Monitoring Forefront TMG

eG Enterprise v6

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows NT, Windows 2000, Windows 2003 and Windows 2008 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

Table of Contents

Table of Figures

1

Monitoring Forefront TMG

Forefront TMG is a comprehensive secure web gateway solution that helps to protect networks in an organization against web-based threats. Forefront TMG also delivers simple, unified perimeter security, with integrated firewall, VPN, intrusion prevention, malware inspection and URL filtering, thus securing the network of the target environment without degrading its performance. Even a small glitch in the performance of the Forefront TMG can expose the target environment to malicious virus attacks and unauthorized access, which may cause significant data loss. To avoid this, the availability and performance of the firewall should be continuously monitored.

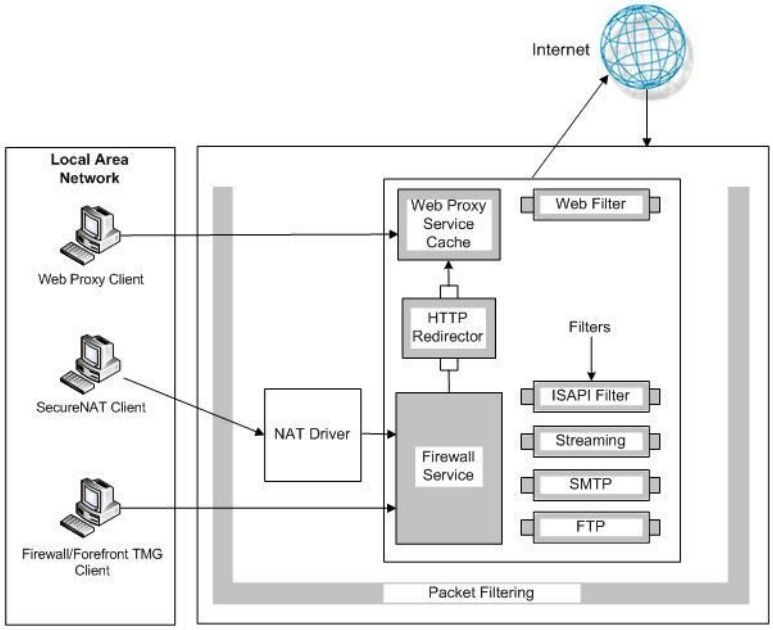

Figure 1.1: The Forefront TMG architecture

eG Enterprise offers a specialized model for monitoring the *Forefront TMG* and reporting the following key statistics that aid the administratrors to proactively detect potential firewall problems in their network.

- How well the content caching capability is utilized by Forefront TMG?
- What is the rate at which data is retrieved from the disk/memory that is allocated for content caching?
- How many times the disk has failed?
- What is the rate at which data is written to the disk?
- How many URLs were retrieved per second from the disk/memory?
- How many packets of data were sent? and how many packets were allowed through this firewall?
- Are there any packets that were backlogged and dropped?

- How many active connections are created?
- Are there any active sessions for this firewall service?
- How many active SIP registrations are available for this firewall service?
- What is the number of TCP/UDP connections made through this firewall service?
- How well the data is read/written for this firewall service?
- How many H.323 calls are being made? How efficiently the H.323 filter handles the calls?
- Is the SOCKS filter capable of handling active connections/sessions?
- How many DNS resolutions are pending/successful when going through the SOCKS filter?
- How well the Forefront TMG acts as a web proxy?
- What is the time taken to service a web proxy client request? How fast the request can be serviced?

