

gateProtect Performance Test of xUTM Appliances

Introduction

The gateProtect xUTM appliances were tested under loads that as closely as possible reflected market requirements. As the results showed, the xUTM appliances sold by gateProtect offer extremely high levels of performance. gateProtect has purposely chosen powerful hardware in order to have sufficient performance reserves at all times and in any situation. This is particularly important at peak loads.

Test procedure

One sample of each appliance offered by gateProtect was subjected to a load test using version 8.0 of gateProtect's firewall. The tests were carried out between two powerful servers, one of which supplied web pages, provided a mail server and responded to layer 3 packets that were generated on the other server as a pure load and sent via a VPN tunnel, some of the packets being encrypted and some unencrypted. The other server called up web pages and sent and received mails. The number of helper processes for the mail proxy on the firewall was increased from 15 to 75, in anticipation of this becoming the standard in version 8.1. The test scenario was defined as follows:

1. Measurements

For HTTP and HTTPS traffic we measured the time between the request and the web page being delivered (including all its images, etc.). For POP3 we measured the time between the request and the delivery of all mails currently on the mail server. Once fetched, the mails were always deleted from the server. For SMTP we measured the elapsed time until the mail server had accepted a mail for sending and had confirmed this. These measurements were shown in diagrams where they were plotted against the number of simulated users who were simultaneously using the firewall.

A simulated user is defined in terms of the load it generates over the course of an eight-hour working day.

Load per user:

- a) HTTP: 190 page views per user over a period of eight hours at an average of 400 KB per page
- b) HTTPS: 10 page views per user over a period of eight hours at an average of 400 KB per page
- c) POP3: 50 incoming mails per user over a period of eight hours at an average of 50 KB
- d) SMTP: 25 outgoing mails per user over a period of eight hours at an average of 50 KB

2. Data traffic

While these measurements were being taken the firewalls were also subjected to loads from the following data traffic, which was also scaled with the number of users and constitutes part of the definition of a simulated user.

Load per user:

- a) FTP: One 700 MB^{*1} download per 10 users
- b) VPN: One firewall-encrypted 10 Kbit/s data stream per 10 users via VPN tunnel (36 MB over the eight-hour period)^{*2}
- c) Traffic Layer 3: One unencrypted 100 Kbit/s data stream per user (360 MB over the eight-hour period)

*1 This assumes that within an 8-hour period every tenth user will launch an FTP download of around 700 MB

*2 This assumes that within an 8-hour period every tenth user will set up a VPN tunnel.

gateProtect Performance Test of xUTM Appliances

3. Activated functions

The following firewall functions were activated for the measurements, where permitted by the relevant appliance license:

- a) Antivirus (for HTTP, HTTPS, POP3 and SMTP)
- b) IDS (default setting: 2632 rules)
- c) Web blocking (sex blacklisted)
- d) Logging and monitoring

4. Measurement points

Measurements were determined for the following numbers of users:

25, 50, 75, 100, 150, 200, 250, 500, 750, 1000, 1500, 2000, 2500, 3000,
3500, 4000, 4500, 5000, 6000, 7000, 8000, 9000 und 10000

Each measurement lasted one hour and was then extrapolated up to eight hours.

5. Representation

The measurement points were plotted on the diagrams as colored symbols and interpolated using Bézier curves.

6. VPN throughput

In addition to the measurements defined above, an additional test was carried out to measure the VPN throughput for each hardware system. For these tests a client-to-server VPN SSL tunnel with AES 128 encryption was set up from one of the test computers to the firewall and data was then sent through this tunnel to the other computer on the firewall; no other data traffic passed the firewall while these measurements were being carried out. The maximum throughput achieved for this data was measured, and the results of these measurements are given below in Mbit/s for each appliance together with the relevant diagram. Where a number of users were using the VPN SSL the throughput per user was calculated by simply dividing the measured overall throughput by the number of users.

