.

You must associate any access profiles that will access hosted content with the hosted content repository.

Associating hosted content with access profiles

A user can access hosted content that is associated with that user's access profile. Each access profile that requires hosted content access must be associated with the entire hosted content repository.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. On the **Upload** button, click the right-side arrow to select **Manage Access** from the list.
   The Access Settings popup screen opens.
3. Select the access profiles to associate with hosted content, then click **OK**.

   A user must belong to an associated access profile to access hosted content.

View the hosted content list, and verify that the access policy association was successful.

Creating a webtop link for the client installer

You can create and customize links that you can assign to full webtops. In this context, *links* are defined applications and web sites that appear on a webtop, and can be clicked to open a web page or application. You can customize these links with descriptions and icons.

1. On the Main tab, click **Access** > **Webtops** > **Webtop Links**.
2. Click **Create**.
   The New Webtop Link screen opens.
3. In the **Name** field, type a name for the webtop.
4. From the **Link Type** list, select **Hosted Content**.
5. From the **Hosted File** link, select `public/share/client/BIGIPMacEdgeClient.zip`.
6. In the **Caption** field, type a descriptive caption.

   The **Caption** field is pre-populated with the text from the **Name** field. Type the link text that you want to appear on the web link.
7. If you want to add a detailed description, type it in the **Detailed Description** field.
8. To specify an icon image for the item on the webtop, click in the **Image** field and choose an image, or click the **Browse** button.

   Click the **View/Hide** link to show or hide the currently selected image.
9. Click **Finished**.

The webtop link is now configured, and appears in the list, and on a full webtop assigned with the same action. You can edit the webtop link further, or assign it to an access policy.

Before you can use this webtop link, it must be assigned to an access policy with a full webtop, using either an advanced resource assign action or a webtop,links and sections assign action.

Adding a webtop, links, and sections to an access policy

You must have an access profile set up before you can add a webtop, links, and sections to an access policy.

You can add an action to an access policy to add a webtop, webtop links, and webtop sections to an access policy branch. Webtop links and webtop sections are displayed on a full webtop.

Important: *Do not assign a webtop for a portal access connection configured for minimal patching mode; this configuration does not work.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile for which you want to edit the access policy.
   The properties screen opens for the profile you want to edit.
3. On the menu bar, click **Access Policy**.
4. In the General Properties area, click the **Edit Access Policy for Profile** ***`profile_name`*** link.
   The visual policy editor opens the access policy in a separate screen.
5. On a policy branch, click the **(+)** icon to add an item to the policy.
   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
6. On the Assignment tab, select the **Webtop, Links and Sections Assign** agent and click **Add Item**.
   The Webtop, Links and Sections Assignment screen opens.
7. In the **Name** field, type a name for the policy item.

   This name is displayed in the action field for the policy.

8. For each type of resource that you want assign:
   a) Click the **Add/Delete** link next to the resource type (**Webtop Links**, **Webtop Sections**, or **Webtop**).

      Available resources are listed.
   b) Select from the list of available resources.

      Select only one webtop.
   c) Click **Save**.
9. Click the **Save** button to save changes to the access policy item.

You can now configure further actions on the successful and fallback rule branches of this access policy item.

Click the **Apply Access Policy** link to apply and activate your changes to this access policy.

Note: *To ensure that logging is configured to meet your requirements, verify the log settings for the access profile.*

Implementation result

As a result of these implementation tasks, you have added the client file to a webtop link.

Hosting Files with Portal Access on Access Policy Manager

About using hosted files with a Portal Access resource

You can use hosted content that you have uploaded to the BIG-IP® Access Policy Manager® to provide the resource and resource items for a Portal Access resource.

When you use hosted content for a Portal Access resource, the link on the webtop for the portal access resource opens a file hosted on the system, instead of a URI. You configure the main Portal Access resource as this linked file. You then configure this file, and all related and required files, as resource items of this file.

In this example, a simple web page consisting of an HTML file, a CSS file, a JavaScript file, and an image are uploaded to a directory in the hosted content repository. The files are then specified as a Portal Access resource and resource items.

File	Location	Description
index.html	`/index.html`	The main web page that displays when the link is clicked. This is the Portal Access Resource.
styles.css	`/styles.css`	The CSS file for the page index.html.
test_image.jpg	`/test_image.jpg`	An image that is referenced on the page index.html.
script.js	`/js/script.js`	A JavaScript file that is referenced from the page index.html.

In this example, hosted content is uploaded as a single **ZIP** file, `test.zip`, then extracted to the location `/test` on the server.

Task summary

To add hosted content to a Portal Access link on Access Policy Manager® (APM®), complete these tasks.

Task list

Uploading files to Access Policy Manager for Portal Access
Associating hosted content with access profiles
Creating a portal access configuration with hosted content
Creating a portal access resource item for hosted content

Uploading files to Access Policy Manager for Portal Access

You upload files to Access Policy Manager® to provide content for a Portal Access webtop link.

