BIG-IP® Service Provider: Administration

Version 13.0

Table of Contents

Implementing SCTP Multistreaming and Multihoming

Overview: SCTP multistreaming

Unlike Transmission Control Protocol (TCP), Stream Control Transmission Protocol (SCTP) includes the ability to support *multistreaming functionality*, which permits several streams within an SCTP connection. While a *TCP stream* refers to a sequence of bytes, an *SCTP stream* represents a sequence of data messages. Each data message (or chunk) contains an integer ID that identifies a stream, an application-defined Payload Protocol Identifier (PPI), a Stream sequence number, and a Transmit Serial Number (TSN) that uniquely identifies the chunk within the SCTP connection. Chunk delivery is acknowledged using TSNs sent in selective acknowledgements (ACKs) so that every chunk can be independently acknowledged. This capability demonstrates a significant benefit of streams, because it eliminates head-of-line blocking within the connection. A lost chunk of data on one stream does not prevent other streams from progressing while that lost chunk is retransmitted.

Task list

Creating an SCTP profile for multistreaming
Configuring an SCTP virtual server

Creating an SCTP profile for multistreaming

You can enable and configure an SCTP profile for multistreaming functionality, which permits several streams within an SCTP connection.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Protocol** > **SCTP**.
   The SCTP profile list screen opens.
2. Click **Create**.
   The New SCTP Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select the **Custom** check box.
   The settings become available for change.
5. In the **Out Streams** field, type a value for the number of outbound streams.

   Important: *Ensure that this value equals the value requested by the servers when the server-side connection is established.*

   Note: *A value of 2, or greater, enables SCTP multistreaming functionality.*
6. In the **In Streams** field, type a value for the number of inbound streams.

   Important: *Ensure that this value equals the value requested by the servers when the server-side connection is established.*

   Note: *A value of 2, or greater, enables SCTP multistreaming functionality.*
7. Click **Finished**.

An SCTP profile is configured for multistreaming functionality, permitting several streams within an SCTP connection.

Configuring an SCTP virtual server

You must create an SCTP profile before you can support SCTP on a virtual server.

You can use SCTP with multistreaming as a transport for a 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 for the SCTP client 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:** This destination address supports the initial SCTP control connection, providing the initial handshake and transfer of valid destination addresses.*
5. From the **Protocol** list, select **SCTP**.
6. From the **Protocol Profile (Client)** list, select a predefined or user-defined SCTP profile.
7. Click **Finished**.

The new virtual server supports SCTP with multistreaming.

Overview: SCTP multihoming

Stream Control Transmission Protocol (SCTP) includes the ability to support *multihoming functionality*, which provides path redundancy for an SCTP connection by enabling SCTP to send packets between multiple addresses owned by each endpoint. SCTP endpoints typically configure different IP addresses on different network interfaces to provide redundant physical paths between the peers. For example, a client and server might be attached to separate VLANs. The client and server can each advertise two IP addresses (one per VLAN) to the other peer. If either VLAN is available, then SCTP can transport packets between the peers.

Task list

Creating an SCTP profile for multihoming
Configuring IP addresses for multihoming connections
Creating a SNAT pool for SCTP
Creating an SCTP pool for multihoming
Configuring an SCTP virtual server for multihoming

Creating an SCTP profile for multihoming

You can enable and configure an SCTP profile for multihoming functionality, which provides path redundancy for an SCTP connection by enabling SCTP to send packets between multiple addresses owned by each endpoint.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Protocol** > **SCTP**.
   The SCTP profile list screen opens.

2. Click **Create**.
   The New SCTP Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select the **Custom** check box.
   The settings become available for change.
5. Configure the client-side multihoming settings.
   a) From the **Client Side Multi-homing** list, select **Enabled** to enable SCTP multihoming for clients.

      When enabled, this setting enables SCTP clients to connect to a virtual server over multiple IP interfaces.

      The **Secondary Addresses** setting appears.
   b) For the **Secondary Addresses** setting, in the **Destination Address** field, type a valid destination address for any virtual server that uses this SCTP profile.
   c) Click **Add**.

