



Intel® Technology Journal

Toward The Proactive Enterprise

**Bringing Security Proactively
into the Enterprise**

Bringing Security Proactively Into the Enterprise

Sanjay Rungta, Information Services and Technology Group, Intel Corporation
Anant Raman, Technology and Manufacturing Group, Intel Corporation
Toby Kohlenberg, Information Services and Technology Group, Intel Corporation
Hong Li, Information Services and Technology Group, Intel Corporation
Manish Dave, Information Services and Technology Group, Intel Corporation
Greg Kime, Information Services and Technology Group, Intel Corporation

Index words: firewalls, packet filtering, port security, super-VLAN, sub-VLAN, distribution layer, access layer, patching, hardened systems, policy-enabled network

ABSTRACT

Prevailing network architectures are designed for openness, collaboration, and sharing. The majority of viruses and worms use the network to spread rapidly through the enterprise network, enabling these cyber threats to reach their targets effortlessly. The most common solution available today for cyber security is hardening of systems via “patching” or keeping the operating systems, applications, and anti-virus software current. This option is reactive and time/labor intensive because security updates are available only after exploits are known and already in use. The currency of software does nothing to prevent cyber attacks from reaching their targets. We believe that policy-enabled network security complemented by system hardening, provides a proactive and more comprehensive strategy to deal with security by reducing the likelihood of cyber threats entering the network and by controlling their spread. Typical enterprise network architectures are developed to bring scalability, extensibility, and availability to the Intranet. Security capabilities have not been part of the enterprise network architecture and are typically implemented in reactive fashion. Additionally, current security capabilities require manual and labor-intensive efforts that negatively impact costs and take time to implement. Firstly, we propose a change to the enterprise network architecture by integrating security components such as packet filtering, stateful inspection, port-based access control, and super/sub Virtual Local Area Networks (VLANs). Secondly, we propose a fundamental change in the implementation of the enterprise network architecture by using a security management system referred to as Policy-Enabled Network Security (PENS) that leverages the new security capabilities in an integrated and proactive manner

and reduces unstructured manual, labor-intensive, and error-prone activities.

INTRODUCTION

One of the major security problems in the enterprise network is the “permit all” capabilities of all Local Area Networks (LANs). While the openness was a catalyst to the growth of computer networks, it also presented and continues to allow computers with security issues to freely connect to the network and potentially infect other computers. LANs do not have the capability to determine who is connected to the network and therefore, in an enterprise, cannot separate out company employees from visitors. Moreover, the enterprise network is now required to support flexible connectivity to portable computers such as laptops for company employees as well as visitors such as field service personnel, consultants, customers, partners, etc.

The enterprise network is also increasingly becoming more diverse and complex with the explosive growth in the usage of wireless technology. New business models such as supplier support of equipment over the Internet, remote operations of equipment over the network (with safety precautions), and company guests accessing their Intranets over the Internet, have blurred the lines between the Intranet and Internet. The added flexibilities in the workplace cause a greater security risk than ever before: the source of an infection nowadays can be any host accessing the enterprise. Figure 1 demonstrates the typical enterprise network.