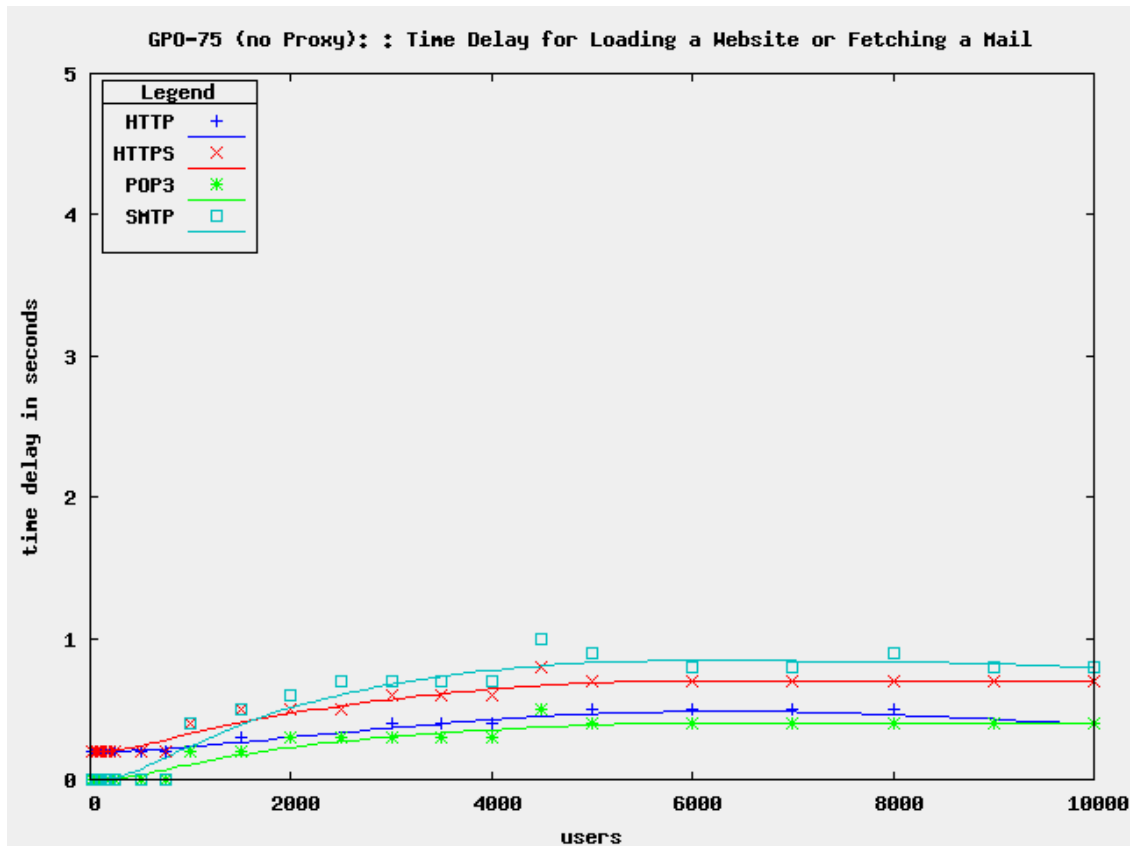
gateProtect Performance Test of xUTM Appliances

Results

1. GPO 75

The GPO 75 does not use a proxy for any services. Since it only needs to filter packets according to the rules of the firewall, all delays are less than one second, even with 10000 simulated users. Alongside these results it should be noted that, for this hardware, the additional load described in section 2 cannot be scaled up with the number of simulated users beyond about 1000 users, because the installed Fast Ethernet network cards do not permit more than 100 Mbit/s.

Measured VPN throughput: 40 Mbit/s.

