Monitoring F5 BIG-IP LTM Analytics

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Analytics (also called Application Visibility and Reporting) is a module on the BIG-IP® system that can be used to analyze the performance of web applications. It provides detailed metrics such as transactions per second, server and client latency, request and response throughput, and sessions. You can view metrics for applications, virtual servers, pool members, URLs, specific countries, and additional detailed statistics about application traffic running through the BIG-IP system.

Transaction counters for response codes, user agents, HTTP methods, countries, and IP addresses provide statistical analysis of the traffic that is going through the system. You can capture traffic for examination and have the system send alerts so you can troubleshoot problems and immediately react to sudden changes.

An Analytics profile is a set of definitions that determines the circumstances under which the system gathers, logs, notifies, and graphically displays information regarding traffic to an application. The Analytics module requires that you select an Analytics profile for each application you want to monitor. The Analytics profile is associated with one or more virtual servers used by the application or with an iApps application service. The BIG-IP® system includes a default Analytics profile called *analytics*. It is a minimal profile that internally logs application statistics for server latency, throughput, response codes, and methods. If the analytics profile is enabled on the F5 traffic manager, then, administrators can figure out the traffic flow between each application load balanced by the F5 traffic manager and the users. These users may access the applications either through a launched application, or a client, or from a client subnet etc. To continuously monitor the traffic flow and figure out the user experience on the applications load balanced by the F5 traffic manager, it is imperative to monitor the F5 traffic manager with the analytics module enabled on it. the forthcoming chapter explains in details on how to monitor the Analytics module of the F5 traffic manager.

Chapter 2: How to Monitor F5 BIG-IP LTM Analytics Using eG Enterprise?

eG Enterprise monitors the F5 BIG-IP LTM Analytics using an eG external agent that is deployed on any remote host. To enable the eG agent to monitor the F5 Analytics of the Big IP-Local Traffic Manager, the *analytics* profile should be enabled on the Big IP-Local Traffic Manager. The section that follows describes how to manage the F5 BIG-IP Load Balancer. Once the above-said requirement is fulfilled, manage the F5 BIG-IP LTM Analytics component for monitoring. The procedure for achieving this is explained in the below section.

2.1 Managing the F5 BIG-IP LTM Analytics

eG Enterprise cannot automatically discover the F5 Big IP-Local Traffic Manager Analytics. This implies that you need to manually add the component for monitoring. To manage a F5 BIG-IP LTM Analytics component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENTS** page that appears next, select *F5 BIG-IP LTM Analytics* as the **Component type**. Then, click the **Add New Component** button. This will invoke Figure 2.1.

