

BIG-IP® TMOS®: Concepts

Version 11.5



Table of Contents

What Is the BIG-IP System?	9
General Configuration Properties	11
TMOS general configuration properties	11
About general device properties	11
Updates to the IP geolocation database	11
About network time protocol (NTP)	12
About DNS configuration	12
About local-traffic properties	12
General local traffic properties	13
SSL Certificates for BIG-IP Devices	15
About SSL digital certificates on the BIG-IP system	15
Supported certificate/key types	15
About RSA certificates	15
About DSA certificates	16
About ECDSA certificates	16
About certificate management	16
Creating a self-signed digital certificate	16
Requesting a certificate from a certificate authority	17
Importing a certificate signed by a certificate authority	18
Exporting a digital certificate	18
Viewing a list of certificates on the system	19
Digital certificate properties	19
External File Management	21
Introduction to external file management	21
Data group files	21
About iFiles	21
External monitor program files	21
SSL certificate files	22
Platform Properties	23
Introduction to platform properties	23
General properties	23
Redundant device properties	24
User administration properties	24
Administrative account passwords	24

SSH access configuration.....	24
Archives.....	25
About archives.....	25
About saving archives.....	26
About archive restoration.....	26
Services.....	27
Managing BIG-IP System Services.....	27
Working with Partitions.....	29
What is an administrative partition?.....	29
About partition Common.....	29
About the current partition.....	30
Relationship of partitions to user accounts.....	30
Object referencing between partitions.....	31
Valid object referencing.....	31
Invalid object referencing.....	32
Working with Folders.....	33
About folders on the BIG-IP system.....	33
Relationship of folders to partitions.....	33
Folder and object naming.....	33
Object referencing between folders.....	33
Users.....	35
Purpose of BIG-IP user accounts.....	35
User account types.....	35
What are user roles?.....	36
User roles on the BIG-IP system.....	36
Default user roles.....	38
Administrative partitions.....	39
Effect of user roles on objects within partitions.....	40
Local user account management.....	40
Admin account configuration.....	40
About secure password policy configuration.....	41
User authentication logout.....	42
Local user account creation.....	43
Local user account view.....	44
Local user account modification.....	45
Delete local user accounts.....	47
Remote user account management.....	48
Introduction to remote BIG-IP user accounts.....	48

Effects of remote user authorization on user accounts.....	52
Authorization for group-based remote user accounts.....	53
About viewing remote user accounts.....	54
About auditing user access to the system.....	55
Logging.....	57
BIG-IP system logging overview.....	57
Types of log messages.....	57
Existing Syslog configurations.....	57
Remote storage of log messages.....	57
Local storage of log messages.....	58
Log level settings for BIG-IP system events.....	59
About local Syslog logging.....	59
Logging system events.....	60
Logging BIG-IP system configuration changes (audit logging).....	60
Code expansion in Syslog log messages.....	60
About enabling and disabling auditing logging.....	61
About remote logging using Syslog-ng.....	61
Interface Concepts.....	63
Introduction to BIG-IP system interfaces.....	63
About link layer discovery protocol.....	63
Interface properties.....	64
Interface naming conventions.....	64
Viewing interface information and media properties.....	65
Interface state.....	65
Fixed Requested Media.....	65
About flow control.....	65
About the LLDP property.....	66
LLDP Attributes.....	66
About interface mirroring.....	66
Neighbor settings.....	66
Related configuration tasks.....	67
Self IP Addresses.....	69
Introduction to self IP addresses.....	69
Types of self IP addresses.....	69
Self IP addresses and MAC addresses.....	70
Self IP addresses for SNATs.....	70
Self IP address properties.....	70
Packet Filters.....	73
Introduction to packet filtering.....	73

Global settings.....	73
Global properties.....	73
Global exemptions.....	74
Protocols.....	75
MAC addresses.....	75
IP addresses.....	75
VLANs.....	75
Order of packet filter rules.....	76
About the action setting in packet filter rules.....	76
Rate class assignment.....	77
One or more VLANs.....	77
Logging.....	77
About filter expression creation.....	77
Spanning Tree Protocol.....	79
Introduction to spanning tree protocols.....	79
About STP protocol.....	79
About the RSTP protocol.....	79
About the MSTP protocol.....	80
About spanning tree with legacy bridges.....	80
Configuration overview.....	81
Spanning tree mode.....	81
Global timers.....	82
About the hello time option.....	82
About the maximum age option.....	82
About the forward delay option.....	82
About the transmit hold count option.....	83
MSTP-specific global properties.....	83
Management of spanning tree instances.....	83
Spanning tree instances list.....	84
About spanning tree instance (MSTP-only) creation.....	84
About instance ID assignment.....	86
Bridge priority.....	86
VLAN assignment.....	86
About viewing and modifying a spanning tree instance.....	86
About deleting a spanning tree instance or its members (MSTP-only).....	87
Interfaces for spanning tree.....	87
Enabling and disabling spanning tree.....	87
STP link type.....	87
STP edge port.....	88
About spanning tree protocol reset.....	88
About managing interfaces for a specific instance.....	88
About viewing a list of interface IDs for an instance.....	89
About port roles.....	89

Port states.....	89
Settings to configure for an interface for a specific instance.....	90
Trunks.....	91
Introduction to trunks.....	91
How do trunks work?.....	91
Overview of LACP.....	92
Trunk name.....	92
Interfaces for a trunk.....	92
About enabling LACP.....	93
LACP mode.....	93
LACP timeout.....	94
Link selection policy.....	94
Automatic link selection.....	94
Maximum bandwidth link selection.....	95
Frame distribution hash.....	95
VLANs, VLAN Groups, and VXLAN.....	97
Introduction to virtual LANs.....	97
VLANs on a BIG-IP system.....	97
Default VLAN configuration.....	97
VLAN name.....	98
VLAN tags.....	98
Interface assignments.....	99
Port-based access to VLANs.....	99
Tag-based access to VLANs.....	99
Source checking.....	100
Maximum transmission units.....	101
VLAN-based fail-safe.....	101
Auto last hop.....	101
CMP hash.....	101
DAG round robin.....	101
Maintaining the L2 forwarding table.....	102
About sFlow polling intervals and sampling rates.....	102
About VLAN groups.....	102
About VLAN group names.....	103
VLAN group ID.....	103
About transparency mode.....	103
About traffic bridging.....	104
About traffic bridging with standby units.....	104
About migration keepalive frames.....	104
About host exclusion from proxy ARP forwarding.....	104
VLAN association with a self IP address.....	105
VLAN assignment to route domains.....	105

About bridging VLAN and VXLAN networks.....	105
About VXLAN multicast configuration.....	106
WCCPv2.....	107
About WCCPv2 redirection on the BIG-IP system.....	107
A common deployment of the WCCPv2 protocol.....	108
Failsafe.....	109
About system fail-safe.....	109
About VLAN fail-safe.....	109
About gateway fail-safe.....	110
Legal Notices and Acknowledgments.....	111
Legal Notices.....	111
Acknowledgments.....	112

What Is the BIG-IP System?

The BIG-IP® system is a set of application delivery products that work together to ensure high availability, improved performance, application security, and access control.

One of the primary functions of the BIG-IP system is to direct different types of protocol and application traffic to an appropriate destination server. The system accomplishes this through its Local Traffic Manager™ module, which can forward traffic directly to a load balancing server pool, or send traffic to a next-hop router, a pool of routers, or directly to a selected node on the network.

Other modules available on the BIG-IP system provide critical functions such as applying security policies to network traffic, accelerating HTTP connections, and optimizing connections across a wide-area network.




Figure 1: Basic BIG-IP system configuration

General Configuration Properties

TMOS general configuration properties

Part of managing the BIG-IP® system involves configuring and maintaining a set of global system properties. These properties allow you to configure:

- General device features, such as NTP and DNS
- General local traffic features, including some global persistence settings
- General global traffic features, including load balancing and metric collection

When you configure general device properties, you are affecting the operation of the BIG-IP system as a whole, rather than just one aspect of it. Similarly, when you configure the general properties related to local traffic or global traffic, you are globally affecting the operation of the local traffic management and global traffic management systems.

About general device properties

The BIG-IP® system general device properties that you can view or configure are:

- The host name
- The BIG-IP software version number
- The number of CPUs available
- The number of CPUs that are active

Other BIG-IP system general device properties that you can configure are:

- Network boot
- Quiet boot

You can also perform operations such as reboot or force the system into an OFFLINE state, and reload the default geolocation data files that the BIG-IP system uses to source the origin of a name resolution request.

Note: Regarding the Always-On Management (AOM) subsystem, AOM now saves the recent contents of serial output from the host across host reboots, even when no serial console is attached. The contents are written to log files on both the AOM and the host, either after a host reboot/power event or at a user's request. This feature is disabled by default. To enable the feature, you must open Secure Shell (*ssh*) and then manually change and restart the *hostconsh* script. For more information, see the F5 Networks knowledge base at <http://support.f5.com>.

Updates to the IP geolocation database

The BIG-IP® system uses an IP geolocation database to determine the origin of a name resolution request. The default database provides geolocation data for IPv4 addresses at the continent, country, state, ISP, and

organization levels. The state-level data is worldwide, and thus includes designations in other countries that correspond to the U.S. state-level in the geolocation hierarchy, for example, provinces in Canada. The default database also provides geolocation data for IPv6 addresses at the continent and country levels.

Note: You can access the ISP and organization-level geolocation data for IPv4 and IPv6 addresses using the *iRules® whereis* command.

Tip: If you require geolocation data at the city-level, contact your F5 Networks® sales representative to purchase additional database files.

You can download a monthly update to the IP geolocation database from F5 Networks.

About network time protocol (NTP)

Network Time Protocol (NTP) is a protocol that synchronizes the clocks on a network. Because, by default, DHCP is enabled for the BIG-IP® system, on the first boot, the BIG-IP system contacts your DHCP server and obtains the IP address of your NTP server. If the DHCP server provides this IP address, the NTP Device Configuration screen displays the NTP server information. If you do not have a DHCP server on your network, or if the DHCP server does not return the IP address of your NTP server, you can manually add the IP address of the NTP server to the BIG-IP system using the BIG-IP Configuration utility.

About DNS configuration

Domain Name System (DNS) is an industry-standard distributed internet directory service that resolves domain names to IP addresses. When you enable DHCP, the system contacts your DHCP server to obtain the IP addresses of your local DNS servers and the domain names that the system searches to resolve local host names. If the DHCP server provides this information, the DNS Device Configuration screen displays the information in the DNS Lookup Server List and the DNS Search Domain List.

If you do not have a DHCP server on your network, or if the DHCP server does not supply the information, you can manually create the two lists. The DNS Lookup Server List allows BIG-IP® system users to use IP addresses, host names, or fully-qualified domain names (FQDNs) to access virtual servers, nodes, or other network objects. The DNS Search Domain List allows BIG-IP system to search for local domain lookups to resolve local host names.

Additionally, you can manually configure the *BIND Forwarder Server List* that provides DNS resolution for servers and other equipment load balanced by the BIG-IP system, that is, for the servers that the BIG-IP system uses for DNS proxy services.

Note: To use DNS Proxy services, you must enable the named service.

About local-traffic properties

The BIG-IP® system includes a set of properties that apply globally to the local traffic management system. These properties fall into two main categories: general local-traffic properties, and persistence properties. You can use the BIG-IP Configuration utility to configure and maintain these properties.

General local traffic properties

This table lists and describes global properties that you can configure to manage the behavior of the local traffic management system.

Property	Default Value	Description
Auto Last Hop	Enabled (checked)	Specifies, when checked (enabled), that the system automatically maps the last hop for pools.
Maintenance Mode	Disabled (unchecked)	Specifies, when checked (enabled), that the unit is in maintenance mode. In maintenance mode, the system stops accepting new connections and slowly completes the processing of existing connections.
VLAN-Keyed Connections	Enabled (checked)	Check this setting to enable VLAN-keyed connections. VLAN-keyed connections are used when traffic for the same connection must pass through the system several times, on multiple pairs of VLANs (or in different VLAN groups).
Path MTU Discovery	Enabled (checked)	Specifies, when checked (enabled), that the system discovers the maximum transmission unit (MTU) that it can send over a path without fragmenting TCP packets.
Reject Unmatched Packets	Enabled (checked)	Specifies that the BIG-IP system sends a TCP RST packet in response to a non-SYN packet that matches a virtual server address and port or self IP address and port, but does not match an established connection. The BIG-IP system also sends a TCP RST packet in response to a packet matching a virtual server address or self IP address but specifying an invalid port. The TCP RST packet is sent on the client-side of the connection, and the source IP address of the reset is the relevant BIG-IP LTM object address or self IP address for which the packet was destined. If you disable this setting, the system silently drops unmatched packets.
Reaper High-water Mark	95	Specifies, in percent, the memory usage at which the system silently purges stale connections, without sending reset packets (RST) to the client. If the memory usage remains above the low-water mark after the purge, then the system starts purging established connections closest to their service timeout. To disable the adaptive reaper, set the high-water mark to 100.
Reaper Low-water Mark	85	Specifies, in percent, the memory usage at which the system starts establishing new connections. Once the system meets the reaper high-water mark, the system does not establish new connections until the memory usage drops below the reaper low-water mark. To disable the adaptive reaper, set the low-water mark to 100. This setting helps to mitigate the effects of a denial-of-service attack.
SYN Check™ Activation Threshold	16384	Specifies the number of new or untrusted TCP connections that can be established before the system activates the SYN Cookies authentication method for subsequent TCP connections.
Layer 2 Cache Aging Time	300	Specifies, in seconds, the amount of time that records remain in the Layer 2 forwarding table, when the MAC address of the record is no longer detected on the network.

Property	Default Value	Description
Share Single MAC Address	Disabled (unchecked)	When this setting is unchecked (disabled), the BIG-IP system assigns to each VLAN a unique MAC address that comes from a pool of available MAC addresses. If you create enough VLANs to exceed the number of MAC addresses available, the system then begins to assign the same MAC address to multiple VLANs. This is the default value and the most common configuration. When this setting is checked (enabled), the BIG-IP system causes all VLANs to share a single MAC address (global). This setting is equivalent to the BigDB variable <code>vlan.macassignment</code> and has two values, <code>unique</code> and <code>global</code> .
SNAT Packet Forwarding	TCP and UDP Only	Specifies the type of traffic for which the system attempts to forward (instead of reject) Any-IP packets, when the traffic originates from a member of a SNAT. There are two possible values: TCP and UDP Only specifies that the system forwards, for TCP and UDP traffic only, Any-IP packets originating from a SNAT member. All Traffic specifies that the system forwards, for all traffic types, Any-IP packets originating from a SNAT member.

SSL Certificates for BIG-IP Devices

About SSL digital certificates on the BIG-IP system

An *SSL digital certificate* is an electronic key pair that allows devices on a network to exchange data securely, using the public key infrastructure (PKI). PKI is based on public and private cryptographic key pairs used to encrypt and decrypt messages sent between two devices.

The BIG-IP® system uses digital certificates with the SSL/TLS protocol to grant authentication to clients on the external network that are generally untrusted. In high-security environments, the BIG-IP system can also use certificates to communicate securely with other systems on the internal network, such as web servers and other BIG-IP systems.

The BIG-IP system can sign a digital certificate in either of two ways:

- By generating and submitting a request to a third-party trusted certificate authority (CA)
- By creating a self-signed certificate. Self-signed certificates are typically used for testing purposes.

Once a certificate is installed or created on the BIG-IP system, other BIG-IP administrative users can specify those certificates in BIG-IP SSL profiles to manage SSL application traffic. Moreover, the BIG-IP system uses digital certificates to establish device trust in device service clustering (DSC™) configurations.

Supported certificate/key types

The BIG-IP® system supports multiple cipher suites when offloading SSL operations from a target server on the network. The BIG-IP system can support cipher suites that use these algorithms:

- Rivest Shamir Adleman (RSA)
- Elliptic Curve Digital Signature Algorithm (ECDSA)
- Digital Signature Algorithm (DSA)

When you generate a certificate request or a self-signed certificate, you specify the type of private key, which determines that specific signing or encryption algorithm that is used to generate the private key.

About RSA certificates

RSA (Rivest Shamir Adleman) is the original encryption algorithm that is based on the concept of a public and a private key. When a public site attempts to communicate with a device such as the BIG-IP® system, the device sends the site a public key that the site uses to encrypt data before sending that data back to the device. The device uses its private key associated with the public key to decrypt the data. Only the device on which the certificate resides has access to this private key.

The RSA encryption algorithm includes an authentication mechanism.

About DSA certificates

DSA (Digital Signature Algorithm) uses a different algorithm for signing key exchange messages than that of RSA. DSA is paired with a key exchange method such as Diffie-Hellman or Elliptical Curve Diffie-Hellman to achieve a comparable level of security to RSA. Because DSA is generally endorsed by federal agencies, specifying a DSA key type makes it easier to comply with new government standards, such as those for specific key lengths.

About ECDSA certificates

When creating certificates on the BIG-IP® system, you can create a certificate with a key type of ECDSA (Elliptic Curve Digital Signature Algorithm). An *ECDSA key* is based on Elliptic Curve Cryptography (ECC), and provides better security and performance with significantly shorter key lengths.

For example, an RSA key size of 2048 bits is equivalent to an ECC key size of only 224 bits. As a result, less computing power is required, resulting in faster, more secure connections. Encryption based on ECC is ideally suited for mobile devices that cannot store large keys. The BIG-IP system supports both the prime256v1 and secp384r1 curve names, although only prime256v1 can be associated with an SSL profile.

About certificate management

You can obtain a certificate for the BIG-IP system by using the BIG-IP® Configuration utility to generate a certificate signing request (CSR) that can then be submitted to a third-party trusted certificate authority (CA). The CA then issues a signed certificate.

In addition to requesting CA-signed certificates, you can create self-signed certificates. You create self-signed certificates primarily for testing purposes within an organization.

When you install the BIG-IP software, the application includes a default self-signed certificate. The BIG-IP system also includes a default CA bundle certificate. This certificate bundle contains certificates from most of the well-known CAs.

Note: To manage digital certificates for the BIG-IP system, you must have a role of Certificate Manager, Administrator, or Resource Administrator assigned to your BIG-IP user account.

Creating a self-signed digital certificate

If you are configuring the BIG-IP® system to manage client-side HTTP traffic, you perform this task to create a self-signed certificate to authenticate and secure the client-side HTTP traffic. If you are also configuring the system to manage server-side HTTP traffic, you must repeat this task to create a second self-signed certificate to authenticate and secure the server-side HTTP traffic.

1. On the Main tab, click **System > File Management > SSL Certificate List**.
The SSL Certificate List screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique name for the SSL certificate.
4. From the **Issuer** list, select **Self**.

5. In the **Common Name** field, type a name.
6. In the **Division** field, type your company name.
7. In the **Organization** field, type your department name.
8. In the **Locality** field, type your city name.
9. In the or **State or Province** field, type your state or province name.
10. From the **Country** list, select the name of your country.
11. In the **E-mail Address** field, type your email address.
12. In the **Lifetime** field, type a number of days, or retain the default, **365**.
13. In the **Subject Alternative Name** field, type a name.
This name is embedded in the certificate for X509 extension purposes.
By assigning this name, you can protect multiple host names with a single SSL certificate.
14. From the **Key Type** list, select a key type.
Possible values are: **RSA**, **DSA**, and **ECDSA**.
15. From the **Size** or **Curve Name** list, select either a size, in bits, or a curve name.
16. If the BIG-IP system contains an internal HSM module, specify a location for storing the private key.
17. Click **Finished**.

Requesting a certificate from a certificate authority

You perform this task to generate a certificate signing request (CSR) that can then be submitted to a third-party trusted certificate authority (CA).

1. On the Main tab, click **System > File Management > SSL Certificate List**.
The SSL Certificate List screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique name for the SSL certificate.
4. From the **Issuer** list, select **Certificate Authority**.
5. In the **Common Name** field, type a name.
6. In the **Division** field, type your company name.
7. In the **Organization** field, type your department name.
8. In the **Locality** field, type your city name.
9. In the or **State or Province** field, type your state or province name.
10. From the **Country** list, select the name of your country.
11. In the **E-mail Address** field, type your email address.
12. In the **Lifetime** field, type a number of days, or retain the default, **365**.
13. In the **Subject Alternative Name** field, type a name.
This name is embedded in the certificate for X509 extension purposes.
By assigning this name, you can protect multiple host names with a single SSL certificate.
14. In the **Challenge Password** field, type a password.
15. In the **Confirm Password** field, re-type the password you typed in the **Challenge Password** field.
16. From the **Key Type** list, select a key type.
Possible values are: **RSA**, **DSA**, and **ECDSA**.
17. From the **Size** or **Curve Name** list, select either a size, in bits, or a curve name.
18. If the BIG-IP system contains an internal HSM module, specify a location for storing the private key.

19. Click **Finished.**

The Certificate Signing Request screen displays.

20. Do one of the following to download the request into a file on your system.

- In the **Request Text** field, copy the certificate.
- For **Request File**, click the button.

21. Follow the instructions on the relevant certificate authority web site for either pasting the copied request or attaching the generated request file.

22. Click **Finished.**

The Certificate Signing Request screen displays.

The generated certificate signing request is submitted to a trusted certificate authority for signature.

Importing a certificate signed by a certificate authority

Before performing this task, confirm that a digital certificate signed by a certificate authority is available.

You can install an SSL certificate that is signed by a certificate authority by importing a certificate that already exists on the system hard drive. You can import a private key, a certificate or certificate bundle, or an archive.

1. On the Main tab, click **System > File Management > SSL Certificate List.**

The SSL Certificate List screen opens.

2. Click **Import.**

3. From the **Import Type list, select **Certificate**.**

4. For the **Certificate Name setting, do one of the following:**

- Select the **Create New** option, and type a unique name in the field.
- Select the **Overwrite Existing** option, and select a certificate name from the list.

5. For the **Certificate Source setting, do one of the following:**

- Select the **Upload File** option, and browse to the location of the certificate file.
- Select the **Paste Text** option, and paste the certificate text copied from another source.

6. Click **Import.**

The SSL certificate that was signed by a certificate authority is installed.

Exporting a digital certificate

You perform this task to export a digital certificate to another device.

1. On the Main tab, click **System > File Management > SSL Certificate List.**

The SSL Certificate List screen opens.

2. Click the name of the certificate you want to export.

The General Properties screen displays.

3. Click **Export.**

The Certificate Export screen displays the contents of the certificate in the **Certificate Text** box.

4. To obtain the certificate, do one of the following:

- Copy the text from the **Certificate Text** field, and paste it as needed into an interface on another system.
- At the **Certificate File** option, click **Download filename** where filename is the name of the certificate file, such as `mycert.crt`.

Viewing a list of certificates on the system

You can perform this task to view a list of existing digital certificates on the BIG-IP® system.

1. On the Main tab, click **System > File Management > SSL Certificate List**.
The SSL Certificate List screen opens.
2. In the Name column, view the list of certificates on the system.

Digital certificate properties

When you use the BIG-IP® Configuration utility to view the list of digital certificates that you have installed on the BIG-IP® system, you can see information for each certificate.

Property	Description
Certificate	The name of the certificate.
Content	The type of certificate content, for example, Certificate Bundle or Certificate and Key.
Common name	The common name (CN) for the certificate. The common name embedded in the certificate is used for name-based authentication. The default common name for a self-signed certificate is <code>localhost.localdomain</code> .
Expiration date	The date that the certificate expires. If the certificate is a bundle, this information shows the range of expiration dates that apply to certificates in the bundle.
Organization	The organization name for the certificate. The organization name embedded in the certificate is used for name-based authentication. The default organization for a self-signed certificate is <code>MyCompany</code> .

External File Management

Introduction to external file management

You can import certain external files for use by iRules®, or you can import or create SSL certificates. The external files that iRules can use are data group files and iFiles. Using the BIG-IP® Configuration utility, you can manage these external files or SSL certificates from a central location.

Data group files

Using the BIG-IP® Configuration utility, you can import an existing file that contains content that you want to reference in an iRule. You import this file from another system to the BIG-IP system.

When you import an existing file to the BIG-IP system, you create an external data group, specifying this information:

- The location of the external file that you want to import to the BIG-IP system.
- A unique name for the imported file.
- The data group type (address, string, or integer).
- The separator for each key/value pair specified in the data group (the default value is : =).
- A unique name for the data group.

About iFiles

Using the BIG-IP® Configuration utility, you can import an existing 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.

To use this feature, you first import an existing file or URL to the BIG-IP system and then assign a new name to the file. To import a file with the BIG-IP Configuration utility and assign it a new name, you use the **System** area of the navigation pane.

External monitor program files

Using the BIG-IP® Configuration utility, you can import an existing external program monitor file to the BIG-IP system, with content that an external monitor can reference.

To use this feature, you first import an existing file from another system to the BIG-IP system and then assign a new name to the file. To import a file and assign it a new name, log in to the BIG-IP Configuration utility, and on the Main tab, expand **System**, and click **File Management**.

