F5® BIG-IQ® Centralized Management: Local Traffic & Network

Version 5.2

Table of Contents

BIG-IQ Local Traffic & Network: Overview

What is Local Traffic & Network?

Local Traffic & Network are two parts of centralized management that you use to manage the local traffic (such as servers, nodes, pools, or pool members) and network (such as interfaces, self IP addresses, or VLANs) configuration objects that move your traffic.

Local Traffic & Network helps the user:

- Create efficient work flows to view the Local Traffic & Network configurations in a relational and dynamic user interface.
- Control access to configuration objects using fine-grained, role-based access control (RBAC). This allows administrators to delegate frequently performed operations (for example, enabling or disabling pool members) to the correct team member.
- Maintain ultimate control of the LTM® configuration by providing a staging option. Delegated team members make all relevant changes, then the administrator can apply them after a quick review.

Understanding BIG-IQ Local Traffic & Network

Local Traffic & Network features include:

- Device discovery with import of local traffic and network objects referenced by discovered devices
- Management of shared objects (such as profiles, monitors, and iRules®)
- Audit log used to record every local traffic or network change and event
- Role-based access control
- Deployment of configurations from snapshots, and the ability to preview differences between snapshots

Local traffic & network provides a centralized management platform so you can perform all these tasks from a single location. Rather than log in to each device to manage the object configuration locally, it is more expedient to use one interface to manage many devices. Not only does this simplify logistics, but you can maintain a common set of configuration objects and deploy a common set of profiles, monitors, and other shared objects to multiple, similar devices from a central interface.

Bringing a device under central management means that its configuration is stored in the local traffic & network database, which is the authoritative source for all configuration entities. This database is also known as the working configuration or working-configuration set.

Once you discover and import services for a device, it is deemed to be under central management. You should not make changes locally (on the BIG-IP device) to a device that is under central management unless there is an exceptional need.

Important: *If changes are made locally for any reason, rediscover and reimport the device to reconcile those changes with the local traffic & network working configuration set. Unless local changes are reconciled, the deployment process overwrites any local changes.*

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

About BIG-IQ Centralized Management configuration sets

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

Current configuration set

The configuration of the BIG-IP device as discovered by BIG-IQ Centralized Management. The *current configuration* is updated during a re-discover and re-import, and before calculating differences during the deployment process.

Working configuration set

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

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

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

Managing Device Resources

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

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

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

Stage 1:

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

Stage 2:

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

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

Adding devices to the BIG-IQ inventory

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

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

If you are running BIG-IP version 11.5.1 up to version 11.6.0, you might need root user credentials to discover and add the device to the BIG-IP devices inventory. You don't need root user credentials for BIG-IP devices running 11.5.0 - 11.5.1 and 11.6.1 - 12.x.

Note: *A BIG-IP device running versions 10.2.0 - 11.5.0 is considered a legacy device and cannot be discovered from BIG-IQ version 5.2. If you were managing a legacy device in previous version of BIG-IQ and upgraded to version 5.2, the legacy device displays as impaired with a yellow triangle next to it in the BIG-IP Devices inventory. To manage it, you must upgrade it to 11.5.0 or later. For instructions, refer to the section titled, Upgrading a Legacy Device.*

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

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

---

1. At the top of the screen, click **Devices**.
2. Click the **Add Device** button.
3. In the **IP Address** field, type the IPv4 or IPv6 address of the device.

4. In the **User Name** and **Password** fields, type the user name and password for the device.
5. If this device is part of a DSC pair, from the **Cluster Display Name** list, select one of the following:
   - For an existing DSC pair, select **Use Existing** from the list and select the name DSC group from the list.
   - To create a new DSC pair, select **Create New** from the list, and type a name in the field.

   For BIG-IQ to properly associate the two devices in the same DSC group, the **Cluster Display Name** must be the same for both members in a group.

   There can be only two members in a DSC group.
6. If this device is configured in a DSC pair, select an option:
   - **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option if this device is part of a DSC pair and you want this device to automatically synchronize configuration changes with the other member in the DSC group.
   - **Ignore BIG-IP DSC sync when deploying configuration changes** Select this option if you want to manually synchronize configurations changes between the two members in the DSC group.
7. 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.*
8. If a framework upgrade is required, in the popup window, in the **Root User Name** and **Root Password** fields, type the root user name and password for the BIG-IP device, and click **Continue**.
9. If in addition to basic management tasks (like software upgrades, license management, and UCS backups) you also want to centrally manage this device's configurations for licensed services, select the check box next to each service you want to discover and then click **Continue**.

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

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

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

Importing service configurations for a device

You must add a device to the BIG-IP Device inventory list, and discover associated services, before you can import services to BIG-IQ for the device.

To manage a device's service configuration from BIG-IQ®, you must import the service configuration from the managed device to BIG-IQ.

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. At the top of the screen, click **Devices**.
2. Click the name of the device you want to import a service configuration from.
3. On the left, click **Services**.
4. For the device's configuration you are importing, 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.
5. Click the **Import** button next to the service you want to import to the BIG-IQ system.

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.

You can now manage the configuration of this service for this device from BIG-IQ.

Managing a device from the device properties screen

You can use a device's Properties screen to manage that device. You can log directly in to the device, remotely reboot it, and create an instant backup of its configuration. You can also view details about the managed device, such as:

- Host name
- Self IP Address
- Build Number
- Software Version
- Status
- Last Contact
- Boot Location
- Cluster Properties

From this screen you can also perform the following tasks:

- Create an instant backup of the device's configuration.
- Change the boot location of the device.
- Edit cluster properties.
- Log directly into the device from BIG-IQ®.
- Reboot the device from BIG-IQ.
- Access details about the health of the device.
- Access statistics for the device (if applicable).
- Access services licensed for the device.

1. At the top of the screen, click **Devices**.
2. Click the name of the device you want to view.
   The device Properties screen opens.

Filtering the BIG-IP device inventory list for specific BIG-IP components

From each BIG-IQ® screen that contains a list of objects, you can easily find specific objects. For example, after you discover several devices, you might want to find a specific device by its name or IP address. To do this, you start by filtering on certain configuration objects. Filtering on specific criteria saves you time because you can view only those objects associated with the criteria you specify.

1. At the top of the screen, click **Devices**.
2. To search for a specific object, in the **Filter** field at the top right of the screen, type all or part of an object's name and click the filter icon.
   BIG-IQ refreshes the screen to show only those devices that contain the object you filtered on.
3. To remove the filter, click the **X** icon next to it.

How do I change object settings on a managed device?

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. Changing the settings is the second step in this process.

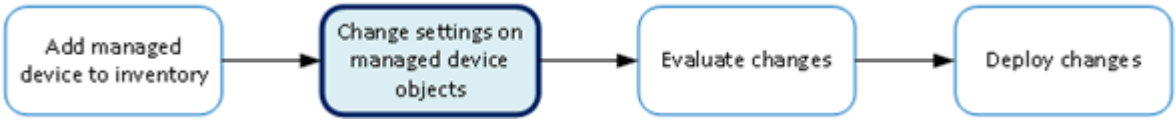

Figure 1: Change managed object workflow

Changing device local traffic objects

Making revisions to the configuration of local traffic objects simplifies managing your devices.

Important: *If you revise configurations on devices that belong to a high availability cluster, the BIG-IQ® synchronizes cluster members automatically when you deploy the change. Do not try to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

---

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select the object type that you want to modify.
   The screen displays a list of objects (of the type you selected) that are defined on this BIG-IQ.
4. Click the name of the object that you want to change.

   If you select **Virtual Servers**, there are a couple of unique operations that you can perform at this point. You can either clone a virtual server to create a new one based on the selected server (see *Cloning a virtual server*), or you can attach iRules to several virtual servers at once (see *Attaching iRules to virtual servers*).

   The Properties screen for the selected object opens.
5. Make changes to the properties you want to modify.
6. When you are satisfied with the changes you have made, click **Save**.
   The revisions you saved are made, and the Properties screen for the selected object closes.

Changes that you make are made only to the pending version. The *pending version* serves as a repository for changes you stage before deploying them to the managed device. Object settings for the pending version are not the same as the object settings on the actual BIG-IP® device until they are deployed or discarded.

Important: *There is an exception to this pattern. When you view properties for a pool member and click* ***Enable****,* ***Disable****, or* ***Force Offline****, you can choose whether you want the change to occur immediately (****Change Now****) or not at all (****Cancel****). The same exception is true when you enable or disable a virtual server.*

To apply the working configuration settings to the BIG-IP device, you now need to deploy the revisions.

Creating a new virtual server

In BIG-IQ® Centralized Management, you can use the Local Traffic interface to add a virtual server to a managed device.

Important: *When you are revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen lists the virtual servers defined on this device.
4. Click **Create**.
   The New Virtual Server screen opens.
5. In the **Name** field, type in a name for the virtual server you are creating.
6. From the **Device** list, select the device on which to create the virtual server.
7. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the virtual server.
8. In the **Description** field, type in a brief description for the virtual server you are creating.
9. If you want the virtual server and its resources to be available for load balancing, for **State (on BIG-IQ)**, select **Enabled**.
10. For the **Source Address**, type an IP address or network from which the virtual server will accept traffic.