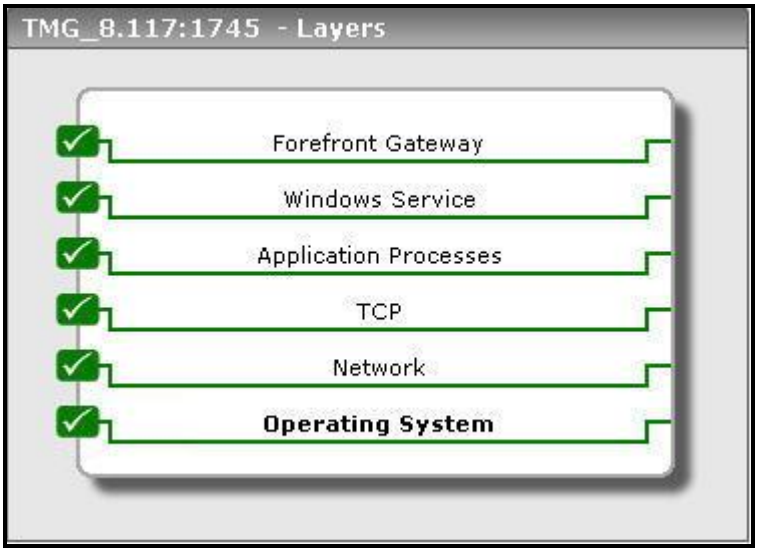

Figure 1.2: The layer model of the Forefront TMG

The **Operating System**, **Network**, **TCP**, **Application Processes** and **Windows Service** layers of the Forefront TMG are similar to that of a *Windows Generic* server model. Since these tests have been dealt with in the *Monitoring Unix and Windows Servers* document, Section 1.1 focuses on the **Forefront Gateway** layer.

1.1 The Forefront Gateway Layer

The tests mapped to this layer reports the critical performance statistics of the *Forefront TMG*.

Figure 1.3: The tests mapped to the Forefront Gateway layer

1.1.1 Forefront TMG Cache Test

Forefront Threat Management Gateway (TMG), when deployed as a web proxy server, can be configured to cache frequently requested web objects in memory and on disk in order to improve web browsing performance and to reduce bandwidth utilization. Web content caching is available for both forward and reverse proxy scenarios. Content caching brings with it different benefits in each of these deployment scenarios.

With content caching enabled, when the TMG firewall receives a web proxy request the firewall will first attempt to fulfill the request from the cache. If the requested content does not exist in the cache, it will make the request of the origin server as normal. When the web server responds, TMG will return the content to the client, and then store any cacheable content in the cache. Subsequent requests for the same content will be served directly from the cache and returned to the client at LAN speeds, eliminating the need to make a trip to the origin server to retrieve the content. This improves page loading speeds for end users and reduces bandwidth consumption on Internet links.

Therefore, if end users complain of slowness when browsing, it would be good practice to check on cache usage first, as an improperly sized cache or an ineffectively utilized cache is a key contributor to poor web browsing experience. Such imperative cache usage checks can be performed using the **Forefront TMG Cache** test. This test monitors how well Forefront TMG utilizes its content cache, promptly captures poor cache usage trends, and indicates whether/not these disturbing trends are owing to insufficient space in the cache. This way, administrators can be forewarned of deficiencies in the caching mechanism of the Forefront TMG, accurately identify where the bottleneck is, and rapidly fix it, to ensure peak web browing performance.

Purpose	Monitors how well Forefront TMG utilizes its content cache, promptly captures poor cache usage trends, and indicates whether/not these disturbing trends are owing to insufficient space in the cache
Target of the test	A Forefront TMG server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT – The port number at which the specified HOST listens to. By default, this is 1745.
4. ISPASSIVE – If this parameter is set to Yes, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Forefront TMG being monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Data from disk cache:
Indicates the rate at which data is retrieved from the disk drive that is allocated for content caching in this firewall.</td><td>KB/Sec</td><td rowspan="2">A high value is desired for these measures. A steep drop in these values is a cause for concern, as it could indicate low cache hits.</td></tr>
<tr><td>Data from memory cache:
Indicates the rate at which data is retrieved from the memory that is allocated for content caching in this firewall.</td><td>KB/Sec</td></tr>
<tr><td>Disk failure rate:
Indicates the rate at which I/O operations failed on the disk that is enabled for content caching since the start of the firewall service.</td><td>Failure/Sec</td><td>An I/O failure occurs when the Forefront TMG fails to read from or write to the disk. A low value is desired for this measure. A steady increase in this value could indicate that the disk does not have enough space to provide caching services – i.e., to service read/write requests. This in turn may cause many requests to the routed to the origin server, thus increasing bandwidth consumption and delaying web access. If this is to be avoided, you need to make sure that the disk cache is properly sized.</td></tr>
<tr><td></td><td>Disk write rate:
Indicates the rate at which data is written to the disk that is allocated for content caching.</td><td>KB/sec</td><td>A high value is desired for this measure.</td></tr>
</table>

