BIG-IP® TMOS®: Implementations

Version 13.0

Table of Contents

Customizing the BIG-IP Dashboard

Overview: BIG-IP dashboard customization

The BIG-IP® dashboard displays system statistics in selectable graphs, gauges, and tables. In addition to the pre-defined views, you can create custom combinations of the dashboard windows, called *views*, and save them in groups, called *view sets*. You can combine windows from different BIG-IP modules in a single view, or use just the windows you want for a single module. Windows are available only for those modules that you have licensed and provisioned.

Note: *The view set name for all pre-defined views is* `standard`.

Customizing the BIG-IP dashboard

You can create custom dashboard displays using the windows for any modules that are available on the BIG-IP® system.

1. On the Main tab, click **Statistics** > **Dashboard**.
   A separate window opens for the BIG-IP dashboard.
2. On the Views control bar, click the Create custom view icon.
   A blank canvas opens in design mode. The Dashboard Windows Chooser displays the available windows, grouped by module. You can click a module to display the available windows.
3. From the Dashboard Windows Chooser, drag and drop the windows you want onto the canvas.

   After you drag a window to the canvas, you can resize it or change it to display the information you want by clicking a tab or filter.

   Note: *The windows are not active when in design mode, so the data does not update in real time.*
4. When you have placed the windows you want onto the canvas, click the Save icon on the Custom Views control bar.
   The Save View popup window opens.
5. Type a name for the view.
6. Type a new name for the view set, or select from the list.
7. Click **OK**.
   The new view is saved, and appears in the **Views** list.
8. Click the double-gear icon on the Custom Views control bar to return to active mode.
   The dashboard displays the custom view you just created, and updates the display with real-time data.

Web Hosting Multiple Customers Using an External Switch

Overview: Web hosting multiple customers using an external switch

You can use the BIG-IP® system to provide hosting services, including application delivery, for multiple customers.

To host multiple web customers, you can incorporate an external switch into the configurations. In this illustration, the BIG-IP system has an interface (5.1) assigned to three VLANs on a network. The three VLANs are **vlanA**, **vlanB**, and **vlanB**. Interface **5.1** processes traffic for all three VLANs. Note that each VLAN contains two servers, and serves a specific customer.

Tip: *An alternate way to implement web hosting for multiple customers is to use the route domains feature.*

Illustration for hosting multiple customers using an external switch

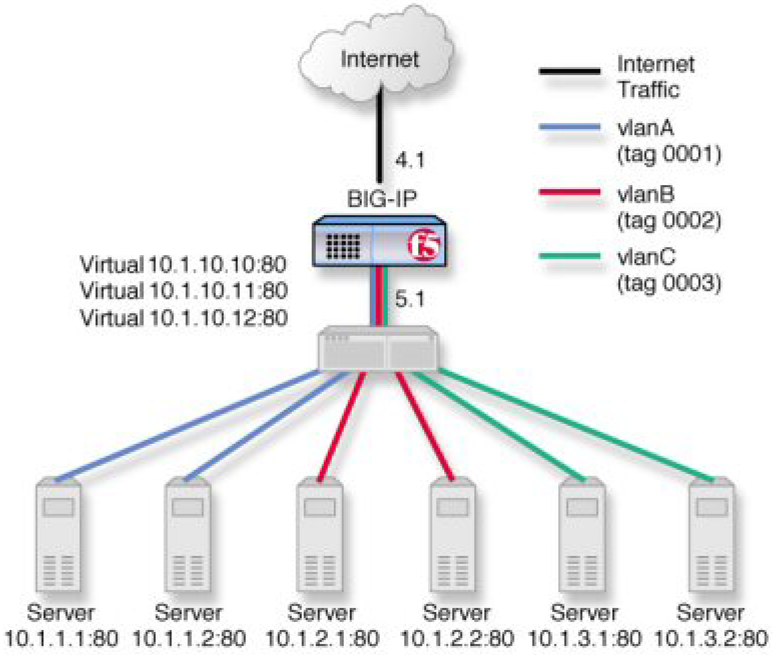

Figure 1: Hosting multiple customers using an external switch

Task summary for hosting multiple customers

Perform these tasks to host multiple customers using an external switch.

Task list

Creating a VLAN with a tagged interface

Creating a load balancing pool
Creating a virtual server for HTTP traffic

Creating a VLAN with a tagged interface

When you create a VLAN with tagged interfaces, each of the specified interfaces can process traffic destined for that VLAN.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. Click **Create**.
   The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.

   The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting:
   a) From the **Interface** list, select an interface number.
   b) From the **Tagging** list, select **Tagged**.
   c) Click **Add**.
6. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. From the **Configuration** list, select **Advanced**.
9. For the **Hardware SYN Cookie** setting, select or clear the check box.

   When you enable this setting, the BIG-IP system triggers hardware SYN cookie protection for this VLAN.

   Enabling this setting causes additional settings to appear. These settings appear on specific BIG-IP platforms only.
10. For the **Syncache Threshold** setting, retain the default value or change it to suit your needs.

    The **Syncache Threshold** value represents the number of outstanding SYN flood packets on the VLAN that will trigger the hardware SYN cookie protection feature.

    When the **Hardware SYN Cookie** setting is enabled, the BIG-IP system triggers SYN cookie protection in either of these cases, whichever occurs first:

    - The number of TCP half-open connections defined in the LTM® setting **Global SYN Check Threshold** is reached.
    - The number of SYN flood packets defined in this **Syncache Threshold** setting is reached.
11. For the **SYN Flood Rate Limit** setting, retain the default value or change it to suit your needs.

    The **SYN Flood Rate Limit** value represents the maximum number of SYN flood packets per second received on this VLAN before the BIG-IP system triggers hardware SYN cookie protection for the VLAN.
12. Click **Finished**.
    The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Creating a load balancing pool

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   ***Tip:** Hold the Shift or Ctrl key to select more than one monitor at a time.*
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
   a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
   b) In the **Address** field, type an IP address.
   c) In the **Service Port** field, type a port number, or select a service name from the list.
   d) (Optional) In the **Priority** field, type a priority number.
   e) Click **Add**.
8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server for HTTP traffic

This task creates a destination IP address for application traffic. As part of this task, you must assign the relevant pool to the virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

   ***Note:** The IP address you type must be available and not in the loopback network.*
5. In the **Service Port** field, type `80`, or select **HTTP** from the list.
6. From the **HTTP Profile** list, select **http**.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. Click **Finished**.

You now have a virtual server to use as a destination address for application traffic.

Web Hosting Multiple Customers Using Untagged Interfaces

Overview: Web hosting multiple customers using untagged interfaces

One way to implement web hosting for multiple customers is to use multiple interfaces on the BIG-IP® system to directly host traffic for multiple customers, without the need for an external switch. With this scenario, you must configure the VLANs with untagged instead of tagged interfaces. As shown in the following illustration, two BIG-IP system interfaces are assigned to each VLAN. For example, interfaces **1.1** and **1.2** are assigned to VLAN **vlanA**. Each interface is assigned to a VLAN as an untagged interface.

Tip: *An alternate way to implement web hosting for multiple customers is to use the route domains feature.*

Illustration for hosting multiple customers using untagged interfaces

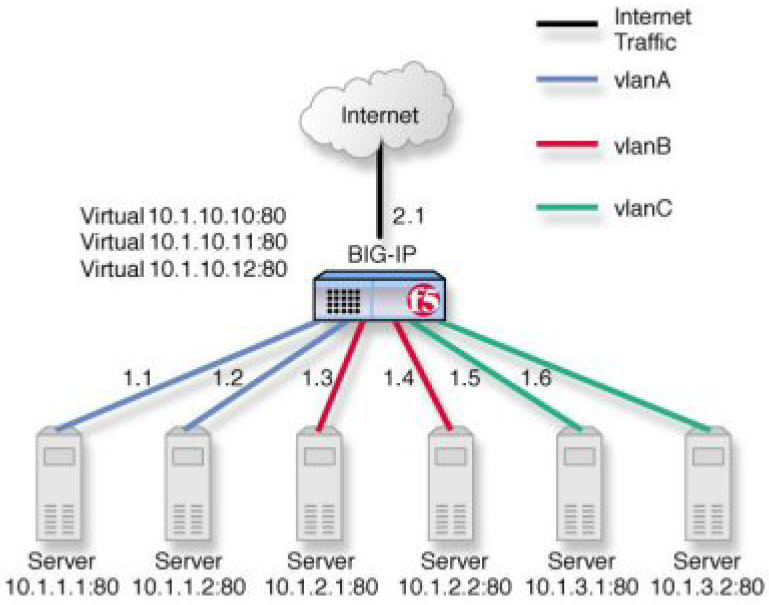

Figure 2: Hosting multiple customers using untagged interfaces

Task summary for hosting multiple customers

Perform these tasks to host multiple customers using tagged interfaces on VLANs.

Task list

Creating a VLAN with an untagged interface
Creating a load balancing pool
Creating a virtual server for HTTP traffic

Creating a VLAN with an untagged interface

You can create a VLAN that uses untagged interfaces.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. Click **Create**.
   The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.

   The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting,
   a) From the **Interface** list, select an interface number.
   b) From the **Tagging** list, select **Untagged**.
   c) Click **Add**.
6. Click **Finished**.
   The screen refreshes, and displays the new VLAN in the list.

The interfaces that you specified in this task process traffic for this VLAN only.

Creating a load balancing pool

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   ***Tip:** Hold the Shift or Ctrl key to select more than one monitor at a time.*
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
   a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
   b) In the **Address** field, type an IP address.
   c) In the **Service Port** field, type a port number, or select a service name from the list.
   d) (Optional) In the **Priority** field, type a priority number.

   e) Click **Add**.

8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server for HTTP traffic

This task creates a destination IP address for application traffic. As part of this task, you must assign the relevant pool to the virtual server.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

   Note: *The IP address you type must be available and not in the loopback network.*
5. In the **Service Port** field, type `80`, or select **HTTP** from the list.
6. From the **HTTP Profile** list, select **http**.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. Click **Finished**.

You now have a virtual server to use as a destination address for application traffic.

Web Hosting Multiple Customers Using Route Domains

Overview: Use of route domains to host multiple web customers on the BIG-IP system

Using the *route domains* feature of the BIG-IP® system, you can provide hosting service for multiple customers by isolating each type of application traffic within a defined address space on the network. This enhances security and dedicates BIG-IP resources to each application.

Using route domains, you can also use duplicate IP addresses on the network, provided that each of the duplicate addresses resides in a separate route domain and is isolated on the network through a separate VLAN. For example, if you are processing traffic for two different customers, you can create two separate route domains. The same node address (such as `10.0.10.1`) can reside in each route domain, in the same pool or in different pools, and you can assign a different monitor to each of the two corresponding pool members.