After importing the file, you use the Local Traffic area of the BIG-IP Configuration utility to create a new external monitor program based on the imported file.

SSL certificate files

Using the BIG-IP® Configuration utility, you can import an existing certificate file from another system to the BIG-IP system, or you can create a new SSL certificate. To import or create an SSL certificate, log in to the BIG-IP Configuration utility, and on the Main tab, expand **System**, and click **File Management**.

An imported certificate file has these attributes:

- Contents
- Common Name
- Organization
- Expiration date
- Partition / Path

Using the BIG-IP Configuration utility, you can view, import, renew, or export a device certificate.

You can also import or export a device key. The properties of a device key are:

- Key type (such as KTYPE_RSA_PRIVATE)
- Key size (such as 1024 bits)
- Security type, either Normal or FIPS (FIPS-enabled systems only)

There are several types of files that you can import using the File Management screens of the BIG-IP Configuration utility. These file types are:

- Key files
- Certificate files
- PKCS 12 (IIS) files
- Archive files
- Certificate Revocation List (CRL) files

Note: Do not attempt to manage certificates by copying SSL certificate files into the `/config/ssl/*` directory and then reloading the system configuration. Certificate-related files are not stored in that location. Therefore, you must use the BIG-IP Configuration utility or `tmsh` to manage certificate files, key files, and CRL files.

Platform Properties

Introduction to platform properties

Part of managing a BIG-IP[®] system involves configuring and maintaining a certain set of system properties. These properties consist of general platform properties such as the BIG-IP system host name, IP address, and passwords for its system administrative accounts.

General properties

You can configure these general properties for the BIG-IP[®] system platform:

The management port and TMM

The BIG-IP system has a management port to handle administrative traffic, and TMM switch interfaces to handle application traffic. *TMM switch interfaces* are those interfaces controlled by the Traffic Management Microkernel (TMM) service.

Management port configuration

By default, DHCP is disabled for the management port on the BIG-IP system. When enabled, DHCP uses UDP ports 67 and 68. On the first boot, the BIG-IP system contacts your DHCP server and obtains a lease for an IP address and default route for the management port, and DNS and NTP servers. You must then configure other system attributes, such as host name and domain name servers. When DHCP is disabled, you manually configure the management port by assigning an IP address and netmask to the port. The IP address that you assign to the management port must be on a different network than the self IP addresses that you assign to VLANs. You can use either an IPv4 or an IPv6 address for the management port. Additionally, if you intend to manage the BIG-IP system from a node on a different subnet of your network, you can specify an IP address for the BIG-IP system to use as a default route to the management port.

Note: *If you do not have a DHCP server on your network, the BIG-IP system assigns a default IP address of 192.168.1.245 to the management port of appliances and virtual systems, and 192.186.1.246 to the management port of VIPRION[®] systems.*

Host name

Every BIG-IP system must have a host name that is a fully qualified domain name. An example of a host name is `bigip-02.win.net`.

Host IP address

Every BIG-IP system must have a host IP address. This IP address can be the same as the address that you used for the management port, or you can assign a unique address. The default value on the screen for this setting is **Use Management Port IP Address**.

Time zone

Another of the general platform properties that you can specify is the time zone. The many time zones that you can choose from are grouped into these categories: Africa, America, Antarctica, Arctic, Asia,

Atlantic, Australia, Europe, Indian, and Pacific. You should specify the time zone region that most closely represents the location of the BIG-IP system you are configuring.

Redundant device properties

A BIG-IP® system is typically part of a device group that synchronizes configuration data across two or more BIG-IP devices and provides high availability (failover and connection mirroring).

To ensure that this operates successfully, you assign a device group (to the `root` folder) to which you want to synchronize configuration data. All folders and sub-folders in the folder hierarchy inherit this device group as a folder attribute.

You also assign a floating traffic group to the `root` folder. All folders and sub-folders in the folder hierarchy inherit this traffic group as a folder attribute.

User administration properties

Part of managing platform-related properties is maintaining passwords for the system account. You can also configure the system to allow certain IP addresses to access the BIG-IP® system through SSH.

Administrative account passwords

When you ran the Setup utility on the BIG-IP® system, you set up some administrative accounts. Specifically, you set up the `root` and `admin` accounts. The `root` and `admin` accounts are for use by BIG-IP system administrators.

Users logging in with the `root` account have terminal and browser access to the BIG-IP system. By default, users logging in with the `admin` account have browser-only access to the BIG-IP system. You can use the general screen for platform properties to change the passwords for `root` and `admin` accounts on a regular basis. To change a password, locate the **Root Account** or **Admin Account** setting, and in the **Password** field, type a new password. In the **Confirm** field, re-type the same password.

SSH access configuration

When you configure SSH access, you enable user access to the BIG-IP® system through SSH. Also, only the IP addresses that you specify are allowed access to the system using SSH.

To configure SSH access, locate the **SSH Access** setting and select the **Enabled** check box. Then use the **SSH IP Allow** setting to select either *** All Addresses** or **Specify Range**, which allows you to specify a range of addresses for which access is allowed.

Archives

About archives

On any BIG-IP® system, you have a set of data that you created when you initially configured the system, using the Setup utility and the BIG-IP Configuration utility or `tmsh`. This data consists of traffic management elements such as virtual servers, pools, and profiles. Configuration data also consists of system and network definitions such as interface properties, self IP addresses, VLANs, and more.

Once you have created the configuration data for the BIG-IP system, you can replicate all of this set of data in a separate file. You can then use this replicated data later, for these reasons:

As an archive for disaster recovery

Using the Archives feature, you can back up the current configuration data, and if necessary, restore the data at a later time. F5 Networks® recommends that you use this feature to mitigate the potential loss of BIG-IP system configuration data. To create an archive, you can use the BIG-IP Configuration utility, which stores the configuration data in a special file known as a user configuration set, or UCS file. You can then use the UCS file to recover from any loss of data, in the unlikely event that you need to do so.

As a way to propagate data to other systems

Using the single configuration file feature, you can easily and quickly propagate the exact configuration of the BIG-IP system to other BIG-IP systems. To create a single configuration file, you export the configuration data to a special file known as an `.scf` file. You can then use the `.scf` file to configure another system in one simple operation.

Before you replace a version of the BIG-IP system with a newer version, you should always create an *archive*, which is a backup copy of the configuration data. This archive is in the form of a user configuration set, or UCS. Then, if you need to recover that data later, you can restore the data from the archive that you created.

A UCS contains the following types of BIG-IP system configuration data:

- System-specific configuration files
- Product licenses
- User accounts and password information
- Domain Name Service (DNS) zone files
- Installed SSL keys and certificates

Each time you back up the configuration data, the BIG-IP system creates a new file with a `.ucs` extension. Each UCS file contains various configuration files needed for the BIG-IP system to operate correctly, as well as the configuration data.

Important: *To create, delete, upload, or download an archive, you must have either the Administrator or Resource Administrator role assigned to your user account.*

About saving archives

By default, the system stores all archives in the directory `/var/local/ucs`. You can specify a different location, but in this case, the BIG-IP Configuration utility does not display the UCS files when you view the list of archives

After you create an archive on the BIG-IP® system, you can download a copy of the UCS file to the system from which you are running the BIG-IP Configuration utility (a secure remote system). This provides an extra level of protection by preserving the configuration data on a remote system. In the unlikely event that you need to restore the data, and a BIG-IP system event prevents you from accessing the archive in the BIG-IP system directory in which you saved the archive, you still have a backup copy of the data.

Important: *If your configuration data includes SSL keys and certificates, be sure to store the archive file in a secure environment.*

About archive restoration

Not only is the `/var/local/ucs` directory the only location on the BIG-IP® system in which you can save an archive, but it is also the only location on the BIG-IP system from which you can restore an archive. However, if you previously downloaded an archive to a remote system, and a BIG-IP system event prevents you from accessing the `/var/local/ucs` directory, you can upload the archive from that remote system.

Services

Managing BIG-IP System Services

The BIG-IP® system includes a number of services that you can start, stop, or restart using the BIG-IP Configuration utility. This ability to start or stop services from within the BIG-IP Configuration utility is useful when you want to run only those services that you need to successfully manage network traffic.

The BIG-IP Configuration utility screen for managing services lists the name of each service and its current status. At a minimum, the services that you can stop or start with the BIG-IP Configuration utility are:

ntpd

Sets and maintains the system time of day.

snmpd

Receives and processes SNMP requests, and sends trap notifications. Note that you must stop this service before updating the SNMP v3 file `/config/net-snmp/snmpd.conf`, which specifies SNMP user names.

sshd

Provides secure remote login between untrusted hosts.

If you have other BIG-IP product modules running on the BIG-IP system, such as Global Traffic Manager™, you might see other services listed on the Services screen.

Important: *You must have either the Administrator or Resource Administrator user role assigned to your user account to stop, start, or restart a service.*

When you view the list of services available on the BIG-IP system, you can also view current status. In the BIG-IP Configuration utility, you can see this status in the History column. Examples of status for a service are: Running for 6 days, Running, and down, Not provisioned.

Working with Partitions

What is an administrative partition?

An *administrative partition* is a logical container that you create, containing a defined set of BIG-IP® system objects. If you have the Administrator or User Manager user role assigned to the BIG-IP system user account, you can create administrative partitions to control other users' access to BIG-IP objects. More specifically, when a specific set of objects resides in a partition, you can give certain users the authority to view and manage the objects in that partition only, rather than to all objects on the BIG-IP system. This gives a finer granularity of administrative control.

The following illustration shows an example of user objects within partitions on the BIG-IP system.



Figure 2: Sample administrative partitions on the BIG-IP system

For every administrative partition on the BIG-IP system, the BIG-IP system creates an equivalent high-level folder with an equivalent name.

About partition Common

During BIG-IP® system installation, the system automatically creates a partition named `Common`. At a minimum, this partition contains all of the BIG-IP objects that the system creates as part of the installation process.

Until you create other partitions on the system, all objects that you or other users create automatically reside in partition `Common`. If you create other partitions later, you cannot move an object in `Common` to one of the new partitions. Instead, you must delete the object from `Common` and recreate it in the new partition.

With respect to permissions, all users on the system except those with a user role of No Access have read access to objects in partition `Common`, and by default, partition `Common` is their current partition.

Some users, such as those with the user role of Administrator, can also create, update, and delete objects in partition `Common`. No user can delete partition `Common` itself.

About the current partition

The *current partition* is the specific partition to which the system is currently set for a logged-in user.

A user who has permission to access all partitions can actively select the current partition, that is, the specific partition he or she wants to view or manage. A user who has permission to access one partition only cannot actively select the current partition; the system sets the current partition for that user when he or she logs in to the system. For example:

- If user `rsmith` has a role of Manager and has access to all partitions on the system, then before creating or managing any object on the BIG-IP® system, she must select the partition that she wants to be the current partition. After selecting the current partition, any object that she creates will reside in that partition, and she can modify or delete only those objects that reside in the current partition.
- Conversely, if user `rsmith` has the role of Manager and is granted access to partition `A` only, then any object that she creates while logged into the BIG-IP system resides in partition `A`. Although she can view objects in partition `Common`, she cannot select `Common` as her current partition because she has read access only for that partition. For user `rsmith`, partition `A` is automatically her current partition, and she cannot change the current partition to create objects in another partition.

Relationship of partitions to user accounts

Partitions have a special relationship to user accounts. With respect to partitions and user accounts, you can:

Assign partition access to user accounts

You can configure a user account to grant the user access to a specific partition. A partition access assignment gives a user some level of access to the specified partition. As an option, for all user roles, you can assign universal access to partitions. This grants users permission to access all partitions instead of one non-Common partition only. Note that assigning partition access to a user does not necessarily give the user full access to all objects in the partition; the user role assigned to the user determines the type of access that the user has to each type of object in the partition.

Create user accounts as partitioned objects

Like other types of objects on the system, user account objects also reside in partitions. Placing user account objects into partitions controls other users' administrative access to those user accounts. Also, like other object types, a BIG-IP® system user account cannot reside in more than one partition simultaneously. Note that when you first install the BIG-IP system, every existing user account (`root` and `admin`) resides in partition `Common`. Note that the partition in which a user account object resides does not affect the partition or partitions to which that user is granted access to manage other BIG-IP objects

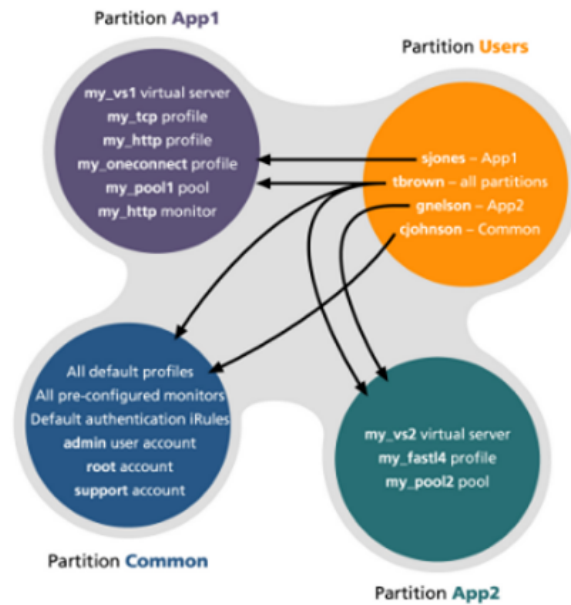


Figure 3: Relationship of partitions to user accounts

Object referencing between partitions

Certain BIG-IP® system objects, such as virtual servers, always reference other objects. Examples of objects that a virtual server can reference are pools, profiles, and iRules®. On BIG-IP system, there are specific validation rules for object referencing with respect to the administrative partitions in which those objects reside.

Valid object referencing

Normally, when you create BIG-IP objects, a referenced object must reside either in the same partition as the object that is referencing it, or in partition **Common**. For example, this figure shows a valid object-referencing configuration where a virtual server and the pool it references reside in the same partition (named **my_app**):



Figure 4: Example of object referencing within a partition

Another valid object referencing case is when the object resides in one partition, while the object it references resides in partition **Common**. This figure shows an example of this configuration, where a virtual server in partition **my_app** references a pool in partition **Common**:




Figure 5: Example of object referencing from another partition to Common

In addition to these two valid object referencing configurations, the system also allows a third type of object referencing, specifically regarding iRules. That is, an iRule can reference any object, regardless of the partition in which the referenced object resides. For example, an iRule that resides in partition `my_app_A` can contain a pool statement that specifies a pool residing in partition `my_app_B`. Neither object is required to reside in `Common`.

Invalid object referencing

This figure shows an example of an invalid object-referencing configuration, where a virtual server resides in partition `Common`, but the pool the virtual server references resides in a different partition. In this case, the virtual server cannot successfully forward traffic to the pool that it is referencing:




Figure 6: Example of object referencing from Common to another partition

Working with Folders

About folders on the BIG-IP system

At the most basic level, a *folder* is a container for BIG-IP® configuration objects and files on a BIG-IP device. Virtual servers, pools, and self IP addresses are examples of objects that reside in folders on the system. Folders resemble standard directories, in that the system includes a root folder (represented by the / symbol) that is the parent for other folders on the system.

A folder can contain other folders.

One of the important ways that you can use folders is to set up full or granular synchronization and failover of BIG-IP configuration data in a device group. You can synchronize and fail over all configuration data on a BIG-IP device, or you can synchronize and fail over objects within a specific folder only.

Relationship of folders to partitions

For each partition on the BIG-IP® system, there is an equivalent high-level folder. For example, for partition `Common`, there is a corresponding high-level folder named `/Common`. This close association of administrative partitions to folders means that when you use the BIG-IP Configuration utility to create objects on a BIG-IP device, the system puts those objects in the current partition, in a folder that you choose. Examples of BIG-IP objects that reside in folders are virtual servers, pools, and self IP addresses.

If you create another administrative partition, such as partition `App_A`, the BIG-IP system automatically creates a high-level folder named `/App_A`. You can then create BIG-IP configuration objects that pertain to application A by changing the current partition to `App_A` and then either creating the objects in folder `/App`, or creating a sub-folder within `/App_A` and then navigating to that sub-folder.

Folder and object naming

The folder in which an object resides automatically becomes part of the object name. For example, if you create a pool in partition `Common`, in folder `/Common`, and then name the pool `my_pool`, the full name of the pool on the system is `/Common/my_pool`. If you create a pool in partition `App_A`, in folder `/App_A`, then the full name of the pool is `/App_A/my_pool`.

Object referencing between folders

When you create two BIG-IP® objects in separate folders, where one object references the other, the referenced object must reside in `/Common` or a sub-folder of `/Common`.

For example, if you create a virtual server in folder `/App_B`, and the virtual server references a load balancing pool, the pool object must reside in folder `/Common` or in a sub-folder of `/Common`.

Users

Purpose of BIG-IP user accounts

An important part of managing the BIG-IP® system is creating and managing user accounts for BIG-IP system administrators. By creating user accounts for system administrators, you provide additional layers of security. User accounts ensure that the system:

- Verifies the identity of users logging into the system (authentication)
- Controls user access to system resources (authorization)

To enable user authentication and authorization, you assign passwords and user roles to your user accounts. Passwords allow you to authenticate your users when they attempt to log in to the BIG-IP system. User roles allow you to control user access to BIG-IP system resources.

You can create and store BIG-IP administrative accounts either locally on the BIG-IP system, or remotely on a separate authentication server.

User account types

There are two types of user accounts on the BIG-IP® system: The system maintenance account and a set of standard user accounts.

The system maintenance account

The *system maintenance account* is a user account that you maintain using the Setup utility. The name of the system maintenance account is `root`. This account resides locally on the BIG-IP system and grants full access to BIG-IP system resources. You configure and maintain this account using the Setup utility and the BIG-IP Configuration utility, respectively.

Standard user accounts

Standard user accounts are user accounts that you create for other BIG-IP system administrators to use. Standard user accounts can reside either locally on the BIG-IP system, or remotely on a remote authentication server. You create and maintain these accounts using the browser-based BIG-IP Configuration utility or the command line interface. Creating standard user accounts allows you to assign various user roles to those accounts as a way to control system administrator access to BIG-IP system resources. A special standard user account is the `admin` account, which automatically exists on any BIG-IP system.

Note: *Excluding the `admin` account, the entire set of standard user accounts that you create for BIG-IP system administrators must reside either locally on the BIG-IP system, or remotely on another type of authentication server.*

You are not required to have any user accounts other than the `root` and `admin` accounts, but F5 Networks® recommends that you create other user accounts, as a way to intelligently control administrator access to system resources.

What are user roles?

User roles are a means of controlling user access to BIG-IP® system resources. You assign a user role to each administrative user, and in so doing, you grant the user a set of permissions for accessing BIG-IP system resources.

The BIG-IP system offers several different user roles that you can choose from when assigning a role to an administrative user. A user role is a property of a user account. Each user role grants a different set of permissions. More specifically, a user role defines:

The resources that a user can manage

User roles define the types of resources, or objects, that a user can manage. For example, a user with the role of Operator can enable or disable nodes and pool members only. By contrast, a user with the Guest role cannot manage any BIG-IP system resources.

The tasks that a user can perform

For example, a user with the role of Operator can enable or disable nodes and pool members, but cannot create, modify, or delete them. Conversely, a user with the Manager role can perform all tasks related to partitioned objects (except for user accounts), including nodes and pool members.

Important: A role defines the type of objects that a user can manage and the tasks that a user can perform on those object types. A role does not define the set of specific, existing objects that the user can access.

User roles on the BIG-IP system

This table lists and describes the various user roles that you can assign to a user account.

User role	Description
Administrator	This role grants users complete access to all partitioned and non-partitioned objects on the system. In addition, accounts with the Administrator role can change their own passwords.
Resource Administrator	This role grants users complete access to all partitioned and non-partitioned objects on the system, except user account objects. In addition, accounts with the Resource Administrator role can change their own passwords.
User Manager	Users with the User Manager role that have access to all partitions can create, modify, delete, and view all user accounts except those that are assigned the Administrator role, or the User Manager role with different partition access. Accounts with the User Manager role that have access to all partitions can also change their own passwords. Users with the User Manager role that have access only to a single partition can create, modify, delete, and view only those user accounts that are in that partition and that have access to that partition only. For example, if your user account has a User Manager role and has access to Partition A only, then you can manage only those user accounts that both reside in and have

User role	Description
Manager	<p>access to Partition A only. User accounts with the <code>User Manager</code> role can change their own passwords.</p> <p>This role grants users permission to create, modify, and delete virtual servers, pools, pool members, nodes, custom profiles, custom monitors, and iRules®. These users can view all objects on the system and change their own passwords.</p>
Certificate Manager	This role grants users permission to manage device certificates and keys, as well as perform Federal Information Processing Standard (FIPS) operations.
iRule Manager	This role grants users permission to create, modify, and delete iRules. Users with this role cannot affect the way that an iRule is deployed. For example, a user with this role can create an iRule but cannot assign it to a virtual server or move the iRule from one virtual server to another. A user with this role can be assigned universal access to administrative partitions.
Application Editor	This role grants users permission to modify nodes, pools, pool members, and monitors. These users can view all objects on the system and change their own passwords.
Acceleration Policy Editor	This role allows users to view, create, modify, and delete all WebAccelerator™™ policy objects in all administrative partitions. Users can also view, create, update, and delete Web Acceleration profiles.
Firewall Manager	This role allows users complete access to all firewall rules and supporting objects, including rules in all contexts, address lists, port lists, and schedules; security logging profiles and supporting objects, including log publishers and destinations; IP intelligence and DoS profiles; association rights for all of the above security profiles to virtual servers; and DoS Device Configuration (the L2-L4 DoS protection configuration). Firewall Managers may be granted access on all partitions or a single partition. Since global and management port rules are defined in <code>Common</code> , only Firewall Managers with rights on <code>Common</code> are allowed to modify global and management port rules. Firewall Managers have no create, update, or delete rights to any other objects, but otherwise have the same read access as the Manager role. Notably, the Firewall Manager role has no permission to create, update, or delete non-network firewall configuration, including Application Security or Protocol Security policies.
Web Application Security Administrator	This role grants users access to Application Security Manager™™ security policy objects, in one or all administrative partitions. These users have read-only permission for these profile types: HTTP, FTP, and

User role	Description
	SMTP. These users have no access to other LTM objects, nor to any TMOS objects. They can, however, change their own passwords. With respect to security policy objects, this role is similar to the Administrator role. You can assign this role only when the BIG-IP system includes the Application Security Manager component.
Web Application Security Editor	These users have no access to other LTM objects, nor to any TMOS objects. They can, however, change their own passwords. You can assign this role only when the BIG-IP system includes the Application Security Manager component.
Operator	This role grants users permission to enable or disable nodes and pool members. These users can view all objects and change their own passwords.
Auditor	This role grants users permission to view all configuration data on the system, including logs and archives. Users with this role cannot create, modify, or delete any data, nor can they view SSL keys or user passwords.
Guest	This role grants users permission to view all objects on the system except for sensitive data such as logs and archives. Users with this role can change their own passwords.
No Access	This role prevents users from accessing the system.

Default user roles

The BIG-IP® system automatically assigns a user role to an account when you create that account. The user role that the system assigns to a user account by default depends on the type of account:

root and admin accounts

The BIG-IP system automatically assigns the Administrator user role to the system maintenance root account and the `admin` account. You cannot change this user-role assignment. Thus, any user who successfully logs into the BIG-IP system using the `root` or `admin` account has full access to system resources and can perform all administrative tasks.

Other user accounts

The BIG-IP system automatically assigns the No Access user role to all standard user accounts other than the `root` and `admin` accounts. If the user account you are using has the Administrator role assigned to it, you are allowed to change another account's user role from the default No Access role to any other user role, including Administrator. For remote user accounts, if you know that most of your administrative users need some amount of access to BIG-IP system resources, you can configure the BIG-IP system to use a role other than No Access as the default user role.

Administrative partitions

When you create configurable objects for the BIG-IP® system, you have the option of putting those objects into administrative partitions. An *administrative partition* is a logical container of BIG-IP system objects such as virtual servers, pools, and monitors. When you first install the BIG-IP system, a default partition already exists named `Common`.

By putting objects into partitions, you establish a finer granularity of access control. Rather than having control over all resources on the BIG-IP system or no resources whatsoever, users with certain permissions can control resources within a designated partition only. For example, users with the role of Operator can mark nodes up or down, but can only mark those nodes that reside within their designated partition.

User accounts are another type of object that you can put into a partition. You put user accounts into administrative partitions strictly for the purpose of giving other users administrative access to those accounts. For example, you can put user accounts into `partition B`, and then assign a set of permissions (known as a user role) to user `Jane` so that she is allowed to modify user accounts in `partition B`.

Each user account on the BIG-IP system has a property known as Partition Access. The Partition Access property defines the partitions that the user can access. A user account can have access to either one partition or all partitions. Access to all partitions is known as *universal access*.