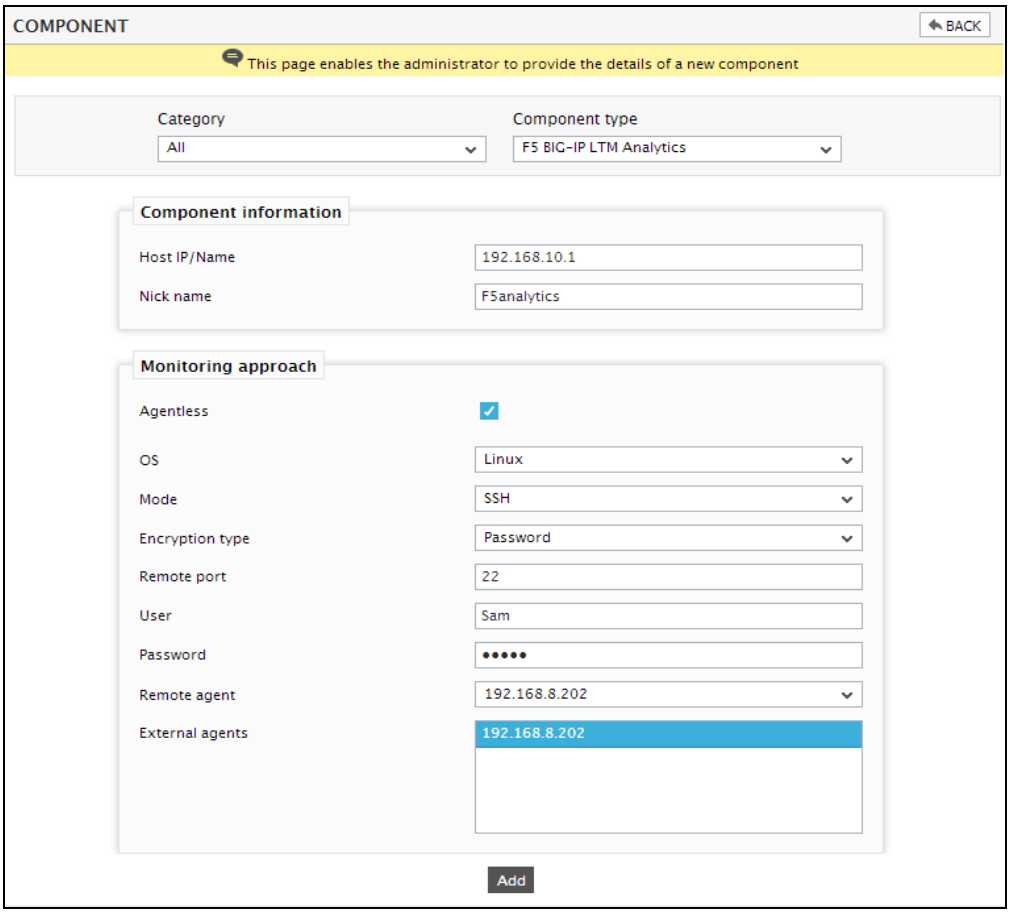

Figure 2.1: Adding a F5 BIG-IP LTM Analytics component

4. Specify the **Host IP/Name** and **Nick name**for the F5 BIG-IP LTM Analytics component. By default, the F5 BIG-IP LTM Analytics component is monitored in an agentless manner. Therefore, the **Agentless** flag will be checked by default.
5. Next, choose **Linux** as the **OS** and **SSH** as the **Mode** for monitoring the F5 BIG-IP LTM Analytics component. Then, choose **Password** as the **Encryption Type**. The **Remote Port** will be 22 by default.
6. Now, specify the credentials of a user who can access the putty in the **User** and **Password** text boxes. Then, select the Remote and External agents.
7. Finally, click **Add** button to register the changes.
8. When you attempt to sign out, a list of unconfigured tests appears (see Figure 2.2). These tests appear only when you want to monitor the F5 BIG-IP LTM Analytics component.

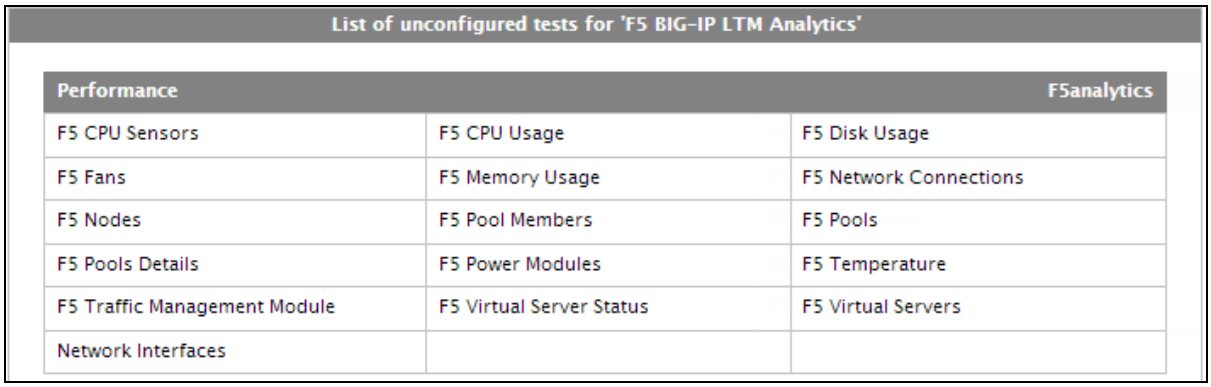

List of unconfigured tests for 'F5 BIG-IP LTM Analytics'

Performance		F5analytics
F5 CPU Sensors	F5 CPU Usage	F5 Disk Usage
F5 Fans	F5 Memory Usage	F5 Network Connections
F5 Nodes	F5 Pool Members	F5 Pools
F5 Pools Details	F5 Power Modules	F5 Temperature
F5 Traffic Management Module	F5 Virtual Server Status	F5 Virtual Servers
Network Interfaces		

Figure 2.2: A list of unconfigured tests of F5 BIG-IP LTM Analytics component

9. Click on any test to configure it. Tho know how to configure these tests, the **Monitoring the F5 BIG-IP LTM Analytics**.
10. Finally, sign out of the eG administrative interface.

