Scaling Suricata to 10Gbps and beyond

Stephen Patton
Senior Product Manager
Telesoft Technologies
Blandford, UK

Jenna Williams
Senior Engineering Manager
Telesoft Technologies
Blandford, UK

Lee Walker
Senior Test Engineer
Telesoft Technologies
Blandford, UK

I. Abstract

The continuing growth of data traffic and the need to protect data and its users is driving the need for ever higher performance cyber threat detection and protection. One threat monitoring capability commonly used is Intrusion Detection System (IDS), available as an appliance or as Open Source software.

This document provides a comparative performance test of a high performance Open Source IDS that is gaining popularity, Suricata, using Telesoft hardware acceleration and Suricata using a standard Intel NIC, with the purpose of understanding how a cyber security integrator could scale a Suricata based IDS beyond 10Gbps throughput.

Both configurations were benchmarked to operate past 8Gbps with the Telesoft MPAC Security card showing no packet loss and a 40% saving in CPU core load at data rates beyond 32Gbps depending on loaded ruleset. Under the same test conditions, a standard 10GbE Intel NIC tuned to support Suricata was found to drop packets as throughput approached 8Gbps.

II. About Telesoft MPAC Security

The Telesoft MPAC Security IDS accelerator works 'out-of-the-box' with Open Source Suricata and offloads CPU intensive signature scanning by looking for rule based patterns across packets captured from 4 x 10GbE. This, coupled with load balancing across up to 64-CPU cores gives significant system performance gains allowing a single server to monitor 2 bi-directional 10G links with zero packet loss.

III. About Suricata

Suricata is a high performance Network IDS, IPS and Network Security Monitoring engine. Open Source and owned by a community run non-profit foundation, the Open Information Security Foundation (OISF). Suricata is developed by the OISF and its supporting vendors.

IV. Test Conditions

In order to measure the performance of the Telesoft MPAC Security card and compare against a standard Intel NIC, two systems were setup as shown in Figure 1, using identical servers commonly used to build a high performance IDS (Dual Socket, HP DL380 Gen 9 containing two Intel Xeon E5-2650 v3 CPUs, each with 10 Cores). One System Under Test (SUT) was fitted with a standard Intel NIC (82599ES 10Gb) and a second SUT fitted with a Telesoft 4x10GbE MPAC Security PCIe card.

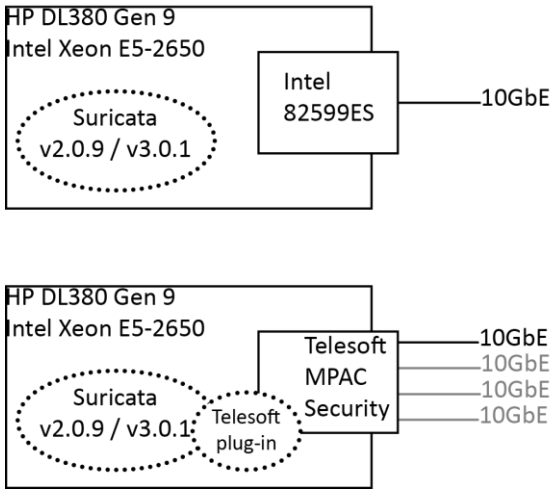

Figure 1 – Test environment

Surciata IDS v3.0.1 was installed on each SUT, configured in worker mode with 16 worker threads with CPU affinity, locking each thread to an individual CPU core. Telesoft contributed code that makes use of the hardware accelerated features added to the system fitted with the Telesoft 4x10GbE MPAC Security PCIe card.

The system was connected to a traffic simulator configured to generate a mix of 90% TCP and 10% UDP as commonly found on mixed Internet traffic. UDP traffic contained randomized ASCII payload with an average packet size of 750 bytes. TCP traffic was synthesized from a streaming video session with an average packet size of 930 bytes.

V. Acceleration functionality

Software analyses the ruleset loaded by Suricata for content tags containing ASCII or hexadecimal strings of interest. These content signatures are then crunched down to 8000 common substrings which are then loaded into the signature matching engine of the MPAC Security Card.

The MPAC Security card scans every offset of every packet for signature matches. Packets are delivered to Suricata via a Telesoft implementation of PF_Ring, where every packet has a metadata header prepended. This metadata contains the signature id and offset of every match found by the hardware.

