

BIG-IP[®] Link Controller[™]: Monitors Reference

Version 11.6



Table of Contents

Legal Notices.....	5
Acknowledgments.....	7
Chapter 1: Monitors Concepts.....	11
Purpose of monitors.....	11
Benefits of monitors.....	11
About iCheck functionality for monitors.....	12
Methods of monitoring.....	12
Comparison of monitoring methods.....	13
Monitor destinations.....	13
About monitor settings.....	14
Transparent and Reverse modes.....	14
Monitors that contain the Transparent or Reverse settings.....	15
The Manual Resume feature.....	15
Resumption of connections.....	16
The Time Until Up feature.....	16
About health and performance monitors.....	16
About address check monitors.....	17
About application check monitors.....	17
About content check monitors.....	17
About path check monitors.....	17
About performance check monitors.....	18
About service check monitors.....	18
About resources and monitor queries.....	18
About the Virtual Location monitor.....	18
Overview of monitor implementation.....	19
Preconfigured monitors.....	19
Custom monitors.....	20
Monitor instances.....	21
Chapter 2: Monitors Tasks.....	23
Creating an SNMP monitor.....	23
Creating a custom monitor.....	23
Deleting a monitor.....	24
Disabling a monitor.....	24
Displaying a monitor.....	25
Enabling a monitor.....	25
Creating an HTTP monitor.....	26
Creating an HTTPS monitor.....	27

Chapter 3: Monitors Settings Reference	29
Health monitor functional categories.....	30
Performance monitor functional category.....	35
BIG-IP monitor settings.....	36
BIG-IP Link monitor settings.....	37
External monitor settings.....	38
FirePass monitor settings.....	39
FTP monitor settings.....	40
Gateway ICMP monitor settings.....	42
HTTP monitor settings.....	43
HTTPS monitor settings.....	46
IMAP monitor settings.....	49
LDAP monitor settings.....	50
MSSQL monitor settings.....	52
MySQL monitor settings.....	54
NNTP monitor settings.....	56
Oracle monitor settings.....	57
POP3 monitor settings.....	59
PostgreSQL monitor settings.....	60
RADIUS monitor settings.....	62
RADIUS Accounting monitor settings.....	63
Real Server monitor settings.....	64
Scripted monitor settings.....	65
SIP monitor settings.....	67
SMTP monitor settings.....	68
SNMP monitor settings.....	69
SNMP Link monitor settings.....	71
SOAP monitor settings.....	72
TCP monitor settings.....	73
TCP Half Open monitor settings.....	75
UDP monitor settings.....	76
WAP monitor settings.....	79
WMI monitor settings.....	80

Legal Notices

Publication Date

This document was published on December 8, 2015.

Publication Number

MAN-0370-04

Copyright

Copyright © 2014-2015, F5 Networks, Inc. All rights reserved.

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

Trademarks

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

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

Patents

This product may be protected by one or more patents indicated at:
<http://www.f5.com/about/guidelines-policies/patents>

Link Controller Availability

This product is not currently available in the United States.

Export Regulation Notice

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

RF Interference Warning

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

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

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

Canadian Regulatory Compliance

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

Standards Compliance

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

Acknowledgments

This product includes software developed by Gabriel Forté.

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

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

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

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

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

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

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

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

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

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

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

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

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

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

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

This product includes software developed by Christos Zoulas.

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

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

Acknowledgments

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

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

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

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

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

This product includes software developed by Jared Minch.

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

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

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

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

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

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

This product includes Hypersonic SQL.

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

This product includes software developed by the Internet Software Consortium.

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

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

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

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

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

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

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

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

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes jxrlib software, copyright ©2009 Microsoft Corp. All rights reserved. Distributed under the new BSD license.

This product includes node-uuid software, copyright © 2010-2012, Robert Kieffer, and distributed under the MIT license.

This product includes opensv software, which is distributed under the Apache 2.0 license.

This product includes owasp-jave-encoder software, copyright © 2014, Jeff Ichnowski, and distributed under the New BSD license.

