F5® BIG-IQ® Centralized Management: ADC

Version 5.1

Table of Contents

Table of Contents

BIG-IQ Application Delivery Controller: Overview

What is Application Delivery Controller?

Application Delivery Controller (ADC) is the part of centralized management that you use to manage the configuration objects (such as servers, nodes, pools, and pool members) that move your traffic.

ADC helps the user:

- Create efficient work flows to view the LTM® 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.

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

If you are running BIG-IP version 11.5.1 up to version 11.6.0, you might need root user credentials to successfully discover and add the device to the BIG-IP devices inventory. Root user credentials are not required for BIG-IP devices running 11.5.0 - 11.5.1 and 11.6.0 - 12.x.

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

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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.

3. At the top of the screen, click **Inventory**.

4. Click the **Add Device** button.

5. In the **IP Address** field, type the IPv4 or IPv6 address of the device.

6. In the **User Name** and **Password** fields, type the user name and password for the device.

7. If this device is part of a DSC group, from the **Cluster Display Name** list, select one of the following:

   • For an existing DSC group, select **Use Existing** from the list and select the DSC group from the list.

   • For a new DSC group, select **Create New** from the list and type a name in the field.

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

8. If this device is configured in a DSC group, select an option:

   • **Initiate BIG-IP DSC sync when deploying configuration changes (Recommended)** Select this option if this device is part of a DSC group and you want this device to automatically synchronize configuration changes with other members 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 members in the DSC group.

9. 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.

10. 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**.

11. 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.

   You can also select these service configuration after you add the BIG-IP device to the inventory.

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

2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.

3. At the top of the screen, click **Inventory**.

4. Click the name of the device you want to import a service configuration from.

5. On the left, click **Services**.

6. 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.

7. 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.

8. 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
• Management IP Address
• Cluster
• Boot Location

From this screen you can also perform the following tasks:

• Log directly into the device from BIG-IQ®.
• Reboot the device from BIG-IQ.
• Create an instant backup of the device's configuration.
• Associate the device to a cluster.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. Click the name of the device you want to view.
   The device Properties screen opens.

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

With BIG-IQ®, you can easily search for specific sets of devices from one central location. 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. This centralized search saves time by displaying only those devices with the search criteria you specify.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.

4. 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.

5. To modify the filter to include or exclude certain objects, click the gear icon next to the **Filter** field and deselect or select objects.

6. 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.
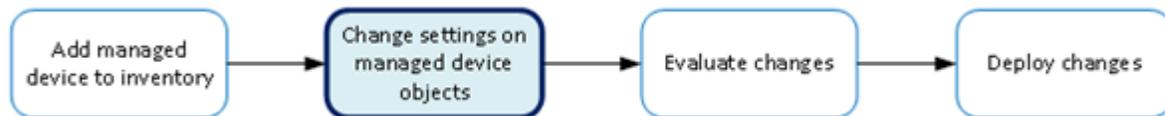

Figure 1: Change managed object workflow

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

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

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. 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.
5. Click the name of the object that you want to change.

   If you select **Virtual Servers**, there are a couple unique operations 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.

6. Make changes to the properties you want to modify.
7. 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 ADC interface to add a virtual server 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on this device.
5. Click **Create**.
   The Virtual Servers - New Item screen opens.
6. In the **Name** field, type in a name for the virtual server you are creating.
7. From the **Device** list, select the device on which to create the virtual server.
8. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the virtual server.
9. In the **Description** field, type in a brief description for the virtual server you are creating.
10. If you want the virtual server and its resources to be available for load balancing, select **Enable** for **State (on BIG-IQ)**.
11. 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.
12. For the **Destination Address**, type the IP address of the destination you want to add to the Destination list.

    The format for an IPv4 address is `I<a>.I<b>.I<c>.I<d>`. For example, `172.16.254.0/24`.

    The format for an IPv6 address is `I<a>:I<b>:I<c>:I<d>:I<e>:I<f>:I<g>:I<h>..`

    For example, `2001:db8:85a3:8d3:1319:8a2e:370:7348`.

    Note: Specifying a netmask is optional for this field.

13. 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.
14. 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.

15. 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.