	Total disk failures: Indicates the number of times the Forefront TMG failed to read from/write to the disk since the start of the firewall service.	Number	A low value is desired for this measure. . A steady increase in this value could indicate that the disk does not have enough space to provide caching services – i.e., to service read/write requests. This in turn may cause many requests to the routed to the origin server, thus increasing bandwidth consumption and delaying web access. If this is to be avoided, you need to make sure that the disk cache is properly sized.
	URL commit rate: Indicates the rate at which the URLs are stored in the disk that is allocated for content caching.	Urls/sec	
	URL retrieve rate from disk cache: Indicates the rate at which the URLs were retrieved from the disk.	Urls/sec	A high value is desired for this measure.
	URL retrieve rate from memory cache: Indicates the rate at which the URLs were retrieved from the memory.	Urls/sec	A high value is desired for this measure.
	Space used for disk cache: Indicates the amount of space that is allocated for content caching in the disk.	KB	
	Space used for memory cache: Indicates the amount of space that is allocated for content caching in the memory.	KB	
	Memory usage ratio: Indicates the ratio of fetches from the memory to the total fetches from the overall cache, expressed as percent.	Percent	This measure indicates how well the memory has been utilized for content caching.
	Stored URLs: Indicates the number of URLs that are currently stored in the cache.	Number	A high value is desired for this measure. A low value can end up increasing cache misses and degrading overall performance. If the value is consistently low, it could indicate that the cache does not have enough space to store many URLs. You may then want to increase cache size.

1.1.2 Forefront TMG Email Test

Network administrators are constantly worried about blocking malware in e-mail and making sure that mail servers don't get flooded by spam. The Microsoft Threat Management Gateway is not only capable of removing dangerous messages and junk, but can also block threatening traffic before delivering it to the mail server. This can be achieved when the Exchange Edge Server and Forefront for Exchange are both installed on the TMG Server, thus making the TMG a truly effective e-mail gateway. Combined with Edge Server and Forefront for Exchange, TMG uses multiple anti-virus engines to scan all e-mails for viruses. When a remote computer tries to establish a connection, a new, reputation-based blacklist feature can block incoming spam before any data is sent to the e-mail server. TMG compares incoming messages against a frequently updated list of spam signatures when looking to block incoming spam messages.

It is evident therefore that the true test of the effectiveness of TMG lies not just in the quantity of messages it scans for viruses, but also the quality of messages it finally delivers to the mail server. If too many infected / spam messages find their way to the mail server, it signifies poor TMG performance! This is why, administrators need to keep a close watch on the number of messages the TMG scans and the number of messages it tags as infected or as spam. To perform this check periodically and understand the level of protection the TMG imparts to their critical email servers, administrators can use the **Forefront TMG Email** test.

This test monitors the Forefront TMG and reports the number of email messages that were scanned for malicious content, the number of messages that were blocked for the malicious content present in it and the number of messages that were categorized as spam.

Purpose	Monitors the Forefront TMG and reports the number of email messages that were scanned for malicious content, the number of messages that were blocked for the malicious content present in it and the number of messages that were categorized as spam		
Target of the test	A Forefront TMG server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** – The port number at which the specified **HOST** listens to. By default, this is 1745. 4. **ISPASSIVE** – If this parameter is set to **Yes**, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for the Forefront TMG that is to be monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Scanned messages:** Indicates the total number of email messages that were scanned/inspected for malicious content by the Forefront TMG during the last 24 hours.	Number	A low value for this measure could indicate either of the following: • A processing bottleneck with the TMG that compels it to take too long to scan messages, resulting in a small number of scanned messages at the end of the day; • Many large messages were scanned by TMG during that day;
	Infected messages: Indicates the total number of infected email messages that were blocked by the Forefront TMG during the last 24 hours.	Number	If the value of this measure is close to the value of the *Scanned messages* measure, it indicates that most of messages to the email server during that day were infected. This could indicate a major virus outbreak, which needs to be immediately investigated.
	Spam messages: Indicates the total number of email messages that were categorized as spam by the Forefront TMG during the last 24 hours.	Number	If the value of this measure is abnormally high, it could be because many valid messages have been wrongly categorized as spam. You may then have to fine-tune TMG to avoid such mishaps.

