BIG-IP® Local Traffic Manager™: Monitors Reference

Version 13.1

Table of Contents

Monitors Concepts

Purpose of monitors

Monitors determine the availability and performance of devices, links, and services on a network. Health monitors check the availability. Performance monitors check the performance and load. If a monitored device, link, or service does not respond within a specified timeout period, or the status indicates that performance is degraded or that the load is excessive, the BIG-IP® system can redirect the traffic to another resource.

Benefits of monitors

Monitors gather information about your network. The information that monitors gather is available for you to view. You can use this information to troubleshoot problems and determine what resources in your network are in need of maintenance or reconfiguration.

About iCheck functionality for monitors

FTP, SMTP, POP3, and IMAP monitors provide inherent iCheck functionality, which reduces the load on BIG-IP systems and improves sustained monitor performance. Additionally, iCheck functionality provides smoother performance characteristics as these monitors approach full capacity.

FTP monitors provide inherent iCheck functionality, which reduces the load on BIG-IP systems and improves sustained monitor performance. Additionally, iCheck functionality provides smoother performance characteristics as these monitors approach full capacity.

Methods of monitoring

The BIG-IP® Local Traffic Manager™, DNS, and Link Controller™ provide three methods of monitoring: simple monitoring, active monitoring, and passive monitoring.

Simple monitoring

Simple monitoring determines whether the status of a resource is up or down. Simple monitors do not monitor pool members (and therefore, individual protocols, services, or applications on a node), but only the node itself. The system contains three simple monitors, **Gateway ICMP**, **ICMP**, and **TCP_ECHO**.

Simple monitors work well when you only need to determine the up or down status of the following:

- A Local Traffic Manager node
- A BIG-IP-DNS or Link Controller server, virtual server, pool, pool member, or link

Active monitoring

Active monitoring checks the status of a pool member or node on an ongoing basis as specified. If a pool member or node does not respond within a specified timeout period, or the status of a node indicates that performance is degraded, the BIG-IP system can redirect the traffic to another pool member or node. There are many active monitors. Each active monitor checks the status of a particular protocol, service, or application. For example, one active monitor is **HTTP**. An **HTTP** monitor allows you to monitor the

availability of the HTTP service on a pool, pool member, or node. A **WMI** monitor allows you to monitor the performance of a node that is running the Windows® Management Instrumentation (WMI) software. Active monitors fall into two categories: Extended Content Verification (ECV) monitors for content checks, and Extended Application Verification (EAV) monitors for service checks, path checks, and application checks.

An active monitor can check for specific responses, and run with or without client traffic.

Note: *An active monitor also creates additional network traffic beyond the client request and server response and can be slow to mark a pool member as down.*

Passive monitoring

Passive monitoring occurs as part of a client request. This kind of monitoring checks the health of a pool member based on a specified number of connection attempts or data request attempts that occur within a specified time period. If, after the specified number of attempts within the defined interval, the system cannot connect to the server or receive a response, or if the system receives a bad response, the system marks the pool member as down. There is only one passive monitor, called an **Inband** monitor.

A passive monitor creates no additional network traffic beyond the client request and server response. It can mark a pool member as down quickly, as long as there is some amount of network traffic.

Note: *A passive monitor cannot check for specific responses and can potentially be slow to mark a pool member as up.*

Comparison of monitoring methods

In the short description, briefly describe the purpose and intent of the information contained in this topic. This element is an F5® requirement.

Monitoring Method	Benefits	Constraints
Simple	• Works well when you only need to determine the up or down status of a node.	• Can check the health of a node only, and not a pool member.
Active	• Can check for specific responses • Can run with or without client traffic	• Creates additional network traffic beyond the client request and server response • Can be slow to mark a pool member as down
Passive	• Creates no additional network traffic beyond the client request and server response • Can mark a pool member as down quickly, as long as there is some amount of network traffic	• Cannot check for specific responses • Can potentially be slow to mark a pool member as up

Monitor destinations

By default, the value for the **Alias Address** setting in the monitors is set to the wildcard `* Addresses`, and the **Alias Service Port** setting is set to the wildcard `* Ports`. This value causes the monitor instance

created for a pool, pool member, or node to take that node's address or address and port as its destination. You can, however, replace either or both wildcard symbols with an explicit destination value, by creating a custom monitor. An explicit value for the **Alias Address** and/or **Alias Service Port** setting is used to force the instance destination to a specific address and/or port which might not be that of the pool, pool member, or node.

The ECV monitor types HTTP, HTTPS, and TCP include the settings **Send String** and **Receive String** for the send string and receive expression, respectively.

The most common **Send String** value is `GET /`, which retrieves a default HTML page for a web site. To retrieve a specific page from a web site, you can enter a **Send String** value that is a fully qualified path name:

```
"GET /www/support/customer_info_form.html"
```

The **Receive String** value is the text string that the monitor looks for in the returned resource. The most common **Receive String** values contain a text string that is included in a particular HTML page on your site. The text string can be regular text, HTML tags, or image names.

The sample **Receive String** value below searches for a standard HTML tag:

```
"<HEAD>"
```

You can also use the default null **Receive String** value `[""]`. In this case, any content retrieved is considered a match. If both the **Send String** and **Receive String** fields are left empty, only a simple connection check is performed.

For HTTP and FTP monitor types, you can use the special values `GET` or `hurl` in place of `Send String` and `Receive String` values. For FTP monitors specifically, the GET value should specify the full path to the file to retrieve.

About monitor settings

Every monitor consists of settings with values. The settings and their values differ depending on the type of monitor. In some cases, the BIG-IP® system assigns default values. This example shows that an HTTP-type monitor has these settings and default values.

The settings specify that an HTTP type of monitor is configured to check the status of an IP address every 5 seconds, and to time out every 16 seconds. The destination IP address that the monitor checks is specified by the Alias Address setting, with the value `* All Addresses`. Thus, in the example, all IP addresses with which the monitor is associated are checked.

```
Name my_http
Type HTTP
Interval 5
Timeout 16
Transparent No
Alias Address * All Addresses
```

The settings specify that an HTTP type of monitor is configured to check the status of an IP address every 30 seconds, and to time out every 5 seconds. The destination IP address that the monitor checks is specified by the Alias Address setting, with the value `* All Addresses`. Thus, in the example, all IP addresses with which the monitor is associated are checked.

```
Name my_http
Type HTTP
Interval 30
Timeout 5
Transparent No
Alias Address * All Addresses
```

Transparent and Reverse modes

The normal and default behavior for a monitor is to ping the destination pool, pool member, or node by an unspecified route, and to mark the node `up` if the test is successful. However, with certain monitor types, you can specify a route through which the monitor pings the destination server. You configure this by specifying the Transparent or Reverse setting within a custom monitor.

Transparent setting

Sometimes it is necessary to ping the aliased destination through a transparent pool, pool member, or node. When you create a custom monitor and set the Transparent setting to Yes, the BIG-IP® system forces the monitor to ping through the pool, pool member, or node with which it is associated (usually a firewall) to the pool, pool member, or node. (That is, if there are two firewalls in a load balancing pool, the destination pool, pool member, or node is always pinged through the pool, pool member, or node specified; not through the pool, pool member, or node selected by the load balancing method.) In this way, the transparent pool, pool member, or node is tested: if there is no response, the transparent pool, pool member, or node is marked as `down`.

Common examples are checking a router, or checking a mail or FTP server through a firewall. For example, you might want to check the router address `10.10.10.53:80` through a transparent firewall `10.10.10.101:80`. To do this, you create a monitor called `http_trans` in which you specify `10.10.10.53:80` as the monitor destination address, and set the Transparent setting to Yes. Then you associate the monitor `http_trans` with the transparent pool, pool member, or node.

This causes the monitor to check the address `10.10.10 53:80` through `10.10.10.101:80`. (In other words, the BIG-IP system routes the check of `10.10.10.53:80` through `10.10.10.101:80`.) If the correct response is not received from `10.10.10.53:80`, then `10.10.10.101:80` is marked `down`.

Reverse setting

With the Reverse setting set to Yes, the monitor marks the pool, pool member, or node down when the test is successful. For example, if the content on your web site home page is dynamic and changes frequently, you may want to set up a reverse ECV service check that looks for the string "`Error`". A match for this string means that the web server was `down`.

Monitors that contain the Transparent or Reverse settings

This table shows the monitors that contain either the Transparent setting or both the Reverse and Transparent settings.

Monitor Type	Settings
TCP	Transparent and Reverse
HTTP	Transparent and Reverse
HTTPS	Transparent and Reverse
TCP Echo	Transparent
Gateway ICMP	Transparent
TCP Half Open	Transparent
ICMP	Transparent
UDP	Transparent

The Manual Resume feature

By default, when a monitor detects that a resource (that is, a node or a pool member) is unavailable, the BIG-IP® system marks the resource as `down` and routes traffic to the next appropriate resource as dictated by the active load balancing method. When the monitor next determines that the resource is available again, the BIG-IP system marks the resource as `up` and immediately considers the resource to be available for load balancing connection requests. While this process is appropriate for most resources, there are situations where you want to manually designate a resource as available, rather than allow the BIG-IP system to do that automatically. You can manually designate a resource as available by configuring the Manual Resume setting of the monitor.

For example, consider a monitor that you assigned to a resource to track the availability of an HTML file, `index.html`, for a web site. During the course of a business day, you decide that you need to restart the system that hosts the web site. The monitor detects the restart action and informs the BIG-IP system that the resource is now unavailable. When the system restarts, the monitor detects that the `index.html` file is available, and begins sending connection requests to the web site. However, the rest of the web site might not be ready to receive connection requests. Consequently, the BIG-IP system sends connection requests to the web site before the site can respond effectively.

To prevent this problem, you can configure the Manual Resume setting of the monitor. When you set the Manual Resume setting to Yes, you ensure that the BIG-IP system considers the resource to be unavailable until you manually enable that resource.

Resumption of connections

If you have a resource (such as a pool member or node) that a monitor marked as `down`, and the resource has subsequently become available again, you must manually re-enable that resource if the monitor's **Manual Resume** setting is set to Yes. Manually re-enabling the resource allows the BIG-IP® system to resume sending connections to that resource.

The procedure for manually re-enabling a resource varies depending on whether the resource is a pool, a pool member, or a node.

The Time Until Up feature

By default, the BIG-IP® system marks a pool member or node as up immediately upon receipt of the first correct response to a `ping` command.

The Time Until Up feature provides a way to adjust the default behavior. This feature allows the system to delay the marking of a pool member or node as up for some number of seconds after receipt of the first correct response. The purpose of this feature is to ensure that the monitor marks the pool member or node as up only after the pool member or node has consistently responded correctly to the BIG-IP system during the defined time period. With this feature, you ensure that a pool member or node that is available only momentarily, after sending one correct response, is not marked as up.

A Time Until Up value of `0` causes the default behavior. When the Time Until Up value is a non-0 value, the BIG-IP system marks a pool member or node as up only when all pool member or node responses during the Time Until Up period are correct.

About health and performance monitors

BIG-IP® systems use two categories of monitors: health monitors and performance monitors. You can associate monitors with the following resources:

- In Local Traffic Manager™: nodes, pools, and pool members
- In DNS: links, servers, virtual servers, pools, and pool members
- In Link Controller™: links, pools, and pool members

Category	Description
Health	Checks resources to determine if they are up and functioning for a given service.
Performance	Gathers information about resources that the system uses to dynamically load balance traffic.

When a virtual server that is being monitored by a health monitor does not respond to a probe from the BIG-IP system within a specified timeout period, the system marks the virtual server down and no longer load balances traffic to that virtual server. When the health monitor determines that the virtual server is once again responsive, the system again begins to load balance traffic to that virtual server. To illustrate, a Gateway Internet Control Message Protocol (ICMP) monitor pings a virtual server. If the monitor does not receive a response from the virtual server, the BIG-IP system marks that virtual server down. When the ping is successful, the system marks the virtual server up.

When a server that is being monitored by a performance monitor displays a degradation in performance, the BIG-IP system redirects traffic to other resources until the performance of the server returns to normal. To illustrate, an SNMP DCA monitor checks the current CPU, memory, and disk usage of a server that is running an SNMP data collection agent, and then dynamically load balances traffic based on the performance of the server.

When a server that is being monitored by a performance monitor displays a degradation in performance, the BIG-IP system redirects traffic to other resources until the performance of the server returns to normal. To illustrate, an SNMP Link monitor checks the current CPU, memory, and disk usage of a server that is running an SNMP data collection agent, and then dynamically load balances traffic based on the performance of the server.

About address check monitors

An *address check monitor* provides a simple verification of an address on a network. This type of monitor sends a request to an IP address. When a response is received, the test is successful.

With BIG-IP® DNS Link Controller™, you can use an address check monitor to monitor a virtual server, a server (which includes all of the virtual servers on a specified server), a pool member, a pool (which includes all of the pool members of a specified pool), or a link. This monitor uses the Gateway Internet Control Message Protocol (ICMP) to perform a simple resource check. The check is successful if the monitor receives a response to an `ICMP_ECHO` datagram.

The following illustration depicts a DNS using a **Gateway ICMP** monitor to verify an IP address for a virtual server.

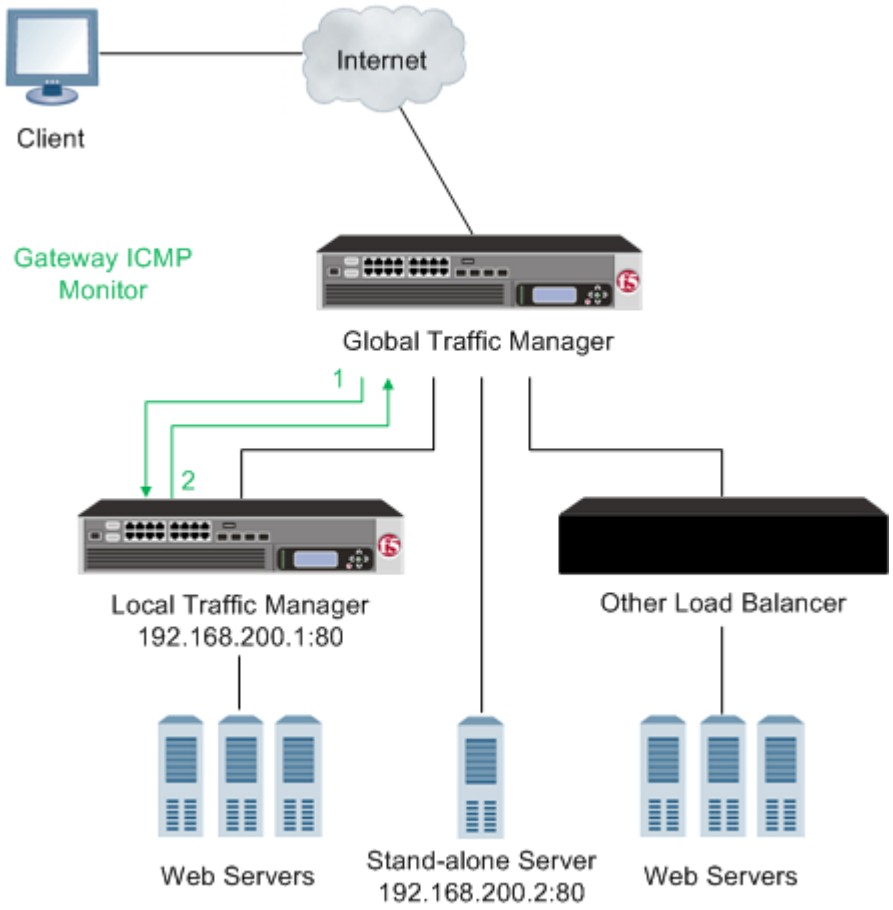

Figure 1: BIG-IP-DNS using a Gateway ICMP monitor

1. BIG-IP-DNS sends an ICMP echo request to a virtual server.
2. An ICMP echo response is received.

An *address check monitor* provides a simple verification of an address on a network. This type of monitor sends a request to a virtual server. When a response is received, the test is successful.

When an address check monitor is associated with a node, it determines the availability of all services associated with that node's IP address. If the monitor is unsuccessful in determining that a node is available, the monitor marks the node and all pool members at that IP address as **Offline**.

The following illustration depicts a Local Traffic Manager™ using a **TCP Echo** monitor to verify an IP address for a virtual server.

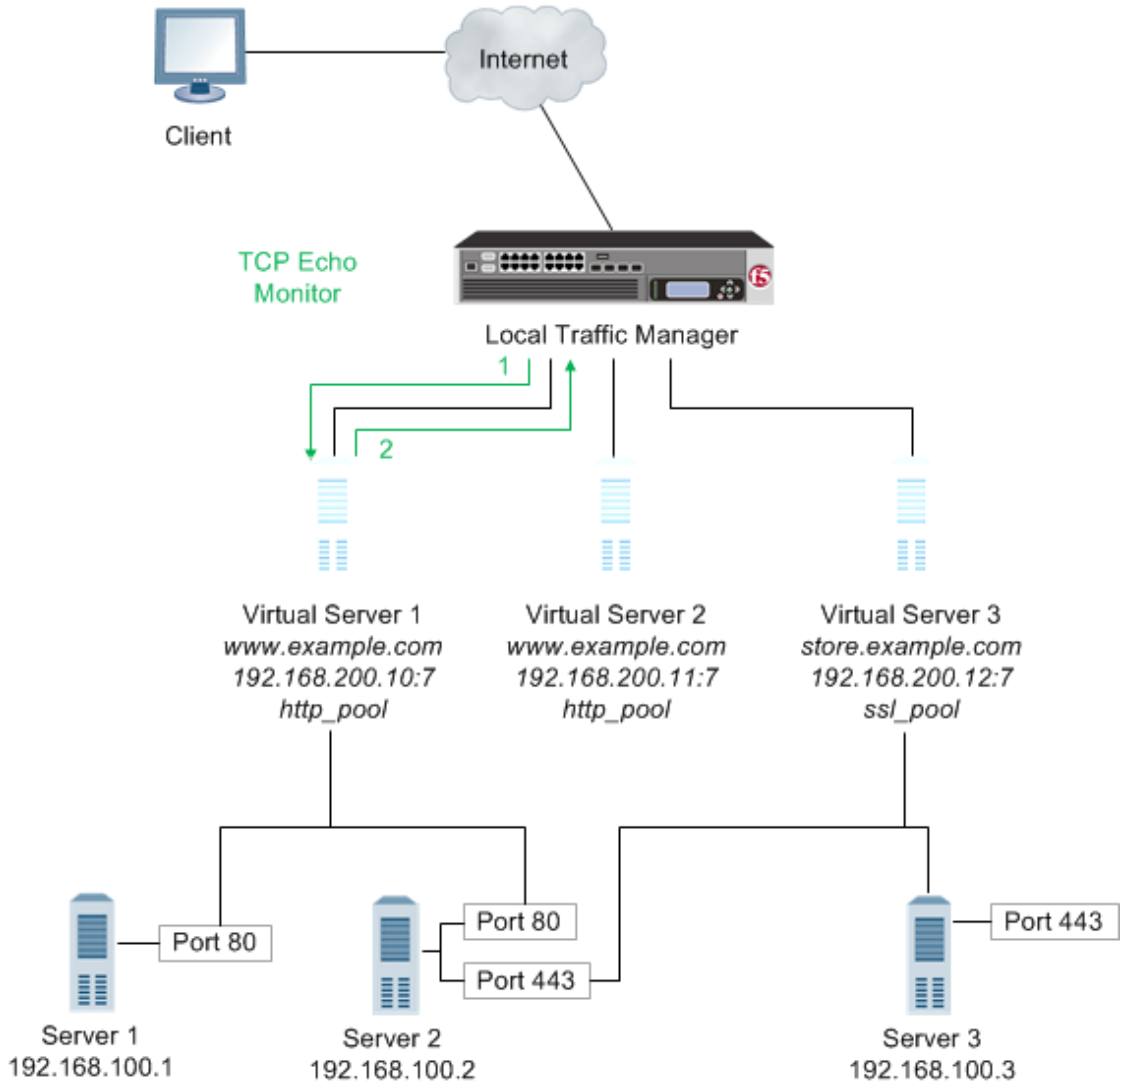

Figure 2: Local Traffic Manager using a TCP Echo monitor

1. Local Traffic Manager sends a TCP echo request to a virtual server.
2. A TCP echo response is received.

About application check monitors

An *application check monitor* interacts with servers by sending multiple commands and processing multiple responses.