Chapter

1

Monitors Concepts

- *Purpose of monitors*
- *Benefits of monitors*
- *Methods of monitoring*
- *Monitor destinations*
- *About monitor settings*
- *Transparent and Reverse modes*
- *The Manual Resume feature*
- *The Time Until Up feature*
- *About health and performance monitors*
- *About address check monitors*
- *About application check monitors*
- *About content check monitors*
- *About path check monitors*
- *About performance check monitors*
- *About service check monitors*
- *About resources and monitor queries*
- *About the Virtual Location monitor*
- *Overview of monitor implementation*
- *Monitor instances*

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 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™, Global Traffic Manager™, 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 Global Traffic Manager 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	<ul style="list-style-type: none"> Works well when you only need to determine the up or down status of a node. 	<ul style="list-style-type: none"> Can check the health of a node only, and not a pool member.
Active	<ul style="list-style-type: none"> Can check for specific responses Can run with or without client traffic 	<ul style="list-style-type: none"> Creates additional network traffic beyond the client request and server response Can be slow to mark a pool member as down
Passive	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> 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 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
Gateway ICMP	Transparent
TCP Half Open	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 Global Traffic Manager™: 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 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 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.

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.

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.

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.

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.

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.

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.

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_head_f5`
- `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 30, the Timeout value is 120, 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 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 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.)

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.

Chapter 2

Monitors Tasks

- *Creating an SNMP monitor*
- *Creating a custom monitor*
- *Deleting a monitor*
- *Disabling a monitor*
- *Displaying a monitor*
- *Enabling a monitor*
- *Creating an HTTP monitor*
- *Creating an HTTPS monitor*

Creating an SNMP monitor

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

1. Click **Create**.
The New Monitor screen opens.
2. Type a name for the monitor.

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

3. 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 **Link Controller > Monitors**.

The Monitor List screen opens.

2. Click **Create**.
The New Monitor screen opens.
3. Type a name for the monitor in the **Name** field.
4. From the **Type** list, select 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 Configuration list, select **Advanced**.
This selection makes it possible for you to modify additional default settings.
7. Configure all settings shown.
8. 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 **Link Controller > Monitors**.
The Monitor List screen opens.
2. Select the **Select** check box for the monitor that you want to delete.
3. Click **Delete**.
A confirmation message appears.
4. Click **Delete**.

The monitor is deleted.

Disabling a monitor

You can disable a monitor to discontinue monitoring a server.

***Note:** 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.*

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

2. Click a monitor name in the list.
The monitor settings and values appear.
3. Click **Instances** on the menu bar.
Any existing monitor instances appear.
4. Select the **Select** check box for the instance you want to manage.
5. Click **Disable**.
6. Click **Update**.

The monitor is disabled and no longer monitoring the server.

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 **Link Controller > Monitors**.
The Monitor List screen opens.
2. Click a monitor name in the list.
The monitor settings and values appear.

You can view the settings and values for the monitor.

Enabling a monitor

You can enable a monitor to begin or resume monitoring a server.

***Note:** 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.*

1. On the Main tab, click **Link Controller > Monitors**.
The Monitor List screen opens.
2. Click a monitor name in the list.
The monitor settings and values appear.
3. Click **Instances** on the menu bar.
Any existing monitor instances appear.
4. Select the **Select** check box for the instance you want to manage.
5. Click **Enable**.
6. Click **Update**.

The monitor is enabled to begin or resume monitoring a server.

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.

1. On the Main tab, click **Link Controller > Monitors**.
The Monitor List screen opens.
2. Type a name for the monitor in the **Name** field.
3. From the **Type** list, select **HTTP**.
The screen refreshes, and displays the configuration options for the **HTTP** monitor type.
4. From the **Parent Monitor** list, select **http**.
The new monitor inherits initial configuration values from the existing monitor.
5. From the Configuration list, select **Advanced**.
This selection makes it possible for you to modify additional default settings.
6. Type a number in the **Interval** field that indicates, in seconds, how frequently the system issues the monitor check. The default is 30 seconds.
7. Type a number in the **Timeout** field 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.
8. Type a regular expression in the **Receive String** field 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.

9. Type a name in the **User Name** field.
10. Type a password in the **Password** field.
11. For the **Reverse** setting, do one of the following:
 - Accept the **No** default option.
 - Select the **Yes** option to make the **Receive Disable String** option unavailable and mark the pool, pool member, or node **Down** when the test is successful.
12. For the **Transparent** setting, do one of the following:
 - Accept the **No** default option.
 - Select the **Yes** option to use a path through the associated pool members or nodes to monitor the aliased destination.