1.1.3 Forefront TMG Firewall Packet Engine Test

In computing, a stateful firewall (any firewall that performs stateful packet inspection (SPI) or stateful inspection) keeps track of the state of network connections (such as TCP streams, UDP communication) travelling across it. The firewall is programmed to distinguish legitimate packets for different types of connections. Only packets matching a known active connection will be allowed by the firewall; others will be rejected/dropped. Packet drops may also occur if the firewall is handling more traffic than it can. To be able to differentiate between these two conditions, administrators should keep track of the packets and connections flowing into the firewall. This is where the **Forefront TMG Firewall Packet Engine** test helps!

The test monitors the traffic flowing through the firewall and reports the rate at which packets are allowed to pass through the firewall. In addition, this test reports the number of dropped, blocked, and backlogged packets, thereby shedding light on what caused the packet drop – genuine packet filtering performed by the firewall or an overload condition on the firewall.

Purpose	Monitors the traffic flowing through the firewall and reports the rate at which packets are allowed to pass through the firewall. In addition, this test reports the number of dropped, blocked, and backlogged packets, thereby shedding light on what caused the packet drop – genuine packet filtering performed by the firewall or an overload condition on the firewall.
Target of the test	A Forefront TMG server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT – The port number at which the specified HOST listens to. By default, this is 1745.
4. ISPASSIVE – If this parameter is set to Yes, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Forefront TMG that is to be monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Packets:
Indicates the rate at which the packets were inspected by this firewall.</td><td>Packets/Sec</td><td></td></tr>
<tr><td>Allowed packets:
Indicates the rate at which the packets were allowed to pass through this firewall.</td><td>Packets/Sec</td><td>A high value is desired for this measure. This measure clearly indicates the load on the firewall.</td></tr>
<tr><td>Backlogged packets:
Indicates the number of packets that are backlogged i.e., the packets that are waiting for the firewall packet engine to create a data pump in the Forefront TMG server.</td><td>Number</td><td>A low value is desired for this measure. This measure can directly have an impact on the Dropped packets measure and vice versa. If there is a steady rise in both the measures simultaneously or if the value of this measure suddenly increases with the immediate rise in the Dropped packets measure, it clearly indicates that the Forefront TMG is not capable of handling the current volume of traffic. If this case occurs consistently even after you observe a constant value in the Active Connections measure, then it is an indication of a bottleneck or capacity constraint with one of the dependent systems of the Forefront TMG such as the DNS or Active Directory.</td></tr>
<tr><td></td><td>Dropped packets:
Indicates the rate at which the packets were dropped by this firewall.</td><td>Packets/sec</td><td>A low value is desired for this measure. If there is a consistent increase in the value of this measure without a corresponding rise in the value of the Backlogged packets measure, it clearly indicates that the Forefront TMG is either processing a lot of malicious traffic or is under attack.</td></tr>
<tr><td></td><td>Data passed rate:
Indicates the rate at which data is allowed to pass through this firewall.</td><td>KB/sec</td><td></td></tr>
</table>

	Created connections: Indicates the rate at which new connections were created on the Forefront TMG server.	Connections/sec	A high value is desired for this measure. A sudden decrease in the value may point to a processing bottleneck with the Forefront TMG.
	Enqueued log items: Indicates the rate at which the logs were enqueued in this firewall.	Packets/sec	
	Packets blocked by NIS: Indicates the rate at which the packets were blocked by the Network Interface service (NIS) in kernel mode.	Packets/sec	
	Active Connections: Indicates the number of active connections through which data is currently passed to this firewall.	Number	Ideally, the value of this measure should be constant over a period of time. If the value of this measure increases suddenly, then it is a clear indicator of an overload condition.
	Avg packets blocked by NIS: Indicates the percentage of packets that were blocked by the NIS in kernel mode.	Percent	
	Dropped Packets ratio: Indicates the percentage of packets that were dropped by this firewall.	Percent	A low value is desired for this measure.