This figure shows how partition access can differ for different user accounts on the BIG-IP system.

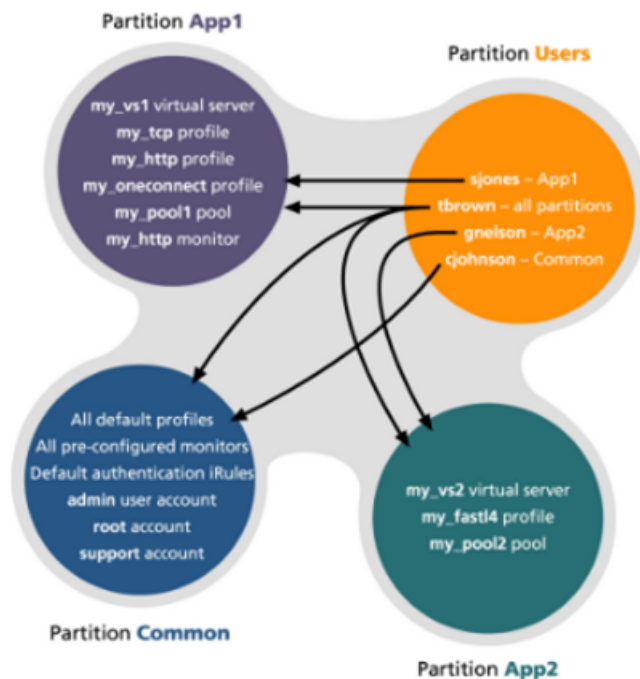


Figure 7: The Partition Access property for user accounts

In this example, the BIG-IP system objects reside in multiple partitions. Note that user accounts are also a type of BIG-IP system object, and as such, reside in a partition named `Users`. (Although you are not required to group user accounts together in a separate partition, for security purposes F5 Networks® highly recommends that you do so.)

To continue with the example, each user account in partition `Users` has access to specific, but different, partitions. Note that user accounts `sjones`, `cjohnson`, and `gnelson` can access one partition only, while the `tbrown` account has universal access.

To summarize, an administrative partition defines a set of objects, including user accounts, that other administrative users can potentially manage. This gives computing organizations greater control over user access to specific objects on the BIG-IP system.

Effect of user roles on objects within partitions

A user role defines the access level that a user has for each object in the user's assigned partition. An *access level* refers to the type of task that a user can perform on an object. Possible access levels are:

Write

Grants full access: that is, the ability to create, modify, enable and disable, and delete an object.

Update

Grants the ability to modify, enable, and disable an object.

Enable/disable

Grants the ability to enable or disable an object.

Read

Grants the ability to view an object.

Local user account management

Managing local user accounts refers to the tasks of creating, viewing, modifying, and deleting user accounts that reside on the BIG-IP® system, using the browser-based BIG-IP Configuration utility.

The BIG-IP Configuration utility stores local user accounts (including user names, passwords, and user roles) in a local user-account database. When a user logs in to the BIG-IP system using one of these locally-stored accounts, the BIG-IP system checks the account to determine the user role assigned to that user account.

Important: *Only users with the role of Administrator or User Manager can create and manage local user accounts. However, users with any role can change their own passwords. Also, if a user with a local user account is logged in to the BIG-IP system, and you subsequently switch the system from local authentication to remote authentication, the local user remains authenticated until the user's login session terminates.*

Admin account configuration

A user account called `admin` resides on every BIG-IP® system. Although the BIG-IP system creates this account automatically, you must still assign a password to the account before you can use it. To initially set the password for the `admin` account, you must run the Setup utility. To change its password later, you use the BIG-IP Configuration utility's Users screens.

The `admin` account resides in the local user account database on the BIG-IP system. By default, the BIG-IP system assigns the Administrator user role, which gives the user of this account full access to all BIG-IP system resources. You cannot change the user role on this account.

About secure password policy configuration

The BIG-IP® system includes an optional administrative feature: a security policy for creating passwords for local BIG-IP system user accounts. A secure password policy ensures that BIG-IP system users who have local user accounts create and maintain passwords that are as secure as possible.

The secure password policy feature includes two distinct types of password restrictions:

Enforcement restrictions

These are, specifically, character restrictions that you can enable or disable. They consist of the minimum password length and the required character types (numeric, uppercase, lowercase, and other kinds of characters). When enabled, the BIG-IP system never enforces restrictions on user accounts that have the Administrator role assigned to them. Consequently, a user with Administrator permissions does not need to adhere to these restrictions when either changing his or her own password, or changing the passwords of other user accounts.

Policy restrictions

These restrictions represent the minimum and maximum lengths of time that passwords can be in effect. Also included in this type of policy restriction are the number of days prior to password expiration that users are warned, and the number of previous passwords that the BIG-IP system should store, to prevent users from re-using former passwords. These restrictions are always enabled, although using the default values provides a minimal amount of restriction.

Note: The value of the **Maximum Duration** setting determines when users receive warning messages to change their passwords. If you change the value of this setting, any subsequent warning messages that users receive indicate the previous maximum duration value, rather than the new value. Once a user changes the password, however, subsequent reminder messages show the new value.

The password policy feature affects passwords for local user accounts only. Passwords for remotely-stored user accounts are not subject to this local password policy, but might be subject to a separate password policy defined on the remote system.

Important: You must have the user role of Administrator assigned to your account to configure this feature.

Configuration settings for a secure password policy

This table lists and describes the settings for a password policy.

Setting	Description	Default value
Secure Password Enforcements	Enables or disables character restrictions, that is, a policy for minimum password length and required characters. When you enable this setting, the BIG-IP Configuration utility displays the Minimum Length and Required Characters settings.	Disabled
Minimum Length	Specifies the minimum number of characters required for a password, and the allowed range of values is 6 to 255 . This setting appears only when you enable the Secure Password Enforcement setting.	6
Required Characters	Specifies the number of numeric, uppercase, lowercase, and other characters required for a password. The allowed range of values is 0 to 127 . This setting appears only when you enable the Secure Password Enforcement setting.	0

Setting	Description	Default value
Password Memory	Specifies, for each user account, the number of former passwords that the BIG-IP system retains to prevent the user from re-using a recent password. The range of allowed values is 0 to 127 .	0
Minimum Duration	Specifies the minimum number of days before a user can change a password. The range of allowed values is 0 to 255 .	0
Maximum Duration	Specifies the maximum number of days that a user's password can be valid. The range of allowed values is 1 to 99999 . This setting applies to all user accounts.	99999
Expiration Warning	Specifies the number of days prior to password expiration that the system sends a warning message to a user. The range of allowed values is 1 to 255 . This setting applies to all user accounts.	7
Maximum Login Failures	Denies access to a user after the specified number of failed authentication attempts. The administrator can then reset the lock to re-enable access for the user.	0

Configuring a password policy for administrative users

Use this procedure to require BIG-IP® system users to create strong passwords and to specify the maximum number of BIG-IP login failures that the system allows before the user is denied access.

1. On the Main tab, click **System > Users**.
2. On the menu bar, click **Authentication**.
3. From the **Secure Password Enforcement** list, select **Enabled**.
Additional settings appear on the screen.
4. For the **Minimum Length** and **Required Characters** settings, configure the default values, according to your organization's internal security requirements.
5. In the **Maximum Login Failures** field, specify a number.
If the user fails to log in the specified number of times, the user is locked out of the system. Therefore, F5 Networks recommends that you specify a value that allows for a reasonable number of login failures before user logout.
6. Click **Update**.

User authentication logout

You can deny access to a user after a specified number of failed authentication attempts. You can then reset the lock to re-enable access for the user.

To set the maximum number of failures before user logout, use the BIG-IP® Configuration utility to locate the Users screen, and then navigate to the Authentication screen. You can then specify a value for the **Maximum Login Failures** setting.

If a user becomes locked out, you can use the **Unlock** button on the User List screen to unlock the user.

Local user account creation

Properties of a local BIG-IP system user account

This table lists and describes the properties that define a local BIG-IP user account.

Property	Description	Default Value
User Name	Specifies the name of the user account. The BIG-IP system is case-sensitive, which means that names such as JONES and Jones are treated as separate user accounts.	No default value
Partition	When viewing the properties of an existing user account, displays the name of the partition in which the user account resides. All partitionable BIG-IP system objects (including user account objects) have the Partition property. Note that you cannot edit the value of this setting.	No default value
Password	Specifies a password that the user will use to log in to the BIG-IP system.	No default value
Role	Specifies the user role that you want to assign to the user account.	No Access
Partition Access	Specifies the partition to which the user has access when logged on to the BIG-IP system. If you have permission to do so, you can assign this value to a new user account, or change this value on an existing user account. This setting appears only when the user role for the account is not Administrator . (Accounts with the Administrator role always have universal partition access, that is, access to all partitions.)	All
Terminal Access	Specifies the level of access to the BIG-IP system command line interface. Possible values are: Disabled and Advanced shell . Users with the Administrator or Resource Administrator role assigned to their accounts can have advanced shell access, that is, permission to use all BIG-IP system command line utilities, as well as any Linux commands.	Disabled

Creating a local user account

You perform this task to create a local user account for BIG-IP administrative users. Only users who have been granted the Administrator or User Manager role can create user accounts. If the user role assigned to your account is Administrator, you can create a user account in any partition on the system. If the user role assigned to your account is User Manager, you can create a user account in any partition to which you have access.

Note: User accounts on the BIG-IP® system are case-sensitive. Thus, the system treats user accounts such as JONES and Jones as two separate user accounts. Note, however, that certain user names, such as admin, are reserved, and are therefore exempt from case-sensitivity. For example, you cannot create a user account named Admin, aDmin, or ADMIN.

1. On the Main tab, click **System > Users**.
2. Using the Partition list in the upper-left corner of the screen, select the name of the partition in which you want the new user account to reside.

Important: The partition you select in this step is not the partition to which you want the new user account to have access. Instead, this selection specifies the partition in which you want the new user account to reside. To grant partition access to a user, you configure the Partition Access property on the New User screen.

3. In the upper right corner of the screen, click **Create**. The New User screen opens.
If the **Create** button is unavailable, you do not have permission to create a local user account. You must have the Administrator or User Manager role assigned to your user account in order to create a local user account.
4. In the **User Name** box, type a name for the user account.
5. For the Password setting, type and confirm a password for the account.
6. To grant an access level other than No Access (the default value), use the Role setting and select a user role.
7. From the **Partition Access** list, select a partition name or **All**.

Note: For user accounts to which you assign the Administrator or Resource Administrator role, this setting is hidden because the value is automatically set to All. You cannot change the Partition Access setting for a user with the Administrator or Resource Administrator role.

8. If you want to allow user access to the command line interface, then from the **Terminal Access** list, select a level of access.

Note: The advanced shell is only available for accounts with the Administrator or Resource Administrator user role.

9. Click **Finished**.

Local user account view

Using the BIG-IP Configuration utility, you can easily display a list of existing local user accounts and view the properties of an individual account. Only users who have been granted the Administrator or User Manager roles can view the settings of other user accounts.

If the user role assigned to your account is Administrator, you can view any user account on the BIG-IP® system, in any partition. If the user role assigned to your account is User Manager, you can view any user account in any partition to which you have access on the BIG-IP system.

To summarize, depending on their own partition access, users with a User Manager role can do some or all of the following:

- Change another user's password
- Change another user's user role
- Change the partition in which the user can access objects (applies only to users who have both a User Manager role and access to all partitions)
- Enable or disable terminal access

Displaying a list of local user accounts

Using the BIG-IP Configuration utility, you can easily display a list of existing local user accounts and view the properties of an individual account. Only users who have been granted the Administrator or User Manager roles can view the settings of other user accounts.

If the user role assigned to your account is Administrator, you can view any user account on the BIG-IP® system, in any partition. If the user role assigned to your account is User Manager, you can view any user account in any partition to which you have access on the BIG-IP system.

To summarize, depending on their own partition access, users with a User Manager role can do some or all of the following:

- Change another user's password
 - Change another user's user role
 - Change the partition in which the user can access objects (applies only to users who have both a User Manager role and access to all partitions)
 - Enable or disable terminal access
1. On the Main tab, click **System > Users**.
 2. View the list of user accounts.

Viewing the properties of a local user account

Using the BIG-IP Configuration utility, you can easily display a list of existing local user accounts and view the properties of an individual account. Only users who have been granted the Administrator or User Manager roles can view the settings of other user accounts.

If the user role assigned to your account is Administrator, you can view any user account on the BIG-IP® system, in any partition. If the user role assigned to your account is User Manager, you can view any user account in any partition to which you have access on the BIG-IP system.

To summarize, depending on their own partition access, users with a User Manager role can do some or all of the following:

- Change another user's password
 - Change another user's user role
 - Change the partition in which the user can access objects (applies only to users who have both a User Manager role and access to all partitions)
 - Enable or disable terminal access
1. On the Main tab, click **System > Users**.
 2. In the user-account list, find the user account you want to view and click the account name. This displays the properties of that user account.

Local user account modification

You use the BIG-IP Configuration utility to modify the properties of any existing local user account, other than the `root` account. When modifying user accounts, consider the following:

- Only users who have been granted either the Administrator or User Manager role can modify user accounts other than their own account.
- A user with the User Manager role can modify only those accounts that reside in the partition to which that user has access. For example, if user `nelson` has a User Manager role and has access to partition `B` only, he can modify only those user accounts that reside in partition `B`. Even in this case, however, for user accounts in partition `B`, user `nelson` cannot modify a user's Partition Access property. If, however, user `nelson` has a User Manager role and has access to all partitions, he can modify all user accounts on the system. This includes changing another user's Partition Access property.
- Users with any role but No Access can modify their own user accounts to change the password. These users cannot modify any other properties of their own user accounts.

Note: When a user changes his own password, the system automatically logs the user off of the BIG-IP Configuration utility. The system then requires the user to use the new password for subsequent logins. This behavior applies even when the new password matches the old password.

- Users with the role of User Manager can modify all of the properties of their own user accounts, except their user role and partition access.

If you have an Administrator user role, you can also change some properties of the `root` account. Specifically, you can change the password of the `root` account, and you can enable or disable access to the BIG-IP[®] system through SSH.

Warning: The Administrator user role provides access to the BIG-IP system prompt. If a user with the Administrator user role is currently logged on to the system, and you change the user role to a role other than Administrator or Resource Administrator, the user can still run commands at the BIG-IP system prompt until he or she logs off of the system.

Modifying the properties of a local user account

You use the BIG-IP Configuration utility to modify the properties of any existing local user account, other than the `root` account. When modifying user accounts, consider the following:

- Only users who have been granted either the Administrator or User Manager role can modify user accounts other than their own account.
- A user with the User Manager role can modify only those accounts that reside in the partition to which that user has access. For example, if user `nelson` has a User Manager role and has access to partition `B` only, he can modify only those user accounts that reside in partition `B`. Even in this case, however, for user accounts in partition `B`, user `nelson` cannot modify a user's Partition Access property. If, however, user `nelson` has a User Manager role and has access to all partitions, he can modify all user accounts on the system. This includes changing another user's Partition Access property.
- Users with any role but No Access can modify their own user accounts to change the password. These users cannot modify any other properties of their own user accounts.

Note: When a user changes his own password, the system automatically logs the user out of the BIG-IP Configuration utility. The system then requires the user to use the new password for subsequent logins. This behavior applies even when the new password matches the old password.

- Users with the role of User Manager can modify all of the properties of their own user accounts, except their user role and partition access.
1. On the Main tab, click **System > Users**.
 2. In the user-account list, click a user account name. This displays the properties of that account.
 3. Change one or more of these settings: Password, role, partition access, or terminal access.
 4. Click **Update**.

Warning: The Administrator user role provides access to the BIG-IP system command prompt. If a user with the Administrator user role is currently logged on to the system, and you change the user role to a role other than Administrator or Resource Administrator, the user can still run commands at the BIG-IP system prompt until he or she logs off of the system.

Modifying the properties of the root account

If you have an Administrator user role, you can also change some properties of the `root` account. Specifically, you can change the password of the `root` account, and you can enable or disable access to the BIG-IP® system through SSH.

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
2. For the **Root Account** setting, type a new password in the **Password** box, and re-type the new password in the **Confirm** box.
3. If you want to grant SSH access, then for the **SSH Access** setting, select the **Enabled** checkbox, and for the **SSH IP Allow** setting, either:
 - Select *** All Addresses**
 - Select **Specify Range** and type a range of IP addresses.
4. Click **Update**.

Important: If you have a redundant system configuration and you change the password on the `admin` account, you must also change the password on the other device group members, to ensure that synchronization of configuration data operates correctly.

Delete local user accounts

If the account you are using has the Administrator or User Manager user role, you can delete other local user accounts. A user with the Administrator role can delete any user account on the BIG-IP® system in any partition. A user with the User Manager role can delete user accounts on the BIG-IP system in only those partitions to which she has access.

When you delete a local user account, you remove it permanently from the local user-account database on the BIG-IP system.

Note: You cannot delete the `admin` user account, nor can you delete the user account with which you are logged in.

Warning: The Administrator user role provides access to the BIG-IP system prompt. If a user with the Administrator user role is currently logged in to the system and you delete the user account, the user can still run commands at the BIG-IP system command prompt until he or she logs off of the system.

Deleting a local user account

When you delete a local user account, you remove it permanently from the local user-account database on the BIG-IP system. If the account you are using has the Administrator or User Manager user role, you can delete other local user accounts. A user with the Administrator role can delete any user account on the BIG-IP® system in any partition. A user with the User Manager role can delete user accounts on the BIG-IP system in only those partitions to which she has access.

Note: You cannot delete the `admin` user account, nor can you delete the user account with which you are logged in.

Warning: The Administrator user role provides access to the BIG-IP system prompt. If a user with the Administrator user role is currently logged in to the system and you delete the user account, the user can still run commands at the BIG-IP system command prompt until he or she logs off of the system.

1. On the Main tab, click **System > Users**.
2. In the user-account list, locate the name of the account you want to delete and select the checkbox to the left of the account name.
3. Click the **Delete** button.
4. Click **Delete** again.

Remote user account management

Rather than store user accounts locally on the BIG-IP[®] system, you can store them on a remote authentication server. In this case, you create all of your standard user accounts (including user names and passwords) on that remote server, using the mechanism supplied by that server's vendor.

Once you have created each user account on the remote server, you can then use the BIG-IP system to assign authorization properties (user role, partition access, and terminal access) for each account, for the purpose of controlling user access to BIG-IP system resources.

Important: You can assign authorization properties to remotely-stored user accounts on a group basis. You can then use the single configuration file (SCF) feature to propagate those properties to other BIG-IP devices on the network.

The BIG-IP Configuration utility stores all local and remote access control information in the BIG-IP system's local user-account database. When a user whose account information is stored remotely logs into the BIG-IP system and is granted authentication, the BIG-IP system then checks its local database to determine the access control properties that you assigned to that user.

Note: The BIG-IP Configuration utility refers to remote user accounts as external users. An external user is any user account that is stored on a remote authentication server.

Important: Only users with the role of Administrator can manage user roles for remote user accounts. Also, if a user with a local user account is logged on to the BIG-IP system, and you subsequently switch the system from local authentication to remote authentication, the local user remains authenticated until the user's login session terminates.

Introduction to remote BIG-IP user accounts

Each BIG-IP system requires one or more administrative user accounts. Rather than store these BIG-IP user accounts locally on the BIG-IP system, you can store BIG-IP user accounts on a remote authentication server. In this case, you create all of your standard BIG-IP user accounts (including user names and passwords) on that remote server, using the mechanism supplied by that server's vendor. The remote server then performs all authentication of those user accounts.

One of the tasks you perform with the BIG-IP Configuration utility is to specify the type of remote user-account server that currently stores your remote user accounts. The available server types that you can specify are:

- Active Directory or Lightweight Directory Access Protocol (LDAP)

- Remote Authentication Dial-In User Service (RADIUS)
- Terminal Access Controller Access-Control System Plus (TACACS+)

To ensure easy management of remotely-stored user accounts, the BIG-IP system automatically creates a single user account named `Other External Users`. This user account represents all of the remotely-stored BIG-IP user accounts that conform to the default access-control properties defined on the BIG-IP system.

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 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 - 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. 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.

11. 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.
12. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
13. 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 <code>tmsh</code> access to the BIG-IP system.

14. Click **Finished**.

You can now authenticate administrative traffic for user accounts that are stored on a remote LDAP or Active Directory 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. 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.
8. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.
9. 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 <code>tmsh</code> access to the BIG-IP system.

10. 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 group-based user authorization, 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 **Servers** setting, type an IP address for the remote TACACS+ server.
The route domain to which this address pertains must be route domain 0.
6. Click **Add**.
The IP address for the remote TACACS+ server appears in the **Servers** list.
7. 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.

8. In the **Confirm Secret** field, re-type the TACACS+ secret.

9. 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.

10. 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`.

11. 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`.

12. 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.

13. From the **Partition Access** list, select the default administrative partition that all remotely-authenticated BIG-IP system user accounts can access.

14. 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 <code>tmsh</code> access to the BIG-IP system.

15. 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 group-based user authorization, your configuration tasks are complete.

Effects of remote user authorization on user accounts

You can sometimes inadvertently affect your own user account, if the BIG-IP system is configured to perform remote user authentication, and you or another system administrator changes the default role or partition assigned to all external user accounts:

- If you log in to the BIG-IP system using one of these remotely-authenticated Administrator accounts, and you or another Administrator user modifies the default role of all external accounts from Administrator to a lesser role, the system modifies the user role of your own account to the lesser role. However, the change to your own account does not actually occur until you log out and log in again to the BIG-IP system.

- Similarly, your user account can be affected if the BIG-IP system is configured to perform remote user authentication, and the default partition assigned to all external user accounts is a specific partition. In this case, if you are logged on to the BIG-IP system through the command line using one of the remotely-authenticated accounts, and another user who is logged on through the BIG-IP Configuration utility modifies the default partition for external users, the BIG-IP system immediately logs you out when you attempt to issue another command.
- When you specify the type of remote server, you can also configure some server settings. For example, you can specify the user role you would like the BIG-IP® system to assign to a remote account if you do not explicitly assign one.
- Once you have configured the remote server, if you want any of the remote accounts to have a non-default user role, you can explicitly assign a user role to those accounts.
- If the remote authentication server is an Active Directory or LDAP server and is set up to authenticate SSL traffic, there is an additional feature that you can enable. You can configure the BIG-IP system to perform the server-side SSL handshake that the remote server would normally perform when authenticating client traffic. In this case, there are some preliminary steps you must perform to prepare for remote authentication using SSL.

Important: *If a BIG-IP system administrator changes the user role or partition assignment (or both) for any remote user account, the BIG-IP system logs out all users immediately. (A remote user account in this case refers to Other External Users.)*

Changing the default remote-account authorization

1. On the Main tab, click **System > Users > Authentication**.
2. Click **Change**.
3. From the **User Directory** list, select **Remote - Active Directory**, **Remote - LDAP**, **Remote - RADIUS**, or **Remote - TACACS+**.
4. From the **Role** list, select a user role.
The BIG-IP system assigns this user role to any remote account to which you have not explicitly assigned a user role.
5. From the **Partition Access** list, select a partition name.
6. From the **Terminal Access** list, select **Enabled** or **Disabled**.
7. Click **Update**.

Authorization for group-based remote user accounts

If you want to assign the same non-default access control properties to a group of remotely-stored user accounts, you can use the remote role groups feature. This feature stores all access control information on a group-wide basis for remotely-stored user accounts.

After using the remote role groups feature, you can propagate that access-control information to all BIG-IP® devices on the network, using the single configuration file (SCF) feature. The remote role groups feature, combined with the SCF feature, removes the need to manually assign access control properties to each individual BIG-IP user within a group, on each BIG-IP device on your network.

You can access the remote role groups feature by logging into the BIG-IP® Configuration utility and navigating to **System > Users > Remote Role Groups**.

Values for the remote role variable

This table lists the allowed values for a variable that you use for defining a role for a remotely-stored user account.

User Role	Value
Administrator	0
Resource Administrator	
User Manager	40
Manager	100
Application Editor	300
Operator	400
Guest	700
Application Security Policy Editor	800
No Access	900

About viewing remote user accounts

Using the BIG-IP Configuration utility, you can display a list of those remote user accounts to which you explicitly assigned a non-default user role. If a remote user account has the default role assigned to it, you cannot see that account in the user account list.

Any users who have access to a partition in which remote accounts reside can view a list of remote user accounts.

Displaying a list of remote user accounts with non-default user roles

1. On the Main tab, click **System > Users**.
2. On the menu bar, click **Authentication**.
3. Verify that the **User Directory** setting specifies a remote authentication server type (Active Directory, LDAP, or RADIUS).
4. On the menu bar, click **User List**.
5. View the list of user accounts. Remote user accounts that are assigned the default user role appear as **Other External Users**.

