The Quarterly Magazine for Digital Forensics Practitioners

DIGITAL FORENSICS / MAGAZINE

COMPETITION!
Win 3 Digital Forensics books from Syngress

ISSUE 09
NOVEMBER 2011

INSIDE

BIG BROTHER FORENSICS

Chad Tilbury takes a look at the rise of Geo Location data and how geo-artifacts can add a crucial dimension to investigations

9 772042 061110 04
Issue 9 / £11.99 TR Media

SIFTER10 PROBES

A successful new approach to building Deep Packet Inspection devices on High-speed networks

by Mark Osborne

/ ADVANCED

This article describes the current strategies for performing Deep Packet Inspection functions for security and network management on high-speed networks; it illustrates the significant drawbacks of these methods. The paper introduces CyberSifts' HANAC architecture and the patented massively parallel search technology Dynamic Parallel Inspection.

You don't need to be a technologist to realise that network usage is escalating rapidly. Smart phones, mobile broadband, WiFi and high-speed, even fibre-optic broadband in every home is the cause of this network phenomenon. This massive demand for cheap network capacity has caused a technology convergence away from expensive, legacy technology so that most large backbone networks within network providers, telecoms companies or large corporate and utilities use TCP/IP on high-speed 10 Gigabit Ethernet (Gbe).

However, as these technologies have become more ubiquitous, so has their abuse, whether from malware/spam/phishing, DDoS attacks, terrorism and or copyright infringement. In the last year or so, there has been an increasing volume of regulation both in US and Europe encouraging network providers to counter these threats to the network economy.

But many practitioners are finding that at the higher network speeds of 10Gb/s, 40Gb/s or even 100 Gb/s, the traditional PC server based SPAM/AV UTM appliances, Web filters and IDS simply can't keep up; they are just too slow.

/ THE PROBLEM

How can this be when 10Gb/s routing and switch hardware is so plentiful and cheap? The answers are the same old story for IP based networking; it is easier to send (or route) a packet in the IP/Ethernet world than it is to secure it. Correspondingly, the devices doing this security work need a lot more muscle.

Most of the security tasks described above, like SPAM protection or Intrusion monitoring require Deep Packet Inspection technology. Fundamentally, this means processing every byte of the transmitted packet and comparing it to a database of a 1000 security vulnerabilities – whilst a router typically has to compare the four bytes of the packet address to usually a much smaller routing database.

Typically, PC based software security products like Web filters or IPS work at speeds of up to 1 Gb/s. After this point, the supervising operating system uses a massive amount of resource simply moving data packets from the NIC into memory where the application can process it. Yet the sad truth is the majority of these "costly packets" are of no interest to the application. In fact, it is highly likely the first act of the system will be to read the packet and then immediately reject it. For example, if our application is a simple web filter, it will only need to process HTTP GET requests on TCP port 80. On an average 10Gb/s link over 90% of the packets will be outside this population. Unfortunately, our application has to sort through this majority of uninteresting packets sequentially to find the packets that meet its processing needs. In doing this, at 10 Gb/s, the typical operating system and application collectively will drop the majority of the packets.

/ STRATEGIES FOR 10GB/S OR 40GB/S

To overcome this problem, vendors usually adopt one of two strategies:

- A Total Hardware solution;
- Use an Enhanced Network card;

/ THE TOTAL HARDWARE SOLUTION

The Total Hardware solution often can cope with the speed and volume of the traffic. Many vendors have developed hardware-based solutions to a number of security and management issues in the 10Gbe or 40Gbe space. Their disadvantage is that they tend to be very expensive as each of them is based on bespoke, unique ASIC architecture, and because of the burnt-in nature of this type of device they are hard to change. This makes them unsympathetic to modern applications as protocols and their exploitation (benevolent or malevolent) generate a fast moving environment.

This type of solution is effective but unattractive; it holds no utility advantage, as it is a "one box, one function" solution. Any CTO that invests in a security solution of this type is unlikely to be rewarded with better usage information (for example) as a reward.

/ ENHANCED NETWORK CARDS

Alternatively several vendors push a 10Gbe enhanced network card, these cards are highly advanced; using a variety of techniques such as polling & zero copy drivers, mutli-channel PCI interfaces, multiple DMA buffers and interface colouring. However, their main objective is to move as many packets from the wire into memory as possible. These can overcome in the short term many of the problems of operating at high speed.

Basically, these cards help the PC server cope with volume of traffic by shifting different traffics types to a number of distinct virtual network interfaces. As shown in figure 1, different processes then can be presented packets at a rate their programs can handle. However, over any period of time the task of balancing traffic

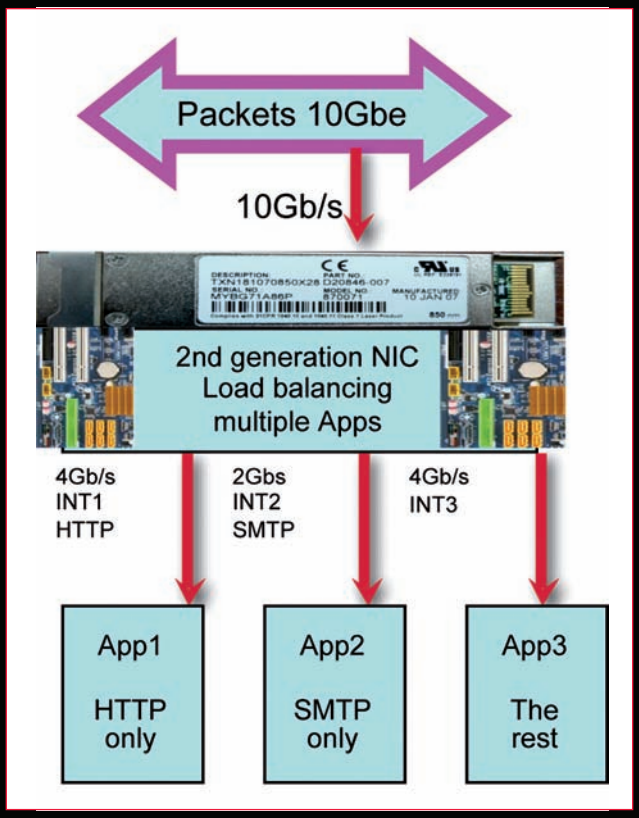

Figure 1. Interface Colouring as a solution

between these different interfaces and programs makes it clumsy and unmanageable

Over the last few years a more successful hybrid approach has been developed where packets can be blocked or forwarded completely within hardware. In other cases, where there is a need for sophisticated server based analysis software, our hardware reduces computational load by either passing only the selected packets to the application (based on layer 2-7 DPI) or by sharing the computational load by producing traffic metrics and counts on behalf of the application. In this way a platform has been developed that offers complex hardware functions to an application, in the same way the Unix kernel offers services to any application.

/ HANAC

The Sifter10 range of probes are advanced appliances combining powerful server technology, state of the art software and revolutionary Hardware Assisted Network Application Co-operation (HANAC) support. HANAC provides full Deep Packet Inspection (Layer2-7) plus Packet filtering, Counting/Classification and Redirection in hardware at full line-rate with an extremely low latency before any resource on the server platform is utilised.

Using this approach the probes break the paradigm of serialisation and allow bespoke packet-processing hardware to run in parallel with complex Intel CPU based applications. This can extend the useful life of your software assets or allow you to develop flexible Linux packet applications without the need to develop special hardware.

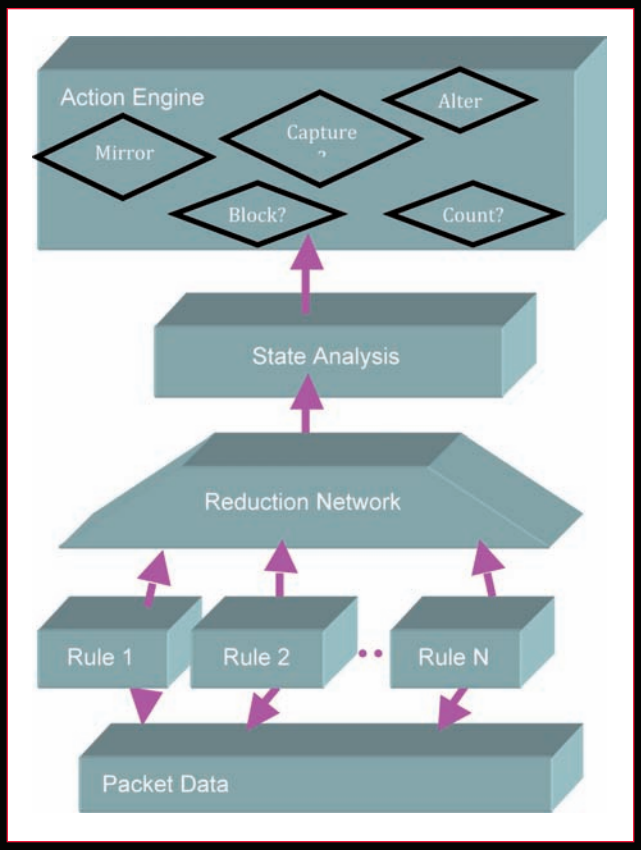

Figure 2. Data Stream Processing

/ THE N-TIER ARCHITECTURE

Most network security and network control applications are extremely parallel in their nature. For example URL filters which check a packet for a specified URL against a blacklist of 1000s or a SPAM filter, which examines an email for a list of 1000s of blocked addresses. None of these tasks need to be done in a serial manner (only split into one or two concurrent tasks), this has been forced upon us because of the way most general-purpose computers function. It works because general-purpose computers are relatively fast but with the emergence of faster networks, the relative speed advantage of general-purpose computers has been eroded. It would be much better if the comparison work were distributed simultaneously amongst numerous simple processors.

This is known as a Multiple Instruction Single Data (MISD) computational model. Using this paradigm, one data record is shared amongst multiple computational units, each executing different instructions on the same data. The HANAC architecture uses this model with a patented inspection process called Dynamic Parallel Inspection. This massively parallel processing technique manipulates data packets into 1024 bits units and distributes them inside multiple separate processors. Thus a large number (thousands) of simple execution units share the data and concurrently implement different packet matching operations.

In Figure 2, the data stream is concurrently presented to a number of execution units (Rule 1, Rule 2, through Rule n). Each unit is responsible for independently performing wire-speed packet processing and outputting a number of signals.

Each of the inspection rules, which are embedded in the execution units, can be changed dynamically. As new needs emerge, new rules can be written and pushed into the units. This can be done online, on the fly, or offline. In fact, these rules can be changed in a production system and are applied in less than 1/1000th of a second. During the application of new rules, the system will maintain all state and continue to apply all existing rules without interruption.

By splitting analysis rules into many discrete engines that can run on the same data in parallel, and by embedding these rules in the gates of an FPGA, we can achieve record-breaking inspection throughputs of 14.88 million packets per second.

By using true hardware separation from the action engines, there is isolation between action processing and signature inspection logic performance. This leads to identical performance, identical throughput, and identical latency with any traffic load, and under full use of the system's analysis policies.

The probes can also track state for each flow through the use of an external memory table. This memory table provides very high performance state memory management to handle up to 300,000 new flows per second (10x better than traditional firewalls), and up to 8 million concurrent stateful flows.

The power of HANACs' processing capability makes it a flexible and utility platform, with a number of deployment strategies that can be used individually for simplicity or combined for sophisticated, near intelligent applications.

/ DEPLOYMENT

The probe can be deployed in three typologies.

- Static in-line
- Passive monitoring
- Advanced Co-operative processing

/STATIC IN-LINE DEPLOYMENT

In this scenario, users deploy the technology mainly as an alternative to a hardware appliance, where the primary requirements are for speed and the filtering functions can be specified in a static rule set and there is no need for interaction with software on the host.

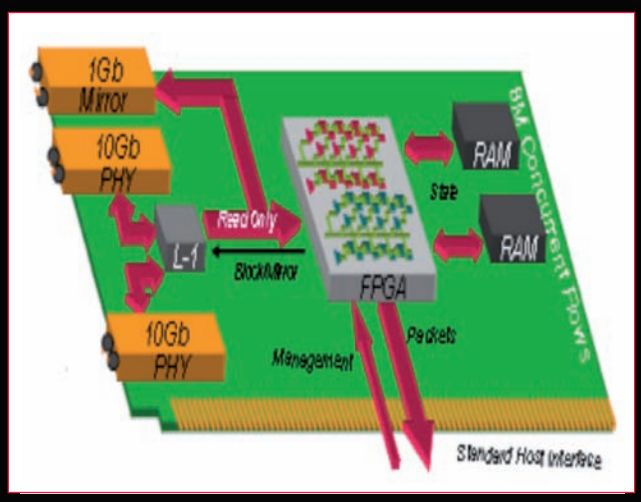

Figure 3. Layer 1 active bridge processing

The physical network interfaces (PHY) on the Sifter10 range are provided in the form of two full duplex 10 Gigabit pluggable XFP "sensing" modules. As shown in Figure 3, data signals received from one PHY are directly transmitted to the other PHY unchanged, creating a Layer 1 bridge with latency less than 1µ second. An FPGA "forwarding engine" controls this process and can block any packet based on the signature patterns.

When configured in this way, the probe is a "stealth mode" device, implemented entirely at the OSI layer 1. There is no MAC address, or need for any layer 2/3-topology changes – making implementation simple and detection practically unlikely.

As the device is implemented as a bump-in-the-wire, the card has been designed to withstand host failure; for example traffic forwarding will continue in the event of a hard drive failure or operating system panic. However, given the modern day requirement for five nines (99.999%) availability, we have developed implementation scenarios that take advantage of Linux clustering technology to provide uninterrupted service.

/ PASSIVE MONITORING

The features of HANAC are impressive, but they have been implemented in Linux as a normal network driver. This means that you don't need to learn a whole book of commands to use it; the standard ifconfig that is used for a normal Ethernet card works just fine (Note – there are also a full set of web and gui based management tools for those that don't like the command line)

Also because the hardware appears as an NIC and uses standard driver module conventions, we don't require a special version of the network capture library (libpcap) and we have no unusual restrictions on its usage or serialisation. This means the probes can run virtually any popular, proprietary Linux or Open-source network applications at these much faster speeds. The bottom-line: you can use free/cheaper/better tools on your carrier class networks.

For example, everyone's favourite open-source software IDS, Snort, is designed primarily for enterprise networks and is a typical example of a high quality monitoring application. Normally, it can monitor a few hundred megabits of traffic with a standard NIC [1]. Using HANAC's pre-emptive selection technique, Snort can monitor a full 10GBps of traffic without modification or the need for clumsy load balancing across the interfaces.

Figure 4 shows attack detection by Snort under increasingly higher loads. Notice that without HANAC, as the packet rate per second increases beyond a few hundred Mbps, Snort loses more and more attacks, quickly becoming ineffective. HANAC insulates Snort's performance from extremely high traffic loads.

This isn't magic; it is because HANAC is using full layer2-7 deep packet inspection to pre-emptively select a population of packets or pre-qualify packets that Snort will be interested in. The other packets that are of no interest are not captured. This keeps the effective data rate at the operating system much lower and is shown in Figure 5.

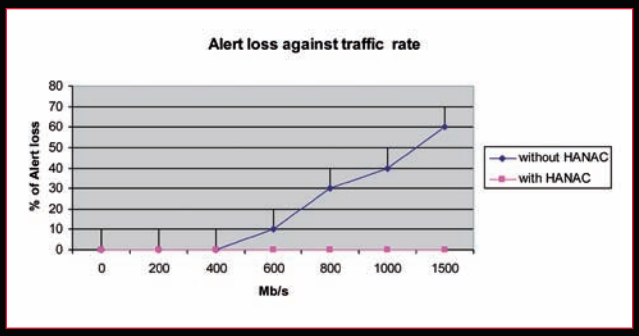

Figure 4. Snort Benefiting from Our Technology

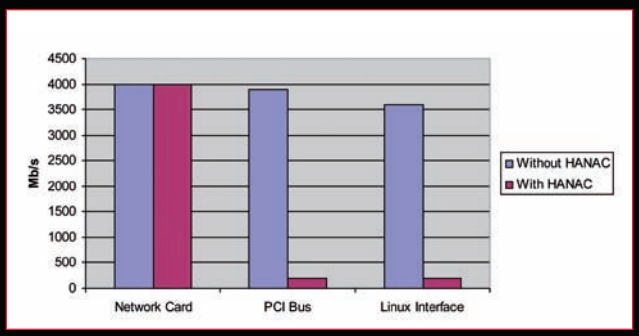

Figure 5. Snort Benefiting from Our Technology

THE SIFTER10 RANGE OF PROBES ARE ADVANCED APPLIANCES COMBINING POWERFUL SERVER TECHNOLOGY, STATE OF THE ART SOFTWARE AND REVOLUTIONARY HARDWARE ASSISTED NETWORK APPLICATION CO-OPERATION (HANAC) SUPPORT

Without HANAC, the effective bit rate measured at the Linux interface is about 3.5 Gb/s because some have already been lost by the operating system and the hardware. With HANAC, the interface only receives the pre-selected packets, which in this case produce a traffic rate below 200Mb/s, well within the safe operating range of most software applications.

/ CO-OPERATIVE PROCESSING

In addition to extending the lifetime of your existing software assets, the Sifter10 has a powerful programming API so advanced server software can task the hardware to collect network meta-data or programmatically modify access-lists to block/forward particular types of traffic.

This cooperative model is exceptionally powerful as it lets traditional software development technology be used to develop high-speed real-time network control applications. Until now, these types of applications had to be developed in hardware or relied on indirect/inaccurate sampling techniques.

As an example, network analysis and reporting software, which detects resource abuse, can be coded with one API call to detect the Top-n subnets sending traffic on a particular link. Using another call to the API, the software can set an access-list to capture all traffic from that subnet for analysis purposes. A third could be used to block that traffic while the server application analyses it and forwards it on another regular Ethernet interface. Examples of these programming techniques are freely available for download.

Using these techniques developing your own version of NTOP would be a few hundred lines of code.

/ PACKAGING

As briefly mentioned above, the Sifter10 comes as a Linux platform. It has Web based and GUI based rule management software. This means straight-out of the box, it can be deployed as a passive monitoring Snort IDS on a high-speed 10Gbe network.

The appliance can also be used to monitor up to about ten 1Gb/s LAN segments using a hierarchy of aggregation switches [2]. Also straight out-of-the-box, the unit can be used as a hardware version of in-line snort. If the uptime is a particular concern, high availability and cluster options are available. For event management, the system is fully integrated with Sguil. As an alternative, many customers prefer the web-based software BASE.

/ CONCLUSION

More and more security and network professionals are using utility hardware in conjunction with standard PC tools to solve complex problems in a parallel manner but also in a cost effective manner. CUDA the use of graphics card hardware is another example of this.

The Sifter10 is the first product that is designed to enhance PC software with hardware assisted massive parallel packet inspection processing. It offers the discriminating user:

- The possibility to extend the life of software assets or allows software vendors to launch their enterprise products into the backbone market;
- The ability to implement high speed filtering and management;
- To deploy advanced value added features into your network, reducing churn and increasing competitive advantage.

These opportunities allow a level of visibility into networks not previously feasible or affordable; an attractive proposition to those with 10Gbe backbones. /

REFERENCES

Kerry Cox and Christopher Gerg, Managing Security with Snort and IDS Tools. O'Reilly, 2004, pp. 226-227

[2] Mark Osborne, How to cheat at Managing Information Security. Syngress, 2006, pp. 212-213

ACKNOWLEDGMENTS

I would like to thank Livio Ricciulli and Ajoy Aswadhati for their sizable contribution to this paper.

/ AUTHOR BIO

Mark Osborne ran the KPMG security practice for many years (1993-2003). He has published several Zero-Day security vulnerabilities (e.g. Fatajack), and has also been an expert witness in the "cash-for-rides" case. Mark has designed the popular open-source wireless IDS/IPS (WIDZ), as well as the largest Cyber Security System in Europe. He is the author of "How To Cheat at Managing Information Security", which reached the Amazon.com Top-500.