Netgate TNSR

Secure Networking Software Platform

Performance Overview

TNSR Performance Overview

Table of Contents

Executive Overview

Netgate TNSR provides the fastest, lowest cost, and most flexible secure networking software platform for today's savvy IT infrastructure providers.

Historically, high-speed packet processing required ASIC/FPGA-centric hardware platforms to achieve performance goals. TNSR breaks that paradigm. TNSR can route up to 1 Tbps (14+ Mpps) in software running on commodity CPUs - which slashes costs, eliminates vendor lock-in, reduces IT infrastructure complexity, and enables elastic scale.

TNSR's engine - based on FD.io's Vector Packet Processing (VPP) - enables extreme packet processing relative to the inherent limits of kernel-based processing.

As this paper will show, TNSR is capable of astounding packet processing performance for the dollar. Applications ranging from simple routing to complex packet processing can now be fulfilled for a fraction of price of previous commercial solutions - opening up a world of secure networking opportunities that would otherwise have been out of reach for all but the largest of enterprises and service providers. TNSR's engine - based on FD.io's Vector Packet Processing (VPP) - enables extreme packet processing relative to the inherent limits of kernel-based processing.

Six Performance Use Cases

Let's take a close look at performance by comparing a set of progressively demanding test cases:

Use Case	Description
1	TNSR Large Packet Routing
2	TNSR Small Packet Routing
3	TNSR IPsec Tunneling / No Encryption
4	TNSR IPsec Tunneling / AES-CBS-128 Encryption
5	TNSR IPsec Tunneling / AES-GCM-128 Encryption
6	TNSR / QuickAssist Technology (Intel® QAT)

Test Case 1 - TNSR Large Packet Routing

At a high level, let's assume a user wants to fill a 10 Gbps link with 1500 byte packets, i.e., a large file download use case. This would require software capable of processing 812,743 packets per second. **Using large packets, TNSR can fill a 10 Gbps with ease.**

Test Case 2 - TNSR Small Packet Routing

Now let's consider a much more demanding use case: one where we need to support hundreds or thousands of individual sessions, each requiring 64 byte "application-sized" packets - but still unencrypted. With 64 byte packets (84 bytes on the wire) we will need to process 14,880,952 packets per second to fill a 10 Gbps link. Using small packets, TNSR fills a 10 Gbps link with aplomb - on a single core of a single Intel Xeon-class processor.

By extrapolation, if a secure networking application requires, say, 100 Gbps routing, the math says we will need to process 14,880,952 packets per second to fill the wire. **With TNSR's ability to pump 14 Mpps per core, we will need a mere 10 CPU cores to achieve 100 Gbps.** Until recently, this would have been unheard of in software on commodity hardware.

Test Case 3 - TNSR IPsec Tunneling / No Encryption

When a user has one or more Virtual Private Clouds (VPCs) that need interconnection to operate as a seamless whole, TNSR is the answer. Simple management through a RESTful API - which integrates to virtually any cloud orchestration management solution - makes installation, configuration, and operation very straightforward.

Test Configuration

In order to minimize the effects of network latency or congestion on throughput measurements, an IPsec tunnel was established between two instances running TNSR in the same region and availability zone. The instances reside in two distinct VPCs. Each VPC has two subnets - one public, one private. The public subnet is used to terminate the IPsec tunnel and forward ESP packets to the IPsec tunnel peer. Management traffic (SSH sessions to configure and monitor the IPsec tunnel) for the IPsec tunnel endpoints is also carried on the public subnet. The private subnet is the source of traffic to be forwarded over the tunnel. An instance on the private subnet of each VPC sends packets to the instance which terminates the IPsec tunnel. Packets are subsequently forwarded through the tunnel and received by the far end instance on the private subnet of the opposing VPC.

The following diagram provides a high-level illustration of the test environment configuration:

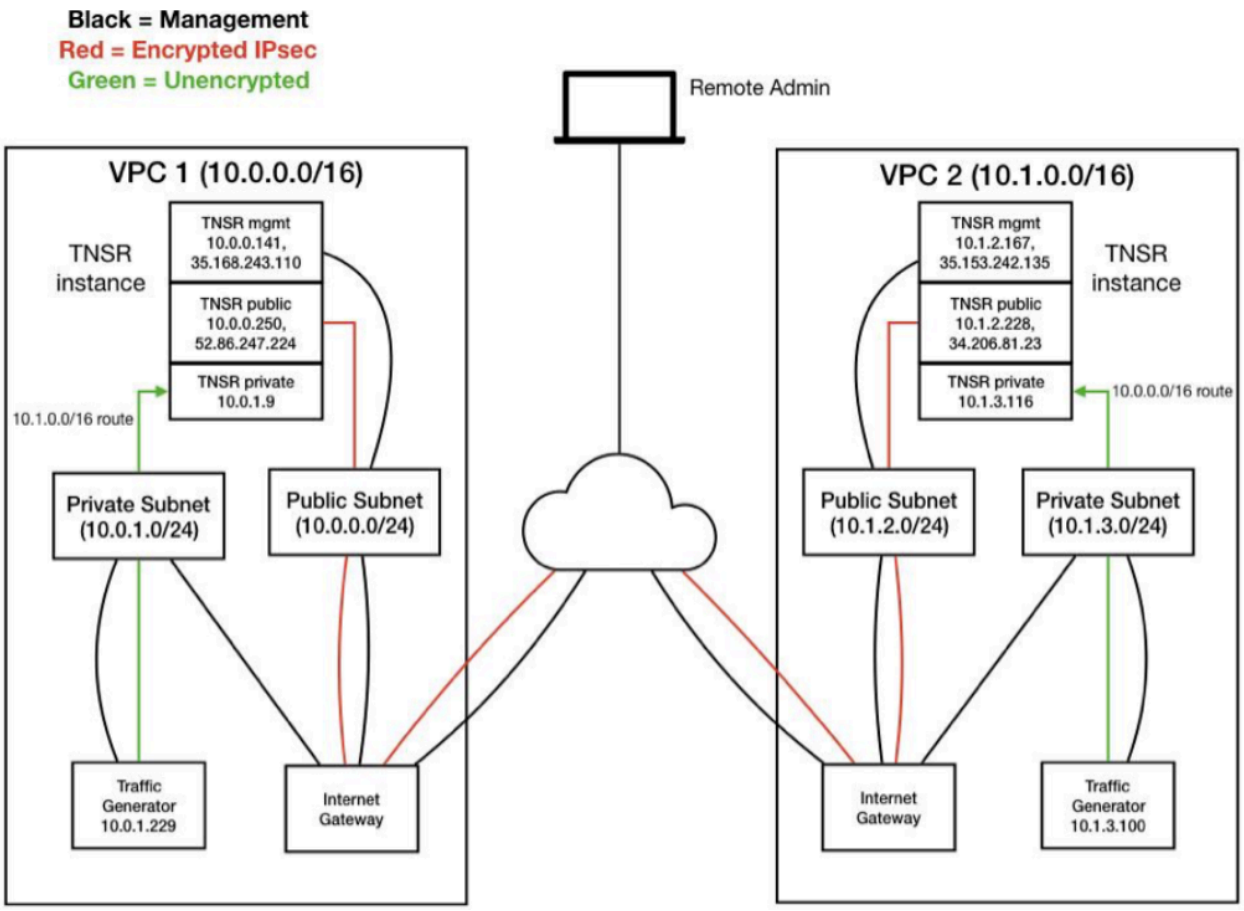

Tests are performed on the AWS Marketplace with private EC2 instances. Each IPsec endpoint and traffic generator host runs on a C5.xlarge instance running CentOS 7.4. Each instance is equipped with four CPUs and 8 GB of RAM and Elastic Network Adapters (ENA).

Note: Only a single core was used for testing.

Both TNSR IPsec endpoints were configured to establish an IPsec tunnel using IKEv2 with a pre-shared key. The IKE security association was configured to use AES-CBC-128 HMAC-SHA1. The ESP security associations were tested both with AES-CBC-128 HMAC-SHA1 and AES-GCM-128-ICV16. The TNSR instances ran a single thread for packet processing.

The traffic generator hosts used an open source bandwidth measurement named iperf3 to send traffic. Tests were run both with a single stream and with 4 streams. The traffic generator hosts had their MTU adjusted down to 1500 (from the default

value of 9001). Tests with iperf3 were invoked with a flag that set the TCP MSS to 1375 in order to ensure that each segment sent would not exceed 1500 bytes once the encapsulation overhead (ESP header, initialization vector, padding, integrity check value, outer IP header) is added.

A baseline measurement was taken between the two IPsec endpoints to determine the maximum bandwidth that could be sent from one host to the other with no encryption or encapsulation applied - which measured at 4.79 Gbps:

Source	Destination	Encryption	Streams	Measurement
VPC 1 IPsec endpoint	VPC 2 IPsec endpoint	None	1	4.79 Gbps
VPC 1 IPsec endpoint	VPC 2 IPsec endpoint	None	4	4.79 Gbps
VPC 2 IPsec endpoint	VPC 1 IPsec endpoint	None	1	4.79 Gbps
VPC 2 IPsec endpoint	VPC 1 IPsec endpoint	None	4	4.79 Gbps

The measurements taken by iperf3 use the amount of data sent in the TCP payload to calculate throughput. The 32 byte TCP header (standard 20 byte header plus 10 bytes for optional field containing timestamps and 2 bytes to pad optional fields to a multiple of 4 bytes) and 20 byte IP header on each packet are not included in the calculation. If 52 bytes from each 1500 byte packet are considered overhead that is not included in the measurement, the maximum result that iperf3 could achieve on a 5 Gbps link would be approximately 4.83 Gbps.