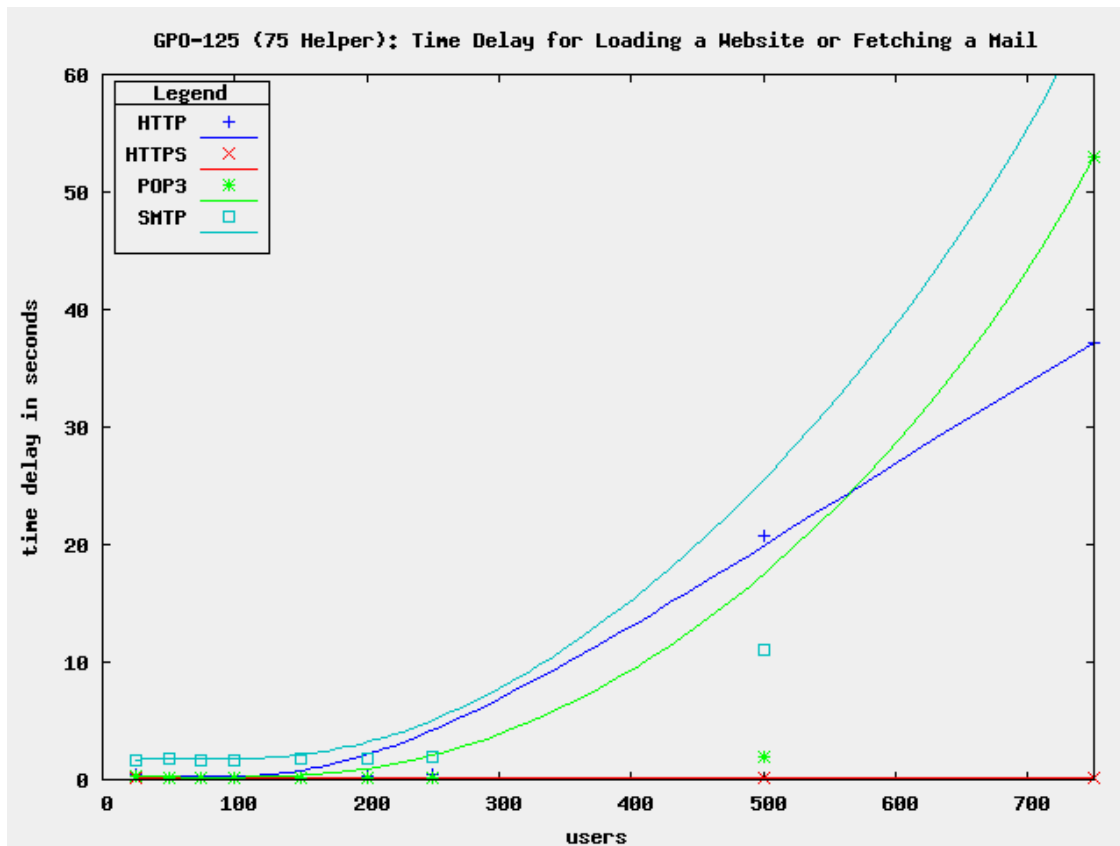


gateProtect Performance Test of xUTM Appliances

2. GPO 125

The license for the GPO 125 does not permit an HTTPS proxy. The delays to HTTPS therefore remain under one second for any number of users, as for the GPO 75. However, when a proxy is used, working via this firewall becomes extremely sluggish from about 300 users, and for more than 500 users it is definitely unsuitable.

Measured VPN throughput: 40 Mbit/s.