Tip: *Before you upload multiple files to Access Policy Manager, you can combine the files in a ZIP archive format. Then, you can upload and extract the files in one step. In this example, four files are uploaded as a single ZIP archive, called* `test.zip`*.*

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. Click the **Upload** button.
   The Create New File popup screen opens.
3. Under **Select File**, click the **Browse** button. Browse and select **test.zip**.
   The **Select File** and **File Name** fields are populated with the file name.
4. In the **File Destination Folder** field, specify the folder path `/test` in which to place the file.
5. From the **File Action** list, select **Upload and Extract**.
6. Click the **OK** button.
   The files appears in the hosted content list, in the folder specified. Any files in subfolders in the archive file also appear in subfolders in the hosted content list.

You must associate any access profiles that will access hosted content with the hosted content repository.

Associating hosted content with access profiles

A user can access hosted content that is associated with that user's access profile. Each access profile that requires hosted content access must be associated with the entire hosted content repository.

1. On the Main tab, click **Access** > **Webtops** > **Hosted Content** > **Manage Files**.
   The Manage Files screen opens.
2. On the **Upload** button, click the right-side arrow to select **Manage Access** from the list.
   The Access Settings popup screen opens.
3. Select the access profiles to associate with hosted content, then click **OK**.

   A user must belong to an associated access profile to access hosted content.

View the hosted content list, and verify that the access policy association was successful.

Creating a portal access configuration with hosted content

1. On the Main tab, click **Access** > **Connectivity / VPN** > **Portal Access** > **Portal Access Lists**.
   The Portal Access List screen opens.
2. Click the **Create** button.
   The New Resource screen opens.
3. Type the name and an optional description.
4. From the **ACL Order** list, specify the placement for the resource.

Option	Description
Last	Select this option to place the new portal access resource last in the ACL list.
After	Select this option to select, from the list of configured ACLs, the ACL that this portal access resource should follow in sequence.
Specify	Select this option to specify an order number, for example, `0` or `631`for the ACL.

5. From **Configuration**, select **Basic** or **Advanced**.

   The **Advanced** option provides additional settings so you can configure a proxy host and port.
6. For the **Match Case for Paths** setting, select **Yes** to specify that portal access matches alphabetic case when matching paths in the portal access resource.
7. From the **Patching Type** list, select the patching type for the web application.

   For both full and minimal patching types, you can select or clear patching methods specific to your selection.

8. If you selected **Minimal Patching** and the **Host Patching** option, type a host search string, or multiple host search strings separated with spaces, and the host replace string, which must be the Access Policy Manager® virtual server IP address or fully qualified domain name.
9. Select the **Publish on Webtop** check box.
10. From the **Link Type** list, select **Hosted Content**.
11. From the **Hosted File** list, select `public/share/test/index.html`.

    This is the filename for this example scenario only. Please select the correct file for your own configuration.
12. In the Customization Settings for English area, in the **Caption** field, type a caption.

    The caption appears on the full webtop, and is required. This field is required even if you do not select the **Publish on webtop** option.
13. Optionally, in the **Detailed Description** field type a description for the web application.
14. In the **Image** field, specify an icon for the web application link. Click the **View/Hide** link to show the current icon.
15. If your application is behind a proxy server, to specify a proxy host and port, you must select **Advanced** for the configuration to display additional fields, and type the proxy host and proxy port.
16. Click the **Create** button.
    The Portal Access resource is saved, and the Portal Access Resource screen now shows a **Resource Items** area.

This completes the portal access resource configuration.

Specify all hosted content files used by this example (all files in the `/test` folder) as resource items.

Creating a portal access resource item for hosted content

You create a portal access resource item in order for hosted content to add a file that is part of a portal access hosted content resource. For example, you might add image files, CSS files, or scripts that are required by the web page or application. You typically use resource items to refine the behavior for web application directories; for example, you might specify `No Compression` and a `Cache All` caching policy for the images for a portal access resource.

Note: *You must add (separately) each hosted file used by the portal access resource, and the resource file itself, as resource items.*

---

1. On the Main tab, click **Access** > **Connectivity / VPN** > **Portal Access** > **Portal Access Lists**.
   The Portal Access List screen opens.
2. Click the name of a portal access resource.
   The Portal Access Properties screen for that resource opens.
3. In the Resource Items area, click the **Add** button.
   A New Resource Item screen for that resource opens.
4. Select that the resource item type is **Hosted Content**.
5. From the **Hosted File** list, select the file to specify as a resource item.

   For purposes of this example, specify `public/share/test/index.html`, `public/share/test/test_image.jpg`, `public/share/test/style.css`, and `public/share/test/js/script.js`.
6. Configure the properties for the resource item.
   - To add headers, select **Advanced** next to New Resource Item.
   - To configure **Session Update**, **Session Timeout**, and **Home Tab**, select **Advanced** next to Resource Item Properties.
7. Click **Finished**.

This creates the portal access resource item.

Implementation result

You have now added a portal access resource and portal access resource items that are based on uploaded hosted content.

Managing Disk Space for Hosted Content

Overview: Managing disk space for hosted content files

By default, the BIG-IP® system allocates 512 MB of disk space to the sandbox for storing hosted content files. If disk space becomes exhausted when you try to upload a hosted content file, an error displays. You can increase the amount of disk space allocated to the sandbox to the maximum of 1024 MB, in addition to deleting any hosted content that you no longer need.