    For this setting to work, you must specify a value other than `0.0.0.0/0` or `::/0` (that is, any/0, any6/0). In order to maximize the utility of this setting, specify the most specific address prefixes that include your customer addresses, but exclude addresses outside of their range.
11. For the **Destination Address**, type the IP address of the destination you want to add to the Destination list.
12. In the **Service Port** field, type a service port number, or select a type from the list.

    When you select a type from the list, the value in the **Service Port** field changes to reflect the associated default, which you can change.
13. To configure the virtual server so that its status contributes to the associated virtual address status, select the check box for **Notify Status to Virtual Address**.

    When this setting is disabled, the status of the virtual server does not contribute to the associated virtual address status. When you enable route advertisement of virtual addresses, this status impacts the behavior of the system.
14. To specify configuration parameters for this virtual server, expand **Configuration** and continue with the next sixteen steps. Otherwise, skip to step 32 in this procedure.
15. For the **Type**, select the type of network service provided by this virtual server. The default is **Standard**.

    Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.
16. For the **Protocol**, select the network protocol name you want the system to use to direct traffic on this virtual server. The default is **TCP**. The Protocol setting is not available when you select **Performance (HTTP)** as the type.

    Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.
17. For the **VLANs and Tunnel Traffic** setting, select the VLANs and tunnels for which the virtual server is enabled or disabled. The default is **All VLANs and Tunnels**. If you select another option, the system presents additional settings.

Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.

18. From the **Source Address Translation** list, select the type of address translation pool used for implementing selective and intelligent source address translation.

Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.

19. In the **Connection Limit** field, type the maximum number of concurrent connections allowed for the virtual server.
20. In the **Connection Rate Limit** field, type the maximum number of connections-per-second allowed for a pool member.

    When the number of connections-per-second reaches the limit for a given pool member, the system redirects additional connection requests. This helps detect Denial of Service attacks, where connection requests flood a pool member. Setting the limit to 0 turns off connection limits.
21. From the **Connection Rate Limit Mode** list, select the scope of the rate limit defined for the virtual server.

Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.

22. If you want the system to translate the virtual server address, select **Address Translation**.

    This option is useful when the system is load balancing devices that have the same IP address.
23. If you want the system to translate the virtual server port, select **Port Translation**.

    This option is useful when you want the virtual server to load balance connections to any service. The default is enabled.
24. From the **Source Port** list, select how you want the system to preserve the connection's source port.

Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation on virtual servers available on* `support.f5.com`.

25. To replicate client-side traffic (that is, prior to address translation) to a member of a specified pool, select that pool from the **Clone Pool (Client)** list.
26. To replicate server-side traffic (that is, prior to address translation) to a member of a specified pool, select that pool from the **Clone Pool (Server)** list, select the device on which to create the virtual server.
27. Use the **Auto Last Hop** list to specify whether you want the system to send return traffic to the MAC address that transmitted the request, even if the routing table points to a different network or interface.
28. From the **Last Hop Pool** list, select the pool the system uses to direct reply traffic to the last hop router.
29. If you want the system to allow IPv6 hosts to communicate with IPv4 servers, select **NAT64**.
30. To specify the virtual server score in percent, type that value in the **VS Score** field.

    Global Traffic Manager™ (GTM™) uses this value to load balance traffic in a proportional manner.
31. To specify additional resource details for this virtual server, expand **Resources** and continue with the next two steps. Otherwise, skip to the last step in this procedure.
32. To specify which iRules® are enabled for this virtual server, use the arrow buttons to move iRules between the **Available** and **Enabled** lists.

    iRules are applied in the order in which they are listed.
33. Use the **Default Pool** list to select the pool name that you want the virtual server to use as the default pool.

A load balancing virtual server sends traffic to this pool automatically, unless an iRule directs the server to send the traffic to another pool.

34. For the **Default Persistence Profile**, select the name of the default profile you want the virtual server to use to maintain session persistence.
35. For the **Fallback Persistence Profile**, select the name of the fallback profile you want the virtual server to use to maintain session persistence.

Note: *You can select* ***Default Persistence Profile*** *alone, or you can select both. That is, if you use* ***Fallback Persistence Profile****, you must also select a* ***Default Persistence Profile****. For additional detail on how fallback persistence profiles work, refer to SOL30483109: Overview of Fallback Persistence on* `AskF5.com`

---

36. Click **Save & Close**.
    The system creates the new virtual server with the settings you specified.

Cloning a virtual server

You can use the BIG-IQ® Local Traffic interface to create a new virtual server based on the specifications for an existing one. This can be a great time saver when you need to create several virtual servers that use a number of similar settings.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on this device.
4. Select the check box associated with the existing virtual server that you want to clone.
5. Click the **Clone** button.
   The BIG-IQ creates a new virtual server using the settings of the one you selected.
6. Modify the parameters for the new virtual server as needed.

Important: *Two virtual servers cannot share the same* ***Destination Address****,* ***Protocol****, and* ***VLAN****.*

---

7. When you are satisfied with the settings for the new virtual server, click **Clone**.
   The system creates the new virtual server with the settings you specified.

How do I manage LTM profiles in BIG-IQ?

You can create or modify custom LTM® profiles in BIG-IQ® Centralized Management and then deploy the profiles to your managed devices. The current release supports a number of profile types.

When you create a profile, you specify a parent profile from which the custom profile inherits its properties. You then specify which of these properties you want to override. You can name any existing profile as a parent profile. When you modify a profile that has child profiles (that is, profiles that name your profile as a parent profile), all of the child profiles inherit any changes you made in the parent profile (except those you choose to override).

Create an LTM profile

You must discover a device and import that device's service configurations before you can add a profile to that device from BIG-IQ® Centralized Management.

Creating a new profile allows you to specify the parameters that define the characteristics you want your virtual servers to use. Each virtual server that references this profile uses the parameters you specify for this profile. Additionally, the parameters you define for this profile are given to the profiles that name this profile as their parent profile.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.

3. Under **LOCAL TRAFFIC**, select **Profiles**.
   The screen displays the list of profiles defined on this device.
4. Click **Create**.
   The Profiles - New Item screen opens.
5. In the **Name** field, type in a name for the LTM profile you are creating.
6. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the profile.
7. For the **Type**, select the type of profile you want to create.
   The **Parent Profile** field displays.
8. From **Parent Profile**, select the parent profile from which you want your profile to inherit settings.

   Note: *The parent profile you select determines the value of the profile parameters for this profile. You can override these values, but if you do not, changes made to parameters in the parent profile propagate to all child profiles.*

   A number of additional fields display, specifying the parameters associated with the parent profile you selected. There are two controls for each field. The first one (a check box) controls whether you want to override the inherited value for that field. The second control (the type varies by field) sets the value you want for the parameter.
9. For any fields you want to override, select the check box and then specify the value you want for the fields you selected.

   Note: *You can select* ***Override All*** *if you want to override all of the parent profile parameter values.*

   Important: *If you override a parent profile parameter, regardless of whether or not you change the parameter's value, then future changes to the parent's parameter value will not be inherited by this profile.*
10. If you are adding a profile that requires a security parameter, specify the passphrase in the corresponding **Passphrase** field.

    Note: *For version 12.0.0 devices, you do not need to supply the pass phase for the profile. For devices prior to version 12.0.0, if you plan to make changes to a Client SSL profile, you need to supply the pass phrase for that profile. If you do not change any of the parameters for the profile or associate the profile with a virtual server or another client SSL profile, then you can leave this field blank. So, if you add a pre-version 12.0.0 device that has a significant number of profile definitions, you do not need to add the pass phrase for every profile, just the ones that you plan to change or associate with an LTM object.*
11. Click **Save & Close**.
    The system creates the new profile you specified and adds it to the list of profiles.

    For detailed information on the impact of using a particular profile parameter value, refer to the BIG-IP documentation on `AskF5.com`.

You can now use the profile you created. You can select it when you configure a virtual server. You can also use it as a parent profile to base new BIG-IP® LTM® profiles on.

You must deploy your changes to the BIG-IP device before you can see these changes on the device.

Edit an LTM profile

By editing a profile, you can revise the parameters that define the characteristics you want your virtual servers to use. Each virtual server that references this profile uses the parameters you specify for this profile. Additionally, the parameters you define for this profile are given to the profiles that name this profile as their parent profile.

1. At the top of the screen, click **Configuration**.

2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Profiles**.
   The screen displays the list of profiles defined on this device.
4. Click the name of the profile you want to edit.
   The screen displays the current settings for the selected profile.
5. In the **Description** field, you can revise the brief description for the selected profile.
6. Under **Referenced by**, note the virtual servers and profiles that refer to this profile.

   Changes you make to this profile impact all of the virtual servers listed here.

   Profiles that are listed here name this profile as their parent profile, so they inherit any changes you make to this profile.
7. Under **Override All**, select the corresponding check box for any fields you want to override, and then specify the value you want for the fields you selected.

   Note: *You can select **All** if you want to override all of the parent profile parameter values.*

   ---
8. If you imported a profile that requires a security parameter, specify the passphrase in the corresponding **Passphrase** field.

   Important: *For imported profiles that use passphrases:*

   - If the profile was imported from a version 12.0.0 or later device, you do not need to re-enter the passphrase.
   - If the profile was imported from a device prior to version 12.0.0 and you plan to make changes to the profile (or if you associate the profile with a virtual server or a child profile), then you must supply the passphrase for the imported profile.
   - If you do not change any of the parameters for the profile or associate the profile with a virtual server or a child profile, then you do not need to re-enter the passphrase.