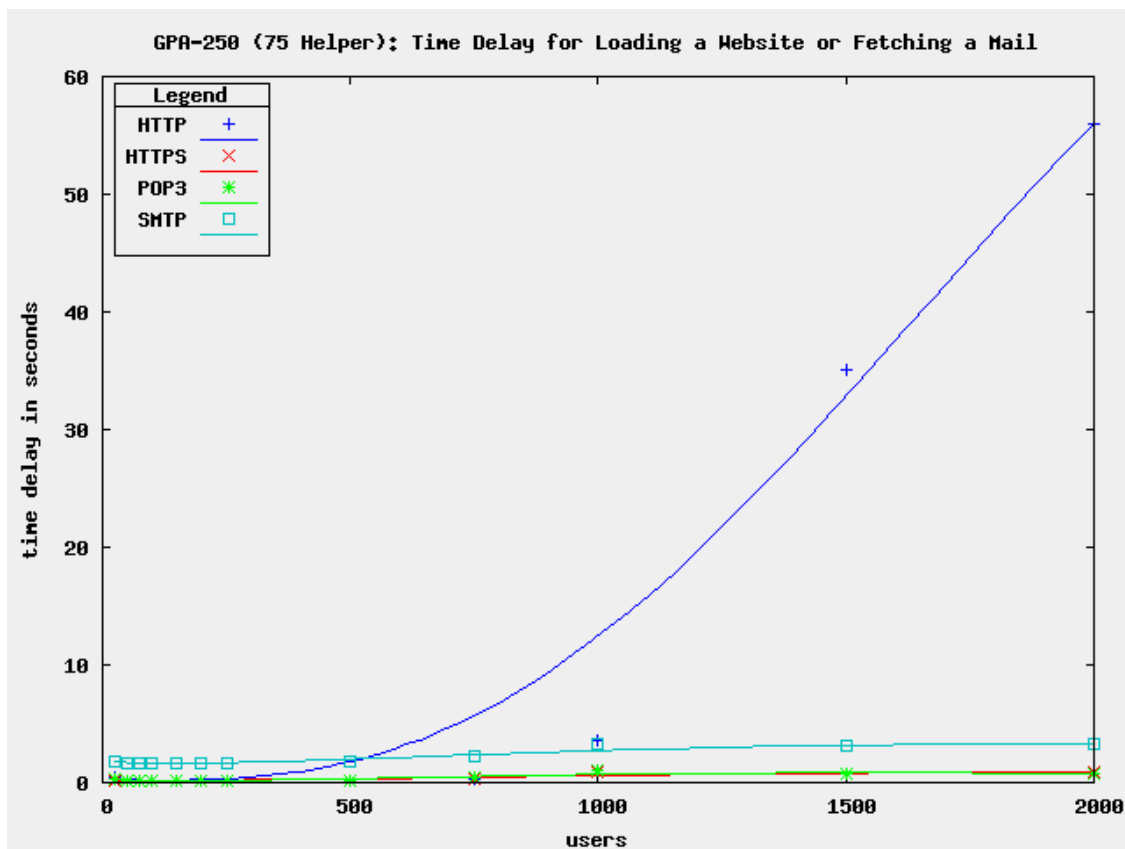


gateProtect Performance Test of xUTM Appliances

3. GPA 250

The license for the GPO 250 also does not permit an HTTPS proxy, so for HTTPS measurements the same applies as for the GPO 125. Up to around 1000 users, all the measured delays are below two seconds. With 750 users the HTTP proxy is struggling and effective surfing is no longer possible, although as regards the volume of mail this hardware could certainly handle more users.

Measured VPN throughput: 140 Mbit/s.

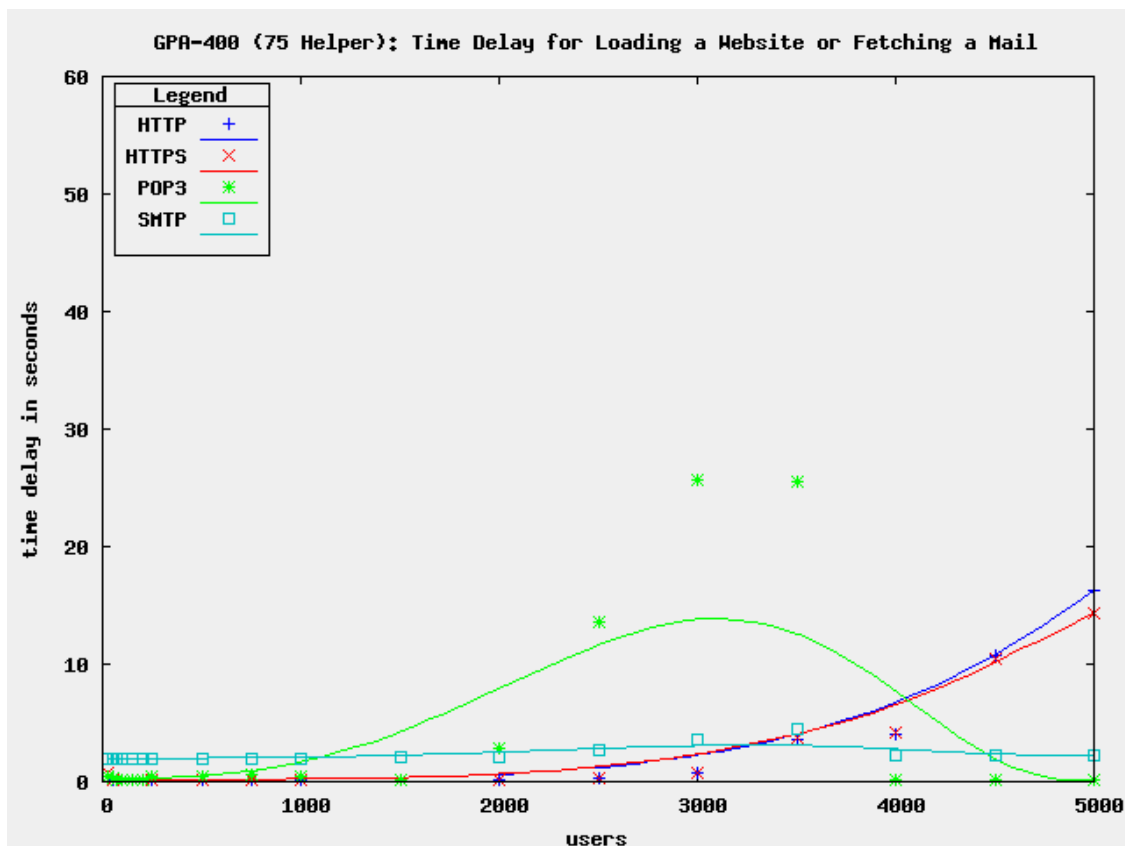


gateProtect Performance Test of xUTM Appliances

4. GPA 400

Although the GPA 400 could handle up to around 5000 users for HTTP, HTTPS and SMTP, considerable delays occurred with more than about 2000 users when fetching mails via POP3. The downward curve for POP3 for higher numbers of users is an artifact: at this point many mail requests were being rejected, and refusing those connection setup requests really was that fast.