Viewing the properties of a remote user account

1. On the Main tab, click **System > Users**.
2. On the menu bar, click **Authentication**.
3. Verify that the **User Directory** setting specifies a remote authentication server type (Active Directory, LDAP, or RADIUS).
4. On the menu bar, click **User List**.

5. View the list of user accounts. Remote user accounts that are assigned the default user role appear as **Other External Users**.
6. In the user-account list, find the user account you want to view and click the account name. This displays the properties of that user account.

***Note:** Note: The only properties displayed for a remote user account are the account name, the user role assigned to the account, the account's partition access, and the account's terminal access.*

About auditing user access to the system

The BIG-IP® system generates a log message whenever a user or an application attempts to log in to or log out of the system. The system logs both successful and unsuccessful login attempts. The system stores these log messages in the `/var/log/secure` file.

When the system logs an authentication message in the `/var/log/secure` file, the message can contain the following types of information:

- The connecting user's ID
- The IP address or host name of the user's interface
- The time of each login attempt
- Successful login attempts for command line interface sessions only
- Failed login attempts for command line interface, BIG-IP Configuration utility, and iControl® sessions
- The time of the logout for command line interface sessions only

This is an example of log messages for both successful and failed login attempts made by user `jsmith`.

```
May 10 16:25:25 jsmith-dev sshd[13272]: pam_audit: user: jsmith(jsmith) from:
/dev/pts/10 at jsmith-dev attempts: 1 in:
[Thu May 10 16:25:23 2007 ] out: [Thu May 10 16:25:25 2007 ]
May 10 16:14:56 jsmith-dev sshd[716]: pam_audit: User jsmith from ssh at
jsmith-dev failed to login after 1 attempts
(start: [Thu May 10 16:14:53 2007 ] end: [Thu May 10 16:14:56 2007 ]).
```


Logging

BIG-IP system logging overview

Viewing and managing log messages is an important part of managing traffic on a network and maintaining a BIG-IP® system. Log messages inform you on a regular basis of the events that are happening on the system.

You can log events either locally on the BIG-IP system or remotely, using The BIG-IP system's high-speed logging mechanism. The recommended way to store logs is on a pool of remote logging servers.

For local logging, the high-speed logging mechanism stores the logs in either the Syslog or the MySQL database on the BIG-IP system, depending on a destination that you define. For remote logging, the high-speed logging mechanism sends log messages to a pool of logging servers that you define.

Types of log messages

Examples of the types of messages that the high-speed logging mechanism can log are:

- BIG-IP® system-level events
- DNS events (for local traffic and global traffic)
- Network Firewall events
- Protocol Security events
- Carrier-grade NAT (CGNAT) events
- Denial of Service (DoS) protection events

Existing Syslog configurations

If you previously configured the BIG-IP® system to log messages locally using the Syslog utility or remotely using the Syslog-ng utility, you can continue doing so with your current logging configuration, without configuring high-speed logging.

Alternatively, however, you can configure local Syslog logging using the high-speed logging mechanism, which is the recommended Syslog configuration. By configuring Syslog using high-speed logging, you can easily switch logging utilities in the future as needs change, without having to perform significant re-configuration.

Remote storage of log messages

The way that you set up remote, high-speed logging is by first defining a pool of logging servers, and then creating an unformatted, remote high-speed log destination that references the pool. If you are using ArcSight,

Splunk, or Remote Syslog logging servers that require a formatted destination, you can also create a formatted log destination for one of those server types. Once those objects are set up, you create a publisher and a custom logging profile pertaining to the type of message you want to log. You then assign the logging profile to a relevant virtual server, and the profile, in turn, references the publisher.

This image shows the BIG-IP® objects that you configure for remote high-speed logging. This figure shows the way that these objects reference one another from a configuration perspective.




Figure 8: BIG-IP object referencing for remote high-speed logging

For an example of configuring remote, high-speed logging, suppose you want to send all Protocol Security messages to a group of remote ArcSight servers. In this case, you would create these objects:

- A load balancing pool for the ArcSight logging servers.
- An unformatted Remote High-Speed Log destination that references the pool of ArcSight logging servers.
- A formatted ArcSight log destination that references an unformatted log destination.
- A publisher that references the formatted and unformatted log destinations.
- A Protocol Security logging profile that references the publisher.
- An LTM® virtual server or GTM™ listener that references the logging profile and the load balancing pool.
- An unformatted Remote High-Speed Log destination that references the pool of ArcSight logging servers.

Local storage of log messages

Although local logging is not recommended, you can store log messages locally on the BIG-IP® system instead of remotely. In this case, you can still use the high-speed logging mechanism to store and view log messages locally on the BIG-IP system.

When you use the high-speed logging mechanism to configure local logging, the system stores the log messages in either the local Syslog data base or the local MySQL data base. The storage database that the BIG-IP system chooses depends on the specific log destination you assign to the publisher:

local-syslog

Causes the system to store log messages in the local Syslog database. When you choose this log destination, the BIG-IP Configuration utility displays the log messages in these categories: System, Local Traffic, Global Traffic, and Audit.

local-db

Causes the system to store log messages in the local MySQL database. When you choose `local-db`, the BIG-IP Configuration utility does not display the log messages.

Log level settings for BIG-IP system events

For each type of system-level process, such as bigdb configuration events or events related to HTTP compression, you can set a minimum log level. The minimum log level indicates the minimum severity level at which the BIG-IP® system logs that type of event. There are many different types of local traffic or global traffic events for which you can set a minimum log level.

The log levels that you can set on certain types of events, ordered from highest severity to lowest severity, are:

- Emergency
- Alert
- Critical
- Error
- Warning
- Notice
- Informational
- Debug

For example, if you set the minimum log level for bigdb events to Error, then the system only logs messages that have a severity of Error or higher for those events.

About local Syslog logging

If you are using the Syslog utility for local logging, whether or not you are using the high-speed logging mechanism you can view and manage the log messages, using the BIG-IP® Configuration utility.

The local Syslog logs that the BIG-IP system can generate include several types of information. For example, some logs show a timestamp, host name, and service for each event. Moreover, logs sometimes include a status code, while the audit log shows a user name and a transaction ID corresponding to each configuration change. All logs contain a one-line description of each event.

For local log messages that the BIG-IP system stores in the local Syslog data base, the BIG-IP system automatically stores and displays log messages in these categories:

- System messages
- Packet filter messages
- Local Traffic messages
- Global Traffic messages
- BIG-IP system configuration (audit) messages

Each type of event is stored locally in a separate log file, and the information stored in each log file varies depending on the event type. All log files for these event types are in the directory `/var/log`.

Logging system events

Many events that occur on the BIG-IP® system are Linux-related events, and do not specifically apply to the BIG-IP system. Using the BIG-IP Configuration utility, you can display these local system messages.

Logging packet filter events

Some of the events that the BIG-IP system logs are related to packet filtering. The system logs the messages for these events in the file `/var/log/pktfilter`.

Logging local traffic events

Many of the events that the BIG-IP system logs are related to local area traffic passing through the BIG-IP system. The BIG-IP system logs the messages for these events in the file `/var/log/audit`.

Logging BIG-IP system configuration changes (audit logging)

Audit logging is an optional feature that logs messages whenever a BIG-IP® system object, such as a virtual server or a load balancing pool, is configured (that is, created, modified, or deleted). The BIG-IP system logs the messages for these auditing events in the file `/var/log/audit`.

There are three ways that objects can be configured:

- By user action
- By system action
- By loading configuration data

Whenever an object is configured in one of these ways, the BIG-IP system logs a message to the audit log.

Code expansion in Syslog log messages

The BIG-IP® system log messages contain codes that provide information about the system. You can run the Linux `zcat` command at the command prompt to expand the codes in log messages to provide more information. In this example, the bold text is the expansion of the log code `012c0012`.

```
Jun 14 14:28:03 sccp bcm56xxd [ 226 ] : 012c0012 : (Product=BIGIP  
Subset=BCM565XXD) : 6: 4.1 rx [ OK 171009 Bad 0 ] tx [ OK 171014 Bad 0 ]
```

About enabling and disabling auditing logging

An optional type of logging that you can enable is audit logging. *Audit logging* logs messages that pertain to configuration changes that users or services make to the BIG-IP® system configuration. This type of audit logging is known as *MCP audit logging*. Optionally, you can set up audit logging for any `tmsh` commands that users type on the command line.

For both MCP and `tmsh` audit logging, you can choose a log level. In this case, the log levels do not affect the severity of the log messages; instead, they affect the initiator of the audit event.

The log levels for MCP logging are:

Disable

This turns audit logging off. This is the default value.

Enable

This causes the system to log messages for user-initiated configuration changes only.

Verbose

This causes the system to log messages for user-initiated configuration changes and any loading of configuration data.

Debug

This causes the system to log messages for all user-initiated and system-initiated configuration changes.

The log levels for `tmsh` logging are:

Disable

This turns audit logging off. This is the default value.

Enable

This causes the system to log messages for user-initiated configuration changes only.

About remote logging using Syslog-ng

If you want to configure remote logging using Syslog-ng, you do not use the high-speed logging mechanism. Configuration of remote logging using Syslog-ng has some key differences compared to a remote, high-speed logging configuration:

- You do not configure log destinations, publishers, or a logging profile or log filter.
- Instead of creating a pool of remote logging servers (as you do with high-speed logging), you specify the IP addresses of the servers using the Remote Logging screen of the BIG-IP® Configuration utility.
- If you want to ensure that the Syslog-ng messages being logged remotely are encrypted, you must first establish a secure tunnel.

Interface Concepts

Introduction to BIG-IP system interfaces

A key task of the BIG-IP® system configuration is the configuration of BIG-IP system interfaces. The interfaces on a BIG-IP system are the physical ports that you use to connect the BIG-IP system to other devices on the network. These other devices can be next-hop routers, Layer 2 devices, destination servers, and so on. Through its interfaces, the BIG-IP system can forward traffic to or from other network devices.

Note: The term *interface* refers to the physical ports on the BIG-IP system.

Every BIG-IP system includes multiple interfaces. The exact number of interfaces that you have on the BIG-IP system depends on the platform type.

A BIG-IP system has two types of interfaces:

A management interface

The *management interface* is a special interface dedicated to performing a specific set of system management functions.

TMM switch interfaces

TMM switch interfaces are those interfaces that the BIG-IP system uses to send or receive application traffic, that is, traffic slated for application delivery.

Each of the interfaces on the BIG-IP system has unique properties, such as the MAC address, media speed, duplex mode, and support for Link Layer Discovery Protocol (LLDP).

In addition to configuring interface properties, you can implement a feature known as *interface mirroring*, which you can use to duplicate traffic from one or more interfaces to another. You can also view statistics about the traffic on each interface.

Once you have configured the properties of each interface, you can configure several other features of the BIG-IP system that control the way that interfaces operate. For example, by creating a virtual local area network (VLAN) and assigning interfaces to it, the BIG-IP system can insert a VLAN ID, or tag, into frames passing through those interfaces. In this way, a single interface can forward traffic for multiple VLANs.

About 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 enables a network device such as the BIG-IP system 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 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 `tmssh`, you can configure the BIG-IP system interfaces to transmit or receive LLDPDU s. More specifically, you can:

- Specify the exact content of LLDPDU s 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 each interface can receive messages. These properties apply to all interfaces on the BIG-IP system.

This figure shows a local LLDP-enabled BIG-IP system, configured to both transmit and receive LLDP messages from neighbor devices on a LAN.




Figure 9: A local BIG-IP system that transmits and receives LLDPDU s

Interface properties

Each interface on the BIG-IP[®] system has a set of properties that you can configure, such as enabling or disabling the interface, setting the requested media type and duplex mode, and configuring flow control. Configuring the properties of each interface is one of the first tasks you do after running the Setup utility on the BIG-IP system. While you can change some of these properties, such as media speed and duplex mode, you cannot change other properties, such as the media access control (MAC) address.

Note: You can configure STP-related properties on an interface by configuring one of the Spanning Tree protocols.

Before configuring interface properties, it is helpful to understand interface naming conventions. Only users with either the Administrator or Resource Administrator user role can create and manage interfaces.

Interface naming conventions

By convention, the names of the interfaces on the BIG-IP[®] system use the format `<s>.<p>` where `s` is the slot number of the network interface card (NIC), and `p` is the port number on the NIC. Examples of interface names are `1.1`, `1.2`, and `2.1`. BIG-IP system interfaces already have names assigned to them; you do not explicitly assign them.

An exception to the interface naming convention is the management interface, which has the special name, MGMT.

Viewing interface information and media properties

Using the BIG-IP Configuration utility, you can display a screen that lists all of the BIG-IP® system interfaces, as well as their current status (UP or DOWN). You can also view other information about each interface:

- MAC address of the interface
- Interface availability
- Media type
- Media speed
- Active mode (such as full)

This information is useful when you want to assess the way that a particular interface is forwarding traffic. For example, you can use this information to determine the specific VLANs for which an interface is currently forwarding traffic. You can also use this information to determine the speed at which an interface is currently operating.

Interface state

You can either enable or disable an interface on the BIG-IP® system. By default, each interface is set to Enabled, where it can accept ingress or egress traffic. When you set the interface to Disabled, the interface cannot accept ingress or egress traffic.

Fixed Requested Media

The Fixed Requested Media property shows that the interface auto-detects the duplex mode of the interface.

About flow control

You can configure the Flow Control property to manage the way that an interface handles pause frames for flow control. *Pause frames* are frames that an interface sends to a peer interface as a way to control frame transmission from that peer interface. Pausing a peer's frame transmissions prevents an interface's First-in, First-out (FIFO) queue from filling up and resulting in a loss of data. Possible values for this property are:

Pause None

Disables flow control.

Pause TX/RX

Specifies that the interface honors pause frames from its peer, and also generates pause frames when necessary. This is the default value.

Pause TX

Specifies that the interface ignores pause frames from its peer, and generates pause frames when necessary.

Pause RX

Specifies that the interface honors pause frames from its peer, but does not generate pause frames.

About the LLDP property

The LLDP property is one of two properties related to LLDP that you can configure for a specific interface. The possible values for this setting are:

Disabled

When set to this value, the interface neither transmits (sends) LLDP messages to, nor receives LLDP messages from, neighboring devices.

Transmit Only

When set to this value, the interface transmits LLDP messages to neighbor devices but does not receive LLDP messages from neighbor devices.

Receive Only

When set to this value, the interface receives LLDP messages from neighbor devices but does not transmit LLDP messages to neighbor devices.

Transmit and Receive

When set to this value, the interface transmits LLDP messages to and receives LLDP messages from neighboring devices.

In addition to the LLDP-related settings that you can configure per interface, you can configure some global LLDP settings that apply to all interfaces on the system.

Moreover, you can view statistics pertaining to any neighbor devices that have transmitted LLDP messages to the local BIG-IP® system.

LLDP Attributes

The LLDP Attributes setting is one of two settings related to LLDP that you can configure for a specific interface. You use this interface setting to specify the content of an LLDP message being sent or received. Each LLDP attribute that you specify with this setting is optional and is in the form of Type, Length, Value (TLV).

About interface mirroring

For reliability reasons, you can configure a feature known as interface mirroring. When you configure *interface mirroring*, you cause the BIG-IP® system to copy the traffic on one or more interfaces to another interface that you specify. By default, the interface mirroring feature is disabled.

Neighbor settings

When a BIG-IP® system interface receives LLDP messages from neighbor devices, the BIG-IP system displays chassis, port, and system information about the content of those messages. Specifically, the system displays values for the standard TLVs for each neighbor. These TLVs are:

Chassis ID

Identifies the chassis containing the IEEE 802 LAN station associated with the transmitting LLDP agent.

Port ID

Identifies the port component of the media service access point (MSAP) identifier associated with the transmitting LLDP agent.

Port description

An alpha-numeric string that describes the interface.

System name

An alpha-numeric string that indicates the administratively-assigned name of the neighbor device.

System description

An alpha-numeric string that is the textual description of the network entity. The system description should include the full name and version identification of the hardware type, software operating system, and networking software of the neighbor device.

System capabilities

The primary functions of the system and whether these primary functions are enabled.

Management address

An address associated with the local LLDP agent used to reach higher layer entities. This TLV might also include the system interface number that is associated with the management address, if known.

Related configuration tasks

After you have configured the interfaces on the BIG-IP® system, one of the primary tasks you perform is to assign those interfaces to the virtual LANs (VLANs) that you create. A *VLAN* is a logical subset of hosts on a local area network (LAN) that reside in the same IP address space. When you assign multiple interfaces to a single VLAN, traffic destined for a host in that VLAN can travel through any one of these interfaces to reach its destination. Conversely, when you assign a single interface to multiple VLANs, the BIG-IP system can use that single interface for any traffic that is intended for hosts in those VLANs.

Another powerful feature that you can use for BIG-IP system interfaces is trunking, with link aggregation. A *trunk* is an object that logically groups physical interfaces together to increase bandwidth. Link aggregation, through the use of the industry-standard Link Aggregation Control Protocol (LACP), provides regular monitoring of link status, as well as failover if an interface becomes unavailable.

Finally, you can configure the BIG-IP system interfaces to work with one of the spanning tree protocols (STP, RSTP, and MSTP). *Spanning tree protocols* reduce traffic on your internal network by blocking duplicate routes to prevent bridging loops.

Self IP Addresses

Introduction to self IP addresses

A *self IP address* is an IP address on the BIG-IP® system that you associate with a VLAN, to access hosts in that VLAN. By virtue of its netmask, a self IP address represents an *address space*, that is, a range of IP addresses spanning the hosts in the VLAN, rather than a single host address. You can associate self IP addresses not only with VLANs, but also with VLAN groups.

Self IP addresses serve two purposes:

- First, when sending a message to a destination server, the BIG-IP system uses the self IP addresses of its VLANs to determine the specific VLAN in which a destination server resides. For example, if VLAN internal has a self IP address of 10.10.10.100, with a netmask of 255.255.255.0, and the destination server's IP address is 10.10.10.20 (with a netmask of 255.255.255.255), the BIG-IP system recognizes that the server's IP address falls within the range of VLAN internal's self IP address, and therefore sends the message to that VLAN. More specifically, the BIG-IP system sends the message to the interface that you assigned to that VLAN. If more than one interface is assigned to the VLAN, the BIG-IP system takes additional steps to determine the correct interface, such as checking the Layer2 forwarding table.
- Second, a self IP address can serve as the default route for each destination server in the corresponding VLAN. In this case, the self IP address of a VLAN appears as the destination IP address in the packet header when the server sends a response to the BIG-IP system.

You normally assign self IP addresses to a VLAN when you initially run the Setup utility on a BIG-IP system. More specifically, you assign one static self IP address and one floating self IP address to each of the default VLANs (internal and external). Later, using the BIG-IP Configuration utility, you can create self IP addresses for other VLANs that you create.

Self IP addresses reside in administrative partitions/folders and are associated with traffic groups. The self IP addresses that you create when you run the Setup utility reside in partition Common (that is folder /Common).

Types of self IP addresses

There are two types of self IP addresses that you can create:

- A *static self IP address* is an IP address that the BIG-IP® system does not share with another BIG-IP system. Any self IP address that you assign to the default traffic group `traffic-group-local-only` is a static self IP address.
- A *floating self IP address* is an IP address that two BIG-IP systems share. Any self IP address that you assign to the default traffic group `traffic-group-1` is a floating self IP address.

Self IP addresses and MAC addresses

For each self IP address that you create for a VLAN, the BIG-IP® system automatically assigns a media access control (MAC) address.

As an alternative, you can globally configure the BIG-IP system to assign the same MAC address to all VLANs. This feature is useful if your network includes a type of switch that does not keep a separate Layer 2 forwarding table for each VLAN on that switch.

Self IP addresses for SNATs

When you configure the BIG-IP® system to manage local area traffic, you can implement a feature known as a secure network address translation (SNAT). A *SNAT* is an object that causes the BIG-IP system to translate the original source IP address of a packet to an IP address that you specify. A SNAT ensures that the target server sends its response back through the BIG-IP system rather than to the original client IP address directly.

When you create a SNAT, you can configure the BIG-IP system to automatically choose a translation address. This ability of the BIG-IP system to automatically choose a translation address is known as *SNAT automapping*, and in this case, the translation address that the system chooses is always an existing self IP address. Thus, for traffic going from the BIG-IP system to a destination server, configuring SNAT automapping ensures that the source IP address in the header of a packet is a self IP address.

When you create an automapped SNAT, the BIG-IP system actually creates a SNAT pool consisting of the system's internal self IP addresses, and then uses an algorithm to select and assign an address from that SNAT pool.

Self IP address properties

It is when you initially run the Setup utility on a BIG-IP® system that you normally create any static and floating self IP addresses and assign them to VLANs. However, if you want to create additional self IP addresses later, you can do so using the BIG-IP Configuration utility.

Note: Only users with either the Administrator or Resource Administrator user role can create and manage self IP addresses.

Note: A self IP address can be in either IPv4 or IPv6 format.

IP address

A self IP address, combined with a netmask, typically represents a range of host IP addresses in a VLAN. If you are assigning a self IP address to a VLAN group, the self IP address represents the range of self IP addresses assigned to the VLANs in that group.

Netmask

When you specify a netmask for a self IP address, the self IP address can represent a range of IP addresses, rather than a single host address. For example, a self IP address of 10.0.0.100 can represent several host IP addresses if you specify a netmask of 255.255.0.0.

VLAN/Tunnel assignment

You assign a unique self IP address to a specific VLAN or a VLAN group:

Assigning a self IP address to a VLAN

The self IP address that you assign to a VLAN should represent an address space that includes the self IP addresses of the hosts that the VLAN contains. For example, if the address of one destination server in a VLAN is 10.0.0.1 and the address of another server in the VLAN is 10.0.0.2, you could assign a self IP address of 10.0.0.100, with a netmask of 255.255.0.0, to the VLAN.

Assigning a self IP address to a VLAN group

The self IP address that you assign to a VLAN group should represent an address space that includes the self IP addresses of the VLANs that you assigned to the group. For example, if the self IP address of one VLAN in a VLAN group is 10.0.20.100 and the address of the other VLAN in a VLAN group is 10.0.30.100, you could assign an address of 10.0.0.100, with a netmask of 255.255.0.0, to the VLAN group.

The VLAN/Tunnel list in the BIG-IP Configuration utility displays the names of all existing VLANs and VLAN groups.

Port lockdown

Each self IP address has a feature known as port lockdown. *Port lockdown* is a security feature that allows you to specify particular UDP and TCP protocols and services from which the self IP address can accept traffic.

You can determine the supported protocols and services by using the `tmsh` command `tmsh list net self-allow defaults`.

If you do not want to use the default setting (**Allow None**), you can configure port lockdown to allow either all UDP and TCP protocols and services (**Allow All**) or only those that you specify (**Allow Custom**).

Note: *High availability-related traffic from configured peer devices in a device group might not be subject to port lockdown settings.*

Traffic groups

If you want the self IP address to be a *floating IP address*, that is, an address shared between two or more BIG-IP devices in a device group, you can assign a floating traffic group to the self IP address. A floating traffic group causes the self IP address to become a floating self IP address.

A floating self IP address ensures that application traffic reaches its destination. More specifically, a floating self IP address enables a source node to successfully send a request, and a destination node to successfully send a response, when the relevant BIG-IP device is unavailable.

If you want the self IP address to be a static (non-floating) IP address (used mostly for standalone devices), you can assign a non-floating traffic group to the self IP address. A non-floating traffic group causes the self IP address to become a non-floating self IP address. An example of a non-floating self IP address is the address that you assign to the default VLAN named HA, which is used strictly to process failover communications between BIG-IP devices, instead of processing application traffic.

Packet Filters

Introduction to 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, using the BIG-IP Configuration utility. 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 BIG-IP system 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. For more information on the `tcpdump` utility, see the online man page for the `tcpdump` command.

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.

Global settings

Global settings for packet filtering are divided into two categories: Properties and Exemptions. The BIG-IP® system applies global settings to all packets coming into the BIG-IP system.

Important: *Note that one of the global settings, **Packet Filtering**, enables packet filtering. When you disable this setting, no packet filter settings or packet filter rules operate, and the BIG-IP system allows all traffic by default.*

Global properties

You can configure three specific global properties for packet filtering.

Packet filter enabling

Before you can implement packet filtering on the BIG-IP® system, you must enable the packet filter feature. You do this by changing the **Packet Filtering** setting to **Enabled**. The default setting for packet filtering is **Disabled**.

Control of unhandled packets

Sometimes a packet does not match any of the criteria that you have specified in the packet filter rules that you have created. For this reason, you must configure the *Unhandled Packet Action* property, which specifies the action that the BIG-IP system should take when the packet does not match packet filter rule criteria.

Possible values for this setting are **Accept**, **Discard**, and **Reject**. The default value is **Accept**.

Warning: *Changing the default value of the *Unhandled Packet Action* property can produce unwanted consequences. Before changing this value to **Discard** or **Reject**, make sure that any traffic that you want the BIG-IP system to accept meets the criteria specified in your packet filter rules.*

Other options

Using the Options property, you can configure two other options:

Filter established connections

When you enable (check) this option, the BIG-IP system filters all ingress packets, even if the packets are part of an existing connection. The default setting is disabled (unchecked). Note that checking this option does not typically enhance security, and can impact system performance.