Task summary

Allocating the maximum amount of disk space for hosted content

You can specify the amount of disk space allocated for hosted content using a database variable.

Note: *The maximum supported disk space is 1024 MB.*

1. Log on to the BIG-IP® system command line and type `tmsh`.
2. Type this command sequence `sys db`.
3. To view the amount of space currently allocated for hosted content, type this command sequence `list total.sandbox.size value`.
4. To specify the amount of disk space allocated for hosted content:
   a) Type this command sequence `modify total.sandbox.size value`.
   This prompt displays. `Values: [enter integer value min:64 max:1024]`
   b) Type a value and press Enter.

Estimating hosted content file disk space usage

To estimate how much disk space hosted content files consume, you can display the sizes of the files in the sandbox from the command line.

1. Log on to the BIG-IP® system command line and type `tmsh`.
2. Type this command sequence `apm resource sandbox list files | grep size`.
   File sizes display.

```
size 397325
size 752662
```

Importing and Exporting Access Profiles

Overview: Importing and exporting access profiles

You can export or import an access profile for the purpose of backing up and restoring a profile or for copying it from one Access Policy Manager® (APM®) system to another. An access profile is exported to a compressed tar file. Import is supported on a system with the same version of APM that created the export file.

Task summary

Exporting an access profile
Importing an access profile

Exporting an access profile

You can export any access profile and later import it to restore it on the same BIG-IP system with Access Policy Manager® (APM®), or import it to another BIG-IP® system with APM (at the same software version).

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Locate the access profile that you want to export.
3. In the Export column, click the **Export** link.
   A compressed archive file (`gz`) downloads. The full filename is `profile-<PartitionName><AccessProfileName>.conf.tar.gz`.

Importing an access profile

You can import an access profile that was previously exported from an Access Policy Manager® (APM®) system to restore the access profile or to add an access profile from another APM system.

Note: *Do not import a policy if you exported it from another version of APM.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click **Import**.
   The Import Profile screen opens.
3. In the **New Profile Name** box, type the name for the new profile.

   Note: *The access profile name must be unique among all access profile and per-request policy names.*

   ---

4. Next to the **Config File Upload** field, click **Choose File**.
   A popup screen opens.
5. Select the file to import and click the **Open** button.

   The full filename for an exported access profile usually follows this pattern `profile-<PartitionName><AccessProfileName>.conf.tar.gz`.
6. Select the **Reuse Existing Objects** check box to reuse objects that exist on the server.

This option reuses APM configuration objects, such as server definitions or resources, instead of recreating them for use with the imported access policy.

Important: *Whether the* ***Reuse Existing Objects*** *check box is enabled or cleared, if these objects exist on the system, they are reused:*

- OAuth client application
- OAuth resource server
- OAuth scope

---

7. Click **Import**.
   The file is imported to the system.

Logging and Reporting

Overview: Configuring remote high-speed APM and SWG event logging

You can configure the BIG-IP® system to log information about Access Policy Manager® (APM®) and Secure Web Gateway events and send the log messages to remote high-speed log servers.

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

Object	Reason
Pool of remote log servers	Create a pool of remote log servers to which the BIG-IP system can send log messages.
Destination (unformatted)	Create a log destination of Remote High-Speed Log type that specifies a pool of remote log servers.
Destination (formatted)	If your remote log servers are the ArcSight, Splunk, or Remote Syslog type, create an additional log destination to format the logs in the required format and forward the logs to a remote high-speed log destination.
Publisher	Create a log publisher to send logs to a set of specified log destinations.
Log Setting	Add event logging for the APM system and configure log levels for it or add logging for URL filter events, or both. Settings include the specification of up to two log publishers: one for access system logging and one for URL request logging.
Access profile	Add log settings to the access profile. The log settings for the access profile control logging for the traffic that comes through the virtual server to which the access profile is assigned.