An FTP monitor, for example, connects to a server, logs in by using a user ID and password, navigates to a specific directory, and then downloads a specific file to the `/var/tmp` directory. If the file is retrieved, the check is successful.

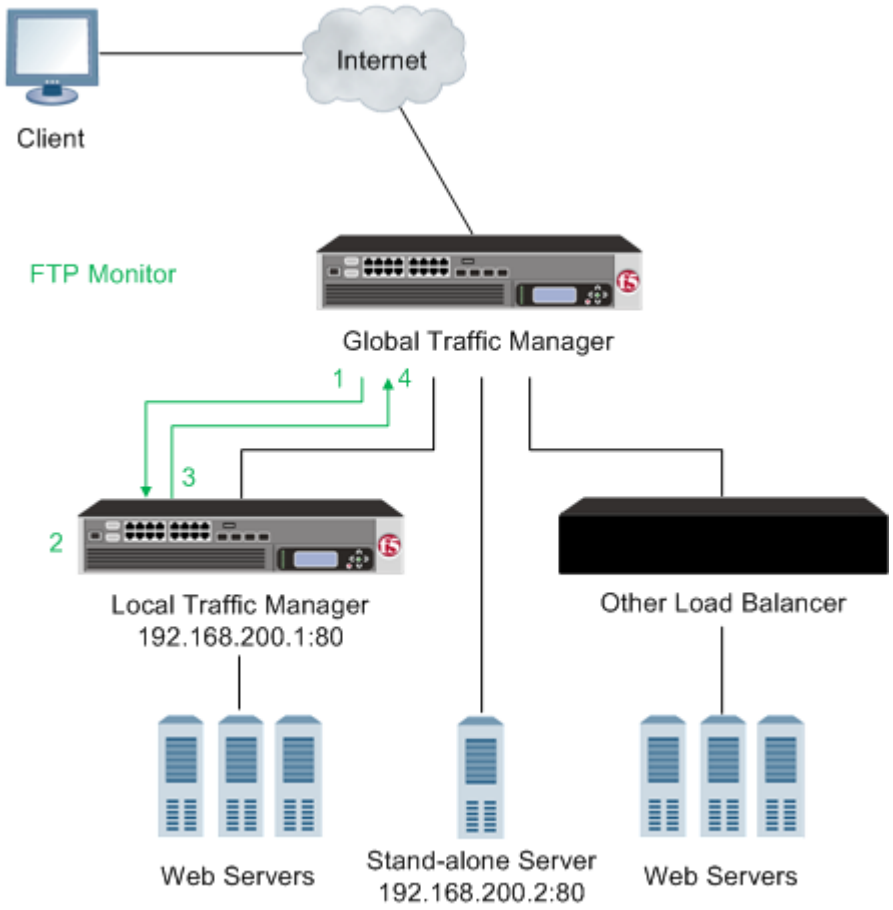

Figure 3: An application check monitor

1. BIG-IP-DNS opens a TCP connection to an IP address and port, and logs in to the server.
2. A specified directory is located and a specific file is requested.
3. The server sends the file to BIG-IP-DNS.
4. BIG-IP-DNS receives the file and closes the TCP connection.

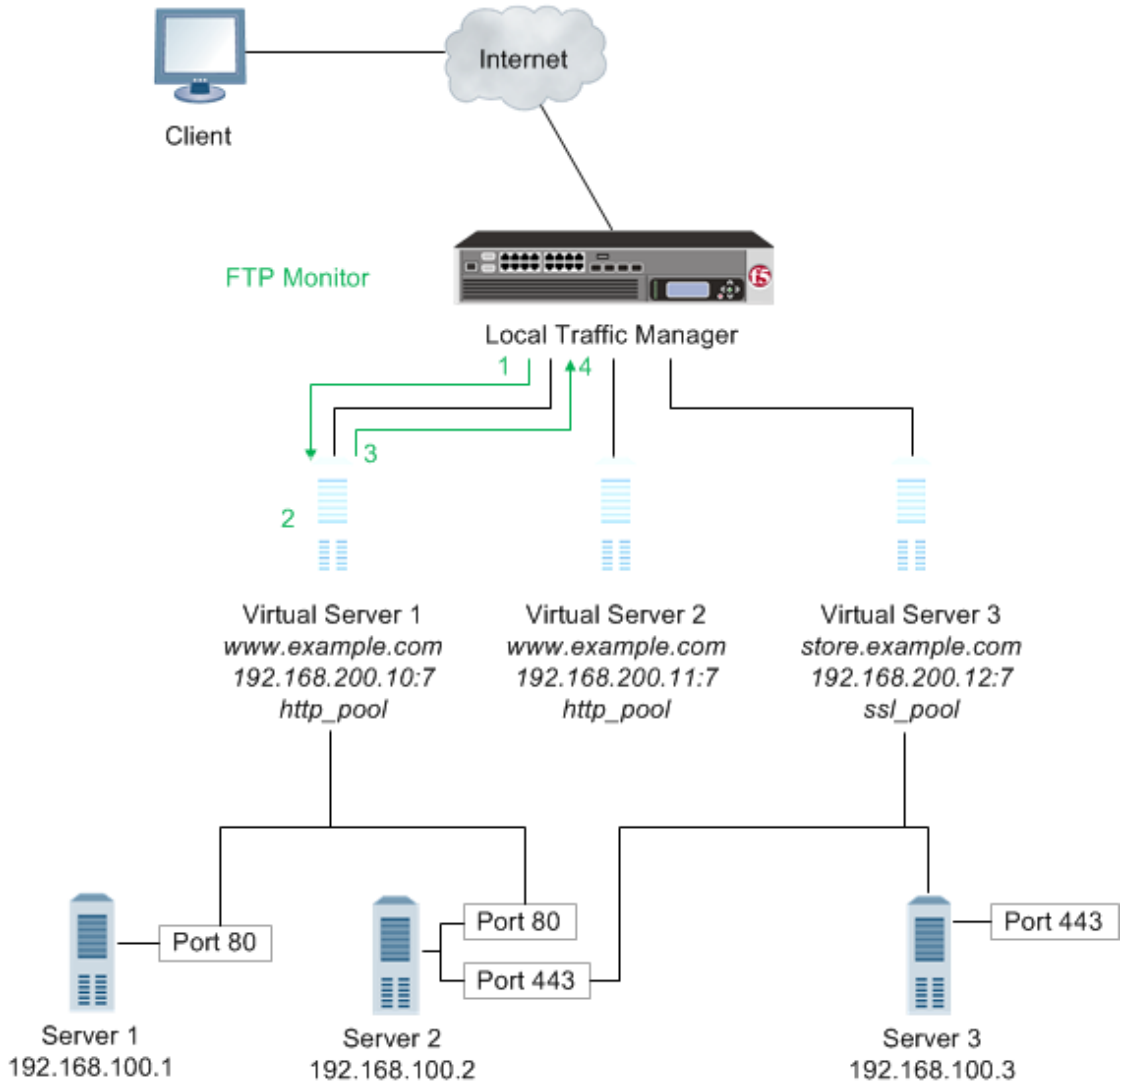

Figure 4: An application check monitor

1. Local Traffic Manager opens a TCP connection to an IP address and port, and logs in to the server.
2. A specified directory is located and a specific file is requested.
3. The server sends the file to Local Traffic Manager.
4. Local Traffic Manager receives the file and closes the TCP connection.

About content check monitors

A *content check monitor* determines whether a service is available and whether the server is serving the appropriate content. This type of monitor opens a connection to an IP address and port, and then issues a command to the server. The response is compared to the monitor's receive rule. When a portion of the server's response matches the receive rule, the test is successful.

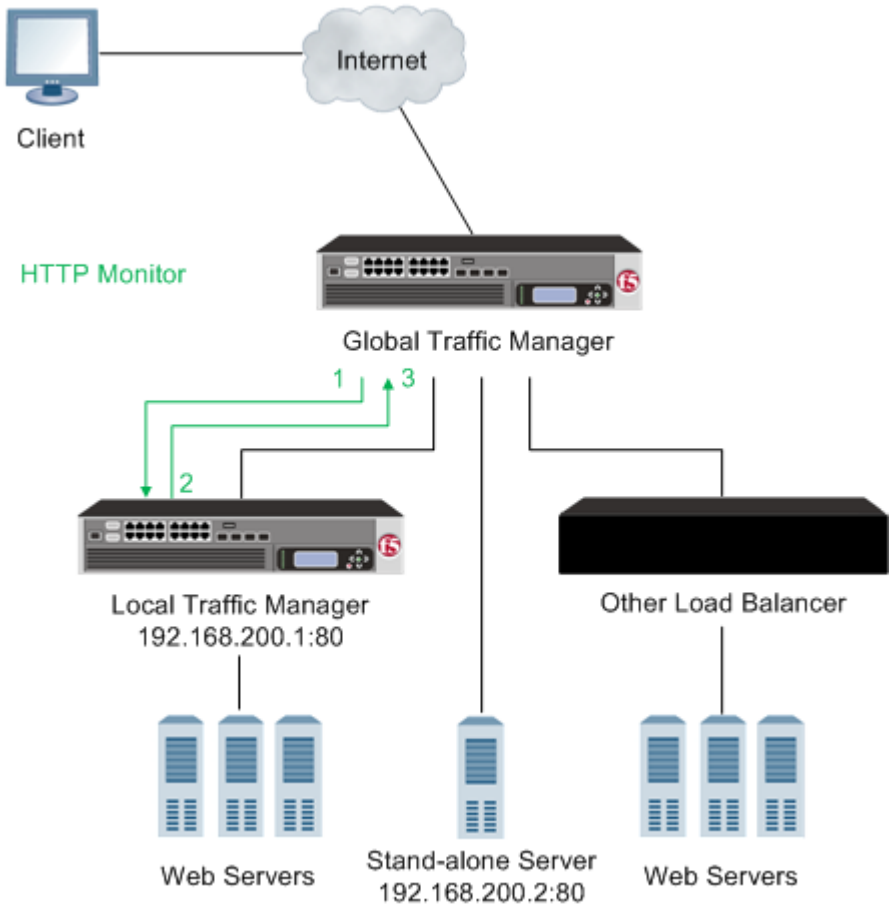

Figure 5: A content check monitor

1. DNS opens a TCP connection to an IP address and port, and issues a command to the server.
2. The server sends a response.
3. BIG-IP-DNS compares the response to the monitor's receive rule and closes the connection.

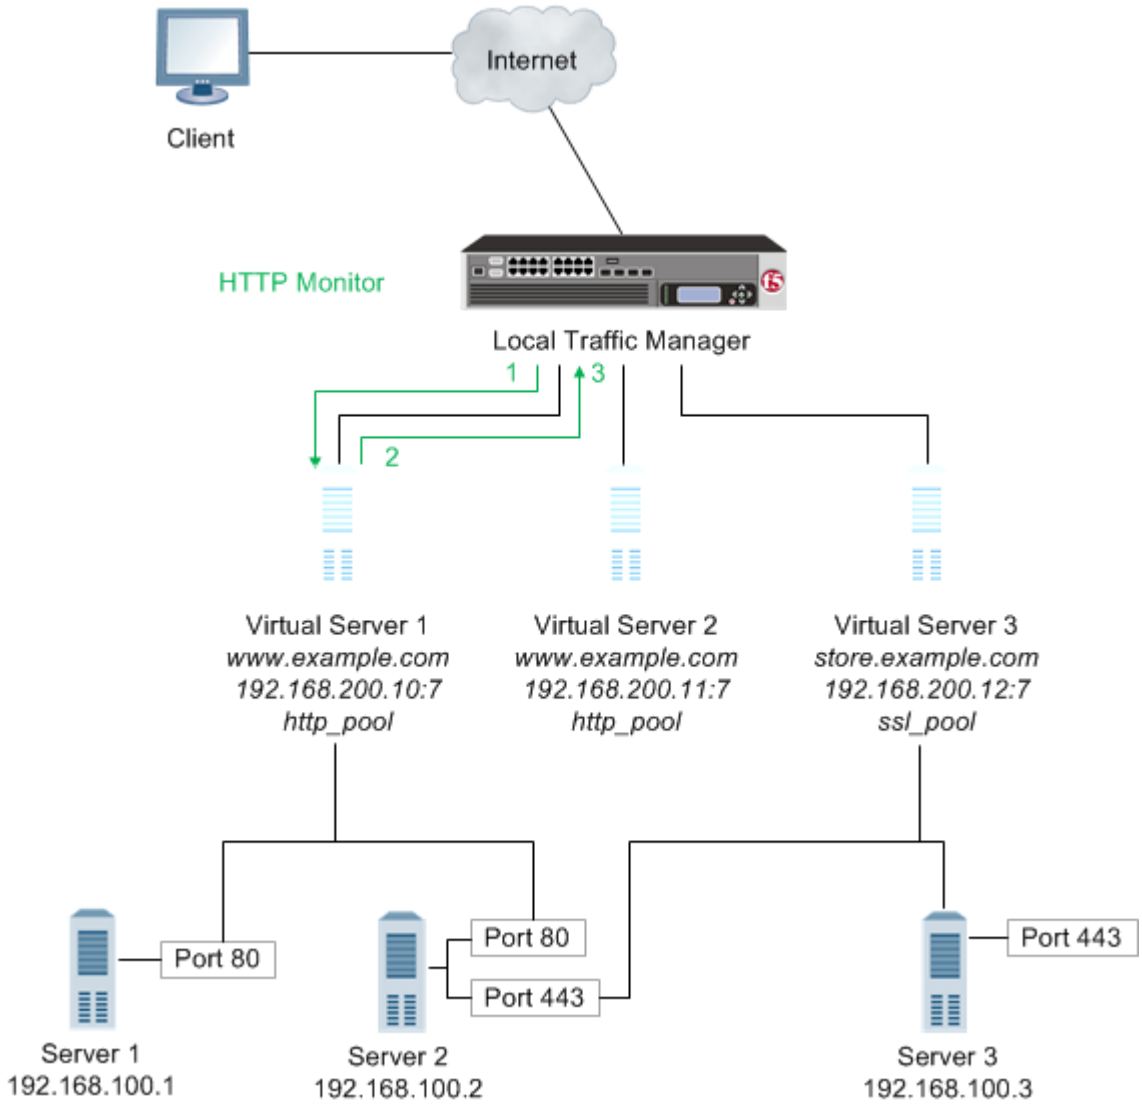

Figure 6: A content check monitor

1. Local Traffic Manager™ opens a TCP connection to an IP address and port, and issues a command to the server.
2. The server sends a response.
3. Local Traffic Manager compares the response to the monitor's receive rule and closes the connection

About path check monitors

A *path check monitor* determines whether traffic can flow through a device to an endpoint. A path check monitor is successful when network paths through firewalls or routers are available.

The following illustration depicts Global Traffic Manager™ (GTM™), which is now known as BIG-IP® DNS, using a **Gateway ICMP** monitor to verify a path to a virtual server.

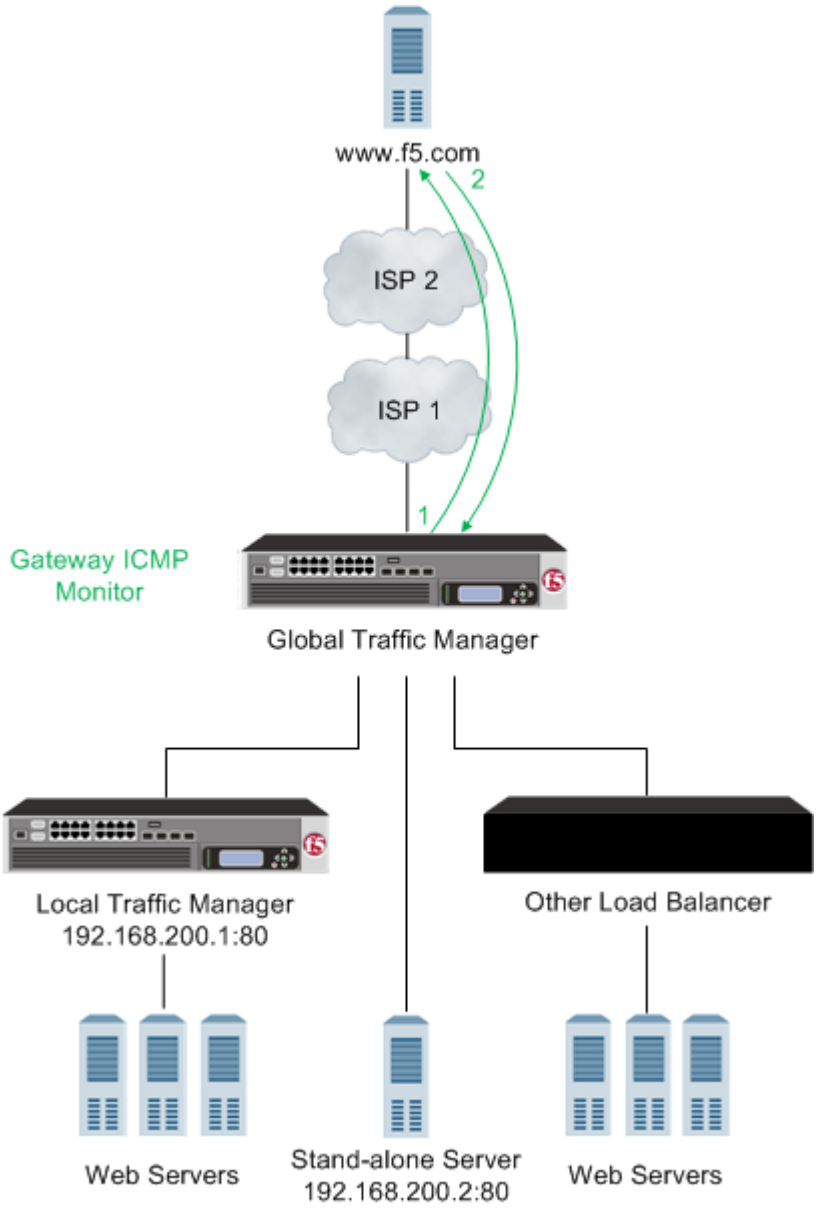

Figure 7: Global Traffic Manager (now known as BIG-IP DNS) using a Gateway ICMP monitor

1. With the **Gateway ICMP** monitor **Transparent** option set to **Yes**, BIG-IP DNS sends an ICMP echo request to a virtual server.
2. An ICMP echo response is received.