16. For the **Type**, select the type of network service provided by this virtual server. The default is **Standard**.

Option	Description
Standard	Specifies a virtual server that directs client traffic to a load balancing pool and is the most basic type of virtual server. When you first create the virtual server, you assign an existing default pool to it. From then on, the virtual server automatically directs traffic to that default pool.
Forwarding (Layer 2)	Specifies a virtual server that shares the same IP address as a node in an associated VLAN.
Forwarding (IP)	Specifies a virtual server like other virtual servers, except that the virtual server has no pool members to load balance. The virtual server simply forwards the packet directly to the destination IP address specified in the client request.
Performance (HTTP)	Specifies a virtual server with which you associate a Fast HTTP profile. Together, the virtual server and profile increase the speed at which the virtual server processes HTTP requests.
Performance (Layer 4)	Specifies a virtual server with which you associate a Fast L4 profile. Together, the virtual server and profile increase the speed at which the virtual server processes Layer 4 requests.
Stateless	Specifies a virtual server that accepts traffic matching the virtual server address and load balances the packet to the pool members without attempting to match the packet to a pre-existing connection in the connection table. New connections are immediately removed from the connection table. This addresses the requirement for one-way UDP traffic that needs to be processed at very high throughput levels, for example, load balancing syslog traffic to a pool of syslog servers. Stateless virtual servers are not suitable for processing traffic that requires stateful tracking, such as TCP traffic. Stateless virtual servers do not support iRules, persistence, connection mirroring, rateshaping, or SNAT automap.
Reject	Specifies that the BIG-IP system rejects any traffic destined for the virtual server IP address.
DHCP Relay	Specifies a virtual server that relays Dynamic Host Control Protocol (DHCP) client requests for an IP address to one or more DHCP servers, and provides DHCP server responses with an available IP address for the client.
Internal	Specifies a virtual server that supports modification of HTTP requests and responses. Internal virtual servers enable usage of ICAP (Internet Content Adaptation Protocol) servers to modify HTTP requests and responses by creating and applying an ICAP profile and adding Request Adapt or Response Adapt profiles to the virtual server.

17. 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.

Option	Description
All Protocols	Specifies that the virtual server supports all network protocols.

Option	Description
TCP	Specifies that the virtual server supports the TCP protocol, defined in RFC 675.
UDP	Specifies that the virtual server supports the UDP protocol, defined in RFC 768.
SCTP	Specifies that the virtual server supports the Stream Control Transmission Protocol (SCTP) protocol, defined in RFC 4960.
Other	Provides the ability to specify another protocol. This setting is not available when you select **Standard** as the type.

18. 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.

Option	Description
All VLANs and Tunnels	Specifies that the virtual server is enabled on all VLANs and tunnels configured on the system.
Enabled on	Specifies that the virtual server is enabled on the VLANs and tunnels specified in the **Selected** list.
Disabled on	Specifies that the virtual server is disabled on the VLANs and tunnels specified in the **Selected** list.

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

- **None**: The system does not use a source address translation pool for this virtual server.
- **SNAT**: The system uses secure network address translation (NAT), as defined in the specified SNAT pool, for address translation.
- **Auto Map**: The system uses all of the self IP addresses as the translation addresses for the pool.

20. In the **Connection Limit** field, type the maximum number of concurrent connections allowed for the virtual server.

21. 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.

22. From the **Connection Rate Limit Mode** list, select the scope of the rate limit defined for the virtual server.

Option	Description
Per Virtual Server	Applies rate limiting to this virtual server.
Per Virtual Server and Source Address	Applies Connection Rate Limit Source Mask to the source IP address of incoming connections to this virtual server, and applies the rate limit to connections sharing the same subnet. The Connection Rate Limit Source Mask specifies the number of bits in the IP address to use as a limit key.
Per Virtual Server and Destination Address	Applies Connection Rate Limit Destination Mask to the destination IP address of outgoing connections from this virtual server, and applies the rate limit to connections sharing the same subnet. The Connection Rate Limit Destination Mask specifies the number of bits in the IP address to use as a limit key.
Per Virtual Server, Destination, and Source Address	Applies Connection Rate Limit Source Mask and Connection Rate Limit Destination Mask to the source and destination IP address of incoming connections to this virtual server, and applies the rate limit to connections

Option	Description
	sharing the same subnet. The Connection Rate Limit Source Mask and Connection Rate Limit Destination Mask specify the number of bits in the IP addresses to use as a limit key.
Per Source Address (All Rate Limiting Virtual Servers)	Applies rate limiting based on the specified source address for all virtual servers that have rate limits specified. **Per Destination Address (All Rate Limiting Virtual Servers)**: Applies rate limiting based on the specified destination address for all virtual servers that have rate limits specified.
Per Source and Destination Address (All Rate Limiting Virtual Servers)	Applies rate limiting based on the specified source and destination addresses for all virtual servers that have rate limits specified.

23. 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.

24. 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.

25. From the **Source Port** list, select how you want the system to preserve the connection's source port.

Option	Description
Preserve	Specifies that the system preserves the value configured for the source port, unless the source port from a particular SNAT is already in use, in which case the system uses a different port.
Preserve Strict	Specifies that the system preserves the value configured for the source port. If the port is in use, the system does not process the connection. Use this setting only for cases that meet at least one of the following conditions: • The port is configured for UDP traffic. • The system is configured for nPath routing or is running in transparent mode (that is, there is no translation of any other Layer 3 or Layer 4 field). • There is a one-to-one relationship between virtual IP addresses and node addresses, or clustered multi-processing (CMP) is disabled.
Change	Specifies that the system changes the source port. This setting is useful for obfuscating internal network addresses.

26. 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.

27. 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.

28. 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.

29. From the **Last Hop Pool** list, select the pool the system uses to direct reply traffic to the last hop router.

30. If you want the system to allow IPv6 hosts to communicate with IPv4 servers, select **NAT64**.

31. 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.