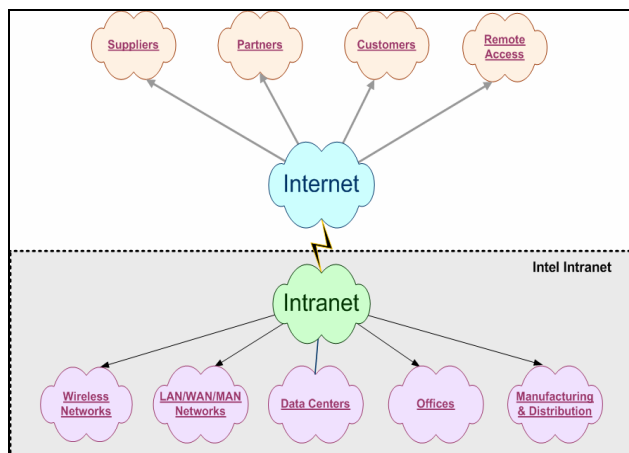


Figure 1: Network connectivity in the enterprise

Unhardened Computer Systems

An enterprise network typically has several operating system and software standards such as UNIX*, Microsoft*, Apple*, OpenVMS*, etc. Most commercial off-the-shelf operating systems provide a general-purpose computing platform for different types of applications and configurations. These systems tend to have more services than are typically needed, and most people tend to keep these services running. Vendors provide frequent software updates and bug fixes. These computer systems, commonly known as unhardened computer systems, have now become the major source of security issues. They can easily become infected and in turn, infect other computers accessible over the LAN.

Inadequate Policies and Management

Traditionally, enterprise policies have been limited to perimeter security resulting in the usage of controls to reduce risks within the enterprise network. With the increase in interactions across the perimeter as we just described, the challenge is to adapt security policies to reflect these new threats. These policies and corresponding operational capabilities have to be able to support complex protection profiles to protect each participant in the entire enterprise network. Each participant contributes to the business and is a challenge to the security of the business based on his/her level of interaction with the environment. This seemingly paradoxical situation increases the complexity of managing security policies and of enforcing them. Traditional processes and methods for implementing security are manual, labor intensive, and error prone.

* Other brands and names are the property of their respective owners.

Intranet No Longer a “Safe Haven”

In the aftermath of virus attacks such as SoBig, Nimda, FunLove, SQL Slammer, etc., as well as a lack of a policy management system, the Intranet, as it exists today, is no longer secure. Enterprises are struggling to keep up the frenetic pace of updating software because the hackers are ready with the next set of attacks before the systems are updated with fixes for the previous attack. Keeping thousands of computers updated is a very reactive process because it can only secure us from problems with known and available solutions.

Solving the Problem

We started our research by reviewing our current network design methods, security capabilities, and management practices and concluded that a change in the network architecture is the most effective way to bring security to the enterprise network proactively. A product-level approach is also possible, but it would require us to find products that can provide enterprise-level security capability for 100,000 to 500,000 computers, 100,000 to 150,000 LANs, and greater than 10,000 subnets. This would take several person years and would not guarantee us a solution to our security problems. This approach is not feasible and is therefore, not addressed in this paper.

ADDING SECURITY TO THE ARCHITECTURE

Today’s enterprise networks are based on a three-tiered architecture called the Hierarchical Model for Internetwork Design [1]. A slightly modified version of this architecture is shown in Figure 2 with the three layers: core, distribution, and access. The core and access layers correspond to Data Link Layer of the Open System Interconnections (OSI) model [2], and the distribution layer corresponds to the Network Layer of the OSI Model, respectively. The core layer separates the enterprise infrastructure backbone (WAN, Internet, etc.) from the end-user areas with workstations, application servers, databases, and other services.

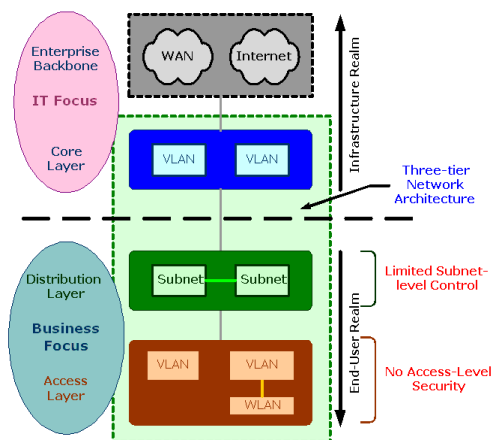


Figure 2: The enterprise network architecture

As pointed out in Figure 2, there are no elements of security in this model as any LAN device can connect and freely communicate with other LAN devices. Additionally, a TCP/IP device can connect and freely communicate with other TCP/IP devices anywhere in the enterprise network.

The Security Capabilities

The two key security additions to the network architecture are as follows:

1. *Access layer capabilities.* These focus on knowledge of the device connecting to the LAN or knowledge of which LAN devices can communicate with each other.
2. *Distribution layer additions.* These focus on subnet-to-subnet-level application-level control.

