F5® BIG-IQ® Centralized Management: Access

Version 5.1

Table of Contents

BIG-IQ Configuration Management: Access Overview

About Access

The BIG-IQ® system offers you centralized management for BIG-IP ®Access Policy Manager® (APM) and F5 Secure Web Gateway (SWG) configurations. Centralized management gives you easy-to-deploy sets of access policies, and access policy configuration objects. This means you don't need to repeat the configuration on each BIG-IP system individually. Access also offers you centralized reporting, which allows you to compare and monitor BIG-IP APM® usage across many groups of devices.

Access configuration workflow

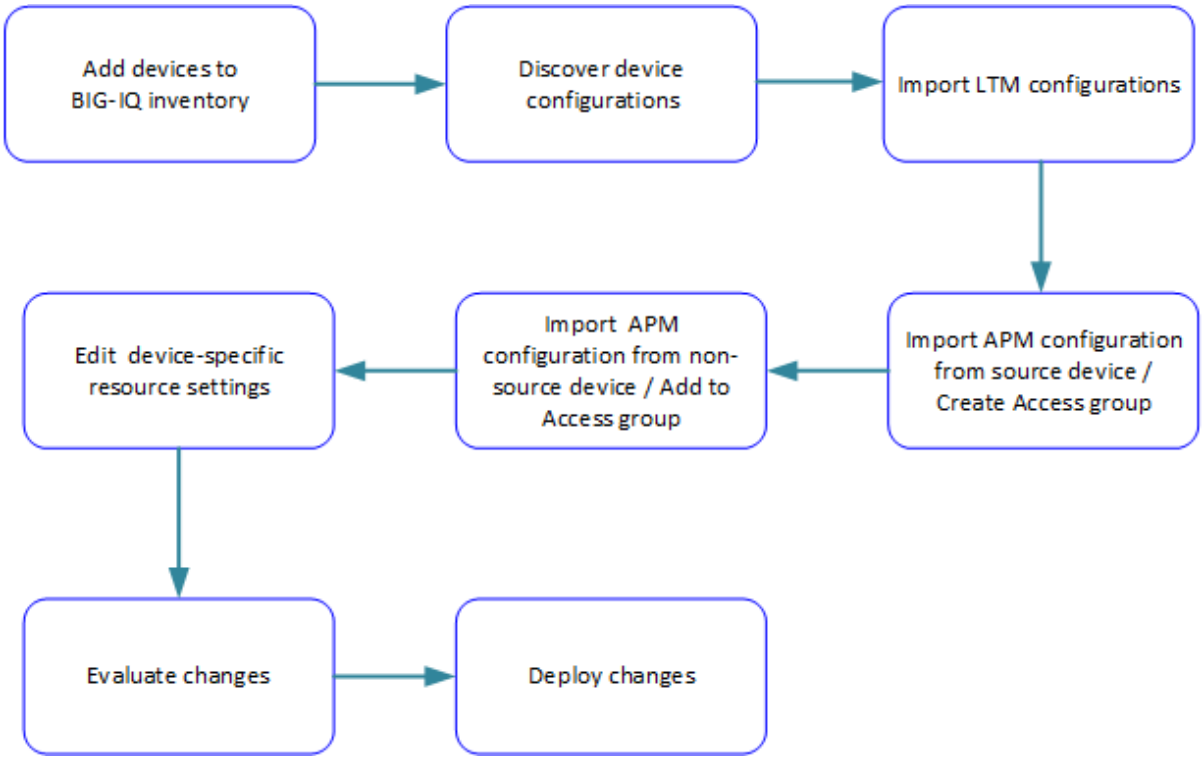

What are Access groups?

Each *Access group* is a group of BIG-IP® devices across which you plan to share the same Access configuration. When you import an APM service configuration from a device, the device must join an Access group.

What is a source device?

A *source device* is the foundation of the shared configuration for other devices in an Access group.

What is a non-source device?

Any *non-source device* is a member of an Access group that accepts the shared configuration from the source device.

About the types of resources that Access imports

When you import an APM® service configuration from a device, the device must join an Access group.

- If the device joins a new Access group, the device becomes the source of the shared configuration for the group; Access imports both shared resources and device-specific resources from the source device.
- If the device joins an existing Access group, Access imports only the device-specific resources from the device.

About shared resources

In an Access group on the BIG-IQ® system, *shared resources* are a set of configuration objects that are expected to be the same on every device in an Access group.

How do shared resources work in the configuration?

Initially, shared resources are imported with the APM® service configuration from the source device. After import, they are read-only on the BIG-IQ® system. The deployment process configures the shared resources on all non-source devices in the Access group. This can result in major configuration changes on the non-source devices, with resources being overwritten, deleted, or added on them.

About device-specific resources

In an Access group on the BIG-IQ® system, *device-specific resources* are a set of configuration objects that are expected to exist on every device in the Access group. However, the properties of these resources can differ from device to device.

For example, an access policy could use an Active Directory server for user authentication. Device `apm_north_america.xyz.com` must use an Active Directory server configured in a North American domain or data center, while device `apm_south_america.xyz.com` must use an Active Directory server configured in a South American domain or data center.

How do device-specific resources work in the configuration?

When you add a device to an Access group, device-specific resources are created from the device's APM® service configuration. Or, if particular resources do not exist on a non-source device, Access creates device-specific resources that match those in the source device configuration. After import, you are instructed to review and change device-specific resources if needed; in addition, you can change them at your option. You can also make a device-specific resource shared, so that its properties can only be configured in the shared resources. At deployment, device-specific resources are configured on the specific devices.

Reporting configuration workflow

BIG-IQ logging nodes are required for Access and SWG reporting. To set up a discovered device so that it sends report data to a logging node, you must run the remote logging configuration. Then, you can run reports.

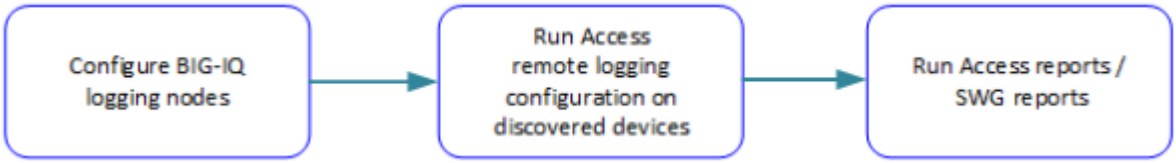

Users, User Groups, and Roles

About users, user groups, and roles

A *user* is an individual to whom you provide resources. You provide access to users for specific BIG-IQ® system functionality through authentication. You can associate a user with a specific role, or associate a user with a user group and then associate the group with a role.

A *role* is defined by its specific privileges.

A *user group* is a group of individuals who have access to the same resources. When you associate a role with a user or user group, that user or user group is granted all of the role's corresponding privileges.

User roles in the Access configuration workflow

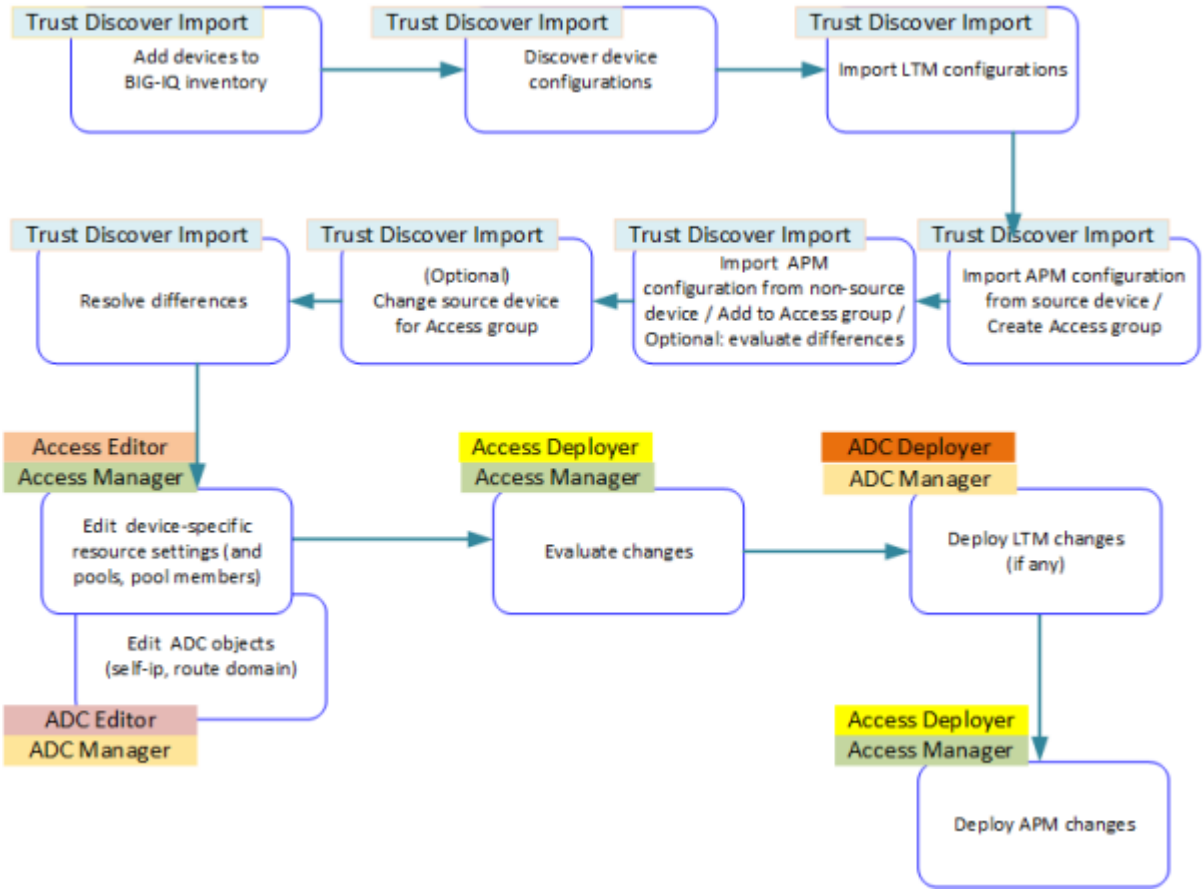

Figure 1: Access configuration workflow with possible user roles

User roles in the reporting configuration workflow

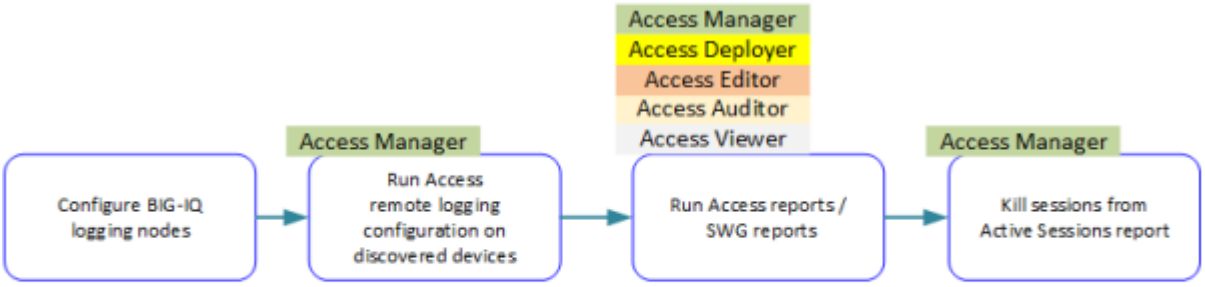

Figure 2: Reporting workflow with possible user roles

Adding a BIG-IQ user

Create a user to provide access to the BIG-IQ system.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, click **USER MANAGEMENT** > **Users**.
4. Click the **Add** button.
5. In the **User Name** field, type the user name for this new user.
6. In the **Full Name** field, type a name to identify this user.
   The full name can contain a combination of symbols, letters, numbers and spaces.
7. In the **Password** and **Confirm Password** fields, type the password for the new user.
8. To associate this user with an existing user group, select the group from the **User Groups** list.
   To associate the user with additional groups, click the plus [+] icon and select another group.
9. From the **User Roles** list, select a user role to associate with this user.
   Each role has a set of unique privileges.
   To associate the user with additional roles, click the plus [+] icon and select another role.
10. Click the **Save** button at the bottom of the screen.

Creating a user group

You create a user group to offer individual users access to the same resources.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the left, click **USER MANAGEMENT** > **User Groups**.
   The User Groups screen opens.
4. Click **Add**.

5. In the **Name** field, type a name for this new user group.
6. From the **Auth Provider** list, select **local (Local)**.
7. From the **Users** list, select a user to add to this group.
   To add additional users, click the plus [+] icon.
8. From the **User Roles** list, select a user role to give its associated system access to this user group.
   To add additional roles, click the plus [+] icon.
9. Click the **Save** button at the bottom of the screen.

User role access descriptions

The table lists standard BIG-IQ® system user roles you might need to assign to your users, depending on their responsibilities in working with Access.

Role	Role Description / Access
Access Auditor	This role provides access to BIG-IQ® Access reports.
Access Deployer	This role has deploy access to Access configuration objects. This role cannot discover and edit devices or policies.
Access Editor	This role has edit access to Access configuration objects. This role cannot discover and deploy devices or policies. This role includes the ability to add, update, and delete pools and pool members from the Access configuration object editor.
Access Manager	This role has deploy and edit access to Access configuration objects, and has access to Access Reports and Dashboard. This role cannot add or remove devices and device groups, and cannot discover, import or delete services.
Access Viewer	This role has view-only access to Access configuration objects and tasks for Access devices that have been discovered. This role cannot edit, discover, or deploy devices or policies.
ADC Deployer	This role has deploy access to ADC configuration objects. This role cannot discover and edit devices or configuration objects. At deployment, Access notifies you if it finds changes in ADC that you must deploy first,
ADC Editor	This role has edit access to ADC configuration objects. A user needs this role to be able to edit or create a self-IP address or a route domain and to view other ADC configuration objects. This role includes the ability to add, update, and delete pools and pool members from ADC; however, you can configure pools and pool members within Access without having this role.
ADC Manager	This role manages the ADC module with full privilege. This role works for a user who needs to: Deploy ADC; edit or create a self-IP address or a route domain; view other ADC configuration objects. This role includes the ability to add, update, and delete pools and pool members from ADC; however, you can configure pools and pool members within Access without having this role.
ADC Viewer	This role permits read-only access to the ADC module. A user who needs to view configuration objects from ADC needs this role.
Trust Discover Import	This role can add and delete devices, discover services and import them, and remove services.
Administrator	This role has access to all aspects of the BIG-IQ system, which can include BIG-IQ Security, BIG-IQ System, and BIG-IQ ADC management. This access includes areas involved in adding individual users, assigning roles,

Role	Role Description / Access
	device discovery, installing updates, activating licenses, and configuring a BIG-IQ high availability (HA) configuration.