32. 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.

33. 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.

34. 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.

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

36. 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`

---

37. Click **Save**.
The system creates the new virtual server with the settings you specified.

Cloning a virtual server

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **LOCAL TRAFFIC**.

4. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
The screen displays the list of virtual servers defined on this device.

5. Select the check box associated with the existing virtual server that you want to clone.

6. From the **Actions** button, select **Clone**.
The BIG-IQ creates a new virtual server using the settings of the one you selected.

7. Modify the parameters for the new virtual server as needed.

---

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

---

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

Attaching iRules to virtual servers

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **LOCAL TRAFFIC**.

4. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
The screen displays the list of virtual servers defined on this device.

5. Select the check boxes associated with the virtual servers to which you want to attach iRules.

6. From the **Actions** button, select **Attach iRules**.
   The Virtual Servers - Attach iRules screen opens.

7. To specify which iRules to attach to the selected virtual servers, click on them in the **Available iRules** list, and click the right arrow to add them to the **iRules to be Attached** list.

8. Specify the order in which you want the iRules to attach using the up and down arrows.

9. 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**.

10. 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**.

11. Click **Save** and then confirm your choice by clicking **Modify**.
    A Modify Items box pops up to show the status of your request.

12. Click **Close** to dismiss the box and complete the process.

Creating a new iRule

You can use the BIG-IQ® ADC 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 ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **LOCAL TRAFFIC**.

4. Under **LOCAL TRAFFIC**, select **iRules**.
   The screen displays a list of iRules® that are known on this device.

5. Click **Create**.
   The iRules - New Item screen opens.

6. For **Name**, type a name for the iRule you are creating.

7. For **Partition**, type the name of the BIG-IP device partition on which you want to create the iRule.

8. For the **Body**, compose the script sequence that defines the iRule.

   For guidance on creating an iRule, consult the AskF5™ (http://www.askf5.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.

9. Click **Save**.
   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.

Creating a new node

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Nodes**.
   The screen displays a list of nodes that are defined on this device.
5. Click **Create**.
   The Nodes - New Item screen opens.
6. In the **Name** field, type in a name for the node you are creating.
7. From the **Device** list, select the device on which to create the node.
8. For the **Address** field, type in the IP address that identifies the new node.
9. For **Partition**, type the name of the BIG-IP device partition on which you want to create the node.
10. In the **Description** field, type in a brief description for the node you are creating.
11. To specify configuration parameters for this node, expand **Configuration** and continue with the next three steps. Otherwise, skip to the last step in this procedure.
12. 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 the monitors from the **Available** list and move the monitor to the **Enabled** list.
13. 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.
14. 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.
15. For the **Connection Limit**, type the maximum number of concurrent connections allowed for this node.
16. 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.
17. Click **Save**.

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

Creating a new pool

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of pools that are defined on this device.
5. Click **Create**.
   The Pools - New Item screen opens.
6. In the **Name** field, type in a name for the pool you are creating.
7. From the **Device** list, select the device on which to create the pool.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the pool.
9. In the **Description** field, type in a brief description for the pool you are creating.
10. To enable a health monitor for this pool, use the **Health Monitors** setting to select the monitor from the **Available** list and move the monitor to the **Enabled** list.
11. 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.

12. In the **Load Balancing Method** field, specify the type of load balancing you want the pool to use. The default is **Round Robin**.
13. 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.

14. To specify advanced properties, expand **Advanced Properties** and continue with the next twelve steps. Otherwise, skip to the last step in this procedure.
15. To automatically enable or disable NATs for connections that use this pool, for the **NAT** setting, select **Allow**.
16. To automatically enable or disable SNATs for connections that use this pool, for the **SNAT** setting, select **Allow**.
17. 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.

18. 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.

19. 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.

20. 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.

21. 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.

22. 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.
Specify	Provides a field in which you can specify a QoS level to apply. Valid values are from 0 to 7.

23. 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.

24. To enable TCP request queueing, select **Request Queueing**.

25. 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.

26. 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.

27. Click **Save**.

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

Creating a new pool member

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

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of pools that are defined on this device.
5. Click the name of the pool to which you are going to add a new member.
   The properties screen for that pool opens.
6. Near the bottom of the screen, click the **New Member** button.
   The Pools - New Item screen opens.
7. 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.
8. For the **Port**, type the service port for the pool member.
9. In the **Description** field, type in a brief description for the pool member you are creating.
10. 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 the monitors from the **Available** list and use the arrow button to move the monitor to the **Enabled** list.
    3. 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.

11. 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. LTM 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.

12. 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.

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

14. 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.

15. Click **Save**.
    The system creates the new pool member with the settings you specified.

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

    Important: You must log in as an Administrator to perform this task.

2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Users**.
5. Click the **Add** button.
6. From the **Auth Provider** list, select the provider that supplies the required credentials for authenticating this user. If you configured BIG-IQ to authenticate using LDAP or RADIUS, you have the option to authenticate this user through one of those methods.
7. In the **User Name** field, type the user name for this new user.
8. 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.
9. In the **Password** and **Confirm Password** fields, type the password for the new user.