Access Layer Security Capabilities

There are three forms of access layer security.

MAC Address Filtering

The simplest of the access layer security capabilities is based on knowing the MAC addresses of all LAN devices and unknown MAC addresses not being permitted to join the VLAN. Most enterprise-level Layer 2 switches support MAC address filtering. This approach is clearly not scalable for the enterprise networks but suitable for very small areas. Additionally, with the prevalence of notebook computers with removable network interfaces, a known MAC address does not provide sufficient access layer authentication.

Port-Based Access Control (802.1x)

The IEEE 802.1X standard [3] offers both wired and wireless devices a method to authenticate the device and the user of the LAN device before connecting to the VLAN. Based on the Extended Authentication Protocol

(EAP), the 802.1X standard routes the EAP network traffic to a RADIUS server [4], the authentication server. Only authenticated users and devices are allowed to connect to the VLAN. All other devices are not allowed to connect to the enterprise. Device authentication can include local security-level checks such as operating system updates and virus signatures. While IEEE 802.1X is more scalable than MAC address filtering and takes network port security to a new level, it is considerably more complex and expensive because it requires the support for the standard from the Layer 2 switches and the operating systems. Legacy operating systems have very limited support for IEEE 802.1X.

Super and Sub VLANs

RFC 3069 [5] introduced the notion of super-VLANs and sub-VLANs to realize the optimization of IP addressing in a switched environment. Each sub-VLAN has its own broadcast domain while using the default gateway IP address from the super-VLAN. A leading network vendor has added two security capabilities in the sub-VLAN and super-VLAN space, called secondary and primary VLANs, respectively [6]. First, the secondary VLANs can be either in an isolated mode where the members cannot communicate with each other or the community mode where the members can communicate with each other in a peer-to-peer fashion, with the primary VLAN providing the default IP gateway access. Super-VLANs and sub-VLANs extend port-level security by creating communities that are allowed to communicate with each other and denying communications to all others. Figure 3 shows all three of the access layer security capabilities.

Distribution Layer Security Capabilities

Packet filtering is the common term for distribution layer security. Packet filtering provides subnet-to-subnet-level network access control. Firewall devices, routers, and Layer 3 devices use the network layer information to allow or deny access to TCP/IP devices. There are two types of packet filtering:

1. Static packet filtering
2. Dynamic/stateful packet filtering

Static Packet Filtering

Static packet filtering provides us the ability to control the source and destination of the network traffic with application access through TCP/IP ports. Static packet filtering rules are applied at network and transport layer headers only. Most routers and Layer 3 switches provide static packet filtering. Typical packet filtering includes “permit” and “deny” rules. Static packet filtering also includes the shielding of internal IP addresses through Network Address Translation (NAT) [7].

Dynamic/Stateful Packet Filtering

Dynamic packet filtering adds additional intelligence to the static packet filtering. Static packet filtering allows the entire dynamic port range (e.g., greater than 1023) during a client-server session. Dynamic packet filtering, on the other hand, knows to look into the data and find the client port. The dynamic packet filtering then allows only that client port in real time for that session, as opposed to static packet filtering, which opens the range of client ports (e.g., > 1023). Dynamic packet filtering prevents security attacks based on hijacking of established sessions. Moreover, dynamic packet filtering can add time sensitiveness by opening and closing ports on an as-needed basis. Various types of dynamic packet filtering are available today. Stateful inspection [8], for example, matches the HTTP protocol-to-protocol headers along with send-receive pairs and data types as specific in the header. Circuit filtering [9] permits inspection of sessions, as opposed to connections or packets, with each session containing a number of connections. Circuit-level filtering takes into account secondary connections such as the data part of the FTP protocol and streaming media. Application filters [10] or “proxies” are application-specific with access to details of the application-level commands. HTTP proxies can be set to deny the GET command but to allow the PUT command instead. Figure 3 shows the distribution layer security capabilities in addition to the access layer security capabilities.