BIG-IP Devices, HA Pairs, and Clusters

Preliminary tips for putting an Access group together

As you start to think about how to group BIG-IP® devices into Access groups that share a configuration, there are a few things you might want to keep in mind. When you select the source device for an Access group, you are selecting the shared configuration for all of the devices in the group. (You can change the source device if needed.)

When you add BIG-IP devices to an Access group, Access evaluates the differences between the source device and the other devices in the group. Access reports the differences for your information. If you need to make configuration changes on any of the devices, Access lets you know which non-source device to change, and which object to update, delete, or add.

Things to know about machine accounts

Machine accounts support Microsoft Exchange clients that use NTLM authentication. An NTLM Auth Configuration object refers to a machine account. If the APM® configurations on the BIG-IP® systems include machine accounts, you might want to be aware of the following information.

In an Access group, the machine accounts on the source and non-source devices must each have been created with the same name. If this is not the case, the deployment fails. The deployment differences will include the names of the devices on which you must reconfigure the machine accounts before you can successfully deploy.

Things to know about bandwidth controller configurations

On a BIG-IP® device, bandwidth controller configuration objects (policies and priority groups) are configured at the system level. In APM ®, they are used to provide traffic shaping for Citrix clients that support MultiStream ICA. In an access policy, a *BWC policy* item refers to a bandwidth controller policy. If the APM configurations on the BIG-IP systems refer to bandwidth controller objects, you should be aware of the following information.

The bandwidth controller configuration objects on the source device are treated as if they were part of the Access shared configuration. That means when you import the APM service configuration from a source device, the bandwidth controller objects are imported and cannot be updated in the BIG-IQ® system. When you deploy the configuration, deployment creates the bandwidth controller objects on the non-source devices.

Access requirements for HA pairs and clusters

For BIG-IP® system high availability, APM ® supports two devices in a Sync-Failover group; these devices can also be referred to as an *HA pair*.

Access has these requirements for HA pairs on BIG-IQ® system configuration:

- If you import a device that is part of an HA pair, you must import the other device in the pair as well. Access must manage the configuration for both devices.

- When you import the devices that are an HA pair, you must place both devices in a cluster that contains only that pair.

  Note: *This is not enforced when you add devices to a cluster. But when you try to deploy the configuration, Access reports errors and deployment fails.*

- When you add devices to an Access group, you must add both members of a cluster to the same Access group. (You can add all clusters to one Access group or add clusters to multiple Access groups.)

  Note: *Access enforces this requirement.*

To avoid problems after you create Access configurations on the BIG-IQ system, you should know which devices constitute each HA pair.

Important: *F5® recommends that you make a list of HA pairs, and keep it available for ready reference while you work in the BIG-IQ system.*

Managing Access Groups

How do I start to centrally manage APM configurations from BIG-IQ?

Here is an overview of your first steps for setting up an Access Policy Manager® (APM®) configuration once, and then being able to deploy that configuration from the BIG-IQ® system to other BIG-IP® devices.

Step 1. Add the BIG-IP device to the inventory list on the BIG-IQ system. You enter the IP address and credentials of the BIG-IP device you're adding, and associate it with a cluster (if applicable.

Step 2. Discover the APM and the Local Traffic Manager™ (LTM) configurations. You must discover LTM first, because APM uses some resources that are managed by LTM.

Step 3. Import the LTM configuration into the BIG-IQ system.

Step 4. Import the APM configuration into the BIG-IQ system. Importing the APM configuration requires that the device be added to an Access Group. You can create a new Access Group with the device as source-device, or you can add the device to another Access Group as non-source device.

What is the best way to create an Access group?

After you add devices to the BIG-IQ® system and discover them, you can create an Access group in either of two ways. Use whichever you prefer, based on your requirements.

- From the Access user interface, you can add multiple devices to an Access group at once. Using this method, you select multiple devices, with one device specified as the source device. Access then imports configurations from the devices, and creates the Access group.
- From the Device Management user interface, you can add one device at a time to an Access group when you import the APM service from each device.

Adding devices to the BIG-IQ inventory

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

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

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

Note: *Access supports BIG-IP system software version 12.1 only.*

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

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. Click the **Add Device** button.
4. In the **IP Address** field, type the IPv4 or IPv6 address of the device.
5. In the **User Name** and **Password** fields, type the user name and password for the device.
6. To add this device to a new cluster:

   Important: *If a device is not a member of a Sync-Failover group that you configured to support an Active-Standby configuration for APM, do not add it to a cluster.*

   If the device is the first member of a Sync-Failover group that you have added to the BIG-IQ system, add it to a new cluster. It does not matter whether this device is the Active or the Standby member of the group.

   a) From the **Cluster Display Name** list, select **Create New**, and then type a new name for this new cluster.

      A cluster name must be unique on the BIG-IQ system. It does not need to match the name of the Sync-Failover group on the BIG-IP device. However, ensuring some similarity between the names might be useful to you, because when you add the second member of the group, you must add it to the same cluster.

   b) Select an option from the **Deployment Settings**:

   - **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option to prompt BIG-IQ to start the DSC synchronization process so that any configuration change made to this device is synchronized with other members of the DSC. This option makes sure all members of the DSC have the most current configuration.
   - **Ignore BIG-IP DSC sync when deploying configuration changes** Select this option to have BIG-IQ deploy any configuration changes for this device to all cluster members. Use this option only if this device is not configured in a DSC Sync-Failover device group, or if any members of the cluster are disabled.

7. To add this device to an existing cluster:

   If the device is the second member of a Sync-Failover group that you have added to the BIG-IQ system, add the device to the existing cluster for that Sync-Failover group.

   a) From the **Cluster Display Name** list, select **Use Existing**, and then select the cluster from the list.
   b) Select an option from the **Deployment Settings**:

   - **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option to prompt BIG-IQ to push any configuration changes to this device to other members of the DSC. This option makes sure all members of the DSC have the most current configuration.
   - **Ignore BIG-IP DSC sync when deploying configuration changes** Select this option to have BIG-IQ deploy any configuration changes for this device to all cluster members. Use this option only if this device is not configured in a DSC Sync-Failover device group, or if any members of the cluster are disabled.

8. Click the **Add** button at the bottom of the screen.
   The BIG-IQ system opens communication to the BIG-IP device, and checks its framework.

   Note: *The BIG-IQ system can properly manage a BIG-IP device only if the BIG-IP device is running a compatible version of the REST framework.*

9. Click the **Add** button at the bottom of the screen.
   When complete, a popup screen displays a status and options to discover service configurations immediately.

10. To discover configurations for APM® and LTM® now, select **Access Policy Manager (APM)**, and the **Local Traffic Manager (LTM)** check box is selected automatically; click **Discover**.

    You can discover service configurations now or do it later.

    BIG-IQ discovers the configurations for the APM and LTM services.

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

Discovering the LTM and APM service configurations

Before you can import configurations from a device, you must first discover them. To prepare to create an Access configuration on the BIG-IQ ®system, you must discover the Local Traffic Manager™ (LTM®) service configuration, and then discover the Access Policy Manager® (APM) service configuration.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. Click the name of the device you want to discover the service configuration from.
4. On the left, click **Services**.
5. For Local Traffic Manager (LTM), click **Discover**.

   You must wait for discovery to complete before you continue.
6. For Access Policy Manager (APM), click **Discover**.

Importing the LTM service configuration

You must discover a service configuration before you can import it.

Before you can import the Access Policy Manager® (APM) service configuration from a discovered device, you must import the Local Traffic Manager™ (LTM®) service configuration.

Important: *You, or any other BIG-IQ system user, cannot perform any tasks on the BIG-IQ system while it is importing a service configuration. Large configurations can take a while to import, so let other BIG-IQ users know before you start this task.*

---

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. Click the name of the device you want to import the service configuration from.
4. On the left, click **Services**.
5. For Local Traffic Manager (LTM), select the **Create a snapshot of the current configuration before importing** check box to save a copy of the device's current configuration.

   You're not required to create a snapshot, but it is a good idea in case you have to revert to the previous configuration for any reason.
6. For Local Traffic Manager (LTM), click **Import**.

The LTM service configuration is imported.

Importing the APM configuration into an Access group

You must discover a service configuration before you can import it.

You import Access Policy Manager® (APM) configuration objects from a device to manage the device configuration from the BIG-IQ® system. As part of the import process, you select an Access group.

Important: *You, or any other BIG-IQ system user, cannot perform any tasks on the BIG-IQ system while it is importing a service configuration. Large configurations can take a while to import, so let other BIG-IQ users know before you start this task.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. Click the name of the device you want to import the service configuration from.
4. On the left, click **Services**.
5. For Access Policy (APM), select the **Create a snapshot of the current configuration before importing.** check box to save a copy of the device's current configuration.

   You're not required to create a snapshot, but it is a good idea in case you have to revert to the previous configuration for any reason.
6. For Access Policy (APM), click **Import**.
7. On the Add to Access Group popup screen, specify either a new or existing Access group:
   - Select **Create New**, in the **Name** field, type a name, and click **Add**.
   - Select **Add to existing**, select a name from the **Name** list, and click **Add**.

   Important: *You must add both members of an HA pair to the same Access group.*

   The device in the **Group Source Device** provides the shared configuration for all devices in the Access group.

   If you add the device to a new Access group, it becomes the source device; its shared resources and device-specific resources are imported. If you add the device to an existing Access group, it becomes a non-source device; its device-specific resources are imported.

The APM service configuration is imported.

Creating an Access group

Before you can create an Access group, you must have at least one device discovered. You must have imported the LTM® service configuration from a device before you can add that device to an Access group

You create an Access group to start to manage the Access configuration for a group of devices.

Note: *When you create an Access group, the service configurations for the devices are imported.*

Important: *You, or any other BIG-IQ system user, cannot perform any tasks on the BIG-IQ system while it is importing a service configuration. Large configurations can take a while to import, so let other BIG-IQ users know before you start this task.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. Click the **Create** button.
   The New Group screen opens.
4. In the **Name** field, type a name for the Access group.
5. From **Source Device**, select the device to be the source of the shared configuration for other devices in the group.
6. If there are devices in the **Managed BIG-IP APM Devices** setting that you want to add to the group now, move them to the **Selected** list.

7. Click **Create & Import**.
   The Access Groups screen opens. Progress information displays in the Status column.
8. If the system discovers differences between a source device and non-source devices, you can see them by clicking the **View Differences** link in the Status column.

Adding a device to an existing Access group

Before you start, you must have at least one device with the APM® service discovered. You must also have imported the LTM® service configuration from the device before you can add that device to an Access group

You add a device to an Access group so you can manage its configuration from Access. When you add a device to an existing Access group, its device-specific configuration resources are imported into Access. Access also creates any device-specific resources that it is missing, from the source device configuration.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. Click the name of the Access group you want to change.
   The properties screen for that group opens, listing the devices in the Access group.
4. Click **Add**.
   An Add Devices popup screen opens.
5. Move the devices you want to add to the **Selected** list.
6. Click **Select**.
   The popup screen closes, showing the Access Groups screen. Progress information displays in the Status column.
7. If the system discovers differences between a source device and non-source devices, you can see them by clicking the **View Differences** link in the Status column.

Changing the source device for an Access group

You might need to make a change when the existing source device is going to be decommissioned. Or, you might do this if the source device is down and a configuration change must be made and deployed in an emergency.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. Click the name of the Access group that you want to change.
   The properties screen for that group opens, listing the devices in the Access group.
4. Select a non-source device.

   An asterisk marks the name of the source device
5. Click **Make Source**.
   A screen displays, prompting you to confirm the source change.
6. Click **Save**.
   The Access Group screen displays while the shared configuration is imported from the newly selected source device. Access evaluates the configuration.
7. If the **Status** field shows that differences were found, click the **View Differences** link to review them, and accept or deny the changes:

   - Click **Accept** to update the Access group with the configuration changes.
   - Click **Deny** to not update the Access group with the configuration changes.

Removing a device from an Access group

You remove a device from an Access group if you no longer want to manage the Access configuration for the device, or if you want to add the device to a different Access group.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. Click the name of the Access group you want to change.
   The properties screen for that group opens, listing the devices in the Access group.
4. Select the check box for that device and click **Remove**.
   A confirmation popup screen opens.
5. Confirm that you want to remove the device.
   The device no longer displays in the Access group. The APM service configuration on the device is no longer managed.

Before you can see new data from the device in Access reports or add the device to another Access group, you must discover the APM service configuration on the device.

Viewing and Editing the Access Configuration

Finding a device-specific resource

In BIG-IQ® Access, you can find a device-specific resource on the left, under **Shared resources** if it has been marked `shared`, or under the specific device to which it belongs.

1. To search for a resource among the shared resources, on the left expand the Access group name and click **Shared resources**.
2. In the Filter field, type all or part of the name of the object, and press Enter.
   The screen displays each shared object type, with the number of matching resources it has found, marked in parentheses. For example, **Access Profiles list (1)**, **CAPTCHA Configuration List (0)**, and so on.
3. To search among device-specific resources, on the left expand the Access group name, click the name of a device, then use the **Filter** field to sort the resources.
4. If you do not know the name of the resource you want to find, to find it you must browse through the shared resource types and device-specific resource types for non-source devices.