      Repeat the addition of each destination address that you want to provide to SCTP clients.
6. From the **Server Side Multi-homing** list, select **Enabled** to enable SCTP multihoming for servers.
7. Click **Finished**.

An SCTP profile is configured for multihoming functionality, providing path redundancy for an SCTP connection by enabling SCTP to send packets between multiple addresses owned by each endpoint.

Configuring IP addresses for multihoming connections

In configuring SCTP multihoming for servers, you can create a server node for each IP address on the servers. For example, to configure two servers with three addresses for each server, you need to create six server nodes.

1. On the Main tab, click **Local Traffic** > **Pools**.
   The Pool List screen opens.
2. Click **Create**.
   The New Pool screen opens.
3. 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**.
4. Click **Finished**.
   The screen refreshes, and you see the new pool in the Pool list.

A pool is now available for use with an SCTP virtual server.

Creating a SNAT pool for SCTP

A virtual server requires a SNAT pool of self-IP addresses so it can provide SCTP multihomed connections to servers.

1. On the Main tab, click **Local Traffic** > **Address Translation** > **SNAT Pool List**.
   The SNAT Pool List screen displays a list of existing SNATs.
2. In the **Name** field, type a name for the SNAT pool.
   An example of a name is `snat-pool-1`.
3. For the **Member List** setting:

a) In the **IP Address** field, type an IP address.

   The BIG-IP system uses this address as a SNAT translation address.

   Important: *This address must NOT be on a directly-connected network.*

b) Click **Add**.

c) Repeat these steps for each IP address that you want to include in the SNAT pool.

4. Click the **Finished** button.

A SNAT pool is available for use in a SCTP virtual server configuration.

Creating an SCTP pool for multihoming

You can create an SCTP pool for multihoming, which includes a pool member for each server's IP address to efficiently manage the SCTP traffic 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 SCTP multihoming pool appears in the Pools list.

Configuring an SCTP virtual server for multihoming

You must prepare an SCTP profile before you can support SCTP on a virtual server. SCTP multihoming also requires the following objects:

- A SNAT Pool containing the server-side IP addresses for the virtual server
- A Node Pool with one node per server-IP address

You can use SCTP as a transport for a virtual server, similar to TCP or UDP.

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 column, click the name of the relevant virtual server.
   This displays the properties of the virtual server.
4. In the **Destination Address** field, type the IP address for the SCTP client 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: *This destination address supports the initial SCTP control connection, providing the initial handshake and transfer of valid destination addresses.*
5. From the **Protocol** list, select **SCTP**.
6. From the **Protocol Profile (Client)** list, select a predefined or user-defined SCTP profile.
7. From the **Source Address Translation** list, select **SNAT**.
8. From the **SNAT pool** list, select the name of an existing SNAT pool.
9. Click **Finished**.

The new virtual server supports SCTP for multihoming functionality.

Load Balancing Diameter Application Requests

Overview: Diameter load balancing

An optional feature of the BIG-IP® system is its ability to load balance and persist requests that applications send to servers running Diameter services. The BIG-IP system can also monitor each server to ensure that the Diameter service remains up and running.

About Diameter profiles

The BIG-IP® system includes a profile type that you can use to manage Diameter traffic. The *Diameter protocol* is an enhanced version of the Remote Authentication Dial-In User Service (RADIUS) protocol.

When you configure a Diameter type of profile, the BIG-IP® system can send client-initiated Diameter messages to load balancing servers. The BIG-IP system can also ensure that those messages persist on the servers.

Task summary

Complete these tasks to configure Diameter load balancing on a BIG-IP® system.

Task list

Creating a custom Diameter profile
Creating a custom Diameter monitor
Creating a pool to manage Diameter traffic
Creating a virtual server to manage Diameter traffic

Creating a custom Diameter profile

The first task in configuring Diameter load balancing on the BIG-IP® system is to create a custom Diameter profile.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Services** > **Diameter**.
   The Diameter profile list screen opens.
2. Click **Create**.
   The New Diameter profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Click **Finished**.

The custom Diameter profile appears in the **New Diameter Profile** list.