A good example of the use of traffic isolation on a network is an ISP that services multiple customers, where each customer deploys a different application. The first illustration shows two route domain objects on a BIG-IP system, where each route domain corresponds to a separate customer, and thus, resides in its own partition. Within each partition, the ISP created the network objects and local traffic objects required for that customer's application (AppA or AppB).

The sample configuration results in the BIG-IP system segmenting traffic for two different applications into two separate route domains. The routes for each application's traffic cannot cross route domain boundaries because cross-routing restrictions are enabled on the BIG-IP system by default. The second illustration shows the resulting route isolation for AppA and AppB application traffic.

Illustration of sample BIG-IP configuration using route domains

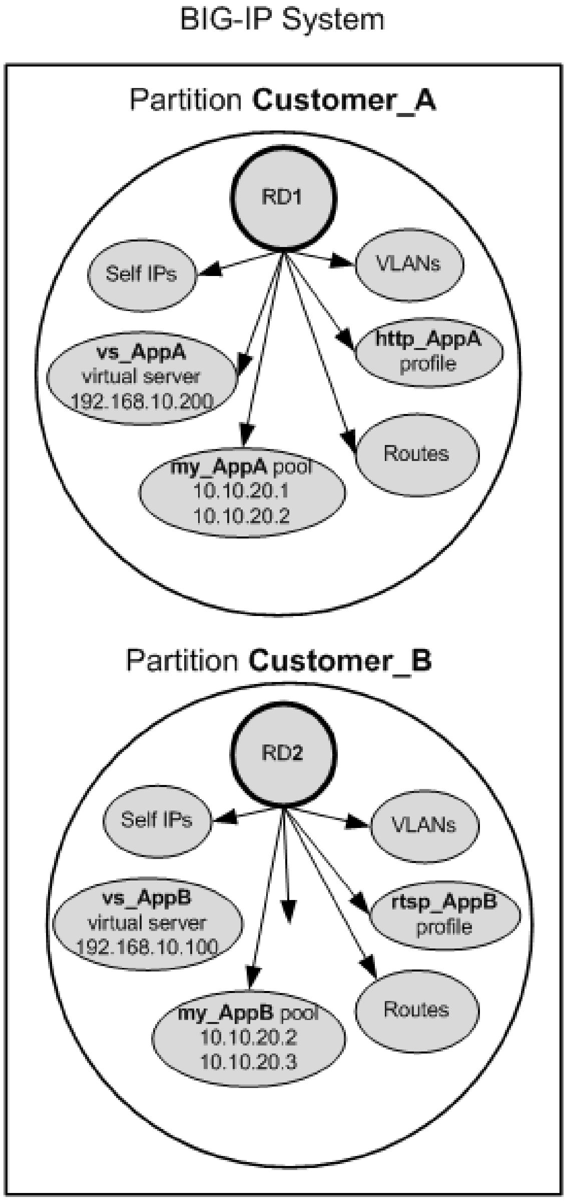

Figure 3: Sample BIG-IP configuration using route domains

Illustration of resulting route domain configuration

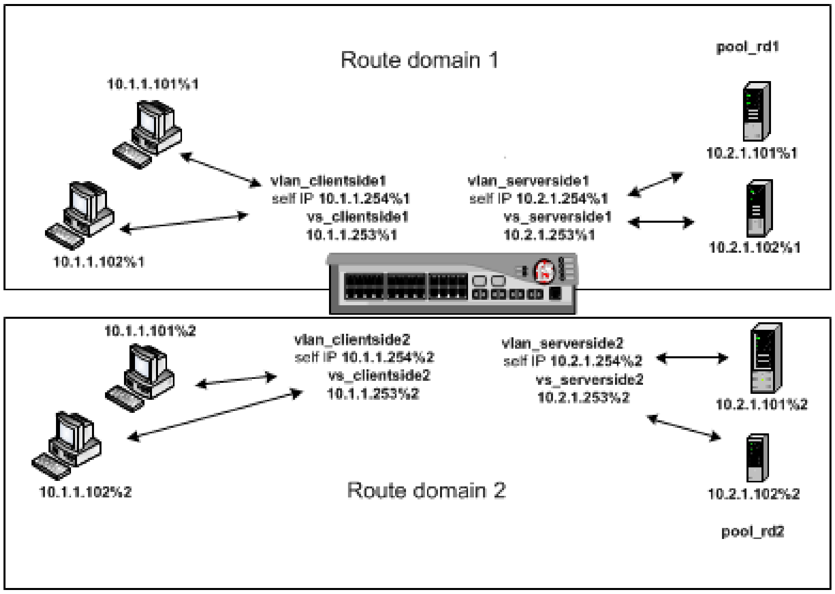

Figure 4: Resulting route domain configuration

Task summary

Perform these tasks to host multiple web customers using route domains.

Task list

Creating an administrative partition
Creating a VLAN with a tagged interface
Creating a self IP address for a default route domain in an administrative partition
Creating a route domain on the BIG-IP system
Creating a load balancing pool
Creating a virtual server
Configuring route advertisement for a virtual address
Adding routes that specify VLAN internal as the resource

Creating an administrative partition

You perform this task to create an administrative partition. An *administrative partition* creates an access control boundary for users and applications.

1. On the Main tab, expand **System** and click **Users**.
   The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. Click **Create**.
   The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.
   An example of a partition name is `Spanned_VIP`.
5. Type a description of the partition in the **Description** field.

This field is optional.

6. For the **Device Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the device group attribute from folder `root`.
Clear the check box and select the name of a device group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the device group attribute from folder `root`.

7. For the **Traffic Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder `root`.
Clear the check box and select the name of a traffic group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder `root`.

8. Click **Finished**.

The new partition appears in the partition list.

Creating a VLAN with a tagged interface

When you create a VLAN with tagged interfaces, each of the specified interfaces can process traffic destined for that VLAN.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. Click **Create**.
   The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.

   The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting:
   a) From the **Interface** list, select an interface number.
   b) From the **Tagging** list, select **Tagged**.
   c) Click **Add**.
6. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
7. In the **MTU** field, retain the default number of bytes (**1500**).
8. From the **Configuration** list, select **Advanced**.
9. For the **Hardware SYN Cookie** setting, select or clear the check box.

   When you enable this setting, the BIG-IP system triggers hardware SYN cookie protection for this VLAN.

   Enabling this setting causes additional settings to appear. These settings appear on specific BIG-IP platforms only.
10. For the **Syncache Threshold** setting, retain the default value or change it to suit your needs.

The **Syncache Threshold** value represents the number of outstanding SYN flood packets on the VLAN that will trigger the hardware SYN cookie protection feature.

When the **Hardware SYN Cookie** setting is enabled, the BIG-IP system triggers SYN cookie protection in either of these cases, whichever occurs first:

- The number of TCP half-open connections defined in the LTM® setting **Global SYN Check Threshold** is reached.
- The number of SYN flood packets defined in this **Syncache Threshold** setting is reached.

11. For the **SYN Flood Rate Limit** setting, retain the default value or change it to suit your needs.

    The **SYN Flood Rate Limit** value represents the maximum number of SYN flood packets per second received on this VLAN before the BIG-IP system triggers hardware SYN cookie protection for the VLAN.
12. Click **Finished**.
    The screen refreshes, and displays the new VLAN in the list.

The new VLAN appears in the VLAN list.

Creating a self IP address for a default route domain in an administrative partition

Before creating a self IP address, ensure that you have created an internal VLAN and an external VLAN on the BIG-IP system.

Using this procedure, you must create two self IP addresses on the BIG-IP system. One self IP address is associated with the internal VLAN, and the other is associated with the external VLAN. Self IP addresses enable the BIG-IP system and other devices on the network to route application traffic through the associated VLAN.

1. On the Main tab, click **Network** > **Self IPs**.
2. Click **Create**.
   The New Self IP screen opens.
3. In the **IP Address** field, type an IP address.

   This IP address should represent the address space of the VLAN that you specify with the **VLAN** setting. Because the route domain that you previously created is the default route domain for the administrative partition, you do not need to append the route domain ID to this IP address.

   The system accepts IP addresses in both the IPv4 and IPv6 formats.
4. In the **Netmask** field, type the network mask for the specified IP address.

   For example, you can type `255.255.255.0`.
5. From the **VLAN/Tunnel** list, select the VLAN to associate with this self IP address.

   - On the internal network, select the internal or high availability VLAN that is associated with an internal interface or trunk.
   - On the external network, select the external VLAN that is associated with an external interface or trunk.
6. Click **Finished**.
   The screen refreshes, and displays the new self IP address.

The BIG-IP system has a self IP address that is associated with the internal or external network.

Creating a route domain on the BIG-IP system

Before you create a route domain:

- Ensure that an external and an internal VLAN exist on the BIG-IP® system.
- If you intend to assign a static bandwidth controller policy to the route domain, you must first create the policy. You can do this using the BIG-IP Configuration utility.

- Verify that you have set the current partition on the system to the partition in which you want the route domain to reside.

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

1. On the Main tab, click **Network** > **Route Domains**.
   The Route Domain List screen opens.
2. Click **Create**.
   The New Route Domain screen opens.
3. In the **Name** field, type a name for the route domain.

   This name must be unique within the administrative partition in which the route domain resides.
4. In the **ID** field, type an ID number for the route domain.

   This ID must be unique on the BIG-IP system; that is, no other route domain on the system can have this ID.

   An example of a route domain ID is `1`.
5. In the **Description** field, type a description of the route domain.
   For example: `This route domain applies to application traffic for Customer A.`
6. For the **Strict Isolation** setting, select the **Enabled** check box to restrict traffic in this route domain from crossing into another route domain.
7. For the **Parent Name** setting, retain the default value.
8. For the **VLANs** setting, from the **Available** list, select a VLAN name and move it to the **Members** list.

   Select the VLAN that processes the application traffic relevant to this route domain.

   Configuring this setting ensures that the BIG-IP system immediately associates any self IP addresses pertaining to the selected VLANs with this route domain.
9. For the **Dynamic Routing Protocols** setting, from the **Available** list, select one or more protocol names and move them to the **Enabled** list.

   You can enable any number of listed protocols for this route domain.
10. From the **Bandwidth Controller** list, select a static bandwidth control policy to enforce a throughput limit on traffic for this route domain.
11. From the **Partition Default Route Domain** list, select either **Another route domain (0) is the Partition Default Route Domain** or **Make this route domain the Partition Default Route Domain**.

    This setting does not appear if the current administrative partition is partition `Common`.

    When you configure this setting, either route domain `0` or this route domain becomes the default route domain for the current administrative partition.
12. Click **Finished**.
    The system displays a list of route domains on the BIG-IP system.

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

Creating a load balancing pool

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.

2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   Tip: *Hold the Shift or Ctrl key to select more than one monitor at a time.*
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:
   a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
   b) In the **Address** field, type an IP address.
   c) In the **Service Port** field, type a port number, or select a service name from the list.
   d) (Optional) In the **Priority** field, type a priority number.
   e) Click **Add**.
8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address in CIDR format.