1.1.4 Forefront TMG Firewall Service Test

Load is a factor that can break a firewall! If the Forefront TMG firewall is overloaded with sessions/connections, it may slow down request processing by the firewall. Under such circumstances, administrators will have to identify the type of connections that are causing the overload – are they TCP connections? VoIP sessions? UDP connections? – and investigate why the count of such connections/sessions are unusually high on the firewall. Sometimes, insufficient worker threads on the firewall can also seriously decapacitate the firewall, rendering the firewall unable to handle its load. Another factor that can influence firewall performance is the ability of the firewall to perform DNS resolutions for its service connections; frequent DNS resolution failures can also delay request processing by the firewall. In the event of a slowdown therefore, administrators should be able to accurately pinpoint the reason for the slowdown – is it an overload condition? Is it because not enough worker threads are free? Or is it because of error conditions such as DNS resolution failures? The **Forefront TMG Firewall Service** test helps administrators in this exercise!

This test monitors the firewall service of the Forefront TMG and reports the following:

- The number active TCP, UDP connections and VoIP sessions.
- The rate at which data is read and written to the Forefront TMG
- The number of active worker threads and the number of worker threads that are currently available
- The number of failed and pending DNS resolutions

This way, network administrators can keep track of the firewall service and be proactively alerted to current/potential disturbances in the performance of the service.

Purpose	monitors the firewall service of the Forefront TMG and reports the following: ➢ The number active TCP, UDP connections and VoIP sessions. ➢ The rate at which data is read and written to the Forefront TMG ➢ The number of active worker threads and the number of worker threads that are currently available ➢ The number of failed and pending DNS resolutions
Target of the test	A Forefront TMG server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** – The port number at which the specified **HOST** listens to. By default, this is 1745. 4. **ISPASSIVE** – If this parameter is set to **Yes**, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.
Outputs of the test	One set of results for the Forefront TMG that is to be monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Accepting TCP connections: Indicates the number of connection objects that were waiting for a TCP connection from the Forefront TMG client after a successful remote connection is established.	Number	A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted.
	Active sessions: Indicates the number of active sessions for this firewall service.	Number	
	Active SIP registrations: Indicates the total number of active SIP (Session Initiation Protocol) registrations.	Number	The Session Initiation Protocol (SIP) is a signaling communications protocol, widely used for controlling multimedia communication sessions such as voice and

	Active SIP sessions: Indicates the total number of active SIP (Session Initiation Protocol) sessions.	Number	video calls over Internet Protocol (IP) networks. A basic VoIP call is based on Session Initiation Protocol (SIP), which is the most common protocol used today. A SIP VoIP call is carried out using User Datagram Protocol (UDP), and incorporates two protocols: Session Initiation Protocol (SIP) for call establishment and termination, and Real Time Protocol (RTP) for media (audio and/or video). A VoIP call requires a minimum of three opened connections, one for SIP and two or more for media. Since the media ports are usually selected dynamically by the phone, the firewall needs to understand SIP in order to open and close the media connections. In Forefront TMG, a SIP filter is provided to manage the opening and closing of the media connections automatically, based on the SIP transactions between allowed endpoints. The filter also checks quota, thus preventing DoS attacks by ensuring that only a configurable number of calls or registrations is allowed by the firewall. Accordingly, if the value of the *Active SIP registrations* measure is equal or close to the maximum registrations allowed by the firewall, it could imply that too many VoIP calls are passing through the firewall. When there is an overload condition, you may want to compare the value of this measure with the *Active TCP connections* and *Active UDP* connections measures to understand the type of connections/sessions that are contributing the most to the overload.
	Active TCP connections: Indicates the number of active TCP connections that are currently passing data through this firewall.	Number	The number of connections that are not established and the pending connections are not counted for this measure. A high value could indicate a TCP connection overload on the firewall.
	Active UDP connections: Indicates the number of active UDP connections for this firewall.	Number	A high value could indicate a UDP connection overload on the firewall.
	Data read rate: Indicates the rate at which data is read by the data pump of the Forefront TMG.	KB/sec	A consistent drop in the value of these measures could indicate a read-write slowdown on the firewall.