The following illustration depicts Local Traffic Manager™ (LTM®) using a **TCP Echo** monitor to verify a path to a virtual server.

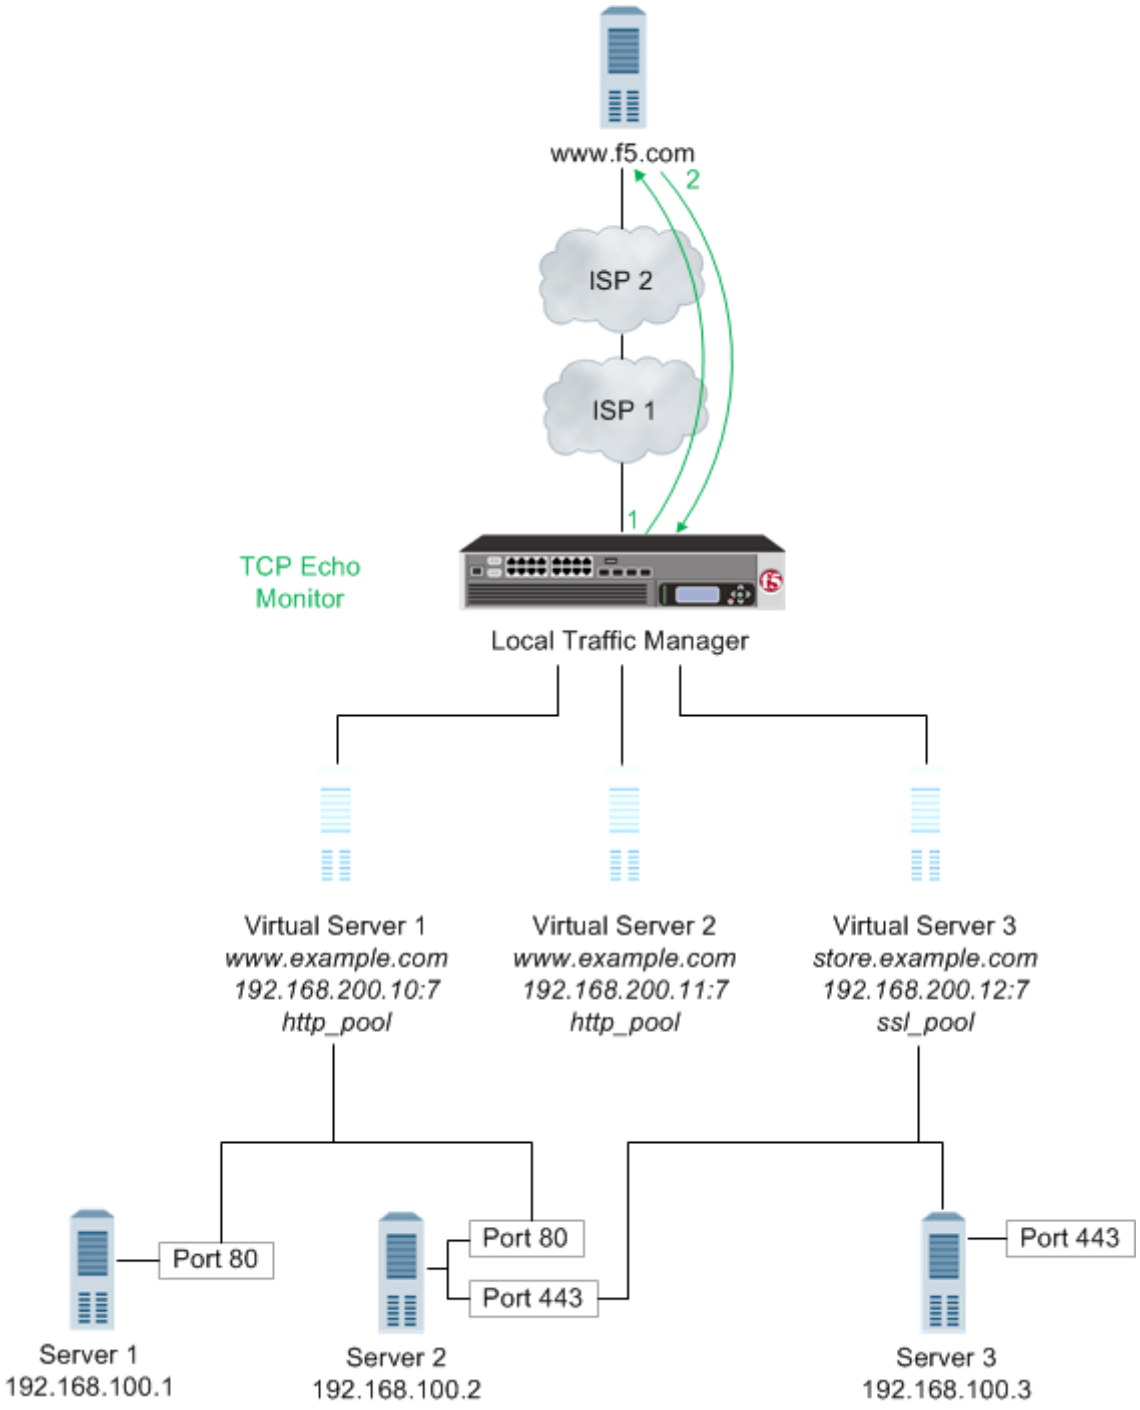

Figure 8: Local Traffic Manager using a TCP Echo monitor

1. With the **TCP Echo** monitor **Transparent** option set to **Yes**, Local Traffic Manager sends a TCP Echo request to a virtual server.
2. A TCP Echo response is received.

About performance check monitors

A *performance check monitor* interacts with a link or server to acquire information about the resource load and the condition of the virtual servers on the server.

On Link Controller™, you assign the BIG-IP Link monitor to link entries. This monitor gathers data from the gateway pool about the flow of the outbound traffic passing through the links.

On DNS, you assign the BIG-IP Link monitor to link entries. This monitor gathers data from the gateway pool about the flow of the outbound traffic passing through the links.

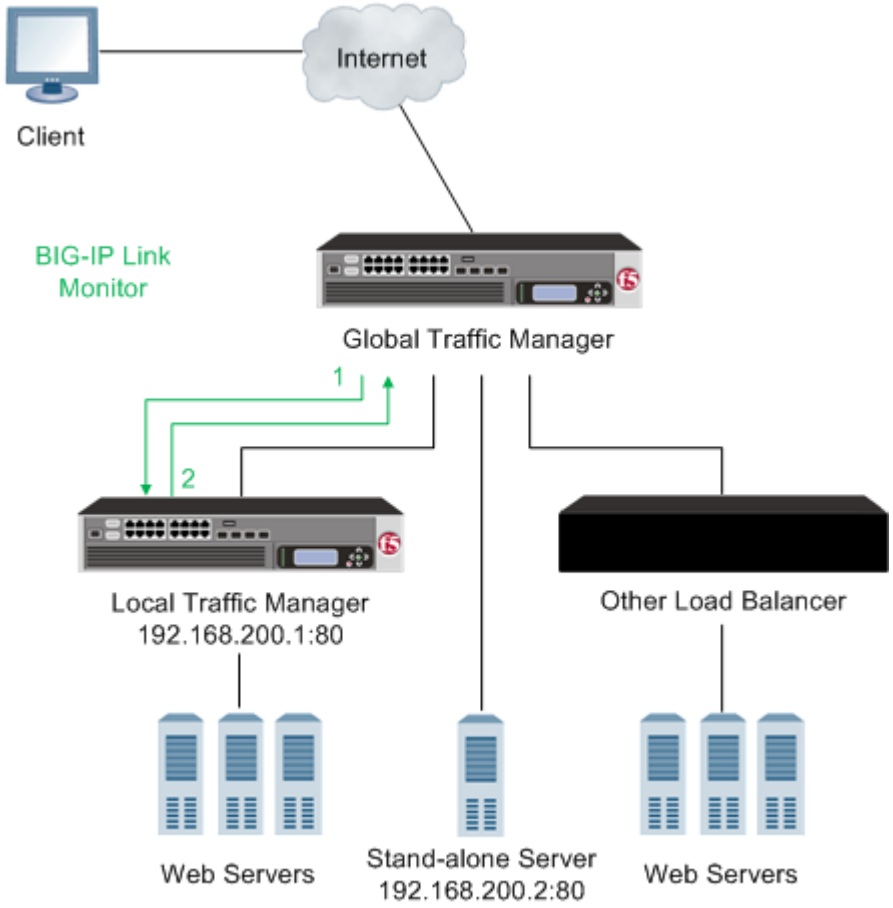

Figure 9: A performance check monitor

1. The BIG-IP Link monitor gathers data from the gateway pool.
2. BIG-IP-DNS evaluates the health of the link and makes a determination about load balancing traffic to the link.

A *performance check monitor* interacts with servers to determine the server load, and to acquire information about the condition of virtual servers.

An SNMP DCA monitor, for example, checks the current CPU, memory, and disk usage of a pool, pool member, or node that is running an SNMP data collection agent, and then dynamically load balances traffic accordingly.

Note: *If you configure a performance monitor, such as the SNMP DCA or WMI monitor type, you should also configure a health monitor. Configuring a health monitor ensures that Local Traffic Manager reports accurate node availability status.*

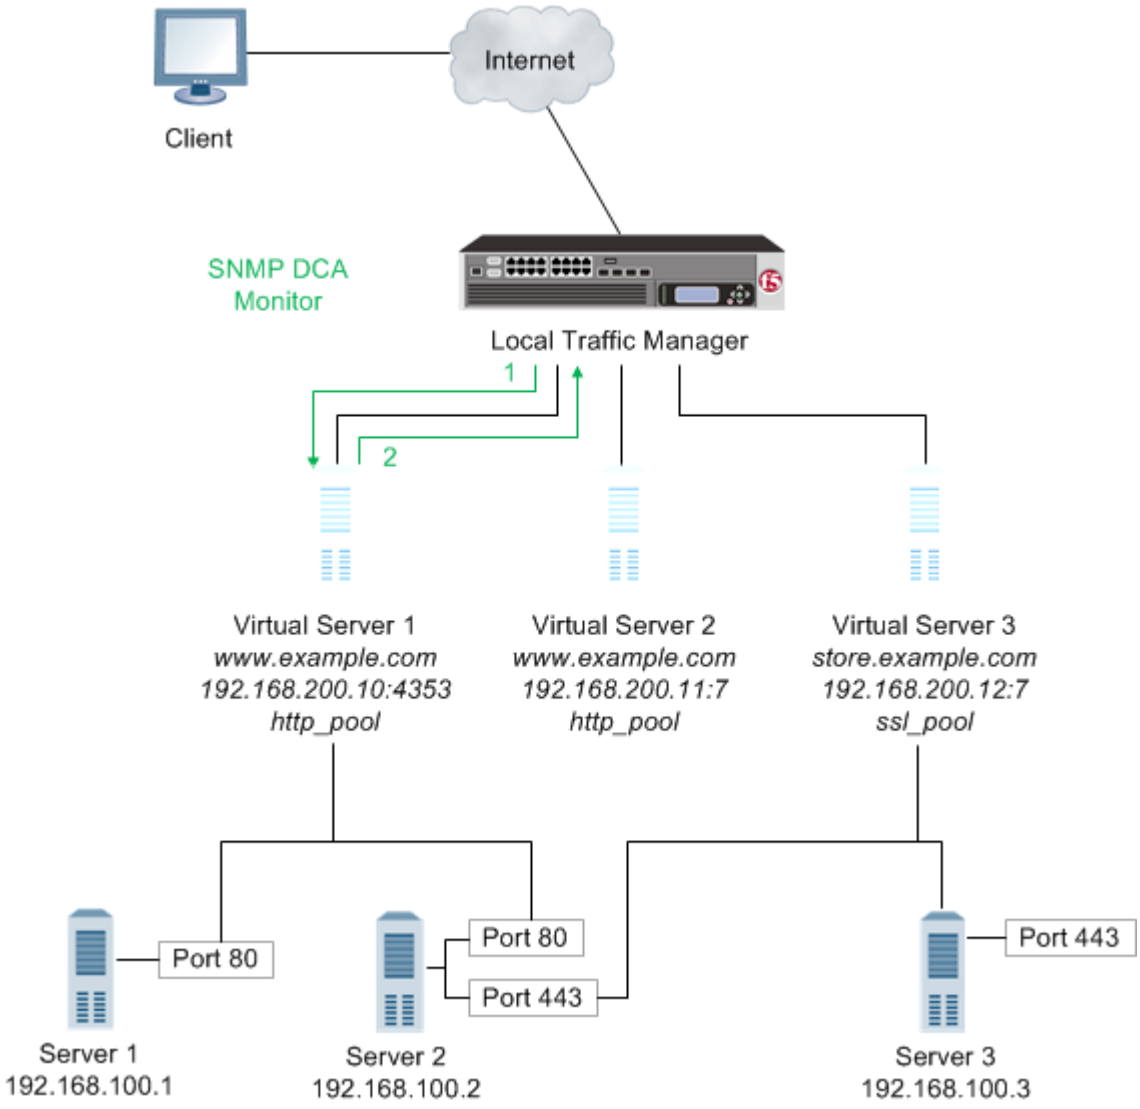

Figure 10: A performance check monitor

1. Local Traffic Manager™ connects with a server to acquire data.
2. The server sends the data to Local Traffic Manager for evaluation and determination of load balancing.

About service check monitors

A *service check monitor* determines whether a service is available. This type of monitor opens a connection to an IP address and port, and then closes the connection. When the TCP connection is established, the test is successful.

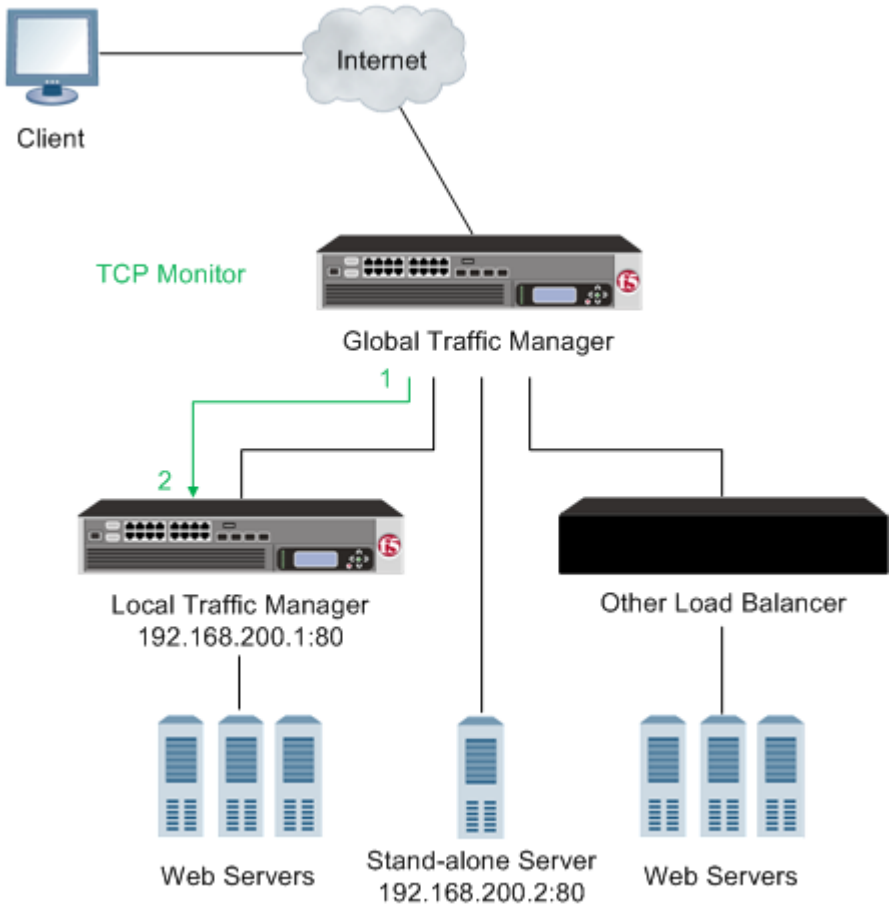

Figure 11: A service check monitor

1. DNS opens a TCP connection to an IP address and port.
2. The TCP connection is closed.

When a service check monitor is associated with pool members, it determines the availability of a service. If the monitor is unsuccessful in determining that a pool member is available, the monitor marks the pool member as **Offline** and no requests are sent to that pool member.