10. Click **Save**.
    The system creates a new user.

11. On the left, click **USER MANAGEMENT** > **Roles**.

12. Click the **Add** button.

13. In the **Name** field, type a name to identify this role.

14. 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.

15. From the **Active Users and Groups** list, select the name of the user you specified in step7.

16. Click **Save**.
    The new role is created.

17. To delegate permissions for a virtual server, complete these next 7 steps.

    a) At the top left of the screen, select **ADC** from the BIG-IQ menu.
    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.

18. To delegate permissions for all of the pool members in a pool, do these next 7 steps.

    a) At the top left of the screen, select **ADC** from the BIG-IQ menu.
    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**.
       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.

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 client SSL 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 make.

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

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

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. In the **Description** field, type in a brief description for the profile you are creating.
8. For the **Type**, select the type of profile you want to create.
   The **Parent Profile** field displays.
9. 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.

10. 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.

11. If you are adding a Client SSL profile and you need to provide the pass phrase for the **Certificate Key Chain**, specify the pass phrase in the **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.

12. Click **Save**.
    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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

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

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **LOCAL TRAFFIC**.
4. Under **LOCAL TRAFFIC**, select **Profiles**.
   The screen displays the list of profiles defined on this device.
5. Click the name of the profile you want to edit.
   The screen displays the current settings for the selected profile.
6. In the **Description** field, you can revise the brief description for the selected profile.
7. 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.
8. 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.*

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.

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 it.

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

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

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **LOCAL TRAFFIC**.

4. Under **LOCAL TRAFFIC**, select **Monitors**.
   The screen displays the list of monitors defined on this device.

5. Click **Create**.
   The Monitors - New Item screen opens.

6. In the **Name** field, type in a name for the monitor you are creating.

7. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the monitor.

8. In the **Description** field, type in a brief description for the monitor you are creating.

9. For the **Type**, select the type of monitor you want to create.
   The **Monitor Template** setting displays.

10. 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**.

11. For **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.

12. 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.

13. 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.

14. 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. When this value is set to 0 (zero), the system uses the interval from the parent monitor. Note that **Timeout** and **Time Until Up** combine to control when a resource is set to up.

15. 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**.

16. 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
```

17. 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.*

18. 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**.*

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

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

20. If the monitored target requires authentication, for the **User Name**, specify the user name.

21. If the monitored target requires authentication, for the **Password**, specify the password.

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

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

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

   The default is **None**.

24. 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**.

25. 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`.

26. 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.`

27. 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.

28. 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.

29. 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).

30. 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.

31. 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.

32. 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.

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

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

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **LOCAL TRAFFIC**.

4. Under **LOCAL TRAFFIC**, select **Monitors**.
The screen displays the list of monitors defined on this device.

5. Select the monitor you want to edit.
The Monitors -screen opens to display the current settings for the selected monitor.

6. In the **Description** field, add or revise a brief description for the monitor you are editing.

7. 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.

8. 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.

Option	Description
Enabled	Enables specification of a different interval to use when checking the health of a resource that is up.

9. 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.

10. 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. When this value is set to 0 (zero), the system uses the interval from the parent monitor. Note that **Timeout** and **Time Until Up** combine to control when a resource is set to up.

11. 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**.

12. 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
```

13. 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.*

14. 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**.*

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

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

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

17. 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.

---

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

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

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

    The default is **None**.

20. 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**.

21. 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**.

22. 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.

23. 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.

24. 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.

25. 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).

26. 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.

27. 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.

28. 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.

29. 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.

Changing network objects

You can make revisions to the configuration of local traffic objects to simplify managing your devices.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. 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®.
5. Click the name of the object you want to change.
   The Properties screen for the selected object opens.
6. Make changes to the properties you want to modify.
7. 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.

Managing a network interface

You can use the BIG-IQ® ADC to enable or disable network interfaces on 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. Under **NETWORK**, select **Interfaces**.
   The screen displays a list of network interfaces defined on devices that are managed by this BIG-IQ.
5. Select the interface you want to change and then select or clear **Enable**.
   The **State** for the selected interface changes on the BIG-IQ.
6. Click **Save**.
   The system creates the new route with the settings you specified.