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 **Link Controller > Monitors**.
The Monitor List screen opens.
2. From the **Type** list, select the type of monitor.
The screen refreshes, and displays the configuration options for the monitor type.
3. From the **Import Monitor** list, select an existing monitor.
The new monitor inherits initial configuration values from the existing monitor.
4. Type a number in the **Interval** field that indicates, in seconds, how frequently the system issues the monitor check. The default is 30 seconds.
5. Type a number in the **Timeout** field 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.
6. Type a number in the **Probe Timeout** field that indicates the number of seconds after which the system times out the probe request to the system. The default is 5 seconds.
7. 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.
8. Type a text string in the **Send String** field 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.
9. Type a regular expression in the **Receive String** field 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.

10. Type a list of ciphers in the **Cipher List** field that match those of the client sending a request, or of the server sending a response.
The default string is `DEFAULT:+SHA:+3DES:+kEDH`.
11. Type a name in the **User Name** field.
12. Type a password in the **Password** field.
13. From the **Client Certificate** list, do one of the following:
 - Accept the default, **None**, to specify no client certificate.
 - Select **ca-bundle** to use the ca-bundle client certificate.
 - Select **default** to use a default client certificate.
14. For the **Reverse** setting, do one of the following:

- Accept the **No** default option.
- Select the **Yes** option to make the **Receive Disable String** option unavailable and mark the pool, pool member, or node **Down** when the test is successful.

15. For the **Transparent** setting, do one of the following:

- Accept the **No** default option.
- Select the **Yes** option to use a path through the associated pool members or nodes to monitor the aliased destination.

16. For the **Alias Address** setting, do one of the following:

- Accept 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.

17. For the **Alias Service Port** setting, do one of the following:

- 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.

Associate the HTTPS monitor with a server, pool, pool member, or node.

Chapter

3

Monitors Settings Reference

- *Health monitor functional categories*
- *Performance monitor functional category*
- *BIG-IP monitor settings*
- *BIG-IP Link monitor settings*
- *External monitor settings*
- *FirePass monitor settings*
- *FTP monitor settings*
- *Gateway ICMP monitor settings*
- *HTTP monitor settings*
- *HTTPS monitor settings*
- *IMAP monitor settings*
- *LDAP monitor settings*
- *MSSQL monitor settings*
- *MySQL monitor settings*
- *NNTP monitor settings*
- *Oracle monitor settings*
- *POP3 monitor settings*
- *PostgreSQL monitor settings*
- *RADIUS monitor settings*
- *RADIUS Accounting monitor settings*
- *Real Server monitor settings*
- *Scripted monitor settings*
- *SIP monitor settings*
- *SMTP monitor settings*
- *SNMP monitor settings*
- *SNMP Link monitor settings*
- *SOAP monitor settings*
- *TCP monitor settings*
- *TCP Half Open monitor settings*
- *UDP monitor settings*
- *WAP monitor settings*
- *WMI monitor settings*

Health monitor functional categories

The following 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™, Global Traffic Manager™, 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 servers 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 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).

Service-check monitor	Description
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, <code>alt.cars.mercedes</code>) 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.
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 <code>rpcinfo</code> command to query the RPC server and verify the availability of a given program.
SASP	Verifies the availability of a IBM® Group Workload Manager. This monitor uses the Server/Application State Protocol (SASP) to communicate with the Group Workload Manager. The monitor queries the Group Workload Manager 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 Group Workload Manager (GWM), it configures the dynamic ratio option for

Service-check monitor	Description
	<p>the resources, allowing the BIG-IP system to select the most appropriate resource to respond to a connection request.</p> <hr/> <p><i>Note:</i> When you assign an SASP monitor, the monitor initially marks the resources as down. This change in status occurs because the GWM 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.</p> <hr/>
SIP	Checks the status of SIP Call-ID services. By default, this monitor type issues an <code>SIP OPTIONS</code> 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 <code>TCP SYN</code> packet to the service. As soon as the monitor receives the <code>SYN-ACK</code> 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
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.
HTTPS	Checks the status of Hypertext Transfer Protocol Secure (HTTPS) traffic. An HTTPS monitor attempts to receive specific content from a web page protected