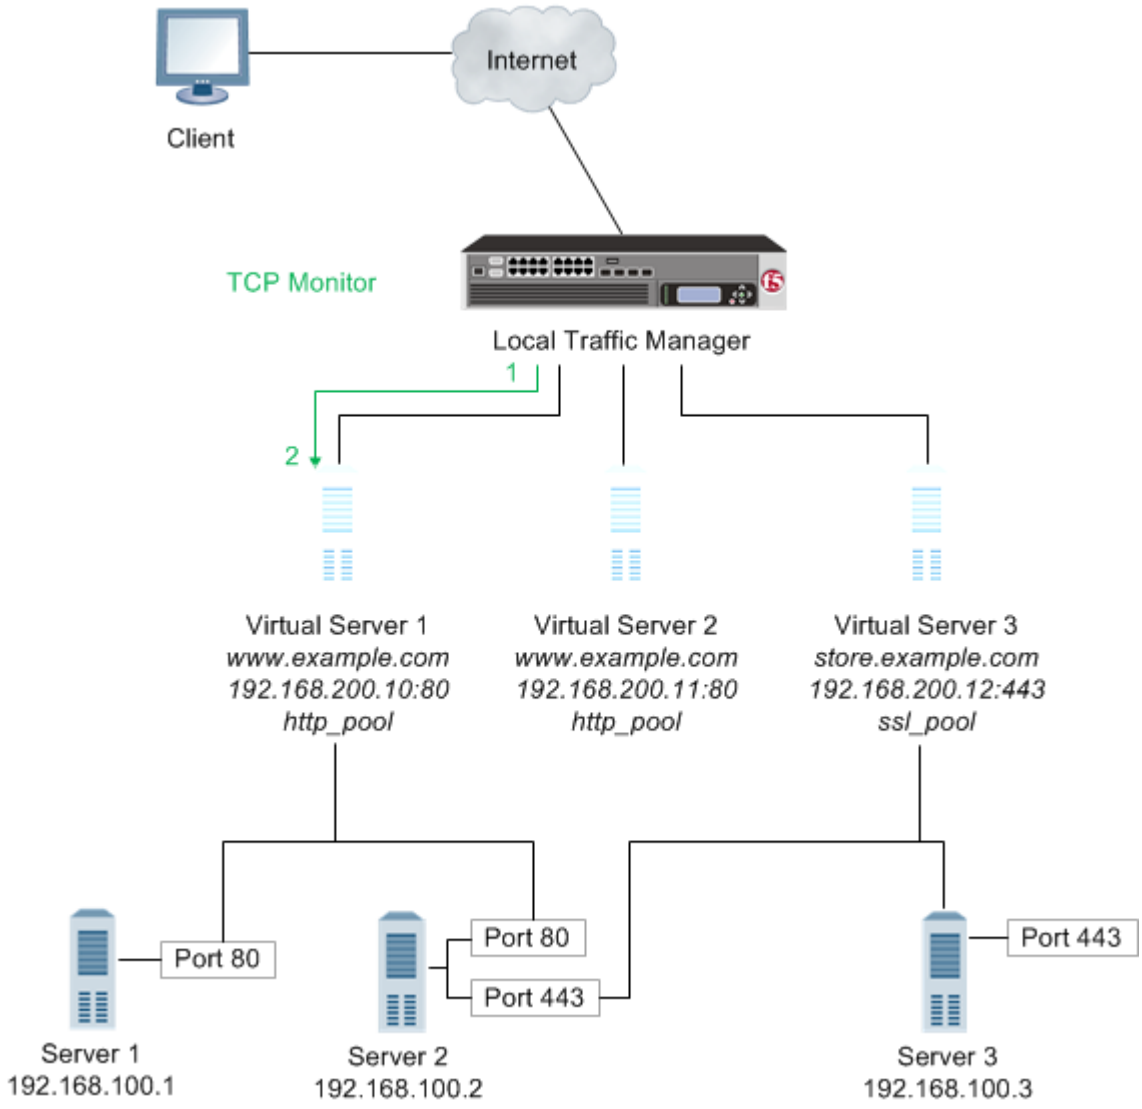

Figure 12: A service check monitor

1. Local Traffic Manager™ opens a TCP connection to an IP address and port.
2. The TCP connection is closed.

About resources and monitor queries

Network resources often perform different functions at the same time. Therefore, it is likely that multiple monitors are checking the availability of a single resource in different ways.

Example:

A BIG-IP® system may monitor a single resource to verify that the connection to the resource is available, that a specific HTML page on the resource can be reached, and that a database query returns an expected result.

About the Virtual Location monitor

The **Virtual Location** monitor optimizes the way that the BIG-IP® system manages connections to pool members by assigning priority groups to local and remote pool members.

The monitor determines whether a pool member is local (residing in the same data center as the BIG-IP system) or remote (residing in a different data center). If a pool member is local, the monitor sets the priority group of the pool member to a higher priority. If a pool member is remote, the monitor sets the priority group of the pool member to a lower priority.

Important: *You must configure Priority Group Activation to specify the minimum number of available members, before the BIG-IP system begins directing traffic to members in a lower priority group.*

About adaptive response time monitoring

Adaptive response time monitoring measures the amount of time between when the BIG-IP® system sends a probe to a resource and when the system receives a response from the resource. It adds an extra dimension to existing monitoring capabilities. A monitor with adaptive response time enabled marks a service as up or down based on the deviation of latency of the monitor probe from the mean latency of a monitor probe for that service. In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down.

About the types of adaptive response time monitoring

There are two types of adaptive response time monitoring:

Absolute
The number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed.

Relative
The percentage of deviation that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed; that is, the running mean latency calculated by the system.

You can enable the adaptive response time monitoring feature on these specific monitors:

- DNS
- Gateway ICMP
- HTTP
- HTTPS
- ICMP
- TCP
- UDP

About calculating the mean latency of a probe

A monitor marks a service down if a response to a probe does not meet the latency requirements of either the absolute limit or the relative limit, that is the running average. By default, the system stores the last five minutes of probe history for each monitor instance in a buffer. The system uses this history to calculate the varying mean latency of the probes for that monitor instance.

How does adaptive response time monitoring work?

This example shows a BIG-IP® Local Traffic Manager™ system configured to handle HTTP traffic using a pool and an HTTP monitor with adaptive response time monitoring enabled (through the **Adaptive** setting).

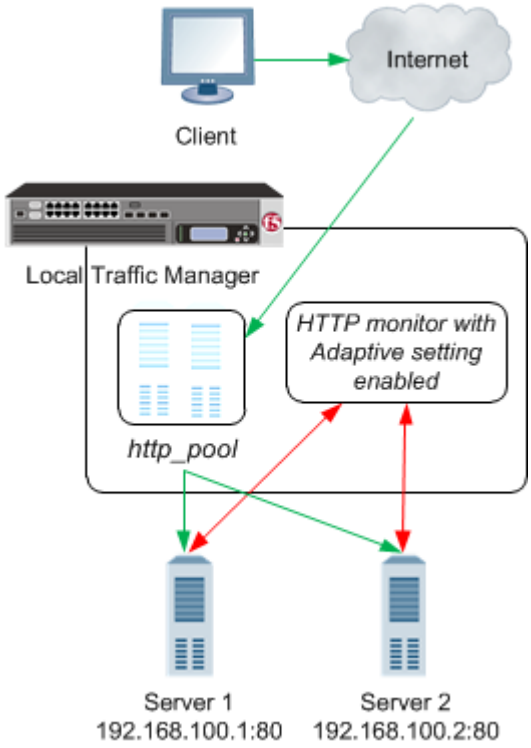

Figure 13: HTTP adaptive response time monitoring

1. A client makes an HTTP request. The HTTP request is represented by the green arrows.
2. The request is routed to an HTTP pool on BIG-IP Local Traffic Manager (LTM).
3. The LTM routes the request to one of two servers in the pool.
4. The HTTP monitor assigned to the pool determines whether the servers are up or down based on the probe latency of each server. The probe is represented by the red arrows.

Using adaptive response time monitoring to optimize a web application

One example of how you can use adaptive response time monitoring is to optimize a moderately configurable web application that is served by several web servers with limited memory capacity. For example, when the web application is overwhelmed with traffic, perhaps at month end, the application may consume excessive amounts of memory and start swapping to disk, substantially degrading performance. Because performance degrades drastically when this condition poccurs, and you do not want the BIG-IP® Local Traffic Manager™ to mark a server down unnecessarily, you can configure the servers in a pool with an HTTP monitor by enabling the **Adaptive** setting.

Using adaptive response time monitoring to mitigate probe attacks

You can use adaptive response time monitoring to mitigate probe attacks. For example, consider the scenario where a popular web application for a financial company receives a a huge number of brute-force logon attempts that cause the web servers to become unresponsive. As the administrator, you can place the web servers in a pool configured for priority-based load balancing and assign an HTTP monitor

with the **Adaptive** setting `Enabled`. When probe latency spikes, the monitor marks the primary servers in the pool down. When all the primary servers are marked down, the system sends requests to a secondary set of servers in the pool that presents a page that does not accept logon attempts.

Overview of monitor implementation

You implement monitors by using either the BIG-IP® Configuration utility or a command line utility. The task of implementing a monitor varies depending on whether you are using a preconfigured monitor or creating a custom monitor. A *preconfigured monitor* is an existing monitor that BIG-IP system provides for you, with its settings already configured. A *custom monitor* is a monitor that you create based on one of the allowed monitor types.

If you want to implement a preconfigured monitor, you need only associate the monitor with a pool, pool member, or node, and then configure the virtual server to reference the relevant pool. If you want to implement a custom monitor, you must first create the custom monitor. Then you can associate the custom monitor with a pool, pool member, or node, and configure the virtual server to reference the pool.

Preconfigured monitors

For a subset of monitor types, the BIG-IP® system includes a set of preconfigured monitors. You cannot modify preconfigured monitor settings, as they are intended to be used as is. The purpose of a preconfigured monitor is to eliminate the need for you to explicitly create a monitor. You use a preconfigured monitor when the values of the settings meet your needs as is.

Preconfigured monitors include the following entries.

- `bigip`
- `bigip_link`
- `gateway_icmp`
- `gtp`
- `http`
- `http_head_f5`
- `https`
- `https_443`
- `https_head_f5`
- `icmp`
- `inband`
- `real_server`
- `snmp_dca`
- `tcp`
- `tcp_half_open`

An example of a preconfigured monitor is the `http` monitor. The example shows the `http` monitor, with values configured for its **Interval**, **Timeout**, and **Alias Address** settings. Note that the Interval value is `530`, the Timeout value is `16120`, the Transparent value is `No`, and the Alias Address value is `* All Addresses`.

If the Interval, Timeout, Transparent, and Alias Address values meet your needs, you simply assign the `http` preconfigured monitor directly to a server, virtual server, pool, pool member, or link. In this case, you do not need to use the Monitors screens, unless you simply want to view the values of the preconfigured monitor settings.

```
Name http
```

```
Type HTTP
Interval 5
Timeout 16
Transparent No
Alias Address * All Addresses
```

If the Interval, Timeout, Transparent, and Alias Address values meet your needs, you simply assign the `http` preconfigured monitor directly to a server, virtual server, pool, pool member, or link. In this case, you do not need to use the Monitors screens, unless you simply want to view the values of the preconfigured monitor settings.

```
Name http
Type HTTP
Interval 30
Timeout 120
Transparent No
Alias Address * All Addresses
```

Important: *All preconfigured monitors reside in partition* `Common`*.*

Custom monitors

You create a custom monitor when the values defined in a preconfigured monitor do not meet your needs, or no preconfigured monitor exists for the type of monitor you are creating.

When you create a custom monitor, you use the BIG-IP® Configuration utility or a command line utility to: give the monitor a unique name, specify a monitor type, and, if a monitor of that type already exists, import settings and their values from the existing monitor. You can then change the values of any imported settings.

You must base each custom monitor on a monitor type. When you create a monitor, the BIG-IP Configuration utility displays a list of monitor types. To specify a monitor type, simply choose the one that corresponds to the service you want to check. For example, if you want to want to create a monitor that checks the health of the HTTP service on a pool, you choose HTTP as the monitor type.

If you want to check more than one service on a pool or pool member (for example HTTP and HTTPS), you can associate more than one monitor on that pool or pool member.

Checking services is not the only reason for implementing a monitor. If you want to verify only that the destination IP address is alive, or that the path to it through a transparent node is alive, use one of the simple monitors, `icmp` or `tcp_echo`. Or, if you want to verify TCP only, use the monitor `tcp`.

Importing settings from a preconfigured monitor

If a preconfigured monitor exists that corresponds to the type of custom monitor you are creating, you can import the settings and values of that preconfigured monitor into the custom monitor. You are then free to change those setting values to suit your needs. For example, if you create a custom monitor called `my_icmp`, the monitor can inherit the settings and values of the preconfigured monitor `icmp`. This ability to import existing setting values is useful when you want to retain some setting values for your new monitor but modify others.

The example shows a custom ICMP-type monitor called `my_icmp`, which is based on the preconfigured monitor `icmp`. Note that the Interval value is changed to `10`, and the Timeout value is `20`. The other settings retain the values defined in the preconfigured monitor.

```
Name my_icmp
Type ICMP
Interval 10
Timeout 20
```

```
Transparent No
Alias Address * All Addresses
```

Checking services is not the only reason for implementing a monitor. If you want to verify only that the destination IP address is alive, or that the path to it through a transparent node is alive, use a simple monitor, such as `gateway_icmp`. Or, if you want to verify TCP only, use the monitor `tcp`.

Importing settings from a preconfigured monitor

If a preconfigured monitor exists that corresponds to the type of custom monitor you are creating, you can import the settings and values of that preconfigured monitor into the custom monitor. You are then free to change those setting values to suit your needs. For example, if you create a custom monitor called `my_gateway_icmp`, the monitor can inherit the settings and values of the preconfigured monitor `gateway_icmp`. This ability to import existing setting values is useful when you want to retain some setting values for your new monitor but modify others.

The example shows a custom ICMP-type monitor called `my_gateway_icmp`, which is based on the preconfigured monitor `gateway_icmp`. Note that the Interval value is changed to `20`, and the Timeout value is `100`. The other settings retain the values defined in the preconfigured monitor.

```
Name my_gateway_icmp
Type Gateway ICMP
Interval 20
Timeout 100
Transparent No
Alias Address * All Addresses
```

Importing settings from a custom monitor

You can import settings from another custom monitor instead of from a preconfigured monitor. This is useful when you would rather use the setting values defined in another custom monitor, or when no preconfigured monitor exists for the type of monitor you are creating. For example, if you create a custom monitor called **my_oracle_server2**, you can import settings from another custom Oracle-type monitor that you created, such as `my_oracle_server1`. Selecting a monitor is straightforward. Like `gateway_icmp`, each of the monitors has a Type setting based on the type of service it checks, for example, `http`, `https`, `ftp`, `pop3`, and a Parent Monitor that is used for importing the custom monitor settings. (Exceptions are port-specific monitors, like the `external` monitor, which calls a user-supplied program.)

Dynamic ratio load balancing

You can configure Dynamic Ratio load balancing for pools that consist of RealNetworks® RealServer™ servers, Microsoft® Windows® servers equipped with Windows Management Instrumentation (WMI), or any server equipped with an SNMP agent such as the UC Davis SNMP agent or Windows® 2000 Server SNMP agent.

You can configure Dynamic Ratio load balancing for pools that consist of RealNetworks® RealServer™ servers or Microsoft® Windows® servers equipped with Windows Management Instrumentation (WMI).

To implement Dynamic Ratio load balancing for these types of servers, BIG-IP® system provides a special monitor plug-in file and a performance monitor for each type of server. The exception is a server equipped with an SNMP agent. In this case, the BIG-IP system provides the monitor only; no special plug-in file is required for a server running an SNMP agent.

To implement Dynamic Ratio load balancing for these types of servers, BIG-IP® system provides a special monitor plug-in file and a performance monitor for each type of server.

You must install the monitor plug-in on each server to be monitored, and you must create a performance monitor that resides on the BIG-IP system. Once you have created a monitor, the monitor communicates directly with the server plug-in.

Monitor plug-ins and corresponding monitor templates

For each server type, this table shows the required monitor plug-in and the corresponding performance monitor types.

Server Type	Monitor plug-in	Monitor Type
RealServer Windows server	`F5RealMon.dll`	Real Server
RealServer UNIX server	`f5realmon.so`	Real Server
Windows server with WMI	`f5isapi.dll` or `F5Isapi64.dll` or `F5.IsHandler.dll`	WMI
Windows 2000 Server server	SNMP agent	SNMP DCA and SNMP DCA Base
UNIX server	UC Davis SNMP agent	SNMP DCA and SNMP DCA Base

Monitor association with pools and nodes

You must associate a monitor with the server or servers to be monitored. The server or servers can be either a pool, a pool member, or a node, depending on the monitor type. You can associate a monitor with a server in any of these ways:

Monitor-to-pool association
This type of association associates a monitor with an entire load balancing pool. In this case, the monitor checks all members of the pool. For example, you can create an instance of the monitor `http` for every member of the pool `my_pool`, thus ensuring that all members of that pool are checked.

Monitor-to-pool member association
This type of association associates a monitor with an individual pool member, that is, an IP address and service. In this case, the monitor checks only that pool member and not any other members of the pool. For example, you can create an instance of the monitor `http` for pool member `10.10.10.10:80` of `my_pool`.

Important: *A monitor associated with an individual pool member supersedes a monitor associated with that pool member's parent pool.*

Monitor-to-node association
This type of association associates a monitor with a specific node. In this case, the monitor checks only the node itself, and not any services running on that node. For example, you can create an instance of the monitor `icmp` for node `10.10.10.10`. In this case, the monitor checks the specific node only, and not any services running on that node. You can designate a monitor as the default monitor that you want the BIG-IP system to associate with one or more nodes. In this case, any node to which you have not specifically assigned a monitor inherits the default monitor.

Some monitor types are designed for association with nodes only, and not pools or pool members. Other monitor types are intended for association with pools and pool members only, and not nodes. Finally, in some instances, some monitor types associated with a node are not mutually exclusive of pools or pool members, and must function in combination in some scenarios.

Node-only monitors specify a destination address in the format of an IP address with no service port (for example, `10.10.10.2`). Conversely, monitors that you can associate with nodes, pools, and pool members specify a destination address in the format of an IP address and service port (for example, `10.10.10.2:80`). Therefore, when you use the BIG-IP Configuration utility to associate a monitor with a pool, pool member, or node, the utility displays only those pre-configured monitors that are designed for association with that server.

For example, you cannot associate the monitor `icmp` with a pool or its members, since the `icmp` monitor is designed to check the status of a node itself and not any service running on that node.

Monitor instances

When you associate a monitor with a server, the BIG-IP® system automatically creates an *instance* of that monitor for that server. A monitor association thus creates an instance of a monitor for each server that you specify. This means that you can have multiple instances of the same monitor running on your servers.

Because instances of monitors are not partitioned objects, a user can enable or disable an instance of a monitor without having permission to manage the associated pool or pool member.