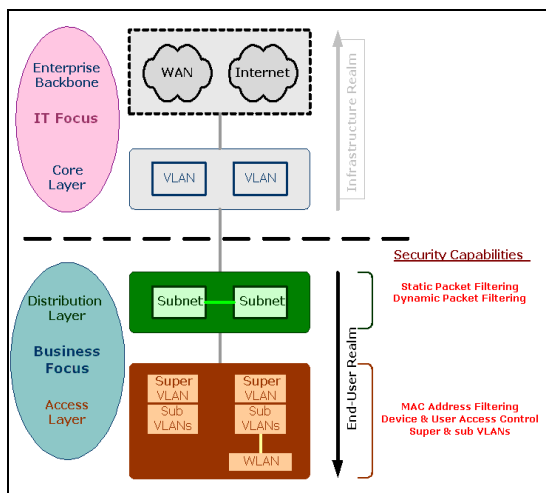


Figure 3: Distributed layer and access layer security

Policy-Enabled Network Security (PENS) Management

PENS [11] is an architecture developed by Intel Information Services and Technology Group (ISTG) that enables a common security policy specification across a heterogeneous enterprise network, and that allows

correlation of abnormal events observed by other network and security-monitoring solutions such as intrusion-detection systems, vulnerability scanning tools, and incident alert services. In this architecture, policies are dynamically linked to the threat environment via an adaptive feedback loop.

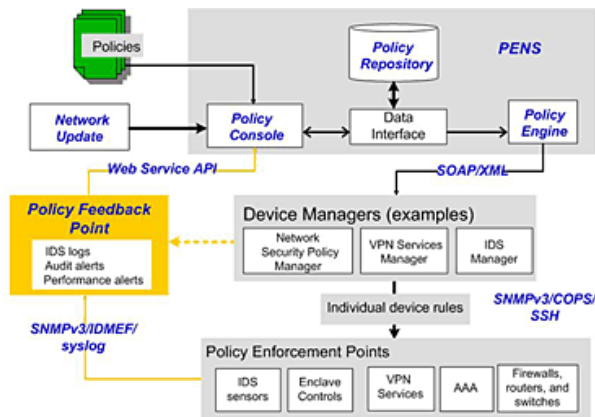


Figure 4: PENS architecture

As shown in Figure 4, the components of this policy-enabled management system include policy server, Policy Enforcement Points (PEPs), and Policy Feedback Points (PFP). As indicated, the main distinction of the PENS architecture from the standard policy-based network management architecture [12] is the introduction of PFP along with the adaptive feedback loop. A PFP collects data on intrusions, security alerts, and other abnormal network behaviors from a variety of systems (e.g., intrusion-detection systems, system performance logs, audit alerts, vulnerability scanners, etc.) and sends such data as feedback to the Policy Decision Point (PDP). The PDP then correlates the feedback data and determines if any policy updates are needed.

The main features for the above PENS architecture are as follows:

- A centralized view of the network from the management console, with real-time updates of the network.
- Automated management with an adaptive feedback loop that can dynamically modify the implemented security controls to address changing security threats.
- The ability to have common policies pushed from a central location to various network devices (PEPs) from different vendors.

Here the PENS server may act as manager of managers where the vendor-specific rules are managed by device managers and are therefore transparent to the PENS administrator.

Resulting Enterprise Network Security Architecture

Figure 5 provides an integrated view of the network and security architecture. Access layer security gives us the ability to authenticate devices and users connecting to the LAN. Device checking will give us the ability to ensure that unhardened systems considered insecure will not be permitted to connect to the enterprise network. Likewise, users unknown to the enterprise will not be permitted to connect to the enterprise network. Additionally, authenticated devices and users will be confined to the required communities of users and systems and denied access to everything else.

Distribution layer security provides subnet-to-subnet-level access control, enabling application-level control over the enterprise network.

The use of super-VLANs and sub-VLANs converts the three-tiered network architecture into a four-tiered network architecture integrated with security, making the Intranet more secure.