Device-specific resources: User interface

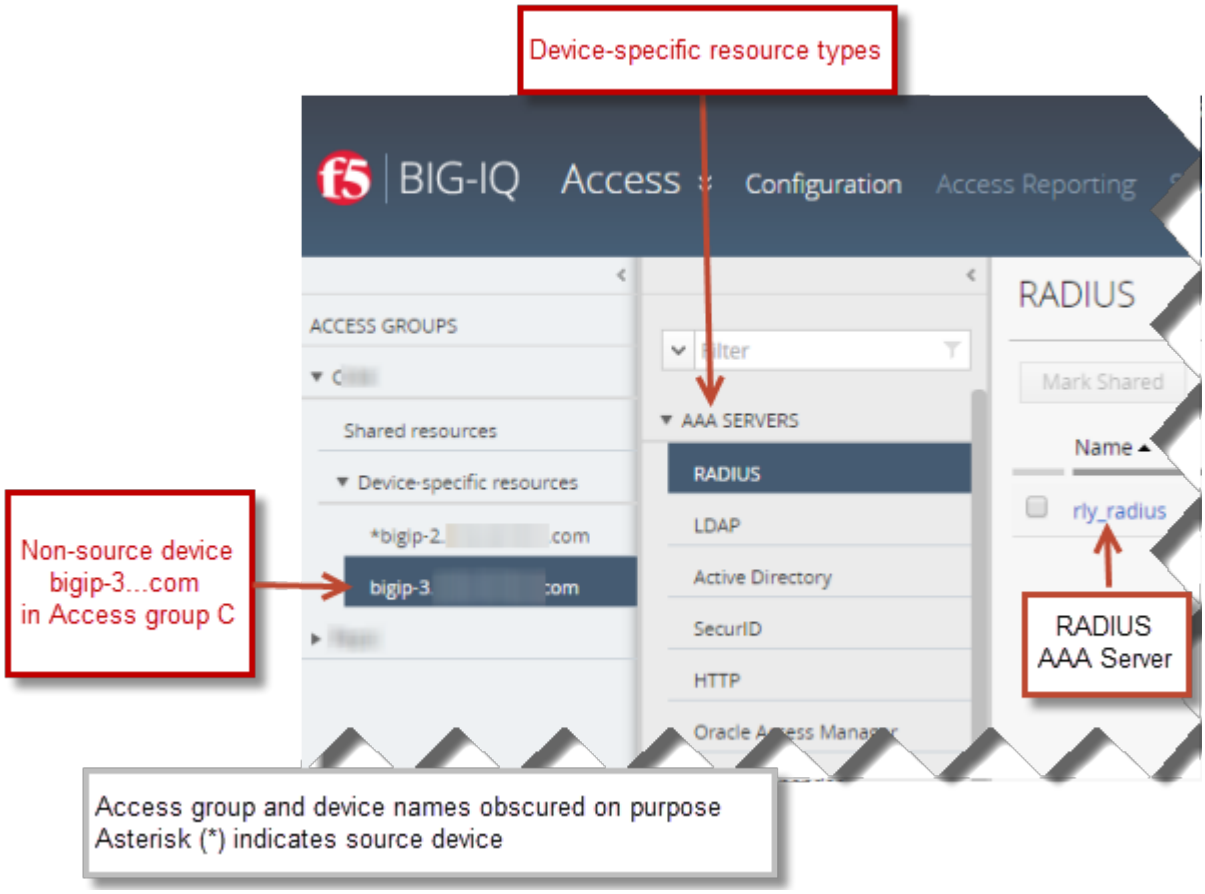

Figure 3: Device-specific resources for a device in an Access group

Shared resources: User interface

Figure 4: Shared resources in an Access group

Editing a device-specific resource

In BIG-IQ® Access, you can update the properties of a device-specific resource in the working configuration.

1. At the top left of the screen, select **Access** from the BIG-IQ menu.
2. On the left, expand an Access group and click the name of the device.
   The screen displays a list of resource types.
3. Expand the resource types and select the particular type of resource that you want to change.
   A the screen displays a list of resources displays.
4. Click the name of the resource that you want to edit.
   The properties screen for that resource opens.
5. Edit the resource properties.

   Note: *Click the question mark (?) icon for help on each property.*

6. Click **Save**.

The change is distributed to the BIG-IP® device when you deploy the configuration.

Sharing a device-specific resource

In BIG-IQ® Access, you can make a device-specific resource act like a shared resource.

Note: *When you make a device-specific resource shared, the resource takes the properties defined for it on the source device*

1. At the top left of the screen, select **Access** from the BIG-IQ menu.
2. On the left, expand the name of the Access group that interests you.
3. Under **Device-specific resources**, click any device name.
   The screen displays a list of resource types.
4. Select the type of resource that you want to change.
   The screen displays a list of resources of that type on the right.
5. From the list, select the check box for the resource that you want to make shared.
6. Click **Mark Shared**.
   The resource no longer displays on the list of device-specific resources.

You can now find the resource on the **Shared resources** list.

Returning a shared resource to device-specific resources

If you made a device-specific resource into a shared resource, you can return it to device-specific resources and configure its properties for each device in the Access group.

Note: *Device-specific resources are a system-defined subset of shared resources. Not all shared resources can be made device-specific.*

1. At the top left of the screen, select **Access** from the BIG-IQ menu.
2. On the left, expand the name of the Access group that interests you.
3. Click **Shared resources**.
   The screen displays a list of resources, with **ACCESS POLICIES** selected.
4. Select the type of resource that you want to change.
   The screen displays a list of resources of that type on the right.
5. From the list, select the resource that you want to return to its device-specific state.
6. Click **Make Device Specific**.
   The resource no longer displays on the list of shared resources.

The resource is now located with the device in **Device-specific resources**. The resource properties now match those from the source device.

You can now change the resource properties to meet the device-specific requirements that you have.

Viewing an access policy

After you've imported a source device, you can view the access policies that are configured on it.

Note: *These are the access policies that will be deployed to all the devices in the Access group. You can view the flow of actions in the policy, but not the properties of the actions.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. On the left, expand the name of the Access group that interests you.
4. In the screen on the right, click the name of an access policy.
   A new screen displays the policy's properties.
5. To view different sections of an access policy, you can scroll left, right, up, and down.
6. To move to another section of a large access policy more quickly than scrolling allows, somewhere in the policy hold the right mouse button down and drag the mouse to move around the policy .
7. To close the screen, click **Cancel**.

About the access policy display

When you view an access policy in BIG-IQ® Access, the items in the policy are of a constant size. If an access policy item name is unusually long and does not include spaces, the name of the policy item will be truncated.

Evaluating and Deploying Changes

How do I evaluate changes made to managed objects?

To change the object settings on a managed device, there are four tasks to perform.

This figure illustrates the workflow you perform to manage the objects on BIG-IP® devices. Evaluating the changes you have made is the third step in this process.

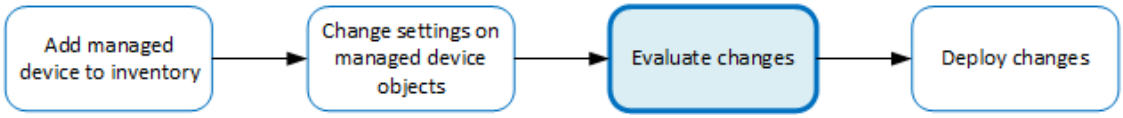

Figure 5: Overview of evaluating changes made to managed objects

Note: *If you need to make an urgent change, you can skip the evaluation step. However, we highly recommend evaluation in all but emergency situations. See Making an urgent deployment for details.*

How do I deploy changes made to managed objects?

Deploying changes applies the revisions that you have made on the BIG-IQ® to the managed BIG-IP® devices.

Note: *Before the BIG-IQ deploys configuration changes, it first reimports the configuration from the managed device to ensure there are no unexpected differences. If there are issues, the default behavior is to discard any changes made on the managed device and then deploy the configuration changes.*

- To accept the default, proceed with the deployment. The settings from the managing BIG-IQ overwrite the settings on the managed BIG-IP device.
- To override the default, rediscover the device and reimport the service. Any changes that have been made using the BIG-IQ are overwritten with the settings from the managed BIG-IP device.

This figure illustrates the workflow you perform to manage the objects on BIG-IP devices. Deploying the settings is the last step in this process.

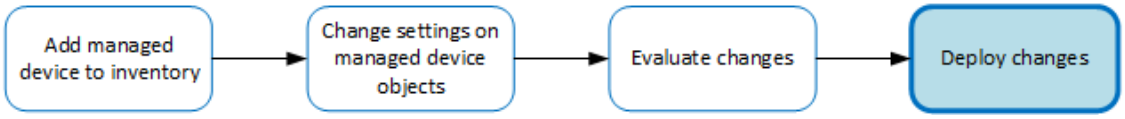

Figure 6: Change managed object workflow

How does deployment to devices in a cluster work?

When you created a cluster in BIG-IQ® inventory, you chose a deployment option for the devices in that cluster.

If you chose to initiate BIG-IP® DSC® sync, and the Sync-Failover group on the BIG-IP system is configured for manual sync, after deployment to either device in the HA pair, Access kicks off manual

sync to the other device. If manual sync succeeds, the deployment is successful. Otherwise, the deployment status shows an error.

If you chose to initiate BIG-IP DSC sync and the Sync-Failover group on the BIG-IP system is configured for automatic sync, after deploying to either device in the HA pair, automatic sync propagates the configuration to the other device. If automatic sync succeeds, the deployment is successful. Otherwise, the deployment status shows an error.

If you chose to ignore BIG-IP DSC sync, you must deploy the configuration from BIG-IQ to both devices in the cluster.

Note: *It is possible that after this, conflicts in DSC sync for these devices will occur.*

Evaluating Access configuration changes

Evaluating your changes gives you a chance to spot critical errors and review your revisions one more time before deploying them.

Note: *Critical errors are issues with a configuration change that cannot be deployed successfully. Verification warnings are less serious in that they may not cause the deployment to fail, but should be reviewed nonetheless.*

Note: *If you have Local Traffic & Network (LTM) changes to deploy, deploy the LTM changes before deploying changes to other components, or those deployments may fail.*

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

   Important: *You must log in as a user with Administrator or Access Manager or Access Deployer access to perform this task.*

2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under **EVALUATE & DEPLOY**, select **Access**.
   The screen opens a list of Access evaluations and deployments that have been created on this device.
5. Under Evaluations, click **Create**.
   The Create Evaluation screen opens.
6. In the **Name** field, type in a name for the evaluation task you are creating.
7. In the **Description** field, type in a brief description for the evaluation task you are creating.
8. For the **Source**, select what you want to evaluate.
   - To compare the object settings currently on the managed device with the object settings in the pending version, select **Current Changes**.
   - To compare the object settings currently on the managed device with the object settings in a stored snapshot, select **Existing Snapshot**, then choose the snapshot you want to use.
9. In the **Target** settings, from the **Group** list, select the Access group that you want to evaluate.
   Devices in the group display in the **Available** field.
10. Move the devices that you want to evaluate to the **Selected** field.

    Note: *If you are evaluating a device that is a member of a cluster set to initiate BIG-IP DSC sync at deployment, you can select either member of the HA pair.*

Note: *If you are evaluating a device that is a member of a cluster set to ignore BIG-IP DSC sync, you should select both devices in the cluster.*

11. If you want to apply access policies on each BIG-IP device after deployment, select **Automatically apply policies after deployment**.
12. Review the evaluation to determine whether you are going to deploy it or not.
    a) If there are critical errors, you cannot deploy these changes. Click each error to see what it is, and then go back to where you made the change to fix it.

       After resolving any critical errors, you can come back and repeat the evaluation.
    b) If there are verification warnings, you can still deploy your changes, but you will probably want to resolve the warnings first. Click each warning to see what it is, and then go back to where you made the change to fix it.

       After resolving any verification warnings, you can come back and repeat the evaluation.
    c) If there are no critical errors or verification warnings, review the changes by clicking the **Difference** link.

       Each change is listed. You can review each one by clicking the name.
13. If the evaluation shows that you must evaluate and deploy Local Traffic configurations, do that before you deploy this evaluation.

To apply the object changes to the managed device, you must deploy them.

Evaluating LTM configuration changes

Evaluating your changes gives you a chance to spot critical errors and review your revisions one more time before deploying them.

Note: *Critical errors are issues with a configuration change that cannot be deployed successfully. Verification warnings are less serious in that they may not cause the deployment to fail, but should be reviewed nonetheless.*

Note: *If you have Local Traffic & Network (LTM) changes to deploy, deploy the LTM changes before deploying changes to other components, or those deployments may fail.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

   Important: *You must log in as an Administrator, ADC Manager, or ADC Deployer to perform this task.*
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under **EVALUATE & DEPLOY**, select **Local Traffic & Network**.
   The screen opens to show a list of LTM evaluations and deployments that have been created on this device.
5. Under Evaluations, click **Create**.
   The Create Evaluation screen opens.
6. In the **Name** field, type in a name for the evaluation task you are creating.
7. In the **Description** field, type in a brief description for the evaluation task you are creating.
8. For the **Source**, select what you want to evaluate.
   - To compare the object settings currently on the managed device with the object settings in the pending version, select **Current Changes**.

- To compare the object settings currently on the managed device with the object settings in a stored snapshot, select **Existing Snapshot**, then choose the snapshot you want to use.

9. For the **Target** setting, identify the devices for which you want to evaluate changes.
   a) If the devices are in a device group, select **Group**, and select the group from the list.
   b) If the devices are not in a device group, select **Device**.
   c) Select the devices from the **Available** list, and use the arrow button to move the devices to the **Selected** list.

   Important: *If you deploy changes to a device that is in a DSC® cluster, you must include both devices before you can create the evaluation.*

   Important: *If the device in the **Selected** list has a filled circle in front of it, a deployment is needed for the BIG-IP device configuration to match the BIG-IQ working configuration for that BIG-IP device. This notification occurs only when creating Web Application Security evaluations.*

10. Click the **Create** button at the bottom of the screen.
    The system adds the new evaluation to the list, and analyzes the changes for errors. When the configuration evaluation completes, you will see how many changes or errors the evaluation found.
11. Review the evaluation to determine whether you are going to deploy it or not.
    a) If there are critical errors, you cannot deploy these changes. Click each error to see what it is, and then go back to where you made the change to fix it.

       After resolving any critical errors, you can come back and repeat the evaluation.
    b) If there are verification warnings, you can still deploy your changes, but you will probably want to resolve the warnings first. Click each warning to see what it is, and then go back to where you made the change to fix it.

       After resolving any verification warnings, you can come back and repeat the evaluation.
    c) If there are no critical errors or verification warnings, review the changes by clicking the **Difference** link.

       Each change is listed. You can review each one by clicking the name.

To apply the object changes to the managed device, you must deploy them.

Deploying LTM configuration changes

When a BIG-IQ® system evaluation of the Access configuration advises you to, you should deploy LTM®before you deploy Access.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

   Important: *You must log in as an Administrator or ADC Deploy to perform this task.*

2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under **EVALUATE & DEPLOY**, select **Local Traffic & Network**.
   The screen displays a list of LTM evaluations and deployments defined on this device.
5. Click the name of the evaluation that you want to deploy.
   The View Evaluation screen opens.
6. Specify whether you want to deploy the changes immediately or schedule deployment for later.
   - To deploy this change immediately:
     1. Select **Deploy Now**.

2. Click **Deploy** to confirm.

- To deploy this change later:

1. Select **Schedule for later**.
2. Select the date and time.
3. Click **Schedule Deployment**.
4. Click **Schedule Deployment** again to confirm.

The process of deploying changes can take some time, especially if there are a large number of changes. During this time, you can click **Cancel** to stop the deployment process.

Important: *If you cancel a deployment, some of the changes may have already deployed.* ***Cancel*** *does not roll back these changes.*

The evaluation you chose is added to the list of deployments on the bottom half of the screen.