	Data write rate: Indicates the rate at which data is written by the data pump of the Forefront TMG.	KB/sec	
	Failed DNS resolutions: Indicates the number of `gethostbyname` and `gethostbyaddr` application programming interface (API) calls that have failed.	Number	The API calls are used to resolve host DNS domain names and IP addresses for Firewall service connections. Ideally, the value of this measure should be minimum. A high value can adversely impact the overall health of the firewall service.
	Log queue size on disk: Indicates the size of the Forefront TMG log queue on disk.	KB	
	Pending DNS resolutions: Indicates the number of `gethostbyname` and `gethostbyaddr` API calls that are currently pending resolution.	KB	Ideally, the value of this mesure should be zero. Generally, the TMG firewall relies heavily on DNS to perform name resolution and authentication. Therefore, it is vital that name resolution be performed quickly and efficiently, especially for TMG firewalls that are joined to a domain. If the value of this measure sustains a non-zero value for a longer period, then the name resolution infrastructure should be investigated closely. These are calls used to resolve host DNS domain names and IP addresses for Firewall service connections.
	Pending TCP connections: Indicates the number of pending TCP connections.	KB	Ideally, the value of this measure should be zero. If the value of this measure increases in accordance with the *PendingDNS* measure, then it indicates that the current workload on the firewall is high and the firewall is incapable of handling such huge workloads.
	Worker threads: Indicates the total number of firewall service worker threads that are currently active.	Number	Higher the value of this measure, the busier the firewall service is. A consistent increase in the value could hint at a potential overload condition.
	Connections blocked by NIS: Indicates the rate at which the connections were blocked by NIS in User mode.	Connections/sec	

	Retrieved percentage of DNS domains: Indicates the percentage of time the DNS domain name was found in the DNS cache of the firewall service.	Percent	A high value is desired for this measure.
	Available worker threads: Indicates the number of Firewall service worker threads that are available or waiting in the completion port queue.	Number	The increase in the number may affect the performance of the host / applications.

1.1.5 Forefront TMG H.323 Filter Test

The Forefront TMG includes a H.323 protocol filter which allows multimedia enriched applications like Microsoft Windows NetMeeting® to place calls through the H.323Gatekeeper filter. NetMeeting allows you to video conference using an electronic white board, exchange files, text chat and have voice conversations with two or more parties. If the firewall is H.323 compliant then you will be able to place these calls through it. Most new video conference systems comply with this standard and have had huge success over Microsoft networks. H.323 protocol filter does not directly allow clients to communicate directly with their peers and acts as a true proxy. This method protects the integrity of your network making it more secure and avoiding personal attacks on unsuspecting users.

The **Forefront TMG H.323 Filter** test helps the administrator to track the number of currently active H.323 calls and the total number of H.323 calls handled since the start of the firewall service.

Purpose	Helps the administrator to track the number of currently active H.323 calls and the total number of H.323 calls handled since the start of the firewall service		
Target of the test	A Forefront TMG server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** – The port number at which the specified **HOST** listens to. By default, this is 1745. 4. **ISPASSIVE** – If this parameter is set to **Yes**, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for the Forefront TMG that is to be monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Active calls:** Indicates the number of H.323 calls that are currently active.	Number	
	Total calls: Indicates the total number of H.323 calls handled by the H.323 filter since the start of the firewall service.	Number	

1.1.6 Forefront TMG Socks Filter Test

Socket Secure (SOCKS) is an Internet protocol that routes network packets between a client and server through a proxy server. Practically, a SOCKS server proxies TCP connections to an arbitrary IP address, and provides a means for UDP packets to be forwarded. The Forefront TMG can perform as a SOCKS Server or a SOCKS proxy. The SOCKS filter provided with Forefront TMG forwards requests from SOCKS applications to the Microsoft Firewall service. Forefront TMG checks the access policy rules to determine if the SOCKS client application can communicate with the Internet.

To understand how well the Forefront TMG filters and processes requests from SOCKS applications, use the **Forefront TMG Socks Filter** test. With the help of this test, you can identify the number of active connections and sessions that are connected using the SOCKS protocol and the rate at which data is read from and written to the client. In addtition, this test reveals the rate of pending DNS resolutions and those DNS resolutions that were successful.