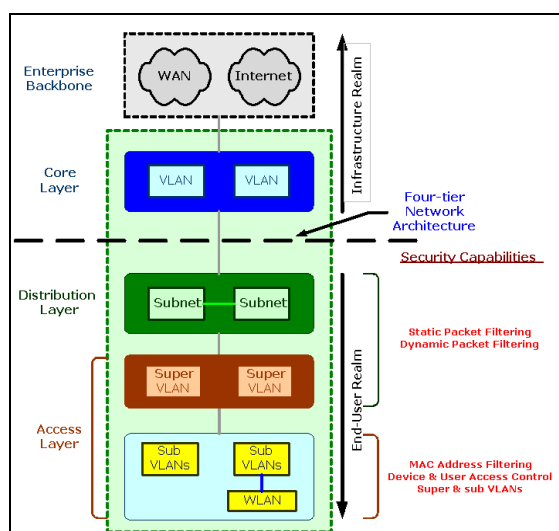


Figure 5: Four-tiered network architecture with security

With the addition of the required key security components to the Enterprise Network Architecture, we now turn our focus to the implementation challenges facing us, among which are significant resource requirements and consistent and comprehensive management of security policies. The following section addresses these challenges in detail.

Integrating Security with PENS

Some of the different ways of integrating the security capabilities proposed earlier with PENS are outlined below.

Location-Based Authentication

We can apply specific security policies to specific networks. For example, the authentication requirements for user access to applications in an office environment will be different from access to applications hosted in an Internet DMZ.

Behavior-Based Re-Authentication

It is possible to alter the authentication time limit after initial authentication is complete. This type of policy allows system administrators to control the authentication time limits when end users misbehave.

PENS can modify the authentication policies of a network and force a re-authentication, particularly when the network is under a cyber attack. It is also possible to specify a different policy for re-authentication dynamically.

Network Status-Based Re-Authentication Requirements

In the case of a serious cyber attack, the policies for infected networks could force all users to log off and disallow reconnection until a compliance scan shows that appropriate security updates are completed.

Authentication-Based User-Specific Policy Delivery

This scenario requires different authentication methods for different users followed by delivering different policies for the users. For example, company visitors have minimal authentication to our network and are allowed access to the Internet with no access to the Intranet. In another scenario a person from the security incident response team may be required to provide strong authentication but would then be able to connect to any network and have full access without being required to run a security scan.

IMPLEMENTATION CHALLENGES

The cost in human resources and technology to add security to network design, based on the new architecture, will not be trivial. The extent of the investment will be determined by the size of the enterprise (number of people, number of computers, number of subnets, number of locations, span of business group across geography, etc.). It is essential to have accurate knowledge of the security weakness in order to control the cost of implementation. The major issues are as follows:

1. Technology or technical issues
2. Business issues
3. Integration with legacy systems

Technical Issues

The availability of sub-VLAN and super-VLANs technology and 802.1x technology is limited. Not all

vendors provide this capability and in addition, not all product lines have the capabilities implemented fully.

Packet filtering requires the identification of all network communications paths. It is imperative to comprehend who needs to communicate with whom and with what applications. Understanding and documenting the various protocol interactions of applications is a very tedious but necessary process in successful implementation of packet filtering. Packet filtering rules can be easily created and managed with the knowledge of sources, destinations, protocols, and ports.

The complexity of the communications model leads to the next issue: the size of the packet filtering rule set. It is important to keep the packet filtering rule size fairly compact and simple. Larger rule sets could lead to over-utilization of the network devices (such as routers and switches) leading to performance degradation of the enterprise network [13].

Packet filtering additionally brings new operational challenges to network management. The Internet Message Control Protocol (ICMP) is considered insecure [14]. Therefore, utility programs such as “ping” and “traceroute” have limited use in this environment. Network operations will require the development of alternate network troubleshooting and debugging methods.

Business Issues

The use of new security capabilities require a significant investment in hardware. Port-level security along with super-/sub-VLAN technology requires the edge network equipment be switched as opposed to the older concentrators. Additionally, only limited products support these capabilities.