- If you chose to deploy immediately, the changes begin to deploy and the Status column updates as it proceeds.
- If you choose to delay deployment, the Status column displays the scheduled date and time.

Deploying the Access configuration

To apply the Access configuration on the BIG-IQ system to your managed devices, you deploy the configuration.

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

Important: *You must log in as an Administrator, Access Manager, or Access Deployer user to perform this task.*

2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under **EVALUATE & DEPLOY**, select **Access**.
   The screen displays a list of Access evaluations and deployments defined on this device.
5. Click the name of the evaluation that you want to deploy.
   The View Evaluation screen opens.
6. Specify whether you want to deploy the changes immediately or schedule deployment for later.

- To deploy this change immediately:

1. Select **Deploy Now**.
2. Click **Deploy** to confirm.

- To deploy this change later:

1. Select **Schedule for later**.
2. Select the date and time.
3. Click **Schedule Deployment**.
4. Click **Schedule Deployment** again to confirm.

The process of deploying changes can take some time, especially if there are a large number of changes. During this time, you can click **Cancel** to stop the deployment process.

Important: *If you cancel a deployment, some of the changes may have already deployed.* ***Cancel*** *does not roll back these changes.*

The evaluation you chose is added to the list of deployments on the bottom half of the screen.

- If you chose to deploy immediately, the changes begin to deploy and the Status column updates as it proceeds.
- If you choose to delay deployment, the Status column displays the scheduled date and time.

Access deployment errors and warnings: causes and resolutions

Problem	Description	Resolution
Access profile type mismatch	The deployment process imports an access profile from the source device to the other devices in the Access group. If an access profile of the same name exists on a non-source device, the access profile types must match. If it does not, a critical error occurs and deployment fails.	On the non-source BIG-IP® device, delete the access profile. Then, redeploy on the BIG-IQ® system.
Sandbox object outside of the `/Common` partition	If partitions exist on the source device in addition to the `/Common` partition, they contain sandbox objects by default. When the deployment process tries to create the sandbox objects, if the same partitions do not exist on the non-source devices, a critical error occurs and deployment fails.	On each non-source BIG-IP device, create the same partitions that exist on the source device. Then, redeploy on the BIG-IQ system.
Machine account	A machine account exists on the source device, but does not exist on a non-source device. A critical error occurs when the deployment process tries to create a machine account on non-source BIG-IP system.	On each non-source BIG-IP device, create a machine account of the same name as the one on the source device. Then, redeploy on the BIG-IQ system.
Non-Access objects	The deployment evaluation process finds that certain virtual servers, SSL profiles, and other objects are used by access policies on the source device but are not present on a non-source device. A critical error occurs because the deployment process cannot create objects not managed by Access.	Create the objects on the non-source BIG-IP devices where needed. Then, redeploy on the BIG-IQ system.
Pools, pool members, self IPs, route domains	Access objects refer to pools, pool members, self IP addresses, and route domains, all of which are managed in ADC. If any of these objects is not present on the source device, evaluation provides a warning that LTM® must be deployed before Access can be deployed. If the warning is ignored, Access deployment fails.	Deploy LTM. Then re-discover LTM before trying to deploy Access.
Adding or updating an OAM server	An Oracle Access Manager (OAM) AAA server exists on the source device. If the deployment process must add or update the OAM server on a non-source device, a message displays advising that the `eam` service on the BIG-IP device must be restarted. The deployment succeeds.	After the deployment completes, restart the `eam` service on the non-source BIG-IP device.

Managing Ongoing Change

How to manage ongoing configuration change

If you make changes on a BIG-IP® device before you have deployed the configuration from the BIG-IQ® system, configuration conflicts can occur. If conflicts do exist, when you deploy the configuration from the BIG-IQ system, you will have to choose between the configuration on the BIG-IQ or on the BIG-IP. You cannot keep both.

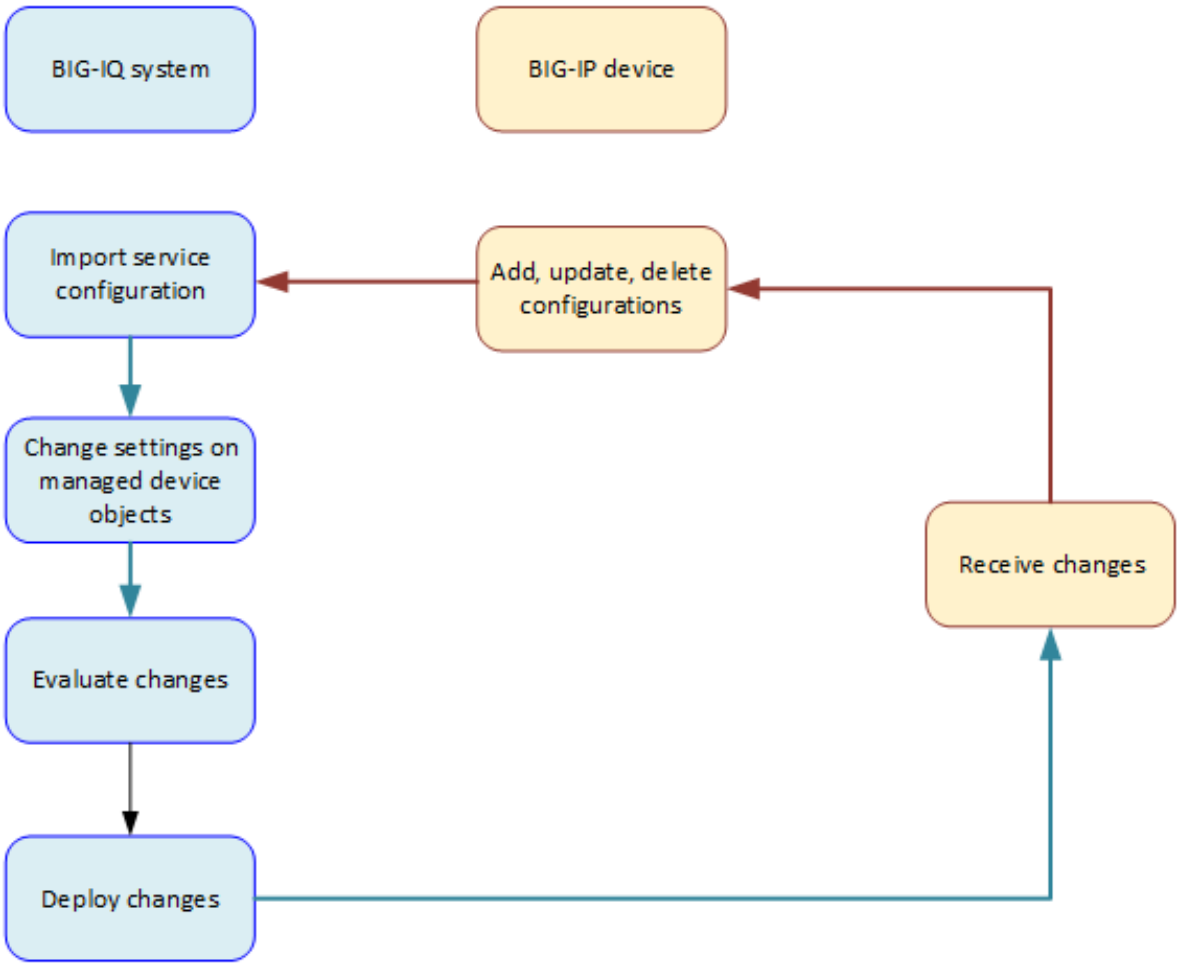

Figure 7: Ongoing change

How does re-import impact the device-specific resources?

When you re-import the APM® service configuration, the process adds and deletes any device-specific resources that were added and deleted on the source device for the Access group. The process, however, does not overwrite any existing device-specific resources on the BIG-IQ® system.

Device-specific resources are processed like this whether you import the APM service configuration from the Device Management user interface, or if you use the Re-import Source option for an Access group.

Guidelines for making changes to the Access configuration

These are general guidelines for updating the configuration:

- You should make any needed change that you can from the Access user interface.
- If you still need to make changes, you should make them on the BIG-IP® source device.

See the table for more specific guidelines.

Resource	Description
Access: Device-specific resource	• Modify device-specific resources on the BIG-IQ® system and deploy the changes. • Add or delete device-specific resources on the source device; then re-import the service configuration into the BIG-IQ system.
Access: Shared resource	Add, modify, and delete shared resources on the source device. Then re-import the service configuration into the BIG-IQ system.
Access: Pools and pool members	You can add and update pools and pool members when you configure some AAA servers in Access. Any changes you make are immediately available in ADC. To deploy these changes, you must deploy ADC before you deploy APM.
ADC: Pools and pool members	If you use ADC to add, update, or delete pools or pool members, you can create conflicts with the Access configuration. If you make changes in ADC, they are not available from Access.
ADC: Route domains and self-IP addresses	To add or edit route domains and self-IP addresses, do so in ADC. To make the changes available in Access, deploy the LTM® working configuration and then reimport the LTM configuration to the BIG-IQ system,
ADC: Virtual servers	Access configuration objects do not refer to virtual servers; however, you probably want to know how to configure them. You can add and edit virtual servers in ADC, but you can configure Access-specific settings, such as specifying an access profile, only on the BIG-IP system. You can add or edit virtual servers in either of these ways:

Resource	Description
	• Add or edit virtual servers in ADC. Deploy the LTM configuration to one or more devices. Edit Access-specific settings on the BIG-IP systems. Reimport the LTM configuration to the BIG-IQ system. • Add or edit a virtual server on the BIG-IP system. Reimport the LTM configuration.
ADC: iRule, nodes, interfaces, routes, VLANs, DNS resolvers	Access configuration objects do not refer to these objects directly. You do not need to worry about conflicts in the Access configuration.

Re-discovering and re-importing the APM service configuration

You can move any changes made to the Access Policy Manager® (APM®) service configuration on the source device into the working configuration for the BIG-IQ® system. You just re-import the source.

Note: *When you use the* ***Reimport Source*** *option for an Access group, it re-discovers and re-imports the APM service configuration. It also detects whether changes were made to the LTM® service configuration and displays a message if you need to re-discover and re-import LTM first.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. In the Access Groups list on the right, click the name of the Access group.
   The Properties screen displays.
4. Click **Reimport Source**.
   A confirmation message displays.

   Important: *Reimporting the source can cause major changes to the working configuration.*
5. To continue with re-discovery and re-import, click **Continue**.

The APM service configuration is imported. Importing the APM service configuration can change objects in the ADC configuration.

You need to move any changes made to the ADC configuration on the source device to the non-source devices too; deploy the LTM service configuration to the non-source devices.

Re-discovering and re-importing the LTM service configuration

You can move any changes made to the Local Traffic Manager™ (LTM®) service configuration on the source device into the working configuration for the BIG-IQ® system. You just re-discover and re-import the LTM service configuration.

Note: *If changes made to Local Traffic configuration objects in ADC dictate that you deploy LTM first, the system displays a message telling you to do that.*

Important: *Do not re-import the LTM service configuration from a non-source device.*

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. Click the name of the device you want to discover a service configuration from.
4. On the left, click **Services**.
5. For Local Traffic (LTM), click **Re-discover**.
   If the current configuration on the BIG-IQ is different than the one on the BIG-IP® device, BIG-IQ displays a screen for you to resolve the conflicts.
6. If there are conflicts, select one of the following options for each object that is different, and then click the **Continue** button:
   - **Use BIG-IQ** to use the configuration settings stored on BIG-IQ.
   - **Use BIG-IP** to override the configuration setting stored on BIG-IQ with the settings from the BIG-IP device.
7. For Local Traffic (LTM), select the **Create a snapshot of the current configuration before importing.** check box to save a copy of the device's current configuration.

   You're not required to create a snapshot, but it is a good idea in case you have to revert to the previous configuration for any reason.
8. For Local Traffic (LTM), click **Re-import**.

The LTM service configuration is imported.

Managing Configuration Snapshots

What is snapshot management?

You can manage configuration snapshots for the configurations you have created on the BIG-IQ® Centralized Management system. A *snapshot* is a backup copy of a configuration. Configuration snapshots are created manually. This type of snapshot does not include events.

Creating a snapshot

You create a configuration snapshot to compare it to another configuration snapshot, or so you can save the current working configuration and then restore from that snapshot if needed.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select **Access**.
   The screen displays a list of Access snapshots that have been created on this device.
4. At the top of the screen, click **Create**.
   The Create Snapshot screen opens.
5. Supply the values on the Create Snapshot screen, and click **Create**.

The system creates the snapshot and adds it to the list of snapshots on the Snapshot and Restore - screen, including information related to the snapshot, including the date it was created, what account created it, and any description.

Comparing snapshots

You can compare two snapshots, or compare a snapshot to the configuration on the BIG-IQ® Centralized Management system to view their differences.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select **Access**.
   The screen displays a list of Access snapshots that have been created on this device.
4. Select the check box to the left of each of the two snapshots to be compared.
5. Click **Compare**.
   The Differences screen opens.
6. Analyze the configuration differences between the two snapshots, When you are finished, click **Cancel** to close the Differences screen, then click **Close**.
   The screen closes and you return to the Snapshot and Restore - screen.

Restoring a snapshot

You can restore a snapshot to change the working configuration to that of the snapshot. Restoring the snapshot merges objects from the snapshot into the BIG-IQ® Centralized Management configuration, and

removes all active locks. No objects in the BIG-IQ configuration are removed when users restore a configuration from a snapshot. Objects are retained in the BIG-IQ configuration and are not deployed to BIG-IP® systems. Once the restore process starts, you cannot modify the BIG-IQ configuration until the process is completed or canceled. If the process is canceled, all configuration settings are rolled back.

Important: *Restoring a snapshot in one component can impact other components that have dependent configuration objects. We recommend that when you restore configurations that involve multiple components, you use snapshots that were created at approximately the same time. Restoring the Access component can require a restore of other dependent modules.*

---

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select **Access**.
   The screen displays a list of Access snapshots that have been created on this device.
4. Select the check box to the left of the snapshot to use to restore the current working configuration to the configuration of the snapshot.
5. Click **Restore**.
   The Restore snapshot to Working Configuration screen opens.
6. Click **Restore** to restore the configuration in the snapshot and have it replace the working configuration.
7. Click **Restore** in the popup screen to confirm that you want to restore the configuration, or click **Cancel** in the popup screen to stop the restore process for this the snapshot.

   You can also click **Cancel** after starting the restore process to roll back the restore.

Managing Event Logs in Access

Managing Event Logs for Access

How do I manage event logs with a Logging Node?