   ---
9. When your edits are complete, click **Save**.
   The system updates the profile with the settings you specified and adds it to the list of profiles.

   For detailed information the impact of using a particular profile parameter value, refer to the BIG-IP documentation on `support.f5.com`.

You must deploy your changes to the BIG-IP® device before you can see these changes on the device.

Attaching iRules to virtual servers

You can use the BIG-IQ® Local Traffic interface to attach iRules® to a set of virtual servers. Adding an iRule sequence to a group of servers at once can save time and help you cut down on errors that result from performing repetitious tasks.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on this device.
4. Select the check boxes associated with the virtual servers to which you want to attach iRules.
5. From the **Actions** button, select **Attach iRules**.
   The Virtual Servers - Attach iRules screen opens.
6. To specify which iRules to attach to the selected virtual servers, select them in the **Available iRules** list, and click the right arrow to add them to the **iRules to be Attached** list.
7. Specify the order in which you want the iRules to attach using the up and down arrows.
8. Specify the list position to attach these iRules.

- To add the rules to the beginning of the existing list, click **Attach to top of each virtual server's iRules list**.
- To add the rules to the end of the existing list, click **Attach to bottom of each virtual server's iRules list**.

9. Specify whether to keep the iRule list order for iRules that are already attached to the virtual servers.
   - To keep the existing list order, click **Keep virtual servers' existing rules list order**.
   - To change the existing list order to what you specified in step 2, click **Reorder virtual servers' existing rules to preserve selected rules order**.
10. Click **Save** and then confirm your choice by clicking **Modify**.
    A Modify Items dialog box pops up to show the status of your request.
11. Click **Close** to dismiss the dialog box and complete the process.

Creating a new iRule

You can use the BIG-IQ® Local Traffic interface to add a new iRule to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

Important: *Rules are different from most other Local Traffic objects in that they associate with virtual servers instead of devices. So to deploy a new iRule to a device, you attach the iRule to a virtual server associated with the target device and then deploy that change.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **iRules**.
   The screen displays a list of iRules® that are known on this device.
4. Click **Create**.
   The New iRule screen opens.
5. For **Name**, type a name for the iRule you are creating.
6. For **Partition**, type the name of the BIG-IP device partition on which you want to create the iRule.
7. For the **Body**, compose the script sequence that defines the iRule.

   For guidance on creating an iRule, consult the AskF5™ (`support.f5.com`) Knowledge Base. You can search the AskF5 website for iRules documentation that provides an overview of iRules, lists the basic elements that make up an iRule, and shows some examples of how to use iRules.
8. Click **Save & Close**.
   The system creates the new iRule with the settings you specified.

To deploy this iRule to a device, attach the iRule to a virtual server associated with the target device and then deploy that change.

Create a new pool

You can use the BIG-IQ® Local Traffic interface to add a pool to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. At the top of the screen, click **Configuration**.

2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of pools that are defined on this device.
4. Click **Create**.
   The New Pool screen opens.
5. In the **Name** field, type in a name for the pool you are creating.
6. From the **Device** list, select the device on which to create the pool.
7. For **Partition**, type the name of the BIG-IP device partition on which you want to create the pool.
8. In the **Description** field, type in a brief description for the pool you are creating.
9. To enable specific health monitors for this pool, select a monitor from the **Health Monitors** list.

   To add additional monitors click the + icon and repeat this step.
10. If you enabled specific monitors for this pool, for the **Availability Requirement** field, specify the minimum number of monitors that must report a pool as being available before the member is defined as being in an up state.
    - If all of the monitors must report the pool available, select **All**.
    - To specify a minimum number, select **At Least**, and then type the minimum number in the **Health Monitors** field.
11. In the **Load Balancing Method** field, specify the type of load balancing you want the pool to use. The default is **Round Robin**.
12. In the **Priority Group Activation** setting, specify how the system load balances traffic. The default is **Disabled**.
    a) To have the system load balance traffic according to the priority number assigned to the pool member, select **Less than**.
    b) If you use a priority number, from the **Available Member(s)** list, select the minimum number of members that must be available in one priority group before the system directs traffic to members in a lower priority group.

       When a sufficient number of members becomes available in the higher priority group, the system again directs traffic to the higher priority group.
13. To specify advanced properties, expand the Advanced Properties area and continue with the next twelve steps. Otherwise, click **Save & Close** now.
14. To automatically enable or disable NATs for connections that use this pool, for the **NAT** setting, select **Allow**.
15. To automatically enable or disable SNATs for connections that use this pool, for the **SNAT** setting, select **Allow**.
16. To specify how the system should respond when the target pool member becomes unavailable, select a value from the **Action On Service Down** list.

Option	Description
None	Specifies that the system takes no action to manage existing connections when a pool member becomes unavailable. The system maintains existing connections, but does not send new traffic to the member.
Reset	Specifies that, if there are no pool members available, the system resets and clears the active connections from the connection table and sends a reset (RST) or Internet Control Message Protocol (ICMP) message. If there are pool members available, the system resets and clears the active connections, but sends newly arriving connections to the available pool member and does not send RST or ICMP messages.
Drop	Specifies that the system simply cleans up the connection.
Reselect	Specifies that the system manages established client connections by moving them to an alternative pool member when monitors mark the original pool member down.

17. To specify the duration during which the system sends less traffic to a newly-enabled pool member, select a value from the **Slow Ramp Time** field.

    The amount of traffic is based on the ratio of how long the pool member has been available compared to the slow ramp time, in seconds. Once the pool member has been online for a time greater than the slow ramp time, the pool member receives a full proportion of the incoming traffic. Slow ramp time is particularly useful for the least connections load balancing mode.

    Important: *Setting this to a non-zero value can cause unexpected Priority Group behavior, such as load balancing to a low-priority member even with enough high-priority servers.*

18. To specify whether the system sets a Type of Service (ToS) level within a packet sent to the client, based on the targeted pool, select a value from the **IP ToS to Client** list.

    Setting a ToS level affects the packet delivery reliability.

Option	Description
Pass Through	The system does not change the ToS level within a packet.
Specify	Provides a field in which you can specify a ToS level to apply. Valid values are from `0` to `255`.
Mimic	Specifies that the system sets the ToS level of outgoing packets to the same ToS level of the most-recently received incoming packet. For example, if the most-recently received packet had a ToS level of 3, the system sets the ToS level of the next outgoing packet to 3.

19. To specify whether the system sets a Type of Service (ToS) level within a packet sent to the server, based on the targeted pool, select a value from the **IP ToS to Server** list.

    Setting a ToS level affects the packet delivery reliability.

Option	Description
Pass Through	The system does not change the ToS level within a packet.
Specify	Provides a field in which you can specify a ToS level to apply. Valid values are from `0` to `255`.
Mimic	Specifies that the system sets the ToS level of outgoing packets to the same ToS level of the most-recently received incoming packet. For example, if the most-recently received packet had a ToS level of 3, the system sets the ToS level of the next outgoing packet to 3.

20. To specify whether the system sets a Quality of Service (QoS) level within a packet sent to the client, based on the targeted pool, select a value from the **Link QoS to Client** list.

    Setting a QoS level determines the packet delivery priority.

Option	Description
Pass Through	The system does not change the QoS level within a packet.
Specify	Provides a field in which you can specify a QoS level to apply. Valid values are from `0` to `7`.

21. To specify whether the system sets a Quality of Service (QoS) level within a packet sent to the server, based on the targeted pool, select a value from the **Link QoS to Server** list.

    Setting a QoS level affects the packet delivery priority.

Option	Description
Pass Through	The system does not change the QoS level within a packet.

Option	Description
Specify	Provides a field in which you can specify a QoS level to apply. Valid values are from `0` to `7`.

22. To specify the number of times the system tries to contact a new pool member after a passive failure, select a value from the **Reselect Tries** field.

    A *passive failure* consists of a server-connect failure, or a failure to receive a data response within a user-specified interval. The default is `0`, which indicates no reselects.

23. To enable TCP request queueing, select **Request Queueing**.
24. To specify the maximum number of connection requests allowed in the queue, type an entry in the **Request Queue Depth** field.

    The default value of `0` permits unlimited connection requests, constrained only by available memory.

25. To specify the maximum number of milliseconds that a connection request can be queued until capacity becomes available, whereupon the connection request is removed from the queue and reset, type an entry in the **Request Queue Timeout** field.

    The default value of `0` permits unlimited time in the queue.

26. Click **Save & Close**.
    The system creates the new pool with the settings you specified.

Create a new pool member

You can use the BIG-IQ® Local Traffic interface to add a pool member to a pool.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of pools that are defined on this device.
4. Click the name of the pool to which you are going to add a new member.
   The properties screen for that pool opens.
5. Near the bottom of the screen, click the **New Member** button.
   The New Pool screen opens.
6. Specify the **Node Type**:
   - If you want the new member to be an existing BIG-IP® node, select **Existing Node** and then select the **Node**.
   - If you want the new member to be identified by an IP address, select **Address** and then type the **Node Name** and **Node Address** for the node.
7. For the **Port**, type the service port for the pool member.
8. In the **Description** field, type in a brief description for the pool member you are creating.
9. Specify the **Health Monitors** for this pool member.
   - To use the settings from the pool, select **Inherit from Pool**
   - To select specific health monitors for this pool member:
     1. Select **Member Specific**.
     2. Select a monitor from the **Health Monitors** list.
     3. To add additional monitors click the + icon and repeat this step
     4. If you activate more than one health monitor, specify the **Availability Requirement**. Either select **All**, or select **At Least**, and then type a number.