For example, a user with the Manager role, who can access partition `AppA` only, can enable or disable monitor instances for a pool that resides in partition `Common`. However, that user cannot perform operations on the pool or pool members that are associated with the monitor. Although this is correct functionality, the user might not expect this behavior. You can prevent this unexpected behavior by ensuring that all pools and pool members associated with monitor instances reside in the same partition.

Monitors Tasks

Creating an SNMP monitor

Create an SNMP monitor that DNS Link Controller™LTM® can use to monitor a third-party server running SNMP.

1. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
2. On the Main tab, click **DNS** > **GSLB** > **Monitors**.
   The Monitor List screen opens.
3. Click **Create**.
   The New Monitor screen opens.
4. Type a name for the monitor.

   Important: *Monitor names are limited to 63 characters.*

5. From the **Type** list, select one of these options:

Option	Description
SNMP DCA	Use this monitor to specify new values for CPU, memory, and disk metrics.
SNMP DCA Base	Use this monitor to specify values for metrics other than CPU, memory, and disk usage.

6. From the **Type** list, select **SNMP**.
7. Click **Finished**.

Creating a custom monitor

Before creating a custom monitor, you must decide on a monitor type.

You can create a custom monitor when the values defined in a pre-configured monitor do not meet your needs, or no pre-configured monitor exists for the type of monitor you are creating.

Important: *When defining values for custom monitors, make sure you avoid using any values that are on the list of reserved keywords. For more information, see solution number 3653 (for version 9.0 systems and later) on the AskF5™ technical support web site.*

1. On the Main tab, click **DNS** > **GSLB** > **Monitors**.
   The Monitor List screen opens.
2. On the Main tab, click **Link Controller** > **Monitors**.
   The Monitor List screen opens.
3. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
4. Click **Create**.
   The New Monitor screen opens.
5. In the **Name** field, type a name for the monitor.
6. From the **Type** list, select the type of monitor.

The screen refreshes, and displays the configuration options for the monitor type.

7. From the **Import Monitor** list, select an existing monitor.
   The new monitor inherits initial configuration values from the existing monitor.
8. From the **Parent Monitor** list, select an existing monitor.
   The new monitor inherits initial configuration values from the existing monitor.
9. From the **Configuration** list, select **Advanced**.
   This selection makes it possible for you to modify additional default settings.
10. Configure all settings shown.
11. Click **Finished**.

Deleting a monitor

Prior to deleting a monitor, you must remove all existing monitor associations.

You can delete obsolete or unused monitors.

Note: *You can manage only those monitors that you have permission to manage, based on your user role and partition access assignment.*

1. On the Main tab, click **DNS** > **GSLB** > **Monitors**.
   The Monitor List screen opens.
2. On the Main tab, click **Link Controller** > **Monitors**.
   The Monitor List screen opens.
3. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
4. Select the **Select** check box for the monitor that you want to delete.
5. Click **Delete**.
   A confirmation message appears.
6. Click **Delete**.

The monitor is deleted.

Displaying a monitor

You can display a monitor and view the settings and values.

Note: *You can manage only those monitors that you have permission to manage, based on your user role and partition access assignment.*

1. On the Main tab, click **DNS** > **GSLB** > **Monitors**.
   The Monitor List screen opens.
2. On the Main tab, click **Link Controller** > **Monitors**.
   The Monitor List screen opens.
3. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
4. Click a monitor name in the list.
   The monitor's properties screen opens, showing the monitor's settings and values.

You can view the settings and values for the monitor.

Creating an HTTP monitor

Before creating a monitor, you must decide on a monitor type.

A custom HTTP monitor enables you to send a command to a server and examine that server's response, thus ensuring that it is serving appropriate content.

Note: *An HTTP monitor can monitor Outlook® Web Access (OWA) in Microsoft® Exchange Server 2007 and Microsoft® SharePoint® 2007 web sites that require NT LAN Manager (NTLM) authentication. NTLM authentication requires a send string that complies with HTTP/1.1, a user name, and a password.*

---

1. On the Main tab, click **DNS** > **GSLB** > **Monitors**.
   The Monitor List screen opens.
2. On the Main tab, click **Link Controller** > **Monitors**.
   The Monitor List screen opens.
3. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
4. In the **Name** field, type a name for the monitor.
5. From the **Type** list, select **HTTP**.
   The screen refreshes, and displays the configuration options for the **HTTP** monitor type.
6. From the **Parent Monitor** list, select **http**.

   The new monitor inherits initial configuration values from the existing monitor.
7. From the **Configuration** list, select **Advanced**.

   This selection makes it possible for you to modify additional default settings.
8. In the **Interval** field, type a number that indicates, in seconds, how frequently the system issues the monitor check. The default is `30` seconds.

   The frequency of a monitor check must be greater than the value of the global-level **Heartbeat Interval** setting. Otherwise, the monitor can acquire out-of-date data.
9. In the **Timeout** field, type a number that indicates, in seconds, how much time the target has to respond to the monitor check. The default is `120` seconds.

   If the target responds within the allotted time period, it is considered up. If the target does not respond within the time period, it is considered down.
10. In the **Probe Timeout** field, type a number that indicates the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
11. For the **Ignore Down Response** setting, do one of the following:
    - Accept the **No** default option.
    - Select the **Yes** option to specify that the monitor accepts more than one probe attempt per interval.
12. In the **Send String** field, type a text string that the monitor sends to the target resource.

    The default string is `GET /`. This string retrieves a default file from the web site.

    Type a fully qualified path name, for example, `GET /www/example/index.html`, if you want to retrieve a specific web site page.
13. In the **Interval** field type a number that indicates, in seconds, how frequently the system issues the monitor check. The default is `5` seconds.

    The frequency of a monitor check must be greater than the value of the global-level **Heartbeat Interval** setting. Otherwise, the monitor can acquire out-of-date data.
14. For the **Up Interval** setting, specify whether to use the up interval:

- If you do not want to use the up interval, Retain the default, **Disabled**.
- To use the up interval, select **Enabled**, and specify how often you want the system to verify the health of a resource that is up.

15. In the **Time Until Up** field, type a number that indicates the number of seconds to wait after a resource first responds correctly to the monitor before setting the resource to up.

The default value is `0` (zero), which disables this option.

16. In the **Timeout** field, type a number that indicates, in seconds, how much time the target has to respond to the monitor check. The default is `30` seconds.

If the target responds within the allotted time period, it is considered up. If the target does not respond within the time period, it is considered down.

17. For **Manual Resume**, specify whether the system automatically enables the monitored resource when the monitor check is successful.

This setting applies only when the monitored resource has failed to respond to a monitor check.

Option	Description
Yes	The system does nothing when the monitor check succeeds, and you must manually enable the monitored resource.
No	The system automatically re-enables the monitored resource after the next successful monitor check.

18. In the **Send String** field, type a text string that the monitor sends to the target resource.

The default string is `GET /\r\n`. This string retrieves a default file from the web site.

Important: *Send string syntax depends upon the HTTP version. Please observe the following conventions.*

Version	*Convention*
HTTP 0.9	*"GET /\n" or "GET /\r\n".*
HTTP 1.0	*"GET / HTTP/1.0\r\n\r\n" or "GET / HTTP/1.0\n\n"*
HTTP 1.1	*"GET / HTTP/1.1\r\nHost: server.com \r\n\r\n" or "GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close\r\n \r\n"*

Type a fully qualified path name, for example, `"GET /www/example/index.html\r\n"`, if you want to retrieve a specific web site page.

19. In the **Receive String** field, type a regular expression that represents the text string that the monitor looks for in the returned resource.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names.

Note: *If you do not specify both a send string and a receive string, the monitor performs a simple service check and connect only.*

---

20. In the **Receive Disable String** field, type a regular expression that represents the text string that the monitor looks for in the returned resource.

Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and

Receive Disable String match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response.

***Note:** If you choose to set the **Reverse** setting to **Yes**, the monitor marks the pool, pool member, or node **Down** when the test is successful.*

---

21. Type a name in the **User Name** field.
22. Type a password in the **Password** field.
23. For the **Reverse** setting, specify whether you want the system to work in reverse mode:
    - If you want the system to work normally, retain the **No** default option.
    - If you want the system to mark the pool, pool member, or node **Down** when the test is successful, select the **Yes** option.
24. For the **Transparent** setting, specify whether you want the monitor to operate in transparent mode:
    - If not, accept the **No** default option.
    - To use a path through the associated pool members or nodes to monitor the aliased destination, select the **Yes** option.
25. For the **Alias Address** setting, specify an alias IP address:
    - Retain the ***All Addresses** default option.
    - Type an alias IP address for the monitor to verify, on behalf of the pools or pool members with which the monitor is associated.

    If the health check for the alias address is successful, the system marks all associated objects `up`. If the health check for the alias address is not successful, then the system marks all associated objects `down`.
26. For the **Alias Service Port** setting, specify an alias port or service for the monitor to check:
    - Accept the ***All Ports** default option.
    - Select an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

The HTTP monitor is configured to monitor HTTP traffic.

Creating an HTTPS monitor

Before creating a monitor, you must decide on a monitor type.

A custom HTTPS monitor enables you to verify the Hypertext Transfer Protocol Secure (HTTPS) service by attempting to receive specific content from a web page protected by Secure Socket Layer (SSL) security.

1. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
2. Click **Create**.
   The New Monitor screen opens.
3. In the **Name** field, type a name for the monitor.
4. From the **Type** list, select **HTTPS**.
   The screen refreshes, and displays the configuration options for the **HTTPS** monitor type.
5. From the **Parent Monitor** list, select an existing monitor.

The new monitor inherits initial configuration values from the existing monitor.

6. From the **Configuration** list, select **Advanced**.

   This selection makes it possible for you to modify additional default settings.

7. In the **Interval** field type a number that indicates, in seconds, how frequently the system issues the monitor check. The default is `5` seconds.

   The frequency of a monitor check must be greater than the value of the global-level **Heartbeat Interval** setting. Otherwise, the monitor can acquire out-of-date data.

8. For the **Up Interval** setting, specify whether to use the up interval:

   - If you do not want to use the up interval, Retain the default, **Disabled**.
   - To use the up interval, select **Enabled**, and specify how often you want the system to verify the health of a resource that is up.

9. In the **Time Until Up** field, type a number that indicates the number of seconds to wait after a resource first responds correctly to the monitor before setting the resource to up.

   The default value is `0` (zero), which disables this option.

10. In the **Timeout** field, type a number that indicates, in seconds, how much time the target has to respond to the monitor check. The default is `16` seconds.

    If the target responds within the allotted time period, it is considered up. If the target does not respond within the time period, it is considered down.

11. For **Manual Resume**, specify whether the system automatically enables the monitored resource when the monitor check is successful.

    This setting applies only when the monitored resource has failed to respond to a monitor check.

Option	Description
Yes	The system does nothing when the monitor check succeeds, and you must manually enable the monitored resource.
No	The system automatically re-enables the monitored resource after the next successful monitor check.

12. In the **Send String** field, type a text string that the monitor sends to the target resource.

    The default string is `GET /\r\n`. This string retrieves a default file from the web site.

Important: *Send string syntax depends upon the HTTP version. Please observe the following conventions.*

Version	*Convention*
HTTP 0.9	*`"GET /\n"` or `"GET /\r\n"`.*
HTTP 1.0	*`"GET / HTTP/1.0\r\n\r\n"` or `"GET / HTTP/1.0\n\n"`*
HTTP 1.1	*`"GET / HTTP/1.1\r\nHost: server.com\r\n\r\n"` or `"GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close\r\n\r\n"`*

Type a fully qualified path name, for example, `"GET /www/example/index.html\r\n"`, if you want to retrieve a specific web site page.

13. In the **Receive String** field, type a regular expression that represents the text string that the monitor looks for in the returned resource.

The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names.

Note: *If you do not specify both a send string and a receive string, the monitor performs a simple service check and connect only.*

---

14. In the **Receive Disable String** field, type a regular expression that represents the text string that the monitor looks for in the returned resource.

    Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and **Receive Disable String** match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response.

Note: *If you choose to set the **Reverse** setting to **Yes**, the monitor marks the pool, pool member, or node **Down** when the test is successful.*

---

15. From the **SSL Profile** list, select an option for the profile:
    - To specify no SSL profile, accept the default, **None**.
    - To use a profile, select an SSL Profile from the list of the available `serverssl` profiles in the BIG-IP® system.
16. Type a name in the **User Name** field.
17. Type a password in the **Password** field.
18. For the **Reverse** setting, specify whether you want the system to work in reverse mode:
    - If you want the system to work normally, retain the **No** default option.
    - If you want the system to mark the pool, pool member, or node **Down** when the test is successful, select the **Yes** option.
19. For the **Transparent** setting, specify whether you want the monitor to operate in transparent mode:
    - If not, accept the **No** default option.
    - To use a path through the associated pool members or nodes to monitor the aliased destination, select the **Yes** option.
20. For the **Alias Address** setting, specify an alias IP address:
    - Retain the ***All Addresses** default option.
    - Type an alias IP address for the monitor to verify, on behalf of the pools or pool members with which the monitor is associated.

    If the health check for the alias address is successful, the system marks all associated objects `up`. If the health check for the alias address is not successful, then the system marks all associated objects `down`.
21. For the **Alias Service Port** setting, specify an alias port or service for the monitor to check:
    - Accept the ***All Ports** default option.
    - Select an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated.

    If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

The HTTPS monitor is configured to monitor HTTPS traffic.

When you are done, associate the HTTPS monitor with a server, pool, pool member, or node.

Configuring a monitor for adaptive response time monitoring

Determine the type of monitor you want to create, and for which custom monitor you want to enable adaptive response time monitoring.

Enable adaptive response time monitoring when you want the BIG-IP® system to update the state of a resource based on the deviation of the latency of the monitor probe from the mean latency of a monitor probe for that service.

1. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
2. Click **Create**.
   The New Monitor screen opens.
3. In the **Name** field, type a name for the monitor.
4. From the **Type** list, select the type of monitor.
   The screen refreshes, and displays the configuration options for the monitor type.
5. From the **Import Monitor** list, select an existing monitor.

   The new monitor inherits initial configuration values from the existing monitor.
6. From the **Parent Monitor** list, select an existing monitor.

   The new monitor inherits initial configuration values from the existing monitor.
7. Select the **Adaptive** check box.
   Additional settings display.
8. From the **Allowed Divergence** list, select one of these options:

Option	Description
Absolute	Type the number of milliseconds that the latency of a monitor probe can exceed the mean latency for the service being probed.
Relative	Type the percentage deviation that the latency of a monitor probe can exceed the mean latency for the service being probed.

9. In the **Adaptive Limit** field, type the maximum length of a monitor probe (in milliseconds), regardless of the calculated mean for the service being monitored.
10. In the **Sampling Timespan** field, type the length (in seconds) of the probe history span that the system uses to calculate the mean latency and standard deviation of a monitor probe.
11. Click **Finished**.

Importing a file for an external monitor

Using the BIG-IP ®Configuration utility, you can import a file from another system to use for creating an external monitor.

1. On the Main tab, click **System** > **File Management** > **External Monitor File List** > **Import**.
2. For the **File Name** setting, click **Browse**.
   The system opens a browse window so that you can locate the file that you want to import to the BIG-IP system.
3. For the **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 file name from the list.
4. Click the **Import** button.

After importing a file onto the system, you must create a local traffic external monitor, specifying the file that you imported.

Configure an SASP monitor

You configure a Server/Application State Protocol (SASP) monitor to verify the availability of resources when your network employs the IBM® Enterprise Workload Manager (EWLM, formerly Group Workload Manager).

1. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
2. Click **Create**.
   The New Monitor screen opens.
3. Type a name for the monitor.

   Important: *Monitor names are limited to 63 characters.*
4. From the **Type** list, select **SASP**.
   The screen refreshes, and displays the configuration options for the SASP monitor type.
5. In the Configuration area, from the **GWM Interval** list, select one of these options:

Option	Description
Automatic (default)	The system uses the interval setting recommended by the Group Workload Manager (GWM).
Specify	Specifies the interval the system uses to query the GWM. Type a number between 10 and 600 seconds.

6. From the **Mode** list, select one of the following:

Option	Description
Push (default)	The Group Workload Manager (GWM) decides on the interval to send Get Weights requests and the SASP monitor listens for messages.
Pull	The SASP monitor uses the GWM interval value to send the Get Weights requests to the workload manager.

7. For the **GWM Primary Address** setting, type the IP address of the Enterprise Workload Manager (formerly Gateway Workload Manager).
8. For the **GWM Secondary Address** setting, type the IP address of the backup Enterprise Workload Manager server (assuming there is a backup server).
9. For the **GWM Service Port** setting, type the number of the port through which the SASP monitor communicates with the Enterprise Workload Manager. The default is `3860`.
10. From the **GWM Protocol** list, select the communications protocol the SASP monitor uses. The default is **TCP**.
11. Click **Finished**.

Associate an SASP monitor with a pool

Before you start, make sure you have an IBM® Enterprise Workload Manager in your enterprise.

You associate a Server/Application State Protocol (SASP) monitor with a pool when configuring a load balancer.

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

The Pool List screen displays.

2. Select the name of a pool to add to an SASP monitor.
3. In the Configuration area, for the **Health Monitors** setting, move the name of an SASP monitor from the **Available** list to the **Active** list.
4. Click **Update**.

Note: *When a monitor is initially added, before it receives an initial health check from the Enterprise Workload Manager (formerly Gateway Workload Manager), on the Members tab, in the General Properties area, the* ***Availability*** *setting is Offline. If the health check is complete, in the Current Members area, the Ratio column displays values set by the workload manager for the pool members.*

Testing a monitor

Before you can test a monitor, you must save a monitor configuration. Monitor testing will not work if the monitor has already been assigned to a pool, pool member, or node.

You can test a monitor to verify a monitor configuration, before applying it to a pool, a pool member, or a node.

1. On the Main tab, click **Local Traffic** > **Monitors**.
   The Monitors List screen opens.
2. Click a monitor name in the list.
   The monitor's properties screen opens, showing the monitor's settings and values.
3. On the menu bar, click **Test**.
   The monitor's test settings display as available for configuring, and the screen shows the results of the last test of the monitor, if any.

   The results persist until either you restart the BIG-IP® system, or a new test is run.
4. Type the address and port of the monitor configuration you want to test in the **Address** fields.

Note: *If either the address or port are already configured for the health monitor itself, the* ***Address*** *field will be pre-populated with these parameters and you cannot configure them. The* ***Address*** *field parameters cannot match an existing node or pool member that is already running the monitor being tested.*

Once the required **Address** field parameters are provided, the **Test** button is available to click to start a test.