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `10.0.0.1` or `10.0.0.0/24`, and an IPv6 address/prefix is `ffe1::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`. When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a `/32` prefix.

   Note: *The IP address you type must be available and not in the loopback network.*
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.

Configuring route advertisement for a virtual address

Before configuring route advertisement on a virtual address, verify that you have enabled one or more dynamic routing protocols on the route domain pertaining to this virtual address. Also verify that you have configured the relevant dynamic routing protocols for route redistribution.

Perform this task to advertise a route for this virtual address to other routers on your network.

Important: *This task pertains only to configurations for which you have enabled dynamic routing protocols on the relevant route domain. If you have not enabled dynamic routing protocols on the relevant route domain, you can skip this task.*

1. On the Main tab, click **Local Traffic** > **Virtual Servers** > **Virtual Address List**.
   The Virtual Address List screen opens.
2. In the Name column, click the virtual address for which you want to advertise a route.
   This displays the properties of that virtual address.
3. Verify that the **ARP** field is selected.
4. From the **Advertise Route** list, choose one of these options:

Option	Description
When any virtual server is available	Specifies that the system advertises a route for this virtual IP address whenever any virtual server associated with this virtual IP address is available.
When all virtual servers(s) are available	Specifies that the system advertises a route for this virtual IP address whenever all virtual servers associated with this virtual IP address is available.
Always	Specifies that the system always advertises a route for this virtual IP address.

5. For the **Route Advertisement** setting, select the box.

   This makes it possible for the BIG-IP system to advertise this virtual IP address when you have enabled any dynamic routing protocols.
6. Click **Update**.
7. Repeat this task for each virtual address for which you want to advertise a route.

The BIG-IP system advertises a route for this virtual address to other routers when one or more dynamic routing protocols are enabled and are configured for route redistribution.

Adding routes that specify VLAN internal as the resource

Ensure that you set the current administrative partition to the partition in which you want a specific customer's configuration to reside.

You must add a route for each destination IP address pertaining to the route domain. A destination address in this case is typically a node address for a pool member.

1. On the Main tab, click **Network** > **Routes**.
2. Click **Add**.
   The New Route screen opens.
3. In the **Name** field, type a unique user name.

   This name can be any combination of alphanumeric characters, including an IP address.
4. In the **Destination** field, type either the destination IP address for the route, or IP address `0.0.0.0` for the default route.

   This address can represent either a host or a network. Also, if you are using the route domains and the relevant route domain is the partition default route domain, you do not need to append a route domain ID to this address.
5. In the **Netmask** field, type the network mask for the destination IP address.
6. From the **Resource** list, select **Use VLAN/Tunnel**.

   A VLAN represents the VLAN through which the packets flow to reach the specified destination.
7. From the **VLAN** list, select **Internal**.

8. Click **Finished**.

The BIG-IP system now includes routes to the nodes in the load balancing pool for a specific route domain.

Implementing the Link Layer Discovery Protocol

Overview: Implementing Link Layer Discovery Protocol

The BIG-IP® system supports Link Layer Discovery Protocol (LLDP). LLDP is a Layer 2 industry-standard protocol (IEEE 802.1AB) that gives a network device such as the BIG-IP system the ability to advertise its identity and capabilities to multi-vendor neighbor devices on a network. The protocol also enables a network device to receive information from neighbor devices.

LLDP transmits device information in the form of LLDP messages known as LLDP Packet Data Units (LLDPDUs).

In general, this protocol:

- Advertises connectivity and management information about the local BIG-IP device to neighbor devices on the same IEEE 802 LAN.
- Receives network management information from neighbor devices on the same IEEE 802 LAN.
- Operates with all IEEE 802 access protocols and network media.

Using the BIG-IP Configuration utility or `tmsh`, you can use this particular implementation to configure BIG-IP system interfaces to transmit LLDPDUs to neighbor devices. More specifically, you can:

- Specify the exact content of LLDPDUs that a BIG-IP system interface transmits to a neighbor device. You specify this content by configuring the **LLDP Attributes** setting on each individual interface.
- Globally specify the frequencies of various message transmittal properties, and specify the number of neighbors from which interfaces can receive messages. These properties apply to all interfaces on the BIG-IP system.

The following illustration shows a BIG-IP system that transmits LLDP messages to three neighbor devices: another BIG-IP system, an external switch, and an external router. Note that LLDP is enabled on all of the devices.

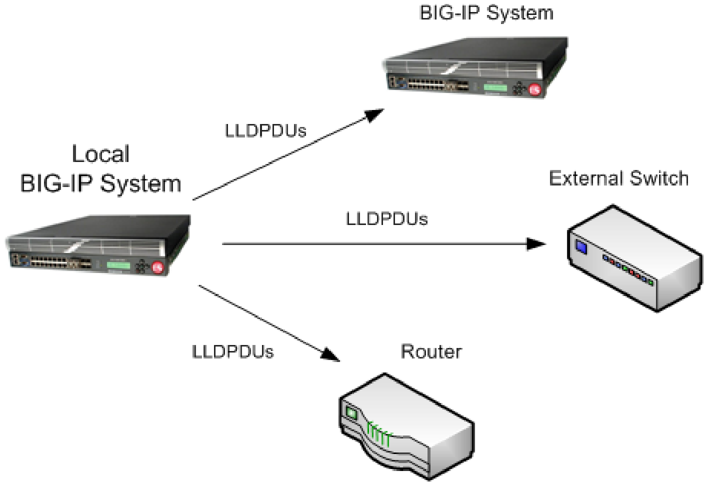

Figure 5: The BIG-IP system and LLDP transmittal

Task summary

Perform these tasks to implement Link Layer Discovery Protocol (LLDP) on selected BIG-IP system interfaces.

Task list

Configuring global LLDP properties

Configuring LLDP settings for an individual interface

Configuring global LLDP properties

You can configure a set of general LLDP properties that apply to all interfaces on the BIG-IP system. These settings pertain to LLDP message transmission frequencies and the maximum number of neighbors to which each interface can send LLDP messages.

***Note:** Although you use this procedure to globally enable the LLDP feature on the BIG-IP system, you can also disable LLDP for any individual interface. You do this by configuring the specific properties of that interface.*

1. On the Main tab, click **Network** > **Interfaces** > **LLDP** > **General**.
   This displays the general LLDP properties that you can configure on the system.
2. From the **LLDP** list, select **Enabled**.
3. For the remainder of the settings, retain or change the default values.
4. Click the **Update** button.

This task activates support for the LLDP protocol on the BIG-IP system, and configures the system to transmit LLDPDUs according to the specified frequencies.

Configuring LLDP settings for an individual interface

You can use this procedure to configure the settings for an individual interface on the BIG-IP system.

1. On the Main tab, click **Network** > **Interfaces** > **Interface List**.
   The Interface List screen displays the list of interfaces on the system.
2. In the Name column, click an interface number.
   This displays the properties of the interface.
3. For the **State** setting, verify that the interface is set to **Enabled**.
4. For the **LLDP** setting, verify that the property is set to **Transmit Only**.
5. For the **LLDP Attributes** setting, verify that the list of attributes in the **Send** field includes all Time Length Values (TLVs) that you want the BIG-IP system interface to send to neighbor devices.
6. Click the **Update** button.

After you perform this task, the interface is configured to send the specified LLDP information to neighbor devices.

Implementation result

This implementation results in this LLDP configuration:

- Support for the LLDP protocol is enabled on the BIG-IP system.
- For all BIG-IP system interfaces, the BIG-IP system attempts to transmit LLDPDUs to neighbor devices every 30 seconds, with a minimum delay between transmissions of 2 seconds.
- The maximum number of neighbors to which each BIG-IP system interface can send LLDPDUs is 10.
- Every BIG-IP system interface can send LLDPDUs to its neighbors.
- No BIG-IP system interface can receive LLDPDUs from its neighbors.

In addition, the content of the LLDPDUs that each BIG-IP system interface sends to its neighbors contains this information:

- Chassis ID
- Port ID
- Time-to-Live value
- Port description
- System name
- System description
- System capabilities
- Port VLAN ID
- Port and protocol VLAN ID
- VLAN name
- Protocol identity
- MAC/PHY config status
- Link aggregation
- Max frame size
- Product model

Using Link Aggregation with Tagged VLANs for a One-network Topology

Overview: Configuring link aggregation using tagged VLANs on one network

You can use the BIG-IP® system in an aggregated two-interface load balancing topology. *Link aggregation* is the process of combining multiple links so that the links function as a single link with higher bandwidth. Aggregating multiple interfaces into a trunk to create a link has the following advantages:

- Link aggregation increases the bandwidth of the individual network interface cards (NICs) in an additive manner.
- If one link goes down, the other link can handle the traffic by itself.

Link aggregation occurs when you create a trunk. A *trunk* is a combination of two or more interfaces and cables configured as one link.

The examples in this implementation show a trunk that includes two tagged interfaces aggregated together. A *tagged interface* is an interface that is configured to process traffic for multiple VLANs. A VLAN tag identifies the specific VLAN and enables traffic to pass through that specific VLAN. To cause traffic for multiple VLANs to be passed through a single trunk, you must assign the same trunk to each VLAN.

In the example, we create a trunk (**trunk1**) that includes two interfaces, **1.1** and **1.2**, and then assign **trunk1** as a tagged interface to both VLAN **external** and VLAN **internal**. Both VLANs (**external** and **internal**) reside on the same network, and are combined to form a VLAN group.

With this configuration, inbound and outbound traffic passing between the BIG-IP system and the vendor switch can use either interface. For example, traffic destined for VLAN **external**l can pass through either interface, **1.1** or **1.2**.