Creating a custom Diameter monitor

After you create a Diameter profile, you can create a custom Diameter monitor. The purpose of the Diameter monitor is to monitor the health of all servers running the Diameter service.

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

2. Click **Create**.
   The New Monitor screen opens.
3. Type a name for the monitor in the **Name** field.
4. From the **Type** list, select **Diameter**.
   The screen refreshes, and displays the configuration options for the **Diameter** monitor type.
5. Configure additional settings based on your network requirements.
6. Click **Finished**.

Creating a pool to manage Diameter traffic

The next step in a basic Diameter load balancing configuration is to define a load balancing pool that contains Diameter servers as its members.

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. 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**.
5. Click **Finished**.

The pool is configured to manage Diameter servers as pool members.

Creating a virtual server to manage Diameter traffic

The final task in configuring Diameter load balancing is to define a virtual server that references the custom Diameter profile and Diameter pool that you created in previous tasks.

Note: *The virtual server to which you assign the Diameter profile must be a Standard type of 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. 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. From the **Configuration** list, select **Advanced**.
6. From the **Diameter Profile** list, select a profile.
7. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
8. Click **Finished**.

The virtual server that references the custom Diameter profile and Diameter pool appears in the Virtual Server list.

Using SIP Profiles

About SIP profiles

The BIG-IP® system includes a services profile that you can use to manage Session Initiation Protocol (SIP) traffic. *Session Initiation Protocol* is an application-layer protocol that manages sessions consisting of multiple participants, thus enabling real-time messaging, voice, data, and video. A session can be a simple two-way telephone call or Instant Message dialogue, or a complex, collaborative, multi-media conference call that includes voice, data, and video.

SIP sessions, which are application level sessions, run through one of three Layer 4 protocols: SCTP, TCP, or UDP. The SIP profile configures how the system handles SIP sessions. The specified Layer 4 protocol profile configures the virtual server to open the required port to allow data to flow through the BIG-IP® system. When you assign a SIP profile to a virtual server, you can also assign an SCTP, TCP, or UDP profile to the server. If you do not assign one of these protocol profiles to the server, the BIG-IP system automatically assigns one for you.

The SIP profile automatically configures the BIG-IP system to handle persistence for SIP sessions using Call-ID. The Call-ID is a globally unique identifier that groups together a series of messages, which are sent between communicating applications. You can customize how the system handles persistence for SIP sessions.

Maximum message size

The BIG-IP system accepts incoming SIP messages that are 65535 bytes or smaller. If a SIP message exceeds this value, the system drops the connection.

Dialog snooping

The BIG-IP system can snoop SIP dialog information and automatically forward SIP messages to the known SIP dialog. To forward these messages, you can specify a SIP proxy functional group.

Community string

You can specify the name of a proxy functional group. You use this setting in the case where you need multiple virtual servers, each referencing a SIP-type profile, and you want more than one of those profiles to belong to the same proxy functional group.

Connection termination criteria

The BIG-IP system terminates a SIP connection when either the application that initiated the session (client) or the application that answered the initiated session (server) issues a BYE transaction. This is appropriate when a SIP session is running on UDP. However, if a SIP session is running on a SCTP or TCP connection, you can prevent the system from terminating the SIP connection.

SIP headers

An optional feature in a SIP profile is header insertion. You can specify whether the BIG-IP system inserts Via, Secure Via, and Record-Route headers into SIP requests. When you assign the configured SIP profile to a virtual server, the BIG-IP system then inserts the header specified in the profile into any SIP request that the BIG-IP system sends to a pool or pool member.

SIP OneConnect

The SIP OneConnect™ feature allows connection flow reuse between inbound and outbound virtual servers for UDP connections. This feature addresses common SIP client behavior where source and destination ports are both 5060.

SIP OneConnect features a built-in dialog-aware behavior that addresses scenarios where the BIG-IP is the intermediary between more than two parties, creating an ambiguity between source and destination for the dialog. For example, in scenarios where an internal client initiates an outbound call using the wildcard virtual server to an external client that already has an existing flow on the inbound virtual server, the SIP OneConnect dialog-aware behavior correctly routes the response traffic.