The next challenge is to group computers based upon the similarity of their communication requirements. As described before, this requires intimate knowledge of all the applications along with their source and destination TCP/IP parameters. One of the many complexities network engineers face is the accommodation of high availability key infrastructure services such as authentication, DNS, Web, database, and middleware within packet filtering rules.

Some of the operational challenges are outlined here:

1. Identifying and separating the servers and workstations of each business unit in the enterprise network because each business unit may have different security requirements.
2. The identification and subsequent separation of mission-critical services from regular services. This is particularly important because mission-critical services from different business units are required to

run continuously on the network even if there is a large enterprise-level cyber attack, such as an SQL Slammer [15].

3. Operational knowledge of the impact of these mission-critical services on the functioning of the enterprise. For example, turning off the enterprise shipping application due to a virus attack may not impact the engineering business unit but may impact the distribution and manufacturing business units.

One of the challenges is keeping the training of IT professionals up to date on new security capabilities. Keeping the operations personnel trained is vital to maintaining service levels. In addition to trouble shooting and diagnostics skills for the various products, IT support personnel need to have working knowledge of the many applications and systems in the enterprise. Deeper knowledge of the system design and components is critical to satisfactory customer service.

Managing Legacy Environments

One of the biggest integration challenges in an enterprise has to be the co-existence with legacy applications and their infrastructures. To minimize the business impacts, changes are typically incremental and are managed through a migration process with well-defined phases. The new security capabilities will present numerous challenges as they will have to be integrated into the technology and customer impact areas of the existing business such as downtime and interruptions.

OUR RECOMMENDATIONS

We favor the divide and conquer approach by dividing the enterprise into several business units, such as Corporate IT, e-Business, Manufacturing, Engineering Design, etc. Each business unit has domains of end users, applications, servers, workstations, and networks, which are typically separate for each business unit. The business units may leverage one or more corporate IT services for authentication, Internet access, DNS, DHCP, NTP, etc. We have compiled the following prerequisites for adding security to the enterprise:

1. Identification of the all networks for the business units.
2. A list of all hardware and software running on the network along with the security configurations.
3. Identification of the mission-critical services in the business unit and the impact of unavailability of network services on the business unit.
4. The communication model for each of the applications running within the business unit including end-user information.

5. Working knowledge of security weakness in the business units.

We recommend that each business unit be in a security enclave [16]. By definition, a security enclave has well-defined boundaries and comprehensive security policies for network operations, based on the requirements of the business unit. For example, one business unit may require the use of corporate authentication whereas another business unit could use local user accounts and groups. In another example, a business unit may require conformance to corporate minimum security specifications for its servers and workstations, whereas another business unit may have legacy computer systems that may never be upgraded. Business-unit-to-business-unit communications can be published and managed through the appropriate configuration of their enclaves.

The usage of enclaves solves operational issue number 1 as described in the Business Issues section. Using PENS with enclaves, as described in Figure 6, solves operational issue number 3 described in the Business Issues section.

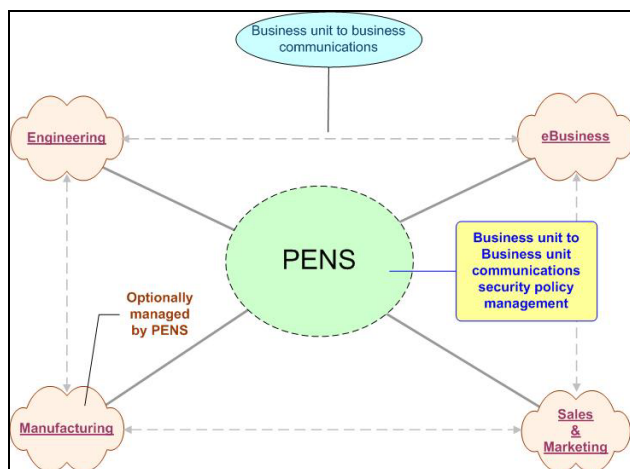


Figure 6: Using PENS to manage enclave-level communications

Not all elements of security described in this paper are required for every network environment. For example, 802.1x port security is a better fit for dynamic environments such as office users with remote access or wireless access. MAC address-based port security is more suitable for data centers and factory environments than 802.1x-based port security. Likewise, we also recommend super-VLAN and sub-VLAN technology for static environments such as data centers and factories but not for office environments. Table 1 lists our recommendations.