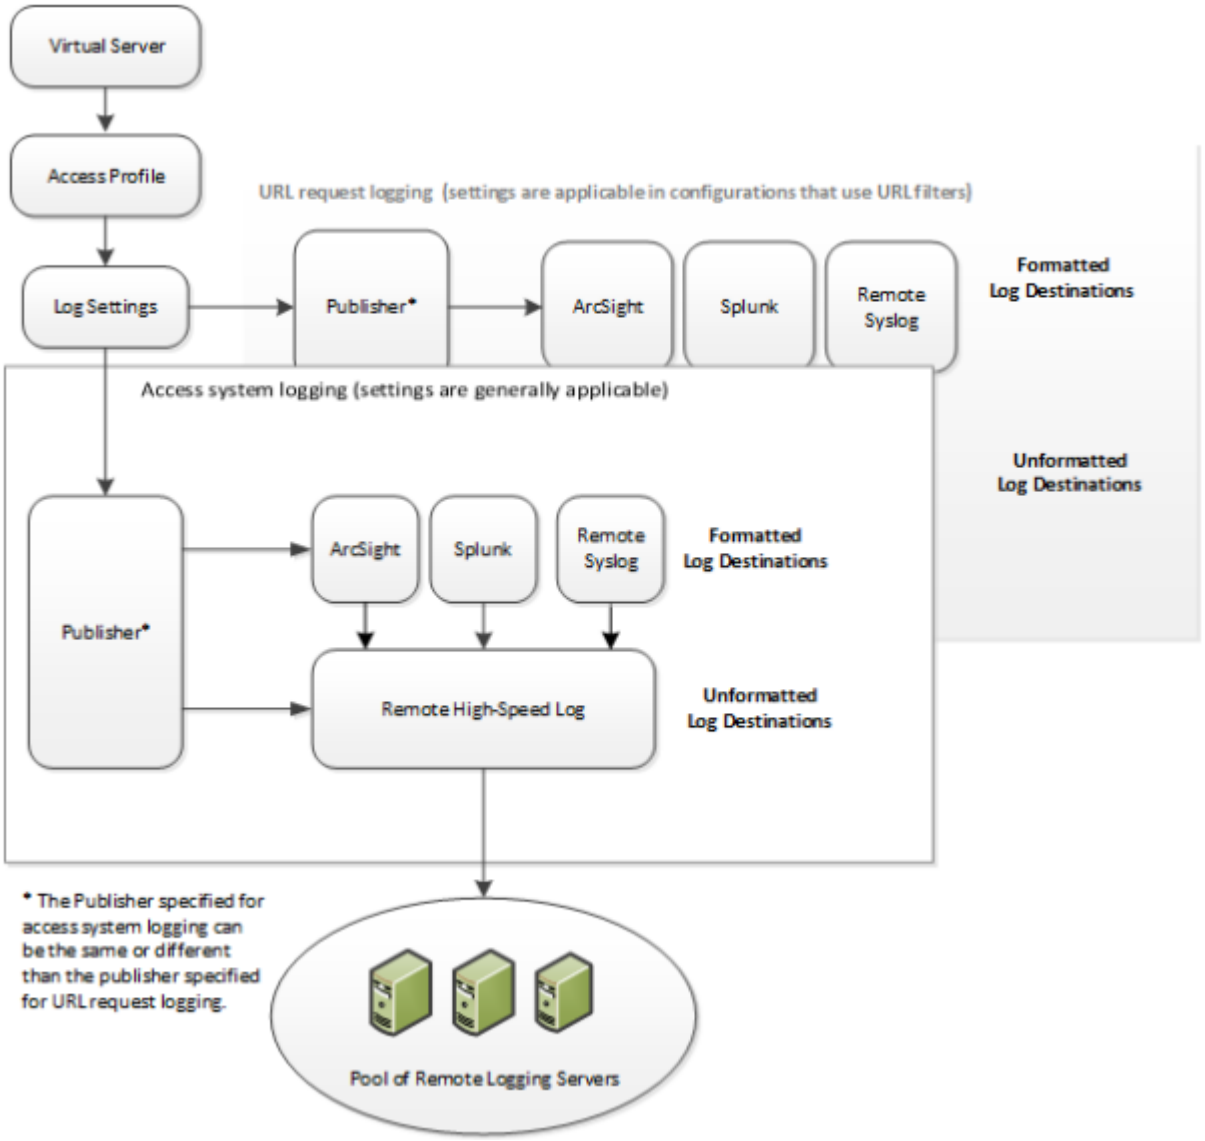

Figure 13: Association of remote high-speed logging configuration objects

Task summary

Perform these tasks to configure remote high-speed APM and SWG event logging on the BIG-IP system.

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

Task list

Creating a pool of remote logging servers
Creating a remote high-speed log destination
Creating a formatted remote high-speed log destination
Creating a publisher
Configuring log settings for access system and URL request events
Disabling logging

About the default-log-setting

Access Policy Manager® (APM®) provides a default-log-setting. When you create an access profile, the default-log-setting is automatically assigned to it. The default-log-setting can be retained, removed, or replaced for the access profile. The default-log-setting is applied to user sessions only when it is assigned to an access profile.

Regardless of whether it is assigned to an access profile, the default-log-setting applies to APM processes that run outside of a user session. Specifically, on a BIG-IP® system with an SWG subscription, the default-log-setting applies to URL database updates.

Creating a pool of remote logging servers

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add the IP address for each remote logging server that you want to include in the pool:
   a) Type an IP address in the **Address** field, or select a node address from the **Node List**.
   b) Type a service number in the **Service Port** field, or select a service name from the list.

   Note: *Typical remote logging servers require port* `514`*.*

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

Creating a remote high-speed log destination

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

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

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

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

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

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

Creating a formatted remote high-speed log destination

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

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

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

   The Splunk format is a predefined format of key value pairs.

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

   Important: *For logs coming from Access Policy Manager® (APM®), only the BSD Syslog format is supported.*
6. If you selected **Splunk**, then from the **Forward To** list, select the destination that points to a pool of high-speed log servers to which you want the BIG-IP system to send log messages.

   The Splunk format is a predefined format of key value pairs.
7. Click **Finished**.

Creating a publisher

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

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

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

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

Configuring log settings for access system and URL request events

Create log settings to enable event logging for access system events or URL filtering events or both. Log settings specify how to process event logs for the traffic that passes through a virtual server with a particular access profile.