     Note: *This setting specifies the number of health monitors that must receive successful responses for the pool member to be considered available.*

10. For the **Ratio**, type the ratio weight you want to assign to the new pool member.

    When you use the ratio load balancing method, you can assign a ratio weight to each pool member in a pool. Local Traffic uses this ratio weight to determine the correct pool member for load balancing. Note that at least one pool member in the pool must have a ratio value greater than 1. Otherwise, the effect equals that of the Round Robin load balancing method.

11. If priority groups are enabled for this pool, type a **Priority Group** number for this member.

    Priority groups must be activated on the pool, if the number of available members for the highest priority group drops below your setting, the traffic is routed to the next highest member. If priority groups are disabled on the pool, this setting is not used.

12. For the **Connection Limit**, type the maximum number of concurrent connections allowed for this pool member.

13. For the **Connection Rate Limit**, type the maximum rate of new connections per second allowed for this pool member.

    When you specify this limit, the system controls the number of allowed new connections per second, thus providing a manageable increase in connections without compromising availability. The default value of `0` specifies that there is no limit on the number of connections allowed per second.

14. Click **Save & Close**.
    The system creates the new pool member with the settings you specified.

Delegate enable and disable permissions

To perform this task, you must log in as an Administrator.

You can assign permission to enable or disable virtual servers or pool members to other users. This allows those users to enable or disable specific virtual servers or pool members immediately, without having to deploy those changes.

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.
3. Click the **Add** button.
4. From the **Auth Provider** list, select **local (Local)** to have BIG-IQ authenticate this user.
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 this new locally-authenticated user.

   You can change the password any time.

8. Click **Save**.
   The system creates a new user.
9. On the left, click **USER MANAGEMENT** > **Roles**.
10. Click the **Add** button.
11. In the **Name** field, type a name to identify this role.
12. From the **Role Type** list, select the kind of role you want to add.

    - To create a role to which you can delegate virtual server permissions to immediately disable or enable virtual servers to which this role is assigned, select **Virtual Server Operator**.
    - To create a role to which you can delegate pool member permissions to immediately disable, enable or force offline pool members of pools to which this role is assigned, select **Pool Member Operator**.

    Permissions for specific virtual servers or pool members are not assigned to this role yet. You need to assign permissions for each object individually.

13. From the **Active Users and Groups** list, select the name of the user you specified in step7.
14. Click **Save**.
    The new role is created.
15. To delegate permissions for a virtual server, complete these sub-steps.
    a) At the top of the screen, click **Configuration**.
    b) On the left, expand **LOCAL TRAFFIC**.
    c) Under **LOCAL TRAFFIC**, select **Virtual Servers**.
    d) Click the name of the virtual server for which you wish to delegate permissions.
    The properties tab for the selected virtual server opens.
    e) Click **Permissions**.
    f) In the **Role** field, type the name of the role you specified in step 13.
    g) Click **Save**.
    The virtual server can now be enabled or disabled by a user logged in with the name you specified in step 7.
16. To delegate permissions for all of the pool members in a pool, do these sub-steps.
    a) Under **LOCAL TRAFFIC**, select **Virtual Servers**.
    b) On the left, expand **LOCAL TRAFFIC**.
    c) Under **LOCAL TRAFFIC**, select **Pools**.
    d) Click the name of the pool to which the pool member belongs.
    The properties tab for the selected pool opens.
    e) Click **Permissions**.
    f) In the **Role** field, type the name of the role you created in steps 13.
    g) Click **Save & Close**.
    Pool members in this pool can now be enabled, disabled, or forced offline by a user logged in with the name you specified in step 7.

Create a new node

You can use the BIG-IQ® Local Traffic interface to add a node to a managed device.

Nodes are the basis for creating a load balancing pool. For any server that you want to be part of a load balancing pool, you must first create a node, that is, designate that server as a node. After designating the server as node, you can add the node to a pool as a pool member. You can also associate a health monitor with the node, to report the status of that server.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

---

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Nodes**.
   The screen displays a list of nodes that are defined on this device.
4. Click **Create**.
   The New Node screen opens.
5. In the **Name** field, type in a name for the node you are creating.
6. From the **Device** list, select the device on which to create the node.
7. For the **Address** field, type in the IP address that identifies the new node.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the node.
9. In the **Description** field, type in a brief description for the node you are creating.

10. To specify configuration parameters for this node, expand **Configuration** and continue with the next steps. Otherwise, click **Save & Close**.
11. Specify the **Health Monitors** for this node.
    - If the BIG-IP® device uses the Node Default setting, select **Node Default**.

      Note: *The default monitor definition is set on the BIG-IP device. You can't revise that definition on the BIG-IQ. Consequently, the definition may well vary from device to device.*
    - To select specific health monitors for this node, select **Node Specific**, then select a monitor from the **Select Monitors** list.

      Note: *To add additional monitors click the + icon and repeat this step.*
12. If you selected **Node Specific**, for **Availability Requirement** specify the number of health monitors that must report a node as being available before the node is defined as being in an up state.
13. For the **Ratio**, type the ratio weight you want to assign to the new node.

    When you are using the Ratio load balancing method, you can assign a ratio weight to each node in a pool. LTM® uses this ratio weight to determine the correct node for load balancing. At least one node in the pool must have a ratio value greater than 1. Otherwise, the effect equals that of the Round Robin load balancing method.
14. For the **Connection Limit**, type the maximum number of concurrent connections allowed for this node.
15. For the **Connection Rate Limit**, type the maximum rate of new connections per second allowed for this node.

    When you specify this limit, the system controls the number of allowed new connections per second, thus providing a manageable increase in connections without compromising availability. The default value of `0` specifies that there is no limit on the number of connections allowed per second.
16. Click **Save & Close**.
    The system creates the new node with the settings you specified.

How do I manage LTM monitors in BIG-IQ?

With HTTP and HTTPS monitors you can track the availability of these services on the nodes, pools, or pool members to which you attach them. To add or edit monitors, you need to log in as an Administrator or ADC Editor.

Note: *You can revise the custom monitors, but you cannot edit the root monitors that ship with the product.*

Create an LTM monitor

You add a new HTTP or HTTPS LTM® monitor so that you can track the availability of these services on the nodes, pools, or pool members to which you attach that monitor.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Monitors**.
   The screen displays the list of monitors defined on this device.
4. Click **Create**.
   The New Monitor screen opens.
5. In the **Name** field, type in a name for the monitor you are creating.
6. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the monitor.
7. In the **Description** field, type in a brief description for the monitor you are creating.

8. For the **Type**, select the type of monitor you want to create.
   The **Monitor Template** setting displays.
9. From **Monitor Template**, select the parent monitor from which you want your monitor to inherit settings.
   A number of additional fields display. The fields that display depend on which monitor template you choose, **HTTP** or **HTTPS**.
10. For **Interval**, either use the default, or specify, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown.
11. From **Up Interval**, specify which interval the system uses to perform the health check when a resource is up.

Option	Description
Disabled	Specifies that the system uses the interval specified in **Interval** to check the health of the resource.
Enabled	Enables specification of a different interval to use when checking the health of a resource that is up.

12. For **Time Until Up**, specify the number of seconds to wait after a resource first responds correctly to the monitor before setting the resource to `up`.

    During the interval, all responses from the resource must be correct. When the interval expires, the resource is marked `up`. The default is 0, meaning that the resource is marked `up` immediately when the first correct response is received.
13. For **Timeout**, specify the number of seconds the target has in which to respond to the monitor request.

    The default is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Note that **Timeout** and **Time Until Up** combine to control when a resource is set to `up`.
14. For **Manual Resume**, specify whether the system automatically changes the status of a resource to `Enabled` at the next successful monitor check.

    If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections. The default is **No**.
15. For **Send String**, specify the text string that the monitor sends to the target object.

    You must include `\r\n` at the end of a non-empty **Send String**. The default setting is `GET /\r\n`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example:

```
GET /www/siterequest/index.html\r\n
```

16. For **Receive String**, specify a regular expression to represent the text string that the monitor looks for in the returned resource.

    The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names.

    Note: *If you do not specify both a **Send String** and a **Receive String**, the monitor performs a simple service check and connect only.*

    ---

17. For **Receive Disable String**, specify a regular expression to represent the text string that the monitor looks for in the returned resource.

    This setting works like **Receive String**, except that the system marks the node or pool member disabled when its response matches **Receive Disable String**.

    Note: *To use this setting, you must specify both **Receive String** and **Receive Disable String**.*

    ---

18. If you selected **HTTPS**, for **Cipher List**, specify the list of ciphers for this monitor.

    The default list is `DEFAULT:+SHA:+3DES:+kEDH`.

19. If the monitored target requires authentication, for the **User Name**, specify the user name.
20. If the monitored target requires authentication, for the **Password**, specify the password.
21. If you selected **HTTPS**, for **Compatibility**, specify the SSL option setting.

    If you select **Enabled**, the SSL option (in OpenSSL) is set to **ALL**.
22. If you selected **HTTPS**, for **Client Certificate**, select the client certificate that the monitor sends to the target SSL server.

    The default is **None**.
23. If you selected **HTTPS**, for **Client Key**, select the key for the client certificate that the monitor sends to the target SSL server.

    The default is **None**.
24. For **Reverse**, specify whether the system marks the target resource down when the test is successful.