Illustration of link aggregation for a one-network topology

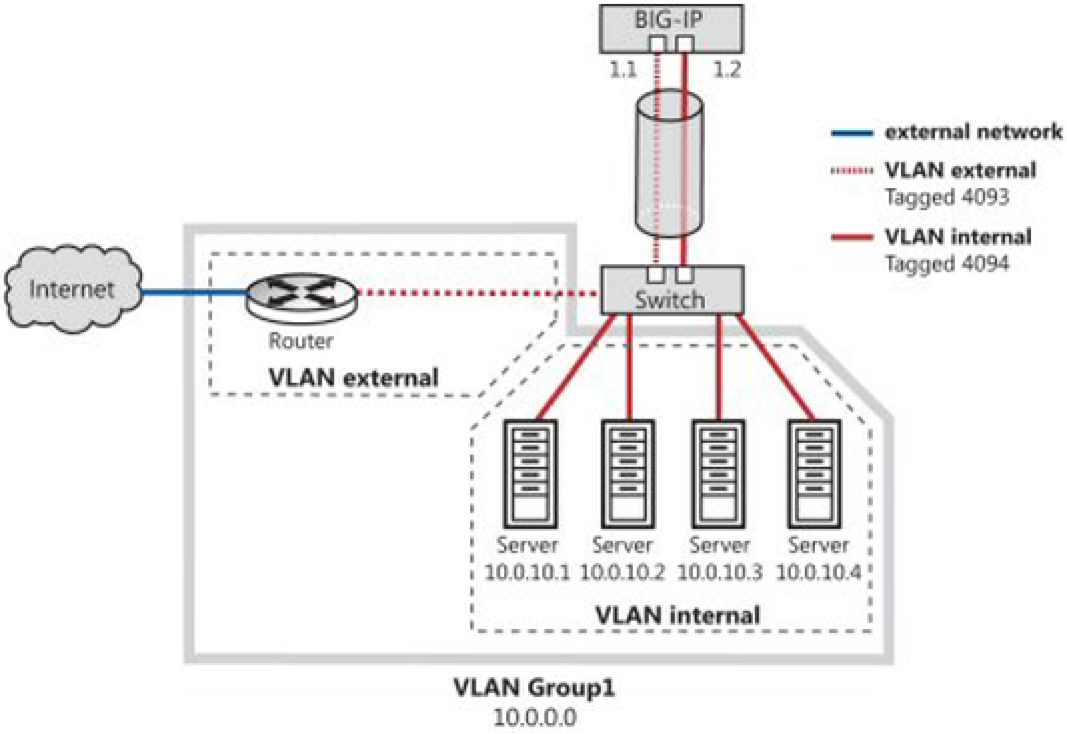

Figure 6: Link aggregation for a one-network topology

Task summary

Perform the following tasks to configure two interfaces (tagged VLANs) to function as a single link with higher bandwidth. In this implementation, you combine the two tagged VLANs into one VLAN group, where the two VLANs are on the same IP network.

Task list

Creating a trunk
Adding a tagged interface to a VLAN
Creating a load balancing pool
Creating a virtual server with source address affinity persistence
Removing the self IP addresses from the default VLANs
Creating a VLAN group
Creating a self IP for a VLAN group

Creating a trunk

You create a trunk on the BIG-IP® system so that the system can then aggregate the links to enhance bandwidth and ensure link availability.

1. On the Main tab, click **Network** > **Trunks**.
   The Trunk List screen opens.
2. Click **Create**.

3. Name the trunk.
4. For the **Interfaces** setting, in the **Available** field, select an interface, and using the Move button, move the interface to the **Members** field. Repeat this action for each interface that you want to include in the trunk.

   Trunk members must be untagged interfaces and cannot belong to another trunk. Therefore, only untagged interfaces that do not belong to another trunk appear in the **Available** list.
5. Select the **LACP** check box.
6. Click **Finished**.

After you create a trunk, the BIG-IP system aggregates the links to enhance bandwidth and prevent interruption in service.

Adding a tagged interface to a VLAN

After you aggregate the links, you assign the trunk to the VLAN as a tagged interface.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
   This displays the properties of the VLAN.
3. For the **Interfaces** setting:
   a) From the **Interface** list, select the trunk name.
   b) From the **Tagging** list, select **Tagged**.
   c) Click **Add**.

The trunk is assigned to the **external** and **internal** VLAN as a tagged interface.

Creating a load balancing pool

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   Tip: *Hold the Shift or Ctrl key to select more than one monitor at a time.*
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:

a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
b) In the **Address** field, type an IP address.
c) In the **Service Port** field, type a port number, or select a service name from the list.
d) (Optional) In the **Priority** field, type a priority number.
e) Click **Add**.

8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server with source address affinity persistence

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address you want to use for the virtual server.

   The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. Locate the relevant profile type for the traffic being managed, and either retain the default value or select a custom profile name.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. For the **Default Persistence Profile** setting, select **source_addr**.

   This implements simple persistence, using the default source address affinity profile.

A client system now has a destination IP address on the BIG-IP system.

Removing the self IP addresses from the default VLANs

Remove the self IP addresses from the individual VLANs. After you create the VLAN group, you will create another self IP address for the VLAN group for routing purposes. The individual VLANs no longer need their own self IP addresses.

1. On the Main tab, click **Network** > **Self IPs**.
2. Select the check box for each IP address and VLAN that you want to delete.
3. Click **Delete**.
4. Click **Delete**.

The self IP address is removed from the Self IP list.

Creating a VLAN group

Create a VLAN group that includes the internal and external VLANs. Packets received by a VLAN in the VLAN group are copied onto the other VLAN. This allows traffic to pass through the BIG-IP® system on the same IP network.

1. On the Main tab, click **Network** > **VLANs** > **VLAN Groups**.
   The VLAN Groups list screen opens.
2. Click **Create**.
   The New VLAN Group screen opens.

3. In the **Name** field, type the name `myvlangroup`.
4. For the **VLANs** setting, select `internal` and `external` VLAN names in the **Available** list and move them to the **Members** list.
5. Click **Finished**.

Creating a self IP for a VLAN group

Before you create a self IP address, ensure that you have created at least one VLAN group.

You perform this task to create a self IP address for a VLAN group. The self IP address for the VLAN group provides a route for packets destined for the network. With the BIG-IP® system, the path to an IP network is a VLAN. However, with the VLAN group feature used in this procedure, the path to the IP network `10.0.0.0` is actually through more than one VLAN. As IP routers are designed to have only one physical route to a network, a routing conflict can occur. With a self IP address on the BIG-IP system, you can resolve the routing conflict by associating a self IP address with the VLAN group.

1. On the Main tab, click **Network** > **Self IPs**.
2. Click **Create**.
   The New Self IP screen opens.
3. In the **IP Address** field, type an IPv4 or IPv6 address.

   This IP address should represent the address space of the VLAN group that you specify with the **VLAN/Tunnel** setting.
4. In the **Netmask** field, type the network mask for the specified IP address.

   For example, you can type `255.255.255.0`.
5. From the **VLAN/Tunnel** list, select the VLAN group with which to associate this self IP address.
6. From the **Port Lockdown** list, select **Allow Default**.
7. For the **Traffic Group** setting, choose one of the following actions:

Action	Result
Retain the default setting, traffic-group-local-only (non-floating).	The system creates a non-floating self IP address that becomes a member of `traffic-group-local-only`.
Select the check box labeled Inherit traffic group from current partition / path.	The system creates a floating self IP address that becomes a member of `traffic-group-1`.
Select a traffic group from the Traffic Group list.	The system creates a floating self IP address that becomes a member of the selected traffic group.

8. Click **Finished**.
   The screen refreshes, and displays the new self IP address.

The BIG-IP system can send and receive traffic through the specified VLAN or VLAN group.

Using Link Aggregation with Tagged VLANs for a Two-network Topology

Overview: Configuring link aggregation of two interfaces using tagged VLANs on two networks

You can use the BIG-IP® system in an aggregated two-interface load balancing topology. *Link aggregation* is the process of combining multiple links so that the links function as a single link with higher bandwidth. Aggregating multiple interfaces into a trunk to create a link has the following advantages:

- Link aggregation increases the bandwidth of the individual network interface cards (NICs) in an additive manner.
- If one link goes down, the other link can handle the traffic by itself.

Link aggregation occurs when you create a trunk. A *trunk* is a combination of two or more interfaces and cables configured as one link.

The examples in this implementation show a trunk that includes two tagged interfaces aggregated together. A *tagged interface* is an interface that is configured to process traffic for multiple VLANs. A VLAN tag identifies the specific VLAN and allows traffic to be passed through that specific VLAN. To cause traffic for multiple VLANs to be passed through a single trunk, you must assign the same trunk to each VLAN.

In the examples, we create a trunk (**trunk1**) that includes two interfaces, **1.1** and **1.2**, and then assign **trunk1** as a tagged interface to both VLAN **external** and VLAN **internal**. One network is connected to VLAN **external**, and a separate network is connected to VLAN **internal**. Consequently, inbound and outbound traffic passing between the BIG-IP system and the vendor switch can use either interface. For example, traffic destined for VLAN **external**l can pass through either interface, **1.1** or **1.2**.

Illustration of link aggregation for a two-network topology

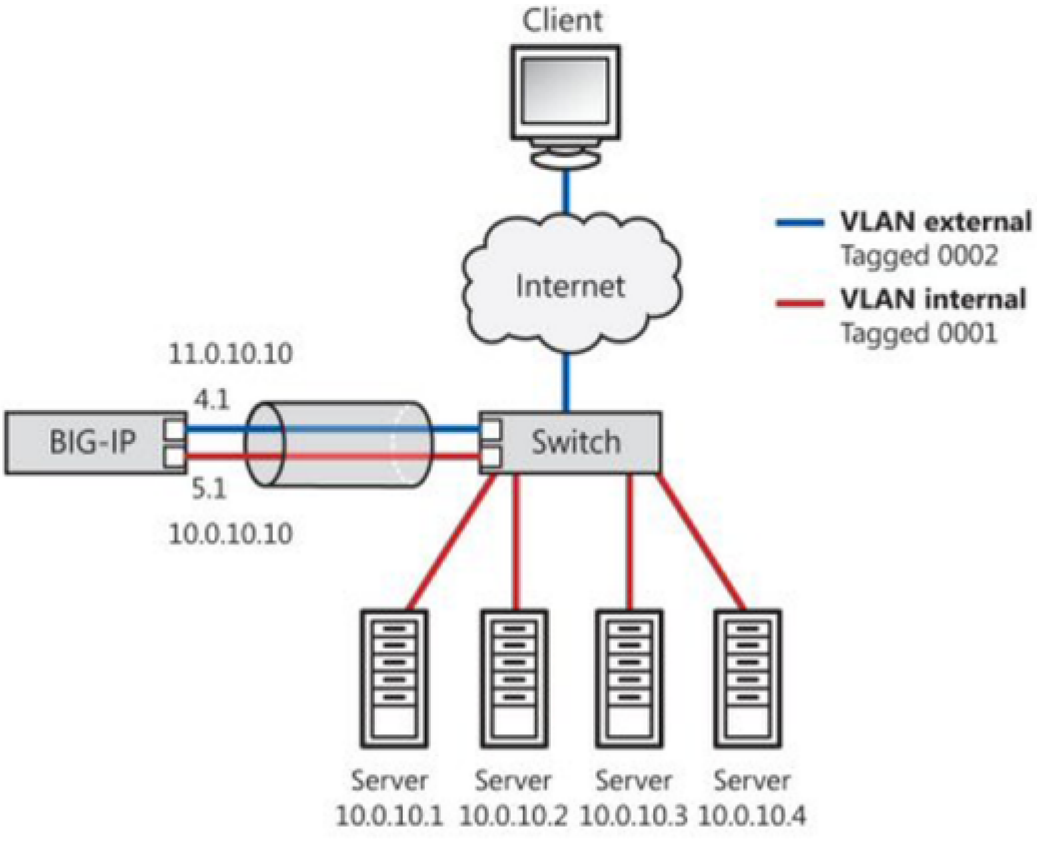

Figure 7: Link aggregation for a two-network topology

Task summary

Perform the following tasks to configure two interfaces (tagged VLANs) to function as a single link with higher bandwidth. In this implementation, each tagged VLAN is on a separate network.