Chapter 3: Monitoring the F5 BIG-IP LTM Analytics

Since application delivery delays, inefficiencies in traffic flow, and failures can cause prolonged service outages, and affect the user experience with the application thus costing an enterprise money and reputation, it is imperative to monitor the Analytics module of the F5 traffic manager. To ensure this, eG Enterprise provides a specialized *F5 Analytics* monitoring model.

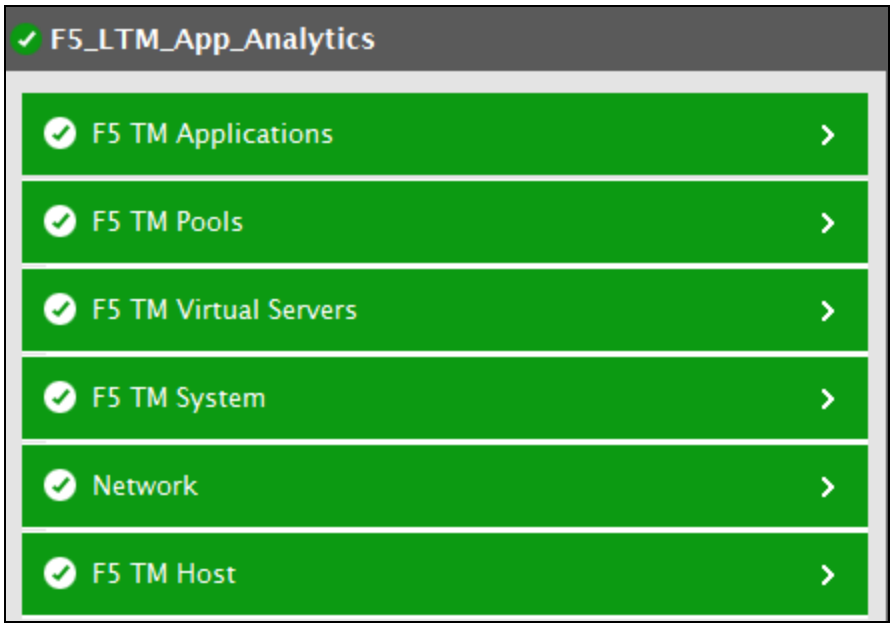

Figure 3.1: Layer model of the F5 BIG-IP LTM Analytics

Each layer of Figure 3.1 is mapped to a variety of tests that provide valuable insights into the overall performance of the F5 Traffic Manager when the *Analytics* module is enabled. With the help of the metrics reported by these tests, administrators can find quick and accurate answers for the following queries:

- How well the I/O operations were performed on each disk of the target F5 Traffic manager?
- What is the average read latency of each disk of the target F5 Traffic manager?
- What is the average percentage of data serviced by the RAM cache?
- What is the percentage of data removed from the RAM cache?
- How many connections are active to the traffic manager?
- How many new server connections are available to the traffic manager?
- What was the average rate at which SSL transactions were made to the traffic manager?

- What was the average rate at which packets were transmitted to and from the traffic manager?
- What was the response time of the applications that were load balanced by the traffic manager when user requests were made through each predefined client subnet?
- How many transactions were processed by the applications that were accessed through each client subnet?
- What is the average page load time of the applications when accessed through each application launched by the user?
- What was the response time of the applications that were load balanced by the traffic manager when user requests were made through each application?
- How many transactions were processed by the applications that were accessed by the user through each application?
- What is the average page load time of the applications when accessed by the user through each application?
- What was the response time of the applications that were load balanced by the traffic manager when user requests were made through each URL?
- How many transactions were processed by the applications that were accessed by the user through each URL?
- What is the average page load time of the applications when accessed by the user through each URL?
- What was the response time of the applications that were load balanced by the traffic manager when user requests were made through each application pool member?
- How many transactions were processed by the applications that were accessed by the user through each application pool member?
- What is the average page load time of the applications when accessed by the user through each application pool member?

Since the tests pertaining to the *F5 TM Pools* layer and the *F5 TM Virtual Servers* layer have already been discussed in the *Monitoring F5 Traffic Managers* document and the *Network* layer has already been discussed in the *Monitoring Unix and Windows Servers* document, the remaining layers are discussed in detail in this document.