Send ICMP error on packet reject

When you enable (check) this option, the system sends an ICMP type 3 (destination unreachable), code 13 (administratively prohibited) packet when an ingress packet is rejected. When you disable (clear) this option, the BIG-IP system sends an ICMP reject packet that is protocol-dependent. The default setting for this option is disabled (cleared).

Global exemptions

There are a number of exemptions you can set for packet filtering. When filtering packets, the BIG-IP[®] system always applies these exemptions, effectively overriding certain criteria you might have previously set within an individual packet filter rule.

VLANs

Using the **VLANs** setting, you can configure the BIG-IP system so that traffic from one or more specified VLANs is exempt from packet filtering. In this case, the system does not attempt to match packets from the specified VLAN or VLANs to any packet filter rule. Instead, the BIG-IP system always accepts traffic from the specified VLAN or VLANs.

For example, if you specify VLAN internal, then no incoming packets from VLAN internal are subject to packet filtering, even if a packet matches the criteria of a packet filter rule.

Possible values are:

Always Accept

When you select this value, a VLAN List setting appears. You can then specify one or more VLANs from which traffic should be exempt from packet filtering.

None

When you select this value, traffic from all VLANs is subject to packet filtering, according to existing packet filter rule criteria. This is the default value.

Protocols

With the **Protocols** setting, you can specify whether ARP and certain ICMP messages are exempt from packet filtering. The individual settings are:

Always accept ARP

When you enable (check) this setting, the system automatically accepts all ARP packets and therefore does not subject them to packet filtering. The default setting is enabled (checked).

Always accept important ICMP

When you enable (check) this setting, the system automatically accepts the following ICMP packet types for IPv4, and therefore does not subject them to packet filtering:

- UNREACH
- SOURCEQUENCH
- REDIRECT
- TIMEXCEED

The default setting is enabled.

MAC addresses

You can use the **MAC Addresses** setting to exempt traffic from certain MAC addresses from packet filtering. Possible values are:

Always Accept

When you select this value, a MAC Address List setting appears. You can then specify one or more MAC addresses from which traffic should be exempt from packet filtering.

None

When you select this value, traffic from all MAC addresses is subject to packet filtering, according to existing packet filter rule criteria. This is the default value.

IP addresses

You can use the **IP Addresses** setting to exempt traffic from certain IP addresses from packet filtering. Possible values are:

Always Accept

When you select this value, an IP Address List setting appears. You can then specify one or more IP addresses from which traffic should be exempt from packet filtering.

None

When you select this value, traffic from all IP addresses is subject to packet filtering, according to existing packet filter rule criteria. This is the default value.

VLANs

Using the **VLANs** setting, you can configure the BIG-IP® system so that traffic from one or more specified VLANs is exempt from packet filtering. In this case, the system does not attempt to match packets from

the specified VLAN or VLANs to any packet filter rule. Instead, the BIG-IP system always accepts traffic from the specified VLAN or VLANs.

For example, if you specify VLAN internal, then no incoming packets from VLAN internal are subject to packet filtering, even if a packet matches the criteria of a packet filter rule.

Possible values are:

Always Accept

When you select this value, a **VLAN List** setting appears. You can then specify one or more VLANs from which traffic should be exempt from packet filtering.

None

When you select this value, traffic from all VLANs is subject to packet filtering, according to existing packet filter rule criteria. This is the default value.

Order of packet filter rules

You use the **Order** setting to specify the order in which you want the BIG-IP® system to apply existing packet filter rules. This setting is required. Possible values for this setting are:

First

Select this value if you want this packet filter rule to be the first rule that the BIG-IP system applies.

Last

Select this value if you want this packet filter rule to be the last rule that the BIG-IP system applies.

After

Select this value, and then select a packet filter rule from the list, if you want the system to apply this packet filter after the packet filter that you select from the list. Note that this setting is most useful when you have more than three packet filter rules configured.

About the action setting in packet filter rules

When a packet matches the criteria that you have specified in a packet filter rule, the BIG-IP® system can take a specific action. You define this action using the **Action** setting. You can choose one of these actions:

Accept

Select **Accept** if you want the system to accept the packet, and stop processing additional packet filter rules, if any exist. This is the default setting.

Discard

Select **Discard** if you want the system to drop the packet, and stop processing additional packet filter rules, if any exist.

Reject

Select **Reject** if you want the system to drop the packet, and also send a rejection packet to the sender, indicating that the packet was refused. Note that the behavior of the system when you select the **Reject** action depends on how you configured the general packet filter Options property, Send ICMP Error on Packet Reject.

Continue

Select **Continue** if you simply want the system to acknowledge the packet for logging or statistical purposes. Setting the **Action** value to **Continue** does not affect the way that the BIG-IP system handles the packet; the system continues to evaluate traffic matching a rule, starting with the next packet filter rule listed.

Rate class assignment

Using the **Rate Class** setting, you can assign a rate class to traffic that matches the criteria defined in a packet filter rule. Note that this setting applies only when you have the rate shaping feature enabled.

The default value for this setting is None. If you previously created rate classes using the rate shaping feature, you can choose one of those rate classes from the **Rate Class** list.

One or more VLANs

You use the **Apply to VLAN** setting to display a list of VLANs and then select a VLAN or VLAN group name. Selecting a VLAN from the list means that the packet filter rule filters ingress traffic from that VLAN only. For example, if you select the value ***All VLANs**, the BIG-IP® system applies the packet filter rule to all traffic coming into the BIG-IP system.

Similarly, if you select the **VLAN internal**, the BIG-IP system applies the packet filter rule to traffic from VLAN internal only. The default value is ***All VLANs**.

If you select the name of a VLAN group instead of an individual VLAN, the packet filter rule applies to all VLANs in that VLAN group.

Logging

If you want to generate a log message each time a packet matches a rule, you can enable logging for the packet filter rule. With this configuration, you can then display the Logging screen in the BIG-IP Configuration utility and view events related to packet filtering.

About filter expression creation

To match incoming packets, the BIG-IP® system must use a filter expression. A *filter expression* specifies the criteria that you want the BIG-IP system to use when filtering packets. For example, the BIG-IP system can filter packets based on the source or destination IP address in the header of a packet.

Using the BIG-IP Configuration utility, you can create a filter expression in either of two ways:

- You can write your own expression, using a Filter Expression box.
- You can specify a set of criteria (such as source or destination IP addresses) that you want the BIG-IP system to use when filtering packets. When you use this method, the BIG-IP system builds a filter expression for you.

Packet Filters

You can have as many rules as you want, limited only by the available memory. Of course, the more statements you have, the more challenging it is to understand and maintain your packet filters.

Spanning Tree Protocol

Introduction to spanning tree protocols

On networks that contain redundant paths between Layer 2 devices, a common problem is bridging loops. Bridging loops occur because Layer 2 devices do not create boundaries for broadcasts or packet floods. Consequently, Layer 2 devices can use redundant paths to forward the same frames to each other continuously, eventually causing the network to fail.

To solve this problem, the BIG-IP[®] system supports a set of industry-standard, Layer 2 protocols known as spanning tree protocols. *Spanning tree protocols* block redundant paths on a network, thus preventing bridging loops. If a blocked, redundant path is needed later because another path has failed, the spanning tree protocols clear the path again for traffic. The spanning tree protocols that the BIG-IP system supports are Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP).

Central to the way that spanning tree protocols operate is the use of bridge protocol data units (BPDUs). When you enable spanning tree protocols on Layer 2 devices on a network, the devices send BPDUs to each other, for the purpose of learning the redundant paths and updating their L2 forwarding tables accordingly, electing a root bridge, building a spanning tree, and notifying each other about changes in interface status.

Note: The term *bridge* refers to a Layer 2 device such as a switch, bridge, or hub.

About STP protocol

STP is the original spanning tree protocol, designed to block redundant paths as a way to prevent bridging loops. The STP algorithm creates one, and only one, spanning tree for the entire network. A *spanning tree* is a logical tree-like depiction of the bridges on a network and the paths that connect them.

Because STP is unable to recognize VLANs and usually exhibits poor performance overall, STP is not the preferred spanning tree protocol to use in VLAN-rich environments. However, all participating interfaces in the spanning tree must use the same spanning tree protocol at any given time. Thus, when you have legacy bridges in your environment that are running STP, interfaces on the BIG-IP[®] system must have the ability to automatically degrade to STP.

Because STP has no knowledge of VLANs, you can have only one spanning tree instance on the BIG-IP system when using STP.

About the RSTP protocol

RSTP is an enhancement to STP, and was designed specifically to improve spanning tree performance. Like STP, RSTP can create only one spanning tree (instance 0), and therefore cannot take VLANs into account

when managing redundant paths. However, RSTP's performance improvements generally make it preferable to STP in non-VLAN environments.

In the case where legacy RSTP bridges are on the network, BIG-IP® system interfaces running MSTP can degrade to RSTP, just as they can degrade to STP.

Like STP, RSTP allows only one spanning tree instance on the BIG-IP system.

About the MSTP protocol

MSTP is an enhancement to RSTP and is the preferred spanning tree protocol for the BIG-IP® system. MSTP is specifically designed to understand VLANs and VLAN tagging (specified in IEEE 802.1q). Unlike STP and RSTP, which allow only one spanning tree instance per system, MSTP allows multiple spanning tree instances. Each instance corresponds to a spanning tree, and can control one or more VLANs that you specify when you create the instance. Thus, for any BIG-IP system interface that you assigned to multiple VLANs, MSTP can block a path on one VLAN, while still keeping a path in another VLAN open for traffic. Neither STP nor RSTP has this capability.

A unique feature of MSTP is the concept of spanning tree regions. A *spanning tree region* is a logical set of bridges on the network that share the same values for certain MSTP configuration settings. These configuration settings are: The MSTP configuration name, the MSTP configuration number, the instance numbers, and the VLAN members of each instance. When the values of these settings are identical on two or more bridges, the spanning tree algorithm considers these bridges to constitute an MSTP region. An MSTP region indicates to the spanning tree algorithm that it can use MSTP for all bridges in that region, and thus take VLANs into account when blocking and unblocking redundant paths.

You do not explicitly create a region. The spanning tree algorithm automatically groups bridges into regions, based on the values you assign to the MSTP configuration name, revision number, instance numbers, and instance members.

MSTP can only operate on bridges that are within a region. However, if the BIG-IP system connects to a bridge in a different MSTP region or outside of an MSTP region, the system still participates in spanning tree. In this case, the system is part of the spanning tree instance 0, also known as the Common and Internal Spanning Tree (CIST).

Note: BIG-IP systems released prior to version 9.0 do not support MSTP.

About spanning tree with legacy bridges

A key concept about spanning tree protocols on the BIG-IP® system is the concept of protocol degradation. *Protocol degradation* occurs when the spanning tree mode on the BIG-IP system is set to MSTP or RSTP, but the system detects legacy bridges (that is, bridges running an older protocol type) on the network. In this case, the BIG-IP system automatically degrades the spanning tree protocol that is running on each applicable interface to match the protocol running on the legacy device.

For example, suppose you set the BIG-IP system to run in MSTP mode. Later, if a bridge running STP is added to the network, the BIG-IP system will detect the legacy device and automatically degrade the protocol running on the BIG-IP system interfaces from MSTP to STP. The mode is still set to MSTP, but the interfaces actually run STP.

If the legacy device is later removed from the network, you can choose, for each BIG-IP system interface, to manually reset the spanning tree protocol back to MSTP.

The basic principle of protocol degradation is that each BIG-IP system interface in a spanning tree runs the oldest protocol that the system detects on the Layer 2 devices of the network. Thus, if a legacy bridge running STP is added to the network, BIG-IP system interfaces running MSTP or RSTP degrade to STP. Similarly, if a legacy bridge is running RSTP (and no bridges are running STP), interfaces running MSTP degrade to RSTP.

Note that when a bridge running MSTP must degrade to RSTP, the spanning tree algorithm automatically puts the degraded bridge into a separate MSTP region.

Configuration overview

Regardless of which spanning tree protocol you choose to use, the BIG-IP® Configuration utility offers a complete set of default configuration settings. Except for choosing a preferred spanning tree protocol to use, there are very few configuration settings that you need to modify to use the spanning tree feature effectively.

Note: *An alternate way to configure spanning tree protocols is to use `tmsh`.*

When you configure spanning tree on a BIG-IP system, you must first decide which protocol, or mode, you want to enable. Because MSTP recognizes VLANs, using MSTP is preferable for the BIG-IP system. However, all bridges in a network environment that want to use spanning tree must run the same spanning tree protocol. If a legacy bridge running RSTP or STP is added to the network, the BIG-IP system must switch to that same protocol.

Fortunately, you do not need to continually reconfigure the BIG-IP system spanning tree mode whenever a legacy bridge is added to the network. Instead, a BIG-IP system interface can detect the addition of a legacy bridge and automatically fall back to either RSTP or STP mode. If the legacy bridge is later removed from the network, you can use the BIG-IP Configuration utility to manually reset the interface back to running MSTP.

Once you have enabled a spanning tree mode, you can configure a set of global options. These options are the same options that are defined in the IEEE standards for the spanning tree protocols. While you can use the default settings in most cases, a few settings require user input.

Spanning tree mode

The Mode option specifies the particular spanning tree protocol that you want to use on the BIG-IP® system. The default value is Pass Through. The possible values are:

Disabled

Specifies that when the BIG-IP system receives spanning tree frames (BPDUs), it discards the frames.

Pass Through

Specifies that when the BIG-IP system receives spanning tree frames (BPDUs), it forwards them to all other interfaces. This is the default setting. When you use Pass Through mode, the BIG-IP system is transparent to spanning tree BPDUs. When set to Pass Through mode, the BIG-IP system is not part of any spanning tree. Note that Pass Through mode is not part of the IEEE spanning tree protocol specifications.

STP

Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the STP protocol. This mode allows for legacy systems on the network.

RSTP

Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the RSTP protocol.

MSTP

Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the MSTP protocol.

When you set the mode to MSTP or RSTP, and a legacy bridge running STP is subsequently added to the spanning tree, the applicable BIG-IP system interface automatically changes to running STP. However, you can manually reset an interface to resume operation in RSTP or MSTP mode if the legacy bridge is later removed from the spanning tree.

Global timers

All three spanning tree protocols, have the same three global timer values that you can specify: Hello Time, Maximum Age, and Forward Delay.

About the hello time option

When you change the value of the Hello Time option, you change the time interval, in seconds, that the BIG-IP[®] system transmits spanning tree information (through BPDUs) to adjacent bridges in the network. The default value for this option is 2.

Warning: *Although valid values are in the range of 1 to 10 seconds, F5 Networks[®] highly recommends that you use the default value (2 seconds). This value is optimal for almost all configurations.*

Note that when running RSTP, you must maintain the following relationship between the Maximum Age and Hello Time options:

$$\text{Maximum Age} \geq 2 * (\text{Hello Time} + 1)$$

About the maximum age option

When you change the value of the Maximum Age option, you change the amount of time, in seconds, that spanning tree information received from other bridges is considered valid. The default value is 20, and the valid range is 6 to 40.

Note that when running RSTP, you must maintain the following relationships between the Maximum Age and the Hello Time and Forward Delay options:

$$\text{Maximum Age} \geq 2 * (\text{Hello Time} + 1)$$
$$\text{Maximum Age} \leq 2 * (\text{Forward Delay} - 1)$$

About the forward delay option

Primarily used for STP, the Forward Delay option specifies the amount of time, in seconds, that the system blocks an interface from forwarding network traffic when the spanning tree algorithm reconfigures a spanning tree. The default value is 15, and the valid range is 4 to 30.

This option has no effect on the BIG-IP® system when running in RSTP or MSTP mode, provided that all bridges in the spanning tree use the RSTP or MSTP protocol. However, if the addition of legacy STP bridges causes neighboring bridges to fall back to running the STP protocol, then the spanning tree algorithm uses the Forward Delay option when reconfiguring the spanning tree.

Note that when running RSTP, you must maintain the following relationship between the Forward Delay and Maximum Age options:

$$\text{Maximum Age} \leq 2 * (\text{Forward Delay} - 1)$$

About the transmit hold count option

When you change the value of the Transmit Hold Count option, you change the maximum number of spanning tree frames (BPDUs) that the system can transmit on a port within the Hello Time interval. This setting ensures that the spanning tree frames do not overload the network, even in unstable network conditions. The default value is 6, and the valid range is 1 to 10.

MSTP-specific global properties

If you are running MSTP, you can configure three additional global properties:

MSTP configuration name

Applicable to MSTP only, the MSTP Configuration Name setting represents a global name that you assign to all bridges in a spanning tree region. A spanning tree region is a group of bridges with identical MSTP configuration names and MSTP configuration revision levels, as well as identical assignment of VLANs to spanning tree instances. All bridges in the same region must have this same configuration name. The name must contain from 1 to 32 characters. This option only appears on the screen when you set the Mode property to MSTP.

MSTP configuration revision

Applicable to MSTP only, the MSTP Configuration Revision setting represents a global revision number that you assign to all bridges in a spanning tree region. All bridges in the same region must have this same configuration revision number. The default value is 0. You can type any value between 0 and 65535. This option only appears on the screen when you set the Mode property to MSTP.

Maximum hop number

Applicable to MSTP only, this global property specifies the maximum number of hops that a spanning tree frame (BPDU) can traverse before it is discarded. The default value is 20. You can specify a value between 1 and 255. This option only appears on the screen when you set the Mode property to MSTP.

Management of spanning tree instances

By default, the spanning tree protocol STP is enabled on all of the interfaces of the BIG-IP® system. The default spanning tree configuration includes a single spanning tree instance, named 0. A spanning tree instance is a discrete spanning tree for a network. While STP and RSTP allow only one spanning tree instance (instance 0), MSTP allows you to create multiple spanning tree instances, to manage redundant paths for specific VLANs on the network.

When running MSTP, instances that you create have instance members. An instance member is a VLAN that you assign to an instance when you create that instance. You can assign as many or as few members to an instance as you deem necessary. By default, all VLANs on the BIG-IP system are members of instance 0.

If you create an instance and attempt to add a VLAN that is already a member of another instance, the BIG-IP system deletes the VLAN from the existing instance and adds the VLAN to the new instance.

Each instance name must be a numeric value that you assign when you create the instance.

Note: Only users with either the Administrator or Resource Administrator role can manage spanning tree instances.

Spanning tree instances list

You can view a list of existing spanning tree instances using the BIG-IP Configuration utility. For STP and RSTP, the only instance listed is instance 0. For MSTP, the list shows instance 0, plus any other instances that you have explicitly created.

When you view a list of instances, you can see the following information for each instance:

- The name of the instance
- The bridge priority number
- The MAC address of the root bridge
- The MAC address of the regional root bridge
- The number of instance members

About spanning tree instance (MSTP-only) creation

The STP and RSTP protocols allow only one spanning tree instance, instance 0, which the BIG-IP® system creates automatically when you enable spanning tree. When running STP or RSTP, you can modify the properties of instance 0, but you cannot create additional instances.

When you are running MSTP, however, the MSTP algorithm can explicitly create instances. The reason that you can create instances is that MSTP recognizes VLANs. By creating an instance and assigning one or more VLANs to it, you can control bridge loops and redundant paths within those VLANs.

For example, suppose you have two interfaces. One interface is assigned to VLAN A, while the other interface is assigned to VLANs A and B. If you are using the STP or RSTP protocol, both of which disregard VLANs, the protocol might block traffic for both VLANs, as shown in this figure.




Figure 10: Using STP or RSTP to block redundant paths

By contrast, the MSTP protocol can make blocking decisions on a per-VLAN basis. In our example, on the interface that carries traffic for two VLANs, you can block traffic for VLAN A, but leave a path open for VLAN B traffic. For example:




Figure 11: A local BIG-IP system that transmits and receives LLDPDUs

Because all BPDUs exchanged within a region always reference instance 0, instance 0 is active on all interfaces. This, in turn, can cause blocking problems. To avoid this, make sure that each VLAN on a BIG-IP system is a member of an instance that you explicitly create, rather than a member of instance 0 only. For example, suppose you create the following:

- Instance 1 with VLAN A as a member, where VLAN A is associated with interface 1.2
- Instance 2 with VLAN B as a member, where VLAN B is associated with interface 1.4

In this case, neither interface will be blocked, because the BPDUs sent from each interface reference a unique instance (either instance 1 or instance 2).

Tip: Because all BPDUs exchanged within a region always reference instance 0, thereby causing instance 0 to be active on all interfaces, unwanted blocking problems can occur. To avoid this, make sure that each VLAN on a BIG-IP system is a member of an instance that you explicitly create, rather than a member of instance 0 only.

About instance ID assignment

When you configure the Instance ID setting, you specify a numeric value for the instance, in the range of 1 to 255. The reason that instance names must be numeric is to handle the requirement that all cooperating bridges agree on the assignment of VLANs to instance IDs. Using numeric values instead of names makes this requirement easier to manage.

Bridge priority

The bridge in the spanning tree with the lowest relative priority becomes the root bridge. A *root bridge* represents the root of a spanning tree, and is responsible for managing loop resolution on the network. F5 Networks® recommends that you configure this setting so that the BIG-IP® system never becomes the root bridge. For this reason, the default value for the **Bridge Priority** setting is 61440, the highest value that you can select. Note that a bridge priority must be in increments of 4096.

VLAN assignment

If you are running MSTP, you can add members to a spanning tree instance. An *instance member* is a VLAN. You add members to an instance by associating one or more VLANs with the instance. The interfaces or trunks associated with each VLAN automatically become part of the spanning tree corresponding to that instance.

For two or more bridges to operate in the same spanning tree, all of those bridges must be in the same region, and therefore must have the same instance numbers, instance members, and VLAN tags.

For example, if a bridge has instance 1, with two VLAN members whose tags are 1000 and 2000, then any other bridges that you want to operate in that spanning tree must also have instance 1 with two VLAN members whose tags are 1000 and 2000.

A particular VLAN cannot be associated with more than one spanning tree instance. For example, if you have two instances named 0 and 1, you can only associate VLAN external with one of those instances, not both. Therefore, before creating an instance, verify that each VLAN you intend to associate with the instance is not a member of another instance.

Tip: *If no VLANs appear in the **Available** list when creating an instance, it is likely that all VLANs on the BIG-IP® system are members of other instances. You can verify this by viewing the members of other instances.*

About viewing and modifying a spanning tree instance

Using the BIG-IP Configuration utility, you can view and modify properties of any instance, including instance 0. If you are running MSTP, you can modify the Bridge Priority and VLANs properties. If you are running RSTP or STP, you can modify only the Bridge Priority property. In no case can you modify the instance ID.

About deleting a spanning tree instance or its members (MSTP-only)

If you are running MSTP, you might have explicitly created some spanning tree instances. If so, you can delete any spanning tree instance except instance 0.

You can also remove VLAN members from an instance. When you remove a VLAN from an instance, the VLAN automatically becomes a member of instance 0. (By default, instance 0 includes any VLAN that is not a member of another instance.)

If you remove all members from an instance, the BIG-IP® system automatically deletes the instance.

Note: *If you are running RSTP or STP, you cannot delete instance 0 or remove members from it.*

Interfaces for spanning tree

Some of the configuration tasks you perform when managing a spanning tree protocol pertain to BIG-IP® system interfaces. The interface-related tasks you perform are:

- Configuring settings on each interface that is to be part of the spanning tree
- Managing interfaces per spanning tree instance

Enabling and disabling spanning tree

When you select the check box for the **STP** setting, you are specifying that the interface can become part of a spanning tree. Once the interface becomes part of the spanning tree, the spanning tree protocol (STP) takes control of all learning and frame forwarding on that interface.

If you disable this setting, the spanning tree protocol treats the interface as non-existent, and does not send BPDUs to that interface. Also, the interface, and not the spanning tree protocol, controls all learning and frame forwarding for that interface.

If you want to enable or disable spanning tree for a trunk, you must do this on the trunk itself, and not on the individual interfaces (including the reference link) that make up the trunk.

Also, you must use the Traffic Management Shell (`tmsh`) to enable or disable spanning tree for a trunk. From within the `tmsh` shell, you can use these commands:

- To enable or disable spanning tree on a trunk, type: `modify net trunk trunk_name stp enabled|disabled`
- To save the change, type: `save /sys config base`
- To view the status of spanning tree on all interfaces, including trunks, type: `show net stp`.