Task list

Creating a trunk
Adding a tagged interface to a VLAN
Creating a load balancing pool
Creating a virtual server with source address affinity persistence

Creating a trunk

You create a trunk on the BIG-IP® system so that the system can then aggregate the links to enhance bandwidth and ensure link availability.

1. On the Main tab, click **Network** > **Trunks**.
   The Trunk List screen opens.
2. Click **Create**.

3. Name the trunk.
4. For the **Interfaces** setting, in the **Available** field, select an interface, and using the Move button, move the interface to the **Members** field. Repeat this action for each interface that you want to include in the trunk.

   Trunk members must be untagged interfaces and cannot belong to another trunk. Therefore, only untagged interfaces that do not belong to another trunk appear in the **Available** list.
5. Select the **LACP** check box.
6. Click **Finished**.

After you create a trunk, the BIG-IP system aggregates the links to enhance bandwidth and prevent interruption in service.

Adding a tagged interface to a VLAN

After you aggregate the links, you assign the trunk to the VLAN as a tagged interface.

1. On the Main tab, click **Network** > **VLANs**.
   The VLAN List screen opens.
2. In the Name column, click the relevant VLAN name.
   This displays the properties of the VLAN.
3. For the **Interfaces** setting:
   a) From the **Interface** list, select the trunk name.
   b) From the **Tagging** list, select **Tagged**.
   c) Click **Add**.

The trunk is assigned to the **external** and **internal** VLAN as a tagged interface.

Creating a load balancing pool

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

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

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, in the **Available** list, select a monitor type, and click **<<** to move the monitor to the **Active** list.

   Tip: *Hold the Shift or Ctrl key to select more than one monitor at a time.*
5. From the **Load Balancing Method** list, select how the system distributes traffic to members of this pool.

   The default is **Round Robin**.
6. For the **Priority Group Activation** setting, specify how to handle priority groups:
   - Select **Disabled** to disable priority groups. This is the default option.
   - Select **Less than**, and in the **Available Members** field type the minimum number of members that must remain available in each priority group in order for traffic to remain confined to that group.
7. Using the **New Members** setting, add each resource that you want to include in the pool:

a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
b) In the **Address** field, type an IP address.
c) In the **Service Port** field, type a port number, or select a service name from the list.
d) (Optional) In the **Priority** field, type a priority number.
e) Click **Add**.

8. Click **Finished**.

The load balancing pool appears in the Pools list.

Creating a virtual server with source address affinity persistence

A virtual server represents a destination IP address for application traffic.

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the **Create** button.
   The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, in the **Address** field, type the IP address you want to use for the virtual server.

   The IP address you type must be available and not in the loopback network.
5. In the **Service Port** field, type a port number or select a service name from the **Service Port** list.
6. Locate the relevant profile type for the traffic being managed, and either retain the default value or select a custom profile name.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. For the **Default Persistence Profile** setting, select **source_addr**.

   This implements simple persistence, using the default source address affinity profile.

A client system now has a destination IP address on the BIG-IP system.

Configuring Packet Filtering

Overview: Setting up packet filtering

Packet filters enhance network security by specifying whether a BIG-IP® system interface should accept or reject certain packets based on criteria that you specify. Packet filters enforce an access policy on incoming traffic. They apply to incoming traffic only.

You implement packet filtering by creating packet filter rules. The primary purpose of a packet filter rule is to define the criteria that you want the BIG-IP system to use when filtering packets. Examples of criteria that you can specify in a packet filter rule are:

- The source IP address of a packet
- The destination IP address of a packet
- The destination port of a packet

You specify the criteria for applying packet filter rules within an expression. When creating a packet filter rule, you can instruct the Configuration utility to build an expression for you, in which case you need only choose the criteria from predefined lists, or you can write your own expression text, using the syntax of the `tcpdump` utility.

Important: *Unlike most IP address configuration settings in the BIG-IP Configuration utility that require the* `%ID` *notation for route domains other than route domain* `0`*, the* ***Source Hosts and Networks*** *and* ***Destination Hosts and Networks*** *settings for packet filter rules accept IP addresses without the* `%ID` *route domain notation. This is because when you apply the packet filter rule to a VLAN, which belongs to a route domain, you are indirectly specifying which route domain's traffic to filter.*

Note: *Packet filter rules are unrelated to iRules®.*

You can also configure global packet filtering that applies to all packet filter rules that you create.

Task summary

By setting up some basic IP routing and configuring packet filtering, specific hosts on the internal VLAN can connect to the internal VLAN's self IP address. These hosts can also use common Internet services such as HTTP, HTTPS, DNS, FTP, and SSH. Traffic from all other hosts in the internal VLAN is rejected.

Task list

Enabling SNAT automap for internal and external VLANs
Creating a default gateway pool
Creating a forwarding virtual server
Enabling packet filtering
Creating a packet filter rule

Enabling SNAT automap for internal and external VLANs

You can configure SNAT automapping on the BIG-IP system for internal and external VLANs.

1. On the Main tab, click **Local Traffic** > **Address Translation**.

The **SNAT List** screen displays a list of existing SNATs.

2. Click **Create**.
3. Name the new SNAT.
4. From the **Translation** list, select **Automap**.
5. For the **VLAN / Tunnel List** setting, in the **Available** list, select **external** and **internal**, and using the Move button, transfer the VLANs to the **Selected** list.
6. Click the **Finished** button.

SNAT automapping on the BIG-IP system is configured for internal and external VLANs.

Creating a default gateway pool

Create a default gateway pool for the system to use to forward traffic.

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. For the **Health Monitors** setting, from the **Available** list, select the **gateway_icmp** monitor and move the monitor to the **Active** list.
5. Using the **New Members** setting, add each router that you want to include in the default gateway pool:
   a) Type the IP address of a router in the **Address** field.
   b) Type an asterisk (`*`) in the **Service Port** field, or select ***All Services** from the list.
   c) Click **Add**.
6. Click **Finished**.

Creating a forwarding virtual server

A virtual server represents a destination IP address for application traffic.

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

   The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `0.0.0.0/0`, and an IPv6 address/prefix is `::/0`.
5. From the **Service Port** list, select ***All Ports**.
6. In the Configuration area of the screen, from the **Type** list, select **Forwarding (IP)**.
7. From the **Protocol** list, select ***All Protocols**.
8. From the **VLAN/Tunnel Traffic** list, select **Enabled On**.
9. For the **VLAN List** setting, from the **Available** box, select **internal**, and click the Move button to move the VLAN name to the **Selected** box.
10.
11. In the Resources area of the screen, locate the **Default Pool** setting and select the pool you created previously.

12. Click **Finished**.

You now have a destination IP address on the BIG-IP system for application traffic.

Enabling packet filtering

Before creating a packet filtering rule, you must enable packet filtering. When you enable packet filtering, you can specify the MAC addresses, IP addresses, and VLANs that you want to be exempted from packet filter evaluation.

1. On the Main tab, click **Network** > **Packet Filters**.
   The Packet Filters screen opens.
2. From the **Packet Filtering** list, select **Enabled**.
3. From the **Unhandled Packet Action** list, select **Accept**.
4. For the **Options** setting, retain the default value or select the check boxes as needed.
5. For the **Protocols** setting, retain the default value or clear the check boxes as needed.
6. From the **MAC Addresses** list, specify a value:

Value	Description
None	When you select this value, all MAC addresses are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the MAC addresses that are exempt from packet filter evaluation, and the BIG-IP Configuration utility displays additional settings.

7. If you directed the **MAC Addresses** setting to always accept specific MAC addresses, provide the details:
   a) In the **Add** field, type a MAC address and click **Add**.
   The MAC address appears in the **MAC Address List** field.
   b) Repeat this step for each MAC address that you want the system to exempt from packet filter evaluation.
8. From the **IP Addresses** list, specify a value:

Value	Description
None	When you select this value, all IP addresses are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the IP addresses that are exempt from packet filter evaluation. The BIG-IP Configuration utility displays additional settings.

9. If you directed the **IP Addresses** setting to always accept specific IP addresses, provide the details:
   a) In the **Add** field, type an IP address and click **Add**.
   The IP address appears in the **IP Address List** field.
   b) Repeat this step for each IP address that you want the system to exempt from packet filter evaluation.
10. From the **VLANs** list, specify a value:

Value	Description
None	When you select this value, all VLANs are exempt from packet filter evaluation.
Always Accept	When you select this value, you can specify the VLANs that are exempt from packet filter evaluation. The BIG-IP Configuration utility displays additional settings.

11. If you configured the **VLANs** setting to always accept specific VLANs, then use the **Move** button to move one or more VLAN names from the **Available** list to the **Selected** list.
12. Click **Update**.

After you enable packet filtering, the BIG-IP® system filters packets according to the criteria in the packet filter rule and the values you configured when enabling the packet filter.

Creating a packet filter rule

When implementing packet filtering, you need to create a packet filter rule.

1. On the Main tab, click **Network** > **Packet Filters**.
   The Packet Filters screen opens.
2. Click **Rules**.
3. Click **Create**.
4. Name the rule.
5. From the **Order** list, select **First**.
6. From the **Action** list, select **Reject**.
7. From the **Rate Class** list, select a rate class if one exists on the system.

   You cannot use this setting if you have bandwidth control policy on the system.
8. From the **Bandwidth Controller** list, select a bandwidth controller policy if one exists on the system.

   You cannot use this setting if you have a rate class on the system.
9. From the **VLAN / Tunnel** list, select **internal**.
10. From the **Logging** list, select **Enabled**.
11. From the **Filter Expression Method** list, select **Enter Expression Text**.
12. In the **Filter Expression** field, choose a value:
    - **Enter Expression Text**. For example: `not dst port 80 and not dst port 443 and not dst port 53 and not dst port 22 and not dst port 20 and not dst port 21 and not dst host` *`internal_self_IP_address`*

      Note: *Replace* `internal_self_IP_address` *with the actual self IP address of VLAN internal.*
    - **Build Expression**. When you select this value, you can build an expression that causes the BIG-IP system to only accept certain protocols, source hosts and networks, destination hosts and networks, and destination ports.

      Important: *Unlike most IP address configuration settings in the BIG-IP Configuration utility that require the* `%ID` *notation for route domains other than route domain* `0`*, the* ***Source Hosts and Networks*** *and* ***Destination Hosts and Networks*** *settings for packet filter rules accept IP addresses without the* `%ID` *route domain notation. This is because when you apply the packet filter rule to a VLAN, which belongs to a route domain, you are indirectly specifying which route domain's traffic to filter.*
13. Click **Finished**.

The packet filter rule is now available for the BIG-IP system to use.