Creating a new route

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. Under **NETWORK**, select **Routes**.
   The screen displays a list of routes defined on devices that are managed by this BIG-IQ.
5. Click **Create**.
   The Routes - New Item screen opens.
6. In the **Name** field, type in a name for the route you are creating.
7. In the **Description** field, type in a brief description for the route you are creating.
8. From the **Device** list, select the device on which to create the route.
9. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route.
10. 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
```

    .

11. 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**.

12. 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.

13. Click **Save**.
    The system creates the new route with the settings you specified.

Creating a new self IP address

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. 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.
5. Click **Create**.
   The Self IPs - New Item screen opens.
6. In the **Name** field, type in a name for the self IP address you are creating.
7. From the **Device** list, select the device on which to create the self IP address.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the self IP.
9. 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.
10. 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).
11. For the **VLAN/Tunnel**, select the VLAN or tunnel to associate with this self IP address.
12. 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 option, where you can specify the ports, protocols, and services to activate on this self IP address.
13. For the **Traffic Group**, select a specific traffic group for the self IP address.
14. Click **Save**.
    The system creates the new self IP address with the settings you specified.

Creating a new route domain

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. Under **NETWORK**, select **Route Domains**.

The screen displays a list of route domains defined on the devices that are managed by this BIG-IQ.

5. Click **Create**.
   The Route Domains - New Item screen opens.

6. In the **Name** field, type in a unique name for the route you are creating.

7. 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.

8. In the **Description** field, type in a brief description for the route domain you are creating.

9. From the **Device** list, select the device on which to create the route domain.

10. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route domain.

11. 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).

12. To specify a VLAN or tunnel for the BIG-IP to use in the route domain, click it in the **Available** list, and click the right arrow to add it to the **Enabled** list.

13. Click **Save**.
    The system creates the new route domain with the settings you specified.

Creating a new VLAN

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.

2. At the top left of the screen, select **ADC** from the BIG-IQ menu.

3. On the left, expand **NETWORK**.

4. Under **NETWORK**, select **VLANs**.
   The screen displays a list of VLANs defined on the devices that are managed by this BIG-IQ.

5. Click **Create**.
   The VLANs - New Item screen opens.

6. In the **Name** field, type a unique name for the VLAN you are creating.

7. In the **Description** field, type a brief description for the VLAN you are creating.

8. 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.

9. From the **Device** list, select the device on which to create the VLAN.

10. For **Partition**, type the name of the BIG-IP device partition on which you want to create the VLAN.

11. In the **MTU** field, specify the maximum transmission unit (MTU) for traffic on this VLAN.

    The default is 1500.

12. To specify which interfaces this VLAN uses for traffic management, select it in the **Interface** list, and then select the **Tagging** for it.

13. Click **Save**.
    The system creates the new VLAN with the settings you specified.

Creating a new DNS resolver

You can use the BIG-IQ® ADC 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **ADC** from the BIG-IQ menu.
3. On the left, expand **NETWORK**.
4. Under **NETWORK**, select **DNS Resolvers**.
   The screen displays a list of DNS resolvers defined on the devices that are managed by this BIG-IQ.
5. Click **Create**.
   The DNS resolvers - New Item screen opens.
6. In the **Name** field, type in a unique name for the DNS resolver you are creating.
7. For **Partition**, type the name of the BIG-IP device partition on which you want to create the DNS resolver.
8. To specify which devices use this DNS resolver for traffic management, in the **Devices** setting, click them in the **Available** list, and use the right arrow to add them to the **Selected** list.
9. Select the **Route Domain Name** that this resolver uses for outbound traffic.

   The default is the default route domain.
10. 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.
11. To specify the Traffic properties, expand the control and then:
    a) If you want the system to answer and issue IPv4-formatted queries, select **Use IPv4**.
    b) If you want the system to answer and issue IPv6-formatted queries, select **Use IPv6**.
    c) If you want the system to answer and issue UDP-formatted queries, select **Use UDP**.
    d) If you want the system to answer and issue TCP-formatted queries, select **Use TCP**.
12. To specify a forward zone used to resolve matching DNS queries, expand the control and then:
    a) Click **Add** to specify a new zone.
       A popup screen opens.
    b) In the **Name** field, type in a unique name for the forward zone you are creating.
    c) In the **Address** field, type in an IP address for the forward zone you are creating.
    d) In the **Service Port** field, type in the port number for the forward zone you are creating.
    e) Click **Add**.
13. Click **Save**.
    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.

Deploying Changes

How do I evaluate changes made to managed objects?

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

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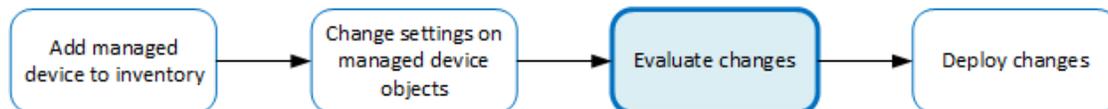

Figure 2: Overview of evaluating changes made to managed objects

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

Evaluating configuration changes

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

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

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

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

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

2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.
3. On the left, expand **EVALUATE & DEPLOY**.
4. Under **EVALUATE & DEPLOY**, select **Local Traffic & Network**.
   The screen displays a list of LTM® evaluations and deployments that have been created on this device.