STP link type

When you specify an STP link type, you ensure that STP uses the correct optimizations for the interface. Possible values are:

auto

When you set the STP link type to `auto`, the BIG-IP® system determines the spanning tree link type, which is based on the Active Duplex interface property.

p2p

When you set the STP link type to p2p, the BIG-IP system uses the optimizations for point-to-point spanning tree links. Point-to-point links connect two spanning tree bridges only. For example, a point-to-point link might connect a 10 Gigabit link to another bridge. For point-to-point links, the Active Duplex property interface should be set to full. Note that p2p is the only valid STP link type for a trunk.

shared

When you set the STP link type to shared, the BIG-IP system uses the optimizations for shared spanning tree links. Shared links connect two or more spanning tree bridges. For example, a shared link might be a 10 Megabit hub. Note that for shared links, the Active Duplex interface property should be set to half.

STP edge port

When you enable the **STP Edge Port** setting, you are explicitly designating the interface as an edge port. An *edge port* is an interface that connects to an end station rather than to another spanning tree bridge. The default setting is disabled (not checked).

If you would rather have the system automatically designate the interface as an edge port, you can enable the STP Edge Port Detection setting instead.

If you enable (check) the **STP Edge Port** setting and the interface subsequently receives STP, RSTP, or MSTP frames (BPDUs), the system disables the setting automatically, because only non-edge interfaces receive BPDUs.

Detection of an STP edge port

When you enable the **STP Edge Port Detection** setting, the system determines whether the interface is an edge port, and if so, automatically designates the interface as an edge port. The system determines edge port status by monitoring the interface and verifying that it does not receive any incoming STP, RSTP, or MSTP frames (BPDUs).

If the system determines that the interface is not an edge port, but you enabled the **STP Edge Port** setting to explicitly designate the interface as an edge port, the system removes the edge port designation from the interface. No interface that receives BPDUs from a bridge can have edge port status, despite the values of the STP Edge Port and STP Edge Port Detection settings.

About spanning tree protocol reset

The spanning tree algorithm automatically detects the presence of legacy STP bridges on the network, and falls back to STP mode when communicating with those bridges. Because legacy STP bridges do not send spanning tree BPDUs periodically in all circumstances, the BIG-IP® system cannot detect when a legacy STP bridge is removed from the network. Therefore, it is necessary to manually notify the BIG-IP system that the algorithm can switch to the RSTP or MSTP protocol again, whenever a legacy bridge has been removed.

About managing interfaces for a specific instance

When you manage an interface for a specific spanning tree instance, you can:

- View a list of interfaces for an instance
- View instance-specific properties of an interface
- Configure instance-specific settings for an interface

About viewing a list of interface IDs for an instance

Using the BIG-IP Configuration utility, you can view a list of the interface IDs associated with a specific spanning tree instance.

If you are using MSTP, the interface IDs that appear in the list are the interfaces assigned to the VLANs that you specified when you created the instance. If you are using STP or RSTP, the interface IDs listed are those that the BIG-IP® system automatically assigned to instance 0.

The list of interface IDs also displays the following information for each interface:

- The STP instance ID
- The priority
- The external path cost
- The port role

About port roles

The Port Role property of a per-instance interface specifies the interface's role in the spanning tree instance. You cannot specify a value for this property; the BIG-IP® system automatically assigns a role to the interface.

The BIG-IP system can assign one of the following roles to an instance interface:

Disabled

The interface has no active role in the spanning tree instance.

Root

The interface provides a path to a root bridge.

Alternate

The interface provides an alternate path to a root bridge, if the root interface is unavailable.

Designated

The interface provides a path away from the root bridge.

Backup

The interface provides an alternate path away from the root bridge, if an interface with a port role of Designated is unavailable. The Backup role assignment is rare.

Port states

The Port State property of an interface specifies the way that the interface processes normal data packets. You cannot specify a value for this property; the BIG-IP® system automatically assigns a state to the interface.

An interface can be in one of the following states at any given time:

Blocking

The interface disregards any incoming frames, and does not send any outgoing frames.

Forwarding

The interface passes frames as needed.

Learning

The interface is determining information about MAC addresses, and is not yet forwarding frames.

Settings to configure for an interface for a specific instance

Selecting an interface priority

Each interface has an associated priority within a spanning tree instance. The relative values of the interface priorities affect which interfaces the system chooses to carry network traffic. Using the **Interface Priority** setting, you can select the interface's priority in relation to the other interfaces that are members of the spanning tree instance.

Typically, the system is more likely to select interfaces with lower numeric values to carry network traffic. A priority value that you assign to an interface can be in the range of 0 to 240, in increments of 16. Thus, the value you assign to an interface can be 0, 16, 32, 64, and so on, up to 240.

The default priority for an interface is 128, the middle of the valid range.

Specifying path cost

Each interface has an associated path cost within a spanning tree instance. The *path cost* represents the relative cost of sending network traffic through that interface. When calculating the spanning tree, the spanning tree algorithm attempts to minimize the total path cost between each point of the tree and the root bridge. By manipulating the path costs of different interfaces, you can steer traffic toward paths that are either faster, more reliable, more economical, or have all of these qualities.

The value of a path cost can be in the range of 1 to 200,000,000, unless you have legacy STP bridges. In that case, because some legacy implementations support a range of only 1 to 65535, you should use this more restricted range when setting path costs on interfaces.

The default path cost for an interface is based on the maximum speed of the interface rather than the actual speed.

For example, an interface that has a maximum speed of 1000 Mb/s (1 Gb/s), but is currently running at a speed of 10 Mb/s, has a default path cost of 20,000.

Link aggregation does not affect the default path cost. For example, if a trunk has four 1 Gb/s interfaces, the default path cost is 20,000.

For MSTP, you can set two kinds of path costs, external and internal. For STP and RSTP, you can set an external path cost only.

External Path Cost

The **External Path Cost** setting is used to calculate the cost of sending spanning tree traffic through the interface to reach an adjacent spanning tree region. The spanning tree algorithm tries to minimize the total path cost between each point of the tree and the root bridge. The external path cost applies only to those interfaces (and trunks) that are members of instance 0.

Internal Path Cost

The **Internal Path Cost** setting allows you to specify the relative cost of sending spanning tree traffic through the interface to adjacent bridges within a spanning tree region. Note that the internal path cost applies only to bridges that support the MSTP mode. The internal path cost applies to those interfaces (and trunks) that are members of any instance, including instance 0.

To summarize, STP and RSTP use external path costs only, and the costs apply to instance 0 interfaces only. MSTP uses both external and internal path costs, and the internal costs apply to interfaces in all spanning tree instances, including instance 0.

Trunks

Introduction to trunks

A *trunk* is a logical grouping of interfaces on the BIG-IP® system. When you create a trunk, this logical group of interfaces functions as a single interface. The BIG-IP system uses a trunk to distribute traffic across multiple links, in a process known as *link aggregation*. With link aggregation, a trunk increases the bandwidth of a link by adding the bandwidth of multiple links together. For example, four fast Ethernet (100 Mbps) links, if aggregated, create a single 400 Mbps link.

With one trunk, you can aggregate a maximum of eight links. For optimal performance, you should aggregate links in powers of two. Thus, you ideally aggregate two, four, or eight links.

The purpose of a trunk is two-fold:

- To increase bandwidth without upgrading hardware
- To provide link failover if a member link becomes unavailable

You can use trunks to transmit traffic from a BIG-IP system to another vendor switch. Two systems that use trunks to exchange frames are known as *peer systems*.

How do trunks work?

In a typical configuration where trunks are configured, the member links of the trunk are connected through Ethernet cables to corresponding links on a peer system.

This figure shows an example of a typical trunk configuration with two peers and three member links on each peer:




Figure 12: Example of a trunk configured for two switches

A primary goal of the trunks feature is to ensure that frames exchanged between peer systems are never sent out of order or duplicated on the receiving end. The BIG-IP® system is able to maintain frame order by using the source and destination addresses in each frame to calculate a hash value, and then transmitting all frames with that hash value on the same member link.

The BIG-IP system automatically assigns a unique MAC address to a trunk. However, by default, the MAC address that the system uses as the source and destination address for frames that the system transmits and receives (respectively), is the MAC address of the lowest-numbered interface of the trunk.

The BIG-IP system also uses the lowest-numbered interface of a trunk as a reference link. The BIG-IP system uses the reference link to take certain aggregation actions, such as implementing the automatic link selection policy. For frames coming into the reference link, the BIG-IP system load balances the frames across all member links that the BIG-IP system knows to be available. For frames going from any link in

the trunk to a destination host, the BIG-IP system treats those frames as if they came from the reference link.

Finally, the BIG-IP system uses the MAC address of an individual member link as the source address for any LACP control frames.

Overview of LACP

A key aspect of trunks is Link Aggregation Control Protocol, or LACP. Defined by IEEE standard 802.3ad, *LACP* is a protocol that detects error conditions on member links and redistributes traffic to other member links, thus preventing any loss of traffic on the failed link. On a BIG-IP® system, LACP is an optional feature that you can configure.

You can also customize LACP behavior. For example, you can specify the way that LACP communicates its control messages from the BIG-IP system to a peer system. You can also specify the rate at which the peer system sends LACP packets to the BIG-IP system. If you want to affect the way that the BIG-IP system chooses links for link aggregation, you can specify a link control policy.

Trunk name

You can use the **Name** setting to specify a unique name for the trunk. This setting is required.

Interfaces for a trunk

Using the **Interfaces** setting, you specify the interfaces that you want the BIG-IP® system to use as member links for the trunk. Once you have created the trunk, the BIG-IP system uses these interfaces to perform link aggregation.

Tip: To optimize bandwidth utilization, F5 Networks® recommends that, if possible, the number of links in the trunk be a power of 2 (for example, 2, 4, or 8). This is due to the frame balancing algorithms that the system uses to map data streams to links. Regardless of the hashing algorithm, a trunk that has 2, 4, or 8 links prevents the possibility of skewing, which can adversely affect data throughput.

The BIG-IP system uses the lowest-numbered interface as the reference link. The system uses the reference link to negotiate links for aggregation.

The interfaces that you specify for the trunk must operate at the same media speed, and must be set at full-duplex mode. Otherwise, the BIG-IP system cannot aggregate the links. Because these media properties are dynamic rather than static (due to auto-negotiation), the `lacpd` service routinely monitors the current status of these properties and negotiates the links for aggregation accordingly. Thus, when the status of these properties qualifies a link to become a working member link, the system adds the link to the aggregation, and the link can begin accepting traffic.

Any interface that you assign to a trunk must be an untagged interface. Furthermore, you can assign an interface to one trunk only; that is, you cannot assign the same interface to multiple trunks. Because of these restrictions, the only interfaces that appear in the Interfaces list in the BIG-IP Configuration utility are untagged interfaces that are not assigned to another trunk. Therefore, before creating a trunk and assigning any interfaces to it, you should verify that each interface for the trunk is an untagged interface.

After creating the trunk, you assign the trunk to one or more VLANs, using the same VLAN screen that you normally use to assign an individual interface to a VLAN.

If you are using one of the spanning tree protocols (STP, RSTP, or MSTP), the BIG-IP system sends and receives spanning tree protocol packets on a trunk, rather than on individual member links. Likewise, use of a spanning tree protocol to enable or disable learning or forwarding on a trunk operates on all member links together, as a single unit.

About enabling LACP

As an option, you can enable LACP on a trunk. Containing a service called `lacpd`, LACP is an IEEE-defined protocol that exchanges control packets over member links. The purpose of LACP is to detect link error conditions such as faulty MAC devices and link loopbacks. If LACP detects an error on a member link, the BIG-IP® system removes the member link from the link aggregation and redistributes the traffic for that link to the remaining links of the trunk. In this way, no traffic destined for the removed link is lost. LACP then continues to monitor the member links to ensure that aggregation of those links remains valid.

By default, the LACP feature is disabled, to ensure backward compatibility with previous versions of the BIG-IP system. If you create a trunk and do not enable the LACP feature, the BIG-IP system does not detect link error conditions, and therefore cannot remove the member link from link aggregation. The result is that the system cannot redistribute the traffic destined for that link to the remaining links in the trunk, thereby causing traffic on the failed member link to be lost.

Important: To use LACP successfully, you must enable LACP on both peer systems.

LACP mode

The **LACP Mode** setting appears on the Trunks screen only when you select the LACP setting. You use the **LACP Mode** setting to specify the method that LACP uses to send control packets to the peer system. The two possible modes are:

Active mode

You specify **Active** mode if you want the system to periodically send control packets, regardless of whether the peer system has issued a request. This is the default setting.

Passive mode

You specify **Passive** mode if you want the system to send control packets only when the peer system issues a request, that is, when the LACP mode of the peer system is set to Active.

If you set only one of the peer systems to Active mode, the BIG-IP® system uses Active mode for both systems. Also, whenever you change the LACP mode on a trunk, LACP renegotiates the links that it uses for aggregation on that trunk.

Tip: We recommend that you set the LACP mode to Passive on one peer system only. If you set both systems to Passive mode, LACP does not send control packets.

LACP timeout

The **LACP Timeout** setting appears on the Trunks screen only when you select the LACP setting. You use the **LACP Timeout** setting to indicate to the BIG-IP® system the interval in seconds at which the peer system should send control packets. The timeout value applies only when the LACP mode is set to Active on at least one of the switch systems. If both systems are set to Passive mode, LACP does not send control packets.

If LACP sends three consecutive control packets without receiving a response from the peer system, LACP removes that member link from link aggregation. The two possible timeout values are:

Short

When you set the timeout value to Short, the peer system sends LACP control packets once every second. If this value is set to Short and LACP receives no peer response after sending three consecutive packets, LACP removes the link from aggregation in three seconds.

Long

When you set the timeout value to Long, the peer system sends LACP control packets once every 30 seconds. A timeout value of Long is the default setting. If set to Long and LACP receives no peer response after sending three consecutive packets, LACP removes the link from aggregation in ninety seconds.

Whenever you change the LACP timeout value on a trunk, LACP renegotiates the links that it uses for aggregation on that trunk.

Link selection policy

In order for the BIG-IP® system to aggregate links, the media speed and duplex mode of each link must be the same on both peer systems. Because media properties can change dynamically, the BIG-IP system monitors these properties regularly, and if it finds that the media properties of a link are mismatched on the peer systems, the BIG-IP system must determine which links are eligible for aggregation.

The way the system determines eligible links depends on a link selection policy that you choose for the trunk. When you create a trunk, you can choose one of two possible policy settings: **Auto** and **Maximum Bandwidth**.

***Note:** The link selection policy feature represents an F5 Networks® enhancement to the standard IEEE 802.3ad specification for LACP.*

Automatic link selection

When you set the link selection policy to Auto (the default setting), the BIG-IP® system uses the lowest-numbered interface of the trunk as a reference link. (A *reference link* is a link that the BIG-IP system uses to make a link aggregation decision.) The system then aggregates any links that have the same media properties and are connected to the same peer as the reference link.

For example, suppose that you created a trunk to include interfaces 1.2 and 1.3, each with a media speeds of 100 Mbps, and interface 1.4, with a different media speed of 1 Gbps. If you set the link selection policy

to Auto, the BIG-IP system uses the lowest-numbered interface, 1.2, as a reference link. The reference link operates at a media speed of 100 Mbps, which means that the system aggregates all links with that media speed (interfaces 1.2 and 1.3). The media speed of interface 1.4 is different (1 Gbps), and therefore is not considered for link aggregation. Only interfaces 1.2 and 1.3 become working member links and start carrying traffic.

If the media speed of interface 1.4 changes to 100 Mbps, the system adds that interface to the aggregation. Conversely, if the media speed of interface 1.4 remains at 1 Gbps, and the speed of the reference link changes to 1 Gbps, then interfaces 1.2 and 1.4 become working members, and 1.3 is now excluded from the aggregation and no longer carries traffic.

Maximum bandwidth link selection

When you set the link selection policy to **Maximum Bandwidth**, the BIG-IP® system aggregates the subset of member links that provide the maximum amount of bandwidth to the trunk.

If interfaces 1.2 and 1.3 each operate at a media speed of 100 Mbps, and interface 1.4 operates at speed of 1 Gbps, then the system selects only interface 1.4 as a working member link, providing 1 Gbps of bandwidth to the trunk. If the speed of interface 1.4 drops to 10 Mbps, the system then aggregates links 1.2 and 1.3, to provide a total bandwidth to the trunk of 200 Mbps. The peer system detects any non-working member links and configures its aggregation accordingly.

***Tip:** To ensure that link aggregation operates properly, make sure that both peer systems agree on the link membership of their trunks.*

Frame distribution hash

When frames are transmitted on a trunk, they are distributed across the working member links. The distribution function ensures that the frames belonging to a particular conversation are neither mis-ordered nor duplicated at the receiving end.

The BIG-IP® system distributes frames by calculating a hash value based on the source and destination addresses (or the destination address only) carried in the frame, and associating the hash value with a link. All frames with a particular hash value are transmitted on the same link, thereby maintaining frame order. Thus, the system uses the resulting hash to determine which interface to use for forwarding traffic.

The **Frame Distribution Hash** setting specifies the basis for the hash that the system uses as the frame distribution algorithm.

The default value is Source/Destination IP address.

Possible values for this setting are:

Source/Destination MAC address

This value specifies that the system bases the hash on the combined MAC addresses of the source and the destination.

Destination MAC address

This value specifies that the system bases the hash on the MAC address of the destination.

Source/Destination IP address

This value specifies that the system bases the hash on the combined IP addresses of the source and the destination.

Important: On certain F5[®] platforms, packets can incorrectly egress on the same BIG-IP trunk member that the external switch ingressed the packets on. You can prevent this by configuring the external switch to use the same algorithm for its frame distribution hash value as you configure on the BIG-IP trunk. For example, if you configure the BIG-IP trunk to base the frame distribution hash value on both source and destination IP addresses, then you must configure the external switch to do the same.

VLANs, VLAN Groups, and VXLAN

Introduction to virtual LANs

A *VLAN* is a logical subset of hosts on a local area network (LAN) that operate in the same IP address space. Grouping hosts together in a VLAN has distinct advantages. For example, with VLANs, you can:

- Reduce the size of broadcast domains, thereby enhancing overall network performance.
- Reduce system and network maintenance tasks substantially. Functionally-related hosts no longer need to physically reside together to achieve optimal network performance.
- Enhance security on your network by segmenting hosts that must transmit sensitive data.

The way that you group hosts into VLANs is by using the BIG-IP Configuration utility to create a VLAN and associate physical interfaces with that VLAN. In this way, any host that sends traffic to a BIG-IP[®] system interface is logically a member of the VLAN or VLANs to which that interface belongs.

VLANs on a BIG-IP system

The BIG-IP[®] system is a port-based switch that includes multilayer processing capabilities. These capabilities enhance standard VLAN behavior, in these ways:

- You can associate physical interfaces on the BIG-IP system directly with VLANs. In this way, you can associate multiple interfaces with a single VLAN, or you can associate a single interface with multiple VLANs.
- You do not need physical routers to establish communication between separate VLANs. Instead, the BIG-IP system can process messages between VLANs.
- You can incorporate a BIG-IP system into existing, multi-vendor switched environments, due to the BIG-IP system's compliance with the IEEE 802.1q VLAN standard.
- You can combine two or more VLANs into an object known as a VLAN group. With a VLAN group, a host in one VLAN can communicate with a host in another VLAN using a combination of Layer 2 forwarding and IP routing. This offers both performance and reliability benefits.

Default VLAN configuration

By default, the BIG-IP[®] system includes two VLANs, named internal and external. When you initially ran the Setup utility, you assigned the following to each of these VLANs:

- A static and a floating self IP address
- A VLAN tag
- One or more BIG-IP system interfaces

A typical VLAN configuration is one in which you create the two VLANs external and internal, and one or more BIG-IP system interfaces assigned to each VLAN. You then create a virtual server, and associate a default load balancing pool with that virtual server. This figure shows a typical configuration using the default VLANs external and internal.




Figure 13: A typical configuration using the default VLANs

Note: *VLANs internal and external reside in partition Common.*

Every VLAN must have a static self IP address associated with it. The self IP address of a VLAN represents an address space, that is, the range of IP addresses pertaining to the hosts in that VLAN. When you ran the Setup utility earlier, you assigned one static self IP address to the VLAN external, and one static self IP address to the VLAN internal. When sending a request to a destination server, the BIG-IP system can use these self IP addresses to determine the specific VLAN that contains the destination server.

For example, suppose the self IP address of VLAN external is 12.1.0.100, and the self IP address of the VLAN internal is 11.1.0.100, and both self IP addresses have a netmask of 255.255.0.0. If the IP address of the destination server is 11.1.0.20, then the BIG-IP system can compare the self IP addresses to the host's IP address to determine that the destination server is in the VLAN internal. This process, combined with checking the ARP cache and a VLAN's L2 forwarding table, ensures that the BIG-IP system successfully sends the request to the destination server.

Important: *In addition to configuring VLAN properties, you must also assign a self IP address to the VLAN.*

VLAN name

When creating a VLAN, you must assign it a unique name. Once you have finished creating the VLAN, the VLAN name appears in the VLAN list.

VLAN tags

A VLAN *tag* is a unique ID number that you assign to a VLAN. If you do not explicitly assign a tag to a VLAN, the BIG-IP® system assigns a tag automatically. The value of a VLAN tag can be between 1 and

4094. Once you or the BIG-IP system assigns a tag to a VLAN, any message sent from a host in that VLAN includes this VLAN tag as a header in the message.

A VLAN tag is useful when an interface has multiple VLANs associated with it; that is, when the interfaces you assigned to the VLAN are assigned as tagged interfaces. In this case, the BIG-IP system can read the VLAN tag in the header of a message to determine the specific VLAN in which the source or destination host resides.

Important: *If the device connected to a BIG-IP system interface is another switch, the VLAN tag that you assign to the VLAN on the BIG-IP system interface must match the VLAN tag assigned to the VLAN on the interface of the other switch.*

Interface assignments

For each VLAN that you create, you must assign one or more BIG-IP® system interfaces to that VLAN, using the **Interfaces** setting. When you assign an interface to a VLAN, you indirectly control the hosts from which the BIG-IP system interface sends or receives messages.

Tip: *You can assign not only individual interfaces to the VLAN, but also trunks.*

For example, if you assign interface 1.11 to VLAN A, and you then associate VLAN A with a virtual server, then the virtual server sends its outgoing traffic through interface 1.11, to a destination host in VLAN A. Similarly, when a destination host sends a message to the BIG-IP system, the host's VLAN membership determines the BIG-IP system interface that should receive the incoming traffic.

Each VLAN has a MAC address. The MAC address of a VLAN is the same MAC address of the lowest-numbered interface assigned to that VLAN.

The BIG-IP system supports two methods for sending and receiving messages through an interface that is a member of one or more VLANs. These two methods are port-based access to VLANs and tag-based access to VLANs. The method used by a VLAN is determined by the way that you add a member interface to a VLAN.

Port-based access to VLANs

With port-based access to VLANs, the BIG-IP® system accepts frames for a VLAN simply because they are received on an interface that is a member of that VLAN. With this method, an interface is an untagged member of the VLAN. Frames sent out through untagged interfaces contain no tag in their header.

Port-based access to VLANs occurs when you add an interface to a VLAN as an untagged interface. In this case, the VLAN is the only VLAN that you can associate with that interface. This limits the interface to accepting traffic only from that VLAN, instead of from multiple VLANs. If you want to give an interface the ability to accept and receive traffic for multiple VLANs, you add the same interface to each VLAN as a tagged interface.

Tag-based access to VLANs

With tag-based access to VLANs, the BIG-IP® system accepts frames for a VLAN because the frames have tags in their headers and the tag matches the VLAN identification number for the VLAN. An interface that accepts frames containing VLAN tags is a *tagged member* of the VLAN. Frames sent out through tagged interfaces contain a tag in their header.

Tag-based access to VLANs occurs when you add an interface to a VLAN as a tagged interface. You can add the same tagged interface to multiple VLANs, thereby allowing the interface to accept traffic from each VLAN with which the interface is associated.

When you add an interface to a VLAN as a tagged interface, the BIG-IP system associates the interface with the VLAN identification number, or tag, which becomes embedded in a header of a frame.

Note: Every VLAN has a tag. You can assign the tag explicitly when creating the VLAN, or the BIG-IP system assigns it automatically if you do not supply one.

Each time you add an interface to a VLAN, either when creating a VLAN or modifying its properties, you can designate that interface as a tagged interface. A single interface can therefore have multiple tags associated with it.

The result is that whenever a frame comes into that interface, the interface reads the tag that is embedded in a header of the frame. If the tag in the frame matches any of the tags associated with the interface, the interface accepts the frame. If the tag in the frame does not match any of the tags associated with the interface, the interface rejects the frame.

This figure shows the difference between using three untagged interfaces (where each interface must belong to a separate VLAN) versus one tagged interface (which belongs to multiple VLANs)




Figure 14: Solutions using untagged (left) and tagged interfaces (right)

The configuration on the left shows a BIG-IP system with three internal interfaces, each a separate, untagged interface. This is a typical solution for supporting three separate customer sites. In this scenario, each interface can accept traffic only from its own VLAN.