    This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you may want to set up a reverse ECV service check that looks for the string Error. A match for this string means that the web server was down. To use this option, you must specify values for **Send String** and `Receive String`.
25. For **Transparent**, specify whether the system operates in transparent mode.

    A monitor in transparent mode directs traffic through the associated pool members or nodes (usually a router or firewall) to the aliased destination (that is, it probes the Alias Address-Alias Service Port combination specified in the monitor). If the monitor cannot successfully reach the aliased destination, the pool member or node through which the monitor traffic was sent is marked `down`.
26. For **Alias Address**, specify an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    The default setting is ***All Addresses**. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
27. For **Alias Service Port**, specify an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    The default setting is ***All Ports**. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
28. For **IP DSCP**, specify the differentiated services code point (DSCP).

    DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service wanted for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default is 0 (zero).
29. For **Adaptive**, specify whether adaptive response time monitoring is enabled for this monitor.

Option	Description
Enabled	The monitor determines the state of a service based on how divergent from the mean latency a monitor probe for that service is allowed to be. When enabled, you can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings.
Disabled	The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.

    If you select **Enabled** for this control, three additional controls are displayed.
30. If you enabled **Adaptive**, for **Allowed Divergence**, specify the type of divergence used for adaptive response time monitoring.

Option	Description
Absolute	The number of milliseconds the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed. Tip: In typical cases, if the monitor detects three probes in a row that miss the latency value you set, the pool member or node is marked `down`.
Relative	The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.

31. If you enabled **Adaptive**, for **Allowed Divergence**, specify the absolute number of milliseconds that may not be exceeded by a monitor probe, regardless of **Allowed Divergence** for a probe to be considered successful.

    This value applies regardless of the value of the **Allowed Divergence** setting.

32. If you enabled **Adaptive**, for **Sampling Timespan**, specify the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe.

Edit an LTM monitor

You revise HTTP or HTTPS LTM® monitors when you want to change the details of how the monitor determines when a service is operational.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Monitors**.
   The screen displays the list of monitors defined on this device.
4. Select the monitor you want to edit.
   The Monitors screen opens to display the current settings for the selected monitor.
5. In the **Description** field, add or revise a brief description for the monitor you are editing.
6. From **Interval**, specify, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown.

   The default is 5 seconds.

7. For **Up Interval**, specify which interval the system uses to perform the health check when a resource is up.

Option	Description
Disabled	Specifies that the system uses the interval specified in **Interval** to check the health of the resource.
Enabled	Enables specification of a different interval to use when checking the health of a resource that is up.

8. For **Time Until Up**, specify the number of seconds to wait after a resource first responds correctly to the monitor before setting the resource to up.

   During the interval, all responses from the resource must be correct. When the interval expires, the resource is marked up. The default is 0, meaning that the resource is marked up immediately when the first correct response is received.

9. From **Timeout**, specify the number of seconds the target has in which to respond to the monitor request.

   The default is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Note that **Timeout** and **Time Until Up** combine to control when a resource is set to up.

10. For **Manual Resume**, specify whether the system automatically changes the status of a resource to `Enabled` at the next successful monitor check.

If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections. The default is **No**.

11. For **Send String**, specify the text string that the monitor sends to the target object.

    You must include `\r\n` at the end of a non-empty **Send String**. The default setting is `GET /\r\n`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example:

```
GET /www/siterequest/index.html\r\n
```

12. For **Receive String**, specify a regular expression to represent the text string that the monitor looks for in the returned resource.

    The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names.

    Note: *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*

    ---

13. For **Receive Disable String**, specify a regular expression to represent the text string that the monitor looks for in the returned resource.

    This setting works like **Receive String**, except that the system marks the node or pool member disabled when its response matches **Receive Disable String**.

    Note: *To use this setting, you must specify both* ***Receive String*** *and* ***Receive Disable String****.*

    ---

14. If you selected **HTTPS**, for **Cipher List**, specify the list of ciphers for this monitor.

    The default list is `DEFAULT:+SHA:+3DES:+kEDH`.

15. If the monitored target requires authentication, for the **User Name**, specify the user name.
16. If the monitored target requires authentication, for the **Password**, specify the password.

    Important: *For imported monitors that use passwords:*

    - If the monitor was imported from a version 12.0.0 or later device, you do not need to re-enter the password.
    - If the monitor was imported from a device earlier than version 12.0.0 and you plan to make changes to the monitor (or if you associate the monitor with an LTM object or child monitor), then you must supply the password for the imported monitor.
    - If you do not change any of the parameters for the monitor or associate the monitor with an LTM object or child monitor, then you do not need to re-enter the password.

    ---

17. If you selected **HTTPS**, for **Compatibility**, specify the SSL option setting.

    If you select **Enabled**, the SSL option (in OpenSSL) is set to **ALL**.

18. If you selected **HTTPS**, for **Client Certificate**, select the client certificate that the monitor sends to the target SSL server.

    The default is **None**.

19. If you selected **HTTPS**, for **Client Key**, select the key for the client certificate that the monitor sends to the target SSL server.

    The default is **None**.

20. For **Reverse**, specify whether the system marks the target resource down when the test is successful.

    This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you may want to set up a reverse ECV service check that looks for the string Error. A match for this string means that the web server was down. To use this option, you must specify values for **Send String** and **Receive String**.

21. For **Transparent**, specify whether the system operates in transparent mode.

A monitor in transparent mode directs traffic through the associated pool members or nodes (usually a router or firewall) to the aliased destination (that is, it probes the Alias Address-Alias Service Port combination specified in the monitor). If the monitor cannot successfully reach the aliased destination, the pool member or node through which the monitor traffic was sent is marked down.

22. For **Alias Address**, specify an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    The default setting is ***All Addresses**. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.

23. For **Alias Service Port**, specify an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    The default setting is ***All Ports**. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

24. For **IP DSCP**, specify the differentiated services code point (DSCP).

    DSCP is a 6-bit value in the Differentiated Services (DS) field of the IP header. It can be used to specify the quality of service desired for the packet. The valid range for this value is 0 to 63 (hex 0x0 to 0x3f). The default is 0 (zero).

25. For **Adaptive**, specify whether adaptive response time monitoring is enabled for this monitor.

Option	Description
Enabled	The monitor determines the state of a service based on how divergent from the mean latency a monitor probe for that service is allowed to be. When enabled, you can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings.
Disabled	The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.

    If you select **Enabled** for this control, three additional controls are displayed.

26. If you enabled **Adaptive**, for **Allowed Divergence**, specify the type of divergence used for adaptive response time monitoring.

Option	Description
Absolute	The number of milliseconds the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed. Tip: In typical cases, if the monitor detects three probes in a row that miss the latency value you set, the pool member or node is marked down.
Relative	The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.

27. If you enabled **Adaptive**, for **Allowed Divergence**, specify the absolute number of milliseconds that may not be exceeded by a monitor probe, regardless of **Allowed Divergence**.

    For a probe to be considered successful, this value applies regardless of the value of the **Allowed Divergence** setting.

28. If you enabled **Adaptive**, for **Sampling Timespan**, length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe.

29. Click **Save & Close**.

Create a new SNAT pool

You can use the BIG-IQ® Local Traffic interface to add a SNAT pool to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same*

changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, click **SNAT Pools**.
   The screen displays a list of SNAT translation members that are defined on this device.
4. Click **Create**.
   The New SNAT Pool screen opens.
5. In the **Name** field, type a name for the SNAT pool you are creating.
6. From the **Device** list, select the device on which to create the SNAT pool.
7. In the **Member List**, type the IP address of the first SNAT translation member you want to include in the SNAT pool.

   Use the + key to add more members, or you can use the **x** key to delete a member.
8. In the **Partition** field, type the name of the partition in which you want to create this SNAT pool.

   An *administrative partition* is a logical container that you create that contains a defined set of BIG-IP system objects. If you enter a partition name that does not exist, you will get an error when you try to deploy this SNAT pool.
9. Click **Save & Close**.
   The system creates the new SNAT pool with the settings you specified.

Changing network objects

You can make revisions to the configuration of Local Traffic objects to simplify managing your devices.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select the object type that you want to modify.
   The screen displays a list of objects of that type that are defined on this BIG-IQ®.
4. Click the name of the object you want to change.
   The Properties screen for the selected object opens.
5. Make changes to the properties that you want to modify.
6. When you are satisfied with the changes you have made, click **Save**.
   The revisions you saved are made, and the Properties screen for the selected object closes.

Changes that you make are made only to the pending version. The *pending version* serves as a repository for changes you stage before deploying them to the managed device. Object settings for the pending version are not the same as the object settings on the actual BIG-IP® device until they are deployed or discarded.

Important: *There is an exception to this pattern. When you view properties for a pool member and click* ***Enable****,* ***Disable****, or* ***Force Offline****, you can choose whether you want the change to occur immediately (****Change Now****), later (****Change Later****), or not at all (****Cancel****). Changes you decide to make later become part of the pending changes for the managed object.*

To apply the pending version settings to the BIG-IP device, you next need to deploy the revisions.

Manage a network interface

You can use the BIG-IQ® Local Traffic component to enable or disable network interfaces on a managed device.

Important: *When you revise configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, click **Interfaces**.
   The screen displays a list of network interfaces defined on devices that are managed by this BIG-IQ.
4. Select the interface you want to change and then select **Enable** or **Disable**.
   The **State** for the selected interface changes on the BIG-IQ.