Note: *The SIP OneConnect dialog-aware feature is independent of the* ***Dialog Aware*** *setting in the SIP profile and is activated as long as two SIP profiles have identical community strings.*

Activating SIP OneConnect

To activate the SIP OneConnect feature, type identical community strings in both SIP profiles used for the two virtual servers responsible for inbound and outbound SIP connections.

To disable the SIP OneConnect dialog-aware behavior and re-enable the default dialog-aware behavior, check the **Dialog Aware** setting when both community strings are set.

Managing GTP Traffic

About GTP profiles

You can create a GPRS Tunneling Protocol (GTP) profile type on the BIG-IP® system to manage Global System for Mobile Communication (GSM), Universal Mobile Telecommunications System (UMTS), and latterly Long-Term Evolution (LTE) subscriber traffic across User Datagram Protocol (UDP) connections. The BIG-IP system supports GTP versions 1 and 2 on UDP connections. When configuring the GTP profile, you can specify the maximum number of messages held in the GTP ingress queue.

Overview: Managing GTP traffic

The BIG-IP® system enables you to manage GTP traffic by configuring the GTP profile for use with a pool and virtual server. When using the GTP profile, you can specify the maximum number of messages held in the GTP ingress queue.

Task summary

Creating a pool
Creating a GTP profile
Creating a virtual server for GTP traffic

Creating a pool

You can create a pool of servers that you can group together to receive and process 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. 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**.
5. Click **Finished**.
6. Repeat these steps for each pool you want to create.

The new pool appears in the Pools list.

Creating a GTP profile

You create a GTP profile to manage GTP traffic.

1. On the Main tab, click **Local Traffic** > **Profiles**.
2. Click **Create**.
   The New GTP Profile screen opens.

3. In the **Name** field, type a unique name for the profile.
4. From the **Parent Profile** list, retain the default value or select another existing profile of the same type.
5. Select the **Custom** check box.
6. (Optional) In the **Ingress Maximum** field, type the maximum number of messages that can be held in the GTP ingress queue.
7. Click **Finished**.

The GTP profile is configured to manage GTP traffic.

Creating a virtual server for GTP traffic

This task creates a GTP destination to manage GTP 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 a wildcard network address in CIDR format, such as `0.0.0.0/0` for IPv4 or `::/0` for IPv6, to accept any traffic.
5. In the **Service Port** field, type the port number used for the GTP connection.

   Note: *Port `2123` is the default GTP-C port, and port `2152` is the default GTP-U port.*

6. From the **Protocol** list, select **UDP**.
7. From the **GTP Profile** list, select **gtp**, or a user-defined GTP profile.
8. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
9. Click **Finished**.

You now have a virtual server to use as a GTP destination to manage GTP traffic.

Configuring Remote RADIUS Authentication

Overview of remote authentication for application traffic

As an administrator in a large computing environment, you can set up the BIG-IP® system to use this server to authenticate any network traffic passing through the BIG-IP system. This type of traffic passes through a virtual server and through Traffic Management Microkernel (TMM) interfaces. Remote authentication servers typically use one of these protocols:

- Lightweight Directory Access Protocol (LDAP)
- Remote Authentication Dial-in User Service (RADIUS)
- TACACS+ (derived from Terminal Access Controller Access Control System [TACACS])
- Online Status Certificate Protocol (OCSP)
- Certificate Revocation List Distribution Point (CRLDP)

To configure remote authentication for this type of traffic, you must create a configuration object and a profile that correspond to the type of authentication server you are using to store your user accounts. For example, if your remote authentication server is an LDAP server, you create an LDAP configuration object and an LDAP profile. When implementing a RADIUS, SSL OCSP, or CRLDP authentication module, you must also create a third type of object. For RADIUS and CRLDP authentication, this object is referred to as a server object. For SSL OCSP authentication, this object is referred to as an OCSP responder.

About RADIUS profiles

The BIG-IP® system includes a profile type that you can use to load balance Remote Authentication Dial-In User Service (RADIUS) traffic.