Viewing the event logs as implemented on BIG-IQ® eases browsing of system event logs, and provides a way to obtain useful insights regarding the activity on applications and/or servers. The BIG-IQ platform enables a single view of all filters and log entries (and details for each entry) from multiple BIG-IP® devices.

It also provides a more intuitive navigation path through the log items.

To properly configure event log viewing:

- Discover and activate a BIG-IQ Logging Node.
- License and provision a BIG-IQ Logging Node.
- Define an external machine to which periodic data snapshots are sent.
- Configure a BIG-IP system to collect event logs and send them to the BIG-IQ Logging Node. Part of this configuration includes a virtual server configured with a logging profile.
- Configure a logging profile on BIG-IQ, assign it to a virtual server, and deploy it to the BIG-IP device that has been configured to collect log events.

  A *logging profile* is used to determine which events the system logs, and where, and the format of these events. It then directs security events to a BIG-IQ Logging Node, and the BIG-IQ system retrieves them from that node.

Logging Node uses a search engine that requires separate services for management and traffic. Keeping those services on separate networks reduces unnecessary congestion. The network designs described here are not required, but considered best practice.

BIG-IQ Networks

- A cluster management network to perform Elasticsearch configuration and status operations
- A cluster traffic network for inter-node communication

Logging Node Networks

- A cluster management network to perform Elasticsearch configuration and status operations
- A cluster traffic network for inter-node communication
- A listener network to handle inbound data traffic

This figure illustrates the network topology required to deploy a logging node for your event logs.

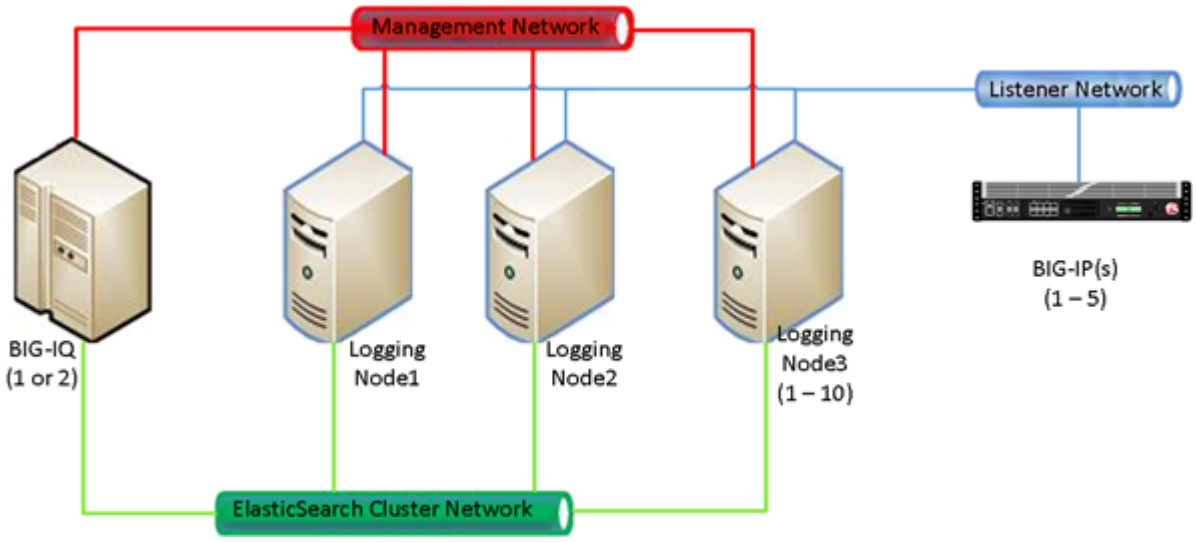

Figure 8: Logging Node network topology

Important: *F5 Networks strongly recommends that the Listener Network and Management Networks be separate. This separation, can help with data protection and management network availability in case the Listener Network is flooded with data.*

What is a BIG-IQ Logging Node?

A *BIG-IQ Logging Node* is a specially-provisioned BIG-IQ® system, that runs the same software version as the BIG-IQ device that you use to manage your security, and the rules that determine your alert responses. After you provision the BIG-IQ Logging Node, you discover it from BIG-IQ and then add the service. Once you configure a logging profile, the Logging Node stores events related to security and application policies from one or more BIG-IP® systems. The BIG-IQ system can then retrieve and manage those logging events.

Note: *The software version on the Logging Node must be the same as the version on its partner BIG-IQ system. If you need to upgrade the Logging Node, follow the instructions in Upgrading BIG-IQ Systems.*

Discover and activate a logging node

Using BIG-IQ® System Management, you can discover a Logging Node and add it to the Logging Group. The BIG-IQ can then access all event on the discovered Logging Node. You can then receive these events from multiple BIG-IP® systems. This unified view makes browsing easier, and provides a complete view of application event activity.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING**.
4. Under **BIG-IQ LOGGING**, select **Logging Nodes**.
5. Click **Add Node**.
6. On the New Logging Device screen, fill in as appropriate:
   a) In **IP Address**, type the management IP address.
   b) In **User Name**, type the user name for an administrator on the Logging Node (for example, `admin`.
   c) In **Password**, type the password for an administrator on the Logging Node (for example, `admin`.

d) In **Transport Address**, type the IP address of the logging node internal self IP address.
e) For **Transport Port**, the default value is 9300. The BIG-IQ uses this port for internal polling and communication with the logging nodes.

7. Click the **Add** button at the bottom of the screen to add the Logging Node to the system. Or, click **Discard** to cancel the operation.

Note: *This operation might take a minute or two.*

8. Repeat these 7 steps for each Logging Node you want to configure.
9. To activate this logging node for the service you want to monitor, in the Services column, click **Add Services**.
   The Logging Node Services screen opens.
10. For the service you want to add, confirm that the **Listener Address** correctly specifies the external self IP address of the Logging Node, and click **Activate**.
    When the service is successfully added, the **Service Status** changes to `Active`.
11. Click **Close**.

Once discovered and activated, this logging node collects the events generated by the configured BIG-IP systems. Thus, BIG-IQ provides a single view of all event entries.

Important: *The **Total Document Count** is not a report of the number of alerts sent to the Logging Node. Instead, it is a sum of various document types sent to the Logging Node. Alerts are included in this list, but this total includes other document types as well.*

Modifying event log indices

Event log indices determine the physical characteristics of what is sent to the Logging Node.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, select **Logging Configuration**.
   The Logging Configuration screen opens to display the current state of the logging node cluster defined for this device.
4. In the Access row in the bottom half of the screen, click the **Configure** button.
   The Access Indices screen opens.
5. For the **Rotation Type**, keep the default setting: **Size Based**.
6. For the **Max Index Size**, type the maximum size of the indices you want to send to the logging node.
   For example, if you type `1000`, when the event log data reaches a size of 1 Gig, it is sent to the logging node.
7. For the **Retained Index Count**, type the total number of indexes you want to store on the logging node.
   The maximum amount of data stored on the Logging Node is the product of the **Max Index Size** and the **Retained Index Count**. When the amount of data reaches this size, the oldest event data is truncated or discarded.
8. Click **Save** to save the indices configuration settings.

Define event snapshot storage locations

Before you can configure the external snapshot storage location, you need the following information on the machine you will use to store the event snapshots:

- storage-machine-IP-address
- storage-file-path
- Read/Write permissions for the storage file path

You need snapshots of your alert data to perform software upgrades, hotfix upgrades, and to restore your .

When event snapshots are created, they need to be stored on a machine other than the Logging Node that stores the events. You define the location for the snapshot by editing the `fstab` file on your Logging Node machines and on the BIG-IQ® and HA peer devices.

Important: *You must perform this task on each Logging Node device, on the BIG-IQ device, and on the BIG-IQ HA peer.*

---

1. On the first device, in the folder `/var/config/rest/elasticsearch/data/`, create a new folder named `essnapshot`.
   `mkdir /var/config/rest/elasticsearch/data/essnapshot`
2. Edit the `/etc/fstab` file to add `/var/config/rest/elasticsearch/data/essnapshot`.
   For example, `//<storage machine ip-address>/<storage-file-path> /var/config/rest/elasticsearch/data/essnapshot cifs iocharset=utf8,rw,noauto,uid=elasticsearch,gid=elasticsearch, 0 0`
3. Run the mount command to mount the snapshot storage location to the new folder.
   For example, from `/var/config/rest/elasticsearch/data` type: `mount essnapshot`.
4. Confirm that the `essnapshot` folder has full read, write, and execute permissions, (specifically `Chmod 777 essnapshot`), and that the owner and group are `elasticsearch` for this folder.
   For example, `ls-l` would yield: `drwxrwxrwx 3 elasticsearch elasticsearch 0 Apr 25 11:27 essnapshot`.
5. Create a test file to confirm that the storage file-path has been successfully mounted.
   For example: `touch testfile`.
   The test file should be created on the storage machine at the location storage file path.
6. Repeat these five steps for each Logging Node, the BIG-IQ, and the BIG-IQ HA peer.

The storage location should now be accessible to the BIG-IQ devices and to the logging node machines.

Define Access snapshot schedules

Before you define snapshot schedules, you must have defined the snapshot storage locations.

Snapshots of the events sent to your Logging Nodes are an essential safeguard for your data. If the machine that stores the events fails, the data can be restored using these snapshots. These snapshots are created based on the snapshot schedules you define. F5 recommends that you schedule snapshots at least every 6 hours and retain at least 4 snapshots.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING**.
4. Under **BIG-IQ LOGGING**, select **Logging Configuration**.
5. For the **Snapshot Schedules** setting, click **Create**.
   The New Logging Snapshot screen opens.
6. For the **Snapshot Name Prefix**, type the string that you want to use to identify the snapshots created by this schedule.

   For example `snapshot_`.
7. In **Snapshots to Keep**, specify the number of snapshots that you want to accumulate before they are deleted for space constraints.

   For example, if you specify 25, then the system will retain a maximum of 25 snapshots before it starts to delete older snapshots as new snapshots are created. You can save up to 100.
8. Define how you want the snapshots scheduled.

Option	Description
Schedule the interval at which you want to create snapshots:	You schedule the system to take snapshots indefinitely. Snapshots are created at the frequency you specify. **1.** Select **Repeat Interval**. **2.** Specify the **Snapshot Frequency**. **3.** Select a time increment. For example, if you set the frequency to **6** and **Hours**, the first log event data snapshot is taken immediately (on **Save**). Subsequent snapshots are taken every 6 hours.
Schedule specific days on which you want to create snapshots:	You schedule the system to take snapshots on specific days. **1.** Select **Days of the Week**. **2.** For the **Days of the Week** setting, select the days on which you want backups to occur. **3.** For the **Start Date**, select the time (date, hour, minute, and AM or PM) on which you want backups to start.

9. Click **Save** to save the new schedule.

How do I license and do the basic setup to start using a Logging Node?

The BIG-IQ® Logging Node runs as a virtual machine in supported hypervisors, or on the BIG-IQ 7000 series platform. You license the Logging Node using the base registration key you purchased. The *base registration key* is a character string that the F5 license server uses to provide access to Logging Node features.

You license Logging Node in one of the following ways:

- If the system has access to the internet, you can have the Logging Node contact the F5 license server and automatically activate the license.
- If the system is not connected to the internet, you can manually retrieve the activation key from a system that is connected to the internet, and transfer it to the Logging Node.
- If your Logging Node is in a closed-circuit network (CCN) that does not allow you to export any encrypted information, you must open a case with F5 support.

When you license the Logging Node, you:

- Specify a host name for the system.
- Assign a management port IP address.
- Specify the IP address of your DNS server and the name of the DNS search domain.
- Specify the IP address of your Network Time Protocol (NTP) servers and select a time zone.
- Change the administrator's default admin and root passwords.

Automatically license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`http://www.f5.com`).

If the BIG-IQ® system is connected to the public internet, you can follow these steps to automatically perform the initial license activation and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://<management_IP_address>`, where `<management_IP_address>` is the address you specified for device management.

2. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

3. Click **Activate**.
The Base Registration Key field is added to the screen.

4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

   Important: *The registration key you use must support a Logging Node capable license.*

5. In the **Add-On Keys** field, paste any additional license key you have.
6. To add another additional add-on key, click the + sign and paste the additional key in the new **Add-On Keys** field.
7. For the **Activation Method** setting, select **Automatic**, and click the **Activate License** button.
   The End User Software License Agreement (EULA) displays.
8. To accept the license agreement, click the **Agree** button.
9. Click the **Next** button at the right of the screen.
   If the license you purchased supports both Logging Node and BIG-IQ Central Management Console, the License Feature Selection popup screen opens. Otherwise the Management Address screen opens.
10. If you are prompted with the License Feature Selection, select **BIG-IQ Logging Node**, and then click **OK**. If you are not prompted, proceed to the next step.

    Important: *This choice cannot be undone. Once you license a device as a Logging Node, you cannot change your mind and license it as a BIG-IQ Management Console.*

    The Management Address screen opens.
11. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.

    You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).
12. In the **Management Port IP Address** field, type the IP address for the management port IP address.

    Note: *The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

13. In the **Management Port Route** field that the system creates, type the IP address for the management port route.
14. Specify what you want the BIG-IQ to use for the **Discovery Address**.
    - To use the management port, select **Use Management Address**.
    - To use the internal self IP address, select **Self IP Address**, and type the IP address.

      Important: *If you are configuring a Logging Node device, you must use the internal self IP address.*

      Note: *The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

15. Click the **Next** button at the right of the screen.
16. In the **DNS Lookup Servers** field, type the IP address of your DNS server.

    You can click the **Test Connection** button to verify that the IP address is reachable.
17. In the **DNS Search Domains** field, type the name of your search domain.

    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
18. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.

    You can click the **Test Connection** button to verify that the IP address is reachable.
19. From the **Time Zone** list, select your local time zone.
20. Click the **Next** button at the right of the screen.

21. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.
22. Click the **Next** button at the right of the screen.

Manually license BIG-IQ and perform initial setup

You must have a base registration key before you can license the BIG-IQ® system. If you do not have a base registration key, contact the F5 Networks sales group (`http://www.f5.com`).

If the BIG-IQ® system is not connected to the public internet, use this procedure to manually activate the license and perform the initial setup.

1. Use a browser to log in to BIG-IQ by typing `https://<management_IP_address>`, where `<management_IP_address>` is the address you specified for device management.
2. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
3. Click **Activate**.
   The Base Registration Key field is added to the screen.
4. In the **Base Registration Key** field, type or paste the BIG-IQ registration key.