Purpose	To help administrators understand how well the Forefront TMG filters and processes requests from SOCKS applications		
Target of the test	A Forefront TMG server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST -** The host for which the test is to be configured. 3. **PORT** – The port number at which the specified **HOST** listens to. By default, this is 1745. 4. **ISPASSIVE** – If this parameter is set to **Yes**, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for the Forefront TMG that is to be monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Active connections:** Indicates the total number of active connections (connected through SOCKS protocol) that are currently passing data through this firewall.	Connections/sec	The value of this measure is incremented by one for each successfully established SOCKS connection and decremented by one if the SOCKS connection is terminated.
	Active sessions: Indicates the total number of active sessions that are connected through SOCKS protocol.	Sessions/sec	This is a good indicator of the load imposed on the firewall by the SOCKS sessions.
	Data read rate: Indicates the rate at which data is read from the client by the server when the connections are established through SOCKS protocol.	KB/sec	
	Data write rate: Indicates the rate at which data is written to the client by the server when the connections are established through SOCKS protocol.	KB/sec	
	Connecting connections: Indicates the number of connections that are currently waiting for a remote computer to connect to using the SOCKS protocol.	Connections/sec	
	Listening connections: Indicates the rate at which the SOCKS filter listens for an incoming connection on a specified port, when a BIND command is issued to the SOCKS filter by a client.	Connections/sec	
	Pending DNS resolutions: Indicates the number of Winsock `getaddrinfo()` requests that are currently pending per second.	Connections/Sec	A low value is desired for this measure. These requests resolve host DNS names and IP addresses for SOCKS connections.

	Successful DNS resolutions: Indicates the number of DNS resolution requests made using SOCKS protocol that are currently resolved per second.	Connections/Sec	A high value is desired for this measure.

1.1.7 Forefront TMG Web Proxy Test

Forefront TMG application filters provide an extra layer of security at the Microsoft Firewall service. Application filters can access the data stream or datagrams associated with a session within the Firewall service. Application filters are registered with the Firewall service and work with some or all of the application-level protocol streams or datagrams. An application filter can perform protocol-specific or system-specific tasks, such as authentication and virus checking. Some of the application filters provided with the Forefront TMG are:

- DNS filter
- FTP access filter
- H.323 filter
- Intrusion detection filters
- RPC filter
- SIP Access Filter
- SMTP filter
- SOCKS filter
- TFTP Access Filter
- Streaming media application filters
- Web Proxy filter

Web Proxy Filter works at the application level on behalf of a client requesting Web-based traffic. Although you cannot disable this filter, you can configure whether the filter applies to specific protocols. By default, it is applied to the Hypertext Transfer Protocol (HTTP), which is configured as follows:

- Direction is Outbound
- Protocol Type is TCP
- Port is 80

When Web Proxy Filter is enabled for a protocol, that protocol can use the following features, if applicable:

- Authentication
- HTTP filtering

To gauge how effectively this filter performs authentication and HTTP filtering, you can use the Forefront TMG Web Proxy Filter test. Using this test, you can proactively detect current or probable bottlenecks or risks in data transfer between web proxy clients and servers.

Purpose	You can proactively detect current or probable bottlenecks or risks in data transfer between web proxy clients and servers
Target of the test	A Forefront TMG server

<table>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT – The port number at which the specified HOST listens to. By default, this is 1745.
4. ISPASSIVE – If this parameter is set to Yes, then it means that, by default, all the Forefront TMG servers being monitored by the eG system are the passive servers of a Forefront TMG cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Forefront TMG that is to be monitored</td></tr>
<tr><td rowspan="5">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Data array received:
Indicates the rate at which data is received from the computers protected by the Forefront TMG within the same array.</td><td>KB/Sec</td><td>The performance of the Forefront TMG is affected when the scanned email messages are too lengthy in terms of size and attachments.</td></tr>
<tr><td>Data array sent:
Indicates the rate at which data is sent from the computers protected by the Forefront TMG within the same array.</td><td>KB/Sec</td><td></td></tr>
<tr><td>Total data array:
Indicates the rate at which data transmission takes place in the computers protected by the Forefront TMG within the same array.</td><td>KB/Sec</td><td>This measure is the sum value of the Data array received and Data array sent measures.</td></tr>
<tr><td>Avg time to service the request:
Indicates the time taken to service a web proxy client request.</td><td>Secs</td><td>The value of this measure does not include the time taken for servicing requests by the SSL tunnel. A high value for this measure indicates a processing bottleneck with the firewall service.</td></tr>
<tr><td>Average request speed:
Indicates the rate at which requests are serviced during the last measurement period.</td><td>KB/Sec</td><td>The value of this measure does not include the time taken for servicing requests by the SSL tunnel. A high value for this measure indicates a processing bottleneck with the firewall service.</td></tr>
</table>

	Active Web Sessions: Indicates the number of web proxy sessions that are currently active.	KB/Sec	The web proxy sessions can either be from different clients or from a client with a single IP address where authentications of the client does not take place.
	Data served from cache in ranges: Indicates the total number of bytes that are serviced from the cache in response to HTTP requests containing range headers.	KB	
	Data served in ranges: Indicates the total number of bytes that are returned in response to HTTP requests containing range headers during the last measurement period.	KB	
	Cache hit ratio: Indicates the percentage of the web proxy client requests to the Forefront TMG that were successfully serviced by the cache.	KB	A high value is desired for this measure and a high value generally indicates that the response time for each service is faster. A value of zero for this measure indicates that the caching capability is not enabled. A low value for this measure generally indicates that either the size of the cache is too small or the requested objects are not avaialbe in the cache.
	Data received rate: Indicates the rate at which data is received from the web proxy clients.	KB/Sec	A high value is desired for this measure. A consistent decrease in the value of this measure clearly indicates that the servicing of the requests is considerably delayed.
	Data sent rate: Indicates the rate at which data is sent to the web proxy clients.	KB/Sec	A high value is desired for this measure. A consistent low value of this measure clearly indicates that the servicing of the requests is considerably delayed.
	Total data transferred: Indicates the overall rate of data transmission between the Forefront TMG and the web proxy clients.	KB/Sec	This measure is the sum of the Data received rate and the Data sent rate measures.
	Failed request rate: Indicates the percentage of requests that failed.	Percent	A low value of this measure is desired.

	Average request processing rate: Indicates the rate at which the web proxy requests were processed.	KB/Sec	This measure takes into account only the HTTPS traffic that is inspected by the Forefront TMG. A high rate is indicative of good health of the firewall service.
	Current compression ratio: Indicates the ratio of the compressed HTTP response body size to that of the uncompressed body size, expressed in terms of percent during the last measurement period.	Percent	This measure takes into account the HTTP responses that are compressed by the Forefront TMG alone.
	Requests from array member: Indicates the ratio of the requests received from another member of the array to the total number of requests that failed during the last measurement period.	Percent	
	Requests to array member: Indicates the ratio of the requests sent to another member of the array to the total number of requests that failed during the last measurement period.	Percent	
	Unknown SSL sessions: Indicates the total number of unknown SSL sessions that were serviced by the SSL tunnel.	Number	
	Connect errors: Indicates the ratio of the errors that occurred while connecting to the total number of failed requests, expressed as percent during the last measurement period.	Number	
	HTTP requests: Indicates the total number of HTTP requests made to the Forefront TMG since the start of the firewall service.	Number	

	HTTPS requests: Indicates the total number of secured HTTPS sessions that were serviced by the SSL tunnel.	Number	
	Outgoing connections: Indicates the rate of outgoing connections that are made from the Forefront TMG.	Connections/sec	
	Incoming connections: Indicates the rate of incoming connections that are made to the Forefront TMG.	Connections/sec	
	Requests: Indicates the rate of incoming requests that were made to the web proxy.	Connections/sec	A higher value indicates that the Forefront TMG would require more resources to service all the incoming requeests. Tbis measure is a clear indicator of the Forefront TMG's load handling ability.
	Reverse data transferred: Indicates the overall rate of data transmitted between the Web proxy and the web publishing servers in response to the incoming requests.	KB/Sec	
	Thread pool active sessions: Indicates the rate at which active sessions are currently serviced by the thread pools.	KB/Sec	A high value is desired for this measure.
	Web proxy authentication queue length: Indicates the number of items that are currently waiting in the web proxy authentication queue.	Number	
	Compression ratio of size reduction: Indicates the ratio of average size reduction of the HTTP response body to the uncompressed body size during the last measurement period.	Percent	

	FTP requests: Indicates the number of File Transfer Protocol (FTP) requests that were made to the web proxy.	Number	A low value for this measure is an indication of the poor caching policy of FTP objects. Try altering the caching policy to get better results.
	Thread pool failures: Indicates the number of requests that were rejected due to the thread pool being full.	Number	

2

Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **Forefront TMG**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.