Table 1: Security recommendations

Environment	Security Model
Dynamic end-user such as office networks	802.1x
Static environment such as business critical systems (e-Business, Shop Floor)	Sub/super VLAN, static MAC address, packet filtering
Lab/development environment	Packet filtering
Partner, supplier, customer	Packet filtering
Remote access	802.1x and dynamic packet filtering

Lastly, we recommend allowing hardened systems to connect to the enterprise. There are several standards for defining hardened systems [17, 18]. Hardened systems contribute to security by controlling it at the source. If a service is not needed and turned off, it cannot be exploited over the network.

CONCLUSION

Until recently, security implementations have been reactive, their modus operandi being to keep current with the latest software updates and virus signatures. We have made security a proactive activity with integration of key capabilities into the network architecture so that networks and systems are delivered with security capabilities in place from the very outset. We have implemented several such enterprise networks in business-critical environments with no virus/worm infections in those areas in 18 months. The approach has been successful in letting our business units perform what is required of them while eliminating unauthorized and random access and probing. Adding new securities to an enterprise has the potential to increase operational costs. We are now focusing on putting together proper adaptive and dynamic policies that enable us to manage the security of the enterprise cost effectively without compromising the security of the enterprise.

ACKNOWLEDGMENTS

We thank Gary Morris, Shane Milburn, Kevin Heine, Amit P. Shah, Colm O’Halloran, Clark Mason, Mark Sokol, Anand Rajavelu, Praveen Sampat, and Richard Phillips for the various security capabilities that have been implemented in manufacturing. We also extend thanks to Ravi Sahita and Satyendra Yadav for their contribution to the development of policy-enabled architecture; and to Jonathan P. Clemens, Greg Tao, and Sridhar Mahankali

for reviewing this paper. We also appreciate the contribution of the ITJ editorial staff.

REFERENCES

- [1] "Internetwork Design Guide," Cisco Systems, <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/>*
- [2] "Open System Interconnection Reference Model," Cisco Systems, <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito/doc/introint.htm>*
- [3] "Port-Based network Access Control," <http://standards.ieee.org/getieee802/download/802.1X-2001.pdf>*
- [4] "RADIUS Protocol Security and Best Practices," Microsoft Corporation, <http://www.microsoft.com/windows2000/techinfo/administration/radius.asp>*
- [5] "RFC 3069–VLAN Aggregation for Efficient IP Address Allocation; Internet RFC/STD/FYI/BCP Archives," <http://www.faqs.org/rfcs/rfc3069.html>*
- [6] "Cisco–Securing Networks with Private VLANs and VLAN Access Control Lists," Cisco Systems, <http://www.cisco.com/warp/public/473/90.shtml>*
- [7] "RFC 1631–The IP Network Address Translator (NAT); Internet RFC/STD/FYI/BCP Archives," <http://computer.howstuffworks.com/framed.htm?parent=nat.htm&url=http://www.faqs.org/rfcs/rfc1631.html>*
- [8] "Stateful Inspection–Webopedia.com," http://networking.webopedia.com/TERM/S/stateful_inspection.html*
- [9] "Computer Security Dictionary: Packet Filtering (screening); ITsecurity.com," <http://www.itsecurity.com/dictionary/packfilt.htm>*
- [10] *Building Internet Firewalls*, Chapman & Zwicky, O'Reilly & Associates, Sebastopol, 1995.
- [11] Hong Li, Ravi Sahita, Greg Kime, Jac Noel, and Satyendra Yadav, "Policy-Enabled Network Security with Adaptive Feedback Loop and Capability-Based Data Model," *Eurescom 2003*, September 2003.
- [12] D. Verma, "Policy-Based Networking, Architecture and Algorithms," *New Riders*, November 2000.
- [13] "Understanding ACL Merger Algorithms and hardware Resources on Cisco Catalyst 6500," Cisco Systems, http://www.cisco.com/en/US/customer/products/hw/switches/ps708/products_white_paper09186a00800c9470.shtml*
- [14] "A summary of DoS/DDoS Prevention, Monitoring, and Mitigation Techniques in Service Provider Environment," Michael Glenn, *SANS Institute*, 2003. <http://www.sans.org/rr/papers/70/1212.pdf>*
- [15] "CERT Advisory CA-2003-04 MS SQL Server Worm," <http://www.cert.org/advisories/CA-2003-04.html>*
- [16] "DoD Electronic Business Architecture," United States Department of Defense, <http://www.amc.army.mil/amc/ci/matrix/documents/dod/jta31e.pdf>*
- [17] "Center for Internet Security," <http://www.cisecurity.org/>*
- [18] "The SANS Institute (SANS)," <http://www.sans.org>*