Telesoft have optimized the Open source Suricata code for integration with the Telesoft enhanced packet format. This accelerated functionality is enabled by setting the mpm-algo in the suricata yaml file to mpac. When enabled, Suricata loads a mapping file containing the signature ids assigned by the MPAC card. This mapping is then used within the Search functions of the Multiple Pattern Matcher allowing the software to perform final validation of the signature matches at the presented offsets within each packet. This ability to jump to specific offsets throughput the packet instead of searching at every byte, provides the CPU acceleration detailed in this document.

VI. Packet Capture and Delivery

To benchmark raw packet capture, traffic was routed to the SUT fitted with the Telesoft MPAC card running Suricata with no rules configured. Traffic was load balanced across all four 10GbE ports. Packet loss and average CPU-core utilization were measured as below.

Figure 2 shows the Average Worker CPU % for a total throughput of 8 to 32 Gbps with a TCP traffic profile, running Suricata v3.0.1 (no rules)

Rate/ Port (Gbps)	Total Rate (Gbps)	Packets Sent (M's)	Loss	Average Worker CPU
2	8	245.76	0.00%	11.02%
4	16	491.52	0.00%	16.86%
6	24	737.28	0.00%	23.50%
8	32	983.04	0.00%	33.34%

Figure 2 – Average Worker CPU % TCP traffic

Figure 3 shows the Average Worker CPU % for a total throughput of 8 to 32 Gbps with a UCP traffic profile, running Suricata v3.0.1 (no rules)

Rate/ Port (Gbps)	Total Rate (Gbps)	Packets Sent (M's)	Loss	Average Worker CPU
2	8	512	0.00%	12.54%
4	16	512	0.00%	19.17%
6	24	1536	0.00%	24.09%
8	32	2048	0.00%	31.19%

Figure 3 – Average Worker CPU % UCP traffic

The results above demonstrate that for both TCP and UDP, packets were delivered from the Telesoft MPAC Security card to Suricata at rates exceeding 32Gbps with no loss of data.

VII. CPU utilization and throughput with rule processing offload

To evaluate CPU utilization with signature scanning offload to the Telesoft MPAC Security, traffic load tests were conducted with different rulesets loaded. Packet loss and CPU utilization were measured for each. The number of rules for each set loaded is shown in Figure 4.

Rule set	Number of rules
Trojan Rule subset	4105
Emerging Rules subset	8180
Open Rules	17480
Emerging Threat Pro Rules	32451

Figure 4 – rulesets used

Figure 5 shows average CPU core utilization and packet loss for the Trojan and Emerging rules.

			Trojan Rules		Emerging Rules	
Rate/Per Port Gbps	Total Rate Gbps	Packets Sent (M's)	Loss	Average Utilization	Loss	Average Utilization
2	8	675.84	0.00%	13.1%	0.00%	14.8%
4	16	675.84	0.00%	21.9%	0.00%	24.8%
6	24	675.84	0.00%	30.8%	0.00%	34.7%
8	32	675.84	0.00%	34.4%	0.00%	40.1%

Figure 5 - utilization (Trojan and Emerging)

Figure 6 shows average CPU core utilization and packet loss for the Open and ET Pro rules.

			Open Rules		ET Pro Rules	
Rate/Per Port Gbps	Total Rate Gbps	Packets Sent (M's)	Loss	Average Utilization	Loss	Average Utilization
2	8	675.84	0.00%	33.7%	0.00%	26.0%
4	16	675.84	0.00%	62.2%	0.00%	46.3%

Figure 6 – utilization (Open and ET Pro)

Hence with the system as specified, a 16Gbps Suricata based IDS can be constructed in a single 20-core server running Emerging Threat Pro. Rates up to and beyond 32Gbps can be achieved by reducing configured rules or increasing thread count.

VIII. CPU utilization and throughput comparison v Intel

A comparison of the Telesoft MPAC Security to Intel NIC was performed by tuning the Intel card as recommended by the OISF Suricata User Guide 'High Performance Considerations'.

A mix of TCP and UDP traffic was routed to each system as before, overlaid with background traffic containing known threats in order that each SUT generated alerts.

Two tests were performed, one using the Emerging Threats ruleset and the other using Emerging Threats Pro ruleset, both at rates from 3Gbps up to 8Gbps. Beyond 8Gbps packet loss on the Intel card caused the system with the Intel NIC to be unusable. The peak CPU usage was measured for each test condition.

Figure 7 shows the average CPU thread utilization for the measured throughput when running the emerging ruleset and Figure 8 the average CPU thread utilization for the measured throughput when running the ET Pro ruleset.