Referencing an External File from within an iRule

Overview: Referencing an external file from an iRule

Using the BIG-IP® Configuration utility or **tmsh**, you can import a file or URL from another system to the BIG-IP system, with content that you want an iRule to return to a client, based on some iRule event. Possible uses for this feature are:

- To send a web page other than the page that the client requested. For example, you might want the system to send a maintenance page instead of the requested page.
- To send an image.
- To use a file as a template and modify the file in the iRule before sending the file.
- To download policy information from an external server and merge that data with a locally-stored policy.

The file that an iRule accesses is known as an *iFile*, and can be any type of file, such as a binary file or a text file. These files are read-only files.

This example shows an iRule that references an iFile named `ifileURL`, in partition `Common`:

```
ltm rule ifile_rule {
   when HTTP_RESPONSE {
   # return a list of iFiles in all partitions
   set listifiles [ifile listall]
   log local0. "list of ifiles: $listifiles"

   # return the attributes of an iFile specified
   array set array_attributes [ifile attributes "/Common/ifileURL"]
   foreach {array attr} [array get array_attributes ] {
   log local0. "$array : $attr"
   }

   # serve an iFile when http status is 404.
   set file [ifile get "/Common/ifileURL"]
   log local0. "file: $file"
   if { [HTTP::status] equals "404" } {
     HTTP::respond 200 ifile "/Common/ifileURL"

  }
}
}
```

iRule commands for iFiles

This list shows the commands available for referencing an iFile within an iRule. All of these commands return a string, except for the command `[ifile attributes IFILENAME]`, which returns an array.

Available iRule commands for referencing an iFile

```
[ifile get IFILENAME]
[ifile listall]
[ifile attributes IFILENAME]
[ifile size IFILENAME]
[ifile last_updated_by IFILENAME]
[ifile last_update_time IFILENAME]
[ifile revision IFILENAME]
```

```
[ifile checksum IFILENAME]
[ifile attributes IFILENAME]
```

Task summary

You can import an existing file to the BIG-IP® system, create an iFile that is based on the imported file, and then write an iRule that returns the content of that file to a client system, based on an iRule event.

Task list

Importing a file for an iRule

Creating an iFile

Writing an iRule that references an iFile

Importing a file for an iRule

Before you perform this task, the file you want to import must reside on the system you specify.

You can import a file from another system onto the BIG-IP® system, as the first step in writing an iRule that references that file.

1. On the Main tab, click **System** > **File Management** > **iFile List** > **Import**.
2. For the **File Name** setting, click **Browse**.
   The system opens a browse window so that you can locate the file that you want to import to the BIG-IP system.
3. Browse for the file and click **Open**.
   The name of the file you select appears in the **File Name** setting.
4. In the **Name** field, type a new name for the file, such as `1k.html`.
   The new file name appears in the list of imported files.
5. Click the **Import** button.

After you perform this task, the file that you imported resides on the BIG-IP system.

Creating an iFile

As a prerequisite, ensure that the current administrative partition is set to the partition in which you want the iFile to reside. Also ensure that the file has been imported to the BIG-IP® system.

You perform this task to create an iFile that you can then reference in an iRule.

1. On the Main tab, click **Local Traffic** > **iRules** > **iFile List**.
2. Click **Create**.
3. In the **Name** field, type a new name for the iFile, such as `ifileURL`.
4. From the **File Name** list, select the name of the imported file object, such as `1k.html`.
5. Click **Finished**.
   The new iFile appears in the list of iFiles.

The result of this task is that you now have a file that an iRule can reference.

Writing an iRule that references an iFile

You perform this task to create an iRule that references an iFile.

***Note:** If the iFile resides in partition `/Common`, then specifying the partition when referencing the iFile is optional. If the iFile resides in a partition other than `/Common`, such as `/Partition_A`, you must include the partition name in the iFile path name within the iRule.*

1. On the Main tab, click **Local Traffic** > **iRules**.
   The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
   The New iRule screen opens.
3. In the **Name** field, type a name, such as `my_irule`.

   The full path name of the iRule cannot exceed 255 characters.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.

   For complete and detailed information iRules syntax, see the F5 Networks DevCentral web site (`http://devcentral.f5.com`).
5. Click **Finished**.
   The new iRule appears in the list of iRules on the system.

Implementation result

You now have an iRule that accesses a file on the BIG-IP®system, based on a particular iRule event.

Configuring Remote User Authentication and Authorization

Overview: Remote authentication and authorization of BIG-IP user accounts

The BIG-IP® system includes a comprehensive solution for managing BIG-IP administrative accounts on your network. With this solution, you can:

Use a remote server to store BIG-IP system user accounts.

The BIG-IP system includes support for using a remote authentication server to store BIG-IP system user accounts. After creating BIG-IP system accounts on the remote server (using the server vendor's instructions), you can configure the BIG-IP system to use remote user authentication and authorization (access control) for that server type.

Assign group-based access.

The BIG-IP system includes an optional feature known as *remote role groups*. With the *remote role groups* feature, you can use existing group definitions on the remote server to define the access control properties for users in a group. This feature not only provides more granularity in assigning user privileges, but also removes any need to duplicate remote user accounts on the BIG-IP system for the purpose of assigning those privileges.

Propagate a set of authorization data to multiple BIG-IP systems.

The BIG-IP system includes a tool for propagating BIG-IP system configuration data to multiple BIG-IP devices on the network. This tool is known as the Single Configuration File (SCF) feature.

Task summary

You can configure the BIG-IP® system to authorize user accounts that are stored on a remote authentication server.

***Important:** If you configure access control settings for group-based accounts (using the remote role groups feature), the BIG-IP system always applies those settings, rather than the default access control settings, to group-based accounts.*

The BIG-IP® system supports several types of authentication servers for storing BIG-IP system administrative user accounts. The actual procedure you use to specify the type of remote server differs, depending on the server type.

Task list

Specifying LDAP or Active Directory server information

Specifying client certificate LDAP server information

Specifying RADIUS server information

Specifying TACACS+ server information

Configuring access control for remote user groups

Saving access control settings to a file

Importing BIG-IP configuration data onto other BIG-IP systems

Specifying LDAP or Active Directory server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.
- If you want to verify the certificate of the authentication server, import one or more SSL certificates.

You can configure the BIG-IP system to use an LDAP or Microsoft® Windows® Active Directory ®server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

***Important:** The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based access control. These values represent the default values that the BIG-IP system applies to any user account that is not part of a remotely-stored user group. Also, for the `Other External Users` user account, you can modify the **Role**, **Partition Access**, and **Terminal Access** settings only when your current partition on the BIG-IP system is set to `Common`. If you attempt to modify these settings when your current partition is other than `Common`, the system displays an error message.*

1. On the Main tab, click **System** > **Users** > **Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - LDAP** or **Remote - Active Directory**.
5. In the **Host** field, type the IP address of the remote server.

   The route domain to which this address pertains must be route domain `0`.
6. For the **Port** setting, retain the default port number (`389`) or type a new port number.

   This number represents the port number that the BIG-IP system uses to access the remote server.
7. In the **Remote Directory Tree** field, type the file location (tree) of the user authentication database on the LDAP or Active Directory server.

   At minimum, you must specify a domain component (that is, `dc=[value]`).
8. For the **Scope** setting, retain the default value (`Sub`) or select a new value.

   This setting specifies the level of the remote server database that the BIG-IP system should search for user authentication.
9. For the **Bind** setting, specify a user ID login for the remote server:
   a) In the **DN** field, type the distinguished name for the remote user ID.
   b) In the **Password** field, type the password for the remote user ID.
   c) In the **Confirm** field, re-type the password that you typed in the **Password** field.
10. In the **User Template** field, type a string that contains a variable representing the distinguished name of the user, in the format `%s`.

    This field can contain only one `%s` and cannot contain any other format specifiers.

    For example, you can specify a user template such as `%s@siterequest.com` or `uxml:id=%s,ou=people,dc=siterequest,dc=com`.
    The result is that when a user attempts to log on, the system replaces `%s` with the user name specified in the Basic Authentication dialog box, and passes that name as the distinguished name for the bind operation. The system also passes the associated password as the password for the bind operation.
11. For the **Check Member Attribute in Group** setting, select the check box if you want the system to check the user's member attribute in the remote LDAP or AD group.
12. To enable SSL-based authentication, from the **SSL** list select **Enabled** and, if necessary, configure these settings:

a) From the **SSL CA Certificate** list, select the name of a chain certificate, that is, the third-party CA or self-signed certificate that normally resides on the remote authentication server.

b) From the **SSL Client Key** list, select the name of the client SSL key.

Use this setting only when the remote server requires that the client present a certificate.

c) From the **SSL Client Certificate** list, select the name of the client SSL certificate.

Use this setting only if the remote server requires that the client present a certificate.

13. In the **Login LDAP Attribute** field, type the account name for the LDAP server.

    The value for this option is normally the user ID. However, if the server is a Microsoft® Windows® Active Directory®server, the value must be the account name `sAMAccountName` (case-sensitive). The default value is none.

14. From the **Client Certificate Name Field** list:

    a) Select either a subject alternate name or the subject name (**Common Name**).

    b) If you select the subject alternate name **Other Name**, then in the **OID** field, type an object identifier (OID).

    The OID indicates the format and semantics of the subject alternate name.

15. For the **Fallback to Local** setting, select the check box when you want to allow configuring remote authentication to fall back to the local authentication when the remote server is unavailable.
16. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
17. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
18. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only `tmsh` access to the BIG-IP system.

19. Click **Finished**.

You can now authenticate administrative user accounts that are stored on a remote LDAP or Active Directory server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Specifying client certificate LDAP server information

Verify that the required user accounts for the BIG-IP® system exist on the remote authentication server.

For authenticating BIG-IP system user accounts (that is, traffic that passes through the management interface [MGMT]), you can configure the BIG-IP system to authenticate certificates issued by a certificate authority's Online Certificate Status Protocol (OCSP) responder.

***Important:** The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based authorization. These values represent the default values or locally configured user accounts (which override the default role) that the BIG-IP system applies to any user account that is not part of a remote role group.*

1. On the Main tab, click **System** > **File Management** > **Apache Certificate List** > **Import**, browse for the certificate file to import, type a name, and click **Import**.
   The certificate will be added to the Apache Certificate list.
2. On the Main tab, click **System** > **Users** > **Authentication**.

3. On the menu bar, click **Authentication**.
4. Click **Change**.
5. From the **User Directory** list, select **Remote - ClientCert LDAP**.
6. In the **Host** field, type the IP address of the remote server.

   The route domain to which this address pertains must be route domain `0`.
7. For the **Port** setting, retain the default port number (`389`) or type a new port number.

   This number represents the port number that the BIG-IP system uses to access the remote server.
8. In the **Remote Directory Tree** field, type the file location (tree) of the user authentication database on the client certificate server.

   At minimum, you must specify a domain component (that is, `dc=[value]`).
9. For the **Scope** setting, retain the default value (`Sub`) or select a new value.