3.1 The F5 TM Host Layer

By monitoring the critical hardware components such as CPUs and the disks of the traffic manager, administrators can proactively be alerted to potential I/O latencies, atencies in read/write operations and abnormalities that crop up in the RAM cache.

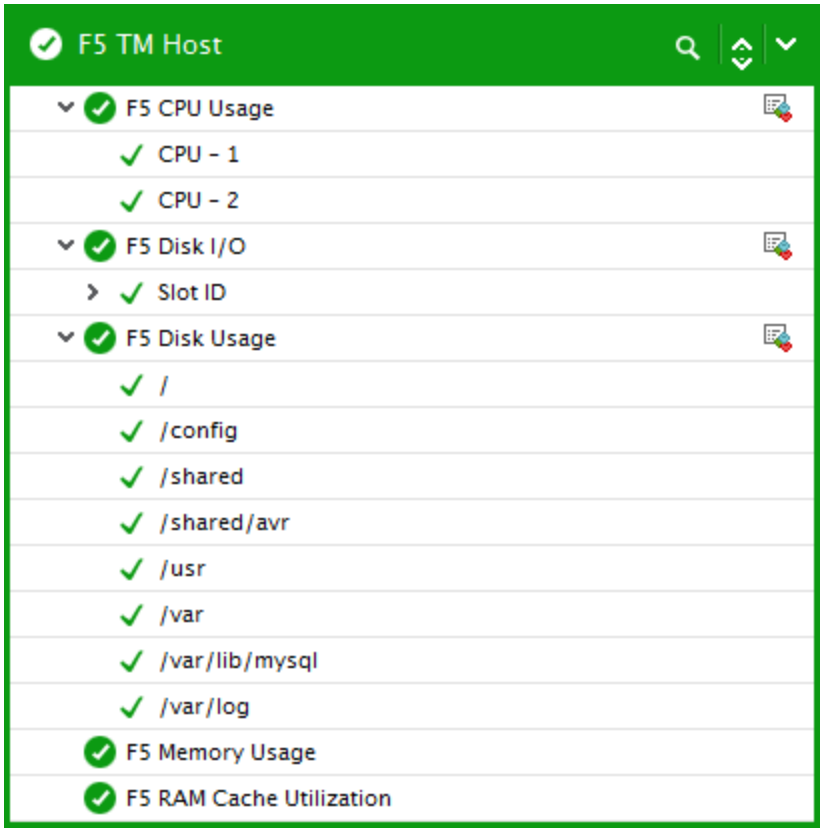

Figure 3.2: The tests pertaining to the F5 TM Host layer

Since most of the tests of this layer have already been discussed in the *Monitoring F5 Traffic Managers* document, let us now discuss the *F5 Disk I/O* and *F5 RAM Cahe Utilization* tests alone in the forthcoming sections.

3.1.1 F5 Disk I/O Test

This test auto-discovers the disks of the F5 traffic manager and reports how well the I/O operations were performed on each disk and in the process, identify delays in I/O operations. In addition, this test also reports the total amount of data transacted when performing read and write operations. This way, this test not only enables administrators to promptly detect current or probable latencies in I/O operations on each disk, but will also help them figure out which I/O activity was most latent – i.e., whether read operations or write operations?

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each disk inserted in the slot of the target F5 traffic manager

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total I/O Operations	Indicates the total number of read and write operations performed on this disk.	Operations	A high value is desired for this measure. A consistent decrease in this value could indicate a processing bottleneck.
Read operations	Indicates the number of read operations performed on this disk.	Operations	If the value of the *Total I/O Operations* measure dips consistently, then, you may want to time-correlate that measure with the value of these two measures to know what is causing the steady decline in disk performance – read operations or write operations? This way, you can figure out when the slowdown actually occurred – when reading or when writing?
Write operations	Indicates the number of write operations performed on this disk.	Operations	
Merged read operations	Indicates the number of read operations that were merged on this disk.	Operations	
Merged write operations	Indicates the number of write operations that were	Operations	

Measurement	Description	Measurement Unit	Interpretation
	merged on this disk.		
Read data	Indicates the amount of data that has been read from this disk.	MB	Ideally, the value of these measures should be high. A consistent decrease in their value is an indication of a processing bottleneck. When users complaint of delays when accessing the disk, you can compare the value of each of these measures across the disks to know which disk is experiencing the slowdown. You can then compare the value of these measures for that disk to know when the slowdown occurred – when reading from or when writing to the disk?
Write data	Indicates the amount of data that has been written to this disk.	MB	
Average read latency	Indicates the average time taken to read the data from this disk.	Milliseconds	Compare the value of this measure across the disks to which disk is catering the read requests at a faster pace.

