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

Version 5.3

Table of Contents

Managing Local Traffic Profiles....5
How do I manage LTM profiles in BIG-IQ?....5
Create an LTM profile....5
Edit an LTM profile....6

Managing Virtual Servers....9
How do I change object settings on a managed device?....9
What virtual server management tasks can I perform?....9
Create a new virtual server....9
Clone a virtual server....12
Attach iRules to virtual servers....12
Change virtual server settings....13

Managing iRules....15
How do I change object settings on a managed device?....15
Create a new iRule....15
Attach iRules to virtual servers....16

Managing Pool & Pool Members....17
How do I change object settings on a managed device?....17
What pool and pool member management tasks can I perform?....17
Create a new pool....17
Create a new pool member....20
Delegate enable and disable permissions....21
Create a new node....22
Change pool or pool member settings....23
Create a new SNAT pool....24

Managing Local Traffic Monitors....27
How do I change object settings on a managed device?....27
What LTM monitor management tasks can I perform?....27
Create an LTM monitor....27
Edit an LTM monitor....30

Managing Network Objects....33
How do I change object settings on a managed device?....33
Change a network object....33
Manage a network interface....33
Create a new route....34
Create a new route domain....34
Create a new VLAN....35
Create a new self IP address....36
Create a new DNS resolver....37

Configure the BIG-IQ to manage an IPsec tunnel....39
How do I start managing an IPsec tunnel?....39

Create a forwarding virtual server for IPsec....39
Create an IKE peer....40
Create a custom IPsec policy....40
Create a bidirectional IPsec traffic selector....41
Configure the IKE daemon....41
Verify IPsec connectivity....41

Configure IPsec event viewing on the BIG-IQ....43
How do I configure viewing IPsec event logs?....43
Create a log publisher pool....43
Create a remote high-speed log destination for IPsec....44
Create a remote Syslog destination for IPsec....44
Configure a log publisher to send IPsec events to the BIG-IQ....45
Import IPsec configuration settings from the BIG-IP device....45
Enable IPsec event collection....46

Troubleshooting an IPsecTunnel....47
Troubleshoot an unhealthy IPsec tunnel using performance statistics....47
Troubleshoot an unhealthy IPsec tunnel using event logs....48

Legal Notices....49
Legal notices....49

Managing Local Traffic Profiles

How do I manage LTM profiles in BIG-IQ?

You can create or modify custom LTM® profiles in BIG-IQ® Centralized Management and then attach them to a virtual server to deploy them to your managed devices.

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

Create an LTM profile

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

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

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

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

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

Note: *For detailed information on the impact of using a particular profile parameter value, refer to the BIG-IP Local Traffic Management: Profiles Reference on* `support.F5.com`.

10. If you are adding a profile that requires a security parameter, specify the passphrase in the corresponding **Passphrase** field.

Note: *For version 12.0.0 devices, you do not need to supply the pass phase for the profile. For devices earlier than version 12.0.0, if you plan to make changes to a Client SSL profile, you need to supply the pass phrase for that profile. If you do not change any of the parameters for the profile or associate the profile with a virtual server or another client SSL profile, then you can leave this field blank. So, if you add a pre-version 12.0.0 device that has a significant number of profile definitions, you do not need to add the pass phrase for every profile, just the ones that you plan to change or associate with an LTM object.*

11. Click **Save & Close**.
    The system creates the new profile you specified and adds it to the list of profiles.

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

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

Edit an LTM profile

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

1. At the top of the screen, click **Configuration**.
2. On the left, expand **LOCAL TRAFFIC**.
3. Under **LOCAL TRAFFIC**, select **Profiles**.
   The screen displays the list of profiles defined on this device.
4. Click the name of the profile you want to edit.
   The screen displays the current settings for the selected profile.
5. 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.

   Any changes you make to this profile are also inherited by all profiles listed here that name this profile as their parent profile.
6. Under the **Override All** check box, select the check box corresponding to any fields you want to override, 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.*

Note: *For detailed information on the impact of using a particular profile parameter value, refer to the BIG-IP Local Traffic Management: Profiles Reference on* `support.F5.com`.

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

Important: *For imported profiles that use passphrases:*

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

---

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

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

Managing Virtual Servers

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.

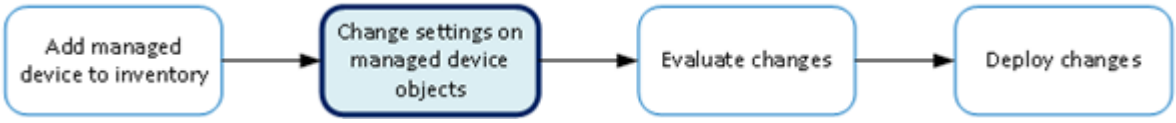

Figure 1: Change managed object workflow

What virtual server management tasks can I perform?

There are a number of ways you can use BIG-IQ® Centralized Management to manage the virtual servers on the managed BIG-IP® devices:

- Create a new virtual server.
- Modify an existing virtual server.
- Clone the settings of an existing virtual server to create a new one.
- Attach a sequence of iRules® to a virtual server.
- View statistics for a virtual server.
- Deploy the virtual server immediately to your managed device.

  Note: *You (or someone else) can also deploy your changes later. For more information about managing changes, look on* `support.F5.com` *in F5 BIG-IQ Centralized Management: Device for the topic: Deploying Changes.*

- Add or remove permissions for a virtual server and assign them to roles that have been defined on this BIG-IQ system. For more information about managing permissions, look on `support.F5.com` in *F5 BIG-IQ Centralized Management: Licensing and Initial Setup* for the topic: *Users, User Groups, Roles, and Authentication*.

Create a new virtual server

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

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.