Create a new route

You can use the BIG-IQ® Local Traffic component to add a route to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select **Routes**.
   The screen displays a list of routes defined on devices that are managed by this BIG-IQ.
4. Click **Create**.
   The New Route screen opens.
5. In the **Name** field, type in a name for the route you are creating.
6. In the **Description** field, type in a brief description for the route you are creating.
7. From the **Device** list, select the device on which to create the route.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route.
9. In the **Destination/Mask** field, type a self IP address and net mask for this route.

   These addresses display in the Destination and Netmask columns of the routing table.

   For example:

```
10.145.193.0/24
```

   .
10. Specify the **Resource** setting.
    - To use a gateway, select **Use Gateway**, and then choose either **IP Address** or **IPv6 Link-Local Address** through which you want the BIG-IQ system to forward packets to the route destination.
    - To use a pool, select **Use Pool**, and then select the pool through which you want the BIG-IQ system to forward packets to the route destination.
    - To use a VLAN or tunnel, select **Use VLAN/Tunnel**, and then select the VLAN or tunnel through which you want the BIG-IQ system to forward packets to the route destination.
    - To use reject packets forwarded to the route destination, select **Reject**.
11. In the **MTU** field, type an optional frame size value for Path Maximum Transmission Unit (MTU). By default, BIG-IP® devices use the standard Ethernet frame size of 1518 bytes (1522 bytes if VLAN tagging is used) with the corresponding MTU of 1500 bytes. For BIG-IP devices that support Jumbo Frames, you can specify another MTU value.
12. Click **Save & Close**.

The system creates the new route with the settings you specified.

Create a new self IP address

You can use the BIG-IQ® Local Traffic component to add a self IP address to a managed device.

Important: *When revising configurations on devices that belong to a high availability cluster, it is important to let the changes synchronize to the cluster members instead of trying to make the same changes to multiple devices. If you try to replicate changes you made on one device in the cluster, the next config sync attempt could fail.*

---

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select **Self IPs**.
   The screen displays a list of self IP addresses defined on the devices that are managed by this BIG-IQ.
4. Click **Create**.
   The New Self IP screen opens.
5. In the **Name** field, type in a name for the self IP address you are creating.
6. From the **Device** list, select the device on which to create the self IP address.
7. For **Partition**, type the name of the BIG-IP device partition on which you want to create the self IP.
8. In the **IP Address** field, type either an IPv4 or an IPv6 address. For an IPv4 address, you should specify a /32 IP address per RFC 3021.
9. In the **Netmask** field, type the netmask for this self IP address. You must type the full netmask.
   Specifying the prefix length in bits is not supported. For example, you could type `255.255.255.255` or `ffff:ffff:ffff:ffff:0000:0000:0000:0000 or ffff:ffff:ffff:ffff::` (with two colons at the end).
10. For the **VLAN/Tunnel**, select the VLAN or tunnel to associate with this self IP address.
11. Specify the **Port Lockdown**.
    - Select **Allow Default** to activate only the default protocols and services. You can determine the supported protocols and services by logging in to the target BIG-IP device and running `tmsh list net self-allow defaults` on the command line.
    - Select **Allow All** to activate all TCP and UDP services on this self IP address.
    - Select **Allow None** to specify that this self IP address accepts no traffic. If you are using this self IP address as the local endpoint for WAN optimization, select this option to avoid potential port conflicts.
    - Select **Allow Custom** or **Allow Custom (Include Default)** to expand the **Custom List** setting, where you can specify the ports, protocols, and services to activate on this self IP address.
12. For the **Traffic Group**, select a specific traffic group for the self IP address.
13. Click **Save & Close**.
    The system creates the new self IP address with the settings you specified.

Create a new route domain

You can use the BIG-IQ® Local Traffic component to add a route domain to a managed device. Using route domains, you can assign the same IP address to more than one device on a network, as long as each instance of the IP address resides in a separate route domain.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select **Route Domains**.
   The screen displays a list of route domains defined on the devices that are managed by this BIG-IQ.

4. Click **Create**.
   The New Route Domain screen opens.
5. In the **Name** field, type in a unique name for the route you are creating.
6. In the **ID** field, type an integer to represent the route domain.

   The integer must be unique on the BIG-IP® device and be between 1 and 65534. The default value (0) indicates that all VLANs on a system pertain to this route domain. When you create new route domains, you can assign VLANs to those route domains which moves the VLANs out of the default route domain.
7. In the **Description** field, type in a brief description for the route domain you are creating.
8. From the **Device** list, select the device on which to create the route domain.
9. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route domain.
10. Select **Strict Isolation** if you want to enforce cross-routing restrictions.

    When selected, routes cannot cross route domain boundaries (so they are strictly isolated to the current route domain). The default is enabled. When disabled, routes can cross route domains. For example, you could add a route to the routing table with a 10.0.0.0%20 (route domain 20) destination and a gateway of 172.27.84.29%32 (route domain 32).
11. To specify a VLAN or tunnel for the BIG-IP to use in the route domain, select it in the **Available** list, and click the right arrow to add it to the **Enabled** list.
12. Click **Save & Close**.
    The system creates the new route domain with the settings you specified.

Create a new VLAN

You can use the BIG-IQ® Local Traffic component to add a VLAN to a managed device. Using VLANs, you can assign the same IP address to more than one device on a network, as long as each instance of the IP address resides in a separate VLAN.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select **VLANs**.
   The screen displays a list of VLANs defined on the devices that are managed by this BIG-IQ.
4. Click **Create**.
   The New VLAN screen opens.
5. In the **Name** field, type a unique name for the VLAN you are creating.
6. In the **Description** field, type a brief description for the VLAN you are creating.
7. In the **Tag** field, type a tag number for the VLAN.

   The tag number can be between 1 and 4094, but must be unique on the target device. If you do not specify a value, the system automatically assigns a tag number.
8. From the **Device** list, select the device on which to create the VLAN.
9. For **Partition**, type the name of the BIG-IP device partition on which you want to create the VLAN.
10. In the **MTU** field, specify the maximum transmission unit (MTU) for traffic on this VLAN.

    The default is 1500.
11. To specify which interfaces this VLAN uses for traffic management, select it in the **Interface** list, and then select the **Tagging** for it.
12. Click **Save & Close**.
    The system creates the new VLAN with the settings you specified.

Creating a new DNS resolver

You can use the BIG-IQ® Local Traffic component to add a DNS resolver to a managed device. Using DNS resolvers, you can assign the same IP address to more than one device on a network, as long as each instance of the IP address resides in a separate DNS resolver.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK**.
3. Under **NETWORK**, select **DNS Resolvers**.
   The screen displays a list of DNS resolvers defined on the devices that are managed by this BIG-IQ.
4. Click **Create**.
   The New DNS Resolver screen opens.
5. In the **Name** field, type in a unique name for the DNS resolver you are creating.
6. For **Partition**, type the name of the BIG-IP device partition on which you want to create the DNS resolver.
7. To specify which devices use this DNS resolver for traffic management, in the **Devices** setting, select them in the **Available** list, and use the right arrow to move them to the **Selected** list.
8. Select the **Route Domain Name** that this resolver uses for outbound traffic.

   The default is the default route domain.
9. To specify the Resolver properties, expand the control and then:
   a) For the **Cache Size**, type the size of the internal DNS resolver cache.

      The default is 5767168 bytes. After the cache reaches this size, when new or refreshed content arrives, the system removes expired and older content and caches the new or updated content.
   b) Select **Answer Default Zones** if you want the system to answer DNS queries for the default zones `localhost`, `reverse`, `127.0.0.1`, `::1`, and `AS112`.

      The default is disabled, meaning that the system passes along the DNS queries for the default zones.
   c) Select **Randomize Query Character Case** if you want the internal DNS resolver to randomize character case in domain name queries issued to the root DNS servers.

      The default is enabled.
10. To specify the Traffic properties, expand the control and select the format or formats for which you want the system to answer and issue queries:
11. To specify a forward zone used to resolve matching DNS queries, expand the control and Click **Add**.

    A popup screen opens.
    a) In the **Name** field, type in a unique name for the forward zone you are creating.
    b) In the **Address** field, type in an IP address for the forward zone you are creating.
    c) In the **Service Port** field, type in the port number for the forward zone you are creating.
    d) Click the **Add** button next to the Service Port.
       The address and port combination is added to the Nameservers box.
    e) To add additional Nameservers, repeat the last two sub-steps.
12. When you are satisfied with the new forward zone, click the **Add** button.
13. If you have specified forward zones, select the check boxes for the zones you want to use.
14. When you are satisfied with the new DNS resolver, click **Save & Close**.
    The system creates the new DNS resolver with the settings you specified.

When the BIG-IP® system receives a query that cannot be resolved from the cache, the system forwards the query to a nameserver associated with the matching forward zone. When the nameserver returns a response, the BIG-IP system caches the response, and returns the response to the resolver making the query.

Users, User Groups, Roles, and Authentication

How do I manage user access to BIG-IQ?

As a network or system manager, you need a way to differentiate between users, and to limit user access based on how they interact with F5® BIG-IQ® Centralized Management and your managed devices.

You can specify how you want users to be authenticated: locally on BIG-IQ, or remotely through your RADIUS or LDAP server. Additional security is provided through bidirectional trust and verification through key and certificate exchange (AuthN and AuthZ).