3.1.2 F5 RAM cache Utilization Test

A RAM cache is a cache of HTTP objects stored in the BIG-IP systems random-access memory (RAM) that subsequent connections can reuse to reduce the amount of load on the backend servers. The RAM cache feature provides the ability to reduce the traffic load to backend servers. This ability is useful if an object on a site is under high demand, if the site has a large quantity of static content, or if the objects on the site are compressed. For compressible data, the RAM cache can store data for clients that can accept compressed data. When used in conjunction with the compression feature on the BIG-IP system, the RAM cache takes stress off the BIG-IP system and the content servers. If the RAM cache is not adequately sized or is not able to cater to the requests, then the backend servers may be overloaded with requests defying the very purpose of the cache. Therefore, it is necessary to constantly keep track of the cache. The **F5 RAM cache Utilization** test helps administrators in this regard!

Using this test, administrators can figure out the average percentage of data that was serviced by the RAM cache, the percentage of bytes that were serviced from the RAM cache etc. In addition, this

test also helps administrators identify the percentage of data that was removed from the RAM cache.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target F5 Local Traffic Manager being monitored

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Cache hit ratio	Indicates the percentage of data that was serviced from the RAM cache.	Percent	Ideally, the value of this measure should be high. A low value indicates that many read requests are serviced by direct disk accesses, which is a more expensive operation in terms of processing overheads.
Average cache hit ratio	Indicates the average percentage of data that was serviced from the RAM cache.	Percent	A high value is desired for this measure.
Cache byte ratio	Indicates the percentage of bytes that were serviced from the RAM cache.	Percent	
Average byte ratio	Indicates the average percentage of bytes that were cached in the RAM cache.	Percent	
Cache eviction ratio	Indicates the percentage of data that was removed	Percent	A low value is desired for this measure. A higher value of this

Measurement	Description	Measurement Unit	Interpretation
	from the RAM cache.		measure may be an indication of outdated entries in the cache which should be evicted from the cache to accommodate new cache entries.
Average Cache eviction rate	Indicates the average percentage of data that was removed from the RAM cache.	Percent	

3.2 The F5 TM System Layer

Using this layer, administrators can determine the number of active client connections and server connections to the traffic manager as well as the throughput of the applications that were load balanced by the traffic manager.

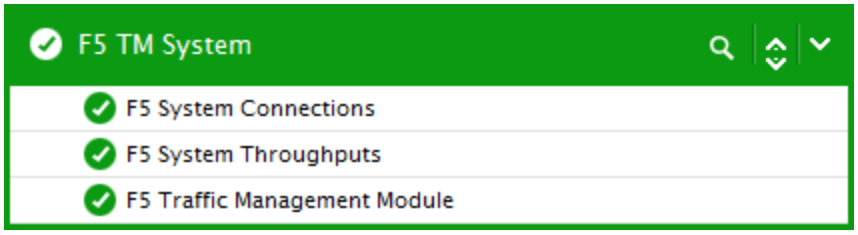

Figure 3.3: The tests pertaining to the F5 TM System layer

Since the *F5 Traffic Management Module* test has already been discussed in the *Monitoring F5 Traffic Managers* document, the remaining tests are discussed in the forthcoming sections.