3. Click **Create**.
   The New Virtual Server screen opens.
4. For **Name**, type in a name for the virtual server you are creating.
5. From **Device**, select the device on which to create the virtual server.
6. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the virtual server.
7. For **Description**, type in a brief description for the virtual server you are creating.
8. If you want the virtual server and its resources to be available for load balancing, for **State (on BIG-IQ)**, select **Enabled**.
9. For **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.
10. For **Destination Address**, type the IP address of the destination you want to add to the Destination list.
11. For **Service Port**, 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.
12. 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.
13. 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.
14. For **Type**, select the type of network service provided by this virtual server. The default is **Standard**.

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

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

    Note: *For details on the significance of choosing one option over another, refer to the BIG-IP documentation about virtual servers available on* `support.f5.com`.
17. From the **Source Address Translation** list, select the type of address translation pool used for implementing selective and intelligent source address translation.

    Note: *For details about the significance of choosing one option over another, refer to the BIG-IP documentation about virtual servers available on* `support.f5.com`.
18. For **Connection Limit** , type the maximum number of concurrent connections allowed for the virtual server.

19. For **Connection Rate Limit** , 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.

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

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

21. 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.

22. 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.

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

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

24. 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.
25. 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.
26. 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.
27. From **Last Hop Pool**, select the pool the system uses to direct reply traffic to the last hop router.
28. If you want the system to allow IPv6 hosts to communicate with IPv4 servers, select **NAT64**.
29. 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.

30. 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.
31. 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.

32. For **Default Pool**, 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.

33. For **Default Persistence Profile**, select the name of the default profile you want the virtual server to use to maintain session persistence.
34. For **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 about how fallback persistence profiles work, refer to SOL30483109: Overview of Fallback Persistence on* `AskF5.com`

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

Clone a virtual server

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on all of the devices managed by this BIG-IQ.
3. Select the check box associated with the existing virtual server that you want to clone.
4. Click the **Clone** button.
   The BIG-IQ creates a new virtual server using the settings of the one you selected and opens the Virtual Servers -Clone screen so you can modify parameters you need to change.
5. Modify the parameters for the new virtual server as needed.

   Important: *Two virtual servers cannot share the same* ***Destination Address****,* ***Protocol****, and* ***VLAN****.*
6. When you are satisfied with the settings for the new virtual server, click **Clone**.
   The system creates the new virtual server with the settings you specified.

Attach iRules to virtual servers

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on this device.
3. Select the check boxes associated with the virtual servers to which you want to attach iRules.
4. Click **Attach iRules**.
   The Bulk Attach iRules screen opens.
5. To specify which iRules to attach to the selected virtual servers, select them in the **Available iRules** list, and move them to the **iRules to be Attached** list.
6. Specify the order in which you want the iRules to attach using the up and down arrows.
7. For **Location**, 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**.
8. Use the **Duplicate Policy** setting to 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 previously, click **Reorder virtual servers' existing rules to preserve selected rules order**.
9. Click **Save & Close**.

Change virtual server settings

Using the BIG-IQ® user interface to make revisions to your virtual server configurations simplifies managing your devices.

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

---

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays a list of the virtual servers that are defined on this BIG-IQ.
3. Click the name of the virtual server that you want to change.

   If you select the check box for the virtual server instead of the name, there are a couple of unique operations that you can perform. 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 virtual server opens.
4. Make changes to the properties you want to modify.

   Note: *For detailed information on the impact of using a particular profile parameter value, refer to the BIG-IP Local Traffic Management: Profiles Reference on* `support.F5.com`.

   ---
5. When you are satisfied with the changes you have made, click **Save & Close**.
   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.

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

Managing iRules

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.

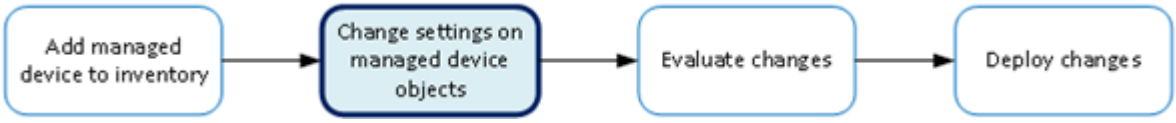

Figure 2: Change managed object workflow

Create a new iRule

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

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

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

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

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

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

Attach iRules to virtual servers

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
   The screen displays the list of virtual servers defined on this device.
3. Select the check boxes associated with the virtual servers to which you want to attach iRules.
4. Click **Attach iRules**.
   The Bulk Attach iRules screen opens.
5. To specify which iRules to attach to the selected virtual servers, select them in the **Available iRules** list, and move them to the **iRules to be Attached** list.
6. Specify the order in which you want the iRules to attach using the up and down arrows.
7. For **Location**, 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**.
8. Use the **Duplicate Policy** setting to 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 previously, click **Reorder virtual servers' existing rules to preserve selected rules order**.
9. Click **Save & Close**.

Managing Pool & Pool Members

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.

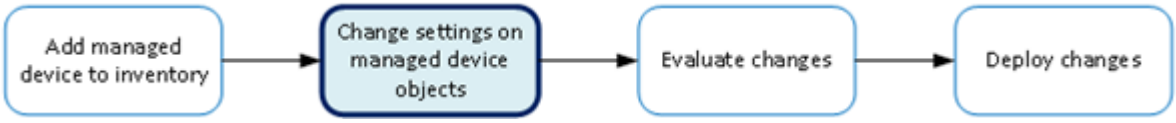

Figure 3: Change managed object workflow

What pool and pool member management tasks can I perform?

There are a number of ways you can use BIG-IQ® Centralized Management to manage the pools and pool members on your managed BIG-IP devices:

- Create a new pool or pool member.
- Modify an existing pool or pool member.
- View statistics for a pool or pool member.
- Deploy the pool and pool member immediately to your managed device; for pool members, you can enable, disable, or force offline immediately.