To help you manage all of this, it's important that you understand the following concepts:

- *Users* - are individuals for whom you are providing access to BIG-IQ resources, including access to managed BIG-IP® devices.
- *User groups* - are a way to organize individuals into groups so that you can grant or change the same privileges to several users at once.
- *Roles* - are associated with specific privileges, which you grant to users, allowing them to do a set of tasks on BIG-IQ, and on your managed devices.

Use my LDAP server to authenticate BIG-IQ users

F5® BIG-IQ® Centralized Management can verify user credentials against your company's LDAP server (LDAP server versions 2 and 3, and OpenLDAP directory, Apache Directory Server, and Active Directory). After you set up BIG-IQ to use your LDAP server, you can add users and user groups that authenticated by your LDAP server.

Set up BIG-IQ to use my RADIUS server for user authentication

Before you can set up authentication, you must have specified your DNS settings. You usually do this when you license F5® BIG-IQ® Centralized Management.

You can set up BIG-IQ to use your company's RADIUS server. You can add two additional backup RADIUS servers in case the primary server is not available for authentication.

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Auth Providers**.
3. Click the **Add** button.
4. From the **Provider Type** list, select **RADIUS.**
5. In the **Name** field, type a name for this new provider.

   This must be a unique name and can be a maximum of 152 characters.
6. In the **Host** and **Port** fields, type the RADIUS server's IP address (or fully qualified domain name) and port number for each of the servers you want to configure.

   The primary server is mandatory. A secondary server and tertiary server, which will be used if the primary or secondary servers fail, are optional.
7. In the **Secret** field, type the case-sensitive text string used to validate communication.
8. In the **Test User** and **Test Password** fields, type a user and password, then click the **Test** button to verify that BIG-IQ can reach the RADIUS server
9. Click the **Save & Close** button at the bottom of the screen.

You can now associate RADIUS server users and groups with BIG-IQ system roles.

Before integrating BIG-IQ with your LDAP server

Before integrating LDAP authentication with the BIG-IQ® system, you must first perform the following tasks:

- Use an LDAP browser to review the groups and users in your directory's structure and where they're located in the hierarchy of organizational units (OUs).
- Decide how you want to map user names.
  - The first option is to map users directly to their Distinguished Name (DN) in the directory with a user bind template in the form of `uid=<username>, ou=people,o=sevenSeas`. For example, when you map John Smith's user name with his DN as `uid=<jsmith>, ou=people,o=sevenSeas` and he logs in as `jsmith`, he is correctly authenticated with his user name in the directory through his DN.
  - The second option is to allow users to log in with names that do not map directly to their DN by specifying a `userSearchFilter` in the form of `(&(uid=%s))` when creating the provider. For example, if John Smith's DN is `cn=John Smith,ou=people,o=sevenSeas`, but you would like him to be able to log in with `jsmith`, specify a `userSearchFilter` in the form of `(&(jsmith=%s))`. If your directory does not allow anonymous binds, you must also specify a `bindUser` and `bindPassword` so that the BIG-I system can validate the user's credentials.
- Decide which groups in your directory to map into BIG-IQ groups.
  - If you configured a `bindUser` and `bindPassword` for users, the BIG-IQ system displays a list of groups from which to choose.
  - If you haven't configured this for your users, you must know the DN for each group.
- Find out the DN where you can for all users and groups. This is the root bind DN for your directory, defined as as `rootDN`, when you create a provider. The BIG-IQ system uses the root bind DN as a starting point when it searches for users and groups.
- Find the host IP address for the LDAP server. The default port is 389, if not specified otherwise.

Set up BIG-IQ to use your LDAP server for user authentication

Before you can set up BIG-IQ to authenticate users against your LDAP server, you have to specify your LDAP server settings on F5® BIG-IQ® Centralized Management and perform all the tasks outlined in the section titled, *Before integrating BIG-IQ with your LDAP server*.

You can configure BIG-IQ to use one or more of your company's LDAP server(s) to authenticate users.

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Auth Providers**.
3. Click the **Add** button.
4. From the **Provider Type** list, select **LDAP**.
5. In the **Name** field, type a name for this new provider.

   This must be a unique name and can be a maximum of 152 characters.
6. In the **Host** field, type the IP address of your LDAP server.
7. For the **Port** setting, select the port your Active Directory server uses.

   If you want BIG-IQ to use an SSL port to communicate with your LDAP server, select port **636**, otherwise leave it at the default port, **389**.
8. To use an SSL port to communicate with the LDAP server, for the **SSL Enabled** setting select the **Enabled** check box.
9. If your LDAP server does not allow anonymous binds, in the **Bind User** and **Bind User Password** fields, type the full distinguished names and passwords for users with query access.
10. In the **Root DN** field, type the root context that contains users and groups.

    The root context must be a full distinguished name.

11. From the **Authentication Method** list, select an option.
    - **Simple** - Select this option to require a user name and password for authentication.
    - **None** - Select this option to prompt the LDAP server to ignore the user name and password.
12. In the **Search Scope** field, type a number to specify the depth at which searches are made.

    Alternatively, you can specify `0` for search only on the named object or `1` for a one-level search scope.
13. In the **Search Filter** field, type the LDAP filter expression that determines how users are found.

    The search filter is determined by your LDAP implementation.
14. In the **Connect Timeout** field, type the number of milliseconds after which the BIG-IP system stops trying to connect to the LDAP server.
15. In the **Read Timeout** field, type the number of seconds the BIG-IP system will wait for a response to a query.
16. In the **User Display Name Attribute** field, type the LDAP field to use for the name that BIG-IQ displays.

    When using Active Directory, this is typically `displayName`.
17. To direct bind to a distinguished name, in the **User Bind Template** field, type the name.
    For example, `cn={username},ou=people,o=sevenSeas`.

    Now, when a user logs in, BIG-IQ inserts the user name into the template in place of the token, and the resulting distinguished name is used to bind to the directory.
18. To prompt the LDAP provider to search for groups based on a specific display name attribute, in the **Group Display Name Attribute** field, type an attribute.

    This attribute is typically `cn`.
19. Leave the **Group Search Filter** at its default query to return all groups under the provided rootDN.

    Alternatively, if you have a large number of groups (more than 100), you can base the search on a specific term by typing a query with a `{searchterm}` token in this field.

    For example: `(&(objectCategory=group)(cn={searchterm}*))`
20. To specify a query for finding a users group, in the **Group Membership Filter** field, type a query string.

    Use the token `{userDN}` anywhere that the user's distinguished name should be supplied in the LDAP query.

    You can use a `{username}` token as a substitute for the user's login name in a query.

    Leave this setting at the default `(|(member={username})(uniqueMember={username}))` unless the provider is Active Directory.
21. To specify a query attribute for finding users in a particular group, in the **Group Membership User Attribute** field, type the attribute.

    When using Active Directory, use `memberof`. For example:
    `(memberOf=cn=group_name,ou=organizational_unit,dc=domain_component)`

    For other LDAP directories, use `groupMembershipFilter`. For example:
    `(groupMembership=cn=group_name,ou=organizational_unit,o=organization)`
22. Select the **Perform Test** check box to test this provider.
23. Click the **Save & Close** button at the bottom of the screen.

Pre-defined RADIUS groups for authentication

You must have root access to the BIG-IQ system's command line through SSH for this procedure.

Some RADIUS deployments include non-standard, vendor-specific attributes in the dictionary files. For these deployments, you must update the BIG-IQ system's default dictionary. Follow these steps if you want to use pre-defined RADIUS user groups on BIG-IQ.

1. Copy the TinyRadius .jar file from the BIG-IQ system.

2. Extract the contents of the TinyRadius `.jar` file.
3. Update the file `org/tinyradius/dictionary/default_dictionary` file, by adding the vendor-specific attributes.
4. Repack the contents into a new `.jar` file.
5. Replace the old TinyRadius `.jar` on each BIG-IQ system with the new TinyRadius .jar file you created in step 4.

For example:

1. From a Linux machine, copy the TinyRadius `.jar` file to your BIG-IQ system by typing: `scp <big-iq-user>@<BIG-IQ-Address>:/usr/share/java/TinyRadius-1.0.jar ~/tmp/tinyrad-upgrade/`
2. Extract the file on your Linux Machine by typing: `jar -xvf TinyRadius-1.0.jar`
3. Edit the `org/tinyradius/dictionary/default_dictionary`, adding the vendor-specific attribute.

```
rm TinyRadius-1.0.jar
jar cvf TinyRadius-1.0.jar *
```

4. Update the jar on the BIG-IQ system by typing: `scp TinyRadius-1.0.jar <your_user>@<BIG-IQ address>:/var/tmp/`
5. SSH to the BIG-IQ system and type the following commands:

```
mount -o remount,rw /usr
cp /var/tmp/TinyRadius-1.0.jar /usr/share/java
mount -o remount,ro /usr
bigstart restart restjavad
```

6. Repeat steps 4 and 5 for each BIG-IQ in a HA configuration.

Now you can use the vendor-specific attributes RADIUS to create your user groups on BIG-IQ.

Add a user and assign them a role

Once you understand exactly who you want to perform certain tasks, you can provide them access to particular areas of F5® BIG-IQ® Centralized Management by adding them as a user and assigning the appropriate standardized role. You can assign as many roles as required to cover the user's responsibilities.