5. Click the **Test** button.

Note: *When the test is running, if you want to cancel, click the* ***Cancel*** *button.*

Test results display in the **Results** field.

Monitors Settings Reference

Health monitor functional categories

These tables describe the functional categories of health monitors, and list the available BIG-IP® monitors within each category. Unless otherwise specified, each monitor is used by Local Traffic Manager™ (LTM), BIG-IP® DNS, and Link Controller™.

Address-check monitors

An *address-check monitor* is a simple monitor that pings an IP address to verify that the address can be reached on a network.

Address-check monitor	Description
Gateway ICMP	Uses Internet Control Message Protocol (ICMP) to make a simple resource check. The check is successful if the monitor receives a response to an `ICMP_ECHO` datagram.
ICMP	Makes a simple node check. The check is successful if the monitor receives a response to an `ICMP_ECHO` datagram.
TCP Echo	Verifies Transmission Control Protocol (TCP) connections. The check is successful if the BIG-IP system receives a response to a TCP Echo message.

Service-check monitors

A *service-check monitor* determines whether a service is available by opening a connection to an IP address and port.

Service-check monitor	Description
Diameter	Monitors the servers that are running the Diameter authentication service. After configuring a Diameter monitor, associate the monitor with a load balancing pool. The BIG-IP system then attempts to establish a TCP connection with a server in the pool. After successfully establishing a connection, the Diameter monitor sends a Capabilities-Exchanging-Request (CER) message to the server. The monitor then waits to receive a Capabilities-Exchanging-Answer (CEA) message, as well as a result code of `DIAMETER_SUCCESS (2001)`.
FirePass	Checks the health of FirePass™ systems.
Inband	Performs passive monitoring as part of client requests. This monitor, when acting as a client, attempts to connect to a pool member. If the pool

Service-check monitor	Description
	member does not respond to a connection request after a user-specified number of tries within a user-specified period, the monitor marks the pool member as `down`. After the monitor has marked the pool member as `down`, and after a user-specified period has passed, the monitor again tries to connect to the pool member (if so configured).
NNTP	Checks the status of Usenet News traffic. The check is successful if the monitor retrieves a newsgroup identification line from the server. An **NNTP** monitor requires a newsgroup name (for example, `alt.cars.mercedes`) and, if necessary, a user name and password.
MSSQL	Performs service checks on Microsoft® SQL Server-based services such as Microsoft® SQL Server versions 6.5 and 7.0.
MQTT	Checks the status of an MQTT server. The check is successful if the monitor is able to connect to the server, log in as the indicated user, and log out.
MySQL	Checks the status of a MySQL™ database server. The check is successful if the monitor is able to connect to the server, log in as the indicated user, and log out.
Oracle	Checks the status of an Oracle® database server. The check is successful if the monitor is able to connect to the server, log in as the indicated user, and log out.
POP3	Checks the status of Post Office Protocol (POP) traffic. The check is successful if the monitor is able to connect to the server, log in as the indicated user, and log out. A POP3 monitor requires a user name and password.
PostgreSQL	Checks the status of a PostgreSQL database server. The check is successful if the monitor is able to connect to the server, log in as the indicated user, and log out.
RADIUS	Checks the status of Remote Access Dial-in User Service (RADIUS) servers. The check is successful if the server authenticates the requesting user. A RADIUS monitor requires a user name, a password, and a shared secret string for the code number.
RADIUS Accounting	Checks the status of Remote Access Dial-in User Service (RADIUS) accounting servers. A RADIUS Accounting monitor requires a user name and a shared secret string for the code number.
RPC	Checks the availability of specific programs that reside on a remote procedure call (RPC) server. This monitor uses the `rpcinfo` command to query

Service-check monitor	Description
	the RPC server and verify the availability of a given program.
SASP	Verifies the availability of a IBM® Enterprise Workload Manager (EWLM, formerly Group Workload Manager). This monitor uses the Server/Application State Protocol (SASP) to communicate with the EWLM. The monitor queries the EWLM for information on the current weights of each managed resource. These weights determine which resource currently provides the best response time. When the monitor receives this information from the EWLM), it configures the dynamic ratio option for the resources, allowing the BIG-IP system to select the most appropriate resource to respond to a connection request. ***Note:*** *When you assign an SASP monitor, the monitor initially marks the resources as* `down`*. This change in status occurs because the EWLM might not yet have information pertaining to its resources. As soon as the monitor receives the results of its query, it changes the status as needed. In most configurations, the monitor receives these results within a few seconds.*
SIP	Checks the status of SIP Call-ID services. By default, this monitor type issues an `SIP OPTIONS` request to a server device. However, you can use alternative protocols instead: TCP, UDP, TLS, and SIPS (that is, Secure SIP).
SMB	Verifies the availability of a Server Message Block/Common Internet File System (SMB/CIFS) server. Use this monitor to check the availability of the server as a whole, the availability of a specific service on the server, or the availability of a specific file used by a service.
SOAP	Tests a web service based on the Simple Object Access Protocol (SOAP). The monitor submits a request to a SOAP-based web service, and optionally, verifies a return value or fault.
TCP Half Open	Monitors the associated service by sending a `TCP SYN` packet to the service. As soon as the monitor receives the `SYN-ACK` packet, the monitor marks the service as `up`.
UDP	Verifies the User Datagram Protocol (UDP) service by attempting to send UDP packets to a pool, pool member, or virtual server and receiving a reply.

Content-check monitors

A *content-check monitor* sends a command to a server and examines that server's response to ensure that it is serving appropriate content.

Content-check monitor	Description
DNS	Checks the status of Domain Name System (DNS) servers, by sending a specific string, and verifying receipt of that string. The check is successful if the DNS server responds with a specified string within a specified period.
HTTP	Checks the status of Hypertext Transfer Protocol (HTTP) traffic. Like a TCP monitor, an HTTP monitor attempts to receive specific content from a web page, and unlike a TCP monitor, might send a user name and password. ***Note:*** *An HTTP monitor can monitor Outlook® Web Access (OWA) in Microsoft® Exchange Server 2007 and Microsoft® SharePoint® 2007 web sites that require NT LAN Manager (NTLM) authentication. NTLM authentication requires a send string that complies with HTTP/1.1, a user name, and a password.*
HTTPS	Checks the status of Hypertext Transfer Protocol Secure (HTTPS) traffic. An HTTPS monitor attempts to receive specific content from a web page protected by SSL security. The check is successful when the content matches the **Receive String** value. ***Note:*** *An HTTP monitor can monitor Outlook® Web Access (OWA) in Microsoft® Exchange Server 2007 and Microsoft® SharePoint® 2007 web sites that require NT LAN Manager (NTLM) authentication. NTLM authentication requires a send string that complies with HTTP/1.1, a user name, and a password.*
https_443	Checks the status of Hypertext Transfer Protocol Secure (HTTPS) traffic, by using port 443.
LDAP	Checks the status of Lightweight Directory Access Protocol (LDAP) servers. A check is successful if entries are returned for the base and filter specified. An LDAP monitor requires a user name, a password, and base and filter strings.
Scripted	Generates a simple script that reads a file that you create. The file contains `send` and `expect` strings to specify lines that you want to send or that you expect to receive.
SMTP	Checks the status of Simple Mail Transport Protocol (SMTP) servers. This monitor type checks only that the server is up and responding to

Content-check monitor	Description
	commands. The check is successful if the mail server responds to the standard `SMTP HELO` and `QUIT` commands.
TCP	Verifies the Transmission Control Protocol (TCP) service by attempting to receive specific content from a resource. The check is successful when the content matches the value of the **Receive String** setting.
WAP	Monitors Wireless Application Protocol (WAP) servers. The common usage for the WAP monitor is to specify the **Send String** and **Receive String** settings only. The WAP monitor functions by requesting a URL and finding the string in the **Receive String** setting in the data returned by the URL response.

Path-check monitors

A *path-check monitor* determines whether traffic can flow through a given device to an arbitrary endpoint. The monitor sends a packet through the network device, or to a remote server, to verify that the traffic can actually pass through the network device, and not just to the device.

Path-check monitor	Description
Gateway ICMP	Uses Internet Control Message Protocol (ICMP) to make a simple resource check. The check is successful if the monitor receives a response to an `ICMP_ECHO` datagram.
ICMP	Makes a simple node check. The check is successful if the monitor receives a response to an `ICMP_ECHO` datagram.
TCP Echo	Verifies Transmission Control Protocol (TCP) connections. The check is successful if the BIG-IP system receives a response to a `TCP Echo` message.

Application-check monitors

An *application-check monitor* is typically a custom monitor or external monitor that tests a specific application. For example, an FTP monitor connects, logs in by using a user ID and password, changes to a specified directory, and requests a specific file. This monitor succeeds when the file is received.

Application-check monitor	Description
BIG-IP	Gathers metrics and statistics information that the Local Traffic Manager (LTM) acquires through the monitoring of its own resources. Typically, it is sufficient to assign only the BIG-IP monitor to a Local Traffic Manager. When you want to verify the availability of a specific resource managed by the LTM, F5 Networks recommends that you first assign the appropriate monitor to the resource through the Local Traffic Manager, and then assign

Application-check monitor	Description
	a BIG-IP monitor to the LTM through the BIG-IP DNS (formerly GTM™). This configuration provides the most efficient means of tracking resources managed by a BIG-IP system.
BIG-IP Link	Gathers metrics and statistics information that the Link Controller™ acquires through the monitoring of its own resources. When you use BIG-IP DNS in a network that contains a Link Controller, you must assign a BIG-IP Link monitor to the Link Controller. This monitor is automatically assigned to the Link Controller if you do not manually assign it.
External	Enables you to create your own monitor type.
FTP	Attempts to download a specified file to the `/var/tmp` directory, and if the file is retrieved, the check is successful. Note that once the file has been successfully downloaded, the BIG-IP system does not save it.
IMAP	Checks the status of Internet Message Access Protocol (IMAP) traffic. An IMAP monitor is essentially a POP3 type of monitor with the addition of the Folder setting. The check is successful if the monitor is able to log into a server and open the specified mail folder.
Module Score	Enables global and local traffic management systems to load balance in a proportional manner to local traffic management virtual servers associated with the BIG-IP® Application Acceleration Manager™ and Application Security Manager™. When you configure a Module Score monitor, the local traffic management system uses SNMP to pull the `gtm_score` values from the downstream virtual servers and set the dynamic ratios on the associated upstream local traffic management pool members or nodes. The Module Score monitor retrieves the `gtm_score` values from the virtual server and the `gtm_vs_score` values associated with the virtual server. Then, if a pool name is not specified, this monitor sets the dynamic ratio on the node that is associated with the virtual server. The BIG-IP system uses the lowest non-zero value of the `gtm_vs_score` values to set the dynamic ratio. If all `gtm_vs_score` values are zero, then the `gtm_score` value is used to set the dynamic ratios. If you specify a pool name in the monitor definition, then the dynamic ratio is set on the pool member.
Virtual Location	Optimizes end-user response time in environments with dynamic distribution of application resources

Application-check monitor	Description
	across multiple data centers. When using the **Virtual Location** monitor, the BIG-IP system sets the **Priority Group** value of all local pool members to `2` (a higher priority). When a member of a load balancing pool migrates to a remote data center the Virtual Location monitor lowers the members **Priority Group** value to `1` (a lower priority). This value adjustment results in subsequent connections being sent to local pool members only if available. If no local pool members are available, connections are sent to the remote pool member.

Performance monitor functional category

This information describes the functional category of performance monitors, and lists the available BIG-IP® monitors. Unless otherwise specified, each type is used by Local Traffic Manager™, BIG-IP DNS (formerly GTM), and Link Controller™.

Performance monitors

A *performance monitor* interacts with the server (as opposed to virtual server) to examine the server load and to acquire information about the condition of virtual servers.

Performance monitor	Description
BIG-IP	Collects data from BIG-IP-DNS (formerly GTM) and Local Traffic Manager. Typically, the Local Traffic Manager probes local pool members and provides the results to BIG-IP-DNS. ***Note:*** *When the BIG-IP monitor fails, all virtual servers for that Local Traffic Manager system are marked unavailable, regardless of the results of individual virtual server probes.*
BIG-IP Link	Gathers metrics and statistics information acquired through the monitoring of BIG-IP-DNS or Link Controller resources.
SNMP	Checks the performance of a server that runs an SNMP agent to load balance to that server. A custom **snmp_gtm** import setting is assigned to servers that are not developed by F5 Networks.
SNMP DCA	Checks the performance of a server running an SNMP agent such as UC Davis, for the purpose of load balancing traffic to that server. With this monitor you can define ratio weights for CPU, memory, and disk use.
SNMP DCA Base	Checks the performance of servers that are running an SNMP agent, such as UC Davis. However, you should use this monitor only when you want the

Performance monitor	Description
	load balancing destination to be based solely on user data, and not CPU, memory, or disk use.
Real Server	Checks the performance of a node that is running the RealSystem Server data collection agent. The monitor then dynamically load balances traffic accordingly.
WMI	Checks the performance of a node that is running the Windows Management Infrastructure (WMI) data collection agent, and then dynamically load balances traffic accordingly. Generally, you would use a WMI monitor with dynamic ratio load balancing. ***Note:*** *When using the* `GetWinMediaInfo` *command with a WMI monitor, Microsoft® Windows Server® 2003 and Microsoft® Windows Server® 2008 require the applicable version of Windows Media® Services to be installed on each server.*

Diameter monitor settings

This table describes the Diameter monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Setting	Value	Description
		Important: *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Origin Host	No default	Specifies the IP address on the BIG-IP system that generates the request to the servers. If you provide no value for this setting, the system uses the self IP address on the VLAN that generates the request to the server.
Origin Realm	`f5.com`	Specifies the realm of the BIG-IP system that generates the request to the servers. By default, this value is `f5.com`.
Host IP Address	No default	Specifies the IP address of the diameter server. If no value is specified, the system uses the BIG-IP system's IP address on the VLAN that the system uses to generate traffic to the server.
Vendor ID	3375	Specifies the vendor identification number assigned to your diameter server by the Internet Assigned Numbers Authority (IANA). The default is `3375`, the IANA ID for F5 Networks.
Product Name	`F5 BIGIP Diameter Health Monitoring`	Specifies the name of the product used to monitor the servers running the Diameter service. By default, this value is `F5 BIGIP Diameter Health Monitoring`.
Auth Application ID	**None**	Specifies the Authentication and Authorization identifier for an application, as described in RFC 3588. The default is **None**. If enabled, any value that you specify must be a 32-bit unsigned value. ***Note:*** *The Auth Application ID must also be present in all Authentication and/or Authorization messages that are defined*

Setting	Value	Description
		in a separate Diameter specification and have an Application ID assigned.
Acct Application ID	**None**	Specifies the Accounting identifier for an application, as described in RFC 3588. The default is **None**. ***Note:*** *The Acct Application ID must also be present in all Accounting messages. Exactly one of the Auth Application ID attribute-value pairs and Acct Application ID attribute-value pairs can be present.*
Vendor Specific Application ID	**None**	Specifies the vendor-specific grouped values for the diameter application, as described in RFC 3588. The default is **None**. ***Note:*** *Exactly one of the Vendor Specific Auth Application ID attribute-value pairs and Vendor Specific Acct Application ID attribute-value pairs can be present. This value must also be present as the first attribute-value pair in all experimental commands defined in the vendor-specific application.*
Vendor Specific Vendor ID	No default	Specifies an attribute-value pair associated with the **Vendor Specific Application ID** monitor setting.
Vendor Specific Auth Application ID	No default	Specifies an attribute-value pair associated with the **Vendor Specific Application ID** monitor setting.
Vendor Specific Acct Application ID	No default	Specifies an attribute-value pair associated with the **Vendor Specific Application ID** monitor setting.

DNS monitor settings

This table describes the DNS monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds.

Setting	Value	Description
		Important: *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Reverse	**No**	Specifies whether the monitor operates in reverse mode. When monitor is in reverse mode, a successful receive string marks the monitored object down instead of up. You can use this mode only if you specify a receive string. The default value is **No**, which specifies that the monitor does not operate in reverse mode.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects

Setting	Value	Description
		up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Query Name	No default	Specifies a query name for the monitor to use in a DNS query, for example, `www.siterequest.com`.
Query Type	**a**	Specifies the type of DNS query that the monitor sends. The default value is **a**. This setting provides the following options. • **a**. Specifies that the monitor will send a DNS query of type A. • **aaaa**. Specifies that the monitor will send a DNS query of type AAAA.
Answer Section Contains	**Query Type**	Specifies the record types required in the answer section of the response in order to mark the status of a node up. The default value is **Query Type**. This setting includes the following options. • **Query Type**. Specifies that the response should contain at least one answer of which the resource record type matches the query type. • **Any Type**. Specifies that the DNS message should contain at least one answer. • **Anything**. Specifies that an empty answer is enough to mark the status of the node up.
Accept RCODE	**No Error**	Specifies the `RCODE` required in the response for an up status. The default value is **No Error**. This setting provides the following options. • **No Error**. Specifies that the status of the node will be marked up if the received DNS message has no error. • **Anything**. Specifies that the status of the node will be marked up irrespective of the RCODE in the DNS message received.
Receive String	No default	Specifies the IP address that the monitor uses from the resource record sections of the DNS response. The IP address should be specified in the dotted-decimal notation or IPv6 notation. The default value is none. If a receive string is not specified, the DNS message is checked against **Accept RCODE** and **Answer Section Contains** settings respectively.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.

Setting	Value	Description
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

External monitor settings

This table describes the External monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.

Setting	Value	Description
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
External Program	No default	Specifies the name of the file for the monitor to use. In order to reference a file, you must first import it using options on the

Setting	Value	Description
		System > **File Management** > **External Monitor Program File List** > **Import** screen. The BIG-IP system automatically places the file in the proper location on the file system.
Arguments	No default	Specifies any command-line arguments that the script requires.
Variables	No default	Specifies any variables that the script requires.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	***All Ports**	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is ***All Ports**. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

FirePass monitor settings

This table describes the FirePass monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value

Setting	Value	Description
		is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	0	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to 0 (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	30	Specifies the number of seconds in which the target must respond to the monitor request. The default is 30 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down.
Timeout	90	Specifies the number of seconds in which the target must respond to the monitor request. The default is 90 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	16	Specifies the number of seconds in which the target must respond to the monitor request. The default is 16 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	5	Specifies the number of seconds after which the system times out the probe request to the system. The default is 5 seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Cipher List	HIGH:!ADH	Specifies the list of ciphers for this monitor. The default list is HIGH:!ADH.
Max Load Average	12.0	Specifies the number that the monitor uses to mark the FirePass system up or down. The system compares the **Max Load Average** setting against a one-minute average of the FirePass system load. When the FirePass system-load average falls within the specified Max Load Average, the monitor marks the FirePass system up. When the average exceeds the setting, the monitor marks the system down. The default is 12.0.
Concurrency Limit	95	Specifies the maximum percentage of licensed connections currently in use under which the monitor marks the **Secure Access Manager** system up. As an example, a setting of 95 percent means that the monitor marks the FirePass system up until 95 percent of licensed connections are in use. When the number of in-use

Setting	Value	Description
		licensed connections exceeds 95 percent, the monitor marks the FirePass system down The default is 95.
User Name	`gtmuser`	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***Username*** *and* ***Password*** *settings.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

FTP monitor settings

This table describes the FTP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.