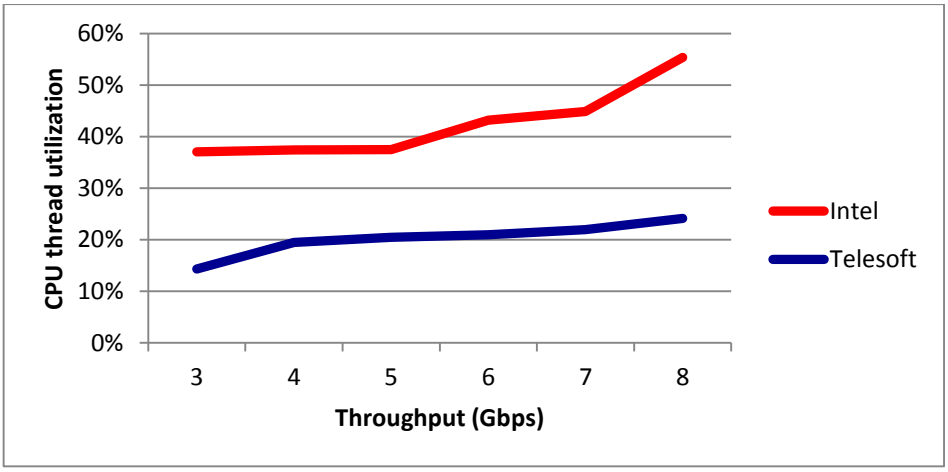

Figure 7 – CPU thread utilization vs Throughput Emerging Threats (8180 signatures)

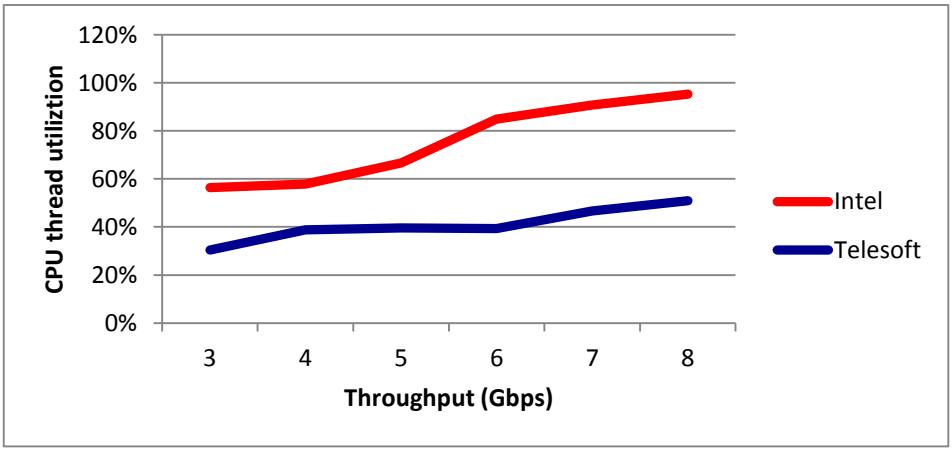

Figure 8 – CPU thread utilization vs Throughput ET Pro (32451 signatures)

The average number of alerts generated during test period was measured as 2300 for the Emerging Threats ruleset and 5880 for the ET Pro ruleset.

These results demonstrate an average of 40% or greater CPU thread utilization when using the Telesoft MPAC Security card. The Telesoft card operates past the Intel 8Gbps limit up to and past 16Gbps on a 20-core system.

Early testing with a 72 CPU core system has shown on average a throughput of 1Gbps per CPU core past 24Gbps with the ET Pro ruleset loaded. Further ongoing testing and optimization is expected to enable throughput closer to 40Gbps to be achieved.

IX. Conclusions

The results obtained demonstrate that by fitting a Telesoft MPAC-Security card to a 20-core system running Suricata 3.0.1, throughput in excess of 16Gbps can be achieved when working with the full Emerging Threats Pro rule set, and beyond 32Gbps if the rule set can be reduced. An average CPU utilization drop of 40% and higher was seen in all test cases, allowing systems to process higher throughputs or run additional processing without impacting IDS performance.

The same 20-core server fitted with a standard Intel NIC can achieve throughputs of up to 8Gbps. Beyond these rates, packets are lost and hence detection accuracy was impacted. CPU utilization at 8Gbps with the full ET Pro ruleset approached 100%, so to process data at 8Gbps and beyond more CPU cores would be required.

For more information about the Telesoft Technologies MPAC-Security 4x10GbE Suricata accelerator, go to: http://telesoft-technologies.com/technologies/accelerator-cards/mpac-security

telesoft
Cyber Defense and Network Visibility