   This setting specifies the level of the remote server database that the BIG-IP system should search for user authentication.
10. For the **Bind** setting, specify a user ID login for the remote server:
    a) In the **DN** field, type the distinguished name for the remote user ID.
    b) In the **Password** field, type the password for the remote user ID.
    c) In the **Confirm** field, re-type the password that you typed in the **Password** field.
11. To enable SSL-based authentication, from the **SSL** list select **Enabled** and, if necessary, configure these settings:
    a) From the **SSL CA Certificate** list, select the name of a chain certificate; that is, the third-party CA or self-signed certificate that normally resides on the remote authentication server.
    b) From the **SSL Client Key** list, select the name of the client SSL key.

       Use this setting only when the remote server requires that the client present a certificate.
    c) From the **SSL Client Certificate** list, select the name of the client SSL certificate.

       Use this setting only if the remote server requires that the client present a certificate.
12. In the **CA Certificate** field, type the absolute folder path of `apache-ssl-cert fileobject` for the CA signing authority.

    The absolute folder path is `/Common/<folder path>/<certificate name>`. To determine the absolute folder path of the `apache-ssl-cert fileobject`, click **System** > **File Management** > **Apache Certificate List** and note the target certificate's partition and path.

    Important: *Apache certificates can only be stored within* `/Common`.
13. In the **Login Name** field, type an LDAP search prefix that will contain the distinguished name (DN) from the user certificate, such as `CN`.

    This specifies the LDAP attribute to be used as a login name. The default is disabled.
14. In the **Login LDAP Attribute** field, type the account name for the LDAP server.

    The value for this option is normally the user ID. However, if the server is a Microsoft® Windows® Active Directory®server, the value must be the account name `sAMAccountName` (case-sensitive). The default value is none.
15. In the **Login Filter** field, type the LDAP attribute that contains the short name of the user.

    This specifies the filter to be applied on the common name (CN) of the client certificate and usually this is the user ID or `sAMAccountName`. The filter is a regular expression used to extract required information from the CN of the client certificate that is matched against the LDAP search results. The default is disabled.
16. For the **Depth** setting, retain the default value (`10`) or type a new value for verification depth.
17. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.

18. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
19. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only `tmsh` access to the BIG-IP system.

20. Click **Finished**.

You can now authenticate administrative traffic for user accounts that are stored on a remote client certificate server. If you have no need to configure group-based user authorization, your configuration tasks are complete.

Specifying RADIUS server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.

You can configure the BIG-IP system to use a RADIUS server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

Important: *The values you specify in this procedure for the* ***Role****,* ***Partition Access****, and* ***Terminal Access*** *settings do not apply to group-based authorization. These values represent the default values that the BIG-IP system applies to any user account that is not part of a role group that is defined on the remote authentication server. Also, for the* `Other External Users` *user account, you can modify the* ***Role****,* ***Partition Access****, and* ***Terminal Access*** *settings only when your current partition on the BIG-IP system is set to* `Common`*. If you attempt to modify these settings when your current partition is other than* `Common`*, the system displays an error message.*

1. On the Main tab, click **System** > **Users** > **Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - RADIUS**.
5. For the **Primary** setting:
   a) In the **Host** field, type the name of the primary RADIUS server.

      The route domain with which this host is associated must be route domain `0`.

   b) In the **Secret** field, type the password for access to the primary RADIUS server.
   c) In the **Confirm** field, re-type the RADIUS secret.
6. If you set the **Server Configuration** setting to **Primary and Secondary**, then for the **Secondary** setting:
   a) In the **Host** field, type the name of the secondary RADIUS server.

      The route domain with which this host is associated must be route domain `0`.

   b) In the **Secret** field, type the password for access to the secondary RADIUS server.
   c) In the **Confirm** field, re-type the RADIUS secret.
7. For the **Fallback to Local** setting, select the check box when you want to allow configuring remote authentication to fall back to the local authentication when the remote server is unavailable.

8. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
9. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
10. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only `tmsh` access to the BIG-IP system.

11. Click **Finished**.

You can now authenticate administrative traffic for BIG-IP system user accounts that are stored on a remote RADIUS server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Specifying TACACS+ server information

Before you begin:

- Verify that the BIG-IP® system user accounts have been created on the remote authentication server.
- Verify that the appropriate user groups, if any, are defined on the remote authentication server.

You can configure the BIG-IP system to use a TACACS+ server for authenticating BIG-IP system user accounts, that is, traffic that passes through the management interface (MGMT).

Important: *The values you specify in this procedure for the **Role**, **Partition Access**, and **Terminal Access** settings do not apply to group-based authorization. These values represent the default values that the BIG-IP system applies to any user account that is not part of a remote role group. Also, for the `Other External Users` user account, you can modify the **Role**, **Partition Access**, and **Terminal Access** settings only when your current partition on the BIG-IP system is set to `Common`. If you attempt to modify these settings when your current partition is other than `Common`, the system displays an error message.*

1. On the Main tab, click **System** > **Users** > **Authentication**.
2. On the menu bar, click **Authentication**.
3. Click **Change**.
4. From the **User Directory** list, select **Remote - TACACS+**.
5. For the **Fallback to Local** setting, select the check box when you want to allow configuring remote authentication to fall back to the local authentication when the remote server is unavailable.
6. For the **Servers** setting, type an IP address for the remote TACACS+ server.

   The route domain to which this address pertains must be route domain `0`.
7. Click **Add**.
   The IP address for the remote TACACS+ server appears in the **Servers** list.
8. In the **Secret** field, type the password for access to the TACACS+ server.

   Warning: *Do not include the symbol `#` in the secret. Doing so causes authentication of local user accounts (such as `root` and `admin`) to fail.*
9. In the **Confirm Secret** field, re-type the TACACS+ secret.
10. From the **Encryption** list, select an encryption option:

Option	Description
Enabled	Specifies that the system encrypts the TACACS+ packets.
Disabled	Specifies that the system sends unencrypted TACACS+ packets.

11. In the **Service Name** field, type the name of the service that the user is requesting to be authenticated to use (usually `ppp`).

    Specifying the service causes the TACACS+ server to behave differently for different types of authentication requests. Examples of service names that you can specify are: `ppp`, `slip`, `arap`, `shell`, `tty-daemon`, `connection`, `system`, and `firewall`.
12. In the **Protocol Name** field, type the name of the protocol associated with the value specified in the **Service Name** field.

    This value is usually `ip`. Examples of protocol names that you can specify are: `ip`, `lcp`, `ipx`, `atalk`, `vines`, `lat`, `xremote`, `tn3270`, `telnet`, `rlogin`, `pad`, `vpdn`, `ftp`, `http`, `deccp`, `osicp`, and `unknown`.
13. From the **Role** list, select the user role that you want the BIG-IP system to assign by default to all BIG-IP system user accounts authenticated on the remote server.
14. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
15. From the **Terminal Access** list, select either of these as the default terminal access option for remotely-authenticated user accounts:

Option	Description
Disabled	Choose this option when you do not want the remotely-stored user accounts to have terminal access to the BIG-IP system.
tmsh	Choose this option when you want the remotely-stored user accounts to have only `tmsh` access to the BIG-IP system.

16. Click **Finished**.

You can now authenticate administrative traffic for BIG-IP system user accounts that are stored on a remote TACACS+ server. If you have no need to configure access control for remotely-stored user groups, your configuration tasks are complete.

Configuring access control for remote user groups

You perform this task to assign a user role, a corresponding administrative partition, and a type of terminal access to a remotely-stored group of user accounts. For a given user group, you can assign as many role-partition combinations as you need, as long as each role is associated with a different partition. If the partition you associate with a role is `All`, this entry might or might not take effect, depending on whether the `All` designation conflicts with other role-partition combinations for that user group. For any conflicts, line order in the configuration is a consideration. To assign multiple role-partition combinations for a user group, you repeat this task for each combination, specifying the same attribute string for each task.

1. On the Main tab, click **System** > **Users**.
2. On the menu bar, click **Remote Role Groups**.
3. Click **Create**.
4. In the **Group Name** field, type the group name that is defined on the remote authentication server. An example of a group name is **BigIPOperatorsGroup**.
5. In the **Line Order** field, type a number.

    This value specifies the order of this access control configuration in the file `/config/bigip/auth/remoterole` for the named group. The LDAP and Active Directory servers read this file line by line.

The order of the information is important; therefore, F5 Networks recommends that you specify a value of `1000` for the first line number. This allows you, in the future, to insert lines before the first line.

6. In the **Attribute String** field, type an attribute.

   An example of an attribute string is `memberOF=cn=BigIPOperatorsGroup,cn=users,dc=dev,dc=net`.

   The BIG-IP system attempts to match this attribute with an attribute on the remote authentication server. On finding a match, the BIG-IP system applies the access control settings defined here to the users in that group. If a match is not found, the system applies the default access control settings to all remotely-stored user accounts (excluding any user account for which access control settings are individually configured).
7. From the **Remote Access** list, select a value.

Option	Description
Enabled	Choose this value if you want to enable remote access for the defined user group.
Disabled	Choose this value if you want to disable remote access for the defined user group. Note that if you configure multiple instances of this remote role group (one instance for each role-partition pair for the attribute string), then choosing a value of **Disabled** disables remote access for all user group members, regardless of the remote role group instance.

8. From the **Assigned Role** list, select a user role for the remote user group.
9. From the **Partition Access** list, select an administrative partition value.

Option	Description
All	Choose this value to give users in the defined group access to their authorized objects in all partitions on the BIG-IP system.
`partition_name`	Choose a specific partition name to give users in the defined group access to that partition only.
Common	Choose this value to give users in the defined group access to partition **Common** only.

10. From the **Terminal Access** list, select the type of command-line access you want to grant users in the group, if any.
11. Click **Finished** or **Repeat**.

After you perform this task, the user group that you specified has the assigned role, partition access, and terminal access properties assigned to it.

Saving access control settings to a file

You can save the running configuration of the system, including all settings for remote user authentication and authorization, in a flat, text file with a specified name and the extension `.scf`.

1. On the BIG-IP® system, access a command-line prompt.
2. At the prompt, open the Traffic Management Shell by typing the command `tmsh`.
3. Type `sys save` *`filename`*.
   `sys save myConfiguration053107` creates the file `myConfiguration053107.scf` in the `var/local/scf` directory.
   `sys save /config/myConfiguration` creates the file `myConfiguration.scf` in the `/config` directory.

You can now import this file onto other BIG-IP devices on the network.

Importing BIG-IP configuration data onto other BIG-IP systems

You can use the `tmsh sys load` command to import a single configuration file (SCF), including access control data, onto other BIG-IP® devices on the network.