Content-check monitor	Description
https_443	by SSL security. The check is successful when the content matches the Receive String value.
LDAP	Checks the status of Hypertext Transfer Protocol Secure (HTTPS) traffic, by using port 443.
Scripted	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.
SMTP	Generates a simple script that reads a file that you create. The file contains <code>send</code> and <code>expect</code> strings to specify lines that you want to send or that you expect to receive.
TCP	Checks the status of Simple Mail Transport Protocol (SMTP) servers. This monitor type checks only that the server is up and responding to commands. The check is successful if the mail server responds to the standard <code>SMTP HELO</code> and <code>QUIT</code> commands.
WAP	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 Receive String value.
	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 <code>ICMP_ECHO</code> datagram.
ICMP	Makes a simple node check. The check is successful if the monitor receives a response to an <code>ICMP_ECHO</code> 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 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 Local Traffic Manager, F5 Networks recommends that you first assign the appropriate monitor to the resource through the Local Traffic Manager, and then assign a BIG-IP monitor to the Local Traffic Manager through the Global Traffic Manager. 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 the Global Traffic Manager 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 <code>/var/tmp</code> 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	<p>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 <code>gtm_score</code> values from the downstream virtual servers and set the dynamic ratios on the associated upstream local traffic management pool members or nodes.</p> <p>The Module Score monitor retrieves the <code>gtm_score</code> values from the virtual server and the</p>

Application-check monitor	Description
<p data-bbox="350 554 545 583">Virtual Location</p>	<p data-bbox="917 201 1451 323">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.</p> <p data-bbox="917 344 1468 537">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.</p> <p data-bbox="917 558 1468 968">Optimizes end-user response time in environments with dynamic distribution of application resources across multiple data centers. When using the Virtual Location monitor, the BIG-IP 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.</p>

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™, Global Traffic Manager™, 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
<p data-bbox="350 1495 441 1524">BIG-IP</p>	<p data-bbox="917 1495 1468 1617">Collects data from Global Traffic Manager and Local Traffic Manager. Typically, the Local Traffic Manager probes local pool members and provides the results to Global Traffic Manager.</p> <hr/> <p data-bbox="917 1650 1451 1772"><i>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.</i></p>
<p data-bbox="350 1814 506 1843">BIG-IP Link</p>	<p data-bbox="917 1814 1468 1908">Gathers metrics and statistics information acquired through the monitoring of Global Traffic Manager or Link Controller resources.</p>

Performance monitor	Description
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 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. <i>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.</i>

BIG-IP monitor settings

This table describes the BIG-IP 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 Settings	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
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.
Ignore Down Response	No	Specifies that the monitor allows more than one probe attempt per interval. 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.
Aggregate Dynamic Ratios	None	Specifies how the system combines the module values to create the proportion (score) for the load balancing operation. The score represents the module's estimated capacity for handling traffic. Averaged values are appropriate for downstream Web Accelerator or Application Security Manager virtual servers. The default is None , meaning that the system does not use the scores in the load balancing operation.

BIG-IP Link monitor settings

This table describes the BIG-IP Link 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 Settings	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.
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

Setting	Value	Description
		up. If the target does not respond within the set time period, the target is considered to be down.
Ignore Down Response	No	Specifies that the monitor allows more than one probe attempt per interval. 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.

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 Settings	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 .
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 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.

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.
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 Settings	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	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.
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

Setting	Value	Description
		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 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. <i>Important: If there is no password security, you must use blank strings [""] for the Username and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 Settings	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.
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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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	<ul style="list-style-type: none"> • 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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> file.

Gateway ICMP monitor settings

This table describes the Gateway 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.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import Settings	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 .
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
Adaptive	Disabled	<p>Specifies whether adaptive response time monitoring is enabled for this monitor.</p> <p>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.</p> <p>Disabled The monitor determines the state of a service based on the Interval, Up Interval, Time Until Up, and Timeout monitor settings.</p>
Allowed Divergence	Relative, 25%	<p>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:</p> <p>Absolute The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed.</p> <p>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.</p>
Adaptive Limit	200 milliseconds	<p>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.</p>
Sampling Timespan	300 seconds (5 minutes)	<p>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.</p>

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.