   Important: *The registration key you use must support a Logging Node capable license.*

5. In the **Add-On Keys** field, paste any additional license key you have.
6. For the **Activation Method** setting, select **Manual** and click the **Generate Dossier** button.
   The BIG-IQ system refreshes and displays the dossier in the **Device Dossier** field.
7. Select and copy the text displayed in the **Device Dossier** field.
8. Click the **Access F5 manual activation web portal** link.
   The Activate F5 Product site opens.
9. Into the **Enter your dossier** field, paste the dossier.

   Alternatively, if you saved the file, click the **Choose File** button and navigate to it.

   After a pause, the license key text displays.
10. Click the **Next** button.

    The Accept User Legal Agreement screen opens.
11. To accept the license agreement, select the **I have read and agree to the terms of this license**, and click **Next**. button.
    The licensing server creates the license key text.
12. Copy the license key.
13. In the **License Text** field on BIG-IQ, paste the license text.
14. Click the **Activate License** button.
15. Click the **Next** button at the right of the screen.
    If the license you purchased supports both Logging Node and BIG-IQ Central Management Console, the License Feature Selection popup screen opens. Otherwise the Management Address screen opens.
16. If you are prompted with the License Feature Selection, select **BIG-IQ Logging Node**, and then click **OK**. If you are not prompted, proceed to the next step.

    Important: *This choice cannot be undone. Once you license a device as a Logging Node, you cannot change your mind and license it as a BIG-IQ Management Console.*

    The Management Address screen opens.
17. In the **Host Name** field, type a fully-qualified domain name (FQDN) for the system.

    You cannot change this name after you add it. The FQDN can consist of letters and numbers, as well as the characters underscore (_), dash (-), or period (.).
18. In the **Management Port IP Address** field, type the IP address for the management port IP address.

Note: *The management port IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.

19. In the **Management Port Route** field that the system creates, type the IP address for the management port route.
20. Specify what you want the BIG-IQ to use for the **Discovery Address**.
    - To use the management port, select **Use Management Address**.
    - To use the internal self IP address, select **Self IP Address**, and type the IP address.

    Important: *If you are configuring a Logging Node device, you must use the internal self IP address.*

    Note: *The self IP address must be in Classless Inter-Domain Routing (CIDR) format. For example:* `10.10.10.10/24`.
21. Click the **Next** button to save your configuration.
22. In the **DNS Lookup Servers** field, type the IP address of your DNS server.

    You can click the **Test Connection** button to verify that the IP address is reachable.
23. In the **DNS Search Domains** field, type the name of your search domain.

    The DNS search domain list allows the BIG-IQ system to search for local domain lookups to resolve local host names.
24. In the **Time Servers** fields, type the IP addresses of your Network Time Protocol (NTP) servers.

    You can click the **Test Connection** button to verify that the IP address is reachable.
25. From the **Time Zone** list, select your local time zone.
26. Click the **Next** button at the right of the screen.
27. In the **Old Password** fields, type the default admin and root passwords, and then type a new password in the **Password** and **Confirm Password** fields.
28. Click the **Next** button at the right of the screen.

Configuring remote logging

BIG-IP® devices that you configure for remote logging send Access reporting and SWG log report data to the BIG-IQ® Logging Node for storage and management.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, click **Access Reporting**.
4. On the left, expand **REMOTE LOGGING CONFIGURATION** and click **Logging Profiles**.
   The Remote Logging Configuration screen opens to display all of the discovered BIG-IP devices that are provisioned with the Access service.
5. Select the BIG-IP devices for which you want to enable remote logging, and then click **Configure**.
   The hostname of the primary logging node is displayed, and the status changes to let you know whether the enable request was successful.

Restore event log snapshots

To submit the REST API calls required by this task, you must provide the administrator user name and password.

The BIG-IQ® user interface does not currently support restoring the event snapshots. However, if a logging node fails, you can manually restore the data up to the last snapshot.

Please note the following:

- The restore operation requires a down time during which no BIG-IQ or Logging Node work is performed.
- During the restore operation, no event data sent to the Logging Node is retained.
- The restore operation restores only the data from the time before the chosen snapshot was created. Data from the time that the chosen snapshot was created to the current time is not restored.
- Before initiating a snapshot restore, make sure that sufficient disk space is allocated to the /var folder on the device to which you are restoring the snapshot.

1. Log in to BIG-IQ system with your administrator user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. On the left, expand **BIG-IQ LOGGING** and select **Logging Configuration**.
   The Logging Configuration screen opens.
4. Click the **View History** button.
   The BIG-IQ Logging Snapshots screen opens.
5. Browse through the list to get an idea of which snapshot you want to restore.
6. On the Logging Configuration screen, next to **Last Snapshot/Time**, click **Restore**.

Logging Node management

There are a number of useful concepts to consider when you manage Logging Nodes for off-box log storage. This reference material might prove helpful in setting up and maintaining your Logging Node configuration.

Index rotation policy

The optimum settings used to configure your logging node indices depend on a number of factors. Some of the key factors are discussed here.

The system provides the ability to dynamically create new indices based either on a specified interval or on a specified size. The primary goal to consider when you make these decisions is how to maintain a maximum disk allocation for the Logging Node data while maintaining capacity for new data that flows in.

Secondary considerations include search optimization, and the ability to optimize old indices to reduce their size.

Generally, the best policy is one that does not create unnecessary indices. The more indices, the lower the overall performance, because your searches have to deal with more shards. For example, if a module knows that it has a low indexing volume (thousands/day) then it makes the most sense to have a large aggregation per rotation (5 days or 30 days). For components like Web Application Security that probably have high indexing volumes, it makes more sense to rotate every 8 hours (which reduces the number of retained indices).

Index rotation also allows changing sharding and replica counts by changing the template on a given index type. New indices created from that template will contain the new shard and replica count properties.

This table shows the default configuration values for each index running on the BIG-IQ®. These values are based on anticipated data ingestion rates and typical usage patterns.

Component	Index Name	Minimum Number of Logging Nodes	Rotation Policy	Retained Index Count	Approximate time window	Size of /var file system
Access	access-event-logs	2	Time/5 days	19	95 days	500 GB
Access	access-stats	2	Time/5 days	19	95 days	500 GB

Component	Index Name	Minimum Number of Logging Nodes	Rotation Policy	Retained Index Count	Approximate time window	Size of /var file system
Web Application Security	asmindex	5	Size/100000 MB	5	N/A	500 GB
FPS	websafe	2	Time/30 days	100	8 years	10 GB

If multiple modules are running on a given Logging Node or if you have higher inbound data rates, you might have to adjust these values to keep the `/var` file system from filling up (there is a default alert to warn of this when the file system becomes 80% full).

The simplest resolution is to revise the retained index count; lowering this value will reduce the disk space requirements but it will also reduce the amount of data available for queries. For details on changing this setting, refer to *Modifying event indices*.

Logging Node sizing guide

Logging Nodes are specialized BIG-IQ® devices designed to provide sufficient CPU, memory, and disk capacity to store and search logging data from BIG-IP® devices. The underlying technology to provide these services is Elasticsearch. (Information about general Elasticsearch comments can be found on their website: `https://www.elastic.co/guide/en/elasticsearch/reference/current/_basic_concepts.html`)

Logging Nodes managed by BIG-IQ provide an Elasticsearch (ES) cluster that can scale horizontally (more nodes = more capacity), but each node in that cluster has limits on disk space. To mitigate that, there are a number of configuration elements that control how much disk is used by the system.

Logging Node Minimum Recommended Configuration

CPU	8 Cores
Memory	32 GB
Disk	10 GB (`/var` file system)

The `/var` file system on the Logging Node (which includes ES data) is only 10GB in size. To store more data on the file system, you need to extend the size. Refer to *Index rotation policy* for details on managing the data requirements. Extending the file system to 500GB is straightforward, assuming overall disk allocation on the BIG-IQ virtual machine is adequate. Log in as `root` to the Logging Node, and run the following commands.

1. `tmsh show sys disk directory`

   The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            10485760        -
/var/log                        7168000         -
```

2. `tmsh modify sys disk directory /var new-size 500000000`

   `tmsh show sys disk directory`

   The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            10485760        500000000
/var/log                        7168000         -
```

3. Reboot the system and then confirm the size disk size.

   ```
   tmsh show sys disk directory
   ```

   The system response will be similar to this:

```
Directory Name                  Current Size    New Size
--------------                  ------------    --------
/config                         1048576         -
/shared                         10240000        -
/var                            500003840       -
/var/log                        7168000         -
```

Logging Node Capacity

The following table is a very rough guide to how much data can be stored on a given Logging Node. The estimate assumes that the Logging Node has been configured to the recommended /var filesystem size. This size is outlined in the *Index rotation policy*. Because all indexes share the same filesystem, the approximate maximum documents per node estimate assumes no other indexes exist on that node.

Module	**Index name**	Average document size (bytes)	Approximate maximum documents per node
Access	access-event-logs	730	500GB / 730 = 700 million
Access	access-stats	730	500GB / 730 = 700 million
ASM	asmindex	1400	500GB / 1400 = 350 million
FPS	websafe	1400	10GB / 1400 = 70 million

Managing Configuration Snapshots

What is snapshot management?

You can manage configuration snapshots for the configurations you have created on the BIG-IQ® Centralized Management system. A *snapshot* is a backup copy of a configuration. Configuration snapshots are created manually. This type of snapshot does not include events.

Comparing snapshots

You can compare two snapshots, or compare a snapshot to the configuration on the BIG-IQ® Centralized Management system to view their differences.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. Under **SNAPSHOT & RESTORE**, select **Access**.
   The screen displays a list of Access snapshots that have been created on this device.
4. Select the check box to the left of each of the two snapshots to be compared.
5. Click **Compare**.
   The Differences screen opens.
6. Analyze the configuration differences between the two snapshots, When you are finished, click **Cancel** to close the Differences screen, then click **Close**.
   The screen closes and you return to the Snapshot and Restore - screen.

Managing Audit Logs in Access

About audit logs

You use audit logs to review changes in the BIG-IQ® system. All BIG-IQ system roles have read-only access to the audit log, and can view and filter entries. Any user with the appropriate privileges can initiate an action.

You view audit logs by selecting **Audit Logging** from the BIG-IQ menu, and then selecting the appropriate service option from the list on the left.

To view or change the audit log archive settings, click the **Archive Settings** button on the Audit Logging screen. Archived audit log files are stored in the `archive-audit.n.txt` file in the appropriate subdirectory of the `/var/config/rest/auditArchive` directory on the BIG-IQ system:

All API traffic on the BIG-IQ system, and every REST service command for all licensed modules, is logged in a separate, central audit log (`restjavad-audit.n.log`) which is located in `/var/log` on the BIG-IQ system.

Considerations when using the audit log

When using the audit log, consider the following:

- The audit log does not record an entry for every generation of a task. It only records an entry when the task status changes.
- When an object is deleted and then recreated with the same name, partition, and other information, the difference between those objects may show the deleted object as being the previous generation of the new object.
- By default, not all columns are displayed by the audit log to conserve space. To review what columns are displayed, click the gear icon in the upper right of the Audit Logging screen.

Actions and objects that generate audit log entries in Access

BIG-IQ® Centralized Management records in the audit log all user-initiated changes that occur on the management system. A change is defined as when certain objects are modified, when certain tasks change state, or when certain user actions are performed. For example, when the admin account is used to log in to the BIG-IQ system, the audit log records the time, the user (admin), the action (New) and the object type (Login). The log does not include changes that occurred on BIG-IP® devices that were imported.

Changes to working-configuration objects generate audit log entries. In addition, these actions generate log entries:

- Creating or deleting a user account.
- Users logging in and logging out, including when the user is logged out due to inactivity.
- Creating or cancelling a device discovery or a device reimport.
- Adding a new device to an access group.
- Creating or deleting an access group.
- Removing all services.
- Reimporting a source device.
- Making a non-source device the source device.

- Reverting to a source device.
- Saving a configurable property in an existing device object.
- Stopping a session.
- Deleting a previously discovered device.
- Creating or deleting a deployment task.
- Creating a difference task.
- Creating, restoring, or deleting a snapshot.
- Editing some system information (such as editing a host name, a root password, a DNS entry, or an SNMP entry).

Audit log entry properties

The audit log displays the following properties for each log entry.

Property	Description
Source	IP address of the client machine that made the change. This property is blank for actions that were initiated by an internal process. For example, when a user invokes a deployment action, the deployment action then invokes a difference task to find the differences between the current configuration and the one to be deployed. The difference task has no Source IP address.
Time	Time that the event occurred. The time is the BIG-IQ system local time and is expressed in the format: mmm dd, yyyy hh:mm:ss (time zone); for example: `Apr 19, 2016 13:09:03(EDT)`.
User	Name of the account that initiated the action, such as an account named `Admin` for an administrative account.
Action	Type of modification. For operation changes, the action types include New, Delete, and Modify. For task changes, the action types include Start, Finish, Failed, and Cancelled.
Object Name	Object identified by a user-friendly name; for example: `newRule1`, `deploy-test`, or `Common/global`. When the name `RootNode` is listed, that indicates that the object is associated with a BIG-IP device. `RootNode` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Changes	Indicates whether there was a change in the object. If **View** occurs in this column, there is a change to the object. To view the detailed differences of the change, click **View**.
Object Type	Classification for this action. When the type `Root Node` is listed, that indicates that the object is associated with a BIG-IP device. `Root Node` is typically seen when creating, deleting or updating log profiles, service policies, or firewall policies.
Parent Type	Class or group of the parent object.

Viewing audit entry differences

In the audit log, when potential changes to an object are logged, the **View** link is shown in the Changes column for that entry. You can click **View** to examine the differences between generations of that object.

1. Log in to BIG-IQ® system with Administrator or Security Manager credentials.

2. Select **Audit Logging** from the BIG-IQ menu, and then from the list on the left, click the option from which to view audit entries.
3. To display differences for an object, click **View** in the Changes column.

   A popup screen opens, showing two columns that compare the differences between the two generations of the object in JSON. In these columns, additions to an object generation are highlighted in green, and differences are highlighted in gold.

   If the system cannot retrieve a generation of an object, the column displays either `Generation Not Available` or `Generation No previous generation`. Object information may not be available if it has been automatically purged from the system to conserve disk space, or if it has been deleted.

   The JSON difference displayed for a delete entry in the audit log shows the JSON difference from the previous operation because the generation identifier is not incremented when an object is deleted.
4. When you are finished, click **Close** on the popup screen to return to the Audit Logging screen.

Filtering entries in the audit log

You can use the Filter field at the top right of the Audit Logging screen to rapidly narrow the scope displayed, and to more easily locate an entry in the audit log.

- Filtering is text-based.
- Filtering is not case-sensitive.
- You can use wild cards, or partial text.
- All BIG-IQ® roles can filter entries.
- To clear the filter, click the **X** to the right of the search string in the **Filtered by** field on the left.