***Note:** This task is optional.*

1. On the BIG-IP system on which you created the SCF, access a command-line prompt.
2. Copy the SCF that you previously created to a location on your network that you can access from the system that you want to configure.
3. Edit the SCF to reflect the management routing and special passwords of the BIG-IP system that you want to configure:
   a) Open the SCF in an editor.
   b) Where necessary, change the values of the management IP address, network mask, management default route, self IP addresses, virtual server IP addresses, routes, default routes, and host name fields to the values for the new system.
   c) If necessary, change the passwords for the `root` and `admin` accounts using the command `user` *`name`* `password none newpassword` *`password`*.

   ***Important:** When configuring a unit that is part of a redundant system configuration and that is using the SCF from the peer unit, do not modify the `root` and `admin` accounts. These accounts must be identical on both units of the redundant system.*

   d) Save the edited SCF.
4. On the BIG-IP system that you want to configure, open the Traffic Management Shell by typing the command `tmsh`.
5. Type `sys load` *`scf_filename`*.
   `sys load myConfiguration053107.scf` saves a backup of the running configuration in the `/var/local/scf` directory, and then resets the running configuration with the configuration contained in the SCF you are loading.

Configuring Administrative Partitions to Control User Access

Overview: Administrative partitions for user access control

The BIG-IP® system includes a powerful authorization feature known as administrative partitions. Using the *administrative partitions* feature, you ensure that BIG-IP system grants administrative users exactly the right type and amount of access to BIG-IP system resources. As a result, you can tailor user access to resources to exactly fit the needs of your organization.

Task summary

There are two main tasks for controlling user access to BIG-IP® system objects.

Task list

Creating an administrative partition
Assigning roles to a user account

Creating an administrative partition

You perform this task to create an administrative partition. An *administrative partition* creates an access control boundary for users and applications.

1. On the Main tab, expand **System** and click **Users**.
   The Users List screen opens.
2. On the menu bar, click **Partition List**.
3. Click **Create**.
   The New Partition screen opens.
4. In the **Partition Name** field, type a unique name for the partition.
   An example of a partition name is `Spanned_VIP`.
5. Type a description of the partition in the **Description** field.

   This field is optional.
6. For the **Device Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the device group attribute from folder `root`.
Clear the check box and select the name of a device group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the device group attribute from folder `root`.

7. For the **Traffic Group** setting, choose an action:

Action	Result
Retain the default value.	Choose this option if you want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder `root`.

Action	Result
Clear the check box and select the name of a traffic group.	Choose this option if you do not want the folder corresponding to this partition to inherit the value of the traffic group attribute from folder `root`.

8. Click **Finished**.

The new partition appears in the partition list.

Assigning roles to a user account

Before performing this task, ensure that you have a user role of Administrator or that you have a role of User Manager for the relevant partition.

You perform this task to change the user roles that are assigned to a user account. You can assign a different role for each partition to which the user has access. By default, the user role that the BIG-IP® system assigns to a user account on each partition is No Access.

Important: *If you are performing this task while the user is logged into the system through* `tmsh`*, the BIG-IP system terminates the user's* `tmsh` *session when the user subsequently issues another* `tmsh` *command. This behavior ensures that the user is notified of the change in permissions and that data integrity is maintained.*

1. Access the BIG-IP ® Configuration utility.
2. In the upper-left corner of the screen, confirm that the **Partition** list is set to the partition in which the user account that you want to modify resides.
3. On the Main tab, click **System** > **Users**.

   The BIG-IP system displays the list of user accounts that reside in the current partition and in partition `Common`. Note that all users except those with a user role of No Access have at least read access to partition `Common`.
4. In the User Name column, click the user account name.
5. For the **Partition Access** setting:
   a) From the **Role** list to select a user role.
   b) From the **Partition** list, select a partition name.
   c) Click the **Add** button.
      A user role pertaining to a partition now appears in the box.
   d) Repeat these steps for each partition to which you want to assign a role for this user.

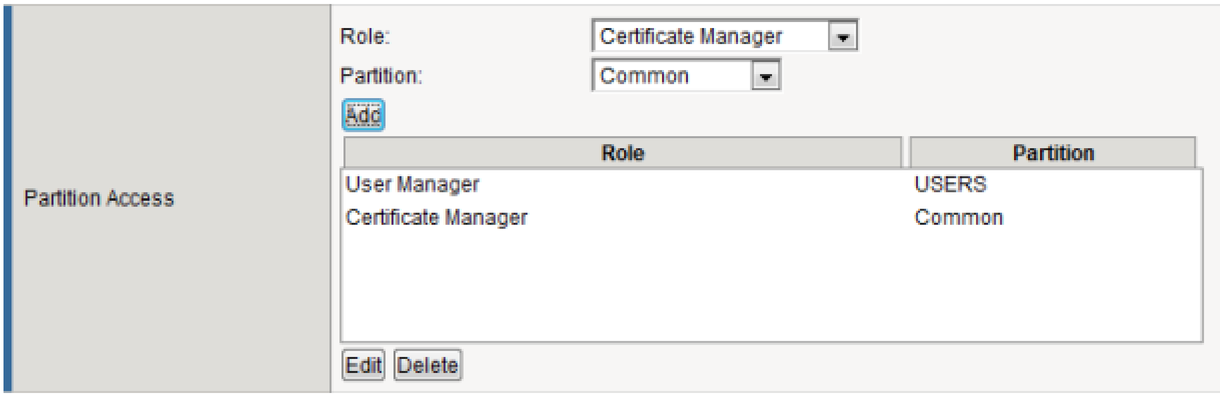

Role: Certificate Manager

Partition: Common

Role	Partition
User Manager	USERS
Certificate Manager	Common

Figure 8: Granting partition access to a BIG-IP user account

After you configure this setting, one or more role-partition combinations are specified for assignment to this user account.

6. Click the **Update** button.

Working with Single Configuration Files

Overview: Working with single configuration files

A *single configuration file (SCF)* is a flat, text file that contains a series of `tmsh` commands, and the attributes and values of those commands, that reflect the configuration of the BIG-IP® system. Specifically, the SCF contains the local traffic management and TMOS® configuration of the BIG-IP system. This figure shows a small part of a sample SCF.

```
  vlan external {
    tag 4093
    interfaces 1.3
}
vlan internal {
    tag 4094
    interfaces 1.10
}
pool dev_https3 {
    members {
        10.60.10.105:https{}
        10.60.10.106:https{}
    }
}
```

The single configuration file feature allows you to save the configuration of a BIG-IP system in a text file. You can then use the text file to easily replicate the configuration across multiple BIG-IP systems. This not only saves you time, but also allows you to create a consistent, secure, comprehensive local traffic management environment on your network.

tmsh commands for single configuration files (SCFs)

You use `tmsh` to manage a single configuration file (SCF). This table lists an overview of `tmsh` commands used to manage SCF files.

tmsh command	Description
`save sys config file [filename]`	Saves a copy of the currently running configuration to an SCF. ***Important:*** *Saving a configuration to an SCF does not affect the running or stored configuration of the BIG-IP® system on which you run the command.*
`load sys config file [filename]`	Replaces or restores an SCF with a saved configuration. When you use this command, the system saves any previously running configuration to the `/var/local/scf/` directory, by default.
`load sys config default`	Restores the factory default settings of the configuration file, while retaining the management IP address and the administrator user name and password.

Task summary

You can perform three main tasks with respect to single configuration files.

Task list

Creating and saving an SCF

Loading an SCF onto a target BIG-IP system

Using an SCF to restore a BIG-IP system configuration

Creating and saving an SCF

You can use `tmsh` to create and save a single configuration file (SCF).

Important: *The system configuration data contained in the text file includes any local device certificate and keys used to establish device trust between this system and the other devices in a BIG-IP device group. These certificates and keys are unencrypted in the text file and are not included in the* `.tar` *file.*

Note: *If you create an SCF file twice (on two different occasions), you can compare the contents of the two files.*

1. Open the TMOS Shell (`tmsh`).
   ```
   tmsh
   ```
2. Create and save an SCF.
   ```
   save sys config file [filename]
   ```
   Note: *If you include the* `.scf` *extension in the file name, the system does not add an additional file extension.*

   The system gathers all of the commands that make up the running configuration, and then saves the configuration to a `.scf` file with the name you specify. The system also creates a corresponding `.tar` file. By default, the system stores these files in the `/var/local/scf` directory, but you can specify a different path if you prefer.

Loading an SCF onto a target BIG-IP system

You can use `tmsh` to load a single configuration file (SCF) on one BIG-IP® system that you created on another BIG-IP system (hereafter referred to as the target BIG-IP system). This saves you from having to recreate the configuration multiple times. Loading an SCF resets the running configuration with the values contained in the stored configuration.

Important: *If you run a* `load` *command or restart the system before you save your changes to the stored configuration, you will lose any changes.*

Note: *To successfully load a configuration that you have replicated, make sure that no line of the configuration is longer than 4096 characters. If there are more than 4096 characters in a single line, the system reverts to the previous running configuration.*

1. Open the TMOS Shell (`tmsh`).
   ```
   tmsh
   ```
2. On the target BIG-IP system, load the saved SCF file.

```
tmsh load sys config file [filename]
```

The system saves the stored configuration to a backup file named `/var/local/scf/backup.scf`, and then uses the configuration stored in the SCF that you are loading.

3. Use a text editor to open the SCF and edit any data that is unique to the target BIG-IP system, such as the management IP address.
4. Save the SCF to the target BIG-IP system.

   ```
   sys save config file [filename]
   ```

   If a backup SCF already exists, the system appends a number to the name of the existing backup file, and then creates a new backup file. In the case of multiple load and save operations:

   - The first time the system backs up the running configuration during a load operation, the system names the backup file `/var/local/scf/backup.scf`.
   - The next time the system backs up the running configuration, the system renames the file from `/var/local/scf/backup.scf` to `/var/local/scf/backup-1.scf` and creates a new file named `/var/local/scf/backup.scf`.
   - If you run the `load` command a third time, the system renames the file from `/var/local/scf/backup-1.scf` to `/var/local/scf/backup-2.scf`, renames the `/var/local/scf/backup.scf` file to `/var/local/scf/backup-1.scf`, and again creates a new file named `/var/local/scf/backup.scf`.

Using an SCF to restore a BIG-IP system configuration

You can use `tmsh` to restore a BIG-IP® system configuration using either a specific single configuration file (SCF) or the factory default configuration.

1. Open the TMOS Shell (`tmsh`).

   ```
   tmsh
   ```

2. Restore the system configuration using one of these options:

   - Restore a system to the factory default configuration by using `tmsh load sys config default`. This command retains the management IP and the assigned root and administrator passwords. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the local traffic management and the operating system configuration to the factory default settings by loading the factory default SCF (`/defaults/defaults.scf`).
   - Restore a system with values defined in the specified SCF by using `tmsh load sys config file [filename]`. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the running configuration to the values contained in the specified SCF.

   Note: *You must run the* `save sys config partitions all` *command to save the running configuration in the stored configuration files.*

Legal Notices

Legal notices

Publication Date

This document was published on March 2, 2017.

Publication Number

MAN-0379-11

Copyright

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

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