Conversely, the configuration on the right shows a BIG-IP system with one internal interface and an external switch. The switch places the internal interface on three separate VLANs. The interface is configured on each VLAN as a tagged interface. In this way, the single interface becomes a tagged member of all three VLANs, and accepts traffic from all three. The configuration on the right is the functional equivalent of the configuration of the left.

Important: If you are connecting another switch into a BIG-IP system interface, the VLAN tag that you assign to the VLAN on the BIG-IP system must match the VLAN tag on the interface of the other switch.

Source checking

When you enable source checking, the BIG-IP[®] system verifies that the return path for an initial packet is through the same VLAN from which the packet originated. Note that the system only enables source checking if the global setting Auto Last Hop is disabled.

Maximum transmission units

The value that you configure for the maximum transmission unit, or MTU, is the largest size that the BIG-IP® system allows for an IP datagram passing through a BIG-IP system interface. By default, the BIG-IP system uses the standard Ethernet frame size of 1518 bytes (1522 bytes if VLAN tagging is used), with a corresponding MTU value of 1500 bytes for a VLAN.

One reason for changing the value of the **MTU** setting is when your BIG-IP platform supports jumbo frames. A *jumbo frame* is an Ethernet frame with more than 1500 bytes, and fewer than 9000 bytes, of payload.

If your BIG-IP platform does not support jumbo frames and a VLAN receives a jumbo frame, the system discards the frame.

VLAN-based fail-safe

VLAN fail-safe is a feature you enable when you want to base redundant-system failover on VLAN-related events. To configure VLAN fail-safe, you specify a timeout value and the action that you want the system to take when the timeout period expires.

Auto last hop

When you create a VLAN, you can designate the VLAN as the last hop for TMM traffic.

CMP hash

The **CMP Hash** setting allows all connections from a client system to use the same set of TMMs. This improves system performance. In configuring the **CMP Hash** value, you can choose the traffic disaggregation criteria for the VLAN, either source IP address, destination IP address, or TCP/UDP source/destination ports. The default value uses TCP/UDP source/destination ports. Note that the **CMP Hash** setting appears only on the properties screen for an existing VLAN.

DAG round robin

You can use the **DAG Round Robin** setting on a VLAN to prevent stateless traffic from overloading a few TMM instances, a condition that can disable an entire BIG-IP system. When enabled, this setting causes the BIG-IP system to load balance the traffic among TMMs evenly, instead of using a static hash. Stateless traffic in this case includes non-IP Layer 2 traffic, ICMP, some UDP protocols, and so on. The default value for this setting is unchecked (disabled).

This feature is particularly useful for firewall and Domain Name System (DNS) traffic. For example, this feature prevents certain types of DDoS attacks, such as an ICMP DDoS attack that can overload the system by sending the same packets repeatedly to a specific subset of TMMs.

To enable the **DAG Round Robin** feature, log in to the BIG-IP Configuration utility, and on the navigation pane, locate the Main tab, expand **Network**, click **VLANs**, and click the **Create** button. Alternatively, instead of creating a new VLAN, you can enable this feature on an existing VLAN by displaying the list of VLANs on the system, and in the Name column, clicking the name of the VLAN you want to modify.

In addition to enabling the **DAG Round Robin** setting, you must also use the Traffic Management Shell (tmsh) to configure a bigdb variable that specifies the relevant destination ports. For example, you can use

the command `tmsch modify sys db dag.roundrobin.udp.portlist value "53:26:19:45"` to specify that the system load balances packets destined for ports 53, 26, 19, or 45. The values that you specify with this BigDB variable apply to all VLANs on which the **DAG Round Robin** setting is enabled.

***Note:** The disaggregation of traffic occurs only to TMMs that are local to a given high-speed bridge (HSB).*

Maintaining the L2 forwarding table

Layer 2 forwarding is the means by which frames are exchanged directly between hosts, with no IP routing required. This is accomplished using a simple forwarding table for each VLAN. The L2 forwarding table is a list that shows, for each host in the VLAN, the MAC address of the host, along with the interface that the BIG-IP® system needs for sending frames to that host. The intent of the L2 forwarding table is to help the BIG-IP system determine the correct interface for sending frames, when the system determines that no routing is required.

The format of an entry in the L2 forwarding table is:

```
<MAC address> -> <if>
```

For example, an entry for a host in the VLAN might look like this:

```
00:a0:c9:9e:1e:2f -> 2.1
```

The BIG-IP system learns the interfaces that correspond to various MAC entries as frames pass through the system, and automatically adds entries to the table accordingly. These entries are known as dynamic entries. You can also add entries to the table manually, and these are known as static entries. Entering static entries is useful if you have network devices that do not advertise their MAC addresses. The system does not automatically update static entries.

The BIG-IP system does not always need to use the L2 forwarding table to find an interface for frame transmission. For instance, if a VLAN has only one interface assigned to it, then the BIG-IP system automatically uses that interface.

Occasionally, the L2 forwarding table does not include an entry for the destination MAC address and its corresponding BIG-IP system interface. In this case, the BIG-IP system floods the frame through all interfaces associated with the VLAN, until a reply creates an entry in the L2 forwarding table.

About sFlow polling intervals and sampling rates

You can change the sFlow settings for a specific VLAN when you want the traffic flowing through the VLAN to be sampled at a different rate than the global sFlow settings on the BIG-IP® system.

About VLAN groups

A *VLAN group* is a logical container that includes two or more distinct VLANs. VLAN groups are intended for load balancing traffic in a Layer 2 network, when you want to minimize the reconfiguration of hosts on that network. This figure shows an example of a VLAN group.




Figure 15: Example of a VLAN group

A VLAN group also ensures that the BIG-IP® system can process traffic successfully between a client and server when the two hosts reside in the same address space. Without a VLAN group, when the client and server both reside in the same address space, the client request goes through the virtual server, but instead of sending its response back through the virtual server, the server attempts to send its response directly to the client, bypassing the virtual server altogether. As a result, the client cannot receive the response, because the client expects the address of the response to be the virtual server IP address, not the server IP address.

Tip: You can configure the behavior of the BIG-IP system so that it always creates a proxy for any ARP requests between VLANs.

When you create a VLAN group, the two existing VLANs become child VLANs of the VLAN group.

VLAN groups reside in administrative partitions. To create a VLAN group, you must first set the current partition to the partition in which you want the VLAN group to reside.

Note: Only users with the Administrator user role can create and manage VLAN groups.

About VLAN group names

When creating a VLAN group, you must assign it a unique name. Once you have finished creating the VLAN group, the VLAN group name appears in the list of existing VLANs groups.

VLAN group ID

A *VLAN group ID* is a tag for the VLAN group. Every VLAN group needs a unique ID number. If you do not specify an ID for the VLAN group, the BIG-IP® system automatically assigns one. The value of a VLAN group ID can be between 1 and 4094.

About transparency mode

The BIG-IP® system is capable of processing traffic using a combination of Layer 2 and Layer 3 forwarding, that is, switching and IP routing. When you set the transparency mode, you specify the type of forwarding that the BIG-IP system performs when forwarding a message to a host in a VLAN. The default setting is

translucent, which means that the BIG-IP system uses a mix of Layer 2 and Layer 3 processing. The allowed values are:

opaque

A proxy ARP with Layer 3 forwarding

translucent

Layer 2 forwarding with a locally-unique bit, toggled in ARP response across VLANs. This is the default setting. When you choose this value and you have a virtual server that references a Fast L4 profile, the BIG-IP system automatically changes the **PVA Acceleration** setting to **None**

transparent

Layer 2 forwarding with the original MAC address of the remote system preserved across VLANs. When you choose this value and you have a virtual server that references a Fast L4 profile, the BIG-IP system automatically changes the **PVA Acceleration** setting to **None**.

About traffic bridging

When you enable the traffic bridging option, you are instructing the VLAN group to forward all non-IP traffic. Note that IP traffic is bridged by default. The default value for this setting is disabled (unchecked).

About traffic bridging with standby units

When enabled, the **Bridge in Standby** setting ensures that the VLAN group can forward packets when the system is the standby device of a redundant system configuration. Note that this setting applies to non-IP and non-ARP frames only, such as Bridge Protocol Data Units (BPDUs).

This setting is designed for deployments in which the VLAN group is defined on a redundant system. You can use the **Bridge in Standby** setting in transparent or translucent modes, or in opaque mode when the global variable `Failover.Standby.LinkDownTime` is set to 0.

Warning: This setting can cause adverse effects if the VLAN group exists on more than one device in a device group. The setting is intended for configurations where the VLAN group exists on one device only. The default setting is enabled (checked).

About migration keepalive frames

The Migration Keepalive setting for a VLAN group, when enabled, instructs the BIG-IP® system to send keepalive frames (that is, TCP keepalives and empty UDP packets, depending on the connection type) when a node is moved from one VLAN group member to another VLAN group member for all existing BIG-IP connections to that node.

About host exclusion from proxy ARP forwarding

A host in a VLAN cannot normally communicate to a host in another VLAN. This rule applies to ARP requests as well. However, if you put the VLANs into a single VLAN group, the BIG-IP® system can perform a proxied ARP request.

A *proxied ARP request* is an ARP request that the BIG-IP system can send, on behalf of a host in a VLAN, to hosts in another VLAN. A proxied ARP request requires that both VLANs belong to the same VLAN group.

In some cases, you might not want a host to forward proxied ARP requests to a specific host, or to other hosts in the configuration. To exclude specific hosts from receiving forwarded proxied ARP requests, you use the BIG-IP Configuration utility and specify the IP addresses that you want to exclude.

Warning: *Although hosts on an ARP exclusion list are specified using their IP addresses, this does not prevent the BIG-IP system from routing traffic to those hosts. A more secure way to prevent traffic from passing between hosts in separate VLANs is to create a packet filter for each VLAN.*

VLAN association with a self IP address

After you create a VLAN or a VLAN group, you must associate it with a self IP address. You associate a VLAN or VLAN group with a self IP address using the New Self IPs screens of the BIG-IP Configuration utility:

Associating a VLAN with a self IP address

The self IP address with which you associate a VLAN should represent an address space that includes the IP addresses of the hosts that the VLAN contains. For example, if the address of one host is 11.0.0.1 and the address of the other host is 11.0.0.2, you could associate the VLAN with a self IP address of 11.0.0.100, with a netmask of 255.255.255.0.

Associating a VLAN group with a self IP address

The self IP address with which you associate a VLAN group should represent an address space that includes the self IP addresses of the VLANs that you assigned to the group. For example, if the address of one VLAN is 10.0.0.1 and the address of the other VLAN is 10.0.0.2, you could associate the VLAN group with a self IP address of 10.0.0.100, with a netmask of 255.255.255.0.

VLAN assignment to route domains

If you explicitly create route domains, you should consider the following facts:

- You can assign VLANs (and VLAN groups) to route domain objects that you create. Traffic pertaining to that route domain uses those assigned VLANs,
- During BIG-IP® system installation, the system automatically creates a default route domain, with an ID of 0. Route domain 0 has two VLANs assigned to it, VLAN internal and VLAN external.
- If you create one or more VLANs in an administrative partition other than `Common`, but do not create a route domain in that partition, then the VLANs you create in that partition are automatically assigned to route domain 0.

About bridging VLAN and VXLAN networks

You can configure Virtual eXtended LAN (VXLAN) on a BIG-IP® system to enable a physical VLAN to communicate with virtual machines (VMs) in a virtual network.




Figure 16: The VXLAN gateway

When you configure a BIG-IP system as an L2 VXLAN gateway, the BIG-IP system joins the configured multicast group, and can forward both unicast and multicast or broadcast frames on the virtual network. The BIG-IP system learns about MAC address and VTEP associations dynamically, thus avoiding unnecessary transmission of multicast traffic.

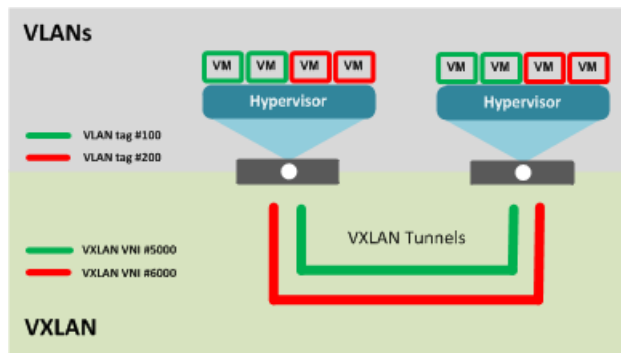


Figure 17: Multiple VXLAN tunnels

About VXLAN multicast configuration

In a VMware vSphere 5.1 environment, you can configure VXLAN without knowing all the remote tunnel endpoints. The BIG-IP[®] system uses multicast flooding to learn unknown and broadcast frames. VXLAN can extend the virtual network across a set of hypervisors, providing L2 connectivity among the hosted virtual machines (VMs). Each hypervisor represents a VXLAN tunnel endpoint (VTEP). In this environment, you can configure a BIG-IP system as an L2 VXLAN gateway device to terminate the VXLAN tunnel and forward traffic to and from a physical network.

WCCPv2

About WCCPv2 redirection on the BIG-IP system

TMOS[®] includes support for Web Cache Communication Protocol version 2 (WCCPv2). *WCCPv2* is a content-routing protocol developed by Cisco[®] Systems. It provides a mechanism to redirect traffic flows in real time. The primary purpose of the interaction between WCCPv2-enabled routers and a BIG-IP[®] system is to establish and maintain the transparent redirection of selected types of traffic flowing through those routers.

To use WCCPv2, you must enable WCCPv2 on one or more routers connected to the BIG-IP[®] system, and configure a service group on the BIG-IP system that includes the router information. The BIG-IP system then receives all the network traffic from each router in the associated service group, and determines both the traffic to optimize and the traffic to which to apply a service.

In configuring WCCPv2 on a network, you define a *service group* on the BIG-IP system, which is a collection of WCCPv2 services configured on the BIG-IP system. A WCCPv2 *service* in this context is a set of redirection criteria and processing instructions that the BIG-IP system applies to any traffic that a router in the service group redirects to the BIG-IP system. Each service matches a service identifier on the router.

The following illustration shows a one-arm configuration on one side of the WAN and an inline (bridge) configuration on the other side.

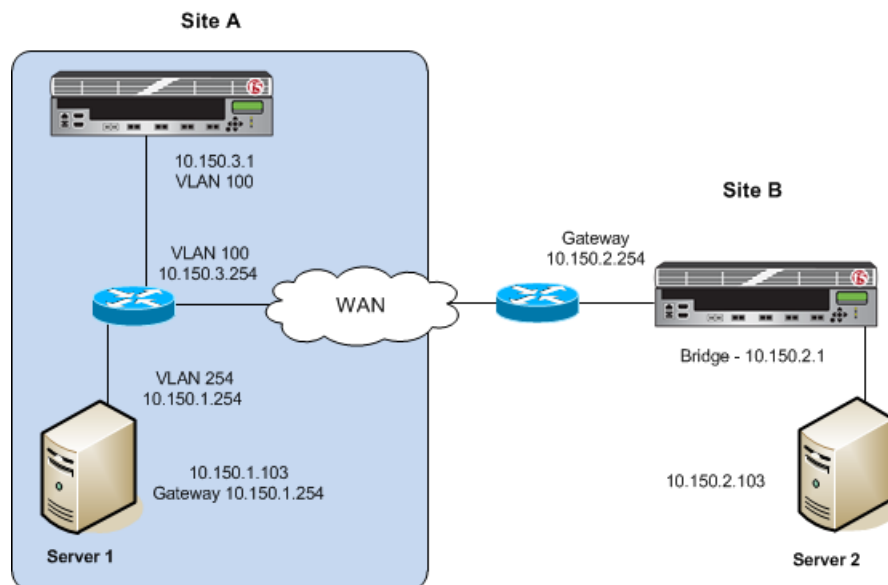


Figure 18: Example of a one-arm configuration

A common deployment of the WCCPv2 protocol

In certain cases, it is not advantageous or even possible to deploy the BIG-IP[®] system inline. For example, in the case of a collapsed backbone where the WAN router and the LAN switch are in one physical device, you might not be able to deploy the BIG-IP system inline.

If you choose not to deploy the BIG-IP system inline, you can use a one-arm deployment. In a *one-arm deployment*, the BIG-IP system has a single (hence, one-arm) connection to the WAN router or LAN switch. The WAN router (or switch) redirects all relevant traffic to the BIG-IP system. In this configuration, the WAN router typically uses Web Cache Communication Protocol version 2 (WCCPv2) to redirect traffic to the BIG-IP system.




Figure 19: Network topology for a one-arm connection

The traffic flow sequence in this illustration is as follows:

1. The client initiates a session.
2. A WAN router redirects traffic to the BIG-IP system.
3. The BIG-IP1 processes traffic and sends it back to the WAN router.
4. The WAN router forwards traffic across the WAN.

Failsafe

About system fail-safe

When you configure system fail-safe, the BIG-IP system monitors various hardware components, as well as the heartbeat of various system services, and can take action if the system detects a heartbeat failure.

You can configure the BIG-IP system to monitor the switch board component and then take some action if the BIG-IP system detects a failure. Using the BIG-IP Configuration utility, you can specify the action that you want the BIG-IP system to take when the component fails. Possible actions that the BIG-IP system can take are:

- Reboot the BIG-IP system.
- Restart all system services.
- Go offline.
- Go offline and cancel the TMM service.
- Fail over and restart TMM.

You configure system fail-safe from the **System > High Availability** screen.

About VLAN fail-safe

For maximum reliability, the BIG-IP[®] system supports failure detection on all VLANs. When you configure the fail-safe option for a VLAN, the BIG-IP system monitors network traffic going through that VLAN. If the BIG-IP system detects a loss of traffic on the VLAN and the fail-safe timeout period has elapsed, the BIG-IP system attempts to generate traffic by issuing ARP requests to nodes accessible through the VLAN. The BIG-IP system also generates an ARP request for the default route, if the default router is accessible from the VLAN. Failover is averted if the BIG-IP system is able to send and receive any traffic on the VLAN, including a response to its ARP request.

For a redundant system configuration, if the BIG-IP system does not receive traffic on the VLAN before the timeout period expires, the system can initiate failover to another device group member, reboot, or restart all system services. For a single device configuration, the system can either reboot or restart all system services. The default action for both configurations is **Reboot**.

Warning: *You should configure the fail-safe option on a VLAN only after the BIG-IP system is in a stable production environment. Otherwise, routine network changes might cause failover unnecessarily.*

Each interface card installed on the BIG-IP system is typically mapped to a different VLAN. Thus, when you set the fail-safe option on a particular VLAN, you need to know the interface to which the VLAN is mapped. You can use the BIG-IP Configuration utility to view VLAN names and their associated interfaces.

There are two ways to configure VLAN fail-safe in the BIG-IP Configuration utility: from the **System > High Availability > VLANs > Add VLAN** screen or from the **Network > VLANs > New VLAN** screen.

About gateway fail-safe

One type of network failure detection is known as gateway fail-safe, which applies to redundant system configurations only. *Gateway fail-safe* monitors traffic between an active BIG-IP® system in a device group and a pool containing a gateway router. You configure the gateway fail-safe feature if you want the BIG-IP system to take an action, such as failover, whenever some number of gateway routers in a pool of routers becomes unreachable.

You can configure gateway fail-safe using the BIG-IP Configuration utility. Configuring gateway fail-safe means designating a pool of routers as a gateway fail-safe pool. When you designate a pool as a gateway fail-safe pool, you provide the following information:

- The name of the pool
- The name of a BIG-IP device in a device group (either the local device or any other device group member)
- The minimum number of gateway pool members that must be available to avoid the designated action
- The action that the BIG-IP system should take when the number of available gateway pool members drops below the designated threshold. The default value is **Failover**.

After you configure gateway fail-safe, specifying an action of **Failover**, the named BIG-IP device (and only that device) fails over to another device group member whenever the number of available pool members falls below the specified threshold. Although all device group members share their pool configurations, each device ignores any gateway fail-safe configuration that does not specify itself as the device associated with the specified gateway pool.

You configure gateway fail-safe from the **System > High Availability > Fail-safe > Gateway** screen.

Legal Notices and Acknowledgments

Legal Notices

Publication Date

This document was published on November 28, 2017.

Publication Number

MAN-0378-04

Copyright

Copyright © 2013-2016, 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

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

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:

<http://www.f5.com/about/guidelines-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.

Acknowledgments

This product includes software developed by Gabriel Forté.

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, <http://www.and.com>.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

In the following statement, *This software* refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. *Similar operating systems* includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at <http://www.perl.com>.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes software with glib library utility functions, which is protected under the GNU Public License.

This product includes software with grub2 bootloader functions, which is protected under the GNU Public License.

This product includes software with the Intel Gigabit Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes software with the Intel 10 Gigabit PCI Express Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes RRDtool software developed by Tobi Oetiker (<http://www.rrdtool.com/index.html>) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (<http://www.nominum.com>).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes software under license from Qosmos (www.qosmos.com).

This product includes software developed by Andrew Tridgell, which is protected under the GNU Public License, copyright ©1992-2000.

This product includes software developed by Jeremy Allison, which is protected under the GNU Public License, copyright ©1998.

This product includes software developed by Guenther Deschner, which is protected under the GNU Public License, copyright ©2008.

This product includes software developed by www.samba.org, which is protected under the GNU Public License, copyright ©2007.

This product includes software from Allan Jardine, distributed under the MIT License.

This product includes software from Trent Richardson, distributed under the MIT License.

This product includes vmbus drivers distributed by Microsoft Corporation.

This product includes software from Cavium.

This product includes software from Webroot, Inc.

This product includes software from Maxmind, Inc.

This product includes software from OpenVision Technologies, Inc. Copyright ©1993-1996, OpenVision Technologies, Inc. All Rights Reserved.

This product includes software developed by Matt Johnson, distributed under the MIT License. Copyright ©2012.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software from NLnetLabs. Copyright ©2001-2006. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of NLnetLabs nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes unbound software from NLnetLabs. Copyright ©2007. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of NLnetLabs nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes GRand Unified Bootloader (GRUB) software developed under the GNU Public License, copyright ©2007.

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes gd-libgd library software developed by the following in accordance with the following copyrights:

- Portions copyright ©1994, 1995, 1996, 1997, 1998, 2000, 2001, 2002 by Cold Spring Harbor Laboratory. Funded under Grant P41-RR02188 by the National Institutes of Health.
- Portions copyright ©1996, 1997, 1998, 1999, 2000, 2001, 2002 by Boutell.Com, Inc.
- Portions relating to GD2 format copyright ©1999, 2000, 2001, 2002 Philip Warner.
- Portions relating to PNG copyright ©1999, 2000, 2001, 2002 Greg Roelofs.
- Portions relating to gdttf.c copyright ©1999, 2000, 2001, 2002 John Ellson (ellson@lucent.com).
- Portions relating to gdf.c copyright ©2001, 2002 John Ellson (ellson@lucent.com).
- Portions copyright ©2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007 2008 Pierre-Alain Joye (pierre@libgd.org).
- Portions relating to JPEG and to color quantization copyright ©2000, 2001, 2002, Doug Becker and copyright ©1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, Thomas G. Lane. This software is based in part on the work of the Independent JPEG Group.
- Portions relating to WBMP copyright 2000, 2001, 2002 Maurice Szmurlo and Johan Van den Brande. Permission has been granted to copy, distribute and modify gd in any context without fee, including a commercial application, provided that this notice is present in user-accessible supporting documentation.

This product includes software developed by Oracle America, Inc. Copyright ©2012.

1. **Java Technology Restrictions.** Licensee shall not create, modify, change the behavior of, or authorize licensees of licensee to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun" or similar convention as specified by Oracle in any naming convention designation. In the event that Licensee creates an additional API(s) which: (a) extends the functionality of a Java Environment; and (b) is exposed to third party software developers for the purpose of developing additional software which invokes such additional API, Licensee must promptly publish broadly an accurate specification for such API for free use by all developer.
2. **Trademarks and Logos.** This License does not authorize an end user licensee to use any Oracle America, Inc. name, trademark, service mark, logo or icon. The end user licensee acknowledges that Oracle owns the Java trademark and all Java-related trademarks, logos and icon including the Coffee Cup and Duke ("Java Marks") and agrees to: (a) comply with the Java Trademark Guidelines at <http://www.oracle.com/html/3party.html>; (b) not do anything harmful to or inconsistent with Oracle's rights in the Java Marks; and (c) assist Oracle in protecting those rights, including assigning to Oracle any rights acquired by Licensee in any Java Mark.
3. **Source Code.** Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of your license. Source code may not be redistributed unless expressly provided for in the terms of your license.
4. **Third Party Code.** Additional copyright notices and license terms applicable to portion of the Software are set forth in the THIRDPARTYLICENSEREADME.txt file.
5. **Commercial Features.** Use of the Commercial Features for any commercial or production purpose requires a separate license from Oracle. "Commercial Features" means those features identified in Table I-I (Commercial Features In Java SE Product Editions) of the Software documentation accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

This product includes utilities developed by Linus Torvalds for inspecting devices connected to a USB bus.