Important: *Since some roles have access only to certain areas or screens in the BIG-IQ user interface, it's important to communicate that to the user. When you assign a role to a user, be sure you outline the responsibilities and restrictions for their role. Clarifying this helps avoid any potential confusion. Also note,these roles do not have access to the global search functionality: Network Security Manager, Network Security Edit, Network Security View, and Trust Discovery Import.*

---

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.
3. Click the **Add** button.
4. From the **Auth Provider** list, select the authentication method you want to use for this user.
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 this new locally-authenticated user.

   You can change the password any time.

8. To associate this user with an existing user group, select the group from the **User Groups** list.

   You aren't required to associate a user group at this point; you can do that later if you want. If you want to associate another user group with this user, click +.

9. From the **User Roles** list, select a user role to associate with this user.

   Each role has a set of unique privileges. If you want to associate another user role with this user, click +.

   Important: *Be sure to let your users know that their access to certain parts of the BIG-IQ user interface depends on which role they are assigned.*

10. Click the **Save & Close** button at the bottom of the screen.

This user now has the privileges associated with the role(s) you selected and BIG-IQ will authenticate this user locally

Let this user know how their BIG-IQ access aligns with their responsibilities and make sure they understand they might not see every screen you or one of their peers does.

Note: *If your BIG-IQ is in an HA pair, you must synchronize this change by refreshing the secondary BIG-IQ.*

How do I limit privileges for users based on their role in the company?

F5® BIG-IQ® Centralized Management provides you the tools you need to customize user access to your managed devices, and to BIG-IQ itself, through the use of role-based privileges. These privileges are based on the responsibilities of your users.

This type of role-specific access also provides you insight into your work flows. You can easily see which user interacted with any given service, and what the interaction was. This can help you quickly troubleshoot any introduced conflicts.

You can set up BIG-IQ to authorize users, giving them access only to the specific information, using these methods:

- **Local authorization** - for this option, BIG-IQ authenticates users.
- **External authorization** - for this option, you can configure BIG-IQ to use your LDAP or RADIUS server to authenticate users.

The responsibilities and roles each of your users has probably depend on the number of people who have access to BIG-IQ.

Assigning more than one role to a user

For example, if you have only two people managing your devices from BIG-IQ, they both most likely need to have full access to all aspects of BIG-IQ at one time or another. For these users, you'd assign them both the Administrator role.

Assigning more granular/specialized privileges to a user

On the other hand, if you're working for a larger company that has specialized roles to manage different services, or different parts of services, you can provide more granular access. For example, if you have two people who manage BIG-IP devices used only for network security purposes, you could assign them both the role of Network Security Manager. Or, if you have two people managing devices used for network security, but you want only one of them to write and edit policies, and the other to (only) deploy the policies, you could assign the first person the Network Security Editor role, and the other person the Network Security Deploy role. In this case, the Network Security Editor can only create, view, and edit

policies, but not deploy them. The Network Security Deploy person can view and deploy policies, but cannot create or edit them.

Standard roles shipped with BIG-IQ

As a system manager, you'll need a way to limit a user's access to certain areas of F5® BIG-IQ® Centralized Management and to its managed devices. The easiest way to do this is to base user access on the responsibilities, or role, the user has in your company. To help you do that, BIG-IQ ships with a set of default roles with certain privileges that you can assign to specific users. Since responsibilities and duties for certain roles are specialized, users assigned to some roles have access to only specific parts of BIG-IQ. These restrictions are outlined in the role description.

Role	This role can:
Administrator	Perform all tasks for setting up and maintaining BIG-IQ and managing devices. This includes discovering devices, adding individual users, assigning roles, installing updates, activating licenses, and so forth.
ADC Deployer	View and deploy ADC configuration objects for managed ADC devices.
ADC Editor	Create and edit all ADC configuration objects.
ADC Manager	Perform all tasks for managing ADC, including creating, viewing, modifying, and deleting Local Traffic and Network objects.
ADC Viewer	Only view all ADC objects and features.
Access Auditor	Only view Access configuration objects and managed Access devices. This role cannot edit, discover, or deploy devices or policies.
Access Deployer	Deploy Access configuration objects. This role cannot discover and edit devices or policies.
Access Editor	View and edit Access configuration objects, including the ability to add, update, and delete pools and pool members from the Access configuration object editor. This role cannot discover or deploy devices or policies.
Access Manager	Deploy and edit Access configuration objects, and view the Access Reporting and dashboard. This role cannot add or remove devices and device groups, and cannot discover, import, or delete services.
Access Viewer	Only view Access configuration objects and discovered Access devices. This role cannot edit, discover, or deploy devices or policies.
Device Manager	Perform all tasks for device management, including device discovery, licensing, software image management, UCS backups, templates, self IP addresses, VLANs, interfaces, and so forth.
Device Viewer	Only view aspects of device management including areas involved in device discovery, group creation, licensing, software image management, UCS backups, templates, self IP addresses, VLANs, interfaces, and so forth.
DNS Viewer	Only view aspects of device management associated with DNS.
Fraud Protection Manager	Perform all tasks for managing the Fraud Protection Service Fraud Protection Service functionality.
Fraud Protection View	Only view Fraud Protection Service Fraud Protection Service objects.
Network Security Deploy	View and deploy Network Security objects.

Role	This role can:
Network Security Manager	Perform all tasks associated with Network Security, including areas involved in creating, viewing, modifying, and deleting shared and firewall-specific security objects. This role does not have access to the global search functionality.
Network Security Edit	Create, view, modify, and delete Network Security objects. This role does not have access to the global search functionality.
Network Security View	Only view Network Security firewall objects. This role cannot edit, discover, or deploy devices or policies. This role does not have access to the global search functionality.
Security Manager	Perform all tasks associated with Network Security, Web Application Security, and Fraud Protection Service, including areas involved in device discovery, creating, viewing, modifying, and deleting Web Application Security, shared and firewall-specific security objects.
Trust Discovery Import	Manage device trust establishment, service discovery, service import, removal of services and removal of trust. This role does not have access to the global search functionality.
Web App Security Deployer	View and deploy ASM configuration objects for managed ASM devices.
Web App Security Editor	Create, view, modify, and delete ASM configuration objects.
Web App Security Manager	View and edit all aspects of Web Application Security, including areas involved in creating, viewing, modifying, and deleting shared and web application-specific security objects.
Web App Security Viewer	Only view ASM configuration objects.

Adding a role

In addition to the standard roles that ship with BIG-IQ®, you can also add some roles that are specific to ADC and device management only.

1. At the top of the screen, click **Devices**.
2. On the left, click **USER MANAGEMENT** > **Roles**.
3. Click the **Add** button.
4. In the **Name** field, type a name to identify this new role.
5. From the **Role Type** list, select the kind of role you want to add.
6. Click the + sign if you want this role to have access to another user or group, and select the device group from the list.
7. From the **Active Users and Groups** list, select the user or group you want to associate with this new role.
8. Click the **Save & Close** button at the bottom of the screen.

Change your BIG-IQ user password

For security reasons, you need to occasionally change your user password.

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.
3. Click your user name.

4. In the **Old Password** field, type the password.
5. In the **Password** and **Confirm Password** fields, type a new password.
6. Click the **Save & Close** button at the bottom of the screen.

Provide a user access to BIG-IQ and assign them a role

When you add a user to a role, you have to specify an authentication method. So, if you want to use LDAP or RADIUS server for authentication, you need to configure that on BIG-IQ® Centralized Management first.

Once you understand exactly who you want to perform certain tasks, you can provide them access to particular areas of BIG-IQ by adding the user and assigning the appropriate standardized role. You can assign as many roles as required to cover the user's responsibilities.

Important: *Since some roles have access only to certain areas or screens in the BIG-IQ user interface, it's important to communicate that to the user. When you assign a role to a user, be sure you outline the responsibilities and restrictions for their role. Clarifying this helps avoid any potential confusion.*

---

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.
3. Click the **Add** button.
4. From the **Auth Provider** list, select **local (Local)** to have BIG-IQ authenticate this user.
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 this new locally-authenticated user.

   You can change the password any time.
8. To associate this user with an existing user group, select the group from the **User Groups** list.

   You aren't required to associate a user group at this point; you can do that later if you want. If you want to associate another user group with this user, click +.
9. From the **User Roles** list, select a user role to associate with this user.

   Each role has a set of unique privileges. If you want to associate another user role with this user, click +.

   Important: *Be sure to let your users know that their access to certain parts of the BIG-IQ user interface depends on which role they are assigned.*

   ---
10. Click the **Save & Close** button at the bottom of the screen.

This user now has the privileges associated with the role you selected.

You can now provide this user with their credentials to access BIG-IQ. Be sure to let them know how their access aligns with their responsibilities and that they might not see every screen you or one of their peers does.

Remove a BIG-IQ user from a role

If a job or responsibilities change for an employee, you can use this procedure to disassociate that BIG-IQ user from an assigned role.

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.

3. On the Users inventory list, click the name of the user.
   The screen refreshes to display the properties for this user.
4. From the **User Roles** list, select the user role to disassociate from this user and click the **X**.
   The selected user role is removed from the list of privileges assigned to this user.
5. Click the **Save & Close** button at the bottom of the screen.

This user no longer has the privileges associated with the role you deleted.

Legal Notices

Legal notices

Publication Date

This document was published on April 14, 2017.

Publication Number

MAN-0577-06

Copyright

Trademarks

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

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

Patents

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

Export Regulation Notice

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

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index

M

N

P

R

S

U

V

W