3.2.1 F5 System Connections Test

This test reports the active client connections and server connections to the target F5 traffic manager. In addition, this test reports the number of active connections to the target F5 traffic manager. By analyzing the average client connections to the F5 traffic manager, administrators can detect sudden surge in the workload of the traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target traffic manager that is being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Active connections	Indicates the total number of active connections to the traffic manager.	Number	
Average active connections	Indicates the average number of active connections to the traffic manager.	Number	
New client connections	Indicates the number of new client connections to the traffic manager.	Number	A high value for this measure is a good indicator on the workload of the traffic manager.
Average new client connections	Indicates the average number of new client connections to the traffic manager.	Number	
New server connections	Indicates the number of new server connections to the traffic manager.	Number	
Average new server connections	Indicates the average number of new server connections to the traffic manager.	Number	

3.2.2 F5 System Throughputs Test

This test tracks the throughput of the applications that were load balanced by the target traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the target traffic manager that is being monitored

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Service data rate	Indicates the rate at which data through the traffic manager was serviced.	Kbps	
Average service data rate	Indicates the average rate at which data through the traffic manager was serviced.	Kbps	
Data received rate	Indicates rate at which data was received by the traffic manager.	Kbps	
Average data received rate	Indicates the average rate at which data was received by the traffic manager.	Kbps	
Data transmitted rate	Indicates the rate at which data was transmitted from the traffic manager.	Kbps	
Average data transmitted rate	Indicates the average rate at which data was transmitted from the traffic manager.	Kbps	
Service packets rate	Indicates the rate at which packets were serviced	Packets/sec	

Measurements	Description	Measurement Unit	Interpretation
	through the traffic manager.		
Average service packets rate	Indicates the average rate at which packets were serviced through the traffic manager.	Packets/sec	
Packets received rate	Indicates the rate at which packets were received by the traffic manager.	Packets/sec	
Average packets received rate	Indicates the average rate at which packets were received by the traffic manager.	Packets/sec	
Packtes transmitted rate	Indicates the rate at which packets were transmitted from the traffic manager.	Packets/sec	
Average packets transmitted rate	Indicates the average rate at which packets were transmitted from the traffic manager.	Packets/sec	
HTTP requests rate	Indicates the rate at which HTTP requests were received by the traffic manager.	Requests/sec	
Average http requests rate	Indicates the average rate at which HTTP requests were received by the traffic manager.	Requests/sec	
SSL transaction rate	Indicates the rate at which SSL transactions were made through the traffic manager.	Transactions/sec	
Average SSL transaction rate	Indicates the average rate at which SSL transactions were made through the traffic manager.	Transactions/sec	

3.3 The F5 TM Applications Layer

Using this layer, administrators can figure out the performance of the web applications that are load balanced by the traffic manager by analyzing the traffic flow when accessed from different applications, application clients, application client subnets etc.