Setting	Value	Description
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import Settings	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.
Send String	GET /	<p>Specifies the text string that the monitor sends to the target object. You must include <code>\r\n</code> at the end of a non-empty send string. The default setting is <code>GET /\r\n</code>, 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: <code>GET /www/siterequest/index.html\r\n</code>.</p> <hr/> <p>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 (<code>\r\n</code>) at the end of a non-empty send string, for example <code>GET /\r\n</code> instead of <code>GET /</code>. If you do not include <code>\r\n</code> 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 (<code>\r\n\r\n</code>), for example, <code>"GET / HTTP/1.1\r\nHost: server.com\r\n\r\n"</code> or <code>"GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close\r\n\r\n"</code>.</p> <hr/>
Receive String	No default	<p>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.</p> <hr/> <p>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</p> <hr/>
User Name	No default	<p>Specifies the user name, if the monitored target requires authentication.</p> <hr/> <p>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</p> <hr/>

Setting	Value	Description
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 <code>ERROR</code> . 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 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.

Setting	Value	Description
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 Settings	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 .
Send String	GET /	Specifies the text string that the monitor sends to the target object. You must include <code>\r\n</code> at the end of a non-empty send string. The default setting is <code>GET /\r\n</code> , 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: <code>GET /www/siterequest/index.html\r\n</code> .

Setting	Value	Description
		<p>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 (<code>\r\n</code>) at the end of a non-empty send string, for example <code>GET /\r\n</code> instead of <code>GET /</code>. If you do not include <code>\r\n</code> 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 (<code>\r\n\r\n</code>), for example, <code>"GET / HTTP/1.1\r\nHost: server.com\r\n\r\n"</code> or <code>"GET / HTTP/1.1\r\nHost: server.com\r\nConnection: close\r\n\r\n"</code>.</p>
Receive String	No default	<p>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.</p> <p>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</p>
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	<p>Specifies the user name, if the monitored target requires authentication.</p> <p>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</p>
Password	No default	<p>Specifies the password, if the monitored target requires authentication.</p> <p>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</p>
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,

Setting	Value	Description
		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 <code>Error</code> . 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	<p>Specifies whether adaptive response time monitoring is enabled for this monitor.</p> <p>Enabled</p> <p>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.</p> <p>Disabled</p> <p>The monitor determines the state of a service based on the Interval, Up Interval, Time Until Up, and Timeout monitor settings.</p>
Allowed Divergence	Relative, 25%	<p>Specifies the type of divergence used when the Adaptive setting is enabled (checkbox 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:</p> <p>Absolute</p> <p>The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed.</p>

Setting	Value	Description
		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 Settings	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.
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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> file.

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 Settings	Selected predefined or	Specifies the selected predefined or user-defined monitor.

Setting	Value	Description
	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.
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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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.

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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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.
Import Settings	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	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 .
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: <code>SELECT * FROM <db_name></code> . 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.

Setting	Value	Description
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Database	No default	Specifies the name of the database that the monitor tries to access, for example, <code>sales</code> or <code>hr</code> .
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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. <i>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.</i>
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 . <i>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.</i>
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: <code>SELECT * FROM <db_name></code> . 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. <i>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</i>
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 this monitor. The Yes setting specifies that the system redirects error messages and additional information to the <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 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

Setting	Value	Description
		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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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. <i>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.</i>
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 .
Send String	No default	Specifies the SQL statement that the monitor runs on the target. A sample is: <code>SELECT * FROM <db_name></code> . 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

Setting	Value	Description
		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. <i>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</i>
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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

Setting	Value	Description
		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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>

Setting	Value	Description
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""], for the User Name and Password settings.</i>
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> file.

PostgreSQL monitor settings

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.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import Settings	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	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

Setting	Value	Description
		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 .
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. <i>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</i>
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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

Setting	Value	Description
		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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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.
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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>

Setting	Value	Description
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import Settings	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.
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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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

Setting	Value	Description
		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 .
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.

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

Name	Escape Sequence
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 bigipl.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 Settings	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.
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 .
File Name	No default	Specifies the name of a file in the <code>/config/eav/</code> 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 <code>*All Addresses</code> . 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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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.
Type	Selected monitor type	Specifies the type of monitor you are creating.
Import Settings	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 .
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 .

Setting	Value	Description
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 . <i>Note: The monitor always marks the target up in response to status code 200 OK.</i>
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 <code>\r\n</code> 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 <code>INVITE</code> , 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 <code>*All Addresses</code> . 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 <code>*All Ports</code> . 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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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 .
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> file.

SNMP monitor settings

This table describes the SNMP 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 Settings	Selected predefined or user-defined monitor	Specifies the selected predefined or user-defined monitor.
Interval	90	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 90 seconds.
Timeout	180	Specifies the number of seconds in which the target must respond to the monitor request. The default is 180 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.
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.
Port	161	Specifies the port number to which this monitor sends SNMP traps.
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.