  Note: *You (or someone else) can also deploy your changes later. For more information about managing changes, look on* `support.F5.com` *in F5 BIG-IQ Centralized Management: Device for the topic: Deploying Changes.*

- Add or remove permissions for a pool or pool member and assign them to roles that have been defined on this BIG-IQ system. For more information about managing permissions, look on `support.F5.com` in *F5 BIG-IQ Centralized Management: Licensing and Initial Setup* for the topic: *Users, User Groups, Roles, and Authentication*.

Create a new pool

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

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of pools that are defined on all of the devices managed by this BIG-IQ.
3. Click **Create**.

The New Pool screen opens.

4. In the **Name** field, type in a name for the pool you are creating.
5. From the **Device** list, select the device on which to create the pool.
6. For **Partition**, type the name of the BIG-IP device partition on which you want to create the pool.
7. In the **Description** field, type in a brief description for the pool you are creating.
8. To enable specific health monitors for this pool, select a monitor from the **Health Monitors** list.

   To add additional monitors click the + icon and repeat this step.
9. If you enabled specific monitors for this pool, for the **Availability Requirement** field, specify the minimum number of monitors that must report a pool as being available before the member is defined as being in an up state.
   - If all of the monitors must report the pool available, select **All**.
   - To specify a minimum number, select **At Least**, and then type the minimum number in the **Health Monitors** field.
10. In the **Load Balancing Method** field, specify the type of load balancing you want the pool to use. .

    The default is **Round Robin**.
11. 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.
12. To add a new pool member for this pool, click **New Member**.

    For details on which values to specify for the fields on the New Pool Member screen refer to *Create a new pool member*.

    Note: *When you create a new pool member while creating a new pool, the new pool member is not actually created until you save the new pool. When you create a new pool member for an existing pool member, the new member is ready to use as soon as you save it.*

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

Option	Description
None	Specifies that the system takes no action to manage existing connections when a pool member becomes unavailable. The system maintains existing connections, but does not send new traffic to the member.
Reset	Specifies that, if there are no pool members available, the system resets and clears the active connections from the connection table and sends a reset (RST) or Internet Control Message Protocol (ICMP) message. If there are pool members available, the system resets and clears the active connections, but sends newly arriving connections to the available pool member and does not send RST or ICMP messages.

Option	Description
Drop	Specifies that the system simply cleans up the connection.
Reselect	Specifies that the system manages established client connections by moving them to an alternative pool member when monitors mark the original pool member down.

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

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

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

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

    Setting a ToS level affects the packet delivery reliability.

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

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

    Setting a ToS level affects the packet delivery reliability.

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

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

    Setting a QoS level determines the packet delivery priority.

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

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

Setting a QoS level affects the packet delivery priority.

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

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

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

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

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

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

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

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

Create a new pool member

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

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

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

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

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

10. 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.

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

12. 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.

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

Delegate enable and disable permissions

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

1. At the top of the screen, click **System**.
2. On the left, click **USER MANAGEMENT** > **Users**.
3. Click the **Add** button.
4. From the **Auth Provider** list, select **local (Local)** to have BIG-IQ authenticate this user.
5. In the **User Name** field, type the user name for this new user.
6. In the **Full Name** field, type a name to identify this user.

   The full name can contain a combination of symbols, letters, numbers and spaces.

7. In the **Password** and **Confirm Password** fields, type the password for this new locally-authenticated user.

   You can change the password any time.

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

    - To create a role to which you can delegate virtual server permissions to immediately disable or enable virtual servers to which this role is assigned, select **Virtual Server Operator**.

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

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

13. From the **Active Users and Groups** list, select the name of the user you specified in step6.
14. Click **Save**.
    The new role is created.
15. To delegate permissions for a virtual server, complete these sub-steps.
    a) At the top of the screen, click **Configuration**.
    b) On the left, expand **LOCAL TRAFFIC** and click **Virtual Servers**.
    c) Click the name of the virtual server for which you want to delegate permissions.
    The properties screen for the virtual server opens.
    d) Click **Permissions**.
    e) In the **Role** field, type the name of the role you specified in step 12.
    f) Click **Save**.
    The virtual server can now be enabled or disabled by a user logged in with the name you specified in step 6.
16. To delegate permissions for all of the pool members in a pool, do these sub-steps.
    a) Under **LOCAL TRAFFIC**, click **Pools**.
    b) Click the name of the pool to which the pool member belongs.
    The properties screen for the selected pool opens.
    c) Click **Permissions**.
    d) In the **Role** field, type the name of the role you created in steps 12.
    e) Click **Save & Close**.
    Pool members in this pool can now be enabled, disabled, or forced offline by a user logged in with the name you specified in step 6.

Create a new node

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

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

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

---

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

8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the node.
9. In the **Description** field, type in a brief description for the node you are creating.
10. To specify configuration parameters for this node, expand **Configuration** and continue with the next steps. Otherwise, click **Save & Close**.
11. Specify the **Health Monitors** for this node.
    - If the BIG-IP® device uses the Node Default setting, select **Node Default**.

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

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

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

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

Change pool or pool member settings

Using the BIG-IQ® user interface to make revisions to your pool or pool member configurations simplifies managing your devices.

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Pools**.
   The screen displays a list of the pools that are defined on this BIG-IQ.
3. Click the name of the pool that you want to change.

   If you select the check box for the pool instead of the name, you can either delete or deploy the pool, or you can view statistics for the pool.

   The Properties screen for the pool opens.
4. Make changes to the pool properties you want to modify.

Note: *For detailed information on the impact of using a particular pool parameter value, refer to the BIG-IP Local Traffic Manager: Implementations on* `support.f5.com`*. For the most comprehensive detail, use the work flow that best matches the purpose of the pool you are configuring.*

5. If you want to edit a pool member:
   a) Click the name of the pool member that you want to change.

      The Properties screen for the pool member opens.
   b) If you select the check box for the member instead of the name, you can enable, disable, or force the member offline. You can also use the **More** button, and then either delete the member, or view statistics for it.
   c) Make changes to the pool member properties that you want to modify.