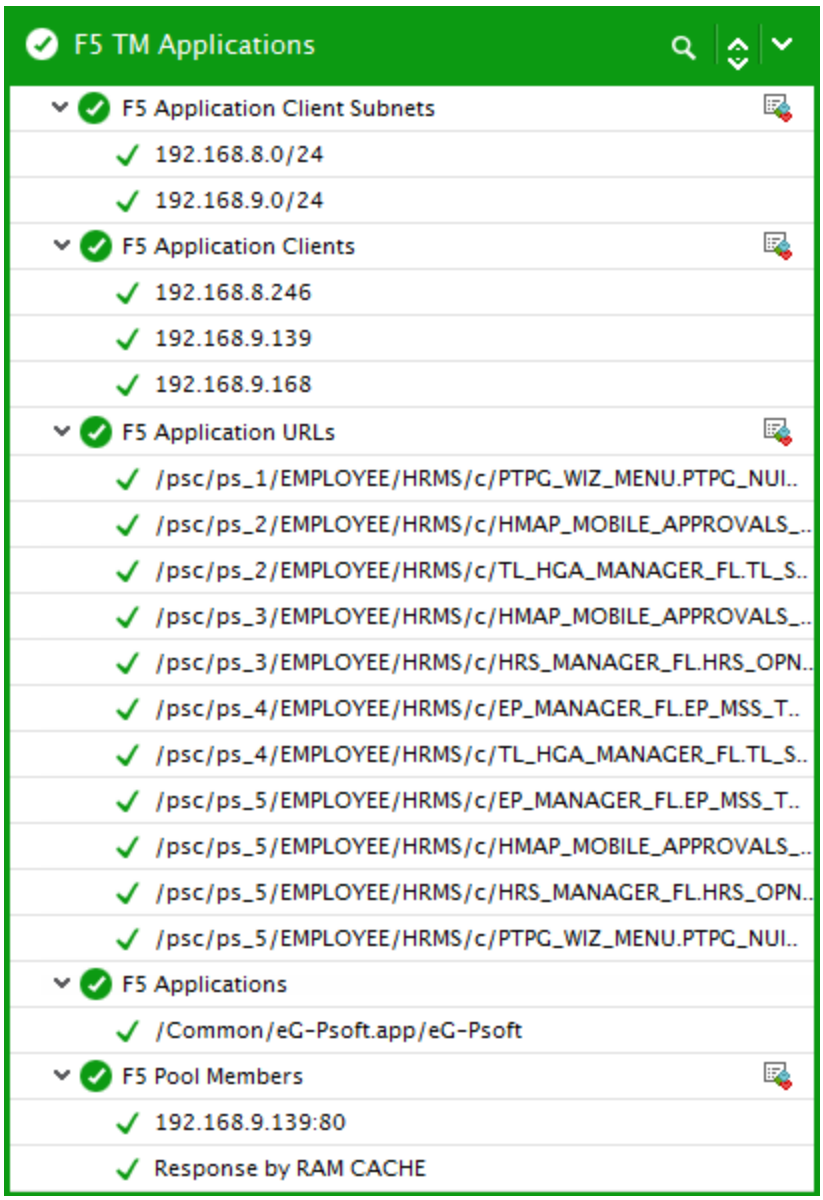

Figure 3.4: The tests pertaining to the F5 TM Applications layer

The sections below discusses each of the test pertaining to this layer in detail.

3.3.1 F5 Applications Test

For each application accessed by the end users, this test reports the time taken by the applications that are load balanced by the target traffic manager to respond to requests, reports the average load time of the application etc. Using this test, administrators can easily identify the applications that were impacting the user experience with the web applications load balanced by the traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each application that is being load balanced by the target traffic manager

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Average new session	Indicates the average number of new sessions used by this application.	Sessions	
Average concurrent session	Indicates the average number of concurrent sessions used by this application.	Sessions	
Average server latency	Indicates the average time taken by the server to process the requests for this application.	Seconds	
Average page load time	Indicates the average time taken by this application to load the pages.	Seconds	

Measurements	Description	Measurement Unit	Interpretation
Average request rate	Indicates the average rate at which requests were received from the client to this application.	Kbps	
Average response rate	Indicates the average rate at which response were received from the back end server to this application.	Kbps	
Transactions	Indicates the total number of transactions processed for this application by the F5 system.	Number	
Average transaction rate	Indicates the rate at which transactions were processed by the F5 system for this application.	Transactions/Sec	
Client sampled transaction	Indicates the total number of client sampled transactions that were processed for this application.	Number	
Maximum server latency	Indicates the maximum time taken by the server to process the requests for this application during the last measurement period.	Seconds	
Maximum page load time	Indicates the maximum time taken by this application to load the web page.	Seconds	

3.3.2 F5 Application URLs Test

For each URL through which applications are accessed by the end users, this test reports the time taken by the applications that are load balanced by the target traffic manager to respond to requests,

reports the average load time of the application etc. Using this test, administrators can easily identify URL segment groups that are adversely impacting user experience with the web application. This way, you can identify the transactions to the web application that are slow.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Application URL of the Application that is load balanced through the target traffic manager

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.
URL Grouping Count	**This test groups URLs based on the URL segments configured for monitoring and reports aggregated response time metrics for every group.** Using this parameter, you can specify a comma- separated list of URL segment numbers based on which the pages are to be grouped. URL segments are the parts of a URL (after the base URL) or path delimited by slashes. So if you had the URL: http://www.eazykart.com/web/shopping/login.jsp , then, http://www.eazykart.com will be the base URL or domain, */web* will be the first URL segment, */shopping* will be the second URL segment, and */login.jsp* will be the third URL segment. By default, this parameter is set to *3*. This default setting, when applied to the sample URL provided above, the test will report metrics for the descriptor */login.jsp*. If you set this value to *2*, then the test will report metrics for the descriptor */shopping*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Average server latency	Indicates the average time taken by the server to process the requests for this application URL.	Seconds	

Measurements	Description	Measurement Unit	Interpretation
Average page load time	Indicates the average time taken by the server to load the pages of this application URL.	Seconds	
Average request rate	Indicates the average rate at which requests were made to access this application URL.	KBps	
Average response rate	Indicates the averagse rate at which response was provided for the requests made to access this application URL.	KBps	
Transactions	Indicates the total number of transactions processed for this application URL.	Number	
Average transaction rate	Indicates the average rate at which transactions were processed for this application URL.	Transactions/Sec	
Client sampled transaction	Indicates the total number of transactions that were sampled by the client for this application URL.	Number	
Maximum server latency	Indicates the maximum time taken to process the requests for this application URL.	Seconds	
Maximum page load time	Indicates the maximum time taken by this application URL to load the page.	Seconds	