Setting	Value	Description
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication.

Setting	Value	Description
		Important: *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Path/Filename	No default	Specifies the full path and file name of the file that the system attempts to download. The health check is successful if the system can download the file.
Mode	**Passive**	• **Passive**. Specifies the data transfer process (DTP) mode. The default is **Passive**. • **Port**. Specifies that the monitor initiates and establishes the data connection with the FTP server.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

Gateway ICMP monitor settings

This table describes the Gateway ICMP monitor configuration settings and default values.

Important:

Use the BIG-IP monitor to monitor a BIG-IP virtual server with a virtual address that overlaps a non-floating IPv6 or IPv4 self IP address. Do not use any other BIG-IP DNS monitor to monitor a virtual server with a virtual address that overlaps a non-floating self IP address.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.

Setting	Value	Description
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Interval	`1`	Specifies, in seconds, the frequency at which the system probes the host server. The default is `1` second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.

Setting	Value	Description
Probe Attempts	`3`	Specifies the number of times that the system attempts to probe the host server, after which the system considers the host server down or unavailable. The default value is `3`.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections.*
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options:

Setting	Value	Description
		Absolute The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

HTTP monitor settings

This table describes the HTTP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater*

Setting	Value	Description
		should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	`GET /`	Specifies the text string that the monitor sends to the target object. You must include `\r\n` at the end of a non-empty send string. The default setting is `GET /\r\n`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example: `GET /www/siterequest/index.html\r\n`. ***Note:***

Setting	Value	Description
		When the send string specifies `HTTP/1.0` or `HTTP/1.1`, the monitor checks the result code before indicating the monitor as up. Additionally, the server response code must include a `200` status code, regardless of the receive-string content, in order for the monitor to mark the server as up. The monitor marks the server as down for any other response, without further processing. When the send string does not specify `HTTP/1.0` or `HTTP/1.1`, the monitor uses `HTTP/0.9` and makes no response code checks. Search string matches on the received reply can further affect the result. ***Important:*** *When you create a new TCP, HTTP, or HTTPS monitor in version 10.2.0 and later, you must include a return and new-line entry (`\r\n`) at the end of a non-empty send string, for example `GET /\r\n` instead of `GET /`. If you do not include `\r\n` at the end of the send string, the TCP, HTTP, or HTTPS monitor fails. When you include a host in a send string, you must duplicate the return and new-line entries (`\r\n\r\n`), for example, `"GET / HTTP/1.1\r\nHost: server.com\r\n\r\n"` or `"GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close\r\n\r\n"`.*
Receive String	No default	Specifies the regular expression representing the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. The only monitors that support regular expression matching are HTTP, HTTPS, TCP, and UDP monitors. ***Note:*** *If you do not specify both a **Send String** and a **Receive String**, the monitor performs a simple service check and connect only.*
Receive Disable String	No default	Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and **Receive Disable String** match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response. ***Note:*** *If you choose to set the **Reverse** setting to **Yes**, the **Receive Disable String** option becomes unavailable and the monitor marks the pool, pool member, or node **Down** when the test is successful.*
User Name	No default	Specifies the user name, if the monitored target requires authentication.

Setting	Value	Description
		Important: *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Reverse	**No**	Instructs the system to mark the target resource down when the test is successful. This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you might want to set up a reverse ECV service check that looks for the string `Error`. A match for this string means that the web server was down. You can use **Reverse** only if you configure both **Send String** and **Receive String**.
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects

Setting	Value	Description
		three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

HTTPS monitor settings

This table describes the HTTPS monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds.

Setting	Value	Description
		Important: *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	`GET /`	Specifies the text string that the monitor sends to the target object. You must include `\r\n` at the end of a non-empty send string. The default setting is `GET /\r\n`, which retrieves a default HTML file for a web site. To retrieve a specific page from a web site, specify a fully-qualified path name, for example: `GET /www/siterequest/index.html\r\n`.

Setting	Value	Description
		Note: When the send string specifies `HTTP/1.0` or `HTTP/1.1`, the monitor checks the result code before indicating the monitor as up. Additionally, the server response code must include a `200` status code, regardless of the receive-string content, in order for the monitor to mark the server as up. The monitor marks the server as down for any other response, without further processing. When the send string does not specify `HTTP/1.0` or `HTTP/1.1`, the monitor uses `HTTP/0.9` and makes no response code checks. Search string matches on the received reply can further affect the result. ***Important:*** *When you create a new TCP, HTTP, or HTTPS monitor in version 10.2.0 and later, you must include a return and new-line entry (`\r\n`) at the end of a non-empty send string, for example `GET /\r\n` instead of `GET /`. If you do not include `\r\n` at the end of the send string, the TCP, HTTP, or HTTPS monitor fails. When you include a host in a send string, you must duplicate the return and new-line entries (`\r\n\r\n`), for example, `"GET / HTTP/1.1\r\nHost: server.com\r\n\r\n"` or `"GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close \r\n\r\n"`.*
Receive String	No default	Specifies the regular expression representing the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. The only monitors that support regular expression matching are HTTP, HTTPS, TCP, and UDP monitors. ***Note:*** *If you do not specify both a **Send String** and a **Receive String**, the monitor performs a simple service check and connect only.*
Receive Disable String	No default	Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and **Receive Disable String** match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response. ***Note:*** *If you choose to set the **Reverse** setting to **Yes**, the **Receive Disable String** option becomes unavailable and the monitor marks the pool, pool member, or node **Down** when the test is successful.*

Setting	Value	Description
Cipher List	`DEFAULT:+SHA:+3DES:+kEDH`	Specifies the list of ciphers for this monitor. The default list is `DEFAULT:+SHA:+3DES:+kEDH`.
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Compatibility	**Enabled**	Specifies, when enabled, that the SSL options setting (in OpenSSL) is set to **ALL**. The default is **Enabled**.
Client Certificate	**None**	For **TLS** and **SIPS** modes only, specifies a client certificate that the monitor sends to the target SSL server. The default is **None**.
Client Key	**None**	For **TLS** and **SIPS** modes only, specifies a key for a client certificate that the monitor sends to the target SSL server. The default is **None**.
Reverse	**No**	Instructs the system to mark the target resource down when the test is successful. This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you might want to set up a reverse ECV service check that looks for the string `Error`. A match for this string means that the web server was down. You can use **Reverse** only if you configure both **Send String** and **Receive String**.
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor.

Setting	Value	Description
		Enabled The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

ICMP monitor settings

This table describes the ICMP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.

Setting	Value	Description
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address-Alias Service Port** combination specified in the monitor). The default is **No**.

Setting	Value	Description
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

IMAP monitor settings

This table describes the IMAP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.

Setting	Value	Description
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Folder	INBOX	Specifies the name of the folder on the IMAP server that the monitor tries to open.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

Inband monitor settings

This table describes the Inband monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.

Setting	Value	Description
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Failures	3	Specifies the number of failed responses that a pool member may send in the **Failure Interval** before the monitor marks the pool member down. The total number of failures can be any combination of failed connection attempts or failures to return data within the interval specified in the **Response Time** box. The default is 3. ***Note:*** *Systems with multiple* tmm *processes use a per-process number to calculate failures, depending on the specified load balancing method. For example, for the Round Robin load balancing method, if any tmm receives* L *failures, the node will be marked down by that tmm.*
Failure Interval	30	Specifies that if the system receives the specified number of **Failures** within this period of time, the monitor marks the pool member down. The default is 30 seconds.
Response Time	10	Specifies the interval in which a pool member must respond with data. If the pool member responds after the specified amount of time, the monitor reports a failure. Specifying a value of 0 (zero) disables this feature. The default is 10 seconds.
Retry Time	300	Specifies the period of time a monitor waits after marking a pool member down, before the monitor requests status from that pool member. If you specify a value of 0 (zero), once the inband monitor marks a pool member down, that pool member is not marked up without outside intervention, either by explicitly marking the pool member up, or by using by using the **Check Until Up** setting in any other monitor (except another **Inband** monitor) configured on the same pool member. (In this case, the other monitor is known as the active monitor, and the **Inband** monitor is known as the passive monitor. If you have this active-passive monitor configuration, do not set **Retry Time** to a value other than 0 (zero). For this active-passive monitor configuration, the active monitor should be the one to mark the pool member up, and setting a value here could result in a possible conflict between two separate processes marking a pool member up at different times.) The default is 300 seconds.

LDAP monitor settings

This table describes the LDAP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**.

Setting	Value	Description
		Note: *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Base	No default	Specifies the location in the LDAP tree from which the monitor starts the health check. A sample value is: `dc=bigip-test,dc=net`
Filter	No default	Specifies an LDAP key for which the monitor searches. A sample value is: `objectclass=*`.
Security	**None**	Specifies the secure protocol type for communications with the target. The default is **None**.
Mandatory Attributes	**No**	Specifies whether the target must include attributes in its response to be considered up. The default is **No**.
Chase Referrals	**Yes**	Specifies whether, upon receipt of an LDAP referral entry, the target follows (or chases) that referral. The default is **Yes**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

Module Score monitor settings

This table describes the Module Score monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`30`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `30` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down.
SNMP Community	`Public`	Specifies the community name that the system must use to authenticate with the host server through SNMP. The default value is `public`. Note that this value is case sensitive.
SNMP Version	**v2c**	Specifies the version of SNMP that the host server uses. The default is **v2c**.
SNMP IP Address	No default	Specifies the IP address the system uses for communicating the module score information.
SNMP Port	`161`	Specifies the port associated with the IP address the system uses for communicating the module score information.
Pool Name	No default	Requires the name of an existing pool.

MQTT monitor settings

This table describes the MQTT monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**.

Setting	Value	Description
		Note: *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
MQTT Version	**3.1.1**	Specifies the protocol version that the monitor will use to communicate with the monitoring object. The default is **3.1.1**.
Client ID	No default	Specifies the Client ID that the monitor will send to communicate with the monitoring object.
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

MSSQL monitor settings

This table describes the MSSQL monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.

Setting	Value	Description
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`91`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `91` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: `SELECT * FROM <db_name>`. This is an optional setting.

Setting	Value	Description
		If you do not specify a send string, the monitor simply tries to establish a connection with the target. If the monitor is successful, the system marks the target up. If the system cannot establish the connection, then it marks the target down.
Receive String	No default	Specifies the response the monitor expects from the target, when the target receives the send string. This is an optional setting, and is applicable only if you configure the **Send String** setting.
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Database	No default	Specifies the name of the database that the monitor tries to access, for example, `sales` or `hr`.
Receive Row	No default	Specifies the row in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Receive Column	No default	Specifies the column in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Count	`0`	Specifies how the system handles open connections for monitor instances. The default is `0` (zero). By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. This functionality allows you to assign multiple instances to the database while reducing the overhead that multiple open connections could cause. The **Count** option allows you to determine the number of instances for which the system keeps a connection open.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system

Setting	Value	Description
		does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

MySQL monitor settings

This table describes the MySQL monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`91`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `91` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: `SELECT * FROM <db_name>`. This is an optional setting. If you do not specify a send string, the monitor simply tries to establish a connection with the target. If the monitor is successful, the system marks the target up. If the system cannot establish the connection, then it marks the target down.
Receive String	No default	Specifies the response the monitor expects from the target, when the target receives the send string. This is an optional setting, and is applicable only if you configure the Send String setting. ***Note:*** *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Database	No default	Specifies the name of the database that the monitor tries to access, for example, `sales` or `hr`.
Receive Row	No default	Specifies the row in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Receive Column	No default	Specifies the column in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Count	`0`	Specifies how the system handles open connections for monitor instances. The default is `0` (zero). By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. This functionality allows you to assign multiple instances to the database while reducing the overhead that multiple open connections could cause. The **Count** option allows you to determine the number of instances for which the system keeps a connection open.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias

Setting	Value	Description
		address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

NNTP monitor settings

This table describes the NNTP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down.

Setting	Value	Description
		Important: *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Newsgroup	No default	Specifies the name of the newsgroup that you are monitoring, for example `alt.car.mercedes`.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the

Setting	Value	Description
		health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

Oracle monitor settings

This table describes the Oracle monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*

Setting	Value	Description
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`91`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `91` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: `SELECT * FROM <db_name>`. This is an optional setting. If you do not specify a send string, the monitor simply tries to establish a connection with the target. If the monitor is successful, the system marks the target up. If the system cannot establish the connection, then it marks the target down.
Receive String	No default	Specifies the response the monitor expects from the target, when the target receives the send string. This is an optional setting, and is applicable only if you configure the Send String setting. ***Note:*** *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Connection String	No default	Specifies the name of the database that the monitor tries to access, for example, `sales` or `hr`.

Setting	Value	Description
		An example for this entry is as follows, where you specify the IP address for the node being monitored, the port for the node being monitored, and the name for the database: `(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=%node_ip%)(PORT=%node_port%)) (CONNECT_DATA=(SID=<db name>)) (SERVER=dedicated))`
Database	`%node_ip%:%node_port%:`	Specifies the name of the database that the monitor tries to access, for example, `sales` or `hr`.
Receive Row	No default	Specifies the row in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Receive Column	No default	Specifies the column in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Count	`0`	Specifies how the system handles open connections for monitor instances. The default is `0` (zero). By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. This functionality allows you to assign multiple instances to the database while reducing the overhead that multiple open connections could cause. The **Count** option allows you to determine the number of instances for which the system keeps a connection open.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

POP3 monitor settings

This table describes the POP3 monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.

Setting	Value	Description
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

PostgreSQL

This table describes the PostgreSQL monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.

Setting	Value	Description
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`91`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `91` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: `SELECT * FROM <db_name>`. This is an optional setting. If you do not specify a send string, the monitor simply tries to establish a connection with the target. If the monitor is successful, the system marks the target up. If the system cannot establish the connection, then it marks the target down.

Setting	Value	Description
Receive String	No default	Specifies the response the monitor expects from the target, when the target receives the send string. This is an optional setting, and is applicable only if you configure the Send String setting. ***Note:*** *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Database	No default	Specifies the name of the database that the monitor tries to access, for example, `sales` or `hr`.
Receive Row	No default	Specifies the row in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Receive Column	No default	Specifies the column in the database where the specified **Receive String** should be located. This is an optional setting, and is applicable only if you configure the **Send String** and the **Receive String** settings.
Count	`0`	Specifies how the system handles open connections for monitor instances. The default is `0` (zero). By default, when you assign instances of this monitor to a resource, the system keeps the connection to the database open. This functionality allows you to assign multiple instances to the database while reducing the overhead that multiple open connections could cause. The **Count** option allows you to determine the number of instances for which the system keeps a connection open.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to

Setting	Value	Description
		this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

RADIUS monitor settings

This table describes the RADIUS monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.

Setting	Value	Description
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Secret	No default	Specifies the secret the monitor needs to access the resource.
NAS IP Address	No default	Specifies the network access server's IP address (NAS IP address) for a RADIUS monitor.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

RADIUS Accounting monitor settings

This table describes the RADIUS Accounting monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.

Setting	Value	Description
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.

Setting	Value	Description
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Secret	No default	Specifies the secret the monitor needs to access the resource.
NAS IP Address	No default	Specifies the network access server's IP address (NAS IP address) for a RADIUS monitor.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

Real Server monitor settings

This table describes the Real Server monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is

Setting	Value	Description
		down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Method	**GET**	Displays the method the monitor uses to contact the server. The setting is **GET**. You cannot modify the method.
Command	**GetServerStats**	Specifies the command that the system uses to obtain the metrics from the resource.
Metrics	`ServerBandwidth: 1.5, CPUPercentUsage, MemoryUsage, TotalClientCount`	Specifies the performance metrics that the commands collect from the target. The default is `ServerBandwidth:1.5, CPUPercentUsage, MemoryUsage, TotalClientCount`.
Agent	**Mozilla/4.0 (compatible: MSIE 5.0; Windows NT)**	Displays the agent for the monitor. The default agent is **Mozilla/4.0 (compatible: MSIE 5.0; Windows NT)**. You cannot modify the agent.

RPC monitor settings

This table describes the RPC monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**.

Setting	Value	Description
		Note: *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Mode	**TCP**	Specifies whether the monitor should verify the availability of the RPC server through TCP or UDP.
Program Number	No default	Specifies the number of the program or application whose availability the monitor needs to verify.
Version Number	No default	Specifies an exact version number of the program identified in the **Program Number** setting. This setting verifies that a version of the given program is available.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

SASP monitor settings

This table describes the SASP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
GWM Interval	**Automatic**	Specifies the frequency at which the system queries Group Workload Manager (GWM). The default is **Automatic**.

Setting	Value	Description
Mode	**Pull**	Specifies whether the load balancer should send Get Weight Request messages (Pull) or receive Send Weights messages (Push) from the GWM server. The default is **Pull**. When configured in the **Pull** mode, the monitor polls the pool member weights by periodically sending a Get Weights Request message to the GWM server. When configured in the **Push** mode, the monitor waits indefinitely to receive pool member weights by means of Send Weights messages from the GWM server. The SASP monitor updates the dynamic ratio for the pool members once it receives the weights.
GWM Primary Address	No default	Specifies the IP address of the primary GWM server.
GWM Secondary Address	No default	Specifies the IP address of the secondary GWM server. When both the GWM primary address and GWM secondary address are configured, but the GWM primary address or GWM secondary address is unreachable, the monitor attempts to reconnect to the unreachable address every 30 seconds. When both the GWM primary address and GWM secondary address are available, only the weights reported by the primary address are used to update the pool-member dynamic ratio. When the GWM primary address is unavailable, the monitor uses the weights reported by the GWM secondary address to update the pool-member dynamic ratio. If the primary address again becomes available, then the monitor uses the weights reported by the primary address to update the pool-member dynamic ratio. When both the GWM primary address and GWM secondary address are unavailable, the monitor uses the weights reported by the first GWM address that becomes available.
GWM Service Port		Specifies the port through which the SASP monitor communicates with the Group Workload Manager. The default is `3860`.
GWM Protocol	**TCP**	Specifies the communications protocol the monitor uses. You can specify **TCP** or **UDP**. The default is **TCP**.