SNMP Link 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 Settings	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.
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.
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 .
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 .
Port	161	Specifies the port number to which this monitor sends SNMP traps.
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.

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 Settings	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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
Password	No default	Specifies the password, if the monitored target requires authentication. <i>Important: If there is no password security, you must use blank strings [""] for the User Name and Password settings.</i>
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.

Setting	Value	Description
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 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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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
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 .
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. <i>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</i>
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 <code>Error</code> . 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.

Setting	Value	Description
		<p>You can set values for the Allowed Divergence, Adaptive Limit, and Sampling Timespan monitor settings.</p> <p>Disabled The monitor determines the state of a service based on the Interval, Up Interval, Time Until Up, and Timeout monitor settings.</p>
Allowed Divergence	Relative, 25%	<p>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:</p> <p>Absolute The number of milliseconds the latency of a monitor probe can exceed the mean latency for the service being probed.</p> <p>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.</p>
Adaptive Limit	200 milliseconds	<p>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.</p>
Sampling Timespan	300 seconds (5 minutes)	<p>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.</p>

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 Settings	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.
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.

Setting	Value	Description
Import Settings	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.
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. <i>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.</i>
Send String	default send string	Specifies the text string that the monitor sends to the target object. The default is default send 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

Setting	Value	Description
		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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 (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.

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 Settings	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.
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 .
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. <i>Note: If you do not specify both a Send String and a Receive String, the monitor performs a simple service check and connect only.</i>
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.

Setting	Value	Description
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.
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 <code>/var/log/<monitor_type>_<ip_address>.<port>.log</code> 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 Settings	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.

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 .
User Name	No default	Specifies the user name, if the monitored target requires authentication. <i>Important:</i> 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. <i>Important:</i> 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. <i>Note:</i> 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.
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.

Index

B

- BIG-IP Link monitor
and settings 37
- BIG-IP monitor
and settings 36

C

- connections
resuming 15–16
- custom monitor
creating 23

E

- External monitor
and settings 38

F

- FirePass monitor
and settings 39
- FTP monitor
and settings 40
- iCheck functionality 12

G

- Gateway ICMP monitor
and settings 42

H

- health monitors
 - about address check 17
 - about application check 17
 - about content check 17
 - about path check 17
 - about performance check 18
 - about service check 18
 - about synchronous queries 18
 - categories 30
- http monitor
creating 26
- HTTP monitor
and settings 43
- https monitor
creating 27
- HTTPS monitor
and settings 46

I

- iCheck functionality
about 12

- IMAP monitor
and settings 49

L

- LDAP monitor
and settings 50

M

- Manual Resume feature 15–16
- monitor
 - deleting 24
 - disabling 24
 - displaying 25
 - enabling 25
- monitor associations 21
- monitor destinations 13
- monitor instances 21
- monitors
 - about benefits 11
 - custom 20
 - health 16
 - methods 12
 - performance 16
 - preconfigured 19
 - purpose 11
 - types of 16, 19
 - Virtual Location 18
- monitor settings
about 14
- MSSQL monitor
and settings 52
- MySQL monitor
and settings 54

N

- NNTP monitor
and settings 56
- nodes
associating monitors with 21

O

- Oracle monitor
and settings 57

P

- performance monitors
categories 35
- pools and pool members
associating monitors with 21
- POP3 monitor
and settings 59

Index

PostgreSQL monitor
and settings *60*

R

RADIUS Accounting monitor
and settings *63*

RADIUS monitor
and settings *62*

Real Server monitor
and settings *64*

resource availability
designating *15–16*

Reverse mode *14–15*

S

Scripted monitor
and settings *65*

server availability
designating *15–16*

SIP monitor
and settings *67*

SMTP monitor
and settings *68*

SNMP DCA Base monitor
and settings *71*

SNMP monitor
and settings *69*

SNMP monitoring
creating monitors *23*

SOAP monitor
and settings *72*

T

TCP Half Open monitor
and settings *75*

TCP monitor
and settings *73*

Time Until Up feature *16*

Transparent mode *14–15*

U

UDP monitor
and settings *76*

V

Virtual Location monitor
about *18*

W

WAP monitor
and settings *79*

WMI monitor
and settings *80*