Shifting from Software to Hardware for Network Security

Ian Land, Manager, Military Business Unit, Altera, Now Part of Intel
Ryan Kenny, Senior Strategic Marketing Manager, Military Business Unit, Altera, Now Part of Intel
Lance Brown, Senior Strategic Marketing Manager, Military Business Unit, Altera, Now Part of Intel
Rob Pelt, Security Product Planning, Altera, Now Part of Intel
Richard Takahashi, CEO, Secturion Systems

WP-01261-1.0 White Paper

Countless publications and articles let us know every day that security is the latest challenge in the next generation of Internet scaling and the information economy. We are also aware of the emerging threats and attacks that simply bypass current security solutions from media reports in the aftermath of large-scale corporate and government data breaches. A multibillion-dollar industry of software-based network security and intrusion detection solutions exists to help address these attacks. However, shortfalls in threat intelligence and modelling and also risk management frameworks that justify the funding of these solutions have inherently kept the sophistication of the attack vectors ahead of the network defense capabilities.

But what can we do about it? And more importantly, does it require a significant paradigm shift in how we develop security approaches and solutions? As Albert Einstein once stated, "The same thinking that produced a problem will not produce a solution."

Cybersecurity as a Political and Technical Component of Modern Business

In the McAfee acquisition announcement in 2010, former Intel CEO Paul Otellini stated that "[we have concluded that] security has become the third pillar of computing." In the time that has passed since that announcement, cybersecurity has become part of the business landscape as both a political and financial issue as much as a technical one.

The risks to both business and nation from failing security policies and products make the investment in better security technologies more than a measure of cost effectiveness, but of necessity. The entire security product industry relies on a set of existing codebases and threat models, without the vision or resources to develop new technologies, nor an evolved understanding of future threats. The barriers to such transformation, according to surveys from Price Waterhouse Coopers, are seen differently by CEOs, CFOs, and CIOs; CEOs bemoan capital resources, CFOs cite lack of leadership, and CIOs complain of lack of education and understanding of the risks in network security planning.

Shortfalls and Failures in Strategic Approach

Multiple news headlines demonstrate why this business component has been critical to modern government and businesses. Enough incidences have occurred to both large businesses and to government offices, such as the Office of Personnel Management, to begin drawing correlations to both the shortfalls in policy and technology that have enabled these breaches.

The analysis of successful and known attacks include some insiders, many front door attacks, and almost all were undetected for long periods of time. The primary question that arises from this data: after 10 to 20 years of software-based malware, firewalls, intrusion detection system (IDS) or intrusion prevention system (IPS), why is it that virus, malware, and network breaches continue to occur without detection?

Research by the SANS Institute, Ponemon Institute, FireEye, and other sources show that persistent cyber attack modelling capabilities developed in Asia and Eastern Europe are able to emulate and characterize release versions of software-based firewall components. As the majority of these products can be purchased on the open market, these products are reverse-engineered and software-simulated. Similar to how we use sandboxes as a network defense technique for isolating intrusive attacks, attackers test thousands of different exploit techniques in their own emulation sandboxes until they can catalog all the successful techniques that work on each variant of available firewall appliances.

Impacts on Enterprise Business

The cost of this cyber intrusion industry to business is measured every year by the Ponemon Institute, and is estimated to be about $3.8 million per network breach, and as much as $450 billion to the global economy.

Impacts on Governments and Politics

Just as impactful but less measurable are the costs of security technology shortfalls in the effectiveness of governance and international politics. Warfare is being transformed from an intermittent event to a constant state of attack and defense due to the lack of data and intrusion definitions and boundaries. This impacts the country's ability to conduct normal political interactions and trade activity. The inability to protect secrets in an open society likewise transforms the ways that the government can execute its primary role in society.

Overview of Network Security Products and Capabilities

Renowned security expert Bruce Schneier calls security a process rather than a product. Despite this characterization, the world financial industry has the tendency to categorize all products and services, creating several classes of products aimed at helping to secure networks and information systems.

Topology of Network Security Solutions

The SANS Institute divides each of the functional elements of network security control by describing these functions as Critical Security Controls. These controls are divided into as many as 20 different categories, with control functions such as Software Inventories, Boundary Defense, Authorization Controls, and Penetration Testing. Each of these controls then resolve into a number of security product solutions by a variety of vendors.

Each of these solutions that are provided by companies such as Intel Security, IBM, Cloudflare, Trend Micro, Skybox Security, HP, Cisco, and Akamai are either software or a combination of hardware and software. Historically, each of these products derive from an existing codebase from early products that are run on a variety of different company server hardware, which often means that old undisclosed weaknesses are perpetuated from product to product.

Figure 1. Network Security Controls

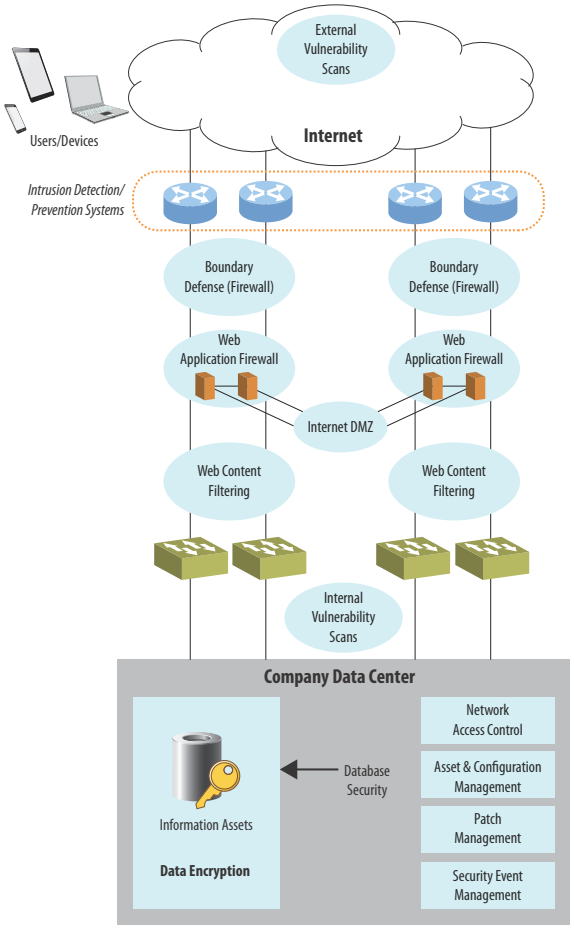

Software as Basis for Updatability and Agility

Using network firewalls and testing sandboxes as representative categories of the network security solutions market, the majority of these solutions are software-based and utilize virtual machines and environments. This is for agility, portability, and economical reasons—very little initial capital is typically needed for software-based products. Most importantly is the ability to update and patch products as new vulnerabilities and threat vectors are discovered.

Even more of the network and packet processing functionality is moving into software and virtual machine structures as a result of research and early product development in software defined networks by a variety of large original equipment manufacturers (OEMs).

Reconfigurable Hardware as Alternative

Reconfigurable hardware such as FPGAs can offer significant advantages to network security appliances, plus offer all of the updatability and agility advantages of software through FPGA system on a chip (SoC) products, new design entry models such as Open Computing Language (OpenCL™), and virtualization support in both the ARM® Cortex®-A9 and ARM Cortex-A53 hard processor subsystems. This virtualization capability is leveraged for software-defined networking (SDN) capabilities on FPGAs and programmable logic SoCs as well.

The data bandwidths and speeds of reconfigurable hardware are a key security advantage. Hardware can enable the network to monitor all activity, while processor and software solutions may enable only partial monitoring until the parallelization model becomes too encumbering. By using reconfigurable hardware instead of dedicated fixed circuits, the necessary ability to update and patch products as new vulnerabilities are detected is maintained. Additionally, because the data processing path is reconfigurable, older hardware can be updated to support new and better security standards.

Brief History of Research into Programmable Hardware for Network Security

In a thesis on reconfigurable hardware for network security, Sascha Muhlbach summarizes previous work, study, and demonstrations of security elements on various FPGAs. He breaks these research efforts into the categories of packet classification, pattern matching, anomaly detection, and applications and communication support.

Packet classification is typically used for filtering in network firewalls, and is accelerated in FPGAs in systems today using ternary content addressable memory (TCAM) architectures. Pattern matching includes the rule-based IDS applications, most recently demonstrated on Altera FPGAs. Anomaly detection includes statistical implementations of IDS systems. Applications and communication support refer to a variety of acceleration functions for protocol operations requiring security, including domain name system (DNS), Secure Shell (SSH), and so on.

In edge-of-network, Internet of Things (IoT) applications, hardware security is being recognized as the preferred approach as well. As summarized in the Electronicsofthings.com blog:

"A hardware-first approach with respect to security and implementation of necessary functionality on the SoC level is vital for fully securing devices and platforms such as FPGAs, wearables, smartphones, tablets, and other intelligent appliances."

In each of the applications described above, FPGA hardware has several distinct security advantages above and beyond just performance acceleration goals.

Hardware Advantage: Fewer Data Vulnerabilities

When designed correctly, FPGAs cannot be altered without detection. Tampering of hardware circuits and hardware anti-tamper features requires physical access to the hardware. This is enabled through authenticated partial reconfiguration as well as the OpenCL and high-level synthesis programming models described hereafter. In addition, a key feature not available for software security solutions is fail-safe operation. Correctly implemented, fail-safe designs ensure the operating parameters of the security product are within its operational boundaries. If the security product has a logical or physical error, such as failure of component, power supply surge, and unknown glitches, the product will fail to a safe state. This ensures the product does not continue operating in a defective state.

Also, FPGA-based network products cannot be altered via a network connection or front door attack. If there is an attempt to alter the design, either the anti-tamper mechanisms or fail-safe design feature will detect and shut down the product.

Figure 2. Next-Generation FPGAs Contain Many System Security Support Functions

Hardware Advantage: Hardware Root of Trust

FPGAs and the built-in security features of modern products offer significant advantages in hardware-based roots of trust in systems. These include protected key storage, obfuscation, and anti-tamper techniques, as well as anti-counterfeiting and asymmetric encryption functions such as the Physically Unclonable Function and hardened Elliptic Curve Signature accelerators. All of this can be aided by the ARM Trustzone functionality enabled in the ARM Cortex-A9 and Cortex-A53 hard processor systems.

For more information, refer to the *Stratix® 10 Secure Device Manager Provides Best-in-Class FPGA and SoC* (PDF) white paper.

Hardware Advantage: Custom Processor Architectures

One area of research in programmable logic for network security appliances is in customized Reduced Instruction Set Computing (RISC) processors for specific network security functions, implemented on FPGAs. One in particular from Cambridge University is called Capability Hardware Enhanced RISC Instructions (CHERI), which demonstrates custom memory management units (MMUs) and memory interfaces to address common and known shortfalls in C-based programming languages. This allows security-aware products based on programmable logic to select 'soft' processor architectures customized for the threat environment or attack surface of that particular node in the security architecture.

Known Software Attacks: Well-Documented Attacks on General Purpose Processor Architectures

Attacks and reverse engineering techniques are well understood and documented for traditional computing architectures and products. Several of these techniques are used to put software solutions into non-sequitur or unexpected states. These same attacks, applied in a checklist fashion to processor and software-based products, require significant customization and considerably more skill for dedicated hardware and programmable hardware devices:

- DC power supplies glitches
- CPU attacks or glitches
- Memory and stack attacks
- Various of memory-based and zero-day exploits

An important distinction between attacks on processors and software and attacks on dedicated hardware is that software attacks can expand. If an attacker finds one exploit that provides access, then the attacker can effectively own the system. Hardware exploits provide far more narrow accesses and exploitation opportunities.

Industry Conclusions on the Hardware-Based Root of Trust

Many processor companies including Intel and Freescale have participated in defining the Trusted Platform Module (TPM) as an example of a hardware root of trust to improve operating system (OS) and application security in computing systems. Intel has further moved to develop their Trusted Execution Technology (TXT) built into their modern CPU architectures.

ARM products themselves do not inherently provide hardware root of trust solutions, although some ARM ecosystem partners provide frameworks that can take advantage of hardware roots of trust.

National Institute of Standards and Technology (NIST) is one of many groups looking at standards for moving roots of trust from software into firmware, boot firmware, and hardware for mobile devices and IoT technologies. Their public study results, and that of most other research into root of trust technologies, all take for granted that software-based roots of trust are problematic; firmware is better, but building roots of trust into processing hardware in secure enclaves is ideal.

Programmable logic components provide the opportunity to integrate a variety of hardware root of trust solutions in the form of metal keys and support, fuse keys and support, and authenticated programmable logic hardware roots of trust. This provides the opportunity to implement a hardware root of trust that is hardened outside the code execution paths, while also programmable and can be updated through authenticated update controls.

Example: DarkStor FPGA-Based Enterprise Data Encryptor

Taking one network security component as an example, the data storage component of an enterprise can immediately improve security by ensuring that network breaches provide access only to protected data, and that access to the decryption capabilities is authenticated by a hardware root of trust. Beginning with isolated data encryption and moving towards improving traffic and protocol security with hardware-rooted security solutions provides a redundant protection mechanism approach. Fixed-rate encryption and key management may also be the most rational first network element to address using programmable hardware solutions. Beyond this first and missing capability in traditional data networks, an understanding of the security product market is required.

The DarkStor Enterprise Data Encryptor uses Altera® Arria® 10 FPGAs as the processing component for the 100G bidirectional encryption and key management functions both for the high performance enabled by modern FPGAs, but for security.

Pending Shifts in Security Product Platforms

The basic overall security stance for an entire industry of security products is reactive. This is the process of identifying threats, cataloging and communicating those threats, testing updates and patches, and rolling out threat signatures and updates to security appliances. The most prominent shift in security product capabilities and models in current security research is the ability to both anticipate threats and model the behavior of attackers to anticipate new attack vectors. These high-performance computing-enabled shifts are substantially enhanced with programmable hardware.

Convolutional Neural Networks and Artificial Intelligence

One technology for applied pattern, image, and speech recognition is the use of cellular convolutional neural networks (CNNs). This processing approach has been studied and demonstrated in graphics processing unit (GPU) and FPGA technologies by numerous academic and industry labs. Further applications of this technology include adaptive threat detection and vulnerability assessment.

Massively Parallel Programming Models

The massively parallel programming model that is needed to continue Moore's Law has been approached in three different ways in the last decade: multicore processors and programming models, virtual machine management and parallel threading, and reconfigurable hardware and FPGAs. For this paper, we will just address the reconfigurable hardware model.

The FPGA hardware model is the only one of these three implementations that enables the security advantages regarding published code vulnerabilities and hardware roots of trust. But the massive increase in programmable logic density made available by Stratix 10 devices, as well as tightly coupled Intel processor and FPGA capabilities that will soon be available on the market, will fuel this capability to the advantage of secure system developers.

OpenCL and High-Level Synthesis Models for Cybersecurity

Heterogeneous computing models, such as the CNNs described above or tightly coupled FPGA and CPU appliances, will require heterogeneous programming capabilities to become practical. Altera has invested extensively to enable these new high-performance, predictive, and secure computing applications using the OpenCL design flow. OpenCL compilers from Altera will parse, synthesize, and compile kernel operators into discrete functions that can run separately on a processor and in parallel FPGA hardware.

Examples of Packet Sniffing Applications on Reprogrammable Hardware

Altera has collaborated in two demonstrations of security applications accelerated onto programmable hardware. The first of these include the porting of the Open Source Suricata IDS/IPS application into OpenCL and implementation on both host processor and FPGA. The second includes two demonstrations of the Open Source Snort algorithm likewise accelerated on Altera FPGAs by Clemson University and Lewis Roades Labs. Information on each of these are available in published papers, and in demonstrations at various trade shows hosted by Altera.

Platform Shift and the "Innovator's Dilemma"

A few well-established companies now provide elements of traditional network and home security to include antivirus, firewalls, and IDS/IPS. The established software-focused products, support contracts, software codebases, and codified threat modelling represent an asset for IT managers attempting to address enterprise security, but they can also represent an obstacle to innovation towards reconfigurable hardware and massively parallel threat detection and prediction capabilities.

This "Innovator's Dilemma" provides an opportunity for new entrants into the network security market, both large and small, to provide security differentiated products through the use of FPGAs and supporting heterogeneous computing models.

Conclusion

The software-centric approach towards enterprise and network security has roots in agile development and quick updatability, but have led to a cascading discovery of new vulnerabilities and attack vectors in these appliances.

Replacing this entire industry with fixed and custom secure hardware solutions is not necessarily the right answer because of cost, and the loss of agility and upgradability in development. However, heavily leveraging programmable logic solutions as part of a network appliance security strategy has both well-studied and newer speculative advantages over traditional software and software stack approaches. Network and enterprise security appliances, based on programmable logic solutions, could provide substantial security advantages if adopted by established or new players alike in the network equipment security market.

Further Information

- Intel: 'Security now third pillar of computing', August 19, 2010: www.techradar.com/us/news/software/computing/intel-security-now-third-pillar-of-computing-711244
- Cybersecurity: The New Business Priority, 2012: www.pwc.com/us/en/view/issue-15/cybersecurity-business-priority.html
- Real World Testing of Next Generation Firewalls, October 2013: www.sans.org/reading-room/whitepapers/analyst/real-world-testing-next-generation-firewalls-34955
- Ponemon Institute's 2015 Global Cost of Data Breach Study Reveals Average Cost of Data Breach Reaches Record Levels, 27 May 2015: www.prnewswire.com/news-releases/ponemon-institutes-2015-global-cost-of-data-breach-study-reveals-average-cost-of-data-breach-reaches-record-levels-300089057.html
- 10 Gbps Line Speed Programmable Hardware for Open Source Network Applications, Livio Ricciulli: www.metanetworks.org
- Hardware/Software Support for Securing Virtualization in Embedded Systems, Franck Bucheron, Arnaud Tisserand, Louis Rilling. 1st Symposium on Digital Trust in Auvergne, December 2014
- FPGAs for the Internet of Things, Abishek Mutha, 27 November 2015: http://electronicsofthings.com/2015/11/fpgas-for-the-internet-of-things/
- Capability Hardware Enhanced RISC Instructions (CHERI), Cambridge University, Multiple Publications 2015: www.cl.cam.ac.uk/research/security/ctsrd/cheri/
- Reconfigurable Architectures and Design Automation Tools for Application-Level Network Security, Sascha Muhlback, Thesis 2014, Published 2015

- The New Era of Mega Trends: Hardware Rooted Security. Ahmed Sallam, VP and CTO, Haardware, Security, and Emerging Solutions, Citrix Solutions: www.citrix.com/articles-and-insights/trends-and-innovation/jan-2015/the-new-era-of-mega-trends-hardware-rooted-security.html
- An FPGA Based Implementation of Snort Pattern-Matching and Preprocessing Engines, Nilim Sarma, Keerthan Jaic, Melissa Smith, Holcombe Department of Electrical and Computer Engineering, Clemson University, August 2015
- Secure Embedded Systems: Digging for Roots of Trust. Ron Wilson, 22 July 2015: systemdesign.altera.com/secure-embedded-systems-digging-for-the-roots-of-trust/
- Stratix 10 Secure Device Manager Provides Best-in-Class FPGA and SoC Design Security. Ting Lu, Ryan Kenny, Sean Atsatt: www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/wp/wp-01252-secure-device-manager-for-fpga-soc-security.pdf
- Trusting the Cloud with Intel Trusted Execution Technology: www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusting-the-cloud-with-trusted-execution-technology-video.html
- Computer Security Resource Center: Hardware Roots of Trust, National Institute of Standards and Technology, 2014: csrc.nist.gov/projects/root-trust/
- Accelerating Open Source Security Using OpenCL: www.altera.com/en_US/pdfs/literature/po/open-source-security_ss.pdf

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
February 2016	1.0	Initial release.