Scripted monitor settings

This table describes the Scripted monitor configuration settings and default values.

When using scripts for monitor settings, you will want to observe the following conditions.

- Scripts must use hard-return line endings (`LF`), not soft-return line endings (`CR-LF`).
- Exactly one character space must be used to separate the `send` or `expect` instruction keywords from the text to send or match.
- The text to send or match extends to the end of the line, even when using quotation marks. Any characters that follow a closing quotation mark will break the match.
- Matching text can match the prefix of a response, but cannot match a substring that is not a prefix, that is, a substring that starts other than at the beginning of the response.

Additionally, within scripts, the following escape sequences apply.

Name	Escape Sequence
Bell	\a
Backspace	\b
Form feed	\f
New line	\n
Return	\r
Tab	\t
Vertical tab	\v
Backslash	\\
Single quotation mark	\'

For example, the following script specifies a simple SMTP sequence. Note that the lines of the file are always read in the sequence specified.

```
expect 220
send "HELO bigip1.somecompany.net\r\n"
expect "250"
send "quit\r\n"
```

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	10	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 10 seconds.
Up Interval	Disabled	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	0	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to 0 (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.

Setting	Value	Description
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
File Name	No default	Specifies the name of a file in the `/config/eav/` directory on the system. The user-created file contains the and data that the monitor uses for the monitor check.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

SIP monitor settings

This table describes the SIP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.

Setting	Value	Description
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections.*
Mode	**UDP**	Specifies the protocol that the monitor uses to communicate with the target object. The default is **UDP**.
Client Certificate	**None**	For **TLS** and **SIPS** modes only, specifies a client certificate that the monitor sends to the target SSL server. The default is **None**.
Client Key	**None**	For **TLS** and **SIPS** modes only, specifies a key for a client certificate that the monitor sends to the target SSL server. The default is **None**.
Additional Accepted Status Codes	**None**	Specifies the additional SIP status codes that the monitor uses to determine target status. The default is **None**. ***Note:*** *The monitor always marks the target up in response to status code* `200 OK`.
Additional Rejected Status Codes	**Status Code List**	This list functions identically to the **Additional Accepted Status Codes** list, except that the monitor treats the list items as error codes, rather than success codes, and so marks the target down.
Header List	No default	Specifies one or more headers that the monitor recognizes.
SIP Request	No default	Type the request line of the SIP message, specifying a complete SIP request line minus the trailing `\r\n` characters. The system uses the response code to determine whether the server is up or down. The monitor performs a simple, customized query to a SIP server. The monitor does not establish connections, perform hand-shaking, or process SIP traffic or requests. It only sends a request to a server and looks at the response code and (aside from matching the response to the request) ignores the rest of the response. As a result, this monitor does not support requests such as `INVITE`, because the monitor does not enter into a dialog.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to

Setting	Value	Description
		this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

SMB monitor settings

This table describes the SMB monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Path/Filename	No default	Specifies a specific file associated with a service. The monitor uses the relative path to the service itself when attempting to locate the file. You are not required to specify a value for this option; however, if you elect to use this option you must also specify a value for **Service Name**.
SMB/CIFS Server	No default	Specifies the NetBIOS server name of the SMB/CIFS server for which the monitor checks for availability. You must specify a server for this monitor to function.
Service Name	No default	Specifies a specific service on the SMB/CIFS for which you want to verify availability. You are not required to specify a service name.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

SMTP monitor settings

This table describes the SMTP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this

Setting	Value	Description
		attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections.*
Domain	No default	Specifies the domain name to check, for example, bigipinternal.com.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

SNMP DCA monitor settings

This table describes the SNMP DCA monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.

Setting	Value	Description
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`30`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `30` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down.
Community	**Public**	Specifies the community name that the system must use to authenticate with the host server through SNMP. The default value is **public**. Note that this value is case sensitive.
Version	**v1**	Specifies the version of SNMP that the host server uses. The default is **V1**.
Agent Type	**UCD**	Specifies the SNMP agent running on the monitored server. The default is **UCD** (UC-Davis).
CPU Coefficient	`1.5`	Specifies the coefficient that the system uses to calculate the weight of the CPU threshold in the dynamic ratio load balancing algorithm. The default is **1.5**.
CPU Threshold	`80`	Specifies the maximum acceptable CPU usage on the target server. The default is `80` percent.
Memory Coefficient	`1.0`	Specifies the coefficient that the system uses to calculate the weight of the memory threshold in the dynamic ratio load balancing algorithm. The default is `1.0`.
Memory Threshold	`70`	Specifies the maximum acceptable memory usage on the target server. The default is `70` percent.
Disk Coefficient	`2.0`	Specifies the coefficient that the system uses to calculate the weight of the disk threshold in the dynamic ratio load balancing algorithm. The default is `2.0`.
Disk Threshold	`90`	Specifies the maximum acceptable disk usage on the target server. The default is `90` percent.
Variables	No default	Presents text fields for specifying unique variable names and value pairs (which represent coefficient and threshold values for

Setting	Value	Description
		other types of data, such as user metrics) and a list containing existing variable definitions that the monitor uses.

SNMP DCA Base monitor settings

This table describes the SNMP DCA Base monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`30`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `30` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down.
Community	**Public**	Specifies the community name that the system must use to authenticate with the host server through SNMP. The default value is **public**. Note that this value is case sensitive.
Version	**v1**	Specifies the version of SNMP that the host server uses. The default is **V1**.
Variables	No default	Presents text fields for specifying unique variable names and value pairs (which represent coefficient and threshold values for other types of data, such as user metrics) and a list containing existing variable definitions that the monitor uses.

SOAP monitor settings

This table describes the SOAP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks

Setting	Value	Description
		the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings ["""] for the* ***User Name*** *and* ***Password*** *settings.*
Protocol	**HTTP**	Specifies the protocol that the monitor uses for communications with the target. The default is **HTTP**.
URL Path	No default	Specifies the URL for the web service that you are monitoring, for example, `/services/myservice.aspx`.
Namespace	No default	Specifies the name space for the web service you are monitoring, for example, `http://example.com/`.
Method	No default	Specified the method by which the monitor contacts the resource.
Parameter Name	No default	Specifies, if the method has parameters, the parameter name.
Parameter Type	**Bool**	Specifies the parameter type. The default is **bool** (boolean).
Parameter Value	No default	Specifies the value for the parameter.
Return Type	**Bool**	Specifies the type for the returned parameter. The default is **bool** (boolean).
Return Value	No default	Specifies the value for the returned parameter.
Expect Fault	**No**	Specifies whether the method causes the monitor to expect a SOAP fault message. The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The

Setting	Value	Description
		default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

TCP monitor settings

This table describes the TCP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.

Setting	Value	Description
Interval	5	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	Disabled	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	0	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to 0 (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	16	Specifies the number of seconds in which the target must respond to the monitor request. The default is 16 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the text string that the monitor sends to the target object.
Receive String	No default	Specifies the regular expression representing the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. The only monitors that support regular expression matching are HTTP, HTTPS, TCP, and UDP monitors. ***Note:*** *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*

Setting	Value	Description
Receive Disable String	No default	Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and **Receive Disable String** match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response. ***Note:*** *If you choose to set the* ***Reverse*** *setting to* ***Yes****, the* ***Receive Disable String*** *option becomes unavailable and the monitor marks the pool, pool member, or node* ***Down*** *when the test is successful.*
Reverse	**No**	Instructs the system to mark the target resource down when the test is successful. This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you might want to set up a reverse ECV service check that looks for the string `Error`. A match for this string means that the web server was down. You can use **Reverse** only if you configure both **Send String** and **Receive String**.
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.

Setting	Value	Description
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

TCP Echo monitor settings

This table describes the TCP Echo monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds.

Setting	Value	Description
		Important: *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.

TCP Half Open monitor settings

This table describes the TCP Half Open monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Interval	`1`	Specifies, in seconds, the frequency at which the system probes the host server. The default is `1` second.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Probe Attempts	`3`	Specifies the number of times that the system attempts to probe the host server, after which the system considers the host server down or unavailable. The default value is `3`.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the

Setting	Value	Description
		BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

UDP monitor settings

This table describes the UDP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.

Setting	Value	Description
Interval	30	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 30 seconds.
Timeout	120	Specifies the number of seconds in which the target must respond to the monitor request. The default is 120 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Interval	1	Specifies, in seconds, the frequency at which the system probes the host server. The default is 1 second.
Probe Timeout	5	Specifies the number of seconds after which the system times out the probe request to the system. The default is 5 seconds.
Probe Attempts	3	Specifies the number of times that the system attempts to probe the host server, after which the system considers the host server down or unavailable. The default value is 3.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	5	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is 5 seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	Disabled	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is 0 (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	0	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to 0 (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	16	Specifies the number of seconds in which the target must respond to the monitor request. The default is 16 seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to **Yes**, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	`default send string`	Specifies the text string that the monitor sends to the target object. The default is `default send string`.
Receive String	No default	Specifies the regular expression representing the text string that the monitor looks for in the returned resource. The most common receive expressions contain a text string that is included in an HTML file on your site. The text string can be regular text, HTML tags, or image names, and the associated operation is not case-sensitive. The only monitors that support regular expression matching are HTTP, HTTPS, TCP, and UDP monitors. ***Note:*** *If you do not specify both a **Send String** and a **Receive String**, the monitor performs a simple service check and connect only.*
Receive Disable String	No default	Use a **Receive String** value together with a **Receive Disable String** value to match the value of a response from the origin web server and create one of three states for a pool member or node: **Up (Enabled)**, when only **Receive String** matches the response, or when both **Receive String** and **Receive Disable String** match the response; **Up (Disabled)**, when only **Receive Disable String** matches the response; or **Down**, when neither **Receive String** nor **Receive Disable String** matches the response. ***Note:*** *If you choose to set the **Reverse** setting to **Yes**, the **Receive Disable String** option becomes unavailable and the monitor marks the pool, pool member, or node **Down** when the test is successful.*
Reverse	**No**	Instructs the system to mark the target resource down when the test is successful. This setting is useful, for example, if the content on your web site home page is dynamic and changes frequently, you might want to set up a reverse ECV service check that looks for the string `Error`. A match for this string means that the web server was down. You can use **Reverse** only if you configure both **Send String** and **Receive String**.
Transparent	**No**	Specifies whether the monitor operates in transparent mode. A monitor in transparent mode uses a path through the associated pool members or nodes to monitor the aliased destination (that is, it monitors the **Alias Address**-**Alias Service Port** combination specified in the monitor). The default is **No**.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.

Setting	Value	Description
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.
Adaptive	Disabled	Specifies whether adaptive response time monitoring is enabled for this monitor. **Enabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings, and the divergence from the mean latency of a monitor probe for that service. You can set values for the **Allowed Divergence**, **Adaptive Limit**, and **Sampling Timespan** monitor settings. **Disabled** The monitor determines the state of a service based on the **Interval**, **Up Interval**, **Time Until Up**, and **Timeout** monitor settings.
Allowed Divergence	Relative, 25%	Specifies the type of divergence used when the **Adaptive** setting is enabled (check box selected). In typical cases, if the monitor detects three consecutive probes that miss the latency value you set, the system marks the pool member or node as down. There are two options: **Absolute** The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed. **Relative** The percentage of deviation the latency of a monitor probe can exceed the mean latency of a monitor probe for the service being probed.
Adaptive Limit	200 milliseconds	Specifies the maximum number of milliseconds that the latency of a monitor probe can exceed the mean latency of a monitor probe, for the service being probed. This value applies regardless of the **Allowed Divergence** setting value. For example, when the **Adaptive** setting is enabled (check box selected) with a value set to `500`, the monitor probe latency cannot exceed 500 milliseconds, even if that value is below the value of the **Allowed Divergence** setting.
Sampling Timespan	300 seconds (5 minutes)	Specifies the length, in seconds, of the probe history window that the system uses to calculate the mean latency and standard deviation of a monitor probe. For example, when the **Adaptive** setting is enabled

Setting	Value	Description
		(check box selected) with a value set to `300` seconds (that is five minutes), then the BIG-IP system uses the last five minutes of probe history to determine the mean latency and standard deviation of a probe.

Virtual Location monitor settings

This table describes the Virtual Location monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Pool Name	No default	Requires the name of an existing pool.

WAP monitor settings

This table describes the WAP monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds.
Interval	`10`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `10` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Up Interval	`Disabled`	Specifies, in seconds, the frequency at which the system issues the monitor check when the resource is up. The enabled default value is `0` (zero), which specifies that the system uses the value of the interval option whether the resource is up or down. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`31`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `31` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.

Setting	Value	Description
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Manual Resume	**No**	Specifies whether the system automatically changes the status of a resource to Enabled at the next successful monitor check. The default is **No**. ***Note:*** *If you set this option to* ***Yes****, you must manually re-enable the resource before the system can use it for load balancing connections.*
Send String	No default	Specifies the text string that the monitor sends to the target object.
Receive String	No default	Specifies the response the monitor expects from the target, when the target receives the send string. This is an optional setting, and is applicable only if you configure the Send String setting. ***Note:*** *If you do not specify both a* ***Send String*** *and a* ***Receive String****, the monitor performs a simple service check and connect only.*
Secret	No default	Specifies the secret the monitor needs to access the resource.
Accounting Node	No default	Specifies the RADIUS server that provides authentication for the WAP target. This setting is optional. Note that if you configure the Accounting Port, but you do not configure the Accounting Node, the system assumes that the RADIUS server and the WAP server are the same system.
Accounting Port	No default	Specifies the port that the monitor uses for RADIUS accounting. The default is `0`, which disables RADIUS accounting.
Server ID	No default	Specifies the RADIUS NAS-ID for this system, in the RADIUS server's configuration.
Call ID	No default	Specifies the 11-digit phone number for the RADIUS server. This setting is optional.
Session ID	No default	Specifies the RADIUS session identification number. This setting is optional.
Framed Address	No default	Specifies the RADIUS framed IP address. This setting is optional.
Alias Address	`*All Addresses`	Specifies an alias IP address for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Addresses`. If the health check for the alias address is successful, the system marks all associated objects up. If the health check for the alias address is not successful, then the system marks all associated objects down.
Alias Service Port	`*All Ports`	Specifies an alias port or service for the monitor to check, on behalf of the pools or pool members with which the monitor is associated. The default setting is `*All Ports`. If the health check for the alias port or service is successful, the system marks all associated objects up. If the health check for the alias port or service is not successful, then the system marks all associated objects down.

Setting	Value	Description
Debug	**No**	Specifies whether the monitor sends error messages and additional information to a log file created and labeled specifically for this monitor. The default setting is **No**, which specifies that the system does not redirect error messages and additional information related to this monitor. The **Yes** setting specifies that the system redirects error messages and additional information to the `/var/log/<monitor_type>_<ip_address>.<port>.log` file.

WMI monitor settings

This table describes the WMI monitor configuration settings and default values.

Setting	Value	Description
Name	No default	Provides a name for the monitor.
Description	No default	Provides a description of the monitor.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import SettingsParent Monitor	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	`30`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `30` seconds.
Timeout	`120`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `120` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value.
Probe Timeout	`5`	Specifies the number of seconds after which the system times out the probe request to the system. The default is `5` seconds.
Ignore Down Response	**No**	Specifies that the monitor allows more than one probe attempt per interval. The default is **No**.
Interval	`5`	Specifies, in seconds, the frequency at which the system issues the monitor check when either the resource is down or the status of the resource is unknown. The default value is `5` seconds. ***Important:*** *F5 Networks recommends that when you configure this option and the* ***Up Interval*** *option, whichever value is greater should be a multiple of the lesser value to allow for an even distribution of monitor checks among all monitors.*

Setting	Value	Description
Time Until Up	`0`	Delays the marking of a pool member or node as up for the specified number of seconds after receiving the first correct response. When this attribute is set to `0` (the default value), the BIG-IP system marks the resource as up immediately after receiving the first correct response.
Timeout	`16`	Specifies the number of seconds in which the target must respond to the monitor request. The default is `16` seconds. If the target responds within the set time period, the target is considered to be up. If the target does not respond within the set time period, the target is considered to be down. The Timeout value should be three times the Interval value, plus one second.
User Name	No default	Specifies the user name, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Password	No default	Specifies the password, if the monitored target requires authentication. ***Important:*** *If there is no password security, you must use blank strings [""] for the* ***User Name*** *and* ***Password*** *settings.*
Method	**POST**	Displays the method the monitor uses to contact the server. The setting is **POST**. You cannot modify the method.
URL	`/scripts/F5Isapi.dll`	Specifies the URL that the monitor uses. The default is `/scripts/f5Isapi.dll`.
Command	`GetCPUInfo`, `GetDiskInfo`, `GetOSInfo`	Specifies the command that the system uses to obtain the metrics from the resource. See the documentation for the resource for information on available commands. The default is `GetCPUInfo, GetDiskInfo, GetOSInfo`. ***Note:*** *When using the* `GetWinMediaInfo` *command with a WMI monitor, Microsoft® Windows Server® 2003 and Microsoft® Windows Server® 2008 require the applicable version of Windows Media® Services to be installed on each server.*

Metrics	`LoadPercentage`,
	`DiskUsage`,
	`PhysicalMemoryUsage:`
	`1.5`,
	`VirtualMemoryUsage:`
	`2.0`
		Specifies the performance metrics that the commands collect from the target. The default is `LoadPercentage, DiskUsage, PhysicalMemoryUsage:1.5, VirtualMemoryUsage:2.0`.

Setting	Value	Description
Agent	`Mozilla/4.0 (compatible: MSIE 5.0; Windows NT)`	Displays the agent for the monitor. The default agent is `Mozilla/4.0 (compatible: MSIE 5.0; Windows NT)`. You cannot modify the agent.
Post	`RespFormat=HTML`	Displays the mechanism that the monitor uses for posting. The default is `RespFormat=HTML`. You cannot change the post format for WMI monitors.

Legal Notices

Legal notices

Publication Date

This document was published on November 13, 2017.

Publication Number

MAN-0470-05

Copyright

Trademarks

For a current list of F5 trademarks and service marks, see *http://www.f5.com/about/guidelines-policies/trademarks*.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at: *https://f5.com/about-us/policies/patents*.

Link Controller Availability

This product is not currently available in the U.S.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a

residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index