1. On the Main tab, click **Access** > **Overview** > **Event Logs** > **Settings**.
   A log settings table screen opens.
2. Select a log setting and click **Edit** or click **Create** for a new APM® log setting.
   A popup screen opens with General Information selected in the left pane.
3. For a new log setting, in the **Name** field, type a name.
4. To specify logging, select one or both of these check box options:
   - **Enable access system logs** - This setting is generally applicable. It applies to access policies, per-request policies, Secure Web Gateway processes, and so on. When you select this check box, **Access System Logs** becomes available in the left pane.
   - **Enable URL request logs** - This setting is applicable for logging URL requests when you have set up a BIG-IP® system configuration to categorize and filter URLs. When you select this check box, **URL Request Logs** becomes available in the left pane.

   Important: *When you clear either of these check boxes and save your change, you are not only disabling that type of logging, but any changes you made to the settings are also removed.*
5. To configure settings for access system logging, select **Access System Logs** from the left pane.
   Access System Logs settings display in the right panel.
6. For access system logging, from the **Log Publisher** list select the log publisher of your choice.
   A log publisher specifies one or more logging destinations.

   Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*
7. For access system logging, retain the default minimum log level, **Notice**, for each option.
   You can change the minimum log level, but **Notice** is recommended.

Option	Description
Access Policy	Events that occur while an access policy runs.
Per-Request Policy	Events that occur while a per-request policy runs.
ACL	Events that occur while applying APM access control lists.
SSO	Events that occur during single-sign on.
Secure Web Gateway	Events that occur during URL categorization on a BIG-IP® system with an SWG subscription.
ECA	Events that occur during NTLM authentication for Microsoft Exchange clients.
OAuth	Events that occur while APM, as an OAuth authorization server, processes requests.
PingAccess Profile	Events related to PingAccess authentication. ***Important:*** *For PingAccess authentication, only the log levels defined in default-log-settings apply.*

Option	Description
VDI	Events related to connections to virtual desktop resources.
Endpoint Management System	Events related to connections to an endpoint management system.

8. To configure settings for URL request logging, select **URl Request Logs** from the left pane.
   URL Request Settings settings display in the right panel.
9. For URL request logging, from the **Log Publisher** list, select the log publisher of your choice.
   A log publisher specifies one or more logging destinations.

   Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*

10. To log URL requests, you must select at least one check box option:
    - **Log Allowed Events** - When selected, user requests for allowed URLs are logged.
    - **Log Blocked Events** - When selected, user requests for blocked URLs are logged.
    - **Log Confirmed Events** - When selected, user requests for confirmed URLs are logged.

    Whether a URL is allowed, blocked, or confirmed depends on both the URL category into which it falls, and the URL filter that is applied to the request in the per-request policy.
11. (Optional) To assign this log setting to multiple access profiles now, perform these substeps:

    Note: *Up to three log settings for access system logs can be assigned to an access profile. If you assign multiple log settings to an access profile, and this results in duplicate log destinations, logs are also duplicated.*

    a) Select **Access Profiles** from the left pane.
    b) Move access profiles between the **Available** and the **Selected** lists.

    Note: *You can delete (and add) log settings for an access profile on the Logs page for the access profile.*

    Note: *You can configure the log destinations for a log publisher from the Logs page in the System area of the product.*

12. Click **OK**.
    The popup screen closes. The table displays.

To put a log setting into effect, you must assign it to an access profile. Additionally, the access profile must be assigned to a virtual server.

Disabling logging

Disable event logging when you need to suspend logging for a period of time or you no longer want the BIG-IP® system to log specific events.

Note: *Logging is enabled by adding log settings to the access profile.*

1. To clear log settings from access profiles, on the Main tab, click **Access** > **Profiles / Policies**.
2. Click the name of the access profile.
   Access profile properties display.
3. On the menu bar, click **Logs**.
4. Move log settings from the **Selected** list to the **Available** list.

5. Click **Update**.

Logging is disabled for the access profile.

About event log levels

Event log levels are incremental, ranging from most severe (**Emergency**) to least severe (**Debug**). Setting an event log level to **Warning** for example, causes logging to occur for warning events, in addition to events for more severe log levels. The possible log levels, in order from highest to lowest severity are:

- **Emergency**
- **Alert**
- **Critical**
- **Error**
- **Warning**
- **Notice** (the default log level)
- **Informational**
- **Debug**

Note: *Logging at the* ***Debug*** *level can increase the load on the BIG-IP® system.*

APM log example

The table breaks a typical Access Policy Manager® (APM®) log entry into its component parts.

An example APM log entry

```
Feb  2 12:37:05 site1 notice tmm[26843]: 01490500:5: /Common/for_reports:Common: bab0ff52:
New session from
client IP 10.0.0.1 (ST=/CC=/C=) at VIP 20.0.0.1 Listener /Common/site1_http
(Reputation=Unknown)
```

Information Type	Example Value	Description
Timestamp	**Feb 2 12:37:05**	The time and date that the system logged the event message.
Host name	**site1**	The host name of the system that logged the event message. Because this is typically the host name of the local machine, the appearance of a remote host name could be of interest.
Log level	**notice**	The text value of the log level for the message.
Service	**tmm**	The process that generated the event.
PID	**[26843]**	The process ID.
Log ID	**01490500**	A code that signifies the product, a subset of the product, and a message number.
Level	**5**	The numeric value of the log level for the message.