1. Log in to BIG-IQ system with Administrator or Security Manager credentials.
2. Select **Audit Logging** from the BIG-IQ menu, and then, on the left, click the area from which to view audit entries.
3. Use the Filter field to narrow your search:
   a) Use the arrow key to the left of the field to select the appropriate filter options.
   b) Type the information specific to the object you want to filter on.
   c) Press Enter.

Option	Description
All	Specifies that all objects should be filtered using the filter text. When this option is used, both the user-visible and the underlying data are searched for a match, so you may see matches to your filter text which do not appear to match it.
Source	Type the source IP address in the filter. When this option is used, both the user-visible and the underlying data are searched for a match, so you may see matches to your filter text which do not appear to match it. Using this option is equivalent to using the **All** option.
Time	Type both a date and a time. Displayed times are given in the local time of the BIG-IQ system. Supported time formats are highly Web browser-dependent. Time formats other than those listed might appear to filter successfully but are not supported. Entering a single date and time results in a filter displaying all entries from the specified date and time to the current date and time. For time formats that use letters and numbers, enter the date time in one of the following formats: • mmm dd yyyy hh:mm:ss. Example: `Jan 7 2014 8:30:00`

Option	Description
	• mmm dd, yyyy hh:mm:ss (time zone). Example: `Apr 28, 2016 13:09:03(EDT)` • mmm dd, yyyy. Example: `Apr 28, 2016` • mmm dd, yyyy hh:mm:ss. Example: `Apr 28, 2016 16:09:06` • ddd mmm dd yyyy hh:mm:ss. Example: `Thu Jan 16 2014 11:13:50` For time formats that use only numbers, enter the date time in one of the following formats: • mm/dd/yy hh:mm:ss. Example: `01/01/16 12:14:15` • m/d/yy hh:mm:ss. Example: `1/1/14 12:14:15` • mm/dd/yyyy hh:mm:ss. Example: `1/1/2014 12:14:15`
Node	Type the node name in the filter.
User	Type the user account name in the filter.
Action: Operation	Type the operation action name in the filter. Operation actions include: New, Delete, and Modify.
Action: Task Status	Type the task status action name in the filter. Task status actions include: Start, Finish, Cancelled, and Failed.
Object Name	Type the full or partial name of the object in the filter. If a partition name is displayed, do not include it in the filter. For example, Common/AddressList_4 would be entered as `AddressList_4`. Because the device-specific object name includes the BIG-IP® host name, you can enter a full or partial device name to get all objects for a specific BIG-IP device.
Object Type	Type the object type in the filter.
Parent	Type the parent name in the filter. Only appears for rules to show the rule list, firewall, or policy that contains the rule.
Parent Type	Type the Parent Type name in the filter. Only appears when the Parent field contains a value.
Contains	Specifies that the filter text is contained within the object specified. When you select **Contains**: • If the filter text is a string, the filter text matches an entire string or only a part of a string. • If the filter text is an IP address, the filter text matches an IPV4 or IPV6 address that is the same as the filter text, or matches an IPV4 address range or subnet that includes the filter text. IPV6 addresses can not be found within a range or subnet. • If the filter text is a port number, the filter text matches a port number that is the same as the filter text, or matches a port number range that includes the filter text.
Exact	Specifies that the filter text is exactly contained within the object specified. When **Exact** is selected: • If the filter text is a string, the filter text matches only the entire string. • If the filter text is an IP address, the filter text matches only an IPV4 or IPV6 address that is the same as the filter text. • If the filter text is a port number, the filter text matches only a port number that is the same as the filter text.

The result of a search filter operation is a set of entries that match the filter criteria, sorted by time.

Customizing the audit log display

You can customize the audit log display to assist you in locating information faster.

- To customize the order of columns displayed, click any column header and drag the column to the location you want.
- To sort by column, click the name of the column you want to sort. Not all columns can be sorted. When sorting items in the Object Name column, partition names are ignored. For example, the object name `Common/rule1` would be sorted without the common partition name, as if it were named `rule1`.
- To resize columns, click the column side and drag it to the preferred location.
- To select what columns are displayed, click the gear icon in the upper right of the Audit Logging screen. In the popup screen, select columns you want to display and clear columns you do not want to display. Move your cursor away from the screen to dismiss it.

Managing audit log archive settings

You can view or change the audit archive settings. The archived audit log files are stored in the `/var/config/rest/auditArchive/` directory on the BIG-IQ® system. You can view Access audit logs based on the following Access roles:

- Deployer.
- Editor.
- Viewer
- Manager.

You can view and configure Access archive settings with only the Access Manager role. The roles Auditor, Deployer, and Viewer cannot view or edit archive settings.

1. Log in to BIG-IQ Centralized Management system with Administrator or Security Manager credentials.
2. Select **Audit Logging** from the BIG-IQ menu.
3. Click the **Archive Settings** button in the upper left of the Audit Logging screen to display the audit log settings.
4. Complete or review the properties and status settings, and click **Save**.

Property	Description
Retain Entries	Specifies the number of days after the audit log entries are archived.
Weekly Update	Specifies which days of the week to update the audit log. Select the check box to the left of each day that you want the audit log to be updated. The default is every day.
Start Time	Specifies when the audit archiving should begin. The default is 12:00 am.
Items Expired	Displays the read-only number of entries that have expired.
Last Error	If an error has occurred, displays the read-only error text for any errors found.
Last Error Time	If an error has occurred, displays a read-only value that contains the time the last error was found. The time in the field is the BIG-IQ system local time and is expressed in the format: ddd mmm dd yyyy hh:mm:ss, for example, `Fri Jan 17 2014 23:50:00`.

About archived audit logs

You can view or change how audit logs are archived by clicking the **Archive Settings** button on the Audit Logging screen.

Archived audit log files are stored in the `archive-audit.n.txt` file in the appropriate subdirectory of the `/var/config/rest/auditArchive` directory on the BIG-IQ® Centralized Management system:

- Network Security audit log: `/var/config/rest/auditArchive/networkSecurity/`
- Web Application Security audit log: `/var/config/rest/auditArchive/webAppSecurity/`
- Fraud Protection Service audit log: `/var/config/rest/auditArchive/websafe/`
- Local Traffic and Network audit log: `/var/config/rest/auditArchive/adc/`
- Device Management audit log: `/var/config/rest/auditArchive/device/`
- Access audit log:`/var/config/rest/auditArchive/access/`

Audit entries are appended to the `archive-audit.0.txt` file. When the `archive-audit.0.txt` file reaches approximately 800 MB, the contents are copied to `archive-audit.1.txt`, compressed into the `archive-audit.1.txt.gz` file, and a new empty `archive-audit.0.txt` file is created, which then has new audit entries appended to it.

Up to five compressed archived audit files can be created before those files begin to be overwritten to conserve space. The compressed audit log archive is named `archive-audit.n.txt.gz`, where n is a number from 1 to 5. As the audit log archives are created and updated, the content of the archives is rotated so that the newest archive is always `archive-audit.1.txt.gz` and the oldest is always the highest numbered archive, typically, `archive-audit.5.txt.gz`.

The file content rotation occurs whenever `archive-audit.0.txt` is full. At that time, the content of each `rchive-audit.n.txt.gz` file is copied into the file with the next higher number, and the content of `archive-audit.0.txt` is copied into `archive-audit.1.txt` and then compressed to create `archive-audit.1.txt.gz`. If all five `archive-audit.n.txt.gz`files exist, during the rotation the contents of `archive-audit.5.txt.gz` are overwritten, and are no longer available.

About audit logs in high-availability configurations

In high-availability (HA) configurations, there is a primary and secondary BIG-IQ® system. During failover, the audit log entries and the audit archive settings are copied from the primary to the secondary system before the secondary system becomes the new primary system.

However, archived audit logs are not copied from the primary to the secondary BIG-IQ system.

About the REST API audit log

The REST API audit log records all API traffic on the BIG-IQ® system. It logs every REST service command for all licensed modules in a central audit log (`restjavad-audit.n.log`) located on the system.

Note: *The current iteration of the log is named `restjavad-audit.0.log`. When the log reaches a certain user-configured size, a new log is created and the number is incremented. You can configure and edit settings in `/etc/restjavad.log.conf`.*

Any user who can access the BIG-IQ system console (shell) has access to this file.

Managing the REST API audit log

The REST API audit log contains an entry for every REST API command processed by the BIG-IQ® system, and is an essential source of information about the modules licensed under the BIG-IQ system. It can provide assistance in compliance, troubleshooting, and record-keeping. With it, you can review log contents periodically, and save contents locally for off-device processing and archiving.

1. Using SSH, log in to the BIG-IQ Access system with administrator credentials.
2. Navigate to the `restjavad` log location: `/var/log`.
3. Examine files with the naming convention: `restjavad-audit.n.log`.

   The letter *n* represents the log number.
4. Once you have located it, you can view or save the log locally through a method of your choice.

Reporting

About Access and SWG reports

Access reports focus on session and logging data from Access devices (managed devices with APM licensed and provisioned). SWG reports focus on user requests (for URLs or applications, for example) from Access devices with SWG provisioned.

Access reports and SWG reports provide the following features.

- Reports on any combination of discovered devices, Access groups, and clusters.
- Graphs for typical areas of concern and interest, such as cross-geographical comparisons or top 10 issues.
- Tabular data to support the graphs.
- Ability to drill down from summarized data to details.
- Ability to save data to CSV files.

Setup requirements for Access and SWG reports

To produce Access reports and SWG reports, these tasks must already be complete.

- Set up the BIG-IQ® logging nodes.
- Add the BIG-IP® devices to BIG-IQ inventory.
- Discover the devices. Devices with the APM® service configuration are what you need.
- Run the remote logging configuration setup on the devices from the Access Reporting screen.

What data goes into Access reports for the All Devices option?

The **All Devices** option for Access reports includes data from the devices that are currently managed (discovered) in the BIG-IQ® system. This is in addition to data from devices that were managed at some point during the report timeframe, but that are not currently managed. With **All Devices** selected, if data from unmanaged devices exists, it displays in reports.

An unmanaged device might be unmanaged temporarily or permanently. Any time a configuration management change causes APM® to be undiscovered, the device and its data are moved to **All Devices** until APM is re-discovered on the device.

You cannot generate a report for an unmanaged device. However, you can generate a report for the timeframe when the device was managed, and then search the report for the unmanaged device name. In the Summary report, All Active Sessions includes the number of sessions that were active on the device when it became unmanaged. Those sessions stay in the Summary and in the Active sessions reports until the next session status update, which occurs every 15 minutes.

Running Access reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.

You can create Access reports for any device with the APM® service configuration on it that has been discovered on the BIG-IQ system, whether or not the device is a member of Access group. To create a report, you can select any combination of Access groups, clusters, and devices.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, select **Access Reporting**.
   A Summary report (for all devices and a default timeframe) starts to generate and display.
4. From the left, select any report that you want to run.
5. At the top left of the screen, from the **ACCESS GROUP/DEVICES** list, either select one of the first two options (**All Devices** and **All Managed Devices**) or, select one or more of the other options (**<*Access group name*>**, **<*Cluster display name*>**, and **<*Device name*>**).
   - **All Devices** Includes Access devices that are currently managed, and Access devices that were managed at one time but are not managed now. (A managed device is one that has been discovered with the APM service configuration.)
   - **All Managed Devices** Includes all Access devices that are currently discovered.
   - **<*Access group name*>** - Select to include all devices in the Access group.
   - **<*Cluster display name*>** - Select to include the devices in the cluster.
   - **<*Device name*>** - Select to include the device. You can select any device from **Managed Devices**, **<*Access group name*>**, or **<*Cluster display name*>**.
6. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
7. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Getting the details that underlie an Access report

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.

From the Summary report, and from most session reports, the initial display includes graphs that summarize the report data. You can get successively more detailed information by clicking a bar or a point on a graph or clicking a link if one is displayed on the screen.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, click **Access Reporting**.

   The Summary report is an example of the type of report that presents high-level data, and provides access to underlying data.

   The Summary starts to generate and display. A timeline and some summaries display across the top of the screen. Graphs display under the summaries. Each graph provides different views of the data.
4. Click anywhere in a summary to get more information.

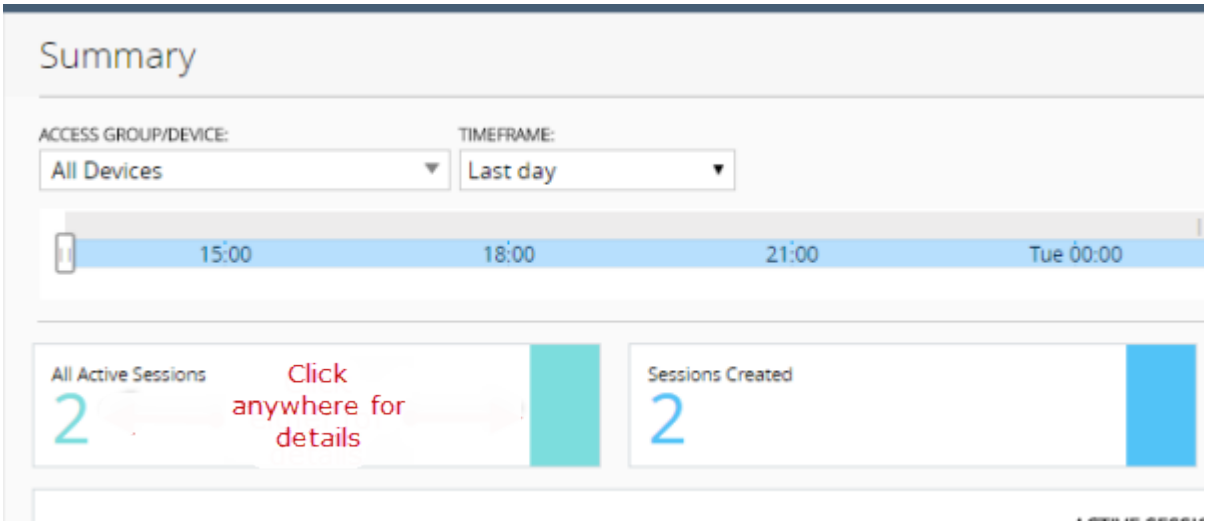

Figure 9: Top left portion of the Summary report display

Additional graphs display, and supporting data displays in a table at the bottom of the screen.

5. If more details are available, click the bars in the graphs to display more details.
6. Scroll down to the table to view the supporting data.
7. If the table includes a **Session ID** field, click the link in that field to open the session details.

