White Paper

VPN Performance Validation

How secure are secure tunnels?

Validation of security platforms should cover functionality, scalability and most importantly performance of both tunnel establishment and the actual application performance in the secure tunnel.

VPN Performance Validation

How secure is your data?

Perhaps the introductory question should read – How is sensitive data being accessed and transferred between sites?

The challenge many of us face today is understanding how secure our data is, especially when it is stored offshore on a cloud managed platform. The communication of data securely between two sites is one of the major issues facing IT systems today. The challenge is further compounded when the applications in use are latency sensitive.

Recent online events have shown that the time needed to gain access to sensitive data is minutes! Reports show that a dictionary type password attack or brute force attack can forcefully provide access to a protected network and data storage devices in less than quarter of an hour. Sadly, the explosion in cheap computing power and personal data on social networking communities is making the criminal's job far easier.

In terms of data access, the most common means of securely accessing/transferring data is through secure Virtual Private Networks (VPNs) of which there are two types: client (IPSec VPN) and/or clientless mode (TLS/SSL/DTLS VPN).

VPN Performance Validation

Unfortunately, it has been shown, that both methods have a number of vulnerabilities. For example a simple flaw with an IPSec VPN security deployment; is the fact that a number of systems are generally interconnected behind the security device. In the event of a security breach, all of the interconnected systems can become accessible in the event of a successful VPN client exploit.

As with all security mechanisms the benefits tend to outweigh the flaws; IPSec continues to be used because of the stronger security principles enabled through certificate exchanges. Certificate security offers greater protection when compared with the process of simple username and password login of the clientless model.

So which model do you choose when it comes to enabling secure information exchange or why chose one over the other? Before ruling either model out, a number of performance validations should be considered. The goal is to not limit the performance validation to solely assessing device vulnerabilities, but to also use the validation process to determine the quality of experience (QoE) from a VPN end-user's perspective.

Validation of security platforms should cover functionality, scalability and most importantly performance of both tunnel establishment and the actual application performance in the secure tunnel.

Choosing the right VPN validation model

As with all performance validation strategies; user Quality of Experience (QoE) should act as the primary motivation to measuring the performance of a secure VPN. When it comes to performance validation of secure VPNs, understanding how an application is behaving inside the tunnel is critical. Simply put, if the end user has a bad experience they will refrain from using the secure service and may in fact attempt data transfer via unsecure paths.

A comprehensive performance validation of secure VPNs for QoE will incorporate various application flow types traversing the secure tunnel. It is recommended to include more than one application type inside the secure tunnel, realistically matching the end-user usage scenario. Applications should include both data and latency sensitive applications of voice and video. It's also recommended to include sinister flows such as a DDoS type attacks e.g. ping floods in the secure VPN.

To date a limited or macro view has been taken when defining a test strategy for secure VPNs. Generally the performance of the security device has been shown to be a measurement of the number of tunnels established and maintained by a security device.

This blinkered view provides a limited perspective on performance i.e. "X" users can login! The following is an example showing why this approach is weak, consider the following:

- What happens when the security device suddenly reboots?
- What happens when each tunnel or end-point has to renegotiate a secure session?

In both cases the end-users are going to experience a disruption to service. How quickly each tunnel re-establishes is an important validation for robustness, but more importantly how quickly the data exchange can re-commence is key.

The importance of stateful per flow emulation and analysis

Reflecting further on the security device reboot scenario, consideration must be given to a number of aspects relating to the tunnel establishment process. Firstly, each secure tunnel session establishment is unique! No two sessions are the same. This means part of the validation requirement is to create unique secure tunnel requests. Secondly, to reliably repeat validation e.g. validate performance when a new update/upgrade patch is applied, an emulation platform is recommended such as VIAVI TeraVM.

As each secure tunnel flow establishment is unique, analysis must be done on a per flow basis, in real time.

Consider the following corporate IT mission statement, as to why per flow analysis is important:

"It's critical that our company's IT team reacts instantly when problems occur. In the event of a crisis, if it takes one member of the team longer to establish a secure connection we lose the ability to respond effectively. In addition, if the delays in accessing mission critical data are prolonged, the team's inability to react as a whole has the potential to seriously damage the company's reputation."

The statement highlights the importance of per flow, in which success is built on the sum of each member's contribution. The other aspect to per flow emulation and performance validation is the use of stateful endpoints. That is each endpoint must negotiate its own secure session and must have the ability to establish multiple application flow types inside the secure VPN.

In the event of an internal network issue e.g. badly configured network QoS, the connected endpoint handling the secured application will also be impacted. This application focused performance validation is independent of the VPN performance validation, but necessary to ensure the smooth delivery of service in secure communications.

In the real world scenario e.g. the company IT team, the lack of visibility in determining performance at an application level could have serious knock on effects which may in fact impact the business.

In addition, per flow emulation and performance validation is necessary to assess the security device's ability to correctly identify illegal flows. In this situation, per flow emulation is used to mix endpoints with good and bad flows attempting to connect with the secure zones.

Finally, a comprehensive performance validation assessment for VPN, can be separated into two key sections:

- Performance of secure VPN session establishment, which includes:
  - Ability to establish a tunnel
  - Time to establish a tunnel
  - Bandwidth of established tunnel
- Performance of applications inside secure tunnels, which includes:
  - Ability to establish connection to application servers
  - Time to first usable data, voice or video flows
  - Available bandwidth for application traffic

When assessing secure VPNs, real time analysis provides insight into how much management bandwidth is required to maintain the secure connection. Evidently, the more tunnel management packages on the line, the less space for the application packets.

Per flow performance validation for secure VPNs

As highlighted a successful validation strategy for secure VPNs requires a per flow approach, which includes the performance measurements on each and every secure VPN flow in real time, with a granular view of quality on each and every application running over the tunnel.

Endpoints must be fully stateful, which encapsulates both the tunnel establishment and the applications inside the secure tunnel. Quality of Experience (QoE) performance measurements are a set of granular measurements on the secure tunnel and the application performance been secured by the tunnel.

Migrating to public/private cloud managed platforms

Virtualization is creating richer networking environments, as cloud managed platforms grow in strength so too does the number of available virtual network functions (VNFs) which include virtual firewalls with secure VPN access.

As with the physical form, the virtual network secure VPNs must maintain the high level of reliability and performance, along with a consistent level of QoE.

Validating the reliability, performance and QoE is a challenge as traditional methods of connectivity no longer apply. The accessibility of certain zones on the cloud managed platform may result in the lack of physical connectivity e.g. tenants within various geo-locations. In many cases performance validation must be conducted via a service chaining principle in the cloud.

For specific details on validating virtual or cloud infrastructure see Cobham's website, which includes a number of application notes for network function virtualized system assessment.

In general, by choosing to utilize Cobham's TeraVM a per flow traffic emulation solution, users are also future proofing their investment, as it will enable them to assess secure VPN appliances as part of any cloud migration strategy. The approach to validating of the secure VPNs as covered by this application note are relevant to the cloud as well.

Why choose TeraVM to validate VPN performance?

The simplest explanation is to take the example of a mobile employee, it is clear there is a need to be continuously connected to data from external locations and indeed use a number of applications to access/transfer data.

The applications in use in the secure connection can vary and in most cases will pass over a corporate VPN and network, the secure connection will enable access to data, voice, and potentially video files.

Using TeraVM's per flow emulation it's possible to validate the performance of the individual VPN, but more importantly scale to the representational load of modern mobile employees at scale i.e. thousands of VPNs running multiple application flows on the individually connected VPNs.

TeraVM's per flow architecture is used to represent real world usage scenarios to determine the performance impact on application and services being delivered in secure tunnels from a number of leading security vendors.

Client versus Clientless VPNs

Defining quantifiable VPN performance measurements

In security administration, is it good enough to say that a particular flavour of security policy enables a greater number of secure tunnels? Or indeed, by loosely stating that by varying the encryption levels i.e. 128 bits vs 256 bits, there is a performance impact, without validation?

When it comes to client versus clientless policies a commonly stated outcome to the question: encryption rate versus the available number of usable VPN seat licences is that "the number of tunnels established will vary".

This vague answer is not a valid QoE assessment as no performance validation has been put on the actual application being secured on a per secure flow basis.

Using a simple delay sensitive application such as VoIP, is it accurate to state that the variation in encryption levels i.e. from 128 bits to 256 bits impacts call quality? The following questions need to be answered:

- How much does the encryption rate impact a latency sensitive application?
- Should the cipher protocol be clientless (SSL/TLS) instead of client (IPSec) based?

As part of a wider secure VPN performance validation strategy the aim is to include a mix of applications using a mix of ciphers, with varying levels of encryption. By doing so it's now possible to derive the best possible combination of security for the application type of video, voice and data.

Illegal flows inside secure VPNs

In the unfortunate event where a security breach occurs, the criminal now has the ability to establish secure tunnels to the inside network, Enterprises must be aware of the potential damage and disruption that this can cause.

One of the major challenges of defending against such attacks is to correctly identify illegal flows inside secure VPNs. This is not as easy as it sounds, for example not all flows are illegal, there may be legitimate requests to services with incorrect privileges, or a genuine user starts downloading a large volume of data in a secure session.

However, a more sinister attack is the use of a VPN to implement a denial of service attack, where an attacker ties up an internal server hosting a critical intranet site or database.

Security attack mitigation testing must include both the inconspicuous and the more easily identifiable flows made up from the Common Vulnerabilities & Exploits (CVEs) cybersecurity threat library and DDoS type attacks.

Security attack mitigation testing in secure VPNs further highlights the need for per flow test functionality with stateful flows inside the secure VPN.

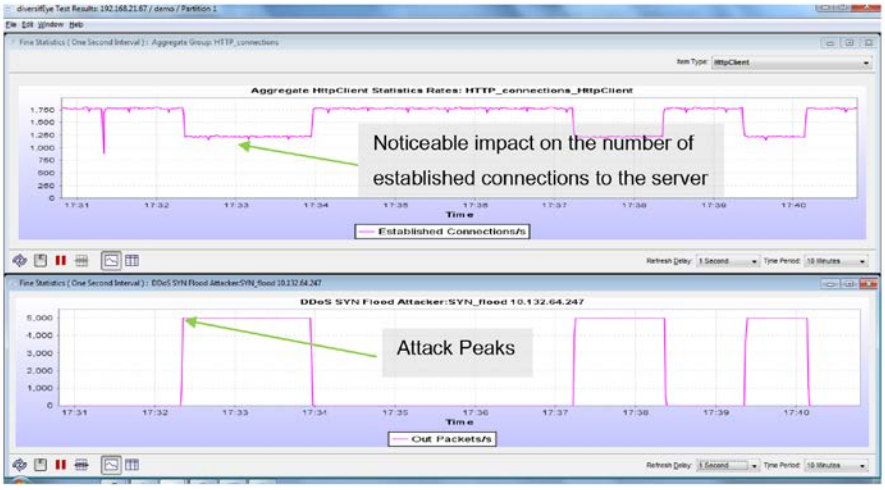

Figure 1 : The above figure highlights the impact a SYN flood attack has on a server on the inside of a firewall.

In figure 1, the attack is being carried to the inside via a legitimate secure VPN. Whilst the firewall stopped SYN attacks against its own base address, it shows that it is powerless against this type of hidden inside attack. It's also possible to see the impact to users connecting to the server as the number of session connections drop.

In the instance above the severity of the attack against the server is small, which may go unnoticed. An interesting side note in relation to the attack test above is the CPU utilization on the security device, it remained steady showing no signs of stress. The device simply passed the SYN flood along. In addition the threat logs show nothing strange or abnormal.

This type of attack can be disruptive and largely go unchecked. The only indication that there may be something happening will be seen at the server. When the attack is mild, the impact may not disrupt all end-users, however the frustration of using the service will force some users to seek alternative means of communication, which may not always be secure.

As the SYN attack becomes more outrageous the eventual outcome is no connectivity to the application server on the inside. In the event of such an attack the security device maintains all the connected VPNs i.e. connection to the enterprise is maintained.

In the figure 2 below, the attack is shown to be of varying ferocity and duration, which impacts the server's connection rate accordingly. The server finally succumbs to the attack with no possibility of a valid VPN user connecting to the service, after the attack ends it takes the server a period of time to recover.

For all users of the secure VPN tunnels, connection to the application server under attack is re-established only when the attack rate eases or stops.

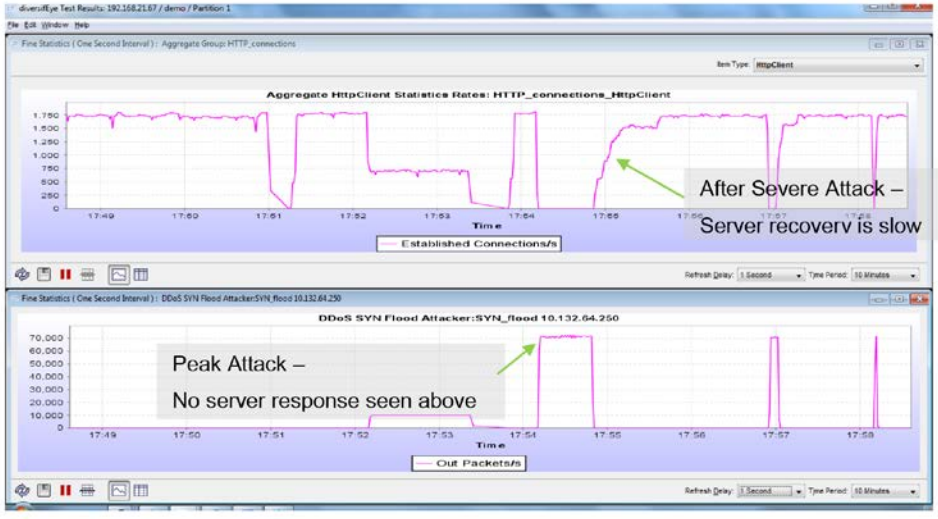

Figure 2 : The impact of a DDoS type attack in a secure VPN connection.

VPN performance validation strategy

VPN seat license availability and reallocation latency

One of the basic validation tests when connecting to a VPN security device is to determine the max number of secure VPNs the device can establish and maintain. Further performance validation should confirm the security device's ability to re-allocate a VPN seat licence when one becomes available.

VIAVI TeraVM per flow test architecture is ideal for these dynamic type tests. Each VPN request entity can be dynamically brought in and out service during live tests. Alternatively, the TeraVM can be given a sample VON user profile range to cycle through active states.

An example performance validation scenario for a firewall with a 50 VPN seat licence is to emulate 100 users. During a validation test run the 50 licences will be allocated. In TeraVM, during live test runs, it's possible to easily set any emulated endpoint into an "out of service" state, meaning a VPN licence will become available in the security device. This licence should then be allocated to one of the waiting emulated end-points.

The latency in which licenses get re-allocated to the awaiting VPN users is a key performance measurement. TeraVM automation suite can be used to automatically detect and log when the re-allocation of the licence occurred.

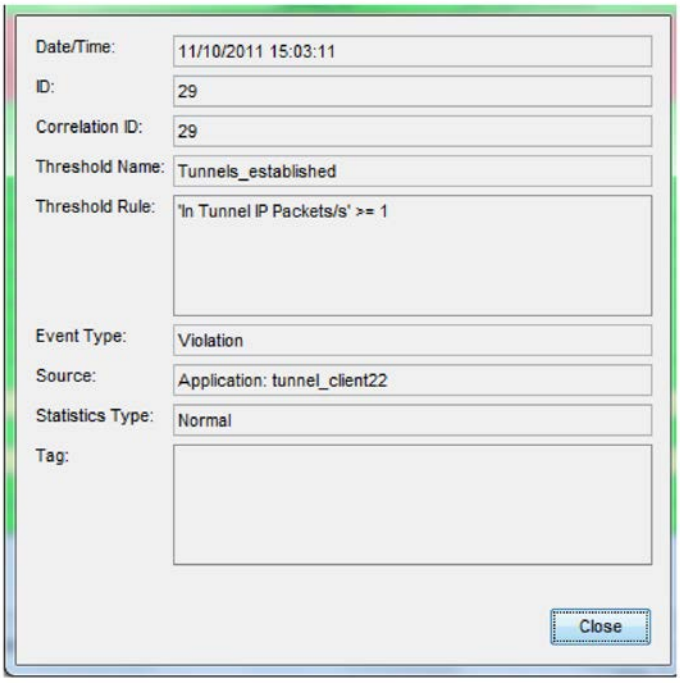

Figure 3: TeraVM can be configured to capture the exact time the emulated VPN client can pass traffic through the secure VPN tunnel.

True QoE measurements for VPN latency must focus on when an end user can actually use a secure tunnel and not be limited to stateless time such as the time it takes to establish a tunnel.

Building a Benchmark

Before commencing any type of performance validation scenario it's worth defining a benchmark on a per test methodology basis. An example validation scenario is to define the max number of connections that clients or end-users can establish with an application server. For this application note the simplest form is to measure connectivity against a HTTP server.

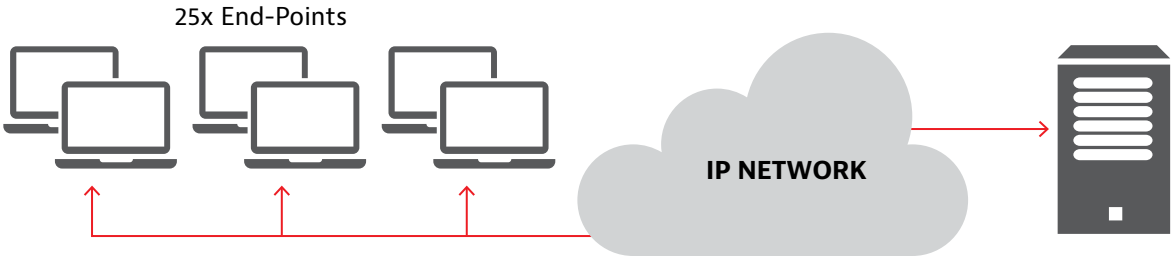

Figure 4 : Emulating user and server side scenarios

A simple validation is to emulate a defined number of end-points attempting connection with a single server. Hence the valid metric at this stage (external to the VPN), is the number of connections established by the server. Using this as a benchmark it's now possible to validate how the VPN security appliance will behave when placed in the path between the emulated client and server endpoints.

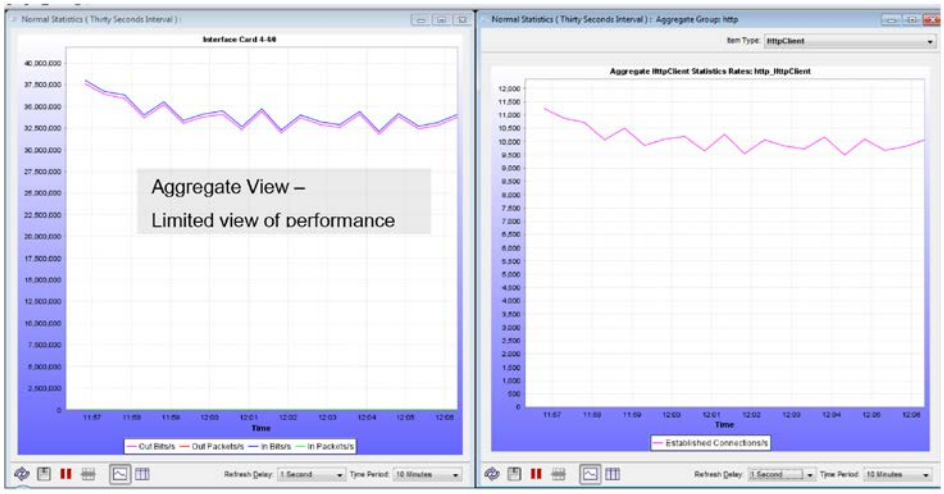

Figure 5 : The above is an aggregate performance overview, note the large number of connections capable. The result on its own, offers very little in terms of a VPN performance insight. Its main purpose will be to show how the security device limits the number of connections per second.

Introducing the security appliance

Once the basic benchmark has been established, the performance validation now focuses on the security performance with the inclusion of the VPN tunneling appliance. As per the benchmark use case, the endpoints will request content from a single server, accept now, the connections will be implemented in secure tunnels.

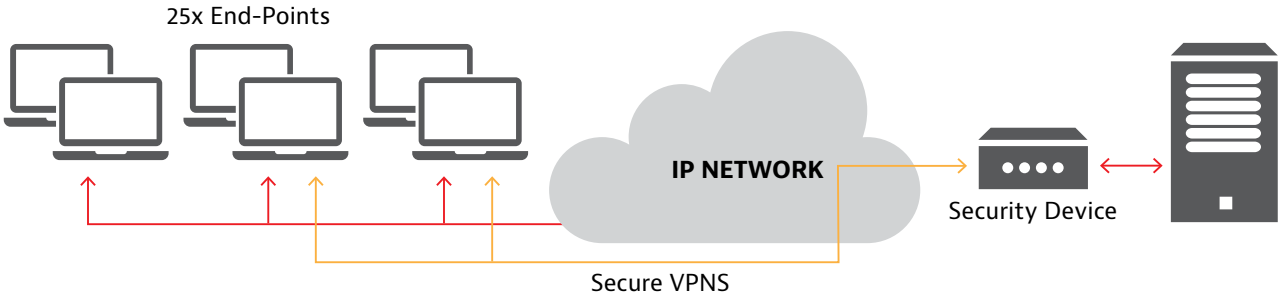

Figure 6 : The inclusion of a VPN tunneling appliance

Analysis of VPN tunnel performance in real time

Using TeraVM it's possible to emulate endpoints using the various types of security encryption protocols (SSL, TLS, DTLS, IPSec – IKEv1/v2).

As suggested a sample performance metric used when analysing QoE is the tunnel establishment time.
A significant aspect to measuring performance of the tunnel establishment time, is the ability to capture the result in real-time. The figure below shows a single secure connection being established in real time.

Using TeraVM's per flow live performance result analysis it's possible to see the individual IPSec enabled tunnel establishment times, below shows the secure connection takes less than 0.5 seconds to establish. TeraVM provides these measurement points on each and every flow.

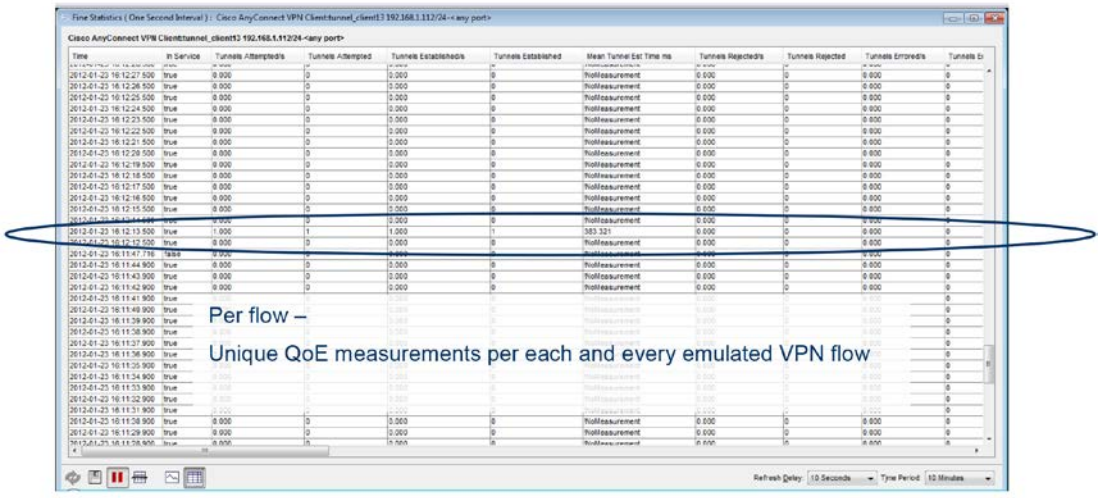

Figure 7 : Real time analysis is pivotal in understanding performance during live test runs. An example measurement use case - many security devices enable re-establishment features or connection termination on the fly, this dynamic behavior can only be captured through real time analysis.

The performance validation methodology can be further enhanced by adding in further network elements. Devices such as various switches/routers which will add latency and their inclusion will further add to the tunnel establishment time.

When assessing a secure VPNs performance it's always worth noting the available bi-directional bandwidth and the overhead associated with secure VPN management packets.

The following are a valid set of metrics for both client and clientless type VPNs, the outlined metrics should provide performance insight on the secure VPN tunnel connection. These VPN specific metrics are readily available within TeraVM.

Secure VPN Metrics	Description
In service	Client active as part of test
Out Tunnel IP Bits/s	Number of outgoing IP bits/second in tunnel
Out Tunnel IP Packets/s	Number of outgoing IP packets/second in tunnel
In Tunnel IP Bits/s	Number of incoming IP bits/second in tunnel
In Tunnel IP Packets/s	Number of incoming IP packets/second in tunnel
Mean Tunnel Est Time ms	Time tunnel connected
Tunnels Attempted/s	Number of attempted tunnels per second
Tunnels Attempted	Number of attempted
Tunnels Established/	Number of established tunnels per second
Tunnels Established	Number of established tunnels
Tunnels Rejected/s	Number of rejected tunnels per second
Tunnels Rejected	Number of rejected tunnels
Tunnels Erred/s	Number of erred tunnels per second
Tunnels Erred	Number of erred tunnels
Tunnels Completed/s	Number of completed tunnels per second
Tunnels Completed	Number of completed tunnels
Out CSTP Control Frames	Number of CSTP Out control frames
In CSTP Control Frames	Number of CSTP In control frames

CDTP Out Packets/s	Number of CDTP Out packets per second
CDTP In Packets/s	Number of CDTP In packets per second
DTLS Sessions Attempted/s	DTLS Sessions attempted per second
DTLS Sessions Accepted/s	DTLS Sessions accepted per second
DTLS Sessions Errored	DTLS Sessions errored per second
DTLS Fall Backs	DTLS Number of fall backs
DTLS Fall Forwards	DTLS Number of fall forward
DTLS DPD Failures	DTLS Number of DPD failures
TLS DPD Failures	TLS Number of DPD failures

Table 1 : Sample VPN SSL/TLS tunnel performance metrics

3rd Party VPNs	Description
Client Out Tunnel IP Bits/s	The number of bits/second sent out by this Client.
Client Out Tunnel IP Packets/s	The number of packets/second sent out by this Client.
Client In Tunnel IP Packets/s	The number of packets/second received in by this Client.
Client In Tunnel IP Bits/s	The number of bits/second received in by this Client.
Client Tunnels Attempted/s	The number of attempted tunnels/second.
Client Tunnels Attempted	The number of attempted tunnels.
Client Tunnels Established/s	The number of tunnels/second established by this Client.
Client Tunnels Established	The number of tunnels established by this Client.
Client Tunnels Errored/s	The number of tunnels/second that failed to be established by this Client..
Client Tunnels Errored	The number of tunnels that failed to be established by this Client.
Client Tunnels Rejected/s	The number of tunnels/second that were rejected for this Client.
Client Tunnels Rejected	The number of tunnels that were rejected for this Client.
Client Tunnels Completed/s	The number of tunnels/second that were completed by this Client.
Client Tunnels Completed	The number of tunnels that were completed by this Client.
Client Total Time Tunnel Established (ms)	The time in ms for which the tunnel was established.
Out Tunnel Control Packets	The number of Tunnel Control Packets sent by the client.
In Tunnel Control Packets	The number of Tunnel Control Packets received in by the client.
IPsec SAs Created	The number of IPsec SAs created by this client.
IPsec SAs Deleted	The number of IPsec SAs deleted by this client.
IPsec SA Failures	The number of IPsec SA failures detected by this client.
In IPsec Bit/s	The number of IPsec bits/second received by the client.
In IPsec Packet/s	The number of IPsec packets/second received by the client.
Out IPsec Bit/s	The number of IPsec bits/second sent out by the client.
Out IPsec Packet/s	The number of IPsec packets/second sent out by the client.
IPsec Bad SA Errors	The number of Bad SA Errors detected by this client.
IPsec Auth Errors	The number of IPsec Authorisation errors detected by this client.
IPsec Decrypt Errors	The number of IPsec Decrypt Errors detected by this client.
IPsec Other Errors	The number of all other IPsec Errors detected by this client.

Table 2 : Sample IPSec 3rd party VPN metrics

Using a consistent set of metrics as outlined, it's possible to review a number of secure VPN protocols and settings.

Reverting to the initial validation suggestions, Figure 8 compares the benchmark application throughput/ connection rate established versus the limitations applied when a security appliance is in the mix.

Some noteworthy application metrics are:

- Connections established per second
- Download speeds

A sample use of this low level test is to compare the various security devices from vendors to determine performance impact.

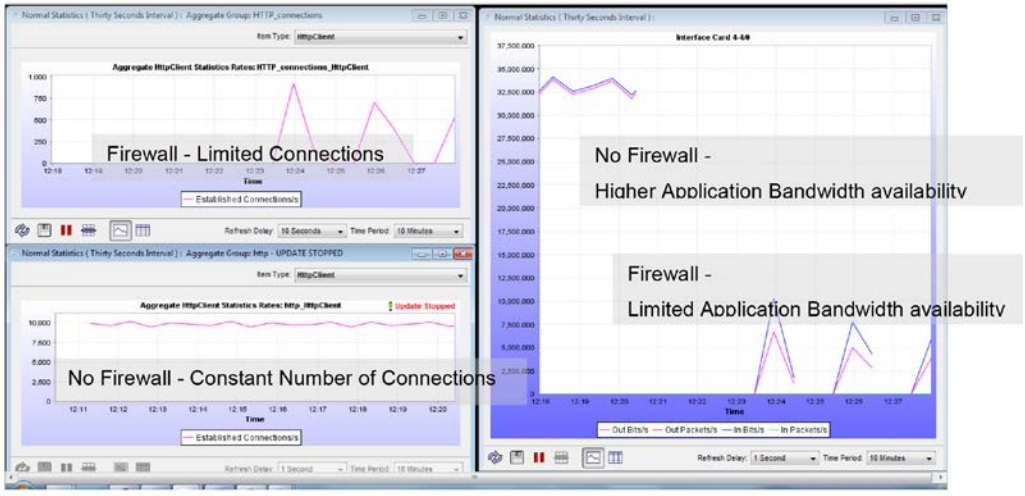

Figure 8 : SSL enabled tunnel performance versus application throughput rate outside of the secure tunnel.

Assessing performance of varying the IPSec algorithm formulae

When assessing the various vendors implementation of secure VPNs, it is worth considering the impact that the IPSec proposal and transform types used in the negotiation has on the application quality. For example -

- Proposal: IKEv2
- Encryption : AES 256
- Pseudo-Random Function (PRF) : SHA-1
- Diffie-Hillman : 2

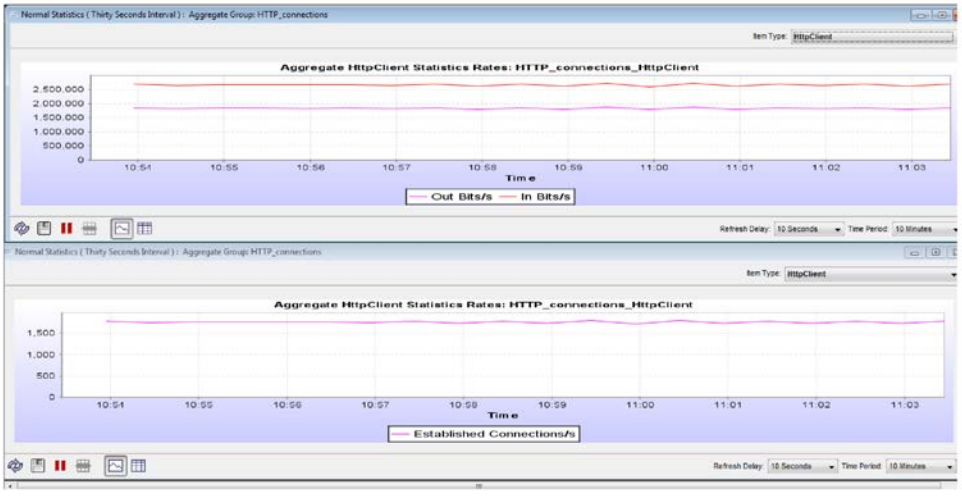

Figure 9: Aggregate view of performance, showing established number of connections over an IPSec (256 bit encryption) tunnel.

By using an emulation system such as TeraVM it's easier and quicker to make updates or changes to all the stateful endpoints with a common value. This enhancement saves time when evaluating the various security devices. Continuing the sample performance validation use case, the next step is to decrease the encryption levels to 128 bits, so the IPSec proposal and transform now look as follows:

- Proposal: IKEv2
- Encryption : AES 128
- Pseudo-Random Function (PRF) : SHA-512
- Diffie-Hillman : 2

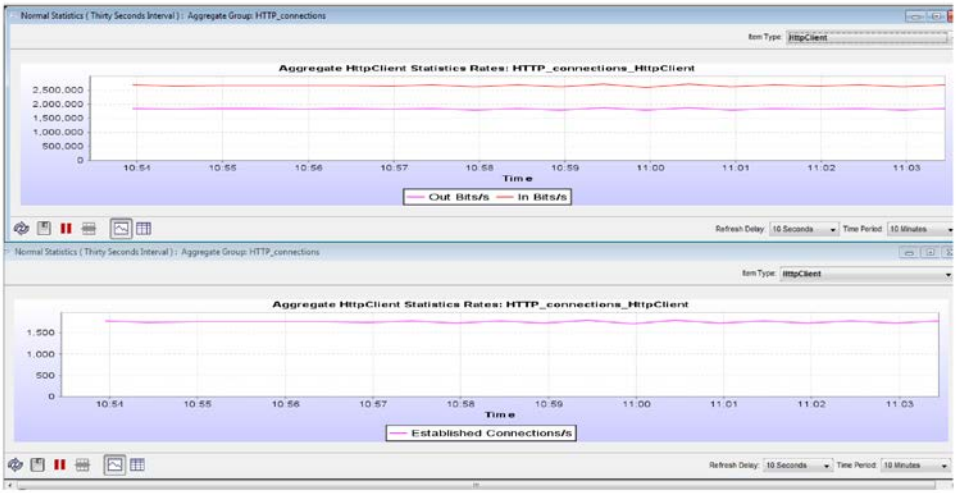

Figure 10 : Security algorithm has been changed for IPSec performance to 128 bit encryption.

TeraVM makes it possible to compare any number of IPSec configurations, enabling users analyze in detail the performance impact a security appliance has on the individual applications that need to be secured. The purpose should be to find a suitable configuration for the individual applications of voice, video and data.

Latency sensitive application performance inside secure tunnels

This next validation use case measures performance of delay sensitive applications such as voice in the secure VPN. When it comes to testing the performance or quality of the voice there are a multitude of things to cover, such as device configuration i.e. the buffer sizes, device registration performance and inevitably the codec types.

The aim is to determine QoE with a wide variety of device configurations and voice codec settings across a number of secure VPN types. This type of testing ensures that the secure VPN device administrator understands how varying settings can impact the QoE on the delay sensitive application.

Through emulation with TeraVM users can easily configure multiple endpoints with varying buffer sizes, using multiple codecs all running in a single test case. By using a per flow tool such as TeraVM it's further possible to choose from a wide variety of codecs including G.729, GSM or even implement a custom codec. So in a single test case all variants of codec can be assessed saving valuable time.

Building on from the previous test methodology the simple test case is to examine a single voice call and use the quality scoring as a benchmark.

A globally accepted measurement for voice quality is Mean Opinion Score (MOS) with a range 0-5, where 5 is superior quality and 1 poorest. Shown below is the MOS achieved for the application in the secure tunnel. Compared with the benchmark, the high score indicates that the security appliance is having a minimum impact on quality.

The configuration on the security device is configured as follows:

- Proposal: IKEv2
- Encryption : 3-des
- Pseudo-Random Function (PRF) : SHA-1
- Diffie-Hillman : 2

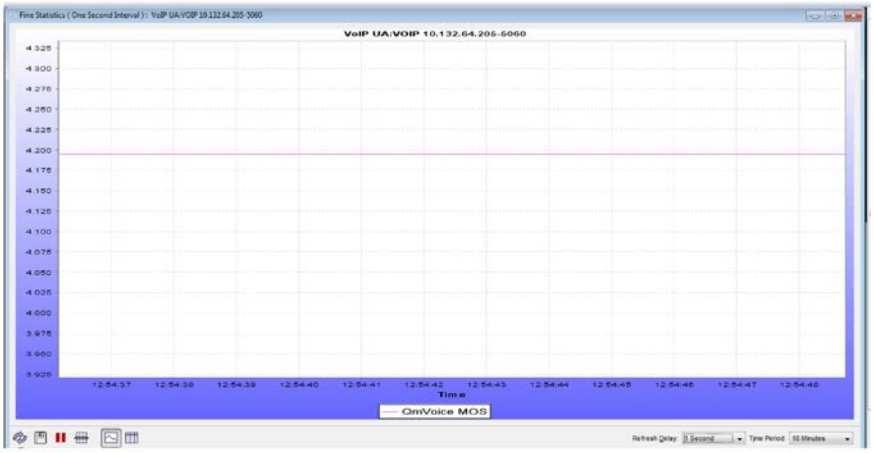

Figure 11 : Near flawless delivery of a latency sensitive application in a secure VPN

Impact of packet loss in secure VPNs

Next validation methodology is to examine the impact that a lossy network with various rates of dropped packets has on both the secure tunnel and unsecure application flows.

The following test results show the impact of 5% and an extreme condition of 10% packet loss on unsecure and secure applications.

Figure 12: Using TeraVM's per flow analysis it's possible to compare VoIP calls with packet loss on both secure versus unsecure connections.

Clearly packet loss has an impact on any delay sensitive applications. However, from the above figure it is possible to see that packet loss has a greater impact on the applications that are running in the secure tunnels.

Conclusion

How data is stored and the encryption mechanisms deployed is only part of the security protection conundrum. Many Enterprises fail to see that another real and very vulnerable component of their data security is the path to it and how end-users access the data.

Many Enterprises now offer remote employees the ability to access data live in their networks via secure VPNs. But how secure is a secure VPN, why choose one form over another? The application note primarily looks into application performance inside secure VPNs and in particular looks at client versus clientless mode arguments.

Selection of the correct VPN protocol is critical in protecting an organisation's sensitive data. However, for whatever reason, if a single application's performance is poor over the secure tunnel, the end-user will most likely shy away from using the secure tunnel. The obvious consequence is that the data is now transmitted in an unsecure way. It's imperative that all applications or services been used in the secure tunnel are fully tested for performance limitations and that IT security teams are aware of any quality issues or scale issues on a per end-user basis.

A breach in security no matter how small is a significant event, in the application note a sample security breach of a Distributed 'Denial of Service' against an internal web server was implemented. The scenario attack happens from inside a secure tunnel! A small, short lived attack had an immediate impact on service delivery. This poor experience or the persistent hampering of the external employee to access data will result in alternative transmission techniques being used to communicate the data. As part of an overall test strategy of a secure access applaince, determining performance with legal and illegal flows is a must.

The price point for high end super computer platforms is as low as the cost of a gaming console, plus the increase of personal information to social networks is a worry as for between the two; cracking passwords and accessing secure networks has never been easier. The trade-off for increased security should not be at the expense of the application performance.

The challenge in maintaining application quality in secure VPN tunnels is greatly simplified by using TeraVM. Security deployments can be successfully validated for the strength of their security attack mitigation strategy plus the overall performance of the multitude of applications that will be carried in the secure VPN tunnels. With per flow test solutions such as TeraVM, with the widest range of applications and third party secure VPNs, security can be deployed with confidence.

VIAVI Solutions

To reach the VIAVI office nearest you,
visit viavisolutions.com/contacts

vpn-perf-validation-wp-wir-nse-ae
30187561 900 1118