3.3.3 F5 Application Clients Test

For each client accessing the applications that are load balanced using the target traffic manager, this test reports the time taken by each client to respond to the requests, reports the average load time of the application accessed by the client etc. Using this test, administrators can easily identify the application clients that are instrumental in impacting the user experience with the web applications that are load balanced by the traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each application client through which traffic flow to the applications load balanced by the target traffic manager is monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Average server latency	Indicates the average time taken by the server to process the requests for this application client.	Seconds	
Average page load time	Indicates the average time taken by this application client to process the requests.	Seconds	
Average request rate	Indicates the average rate at which requests were made by this application client.	Kbps	

Measurements	Description	Measurement Unit	Interpretation
Average response rate	Indicates the average rate at which response was provided for the requests made by this application client.	Kbps	
Transactions	Indicates the total number of transactions processed for this application client during the last measurement period.	Number	
Average transaction rate	Indicates the average rate at which transactions were processed for this application client.	Transactions/Sec	
Client sampled transaction	Indicates the total number of transactions that were sampled for this application client during the last measurement period.	Number	
Maximum server latency	Indicates the maximum time taken to process the requests for this application client.	Seconds	
Maximum page load time	Indicates the maximum time taken by this application client to load the page.	Seconds	

3.3.4 F5 Application Client Subnets Test

A client subnet is a group of client IPs that are within the same network.

For each pre-defined client subnet, this test reports the time taken by the applications that are load balanced by the target traffic manager to respond to requests when requested through each client subnet, reports the average load time of the application accessed through each client subnet etc.

Using this test, administrators can easily identify the application client subnets that were impacting the user experience with the web applications load balanced by the traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each application client subnet through which traffic flow to the applications load balanced by the target traffic manager is monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Description	Measurement Unit	Interpretation
Average server latency	Indicates the average time taken by the server to process the requests for this client subnet.	Seconds	
Average page load time	Indicates the average time taken by this client subnet to process the requests.	Seconds	
Average request rate	Indicates the average rate at which requests were made by this client subnet.	Kbps	
Average response rate	Indicates the average rate at which response was provided for the requests made by this client subnet.	Kbps	
Transactions	Indicates the total number of transactions processed	Number	

Measurements	Description	Measurement Unit	Interpretation
	for this client subnet during the last measurement period.		
Average transaction rate	Indicates the average rate at which transactions were processed for this client subnet during the last measurement period.	Transactions/Sec	

3.3.5 F5 Application Pool Members Test

For each Application pool member, this test reports the time taken by the applications to respond to requests when requested through each application pool member, reports the average load time of the applications accessed through each application pool member etc. Using this test, administrators can easily identify the pool members that were impacting the user experience with the web applications load balanced by the traffic manager.

Target of the test : A Big-IP/F5 Local Traffic Manager

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each pool member in the target traffic manager that is to be monitored

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the traffic manager
Port	The port number on which the traffic manager is listening. By default, this is *NULL*.

Measurements made by the test

Measurements	Descriptions	Measurement Unit	Interpretation
Average server latency	Indicates the average time taken by the server to	Seconds	

Measurements	Descriptions	Measurement Unit	Interpretation
	process the requests for this pool member.		
Average page load time	Indicates the average time taken by this pool member to process the requests.	Seconds	
Average request rate	Indicates the average rate at which requests were made by this pool member.	KBps	
Average response rate	Indicates the average rate at which response was provided for the requests made by this pool member.	KBps	
Transactions	Indicates the total number of transactions processed for this pool member during the last measurement period.	Number	
Average transaction rate	Indicates the average rate at which transactions were processed for this pool member.	Transactions/Sec	
Client sampled transaction	Indicates the total number of client sampled transactions that were processed for this pool member during the last measurement period.	Number	
Maximum server latency	Indicates the maximum time taken by the server to process the requests for this pool member.	Seconds	
Maximum page load time	Indicates the maximum time taken by this application to load the page.	Seconds	

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.