Information Type	Example Value	Description
Partition	**/Common/for_reports:Common**	The partition.to which configuration objects belong.
Session ID	**bab0ff52**	The ID associated with the user session.
Log message	**New session from client IP 10.0.0.1 (ST=/CC=/C=) at VIP 20.0.0.1 Listener /Common/ site1_http (Reputation=Unknown)**	The generated message text.

About local log destinations and publishers

The BIG-IP® system provides two local logging destinations:

local-db
Causes the system to store log messages in the local MySQL database. Log messages published to this destination can be displayed in the BIG-IP Configuration utility.

local-syslog
Causes the system to store log messages in the local Syslog database. Log messages published to this destination are not available for display in the BIG-IP Configuration utility.

Note: *Users cannot define additional local logging destinations.*

The BIG-IP system provides a default log publisher for local logging, sys-db-access-publisher; initially, it is configured to publish to the local-db destination and the local-syslog destination. Users can create other log publishers for local logging.

Configuring a log publisher to support local reports

APM® provides preconfigured reports that are based on log data. To view the reports and to display log data from the BIG-IP® Configuration utility, configure a publisher to log to the local-db destination.

Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Select the log publisher you want to update and click **Edit**.
3. For the **Destinations** setting, select **local-db** from the **Available** list, and move the destination to the **Selected** list.
4. Click **Finished**.

To use a log publisher, specify it in an access policy log setting, ensure that the access profile selects the log setting, and assign the access profile to a virtual server.

Note: *Log settings are configured in the **Access** > **Overview** > **Event Log** > **Settings** area of the product.*

Viewing an APM report

If Access Policy Manager® (APM®) events are written to the local database on the BIG-IP® system, they can be viewed in APM reports.

Create a report to view event log data.

1. On the Main tab, click **Access** > **Overview** > **Access Reports**.
   The Reports Browser displays in the right pane. The Report Parameters popup screen opens and displays a description of the current default report and default time settings.
2. (Optional) Select the appropriate **Restrict by Time** settings.
3. Click **Run Report**.
   The popup screen closes. The report displays in the Reports Browser.

You can select and run various system-provided reports, change the default report, and create custom reports.

Viewing URL request logs

To view URL request logs from the user interface, your access profile log setting must enable URL request logs. The log setting must also specify a log publisher that publishes to the local-db log destination.

You can display, search, and export URL request logs.

1. On the Main tab, click **Access** > **Overview** > **Event Logs** > **URL Request Logs**.
   Any logs for the last hour are displayed.

   Note: *APM® writes logs for blocked requests, confirmed requests, allowed requests, or all three, depending on selections in the access profile log setting.*
2. To view logs for another time period, select it from the list.
3. To search the logs, type into the field and click **Search** or click **Custom Search** to open a screen where you can specify multiple search criteria.
4. To export the logs for the time period and filters, click **Export to CSV**.

Configuring a log publisher to supply local syslogs

If you must have syslog files available on the local device, configure a publisher to log to the local-syslog destination.

Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Select the log publisher you want to update and click **Edit**.
3. For the **Destinations** setting, select **local-syslog** from the **Available** list, and move the destination to the **Selected** list.
4. Click **Finished**.

To use a log publisher, specify it in an access policy log setting, ensure that the access profile selects the log setting, and assign the access profile to a virtual server.

Note: *Log settings are configured in the **Access** > **Overview** > **Event Log** > **Settings** area of the product.*

Preventing logging to the /var/log/apm file

To stop logs from being written to the /var/log/apm file, remove the local-syslog destination from log publishers that are specified for access system logging in APM® log settings.

Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Select the log publisher you want to update and click **Edit**.
3. For the **Destinations** setting, if the **Selected** list contains **local-syslog**, move it to the **Available** list.
4. Click **Finished**.

To use a log publisher, specify it in an APM log setting, ensure that the log setting is assigned to an access profile, and assign the access profile to a virtual server.

Note: *Log settings are configured in the **Access** > **Overview** > **Event Log** > **Settings** area of the product.*

About local log storage locations

The BIG-IP® system publishes logs for portal access traffic and for connections to virtual desktops (VDI) to the `/var/log/rewrite*` files. APM® cannot publish these logs to remote destinations.

APM can publish URL request logs to remote or local destinations. Logs published to the local-db destination are stored in the local database and are available for display from the Configuration utility. Logs published to the local-syslog destination are stored in the `/var/log/urlfilter.log` file.

APM can publish access system logs to remote or local destinations. Logs published to the local-db destination are stored in the local database. Logs in the local database are available for display in APM reports. Logs published to the local-syslog destination are stored in the `/var/log/apm` file.