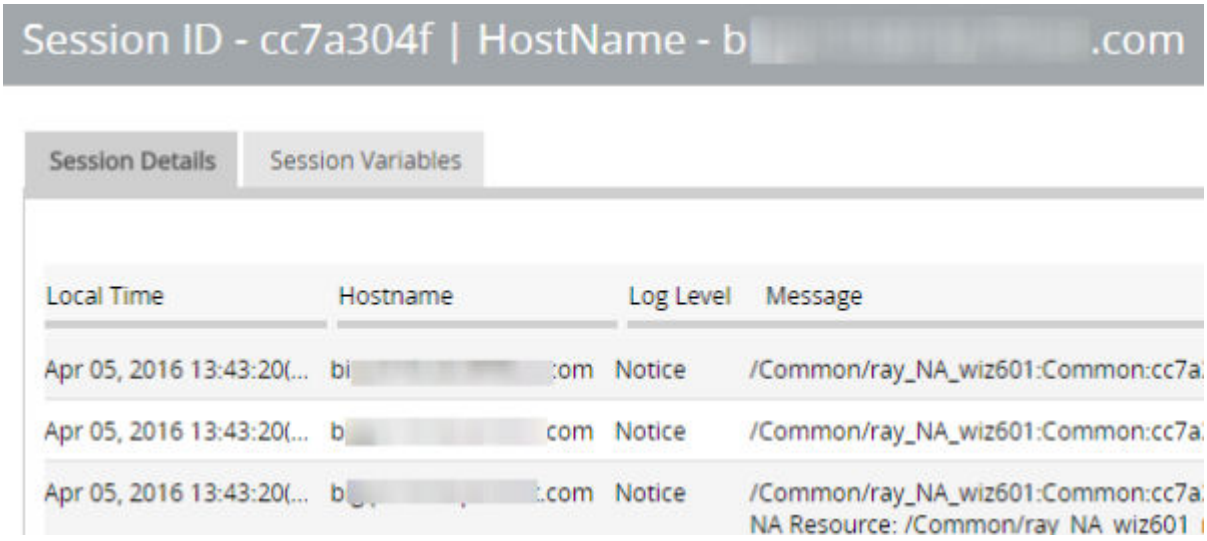

Figure 10: Session details popup screen (with addresses and host names blurred)

8. To change which records display on this screen, select a log level from the **LOG LEVEL** list at the top of the screen.

Stopping sessions on BIG-IP devices from Access

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.

You can stop currently active sessions on BIG-IP® devices, using the Active sessions report on the BIG-IQ system.

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

2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, select **Access Reporting**.
   A SUMMARY report starts to generate and display.
4. On the left, from **Sessions**, select **Active**.
   The screen displays a list of active sessions for all devices.
5. To display sessions for particular devices, groups, or clusters only, select them from the **ACCESS GROUP/DEVICE** list at upper left.
   The screen displays the active sessions for the selected devices.
6. To stop specific sessions only, select the sessions that you want to end and click **Kill Selected Sessions**.
7. To stop all sessions, click **Kill All Sessions**.

Running SWG reports

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.Only a device with SWG provisioned on it can provide data for SWG reports.

You can create SWG reports for Access groups, clusters (in Access groups), or devices that you select from the Access groups and clusters (in Access groups) on the BIG-IQ system.

1. Log in to the BIG-IQ system with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, select **SWG Reporting**.
   A Summary report (for the managed devices and a default timeframe) starts to generate and display.
4. From the left, select any report that you want to run.
5. From the **ACCESS GROUP/DEVICE** list at upper left, select **Managed Devices** or select one or more of these options:
   - ***<Access group name>*** - Select to include all devices in the Access group.
   - ***<Cluster display name>*** - Select to include the devices in the cluster.
   - ***<Device name>*** - Select to include the device. You can select any device from **Managed Devices**, ***<Access group name>***, or ***<Cluster display name>***.
6. From the **TIMEFRAME** list, specify a time frame:
   - Select a predefined time period - These range from **Last hour** to **Last 3 months**.
   - Set a custom time period - Select **Between**, **After**, or **Before**, and click the additional fields that display the set dates and times that support your selection.
7. To save report data in a comma-separated values file, click the **CSV Report** button.
   A CSV file downloads.

Getting the details that underlie an SWG report

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.Only a device with SWG provisioned on it can provide data for SWG reports.

From the Summary report, the initial display includes graphs that summarize the report data. You can get more detailed information by clicking a bar or a point on a graph to see additional graphs and tables with supporting entries.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, select **SWG Reporting**.
   The Summary starts to generate and display. A timeline and some summaries display across the top of the screen. Graphs display under the summaries. Each graph provide different views of the data.
4. Click any bar in a graph on the display to get more information.
   Additional graphs provide different views of the data, and supporting data displays in a table at the bottom of the screen.
5. If more details are available, click the bars in the graphs to display them.
6. Scroll down to the table to view the supporting data.

About the maximum number records for Access and SWG reports

When you run an Access report or an SWG report, Access can get up to 10,000 records to display to you. After you scroll to the end of those 10,000 records, Access displays a message. At that point, all you can do is select fewer devices or select a shorter timeframe.

Setting the timeframe for your Access or SWG report

For Access to have report data for a device, the device must have been added to the BIG-IQ® system, discovered, and had the Access remote logging configuration run for it.

Use the **TIMEFRAME** list at the top of any Access or SWG report to change the report time period.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Access** from the BIG-IQ menu.
3. At the top of the screen, select **Access Reporting** or **SWG Reporting**.
4. To set a predefined timeframe, select one of these from the **TIMEFRAME** list: **Last hour**, **Last day**, **Last week**, **Last 30 days**, **Last 3 months**.
5. To set a custom timeframe, select one of these from the**TIMEFRAME** list:
   - **Between**: Click each of the additional fields that display to select dates and times. The report displays the records between those dates and times.
   - **Before**: Click the additional fields that display to select a date and a time. The report displays the records before that date and time.
   - **After**: Click the additional fields that display to select a date and a time. The report displays the records after that date and time.

Access report problems: causes and resolutions

Problem	Resolution
A session is over, but it continues to display in the Active sessions report.	If a session starts when logging nodes are up and working, but terminates during a period when logging modes are unavailable, the session remains in the Active sessions report for 15 minutes. After 15 minutes, the session status is updated and the session is dropped from the report.
Active sessions are included in the Summary and Active sessions reports for a	Sessions were active on a device when it was removed from an Access group and became unmanaged. Sessions that were active when the device became unmanaged remain counted in All Active Sessions on the Summary screen and stay in the Active sessions report until the next session status update, which occurs every 15 minutes.

Problem	Resolution
device that is no longer managed.	
A session is over, but **Session Termination** and **Session Duration** are blank in a session report.	If a session starts when logging nodes are up and working but terminates during a period when logging nodes are unavailable, the session termination is not recorded and the session duration cannot be calculated.

What can cause logging nodes to become unavailable?

Logging nodes are highly available, but it is still possible for them to become unavailable. This could occur, for example, if all logging nodes are on devices in the same rack in a lab, and the power to the lab shuts down.

Reference

About iApps and Access

On a BIG-IP® system, a configuration that is created using an iApp can be updated only by using the same iApp. Access does not support iApps®. Access does not import, manage, or deploy resources that were created using an iApp.

Shared configuration resources

The tables list configurations that are shared or can be made shared.

Table 1: Access policies and related resources

Resource	Description
Policies	Access policies
Profiles	Properties for the session
CAPTCHA configurations	Specifies the CAPTCHA service
NTLM Auth Configuration	Used to authenticate Exchange applications

Table 2: AAA servers

Resource	Description
RADIUS*	RADIUS accounting and RADIUS authentication
LDAP*	LDAP and LDAPs authentication; LDAP queries
Active Directory*	Active Directory authentication and query
Active Directory Trusted Domains	Authenticate users across all trusted domains or forests for a customer
SecurID*	RSA SecurID authentication
HTTP*	HTTPS authentication; HTTP Basic/NTLM authentication
Oracle Access Manager*	Native integration with Oracle Access Manager
OCSP Responder*	Machine certificate revocation status; user certificate revocation status
CRLDP*	Retrieve Certificate Revocation Lists from network locations (Distribution Points)
TACACS+*	TACACS+ authentication and accounting
Kerberos*	Kerberos end-user login; basic or Kerberos authentication

Resource	Description
SAML*	External SAML Identity Provider for the BIG-IP® system, as a SAML service provider, to communicate with
Endpoint Management Systems*	Server properties *This resource is device-specific but can be made shared

Table 3: ACLs

Resource	Description
User-defined ACLs*	ACLs that users create
All ACLs*	The order of system-defined and user-defined ACLs *This resource is device-specific but can be made shared

Table 4: SSO Configurations

Resource	Description
HTTP Basic	Single sign-on (SSO) using cached user identity and authorization header
NTLMV1	Challenge-response; proves user identity without sending password to server
NTLMV2	Challenge-response; proves user identity without sending password to server
Kerberos	Transparent authentication of users to Windows Web application servers (IIS) joined to Active Directory domain
Forms	Detects start URL match and uses cached user identity to construct and send HTTP form-based post request on behalf of the user
Forms - Client Initiated	Detects login page request, puts generated JavaScript code into login page, and returns it to client, where it is automatically submitted by the inserted JavaScript
SAML	SAML local Identity Provider (IdP) service is a type of SSO service that BIG-IP, configured as an IdP, provides

Table 5: SAML

Resource	Description
Local SP Services	BIG-IP system as Service Provider (SP) provides SP services
External IdP Connectors	BIG-IP system as SP relies on external Identity Providers (IdPs) for authentication

Resource	Description
Local IdP Services	BIG-IP system as IdP provides SSO authentication services
External SP Connectors	BIG-IP system as IdP works with external SPs
Artifact Resolution Services*	Supports SAML artifacts on a BIG-IP system configured as a SAML IdP
BIG-IP IdP Automation	Supports configuration automation
SAML Resources	Resources to support the SAML configuration
	*This resource is device-specific but can be made shared

Table 6: Local User DB

Resource	Description
Manage Instances	Local user database instances

Table 7: Hosted Content

Resource	Description
Manage Files	Hosted content files
Manage Profile Access	Access control for hosted content files using access profiles

Table 8: Webtops

Resource	Description
Webtops	Webtop used in Portal Access or Network Access
Webtop Links	Links for inclusion on a webtop
Webtop Sections	Sections to organize content on a webtop

Table 9: Secure Web Gateway

Resource	Description
URL Categories	URL categories
URL Filters	URL filters
Applications	System-defined list of applications
Application Filters	User-defined application filters
Report Settings	Sets up statistics (for use with SWG subscription service)

Table 10: Network Access

Resource	Description
Network Access resource*	A Network Access resource allows user access to the local network through a secure VPN tunnel

Resource	Description
Lease Pools*	IPV4 or IPV6 lease pools associate a group of IP addresses with a Network Access resource
Client Traffic Classifiers*	Used to shape traffic for Network Access client connections from Windows
Client Rate Classes	Base and peak rates for traffic; associated with a client traffic classifier
	*This resource is device-specific but can be made shared

Table 11: Application Access

Resource	Description
App Tunnels*	Provide secure, application-level TCP/IP connections from a client to the network
Remote Desktops*	Allow users to access internal servers (Citrix, VMware View Connection, or Microsoft Remote Desktop) in virtual desktop sessions
VDI Profiles	Virtual desktop interface profile for a remote desktop configuration
Citrix Bundles	Hosted content used to deliver a Citrix Receiver client to a user's Windows computer
Microsoft Exchange	Profile for Microsoft Exchange application authentication
	*This resource is device-specific but can be made shared

Table 12: Portal Access

Resource	Description
Portal Access resources*	Provide user access to internal web applications with a web browser from outside the network
Rewrite profiles	An LTM profile treated as a shared resource
	*This resource is device-specific but can be made shared

Table 13: Resources that are not grouped in the user interface

Resource	Description
Per-Request Policies	Policies that run for requests made after a session is established
Secure Connectivity	Connectivity profile for remote access
Event Logs Settings	Log settings for APM, components within APM, and SWG

Table 14: Bandwidth Controllers

Resource	Description
Policies	A resource configured outside of APM at the system level and that is treated as a shared resource in Access.
Priority Groups	A resource configured outside of APM at the system level and that is treated as a shared resource in Access.

Device-specific configuration resources

Theses tables list device-specific resources.

Table 15: AAA servers

Resource	Description
RADIUS*	RADIUS accounting and RADIUS authentication
LDAP*	LDAP and LDAPs authentication; LDAP queries
Active Directory*	Active Directory authentication and query
SecurID*	RSA SecurID authentication
HTTP*	HTTPS authentication; HTTP Basic/NTLM authentication
Oracle Access Manager*	Native integration with Oracle Access Manager
OCSP Responder*	Machine certificate revocation status; user certificate revocation status
CRLDP*	Retrieve Certificate Revocation Lists from network locations (Distribution Points)
TACACS+*	TACACS+ authentication and accounting
Kerberos*	Kerberos end-user login; basic or Kerberos authentication
SAML*	External SAML Identity Provider for the BIG-IP® system, as a SAML service provider, to communicate with
Endpoint Management Systems*	Server properties
	*This resource is device-specific but can be made shared

Table 16: ACLs

Resource	Description
User-defined ACLs*	ACLs that users create
All ACLs*	The order of system-defined and user-defined ACLs

Resource	Description
	*This resource is device-specific but can be made shared

Table 17: SAML

Resource	Description
Artifact Resolution Services*	Supports SAML artifacts on a BIG-IP system configured as a SAML IdP *This resource is device-specific but can be made shared

Table 18: Network Access

Resource	Description
Network Access resource*	A Network Access resource allows user access to the local network through a secure VPN tunnel
Lease Pools*	IPV4 or IPV6 lease pools associate a group of IP addresses with a Network Access resource
Client Traffic Classifiers*	Used to shape traffic for Network Access client connections from Windows *This resource is device-specific but can be made shared

Table 19: Application Access

Resource	Description
App Tunnels*	Provide secure, application-level TCP/IP connections from a client to the network
Remote Desktops*	Allow users to access internal servers (Citrix, VMware View Connection, or Microsoft Remote Desktop) in virtual desktop sessions *This resource is device-specific but can be made shared

Table 20: Portal Access

Resource	Description
Portal Access resources*	Provide user access to internal web applications with a web browser from outside the network *This resource is device-specific but can be made shared

Table 21: Portal Access resources that can be device-specific or made shared

Resource	Description
Machine Account	For Microsoft Exchange clients that use NTLM authentication

Legal Notices

Legal notices

Publication Date

This document was published on September 14, 2016.

Publication Number

MAN-0615-01

Copyright

Trademarks

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

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

Patents

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

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

Index

A

B

C

D

E

H

I

L

M

N

O

P

R

S

T

U