When you configure a RADIUS type of profile, the BIG-IP system can send client-initiated RADIUS messages to load balancing servers. The BIG-IP system can also ensure that those messages are persisted on the servers.

Task summary for RADIUS authentication of application traffic

To configure remote authentication for RADIUS traffic, you must create a configuration object and a profile that correspond to the RADIUS authentication server you are using to store your user accounts. You must also create a third type of object. This object is referred to as a server object.

Task list

Creating a RADIUS server object for authenticating application traffic remotely

Creating a RADIUS configuration object for authenticating application traffic remotely

Creating a custom RADIUS profile

Modifying a virtual server for RADIUS authentication

Creating a RADIUS server object for authenticating application traffic remotely

A *RADIUS server object* represents the remote RADIUS server that the BIG-IP system uses to access authentication data.

1. On the Main tab of the navigation pane, click **Local Traffic** > **Profiles**.
2. From the Authentication menu, choose **RADIUS Servers**.
3. Click **Create**.
4. In the **Name**field, type a unique name for the server object, such as`my_radius_server`.
5. In the **Host** field, type the host name or IP address of the RADIUS server.
6. In the **Service Port** field, type the port number for RADIUS authentication traffic, or retain the default value (`1812`).
7. In the **Secret** field, type the secret key used to encrypt and decrypt packets sent or received from the server.
8. In the **Confirm Secret** field, re-type the secret you specified in the **Secret** field.
9. In the **Timeout** field, type a timeout value, in seconds, or retain the default value (`3`).
10. Click **Finished**.

You now have a RADIUS server object that the RADIUS configuration object can reference.

Creating a RADIUS configuration object for authenticating application traffic remotely

The BIG-IP system configuration must include at least one RADIUS server object.

You use a RADIUS authentication module when your authentication data is stored on a remote RADIUS server. A *RADIUS configuration object* specifies information that the BIG-IP system needs to perform the remote authentication.

1. On the Main tab of the navigation pane, click **Local Traffic** > **Profiles**.
2. From the Authentication menu, choose **Configurations**.
3. Click **Create**.
4. In the **Name** field, type a unique name for the configuration object, such as`my_radius_config`.
5. From the **Type** list, select **RADIUS**.
6. For the **RADIUS Servers**setting, select a RADIUS server name in the **Available** list, and using the Move button, move the name to the **Selected** list.
7. In the **Client ID** field, type a string for the system to send in the **Network Access Server (NAS)-Identifier** RADIUS attribute.
8. Click **Finished**.

You now have a RADIUS configuration object that a RADIUS profile can reference.

Creating a custom RADIUS profile

The next task in configuring RADIUS-based remote authentication on the BIG-IP® system is to create a custom RADIUS profile.

1. On the Main tab, click **Local Traffic** > **Profiles** > **Authentication** > **Profiles**.
   The Profiles list screen opens.
2. Click **Create**.
   The New Authentication Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. Select **RADIUS** from the **Type** list.
5. Select **radius** in the **Parent Profile** list.
6. Select the RADIUS configuration object that you created from the **Configuration** list.
7. Click **Finished**.

The custom RADIUS profile appears in the **Profiles** list.

Modifying a virtual server for RADIUS authentication

The final task in the process of implementing authentication using a remote RADIUS server is to assign the custom RADIUS profile to a virtual server that is configured to process HTTP traffic (that is, a virtual server to which an HTTP profile is assigned).

1. On the Main tab, click **Local Traffic** > **Virtual Servers**.
   The Virtual Server List screen opens.
2. Click the name of a virtual server.
3. From the **Configuration** list, select **Advanced**.
4. For the **Authentication Profiles** setting, in the **Available** field, select a custom RADIUS profile, and using the **Move** button, move the custom RADIUS profile to the **Selected** field.
5. Click **Update** to save the changes.

The virtual server is assigned the custom RADIUS profile.

Legal Notices

Legal notices

Publication Date

This document was published on March 27, 2018.

Publication Number

MAN-0652-00

Copyright

Copyright © 2018, F5 Networks, Inc. All rights reserved.

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.

Index