Code expansion in Syslog log messages

The BIG-IP® system log messages contain codes that provide information about the system. You can run the Linux command `cat` *`log`* `|bigcodes |less` at the command prompt to expand the codes in log messages to provide more information. For example:

```
   Jun 14 14:28:03 sccp bcm56xxd [ 226 ] : 012c0012 : (Product=BIGIP Subset=BCM565XXD) :
6: 4.1 rx [ OK 171009 Bad 0 ] tx [ OK 171014 Bad 0 ]
```

About log level configuration

Log levels can be configured in various ways that depend on the specific functionality. Log levels for access portal traffic are configured in the System area of the product. The log level for the URL database download is configured in the default-log-setting in the **Access** > **Overview** > **Event Log** > **Settings** area of the product. The log level for NTLM authentication of Microsoft Exchange clients is configured using the ECA option in any log setting. Other access policy (and Secure Web Gateway) log levels are configured in any log setting.

Updating the log level for NTLM for Exchange clients

Before you follow these steps, you must have an access profile that you configured to use for NTLM authentication of Microsoft Exchange clients. You must know the name of the log setting that is assigned to that access profile. (The default-log-setting is assigned by default, but your access profile configuration might be different.)

You can change the level of logging for NTLM authentication for Microsoft Exchange clients.

Note: *Logging at the default level,* ***Notice****, is recommended.*

1. On the Main tab, click **Access** > **Overview** > **Event Logs** > **Settings**.
   A log settings table screen opens.
2. Select the check box for the log setting that you want to update and click **Edit**.
   A popup screen opens.
3. To configure settings for access system logging, select **Access System Logs** from the left pane.
   Access System Logs settings display in the right panel.
4. For the **ECA** setting, select a log level.

   Note: *Setting the log level to* ***Debug*** *can adversely impact system performance.*
5. Click **OK**.
   The popup screen closes.

Configuring logging for the URL database

Configure logging for the URL database so that log messages are published to the destinations, and at the minimum log level, that you specify. (Logging for the URL database occurs at the system level, not the session level, and is controlled using the default-log-setting log setting.)

Note: *A URL database is available only on a BIG-IP® system with an SWG subscription.*

1. On the Main tab, click **Access** > **Overview** > **Event Logs** > **Settings**.
   A log settings table screen opens.
2. From the table, select **default-log-setting** and click **Edit**.
   A log settings popup screen displays.
3. Verify that the **Enable access system logs** check box is selected.
4. To configure settings for access system logging, select **Access System Logs** from the left pane.
   Access System Logs settings display in the right panel.
5. From the **Log Publisher** list, select the log publisher of your choice.

   A log publisher specifies one or more logging destinations.

   Important: *The BIG-IP® system is not a logging server and has limited capacity for storing, archiving, and analyzing logs. For this reason a dedicated logging server is recommended.*
6. To change the minimum log level, from the **Secure Web Gateway** list, select a log level.

   Note: *Setting the log level to* ***Debug*** *can adversely impact system performance.*

   The default log level is **Notice**. At this level, logging occurs for messages of severity Notice and for messages at all incrementally greater levels of severity.
7. Click **OK**.
   The popup screen closes. The table displays.

Setting log levels for Portal Access events

Change the logging level for access policy events when you need to increase or decrease the minimum severity level at which Access Policy Manager® (APM®) logs that type of event. Follow these steps to change the log level for events that are related to portal access traffic.

Note: *You can configure log levels for additional APM options in the Event Logs area.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Options**.
2. Scroll down to the Access Policy Logging area.

   Note: *The log settings that you change on this page impact only the access policy events that are logged locally on the BIG-IP® system.*

3. For **Portal Access**, select a logging level from the list.

   Warning: *F5® recommends that you do not set the log level for* ***Portal Access*** *to* ***Debug****. Portal Access can stop working. The BIG-IP system can become slow and unresponsive.*

4. Click **Update**.

APM starts to log events at the new minimum severity level.

Returning HTTP Status 503 to Web Applications

Overview: Returning HTTP status code 503 for APM-generated error pages

Access Policy Manager® (APM®) sends an error page to a web application when, for example, the application sends a request to a non-existent domain name or uses a URL that is blocked by the configuration. By default, the error page is a human-readable HTML page with HTTP status 200. This can lead to errors if the web application expects an error status or error notifications, but receives a normal reply. Instead, proxy servers can reply with HTTP status 503. Current browsers display the content of HTTP 503 replies to the user, if possible, and client code can recognize an abnormal situation.

To change the default behavior to reply with HTTP status code 503 for an APM-generated error page, you can configure an access profile or an access policy.

- To change the default behavior for any client session created through the access profile, configure a setting in the access profile.
- To change the default behavior for a client session based on any criteria that you want, configure logic (and a session variable) in the access policy.

Configuring an access profile to return HTTP status code 503

For Access Policy Manager® (APM®) to return an HTTP status code 503 for APM-generated error pages, you can select **Use HTTP Status 503 for Error Pages** in the access profile.

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. Click the name of the access profile that you want to edit.
   The properties screen opens.
3. For **Settings**, select the **Custom** check box.
   The settings become available to edit.
4. Select the **Use HTTP Status 503 for Error Pages** check box.
5. Click **Finished**.