This product includes perl-PHP-Serialization software, developed by Jesse Brown, copyright ©2003, and distributed under the Perl Development Artistic License (<http://dev.perl.org/licenses/artistic.html>).

This product includes software developed by members of the CentOS Project under the GNU Public License, copyright ©2004-2011 by the CentOS Project.

This product includes software developed by members of the OpenJDK Project under the GNU Public License Version 2, copyright ©2012 by Oracle Corporation.

This product includes software developed by The VMWare Guest Components Team under the GNU Public License Version 2, copyright ©1999-2011 by VMWare, Inc.

This product includes software developed by The Netty Project under the Apache Public License Version 2, copyright ©2008-2012 by The Netty Project.

This product includes software developed by Stephen Colebourne under the Apache Public License Version 2, copyright ©2001-2011 Joda.org.

This product includes software developed by the GlassFish Community under the GNU Public License Version 2 with classpath exception, copyright ©2012 Oracle Corporation.

This product includes software developed by the Mort Bay Consulting under the Apache Public License Version 2, copyright ©1995-2012 Mort Bay Consulting.

This product contains software developed by members of the Jackson Project under the GNU Lesser General Public License Version 2.1, ©2007 – 2012 by the Jackson Project”.

This product contains software developed by QOS.ch under the MIT License, ©2004 – 2011 by QOS.ch.

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Daniel Stenberg. Copyright ©1996 - 2012, Daniel Stenberg, (daniel@haxx.se). All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

This product includes software licensed from Rémi Denis-Courmont under the GNU Library General Public License. Copyright ©2006 - 2011.

This product includes software developed by jQuery Foundation and other contributors, distributed under the MIT License. Copyright ©2014 jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Trent Richardson, distributed under the MIT License. Copyright ©2012 jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Allan Jardine, distributed under the MIT License. Copyright ©2008 - 2012, Allan Jardine, all rights reserved, jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Douglas Gilbert. Copyright ©1992 - 2012 The FreeBSD Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE FREEBSD PROJECT "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FREEBSD PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The views and conclusions contained in the software and documentation are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the FreeBSD Project.

This product includes software developed as open source software. Copyright ©1994 - 2012 The FreeBSD Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). Copyright ©1998 - 2011 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software licensed from William Ferrell, Selene Scriven and many other contributors under the GNU General Public License, copyright ©1998 - 2006.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the

software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product includes software developed by Brian Gladman, Worcester, UK Copyright ©1998-2010. All rights reserved. The redistribution and use of this software (with or without changes) is allowed without the payment of fees or royalties provided that:

- source code distributions include the above copyright notice, this list of conditions and the following disclaimer;
- binary distributions include the above copyright notice, this list of conditions and the following disclaimer in their documentation.

This software is provided 'as is' with no explicit or implied warranties in respect of its operation, including, but not limited to, correctness and fitness for purpose.

This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory. Copyright ©1990-1994 Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory.
4. Neither the name of the University nor of the Laboratory may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software developed by Sony Computer Science Laboratories Inc. Copyright © 1997-2003 Sony Computer Science Laboratories Inc. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY SONY CSL AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL SONY CSL OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This software incorporates JFreeChart, ©2000-2007 by Object Refinery Limited and Contributors, which is protected under the GNU Lesser General Public License (LGPL).

This product contains software developed by the Mojarra project. Source code for the Mojarra software may be obtained at <https://jaserverfaces.dev.java.net/>.

This product includes software developed by McAfee®.

This product includes software developed by Ian Gulliver ©2006, which is protected under the GNU General Public License, as published by the Free Software Foundation.

This product contains software developed by the RE2 Authors. Copyright ©2009 The RE2 Authors. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes the Zend Engine, freely available at <http://www.zend.com>.

This product includes software developed by Digital Envoy, Inc.

This product contains software developed by NuSphere Corporation, which is protected under the GNU Lesser General Public License.

This product contains software developed by Erik Arvidsson and Emil A Eklund.

This product contains software developed by Aditus Consulting.

This product contains software developed by Dynarch.com, which is protected under the GNU Lesser General Public License, version 2.1 or later.

This product contains software developed by InfoSoft Global (P) Limited.

This product includes software written by Steffen Beyer and licensed under the Perl Artistic License and the GPL.

This product includes software written by Makamaka Hannyaharamitu ©2007-2008.

Rsync was written by Andrew Tridgell and Paul Mackerras, and is available under the GNU Public License.

This product includes Malloc library software developed by Mark Moraes. (©1988, 1989, 1993, University of Toronto).

This product includes open SSH software developed by Tatu Ylonen (ylo@cs.hut.fi), Espoo, Finland (©1995).

This product includes open SSH software developed by Niels Provos (©1999).

This product includes SSH software developed by Mindbright Technology AB, Stockholm, Sweden, www.mindbright.se, info@mindbright.se (©1998-1999).

This product includes free SSL software developed by Object Oriented Concepts, Inc., St. John's, NF, Canada, (©2000).

This product includes software developed by Object Oriented Concepts, Inc., Billerica, MA, USA (©2000).

This product includes free software developed by ImageMagick Studio LLC (©1999-2011).

This product includes software developed by Bob Withers.

This product includes software developed by Jean-Loup Gailly and Mark Adler.

This product includes software developed by Markus FXJ Oberhumer.

This product includes software developed by Guillaume Fihon.

This product includes QPDF software, developed by Jay Berkenbilt, copyright ©2005-2010, and distributed under version 2 of the OSI Artistic License (<http://www.opensource.org/licenses/artistic-license-2.0.php>).

This product includes JZlib software, Copyright © 2000-2011 ymnk, JCraft, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes Apache Lucene software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes Apache MINA software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes OData4J software, distributed under the Apache License version 2.0.

This product includes software developed by the Visigoth Software Society (<http://www.visigoths.org/>).

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes software developed by Addy Osmani, and distributed under the MIT license. Copyright © 2012 Addy Osmani.

This product includes software developed by Charles Davison, and distributed under the MIT license. Copyright © 2013 Charles Davison.

This product includes software developed by The Dojo Foundation, and distributed under the MIT license. Copyright © 2010-2011, The Dojo Foundation.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes software developed by Douglas Crockford, douglas@crockford.com.

This product includes ec2-tools software, copyright © 2008, Amazon Web Services, and licensed under the Amazon Software License. A copy of the License is located at <http://aws.amazon.com/asl/>.

This product includes the ixgbev Intel Gigabit Linux driver, Copyright © 1999 - 2012 Intel Corporation, and distributed under the GPLv2 license, as published by the Free Software Foundation.

This product includes Apache Ant software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes libwebp software. Copyright © 2010, Google Inc. All rights reserved.

This product includes isc-dhcp software. Copyright © 2004-2013 by Internet Systems Consortium, Inc. ("ISC"); Copyright © 1995-2003 by Internet Software Consortium.

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

This product includes jQuery Sparklines software, developed by Gareth Watts, and distributed under the new BSD license.

This product includes jsdiff software, developed by Chas Emerick, and distributed under the BSD license.

This product includes winston software, copyright © 2010, by Charlie Robbins.

This product includes Q software developed by Kristopher Michael Kowal, and distributed under the MIT license. Copyright © 2009-2013 Kristopher Michael Kowal.

This product includes SlickGrid software developed by Michael Liebman, and distributed under the MIT license.

This product includes JCraft Jsch software developed by Atsuhiko Yamanaka, copyright © 2002-2012 Atsuhiko Yamanaka, JCraft, Inc. All rights reserved.

This product includes DP_DateExtensions software developed by Jim Davis, Copyright © 1996-2004, The Depressed Press of Boston (depressedpres.com). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the DEPRESSED PRESS OF BOSTON (DEPRESSEDPRESS.COM) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

All code not authored by the Depressed Press is attributed (where possible) to its rightful owners/authors, used with permission and should be assumed to be under copyright restrictions as well.

This product includes Boost libraries, which are distributed under the Boost license (http://www.boost.org/LICENSE_1_0.txt).

This product includes Angular software developed by Google, Inc., <http://angularjs.org>, copyright © 2010-2012 Google, Inc., and distributed under the MIT license.

This product includes node.js software, copyright © Joyent, Inc. and other Node contributors. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,

FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes crypto.js software, copyright © 2009-2013, Jeff Mott, and distributed under the BSD New license.

This product includes the epoxy.js library for backbone, copyright © 2012-2013 Greg MacWilliam. (<http://epoxyjs.org>)

This product includes Javamail software, copyright © 1997-2013 Oracle and/or its affiliates, all rights reserved; and copyright © 2009-2013 Jason Mehrens, all rights reserved. This software is distributed under the GPLv2 license.

This product includes Leaflet software, copyright © 2010-2014, Vladimir Agafonkin, and copyright © 2010-2011, CloudMade; all rights reserved. This software is distributed under the BSD license.

This product includes underscore software, copyright © 2009-2014 Jeremy Ashkenas, DocumentCloud, and Investigative Reporters & Editors.

This product MAY include Intel SSD software subject to the following license; check your hardware specification for details.

1. LICENSE. This Software is licensed for use only in conjunction with Intel solid state drive (SSD) products. Use of the Software in conjunction with non-Intel SSD products is not licensed hereunder. Subject to the terms of this Agreement, Intel grants to You a nonexclusive, nontransferable, worldwide, fully paid-up license under Intel's copyrights to:

- copy the Software onto a single computer or multiple computers for Your personal, noncommercial use; and
- make appropriate back-up copies of the Software, for use in accordance with Section 1a) above.

The Software may contain the software or other property of third party suppliers, some of which may be identified in, and licensed in accordance with, any enclosed "license.txt" file or other text or file.

Except as expressly stated in this Agreement, no license or right is granted to You directly or by implication, inducement, estoppel or otherwise. Intel will have the right to inspect or have an independent auditor inspect Your relevant records to verify Your compliance with the terms and conditions of this Agreement.

2. RESTRICTIONS. You will not:

- a. copy, modify, rent, sell, distribute or transfer any part of the Software, and You agree to prevent unauthorized copying of the Software; and,
- b. reverse engineer, decompile, or disassemble the Software; and,
- c. sublicense or permit simultaneous use of the Software by more than one user; and,
- d. otherwise assign, sublicense, lease, or in any other way transfer or disclose Software to any third party, except as set forth herein; and,
- e. subject the Software, in whole or in part, to any license obligations of Open Source Software including without limitation combining or distributing the Software with Open Source Software in a manner that subjects the Software or any portion of the Software provided by Intel hereunder to any license obligations of such Open Source Software. "Open Source Software" means any software that requires as a condition of use, modification and/or distribution of such software that such software or other software incorporated into, derived from or distributed with such software:
 - a. be disclosed or distributed in source code form; or
 - b. be licensed by the user to third parties for the purpose of making and/or distributing derivative works; or

- c. be redistributable at no charge.

Open Source Software includes, without limitation, software licensed or distributed under any of the following licenses or distribution models, or licenses or distribution models substantially similar to any of the following:

- a. GNU's General Public License (GPL) or Lesser/Library GPL (LGPL),
- b. the Artistic License (e.g., PERL),
- c. the Mozilla Public License,
- d. the Netscape Public License,
- e. the Sun Community Source License (SCSL),
- f. vi) the Sun Industry Source License (SISL),
- g. vii) the Apache Software license, and
- h. viii) the Common Public License (CPL).

3. **OWNERSHIP OF SOFTWARE AND COPYRIGHTS.** Title to all copies of the Software remains with Intel or its suppliers. The Software is copyrighted and protected by the laws of the United States and other countries, and international treaty provisions. You may not remove any copyright notices from the Software. Intel may make changes to the Software, or to materials referenced therein, at any time and without notice, but is not obligated to support or update the Software. Except as otherwise expressly provided, Intel grants no express or implied right or license under Intel patents, copyrights, trademarks, or other intellectual property rights.
4. **Entire Agreement.** This Agreement contains the complete and exclusive statement of the agreement between You and Intel and supersedes all proposals, oral or written, and all other communications relating to the subject matter of this Agreement. Only a written instrument duly executed by authorized representatives of Intel and You may modify this Agreement.
5. **LIMITED MEDIA WARRANTY.** If the Software has been delivered by Intel on physical media, Intel warrants the media to be free from material physical defects for a period of ninety (90) days after delivery by Intel. If such a defect is found, return the media to Intel for replacement or alternate delivery of the Software as Intel may select.
6. **EXCLUSION OF OTHER WARRANTIES.** EXCEPT AS PROVIDED ABOVE, THE SOFTWARE IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND, INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. Intel does not warrant or assume responsibility for any errors, the accuracy or completeness of any information, text, graphics, links or other materials contained within the Software.
7. **LIMITATION OF LIABILITY.** IN NO EVENT WILL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION OR LOST INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS PROHIBIT EXCLUSION OR LIMITATION OF LIABILITY FOR IMPLIED WARRANTIES OR CONSEQUENTIAL OR INCIDENTAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU MAY ALSO HAVE OTHER LEGAL RIGHTS THAT VARY FROM JURISDICTION TO JURISDICTION.
8. **TERMINATION OF THIS AGREEMENT.** Intel may terminate this Agreement at any time if You violate its terms. Upon termination, You will immediately destroy the Software or return all copies of the Software to Intel.
9. **APPLICABLE LAWS.** Claims arising under this Agreement will be governed by the laws of Delaware, excluding its principles of conflict of laws and the United Nations Convention on Contracts for the Sale of Goods. You may not export the Software in violation of applicable export laws and regulations. Intel is not obligated under any other agreements unless they are in writing and signed by an authorized representative of Intel.

- 10. GOVERNMENT RESTRICTED RIGHTS.** The Software is provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR52.227-14 and DFAR252.227-7013 et seq. or their successors. Use of the Software by the Government constitutes acknowledgment of Intel's proprietary rights therein. Contractor or Manufacturer is Intel Corporation, 2200 Mission College Blvd., Santa Clara, CA 95054.

Index

A

- access level
 - definition 40
 - enable/disable 40
 - read 40
 - update 40
 - write 40
- action setting
 - for packet filter rules 76
- Active Directory server information 49
- address space
 - definition 69
- admin account
 - configuration of 40
 - password 40
- administrative account passwords 24
- administrative accounts
 - creation of 35
- administrative boundaries
 - setting 30
- administrative partition
 - assign partition access 30
 - create user accounts 30
 - definition 39
- administrative partition Common 29
- administrative partitions
 - definition 29
 - setting current 30
- administrative traffic
 - authenticating 49
- AOM
 - and console output 11
- archive
 - backup 25
 - propagate data 25
 - restoration 26
 - saving 26
- audit logging
 - 60
 - disable 61
 - enable 61
- audit user access 55
- authentication lockout 42
- authorization
 - for groups of user accounts 53
- auto last hop 101
- automatic link selection 94

B

- BIG-IP system
 - description 9
- BIG-IP system services
 - managing 27
- boundaries, administrative
 - setting 30
- bridge 81

- bridge priority
 - for spanning tree 86

C

- certificate properties
 - list of 19
- certificates
 - creating 16
 - exporting SSL 18
 - importing 18
 - requesting from CAs 17
- CMP Hash setting
 - for VLANs 101
- code expansion
 - syslog messages 60
- configuration properties 11
- console output
 - saving 11
- current partition 30

D

- data group files
 - address 21
 - example 21
 - integer 21
 - string 21
- delete local user account 47
- device properties 11
- DHCP protocol 12
- DHCP server 12
- digital certificates
 - about 15
 - importing 18
- distribution function
 - for frames 95
- DNS 12
- Domain Name System 12
- DSA encryption algorithm 16
- DSA signature algorithm 15

E

- ECDSA encryption algorithm 15–16
- edge port 88
- Elliptic Curve Digital Signature Algorithm 15–16
- enforcement restriction 41
- external file management
 - data group files 21
 - iFiles 21
- external monitor program file 21

F

- fail-safe
 - and VLANs 101, 109

fail-safe (*continued*)
 for gateway pools 110
 for system daemons 109

filter expression
 about creation 77
 definition 77

flow control property 65

folders
 defined 33
 object naming 33
 object referencing 33
 relationship with partitions 33

forward delay option 82

frames
 distributing 95
 transmitting 95

G

Gateway fail-safe 110

general properties
 host IP address 23
 host name 23
 management port and TMM 23
 management port configuration 23
 time zone 23

global exemptions
 for IP address 75
 for MAC address 75
 for protocol 75
 for VLAN 74–75

global properties
 controlling unhandled packets 73
 enabling packet filter 73
 options 73

global timer 82

global timers 82

granular synchronization
 with folders 33

H

hash algorithms 95

hello time option 82

I

iFiles 21

instance ID 86

instance member
 defined 86

interface ID
 viewing 89

interface mirroring
 definition 66

interface name
 example 64

interface naming conventions 64

interface priority 90

interface setting 92

invalid object referencing 31

IP address 75

IP geolocation 11

J

jumbo frames 101

L

LACP

 and timeout 94
 enabling 93
 overview 92

LACP mode 93

last hops

 automatically mapping 13

Layer 2 forwarding table

 and cached records 13
 maintaining 102

LDAP server information

 specifying 49

link aggregation control protocol 92

link layer discovery protocol
 about 63

link selection policy

 and auto option 94
 and maximum bandwidth 94–95
 automatic 94

LLDP

 about 63
 neighbor settings 66
 property 66

LLDP attributes

 about 66

local area network 97

local syslog logging 59

local traffic

 properties 12

local user account

 creating 43
 creation 43
 delete 47
 deleting 47
 management of 40
 modification 45
 modifying properties 46
 view 44

local user accounts

 displaying 44
 properties 45

logging

 audit 60–61
 code expansion 60
 level setting 59
 local storage 58
 local syslog 59
 local traffic events 60
 message types 57
 overview 57
 packet filter events 60
 remote 61
 remote storage 57
 syslog 60

- logging (*continued*)
 - syslog-ng 61
 - system events 60
- login failures 42
- log level
 - setting 59
- log level setting 59
- log message
 - remote storage 57
- log messages
 - 77
 - local storage 58

M

- MAC address
 - and packet filtering 75
 - and self IP addresses 70
- maintenance mode
 - and connections 13
- maximum age option
 - default 82
 - valid range 82
- maximum bandwidth
 - for link selection policy 95
- maximum transmission units
 - and default value 101
- MCP audit logging
 - definition 61
- MSTP 80
- MSTP global properties 83
- MTU discovery
 - and TCP packet fragmentation 13
- MTU setting
 - and default value 101

N

- neighbor settings
 - LLDP 66
- network
 - configuring one-arm deployment 108
- Network Time Protocol 12
- NTP 12

O

- object name 33
- object referencing
 - 31
 - between folders 33
- one-arm deployment
 - overview 108
- overlay networks
 - using VXLAN 106

P

- packet filtering
 - about 73
 - global settings 73

- packet filter rules
 - about 73
 - and action setting 76
 - for VLANs 77
 - logging 77
 - ordering 76
- partition
 - Common 33
 - relationship with folders 33
- partition Common 29, 31
- partitions
 - setting current 30
- password
 - administrative account 24
- password policies 42
- password policy restrictions 41
- path cost 90
- pause frame
 - definition 65
- platform properties 23
- policy
 - for link selection 94
- policy restriction 41
- polling interval
 - on VLANs for sFlow 102
- port-based access
 - to VLANs 99
- port roles
 - for spanning tree instance 89
- port states
 - blocking 89
 - forwarding 89
 - learning 89
- private keys
 - types of 15
- protocol degradation 80
- protocols setting 75
- proxied ARP requests
 - and VLAN groups 104

R

- RADIUS protocol
 - for remote server authentication 50
- rate class setting 77
- reaper high-water mark
 - described 13
- reaper low-water mark
 - described 13
- redundant device properties 24
- remote account
 - changing default 53
- remote authentication server 48
- remote user account
 - displaying 54
 - displaying properties 54
 - management 48
 - non-default user roles 54
- remote user accounts
 - authorization 52
 - default 52
 - viewing 54

- root account
 - properties, modifying 47
- root bridge 86
- route domains
 - and VLAN assignment 105
- routing
 - one-arm mode 107
- RSA encryption algorithm 15
- RSTP 79

S

- sampling rate
 - on VLANs for sFlow 102
- secure network address translation 70
- secure password policy configuration
 - enforcement 41
 - policy 41
 - restrictions 41
- secure sockets layer 15
- self IP address
 - and MAC addresses 70
 - and VLAN 69
 - definition 69
 - properties 70
 - purpose 69
 - types defined 69
- self-signed certificates
 - creating 16
- serial console output
 - saving 11
- sFlow
 - and VLANs 102
- SNAT
 - automapping 70
- source checking
 - for VLANs 100
- spanning tree
 - deleting an instance for MSTP 87
- spanning tree instance
 - modifying 86
 - viewing 86
- spanning tree instances
 - creating MSTP 84
 - listing 84
 - managing 83
- spanning tree mode
 - disabled 81
 - MSTP 81
 - pass through 81
 - RSTP 81
 - STP 81
- spanning tree protocol reset
 - about 88
- spanning tree protocols
 - and global timers 82
 - and legacy bridges 80
 - and link type 87
 - and MSTP 80
 - and RSTP 79
 - and STP edge port 88
 - configuration 81

- spanning tree protocols (*continued*)
 - definition 79
 - managing interface 88
- SSH access configuration 24
- SSL certificate file 22
- SSL certificates
 - about 15
 - importing 18
 - managing 16
- standard user account 35
- STP
 - managing 87
- STP edge port
 - detecting 88
 - setting 88
- syslog
 - existing configuration 57
 - local logging 59
 - log messages 60
- syslog-ng
 - remote logging 61
- system fail-safe 109
- system interface management
 - about 63
- system interface properties
 - about 64
 - about LLDP 66
 - Fixed Requested Media 65
 - flow control 65
 - interface naming conventions 64
 - interface state 65
 - LLDP Attributes 66
 - viewing 65
- system interfaces
 - and related configuration tasks 67
 - and TMM switch 63
 - and trunk feature 67
 - assigning to VLAN 67
 - definition 63
 - for STP 87
 - for VLAN 99
 - mirroring 63
 - viewing 65
- system interface state 65
- system maintenance account 35

T

- TACACS+ protocol
 - for remote server authentication 51
- tag-based access
 - to VLANs 99
- timeout
 - value for LACP 94
- TMOS
 - configuration properties 11
- traffic bridging
 - and standby units 104
 - and VLAN groups 104
- traffic redirection
 - about WCCPv2 107
- transmit hold count option 83

- transparency mode
 - opaque [103](#)
 - translucent [103](#)
 - transparent [103](#)
- trunks
 - and interface [92](#)
 - and LACP [92](#)
 - and link aggregation [91](#)
 - and name [92](#)
 - explained [91](#)

U

- unmatched packets
 - and system behavior [13](#)
- user access
 - audit [55](#)
- user accounts
 - administrative access [39](#)
 - authenticating [49](#)
 - authorization [52](#)
 - authorization for groups [53](#)
 - creation of [35](#)
 - group authorization [53](#)
 - management of [35](#)
 - operator role [39](#)
 - partition access property [39](#)
 - partitioned object [30](#)
 - purpose of [35](#)
 - remote [48, 52](#)
 - standard [35](#)
 - system maintenance [35](#)
 - to control access [36](#)
 - types [35](#)
- user administration properties [24](#)
- user authentication lockout [42](#)
- user lockout [42](#)
- user manager role [44](#)
- user roles
 - access level [40](#)
 - administrative [36](#)
 - administrator [38](#)
 - default [38](#)
 - definition [36](#)
 - described [36](#)
 - for local accounts [43](#)
 - no access [38](#)
 - permissions [36](#)
 - root [38](#)

V

- valid object referencing [31](#)
- Virtual eXtended LAN, *See* VXLAN
- virtual local area network [97](#)
- VLAN ++
 - and system interfaces [99](#)
- VLAN-based fail-safe [101](#)
- VLAN behavior
 - on the BIG-IP system [97](#)
- VLAN configuration
 - and external [97](#)
 - and internal [97](#)
- VLAN fail-safe [109](#)
- VLAN group ID [103](#)
- VLAN group names [103](#)
- VLAN groups
 - about [102](#)
 - and keepalives [104](#)
 - and proxied ARP requests [104](#)
- VLAN-keyed connections
 - enabling [13](#)
- VLAN name [98](#)
- VLANs
 - and fail-safe [101](#)
 - and group self IP address [105](#)
 - and name [98](#)
 - and packet filtering [77](#)
 - and port-based access [99](#)
 - and route domain [105](#)
 - and self IP address [105](#)
 - and source checking [100](#)
 - and tag-based access [99](#)
 - and tags [98](#)
 - as last hop [101](#)
 - as STP instance [86](#)
- VLAN setting
 - and packet filter [75](#)
- VLAN tags [98](#)
- VXLAN
 - about [106](#)
 - about configuring BIG-IP system as gateway [105](#)
 - bridging with L2 VLAN network [105](#)

W

- WCCPv2
 - configuring one-arm deployment [108](#)
 - description [107](#)