AUTHORS' BIOGRAPHIES

Sanjay Rungta is a staff network engineer with Intel's Information Services and Technology Group. He received his B.S.E.E. degree from Western New England College and his M.S. degree from Purdue University in 1991 and 1993, respectively. He is lead architect and designer for the Local Area Network for Intel. He has over 11 years of network engineering experience with three years of experience in Internet web hosting. He holds one United States patent in the area of Network Engineering. His e-mail is sanjay.rungta at intel.com.

Anant Raman is a staff engineer in Components Automation Systems with Intel's Technology and Manufacturing Group. He received his B.Tech degree from the Indian Institute of Technology, Bombay in 1981, a Master of Science in Mechanical Engineering in 1984, and a Master of Computer Science in 1991 from Arizona State University. He is the lead architect and designer for the e-Diagnostics infrastructure and manufacturing security initiatives within Intel. He also chairs the ISMT e-Diagnostics security sub team that is responsible for developing e-Diagnostics security guidelines for the industry. He holds three United States patents in the areas of software, network engineering, and security. His e-mail is anant.raman at intel.com.

Hong Li is a senior researcher with Intel's Information Services and Technology Group, responsible for trustworthy and survivable systems research. She led the development of several IT security strategies and architectures. She is also active within the Intel and external research communities. She is a 2004 Santa Fe Institute Business Network Fellow. Hong holds a Ph.D. degree in Electrical Engineering from Penn State University. She is also a certified information systems security professional (CISSP). Her e-mail is hong.c.li at intel.com.

Manish Dave is a staff network engineer with Intel's Information Services and Technology Group. He is lead engineer and designer for the Internet Connectivity and external network connectivity for Intel. He has over ten years of network engineering experience and network security experience. His e-mail is manish.dave at intel.com.

Greg Kime is a security architect with Intel's Information Services and Technology Group. He is responsible for the development and fostering of long-term security strategy and architecture. His focus is on wired, wireless, and platform security architectures along with process and governance development in a program called Enclaves. He is a CISSP. His e-mail address is greg.kime at intel.com.

Toby Kohlenberg is a senior information security technologist for Intel's ISTG Risk Management group. He has extensive experience in penetration testing, incident response, architecture design and review, and IDS, among others. In the last couple of years he has been responsible for developing security architectures for Intel's deployment of secure WLANs and Windows* 2000/Active Directory, and for the overall IDS strategy including the Security Operations Center. He is a handler for the Internet Storm Center and a co-author of the book *Snort 2.1* from Syngress. He currently is responsible for providing information security consulting to Intel product groups as well as evaluating new and emerging technologies. He currently holds the CISSP GIAC Certified Firewall Analyst (GCFW), Certified Incident Handler (GCIH), and (GIAC Certified Intrusion Analyst (GCIA) certifications. His e-mail is toby.kohlenberg at intel.com.

Legal notices at
<http://www.intel.com/sites/corporate/tradmarx.htm>.

Copyright © Intel Corporation 2004. This publication was downloaded from <http://developer.intel.com/>.

* Other brands and names are the property of their respective owners.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm