Internet Worm and Virus Protection for Very High-Speed Networks

John Lockwood
Washington University in Saint Louis
lockwood@arl.wustl.edu - (314) 935-4460 - http://www.arl.wustl.edu/~lockwood

Abstract
The security of the Internet can be improved using reconfigurable hardware. A platform has been implemented that actively scans and filters Internet traffic at multi-Gigabit/second rates using reconfigurable hardware. Modular components implemented in FPGA logic process packet headers and scan for signatures of malicious software (malware) carried in packet payloads. Additional FPGA circuits track the state of Transmission Control Protocol (TCP) flows. Regular Expressions and fixed-string scanning circuits are implemented in parallel hardware. Dynamic reconfiguration enables remote modules to be reconfigured to scan for new signatures. Network-wide protection is achieved by the deployment of multiple systems throughout the Internet.

Introduction
Computer viruses and Internet worms cause billions of dollars in lost productivity. Well-known Internet worms like Nimda, Code Red and Slammer contain strings of malicious code that can be detected as they flow through the network. By processing the content of Internet traffic in real-time, a computer virus or Internet worm can be detected and prevented from propagating. Our system scans the full payload of packets to route, block, and account for the content in the flow. One challenge in implementing the system was that the location of a signature in the packet payload was not deterministic--it could appear at any position within the traffic flow. Another challenge to implementing the system was that signatures could span multiple packets and be interleaved among multiple traffic flows. The paper will describe how these challenges were met and overcome.

Related Work
A common requirement for network intrusion detection and prevention systems is the requirement to search for predefined signatures in the packet payload. Since conventional software-based algorithms for deep packet inspection have not kept pace with high-speed networks, hardware-based solutions are desirable. Hence, important building blocks of these systems include fast signature matching and protocol processing circuits. Most systems in this class have a common requirement for string matching. For example, a media file can be characterized by the presence of a string of bytes (for the rest of the paper, a string is synonymous to a signature) and its transmission across a link can be monitored by looking for the presence of this string on the link.

Key Contribution
Our key contribution is to envision, design and develop a cohesive malware protection system that includes an FPGA-based network platform, Internet protocol processing circuits, content matching modules, and automated design tools to enable the implementation and timely updating of network security applications in reconfigurable hardware. The system allows for the immediate blocking of known viruses and may be rapidly reprogrammed to recognize and block new threats. These upgrades are system-driven, and are not dependant upon actions by the end users to assure that the protection remains up to date.

The system's foundation is the Field-programmable Port Extender (FPX), which is implemented with two FPGAs, five banks of memory and two high-speed (OC-48 rate) network interfaces. The network interfaces connect to one of several types of Gigabit-speed line card interface cards, including several types of Gigabit Ethernet and ATM interfaces. On the FPX, one FPGA is used to route individual traffic flows through the device, while the other is dynamically reconfigured over the network to perform customized packet processing functions. Using the latest FPGA technology, the system could easily scale to process 10 Gigabit/second OC-192 flows.

A TCP/IP wrapper, implemented in FPGA logic, reconstructs the flow of transmitted data by tracking sequence numbers of consecutive packets to provide a byte-ordered data stream to the content scanning engines. This means that even if a malware signature has been fragmented across multiple packets, it still will be detected and blocked. In order to maintain the state of multiple traffic flows, the system architecture has been designed to store the state of a TCP/IP flow in memory. Given that each flow occupies 64 bytes of memory, one 512 Mbyte SDRAM (about half of the memory on the FPX) module can track 8 million simultaneous traffic flows.

Two methods are used to search for signatures: a finite automata scans for regular expressions and a Bloom filter scans for fixed strings. The number of regular expressions that can be searched grows with the amount of the FPGA logic on the device, while the number of fixed strings that can be searched grow with the size of on-chip RAM. A Bloom filter allows a scanning engine to identify up to 1,700 fixed-length strings. Both types of our engines can scan traffic at traffic at 600 Mbps. By implementing four engines that run in parallel, the FPX can process data at a rate of 2.4 Gigabits per second using a single Xilinx Virtex 2000E FPGA.

An automated design flow builds packet scanning circuits in hardware. Custom circuits are built by an automated program that reads a list of signatures from a database table, optimizes each finite automata, integrates Internet protocol processing hardware, compiles the circuit into gates, routes and places the circuit into a FPGA, and then reconfigures remote devices over the network.

Conclusions

We have designed and developed a system that blocks the spread of Internet worms and computer viruses. Our system uses reconfigurable hardware to scan Internet traffic for malware. Malware is identified by signatures that may consist of either fixed strings or regular expressions. TCP/IP flows are tracked so that signatures spanning multiple packets can be detected. An automated design flow allows new circuits to be rapidly deployed to protect the network against new attacks.

References

- J. W. Lockwood. An open platform for development of Network processing modules in reprogrammable hardware. In IEC DesignCon'01, pages WB–19, Santa Clara, CA, Jan. 2001.
- R. Sidhu and V. K. Prasanna. Fast Regular Expression Matching using FPGAs. Field-Programmable Custom Computing Machines (FCCM), Rohnert Park, CA, Apr. 2001.
- R. Fanklin, D. Caraver, and B. Hutchings. Assisting network intrusion detection with reconfigurable hardware. Field Programmable Custom Computing Machines (FCCM), Apr. 2002.
- M. Fisk and G. Varghese. Fast content-based packet handling for intrusion detection. Technical Report CS2001-0670, University of California, San Diego, 2001.
- J. W. Lockwood, N. Naufel, J. S. Turner, and D. E. Taylor. Reprogrammable Network Packet Processing on the Field Programmable Port Extender (FPX). In ACM International Symposium on Field Programmable Gate Arrays (FPGA), pages 87–93, Monterey, CA, USA, Feb. 2001.
- J. Moscola, J. Lockwood, and R. P. Loui. Implementation of a Content-Scanning Module for an Internet Firewall. Field-Programmable Custom Computing Machines (FCCM), Apr. 2003.
- M. Necker, D. Contis, and D. Schimmel. TCP-Stream Poster on Reassembly and State Tracking in Hardware. Field-Programmable Custom Computing Machines (FCCM), Apr 2002.
- D. V. Schuehler and J. W. Lockwood. TCP-Splitter: A TCP/IP Flow Monitor in Reconfigurable Hardware. Symposium on High Performance Interconnects (HotI), pages 127–131, Stanford, CA, USA, Aug. 2002.