5. Under Evaluations, click **Create**.
   The Create Evaluation screen opens.
6. In the **Name** field, type in a name for the evaluation task you are creating.
7. In the **Description** field, type in a brief description for the evaluation task you are creating.
8. For the **Source**, select what you want to evaluate.

   • To compare the object settings currently on the managed device with the object settings in the pending version, select **Current Changes**.

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

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

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

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

10. Click the **Create** button at the bottom of the screen.
    The system adds the new evaluation to the list, and analyzes the changes for errors. When the configuration evaluation completes, you will see how many changes or errors the evaluation found.

11. Review the evaluation to determine whether you are going to deploy it or not.
    a) If there are critical errors, you cannot deploy these changes. Click each error to see what it is, and then go back to where you made the change to fix it.

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

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

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

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

How do I deploy changes made to managed objects?

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

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

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

This figure illustrates the workflow you perform to manage the objects on BIG-IP devices. Deploying the settings is the last step in this process.
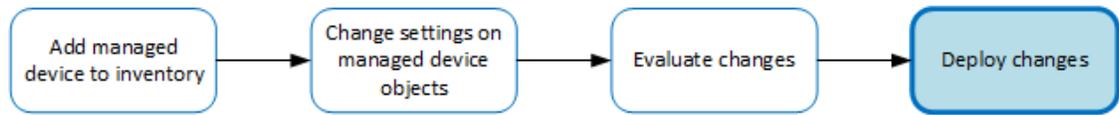
Figure 3: Change managed object workflow

Deploying configuration changes

To apply the changes you made on the BIG-IQ to your managed device, you must deploy those changes to the managed device.

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

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

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

   • To deploy this change immediately:

     1. Select **Deploy Now**.
     2. Click **Deploy** to confirm.
   • To deploy this change later:

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

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

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

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

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

Making an urgent deployment

If you need to make urgent changes, you can skip the evaluation task and deploy changes right now.

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

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

2. At the top left of the screen, select **Change Management** from the BIG-IQ menu.

3. On the left, expand **EVALUATE & DEPLOY**.

4. Under **EVALUATE & DEPLOY**, select **Local Traffic & Network**.
   The screen displays a list of LTM® evaluations and deployments that are defined on this device.

5. Under Deployments, click **Create**.
   The Create Deployment screen opens.

6. In the **Name** field, type in a name for the deployment task you are creating.

7. In the **Description** field, type in a brief description for the deployment task you are creating.

8. For the **Source** setting, select what you want to deploy.

   • To deploy your changes to the managed device, select **Current Changes**.
   • To deploy the object settings from a stored snapshot, select **Existing Snapshot**, then choose the snapshot you want to use.

9. Using the **Target** settings, identify the devices for which you want to deploy changes.

   a) If the devices are in a device group, select **Group**, and select the group.
   b) If the devices are not in a device group, select **Device**.
   c) Select the devices from the **Available** list and use the arrow button to move the devices to the **Enabled** list.

10. Consider one more time how you want to deploy these changes.

    • To make the changes right now, click **Deploy immediately**.
    • If you want to review the changes, click **Create evaluation**.

11. Click **Create**.

    • If you selected **Deploy immediately**, the changes begin to deploy and the Status column updates as it proceeds.
    • If you selected **Create evaluation**, the new evaluation is added to the list and the changes are analyzed for errors. When the evaluation completes, you will see how many changes or errors the evaluation found.

Deploying to one device when a cluster member is down

Deploying changes to a device in a cluster that has a device offline will generally fail. Normally, all device members must be available before you deploy changes to a cluster member. However, if you need to deploy changes before all cluster members are available you can do so.

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

   Important: You must log in as an Administrator to modify cluster properties (step 5). Administrators, ADC Deploy, or an ADC Manager can deploy changes to individual cluster members.

2. At the top left of the screen, select **Device Management** from the BIG-IQ menu.

3. On the left, click **BIG-IP DEVICES**.
   The screen displays a list of managed devices for this BIG-IQ Centralized Management system.

4. Under Device Name, click the cluster member to which you want to deploy changes.
   The properties screen for this member opens.

5. Under Cluster Properties, click **Edit**.
   The Cluster Properties popup screen for this cluster opens.

6. For Deployment Settings, select **Ignore BIG-IP DSC sync when deploying configuration changes**.

7. Click **OK**, and then click **Close**.

With the "ignore sync" option selected, you can now deploy changes to the member that is available, and BIG-IQ will not attempt to sync those changes to the member that is unavailable.

Use the *Deploying configuration changes* task to deploy changes to the available member. When you select the target device for deployment, do not select the unavailable device.

Managing Configuration Snapshots

Creating a snapshot

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

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

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

Comparing snapshots

You can compare two snapshots to view their differences.

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

Restoring a snapshot

You can restore a snapshot to change the working configuration to that of the snapshot. Restoring the snapshot merges objects from the snapshot into the BIG-IQ® configuration and removes all active locks. No objects in the BIG-IQ configuration are removed. Once the restore process starts, you cannot modify the BIG-IQ configuration until the process is completed or canceled. If the process is canceled, all configuration settings are rolled back.

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

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

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

Users, User Groups, Roles, and Authentication

How do I manage access to BIG-IQ and my managed BIG-IP devices?

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.

Can I 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 a 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.

Before you can set up authentication, you must have specified your DNS settings. You usually do this when you license the F5 BIG-IQ Centralized Management system. You can set up F5 BIG-IQ Centralized Management 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. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Auth Providers**.
5. Click the **Add** button.
6. From the **Provider Type** list, select **RADIUS.**
7. 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.
8. 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.
9. In the **Secret** field, type the case-sensitive text string used to validate communication.

10. 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

11. Click the **Save** 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 an LDAP 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 F5 BIG-IQ Centralized Management to user your company's LDAP server to authenticate users. You can specify multiple LDAP servers for user authentication.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Auth Providers**.
   The Auth Providers screen opens.
5. Click the **Add** button.
6. From the **Provider Type** list, select **LDAP**.
7. 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.
8. In the **Host** field, type the IP address of your LDAP server.
9. If your Active Directory server uses a port other than the default, 389, in the **Port** field, type the number of the alternative port.

10. If you want BIG-IQ to use an SSL port to communicate with the LDAP server, for the **SSL Enabled** setting, select the **Enabled** check box.

    Note that the **Port** setting automatically changes to **636**.

11. 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.

12. In the **Root DN** field, type the root context that contains users and groups.

    The root context must be a full distinguished name.

13. 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.

14. 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.

15. 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.

16. In the **Connect Timeout** field, type the number of milliseconds after which the BIG-IP system stops trying to connect to the LDAP server.

17. In the **Read Timeout** field, type the number of seconds the BIG-IP system will wait for a response to a query.

18. 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`.

19. 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.

20. 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`.

21. 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}*))`

22. 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.

23. 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)`

24. Select the **Perform Test** check box to test this provider.

25. Click the **Save** button at the bottom of the screen.

BIG-IQ Centralized Management now authenticates users against the configured LDAP server.

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 locally-authenticated user

Create a user to provide access to F5® BIG-IQ® Centralized Management.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Users**.
5. Click the **Add** button.
6. From the **Auth Provider** list, select **Local**.
7. In the **User Name** field, type the user name for this new user.
8. 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.
9. In the **Password** and **Confirm Password** fields, type the password for the new user.
10. 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.

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

Each role has a set of unique privileges.

12. Click the **Save** button at the bottom of the screen.

How do I limit privileges for users?

As a system manager, you need a way to limit user privileges based on their responsibilities. To help you do that, F5® BIG-IQ® Centralized Management ships with a set of default roles that you can assign to users. Roles are shared between BIG-IQ systems in a high availability pair, so they remain assigned to users even if the primary BIG-IQ system fails over.

Standard roles shipped with BIG-IQ

F5® BIG-IQ® Centralized Management ships with several standard roles, which you can assign to individual users, or to a user group. Roles are shared between BIG-IQ systems in a high availability pair, so they remain assigned to users even if the BIG-IQ system fails over.

Role	Role Description / Access
Administrator	This role has access to all licensing aspects of System Management and Device Management. This includes access for adding individual users, assigning roles, discovering BIG-IP® systems, installing updates, activating licenses, and setting up BIG-IQ® in a high availability (HA) configuration.
ADC Deployer	This role has access to deploy and view ADC configuration objects for managed ADC devices.
ADC Editor	This role has access to edit all ADC configuration objects.
ADC Manager	This role has access to all aspects of ADC, including areas involved in creating, viewing, modifying, and deleting Local Traffic and Network objects.
ADC Viewer	This role has view-only access for all ADC objects and features.
Access Auditor	This role has access to all Access reports and dashboard.
Access Deployer	This role has deploy access to Access configuration objects. This role cannot discover and edit devices or policies.
Access Editor	This role has edit access to Access configuration objects. This role cannot discover and deploy devices or policies. This role includes the ability to add, update, and delete pools and pool members from the Access configuration object editor.
Access Manager	This role has deploy and edit access to Access configuration objects, and has access to Access Reports and Dashboard. This role cannot add or remove devices and device groups, and cannot discover, import, or delete services.
Access Viewer	This role has view-only access to Access configuration objects and tasks for Access devices that have been discovered. This role cannot edit, discover, or deploy devices or policies.
Device Manager	This role has access to all aspects of Device Management, including areas involved in device discovery, group creation, licensing, software

Role	Role Description / Access
	image management, UCS backups, templates, connectors, certificates, self IP addresses, VLANs, and interfaces.
Device Viewer	This role has read-only access to all aspects of Device Management, including areas involved in device discovery, group creation, licensing, software image management, UCS backups, templates, connectors, certificates, self IP addresses, VLANs, and interfaces.
Fraud Protection Manager	This role has access to all aspects of the Fraud Protection Service functionality for Web Client Security.
Fraud Protection View	This role has view-only access to all Fraud Protection Service objects for Web Client Security .
Network Security Deploy	This role has access to view and deploy Network Security objects.
Network Security Manager	This role has access to all aspects of Network Security, including areas involved in creating, viewing, modifying, and deleting shared and firewall-specific security objects.
Network Security Edit	This role has access to create, view, and modify objects for Network Security.
Network Security View	This role has view-only access to firewall objects for Network Security. This role cannot edit, discover, or deploy devices or policies.
Security Manager	This role has access to all aspects of Network Security, Web Application Security, and Web Client Security, including areas involved in device discovery, creating, viewing, modifying, and deleting Web Application Security, shared and firewall-specific security objects.
Trust Discovery Import	This role manages device trust establishment, service discovery, service import, removal of services and removal of trust.
Web App Security Deployer	This role can deploy and view ASM configuration objects for managed ASM devices.
Web App Security Editor	This role manages config objects within the ASM module.
Web App Security Manager	This role has access to 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	This role permits read-only access to the ASM module.

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 left of the screen, select **Device Management** from the BIG-IQ menu.
2. At the top of the screen, click **Inventory**.
3. On the left, click **USER MANAGEMENT** > **Roles**.
4. Click the **Add** button.
5. In the **Name** field, type a name to identify this new role.
6. From the **Role Type** list, select the kind of role you want to add.

   This role has no permissions when you first create it. You have to add permissions after you save the role.

7. Click the **+** sign if you want this role to have access to another user or group, and select the device group from the list.
8. From the **Active Users and Groups** list, select the user or group you want to associate with this new role.
9. Click the **Save** button at the bottom of the screen.

Changing the default password for the administrator user

When you license and do the initial setup, F5® BIG-IQ® Centralized Management system prompts you to automatically create the administrator user.

For security reasons, it is important to change the administrator role password from the default, admin.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Users**.
5. In the User Name column, click **admin**.
   The Admin User properties screen opens.
6. In the **Old Password** field, type the password.
7. In the **Password** and **Confirm Password** fields, type a new password.
8. Click the **Save** button at the bottom of the screen.

Associating a user or user group with a role

Before you can associate a user or user group with a role, you must create a user or user group.

When you associate a user or user group with a role, you define the resources users can view and modify. You can associate multiple roles with a given user.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. From the **User Roles** list, select a user role to associate with this user.
   Each role has a set of unique privileges.
5. From the **Active Users and Groups** list, select the users or user groups to add to this role.
6. Click the **Save** button at the bottom of the screen.

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

Disassociating a user from a role

Use this procedure to disassociate a user from an assigned role.

1. Log in to F5® BIG-IQ® Centralized Management with your user name and password.
2. At the top left of the screen, select **System Management** from the BIG-IQ menu.
3. At the top of the screen, click **Inventory**.
4. On the left, click **USER MANAGEMENT** > **Users**.
5. On the Users inventory list, click the name of the user.
   The screen refreshes to display the properties for this user.
6. 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.

7. Click the **Save** button to save your changes.

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

Legal Notices

Legal notices

Publication Date

This document was published on September 14, 2016.

Publication Number

MAN-0577-04

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

A

Access Administrator role
defined *45*
Access Auditor role
defined *45*
ADC
about *5*
ADC Deploy role
defined *45*
ADC Editor role
defined *45*
ADC Manager role
defined *45*
ADC Viewer role
defined *45*
admin*, See* administrator
Administrator role
defined *45*
administrator user
changing password for *47*
administrator user password
changing *47*
Application Delivery Controller*, See* ADC
authenticate user
using an LDAP server *41*
authentication
before configuring LDAP
user authentication
before configuring through LDAP *42*
configuring BIG-IQ to use LDAP *42*
configuring with RADIUS *41*
authorization
for user access *41*

B

BIG-IP devices
rebooting *9*
BIG-IQ Device
about ADC *5*
BIG-IQ inventory
adding devices to *7*
BIG-IQ users
authenticating with LDAP *41*

C

centralized management
of BIG-IP devices *7*, *10*
changes
about evaluating before deploying *35*
evaluating before deploying *35*
compare snapshots *39*
configuration changes
about deploying *36*
deploying to a device *37*
deploying urgent *38*

configuration changes *(continued)*
evaluating *35*
making urgent *37*
configuration deployment
about *36*
configurations
filtering for devices *9*
importing for services *8*

D

default users *41*
delegate
assigning enable disable permissions *21*
deployment
of configuration changes *36*, *37*
device configurations
filtering *9*
device inventory
about *7*, *10*
device management
about *7*, *10*
searching for BIG-IP components *9*
Device Manager role
defined *45*
device managers
how to manage *24*
device profile
creating for LTM *22*
editing for LTM *24*
device profiles
how to manage *22*
devices
about discovering *7*, *10*
adding to BIG-IQ inventory *7*
discovering *7*
viewing details *9*
discovery
defined *7*, *10*
DNS resolver
creating *34*

E

emergency deployment
of configuration changes *37*, *38*
enable disable permission
delegating *21*
evaluation of changes
before deploying *35*

F

Fraud Protection Manager role
defined *45*
Fraud Protection View role
defined *45*

Index

V

virtual servers
    attaching iRules *15*
    cloning *15*
    creating *11*
    delegating permissions *21*
VLAN
    creating *33*

Index