Measured VPN throughput: 200 Mbit/s.

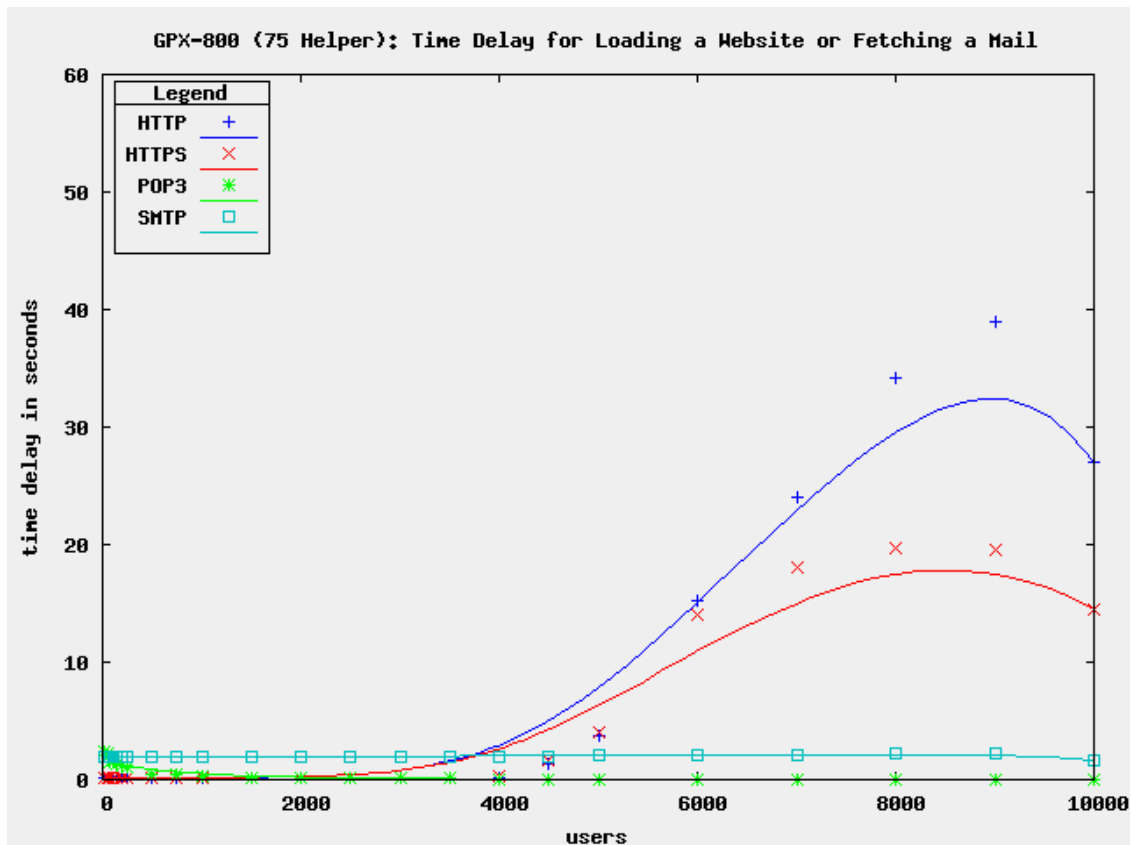


gateProtect Performance Test of xUTM Appliances

5. GPX 800

From 5000 users the GPX 800 exhibits delays for HTTP and HTTPS. On this machine the mail proxy shows no further weaknesses up to 10000 users.

Measured VPN throughput: 250 Mbit/s.



gateProtect Performance Test of xUTM Appliances

6. GPX 1000

On the GPX 1000, working becomes slightly sluggish from around 8000 users, but is still entirely possible right up to 10000 users. It should be taken into account, however, that in networks with so many users the IDS, antivirus and even VPN are generally provided by external systems, so that a major part of this firewall's system resources are left unused.

Measured VPN throughput: 300 Mbit/s.