Note: The 5 Gbps link constraint per EC2 instance is imposed by Amazon.

Test Case 4 - TNSR AES-CBC-128 HMAC-SHA1 Encryption

Applying AES-CBC-128 encryption reduces throughput anywhere from 4.79 Gbps with no encryption to 3.30-3.62 Gbps for single stream processing, a decline of roughly 31%:

Source	Destination	Encryption	Streams	Measurement
VPC 1 traffic generator	VPC 2 traffic generator	AES-CBC-128 HMAC-SHA1 1	1	3.62 Gbps
VPC 1 traffic generator	VPC 2 traffic generator	AES-CBC-128 HMAC-SHA1 1	4	4.58 Gbps
VPC 2 traffic generator	VPC 1 traffic generator	AES-CBC-128 HMAC-SHA1 1	1	3.30 Gbps
VPC 2 traffic generator	VPC 1 traffic generator	AES-CBC-128 HMAC-SHA1 1	4	4.59 Gbps

Test Case 5 - TNSR AES-GCM-128 ICV16 Encryption

Using the newer, more efficient GCM encryption schema, we observe a slight gain in throughout relative to CBC —based encryption:

Source	Destination	Encryption	Streams	Measurement
VPC 1 traffic generator	VPC 2 traffic generator	AES-GCM-128 ICV16	1	3.93 Gbps
VPC 1 traffic generator	VPC 2 traffic generator	AES-GCM-128 ICV16	1	4.59 Gbps
VPC 2 traffic generator	VPC 1 traffic generator	AES-GCM-128 ICV16	1	3.60 Gbps
VPC 2 traffic generator	VPC 1 traffic generator	AES-GCM-128 ICV16	1	4.59 Gbps

Note the additional overhead from Encapsulating Security Payload (ESP) includes:

- 20 bytes for an outer IP header
- 8 bytes for an ESP header
- 2 bytes for padding length & next header type
- 16 (AES-CBC) or 8 (AES-GCM) bytes for an initialization vector
- 12 (HMAC-SHA1) or 16 (AES-GCM) bytes for an integrity check value

The total extra overhead is 58 bytes (AES-CBC HMAC-SHA1) or 54 bytes (AES-GCM). Thus, the maximum throughput possible using iperf3 on a 5 Gbps link is 4.63 Gb/s for AES-CBC-128 HMAC-SHA1, and 4.65 Gb/s for AES-GCM-128 ICV16.

Test Case 6 - VPP + QAT

As shown above, even with the power of VPP brought to bear, traffic encryptIon still imposes a severe throughput tax. Intel's QuickAssist Technology (Intel® QAT) can be used to enhance security and compression performance in cloud, networking, big data, and storage applications. Netgate measured the impact of offloading crypto-processing to a small form factor cryptographic accelerator card running QAT.

Test Configuration

The test setup was a classic "4 box" architecture: Source-Router-Router-Sink. Each router was a water-cooled, i7-6750X, 10 core, 3.5 GHz-based system, equipped with a dual-port 40Gbps ix710 card, and a CPIC card - effectively intended to

represent 'ultimate white boxes'. Source and Sink instances were more pedestrian quad-core Xeon boxes, each with a 40 Gbps NIC.

In the end, neither high core count nor water-cooled routers were required for this level of testing.

Multiple test scenarios are shared below. The clearest comparative to prior test scenarios covered earlier in this paper is the last line, VPP packet processing accompanied by QAT crypto processing. **In this case, we observed 32.730 Gbps of throughput, a 7x improvement over VPP with no crypto offload.**

Packet Processing	Crypto Processing	Encryption	Integrity Algorithm	Iperf TCP1500 (Gbps)
VPP	QAT	AES-CBC-128	SHA1	8.740
VPP	VPP Native (Open SSL 1.0.1)	AES-CBC-128	SHA1	2.030
VPP	VPP AESNI MB	AES-CBC-128	SHA1	7.420
VPP	VPP AESNI GCM	AES-GCM-128-16		13.700
VPP	VPP QAT	AES-CBC-128	SHA1	32.680
VPP	VPP QAT	AES-GCM-128-16		32.730

Summary

Irrespective of the size or location (cloud, premises, hybrid) or complexity of your organization's IT infrastructure, the one constant is the never ending need for increased network throughput, and under increasingly demanding traffic mix and security conditions. Netgate's TNSR Secure Networking Software Platform breaks the mold and delivers the fastest, lowest cost, and most flexible secure networking software platform available for today's savvy IT infrastructure needs.