For APM-generated error pages, APM returns HTTP status code 503 instead of HTTP 200 for traffic on the virtual server associated with this access profile.

Configuring an access policy to return HTTP status code 503

You can set the value of the `session.error.use503` session variable in an access policy to change the HTTP status code that Access Policy Manager® (APM®) returns for error pages that APM generates. If you enable the session variable, APM responds with HTTP status 503. If you disable the session variable later on the branch, APM goes back to returning HTTP status 200, which is the default behavior.

Note: *If the* ***Use HTTP Status 503 for Error Pages*** *check box is selected in the access profile, changing the value of the* `session.error.use503` *session variable in an access policy has no effect.*

---

1. On the Main tab, click **Access** > **Profiles / Policies**.
   The Access Profiles (Per-Session Policies) screen opens.
2. In the Per-Session Policy column, click the **Edit** link for the access profile you want to configure.
   The visual policy editor opens the access policy in a separate screen.

3. Add items to the policy to specify the criteria under which you want APM to return HTTP status 503.
4. On the policy branch where you want APM to return HTTP status 503, click the **(+)** icon to add an item to the policy.
   A popup screen displays actions on tabs, such as General Purpose and Authentication, and provides a search field.
5. On the Assignment tab, select **Variable Assign** and click **Add Item**.
   A properties screen opens.
6. Click **Add new entry**.
   An **Empty** entry displays.
7. Click the **change** link next to the empty entry.
   A dialog box, where you can enter a variable and an expression, opens.
8. On the left side, retain the default setting, **Custom Expression**, and in the field type `session.error.use503`.
9. On the right side, select **text** from the list and set the value of the session variable:
   - To use HTTP status 503 for APM-generated error pages, type any letter or word in the field, such as `t`.
   - To restore the default behavior (use the HTTP status 200 for error pages), remove all text from the field.
10. Click **Finished** to save the variable and expression and return to the Variable Assign action popup screen.
11. Click **Save**.
    The properties screen closes and the policy displays.
12. Click the **Apply Access Policy** link to apply and activate the changes to the policy.

Resources and Documentation

Additional resources and documentation for BIG-IP Access Policy Manager

You can access all of the following BIG-IP® system documentation from the AskF5™ Knowledge Base located at `http://support.f5.com/`.

Document	Description
BIG-IP® Access Policy Manager®: Application Access	This guide contains information for an administrator to configure application tunnels for secure, application-level TCP/IP connections from the client to the network.
BIG-IP® Access Policy Manager®: Authentication and Single-Sign On	This guide contains information to help an administrator configure APM for single sign-on and for various types of authentication, such as AAA server, SAML, certificate inspection, local user database, and so on.
BIG-IP® Access Policy Manager®: Customization	This guide provides information about using the APM customization tool to provide users with a personalized experience for access policy screens, and errors. An administrator can apply your organization's brand images and colors, change messages and errors for local languages, and change the layout of user pages and screens.
BIG-IP® Access Policy Manager®: Edge Client and Application Configuration	This guide contains information for an administrator to configure the BIG-IP® system for browser-based access with the web client as well as for access using BIG-IP Edge Client® and BIG-IP Edge Apps. It also includes information about how to configure or obtain client packages and install them for BIG-IP Edge Client for Windows, Mac, and Linux, and Edge Client command-line interface for Linux.
BIG-IP® Access Policy Manager®: Implementations	This guide contains implementations for synchronizing access policies across BIG-IP systems, hosting content on a BIG-IP system, maintaining OPSWAT libraries, configuring dynamic ACLs, web access management, and configuring an access policy for routing.
BIG-IP® Access Policy Manager®: Network Access	This guide contains information for an administrator to configure APM Network Access to provide secure access to corporate applications and data using a standard web browser.
BIG-IP® Access Policy Manager®: Portal Access	This guide contains information about how to configure APM Portal Access. In Portal Access, APM communicates with back-end servers, rewrites links in application web pages, and directs additional requests from clients back to APM.
BIG-IP® Access Policy Manager®: Secure Web Gateway	This guide contains information to help an administrator configure Secure Web Gateway (SWG) explicit or transparent forward proxy and apply URL categorization and filtering to Internet traffic from your enterprise.
BIG-IP® Access Policy Manager®: Third-Party Integration	This guide contains information about integrating third-party products with Access Policy Manager (APM®). It includes

Document	Description
	implementations for integration with VMware Horizon View, Oracle Access Manager, Citrix Web Interface site, and so on.
BIG-IP® Access Policy Manager®: Visual Policy Editor	This guide contains information about how to use the visual policy editor to configure access policies.
Release notes	Release notes contain information about the current software release, including a list of associated documentation, a summary of new features, enhancements, fixes, known issues, and available workarounds.
Solutions and Tech Notes	Solutions are responses and resolutions to known issues. Tech Notes provide additional configuration instructions and how-to information.

Legal Notices

Legal notices

Publication Date

This document was published on November 15, 2017.

Publication Number

MAN-0508-05

Copyright

Trademarks

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

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

Patents

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

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

B

C

D

E

F

G

H

I

L

M

N

O

P

R

S

T

U

V

W

X