      Note: *For detailed information on the impact of using a particular pool member parameter value, refer to the BIG-IP Local Traffic Manager: Implementations on* `support.f5.com`*. For the most comprehensive detail, use the work flow that best matches the purpose of the pool member you are configuring.*
   d) When you are satisfied with the changes you have made to the pool member, click **Save & Close**.
6. You can edit another pool member, or expand the Advance Properties area and make additional pool parameter changes.

   Note: *For detailed information on the impact of using a particular pool member parameter value, refer to the BIG-IP Local Traffic Manager: Implementations on* `support.f5.com`*. For the most comprehensive detail, use the work flow that best matches the purpose of the pool member you are configuring.*
7. To make revisions to the permissions associated with this pool, on the left, click **Permissions**.

   Note: *For detailed information about managing permissions, refer to Users User Groups Roles and Authentication in F5 BIG-IQ Centralized Management: Licensing and Initial Setup on* `support.f5.com`*.*
8. When you are satisfied with the changes you have made to the pool, click **Save & Close**.
   The revisions you saved are made, and the Properties or Permissions screen for the pool closes.

Changes that you make to pools or pool members 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.

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

Note:

Create a new SNAT pool

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

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

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

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

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

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

Managing Local Traffic Monitors

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.

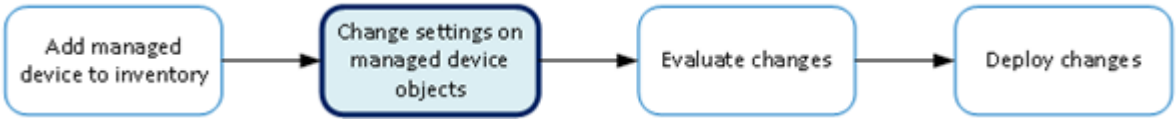

Figure 4: Change managed object workflow

What LTM monitor management tasks can I perform?

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

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

Create an LTM monitor

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

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Monitors**.
3. Click **Create**.
   The New Monitor screen opens.
4. In the **Name** field, type in a name for the monitor you are creating.
5. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the monitor.
6. In the **Description** field, type in a brief description for the monitor you are creating.
7. For **Type**, select the type of monitor you want to create.
   The **Monitor Template** setting displays.
8. 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**.
9. For **Interval**, either use the default, or specify, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown.

10. 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.

11. 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.

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

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

13. 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**.

14. 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
```

15. 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.*

    ---

16. 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**.*

    ---

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

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

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

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

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

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

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

    The default is **None**.

22. 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**.

23. 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 might 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`.

24. 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`.

25. 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.

26. 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.

27. 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).

28. 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.

29. 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.

30. 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.

31. 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.

Note: *You can not edit monitors root monitors.*

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Monitors**.
3. Select the monitor you want to edit.
   The Monitors Properties screen opens to display the current settings for the selected monitor.
4. In the **Description** field, if this is not an imported profile you can add or revise a brief description for the monitor you are editing.
5. 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.
6. For **Up Interval**, specify which interval the system uses to perform the health check when a resource is up.

Option	Description
Disabled	Specifies that the system uses the interval specified in **Interval** to check the health of the resource.
Enabled	Enables specification of a different interval to use when checking the health of a resource that is up.
7. 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.
8. From **Timeout**, specify the number of seconds the target has in which to respond to the monitor request.

   The default is 16 seconds. If the target responds within the set time period, it is considered up. If the target does not respond within the set time period, it is considered down. Note that **Timeout** and **Time Until Up** combine to control when a resource is set to up.
9. 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**.
10. 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
```

11. 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.*

    ---

12. 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****.*

    ---

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

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

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

    ---

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

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

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

    The default is **None**.

18. 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**.

19. 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 might 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**.

20. 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.

21. 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.

22. 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.

23. 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).

24. 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.

25. 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.

26. 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.

27. 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.

28. Click **Save & Close**.

Managing Network Objects

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.

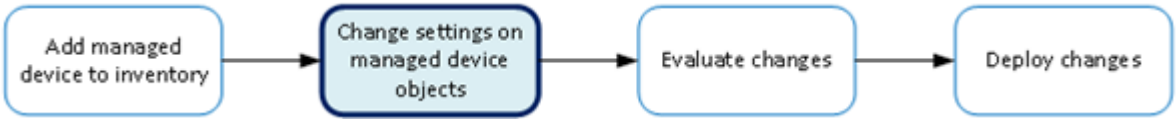

Figure 5: Change managed object workflow

Change a network object

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

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

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

Manage a network interface

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

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

---

1. At the top of the screen, click **Configuration**.
2. Under **NETWORK**, click **Interfaces**.

The Interfaces screen displays a list of network interfaces defined on devices that are managed by this BIG-IQ.

3. Select the check box for the interface you want to change, and then click **Enable** or **Disable**.
   The State for the selected interface changes on the BIG-IQ.

Create a new route

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

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

1. At the top of the screen, click **Configuration**.
2. Under **NETWORK**, click **Routes**.
3. Click **Create**.
   The New Route screen opens.
4. In the **Name** field, type in a name for the route you are creating.
5. In the **Description** field, type in a brief description for the route you are creating.
6. From the **Device** list, select the device on which to create the route.
7. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route.
8. 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
```

9. Specify the **Resource** setting.
   - To use a gateway, select **Use Gateway**, and then from **Gateway Address**, choose either **IP Address** or **IPv6 Link-Local Address**. This is the method 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 reject packets forwarded to the route destination, select **Reject**.
10. 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.
11. Click **Save & Close**.
    The system creates the new route with the settings you specified.

Create a new route domain

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

1. At the top of the screen, click **Configuration**.
2. Under **NETWORK**, select **Route Domains**.
3. Click **Create**.
   The New Route Domain screen opens.
4. In the **Name** field, type in a unique name for the route you are creating.
5. 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.

   Important: *When you assign a VLAN to a route domain, keep in mind that any self IP addresses that use that VLAN must use the same route domain. For example, if self IP* `10.0.0.0%20` *(route domain with ID 20) is assigned to VLAN-1, you cannot assign VLAN-1 to any route domain except a route domain with ID 20.*

   ---
6. In the **Description** field, type in a brief description for the route domain you are creating.
7. From the **Device** list, select the device on which to create the route domain.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the route domain.
9. Select **Strict Isolation** if you want to enforce cross-routing restrictions.

   When **Enabled** is selected, routes cannot cross route domain boundaries (so they are strictly isolated to the current route domain). The default is enabled. When this setting is 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).
10. To specify a VLAN or tunnel for the BIG-IP device to use in the route domain, select it in the **Available** list, and use the arrow to add it to the **Selected** list.

    Note: *A VLAN and tunnel can only be referenced by one route domain at a time, so if the VLAN or tunnel you select is currently referenced by another route domain, it will be removed from that route domain when you attach it to this route domain.*

    Note: *Before removing a VLAN from a route domain, recall that every self IP address must use the route domain as its VLAN. So if a self IP address uses a VLAN named* `VLAN-2` *and also uses the default route domain 0, do not remove* `VLAN-2` *from route domain 0.*

    ---
11. Click **Save & Close**.
    The system creates the new route domain with the settings you specified.

Create a new VLAN

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

1. At the top of the screen, click **Configuration**.
2. Under **NETWORK**, select **VLANs**.
3. Click **Create**.
   The New VLAN screen opens.
4. In the **Name** field, type a unique name for the VLAN you are creating.
5. In the **Description** field, type a brief description for the VLAN you are creating.

6. 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.
7. From the **Device** list, select the device on which to create the VLAN.
8. For **Partition**, type the name of the BIG-IP device partition on which you want to create the VLAN.
9. In the **MTU** field, specify the maximum transmission unit (MTU) for traffic on this VLAN.

   The default is 1500.
10. To specify which interfaces this VLAN uses for traffic management, select one from the **Interface** list, and then select the **Tagging** for it.

    You can add more than one interface by clicking the Add + button.
11. Click **Save & Close**.
    The system creates the new VLAN with the settings you specified.

Create a new self IP address

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

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

1. At the top of the screen, click **Configuration**.
2. Under **NETWORK**, click **Self IPs**.
3. Click **Create**.
   The New Self IP screen opens.
4. In the **Name** field, type in a name for the self IP address you are creating.
5. From the **Device** list, select the device on which to create the self IP address.
6. For **Partition**, type the name of the BIG-IP device partition on which you want to create the self IP.
7. 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.
8. 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).
9. For **VLAN/Tunnel**, select the VLAN or tunnel to associate with this self IP address.

   Important: *When you assign a VLAN to this self IP address, keep in mind that the self IP address and VLAN must use the same route domain. For example, if you assign self IP `10.0.0.0%20` (route domain with ID 20) to VLAN-1, you will not be able to assign VLAN-1 to any route domain except a route domain with ID 20.*
10. Specify the **Port Lockdown**.
    - Select **Allow Default** to activate only the default protocols and services. You can determine the supported protocols and services by logging in to the target BIG-IP device and running `tmsh list net self-allow defaults` on the command line.
    - Select **Allow All** to activate all TCP and UDP services on this self IP address.

- Select **Allow None** to specify that this self IP address accepts no traffic. If you are using this self IP address as the local endpoint for WAN optimization, select this option to avoid potential port conflicts.
- Select **Allow Custom** or **Allow Custom (Include Default)** to expand the **Custom List** setting, where you can specify the ports, protocols, and services to activate on this self IP address.

11. For the **Traffic Group**, select a specific traffic group for the self IP address.
12. Click **Save & Close**.
    The system creates the new self IP address with the settings you specified.

Create a new DNS resolver

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

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

   The default is the default route domain.
8. To specify the Resolver properties, expand the Resolver area.
9. 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.
10. 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.
11. 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.
12. To specify the Traffic properties, expand the area and select the format or formats for which you want the system to answer and issue queries.
13. To specify a forward zone used to resolve matching DNS queries, expand the Forward Zones area and click **Add**.

    A popup screen opens.
    a) In the **Name** field, type in a unique name for the forward zone you are creating.
    b) In the **Address** field, type in an IP address for the forward zone you are creating.
    c) In the **Service Port** field, type in the port number for the forward zone you are creating.
    d) Click the **Add** button next to the Service Port.
       The address and port combination is added to the **Nameservers** box.
    e) To add additional nameservers, repeat the last two sub-steps.
14. When you are satisfied with the new forward zone, click the **Add** button.

15. If you have specified forward zones, select the check boxes for the zones you want to use.
16. When you are satisfied with the new DNS resolver, click **Save & Close**.
    The system creates the new DNS resolver with the settings you specified.

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

Configure the BIG-IQ to manage an IPsec tunnel

How do I start managing an IPsec tunnel?

You can use BIG-IQ® Centralized Management to manage an IPsec tunnel. To set up IPsec tunnel management, you need to:

- Configure a data collection device.
- Configure the BIG-IQ system to manage the IPsec tunnel.
  - Create a forwarding virtual server for IPsec.
  - Create an IKE peer.
  - Create a custom IPsec policy.
  - Create a bidirectional IPsec traffic selector.
  - Configure the IKE daemon.
  - Verify IPsec connectivity.

After you complete these initial configuration tasks, you can manage the settings that control your IPsec tunnel traffic. You can also use the BIG-IQ statistics to troubleshoot the tunnel health.

Create a forwarding virtual server for IPsec

For IPsec, you create a forwarding (IP) type of virtual server to intercept IP traffic and direct it over the tunnel. With a forwarding (IP) virtual server, destination address translation and port translation are disabled.

1. At the top of the screen, click **Configuration**.
2. Under **LOCAL TRAFFIC**, select **Virtual Servers**.
3. Click **Create**.
   The New Virtual Server screen opens.
4. For **Name**, type in a name for the virtual server you are creating.
5. From **Device**, select the device on which to create the virtual server.
6. For **Partition**, type the name of the BIG-IP® device partition on which you want to create the virtual server.
7. For **Description**, type in a brief description for the virtual server you are creating.
8. For **Destination Address**, type a wildcard network address in CIDR format, such as `0.0.0.0/0` for IPv4 or `::/0` for IPv6, to accept any traffic.
9. From **Service Port**, select ***All Ports**.
10. From **Protocol**, select ***All Protocols**.
11. For **VLANs and Tunnel Traffic**, retain the default selection, **All VLANs and Tunnels**.
12. Leave all other fields at their default settings.
13. Click **Save & Close**.
    The system creates the new virtual server with the settings you specified.

Create an IKE peer

The IKE peer object identifies to the system you are configuring the other device that it communicates with during Phase 1 negotiations. The IKE peer object also specifies the specific algorithms and credentials to use for Phase 1 negotiation.

Important: *You must configure the devices at both ends of the IPsec tunnel.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK** > **IPsec** and then click **IKE Peers**.
3. Click **Create**.
   The New IKE Peer screen opens.
4. For **Name**, type a unique name for the IKE peer.
5. For **Description**, type a brief description of the IKE peer.
6. From **Device**, select the hostname of the device for which you are creating the new peer.
7. For the remainder of the fields on this screen, configure the values as you would if you were configuring an IKE peer on a BIG-IP® device.

   Note: *For details on configuring an IKE peer, refer to the BIG-IP TMOS: Tunneling and IPsec documentation on* `support.f5.com`
8. Click **Save & Close**.
   The system creates the new IKE peer with the settings you specified.

Create a custom IPsec policy

You can create a custom IPsec policy so that you can use a policy other than the default IPsec policy (`default-ipsec-policy` or `default-ipsec-policy-isession`). A typical reason for creating a custom IPsec policy is to configure IPsec to operate in Tunnel rather than Transport mode. Another reason is to add payload compression before encryption.

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK** > **IPsec** and then click **IPsec Policies**.
3. Click **Create**.
   The New IPsec Policy screen opens.
4. For **Name**, type a unique name for the policy.
5. For **Description**, type a brief description of the policy.
6. For the remainder of the fields on this screen, configure the values as you would if you were configuring an IKE peer on a BIG-IP® device.

   Note: *For details on configuring a IPsec security policy, refer to the BIG-IP TMOS: Tunneling and IPsec documentation on* `support.f5.com`.
7. Click **Save & Close**.
   The system creates the new security policy with the settings you specified.

Create a bidirectional IPsec traffic selector

A traffic selector filters traffic based on the IP addresses and port numbers that you specify, as well as the custom IPsec policy you assign.

Important: *You must configure the devices at both ends of the IPsec tunnel.*

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK** > **IPsec** and then click **Traffic Selectors**.
3. Click **Create**.
   The New Traffic Selector screen opens.
4. For **Name**, type a unique name for the traffic selector.
5. For **Description**, type a brief description of the traffic selector.
6. From **Device**, select the hostname of the device for which you are creating the new traffic selector.
7. For the remainder of the fields on this screen, configure the values as you would if you were configuring a traffic selector on a BIG-IP® device.

   Note: *For details on configuring a traffic selector, refer to the BIG-IP TMOS: Tunneling and IPsec documentation on* `support.f5.com`.
8. Click **Save & Close**.
   The system creates the new traffic selector with the settings you specified.

Configure the IKE daemon

To complete the configuration sequence for managing an IPsec tunnel on the BIG-IQ®, you need to configure the IKE daemon

1. At the top of the screen, click **Configuration**.
2. On the left, expand **NETWORK** > **IPsec** and then click **IKE Daemon**.
3. In the Name column, select the **ikedaemon** link that corresponds to the host name of the BIG-IP® device from which you imported the IPsec tunnel configuration.
   The IKE daemon properties screen for that BIG-IP device opens.
4. For External Log Publisher, select **default-ipsec-log-publisher**.
5. Click the **Save & Close** button at the bottom of the screen.

Verify IPsec connectivity

After you have configured an IPsec tunnel and before you configure additional functionality, you can verify that the tunnel is passing traffic.

Note: *Only data traffic matching the traffic selector triggers the establishment of the tunnel.*

1. At the top of the screen, click **Monitoring**.
2. On the left, expand **EVENTS** > **IPsec** and click **Events**.
   The IPsec Event Logs screen opens.

3. Examine the screen, looking for event logs that relate to successful IPsec tunnel creation, to confirm IPsec connectivity.

Configure IPsec event viewing on the BIG-IQ

How do I configure viewing IPsec event logs?

You can use BIG-IQ® Centralized Management to view IPsec events. To set up IPsec event log viewing, you need to:

- Configure the BIG-IP® devices that comprise the IPsec tunnel to send events to the data collection device.
  - Create a log publisher pool.
  - Create a remote high-speed log destination for IPsec.
  - Create a remote Syslog destination for IPsec.
  - Configure a log publisher to send IPsec events to the BIG-IQ.
- Configure the BIG-IQ system to view IPsec events.
  - Import IPsec configuration settings from the BIG-IP device.
  - Enable IPsec event collection.

After you complete these initial configuration tasks, you can view IPsec events on the BIG-IQ.

Create a log publisher pool

Creating a log publisher pool is part of the sequence you perform to route IPsec events from the BIG-IP® device to your data collection device so that you can view these events from the BIG-IQ®.

Important: *Perform this task on the BIG-IP devices that comprise the IPsec tunnel; not on the BIG-IQ.*

Important: *You must perform these steps on both of the BIG-IP devices that comprise the IPsec tunnel.*

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.

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

   Important: *The pool name is limited to 63 characters.*
4. Under Resources, create a new member.
   a) For **Node Name**, type a name for the member
   b) For **Address**, type the self IP address of your data collection device.
   c) For **Service Port**, type `9997`.
   d) Click **Add**.
      The system creates the new pool member.
5. Click **Finished**.

The log publisher pool you created is added to the pools list.

Create a remote high-speed log destination for IPsec

Before creating a remote high-speed log destination for IPsec, you must create a log publishing pool.

Creating a remote high-speed log destination is part of the sequence you perform to route IPsec events from the BIG-IP® device to your data collection device so that you can view these events from the BIG-IQ®.

Important: *Perform this task on the BIG-IP devices that comprise the IPsec tunnel; not on the BIG-IQ.*

Important: *You must perform these steps on both of the BIG-IP devices that comprise the IPsec tunnel.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a name to identify the IPsec remote high speed log destination.
4. From the **Type** list, select **Remote High-Speed Log**.
5. From the **Pool Name** list, select the log publisher pool that you defined previously.
6. From the **Protocol** list, select the TCP protocol.
7. Click **Finished**.

Create a remote Syslog destination for IPsec

Before creating a remote Syslog log destination for IPsec, you must create a log publishing pool and a high-speed log destination for IPsec.

Creating a remote Syslog log destination is part of the sequence you perform to route IPsec events from the BIG-IP® device to your data collection device so that you can view these events from the BIG-IQ® system.

Important: *Perform this task on the BIG-IP devices that comprise the IPsec tunnel; not on the BIG-IQ.*

Important: *You must perform these steps on both of the BIG-IP devices that comprise the IPsec tunnel.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Destinations**.
   The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type `IPsec-Syslog` to identify the IPsec Syslog destination.
4. From the **Type** list, select **Remote Syslog**.
5. From the **Syslog Format** list, select a format for the logs.
6. From the **Forward To** list, select the name of the IPsec remote high speed log.
7. Click **Finished**.

Configure a log publisher to send IPsec events to the BIG-IQ

To send the IPsec event logs to the data collection device, you must configure a publisher to send them to the IPsec Syslog destination. This is the last task in the sequence you perform to route IPsec events from the BIG-IP® device to your data collection device so that you can view these events from the BIG-IQ®

Important: *Perform this task on the BIG-IP devices that comprise the IPsec tunnel; not on the BIG-IQ.*

Important: *You must perform these steps on both of the BIG-IP devices that comprise the IPsec tunnel.*

1. On the Main tab, click **System** > **Logs** > **Configuration** > **Log Publishers**.
   The Log Publishers screen opens.
2. Click the log publisher named **default-ipsec-log-publisher**.
   The Log Publisher properties screen opens.
3. For the **Destinations** setting, select **IPsec-Syslog** from the **Available** list, and move it to the **Selected** list.
   Both **local-syslog** (the default) and **IPsec-Syslog** are listed in the **Selected** list.
4. Click **Update**.

IPsec events will now route to the data collection device.

To use the IPsec tunnel configuration that you just completed on the BIG-IQ, you must import the settings for this device to the BIG-IQ.

Import IPsec configuration settings from the BIG-IP device

Before you can import settings from a managed device, you must have completed the configuration task on the BIG-IP® device. See *Configure the BIG-IP device to send IPsec events to your data collection device* for details.

To manage an IPsec tunnel on BIG-IQ®, you need to import the settings configured on the BIG-IP devices that reside on one each end of the tunnel.

Important: *Perform this task on the BIG-IQ for each of the BIG-IP devices that make up the IPsec tunnel.*

1. At the top of the screen, click **Devices**.
2. In the Services column, click the link that lists the currently managed services for the BIG-IP device that you configured with IPsec tunnel settings .
   The Services screen for the selected device opens.
3. For Local Traffic (LTM), click **Re-discover**.
   The system discovers the LTM configuration settings for the BIG-IP device
4. For Local Traffic (LTM), click **Re-import**.
   The system imports the LTM service for the BIG-IP device
5. Click **Cancel**.

The IPsec tunnel settings you configured on the BIG-IP device are imported for the selected device.

Enable IPsec event collection

To view IPsec tunnel events on BIG-IQ®, you need to activate IPsec event collection for your data collection device (DCD) cluster.

1. At the top of the screen, click **System**.
2. On the left, expand **BIG-IQ DATA COLLECTION** and then select **BIG-IQ Data Collection Devices**.
   The BIG-IQ Data Collection Devices screen opens to list the data collection devices in the cluster
3. In the Services column, click **Add Services**.
   The Services screen for this DCD opens.
4. For IPsec, click **Activate**.

   The Listener Address displays the internal self IP address configured for the DCD. The self IP address is currently the recommended address for collecting event log data.

   The system begins collecting IPsec events.
5. Click the **Save & Close** button at the bottom of the screen.

You can now view IPsec event logs using the BIG-IQ user interface.

Troubleshooting an IPsecTunnel

Troubleshoot an unhealthy IPsec tunnel using performance statistics

Before you can troubleshoot the tunnel using statistics:

- You must have configured BIG-IQ® to display statistics for your IPsec tunnel.
- You need to know the IP address or host name of the BIG-IP® devices that form the IPsec tunnel.

When you learn that an IPsec tunnel is unhealthy (for example, your helpdesk might have opened a ticket), you can use the IPsec performance statistics to troubleshoot the tunnel.

Note: *If one end of the tunnel uses a device other than a BIG-IP device, you can troubleshoot only that end of the tunnel.*

---

1. At the top of the screen, click **Devices**.
2. Find one of the BIG-IP devices that form the IPsec tunnel.
   - If you have the IP address of the device, from the **Filter** selector, select **Address** and type the IP address of the BIG-IP device.
   - If you have the host name of the device, from the **Filter** selector, select **Device Name** and type the host name of the BIG-IP device.

   The filter you created displays at the top of the screen and only the BIG-IP device you identified is listed.
3. Click the device name for the BIG-IP device.
   The properties screen for the device opens.
4. On the left, click **Health**.
   A health summary screen displays current usage levels for the device.
5. In the upper right corner, click **View Health Statistics**.
   The Device Health statistics summary page opens, displaying data only for the selected BIG-IP device.
6. Scan the graphs for details about the device's performance that reveal the source of the issue. If you find the issue, skip to step 11.
7. In the upper left corner, click the back arrow.
   The health summary screen for the device opens again.
8. In the upper right corner, click **View Traffic Statistics**.
   The Device Traffic statistics summary page opens, displaying data only for the selected BIG-IP device.
9. Scan the graphs for details about the device's performance that reveal the source of the issue. If you find the issue, skip to step 11.
10. If you don't find the source of the problem after examining the traffic and device health statistics, delete the filter you created in step 2, and then repeat the last 8 steps for the other BIG-IP device in the IPsec tunnel. If only one end of the tunnel is made up of a BIG-IP device, proceed to the task *Troubleshoot an unhealthy IPsec tunnel using event logs*, to see if you can isolate the issue by inspecting the IPsec event logs. If you find the issue, skip to step 11.
11. Fix the issues you discovered with the configuration objects, and then deploy those changes to the relevant BIG-IP devices to resolve the problem.

If you were not able to isolate the cause of the issue, perform the task: *Troubleshoot an unhealthy IPsec tunnel using event logs*.

Troubleshoot an unhealthy IPsec tunnel using event logs

Before you can troubleshoot a tunnel by examining the IPsec event logs, you must have configured IPsec event logging. (See *Configure IPsec event viewing on the BIG-IQ* for details.)

When you learn that an IPsec tunnel is unhealthy (for example, your helpdesk might have opened a ticket), you can troubleshoot the tunnel by examining the IPsec event logs.

1. At the top of the screen, click **Monitoring**.
2. On the left, expand **EVENTS** > **IPsec** and click **Events**.
   The IPsec Event Logs screen opens and displays all of the logs collected from your IPsec tunnel.
3. Use the **DEVICE**, **TIMEFRAME**, and **LOG LEVEL** filters to display the logs that you think will reveal the source of the issue.
4. Analyze the log of events to find the issue that is causing the IPsec tunnel to perform improperly.
5. Fix the issues you discover, and then deploy those changes to the relevant BIG-IP® devices.

Legal Notices

Legal notices

Publication Date

This document was published on July 6, 2017.

Publication Number

MAN-0577-07

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

C

centralized management
  of BIG-IP devices *9*, *15*, *17*, *27*, *33*
configuring
  IPsec event viewing *43*
  IPsec tunnel management *39*
creating
  iRules *15*
  Local Traffic monitors *27*
  network objects *33*
  pools & pool members *17*
  profiles *5*
  virtual servers *9*

D

device inventory
  about *9*, *15*, *17*, *27*, *33*
device management
  about *9*, *15*, *17*, *27*, *33*
device managers
  how to manage *27*
device profile
  creating for LTM *5*
  editing for LTM *6*
device profiles
  how to manage *5*
devices
  about discovering *9*, *15*, *17*, *27*, *33*
disable permissions
  assigning *21*
discovery
  defined *9*, *15*, *17*, *27*, *33*
DNS resolvers
  creating *37*

E

editing
  iRules *15*
  Local Traffic & Network profiles *5*
  Local Traffic monitors *27*
  network objects *33*
  pools & pool members *17*
  virtual servers *9*
enable and disable permissions
  delegating *21*

F

forwarding virtual servers
  creating for IPsec *39*

I

IKE daemon
IKE daemon *(continued)*
  configure *41*
IKE peers
  creating for IPsec *40*
IPsec configuration
  importing from BIG-IP *45*, *46*
IPsec event viewing
  configuring *43*
IPsec events
  how to start viewing *43*
IPsec IKE peers
  creating *40*
IPsec policy
  creating *40*
IPsec traffic selectors
  creating *41*
IPsec tunnel
  how to set up management *39*
  start managing *39*
  Troubleshooting *47*
  troubleshooting using event logs *48*
  troubleshooting using statistics *47*
  verifying connectivity *41*
IPsec tunnel management
  configuring *39*
IPsec Tunnel mode
  verifying connectivity *41*
iRules
  creating new *15*
  managing *15*

L

Local Traffic & Network profiles
  managing *5*
Local Traffic monitors
  managing *27*
log destinations
  creating *44*
log publisher pool
  creating *43*
logging
  and publishers *45*
LTM monitor
  adding *27*
  editing *30*
LTM profile
  creating *5*
  editing *6*

M

managed devices
  about discovering *9*, *15*, *17*, *27*, *33*
  changing objects *33*
  changing objects for *13*, *23*
managing iRules *15*
managing Local Traffic monitors *27*

managing network objects *33*
managing pools & pool members *17*
managing profiles *5*
managing virtual servers *9*

N

network interfaces
    managing *33*
network objects
    managing *33*
nodes
    creating *22*

P

pending version
    defined *13*, *23*, *33*
permissions
    delegating *21*
pool
    creating *43*
pool members
    creating *20*
pools
    creating *17*
pools & pool members
    managing *17*
pools and pool members
    what you can manage *17*
publishers
    creating for logging *45*

R

Remote High-Speed Log
    creating *44*
Remote Syslog
    creating *44*
route domains
    creating *34*
routes
    creating *34*

S

self IP addresses
    creating *36*
servers
    and publishers for log messages *45*
settings
    changing for pool or pool members *23*
SNAT pools
    creating *24*

T

traffic selectors
    creating *41*
troubleshooting
    IPsec tunnel, using event logs *48*
troubleshooting *(continued)*
    IPsec tunnel, using statistics *47*
Troubleshooting
    IPsec tunnel *47*
Tunnel mode
    verifying connectivity *41*

V

virtual servers
    attaching iRules *12*, *16*
    cloning *12*
    creating *9*
    delegating permissions *21*
    managing *9*
    what can you manage *9*
VLANs+
    creating *35*